Universität Bielefeld
Technische Fakultät

Dissertation

zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften (Dr. rer. nat.)

Subliminal Copresence Systems

Christian Leichsenring

June 2014

Examiners:
Dr. Marc Hanheide
Dr. Thomas Hermann
Prof. Dr. Helge Ritter

Handed in for review:	June 2014
Oral examination:	February 2015
Final version for publication:	December 2015

Abstract

Telepresence research has focused on the ideal of recreating face-to-face conversations via remote mediated channels – maximising what has been termed *social presence*. A mostly overlooked aspect of communication is the simple sense of being together; the ability to be close to someone without necessarily having to interact consciously. Goffman [182] described this as *copresence*. We propose a class of systems to specifically support this mode of communication over a distance which we call *subliminal copresence systems* (SCS); they fulfil the definition of *calmness* as coined by Weiser [564]. We think such systems have the potential to fill a gap that was left during the increase in technologically mediated communication as a consequence of the rise in long-distance relationships in recent decades [507] and the rapidly growing importance of computer-based media in social interactions such as e-mail and especially online social networks [154, 196, 323].

In this work, we explore SCS in terms of their user acceptance, their effectiveness to convey information, their potential to influence and create feelings of connectedness and presence, and their ability to do so without being annoying or distracting.

To this end, we implemented two proof-of-concept subliminal copresence systems called FEELABUZZ and upstairs. In FEELABUZZ, two unmodified smartphones are used to constantly transmit one user's movements to another user as vibration of the phone and vice versa. In upstairs, two remote rooms are virtually stacked so that it sounds as if Room A were located above Room B and conversely, Room B were above Room A. For this, contact microphones are used to transduce the structure-borne sounds of the floor.

Three user studies were conducted to evaluate the effectiveness of these systems. Firstly, FEELABUZZ was shown to recognisably confer basic activity types. FEELABUZZ and upstairs were then tested in two related, longitudinal studies to evaluate the systems' capabilities to evoke a sense of copresence over a distance. Both systems were shown to create copresence to a significantly larger extent than social presence. User acceptance was higher for

upstairs than for FEELABUZZ and upstairs incurred a significantly smaller amount of cognitive load on its users.

To address remaining acceptance and cognitive load issues brought up in these studies, we investigate the introduction of automatic filters to reduce the necessity of interpretation by the users. We explore context recognition in the area of subliminal copresence systems and present two such systems – one using instant messaging clients and another using mobile phones. Both achieved good recognition rates but worrying about user acceptance of such black box systems whose subsymbolic models are incomprehensible even for experts, we present a system that extracts symbolic rules from these models. We compare existing methods for the extraction of such rules and also discuss different ways to present rule sets to the users in order for them to modify them and feed them back into the system. A partial prototype of one such system was implemented and is also discussed.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Motivation

The predominant ideal of interpersonal communication as pursued by fields such as telepresence research[1] and computer-supported cooperative work (CSCW) is basically the situation of two or more people sitting or standing together talking [39, 75, 481, 494, 543] which is usually simply called face-to-face communication. The ultimate goal of the research on telepresence systems for connecting people thus is to duplicate this rich and dense interaction for people who are physically remote. We do not want to argue against the importance of pursuing this goal and we acknowledge the need for much further research towards it. We do want to highlight, however, that it misses some maybe less obvious but nevertheless important aspects of what it means to share a physical space – aspects such as simply recognising the other person's presence, getting frequent simple cues about their activities, or being together in what is called *companionable silence* [462, 512, 529]. To illustrate, consider the following scenario (cf. Figure 1.1 on the next page):

> Susan and Geena are sitting in the same room, doing different things on their own. Both are fully immersed in their particular task. However, when one decides to seek to communicate with the other, she can mostly assess the interruptibility of the other quite easily. There are also subtle

[1] *Telepresence* also includes fields like teleoperation; as will be elaborated later on, *social presence* is the less well-known but more precise term.

Figure 1.1. Physical proximity enables subliminal cues while allowing concentrating on separate tasks.

cues to seek attention without interrupting too much [135]. Conversely, when one of the two needs attention, possibly due to an emergency or simply because she desires to communicate, this urge comes to her friend's attention quite easily. Usually, an attempt to interrupt can be ignored. This depends, however, on the urgency of the request and the degree of immersion in the task. Finally, switching from unconscious to deliberate communication is also easier and more seamless for both since the sense of closeness in an emotional way is much stronger when people are just sitting close to each other, even if they do not interact consciously.

Additionally, even at times when they are not in the same room but in nearby rooms within the same house or flat, Susan and Geena still perceive each other's presence through subliminal noises through the walls such as the clattering of plates, footsteps or muffled talking.

We propose that all of these subliminal cues are very important for a feeling of connectedness and presence over extended periods of time as well as to be able to escalate to more elaborate and full-fledged forms of communication as the ones traditional social presence aspires to achieve. We therefore propose a novel class of systems we call *subliminal copresence systems* (SCS).

1.2 Research Context

The mode of communication SCS provide has been largely neglected, not only in research but even in fiction [339]. There is one area of research, however, which set out with a very similar motivation and similar aims. This is the research into so-called *awareness systems* [350, 457]. The older and maybe better known *media spaces* [44] are a subclass of these [457]. Our subliminal copresence systems are a subset of awareness systems by definition but surprisingly, the concrete incarnations of awareness systems built or proposed so far have mostly been using the visual channel or speech-dominated audio, despite an emphasis on multimodality when motivating awareness systems [350]. This is why we emphasised the subliminality when naming our new class of systems.[2] We furthermore named them copresence systems instead of awareness systems because of the various problems with the term awareness and the fact that even in awareness research, connectedness and copresence are acknowledged as a more precise and better fitting term and measure (see Chapter 2 on definitions and discussions of these concepts and an overview of the discussion in the literature).

[2]Considering that we actually almost exclusively use Weiser's term "calmness" in this work, one might ask why we did not choose *calm copresence systems* as a more appropriate name. The answer to this fair question would be that we did not perceive "calm" as a readily understood term out of context, even though it is more well-defined within the scope of this work.

Figure 1.2. Pointing, touching, gesturing . . . and sharing coffee and biscuits [494] – where talking is concerned, there is nothing that comes close to face-to-face interaction. ⓒ Photographed by Chris JL, published under the Creative Commons *by-nc-nd* licence on Flickr.

Similar to calmness, privacy has been emphasised for awareness systems but the very choice of video or voice as channels oppose this goal.[3]

Section 1.5 elaborates on the difference between conventional awareness systems and subliminal copresence systems. The two are close, though, and much can be learnt from the thoughts and investigations that took place in that field.

Markopoulos et al. [351] for example stated about awareness systems that they “can bring about important, if subtle, benefits, such as increasing the effectiveness of collaborative work, fostering social relationships and improving the general well-being of individuals”, which we find equally apt to describe SCS.

While awareness systems research is the field that is closest to the research presented here, similar ideas and concepts had been around before.

Kuwabara et al. [312], for example, suggest *connectedness-oriented communication* as an alternative to what they call *content-oriented communication*. Connectedness-oriented communication is to be understood as focusing on the social relationships instead of on the message content and is characterised as being low-bandwidth. While this, too, sounds very similar to our concepts above, the concrete information to share is symbolic presence and status information.

Nowak and Biocca [401] distinguish telepresence, social presence and copresence. They describe *telepresence* as the sense of immersion in the simulation of a remote place. *Copresence*, as defined by Goffman [182], refers to the ability to perceive others and a sense of someone else's ability to perceive oneself, and is described by Goffman as a very basic form of interaction (twice called “mere copresence” [182]) as opposed to states he calls *focused interaction* or *full scale coparticipation*. For *social presence*, Nowak and Biocca [401] highlight the definition of Short et al. [494] who define it to be “the degree of salience of the other person in the interaction and the consequent salience of the

[3]Although not all researchers go so far as e. g. Tee et al. [528] whose system shares the contents of one's own screen.

interpersonal relationships". As the term "salience" at least suggests a conscious aspect, we could declare that our goal with SCS is to be able to achieve a maximum of copresence and connectedness *without* simultaneously and inevitably increasing social presence (cf. Definition 8). IJsselsteijn et al. [258] argued similarly for awareness systems that the level of social presence will be low, but the "sense of connectedness, the feeling of being in touch, can be strong". Definitions of these terms in the literature vary, though [39], so before we formally define these concepts for the context of our work in Chapter 2, it is maybe safer to just declare our goal to be achieving remote copresence that can work on a subconscious level and focuses on systems that can run in the background over prolonged periods of time. Chapter 2 also takes a deeper look at the presence concepts we just described, as well as at other concepts such as connectedness and awareness.

There have been a number of both academic and commercial systems that are stated to be or could be seen as promoting copresence or the feeling of connectedness. Most approaches are very explicit in that they need deliberate action to communicate through them or they translate implicit actions into more or less artificial symbols or cues to represent them on the other end of the communication [36, 72, 131, 191, 204, 281, 404, 554] (cf. Section 3.2). While there is not necessarily anything wrong with this approach, our focus lies on ways of communication that are as subconscious and thus implicit as possible at both ends of the transmission. This means we want to take something that people do anyway and present it to the other interactant in as natural a manner as possible.

There are many conceivable actions that can be exploited for such systems and there are a number of systems that have already transmitted such actions. Among activities that can be transmitted are

– the interactants' movements, transmitted audibly as if a floor or wall was between the interactants [50] (cf. our *upstairs* system, Chapter 5), or tangibly as in FEELABUZZ [326, 546] (cf. Chapter 4), although there, the translation is at least somewhat artificial due to hardware constraints, or visibly, for example as shadows,[4]

- usage of objects, acoustically for example with connected wardrobes, kitchen closets or desks, or through changing inherent properties of objects [111],
- induced temperature changes, such as the warmth of a chair or sofa from the touch of the remote interactant, and even
- involuntary movements such as breathing [225] or the heartbeat [570].

Section 3.2 gives a more detailed review of previous related systems and Chapter 3 discusses general aspects and guidelines when designing subliminal copresence systems. Section 8.2 goes into more detail regarding ideas for future SCS. We think that due to the fact that each channel on its own only transmits a limited amount of information and most systems have "blind spots" in that there exist situations when they cannot transmit any useful information at all (when the user does not move, does not interact with a smart object, does not trigger a sensor etc.), the most effective strategy might be to use a large number of such channels together. Of course we cannot test this hypothesis since we only actually implemented two such channels. Future smart homes, however, might greatly facilitate the creation of many channels based on the principles stated in Chapter 3 without introducing many new artifacts which would raise problems of user acceptance, obtrusiveness and cost.

1.3 Preliminary Considerations

All examples we mentioned in the previous section fall into one of two categories that might have different implications: 1) systems that try to emulate a presence behind a wall, door, window, floor or ceiling, and 2) those that introduce effects of a sender's presence within the same space as the receiving person without of course providing a simulation of the cause of these effects.Therefore there might be an

[4] A student project called *Window^2*, supervised by Till Bovermann and René Tünnermann implemented a video connection between two rooms in such a way, allowing the signal-to-noise ratio of the link to be regulated using a simple knob.

Box 1: Digressions
Throughout the text, there will be grey boxes like this one. They discuss topics and concepts that expand on things covered by the main body of text but that are not necessary to understand further passages, so feel free to ignore these little digressions.

uncanny poltergeist or intruder kind of effect or simply a more intrusive perception of the systems. On the other hand, the increased spatial proximity might be perceived as favourable for very close relationships. Increased proximity might also increase longing though. We will begin to investigate these questions in Section 6.5.1.7.

The best studied kind of long-distance relationships (LDRs) is the long-distance romantic relationship (LDRR). It is also the one that might profit the most from suchlike systems and it is one that is growing in number [507] with increased mobility and an explosion in media use such as e-mail, instant messaging (IM) and especially online social networks (cf. Foster [154], Gross [196] and Lee [323], and Figure 1.3 on the facing page) and to some degree even traditional media such as the telephone [228]. In any case, such systems offer opportunities for many kinds of investigations.

It is interesting to note that people in LDRRs have been found to report similar levels of satisfaction with their relationship as people in proximal romantic relationships (PRRs) [466, 467]. Although different reasons have been suggested, it is in this regard important to highlight that, generally speaking, PRRs compared to LDRRs trade desirable proximity for freedom.[5] Therefore, systems that aim to deliver a sense of proximity or togetherness but by doing so also bear the danger of limiting personal freedoms, have the potential to fare worse than either unaided LDRRs or PRRs. Conversely, they could – if done right – even improve satisfaction over both types or they might not change the particular measure of satisfaction with the relationship at all. In any case, this example shows how potentially crucial the right kind of

[5]Section 2.5 elaborates on LDRRs and their advantages and disadvantages compared to PRRs.

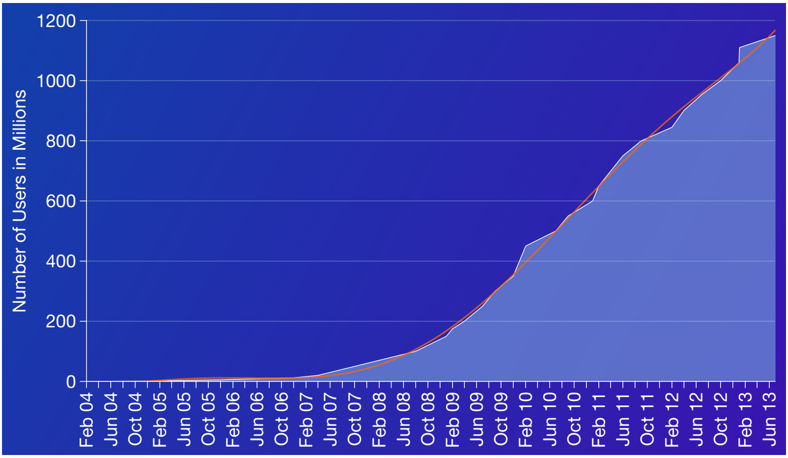

Figure 1.3. Development of monthly active users of Facebook over time, superimposed with an 7th order polynomial trend line. Source of data: Foster [154]

communication channel at the right time might turn out to be. This is another question we set out to investigate in Chapter 6.

There can be situations, however, when a constant stream of subsymbolic information is not desirable, even when respecting the design guidelines detailed in Chapter 3, especially when many such channels are active at the same time as suggested earlier. It might then increase calmness to automatically select subliminal channels based on the status of the users (especially their cognitive load and privacy concerns). The effectiveness of such measures critically depends on how well the system can identify the user status and react to it accordingly. As it is unlikely that a classifier that is not personalised will yield good results other than for the most commonplace situations a lot of initial training data would probably be needed to get good results for the varied everyday situations people find themselves in and therefore calmness would quite likely suffer greatly. It would be very helpful to use the world knowledge of the users about their own routines and preferences but it can hardly be expected that end users manipulate the subsymbolic and intransparent inner parameters on which all high-performing modern machine learning algorithms rely, such as neuron weights or support vectors.

An automatic system would preferably learn from user behaviour but would also take explicit statements in order to produce less errors that had to be corrected. There are many situations that can be easily described by simple rules of thumb but that would require many training examples. More fine-grained approximations of the desired behaviour could then be derived from such rules of thumb through data from cases where they do not apply. Such a "best of both worlds" approach seems very desirable to us for such systems. Additionally, learning systems that control important everyday applications should not be the black box systems most sophisticated trained machine learning models are.

We therefore decided in favour of a closed loop of a subsymbolic state-of-the-art machine learning system and a rule extraction mechanism that produces a *simplified* symbolic rule set from the subsymbolic model. This rule set ought to be comprehensible by non-expert end users and give them an insight into the most important dependencies that govern their personal trained model while at the same time, the system gives them a possibility to edit this rule set or add new rules to anticipate and prevent prediction errors when a user does not follow a daily routine (see Figure 7.21 on page 255 to see how these classification errors occur in case of week-ends after learning the working day routines). Finally, the system closes the loop by having a knowledge insertion part that feeds the changes on the symbolic user interface side back into the subsymbolic system so that the predictor adheres to them in future unless many contrary training data defy the user's maybe overly generic statements. We designed and began to implement one such a system called *Umber* which we will describe in Chapter 7.

The aim of a constant subliminal connection immediately raises questions about privacy. Even though we often enjoy sharing a room with someone we like, we at times just as much enjoy being able to close a door between us and this other person in order to put a wall between us. More generally, we prefer to adjust the amount of information we share [414]. The perfect system therefore would enable us to adjust the kind and amount of information that is being transmitted. It would also allow users of the system to monitor what is transmitted of and by them. In principle, it is thereby possible to exhibit much more fine-grained control over the kind and amount of information transmitted than it is possible by closing a door and this manipulation of outgoing information can be done differently on different channels and modalities.

This could be called an *adjustable door* or maybe an *adjustable wall* and can in its most simple form be imagined as a kind of communication gradient that runs from a full-blown audio-visual connection (which we call high-level channels), perhaps augmented by new supporting channels, down to very basic channels that transmit very little information and are thus more likely to be accepted as running all the time (low-level channels).[6]

Defining such a one-dimensional gradient has the advantage that it could be navigated relatively easily by the users themselves, i. e. with a simple knob or slider. On the other hand, it is a simplified construct that can never hope to actually adequately reflect reality [110]. A more fine-grained control is probably impractical to do manually the whole time though. This is where we envision our Umber system to be useful (cf. Chapter 7).

Another level of privacy concern is also addressed by using low-level channels in that we expect them to depend largely on context information because they must not be complicated to process for the human brain (i. e. have a low cognitive load), even though this does not mean that it is not possible for users to extract complex information. It helps that humans usually have very little problems integrating a very rich context to interpret data at hand. This context dependence means that only knowing a person quite well (for example their schedule or habits) will enable the interpretation of ambiguous data to turn it into useful information. This helps to ameliorate any privacy intrusion in that when the wrong person gains a handle on the data stream, it will not reveal much actual information to such an eavesdropper.

[6]See Definition 10 and its adjoining remarks for definitions and a brief discussion of the terms *high-level* and *low-level channels*.

To take FEELABUZZ as one example, Geena might not perceive more than rapid bursts of acceleration but because she knows a lot about Susan producing the movement, she immediately realises that she is running to catch the bus to be there in time (cf. Section 6.5.2.2 for real-world examples of this). Such examples can be given for most of the described systems – context produces an additional layer of expressiveness and ease of interpretability that the data itself lacks.[7]

Thirdly, the persisting ambiguity helps to maintain the important aspect of plausible deniability. See Chapter 2 for an explanation and a definition of plausible deniability (Definition 13) and of privacy (Definition 12).

The importance of context also means, however, that we do not believe that using subliminal copresence systems exclusively would work for an extended period of time. In order to get the context information one needs to make proper sense of the SCS, it is mandatory to perform explicit, symbolic communication (e. g. language) and richer, more obtrusive forms of communication (audio or audio-visual links or of course real life conversations). We expect this to also be important for the emotional dimension that is an important part of SCS – we actually postulate that both implicit and explicit communication will play their part and that stripping one or the other from an interpersonal relationship will impose serious strain to it although we acknowledge that explicit without implicit communication is arguably the more usual and less devastating form of constriction than the other way around.

Besides the LDRRs we already briefly mentioned, and will discuss to some extent in Section 2.5, some other examples of such constricted interpersonal relationships that are interesting to consider are solitary confinement, annoying neighbours heard through thin walls or through the floor, pen pals or relationships described in epistolary novels, and e-mail and web forum conversations that regularly run out of control and end in so-called flame wars.

[7] For more on context, refer to Section 7.4.

1.4 Scope, Objectives and Method

The aim of this work and the research presented in it is to explain and explore a new class of awareness systems called subliminal copresence systems. Two exemplary systems are presented in concept and as implementations and they are both tested to show that they effectively accomplish their desired functionality. These systems are also used to test some hypotheses regarding the effect of subliminal copresence systems on romantic relationships. For this, a series of questionnaires are developed and the effectiveness of their different scales is measured.

Partly based on the experiences from these evaluations, the need for an automatic but comprehensible filtering is argued for and one such system is detailed and partly implemented.

This work also aims to be a useful guide for the design and development of additional subliminal copresence systems.

The methods used include a systematic survey and consolidation of the available literature on awareness systems and neighbouring fields, case studies including implementation and evaluation, discussions with peers and researchers from other fields as well as study participants in order to profit from their first-hand experience with the systems, and statistical reliability tests to improve our questionnaires for future deployment.

1.5 Contributions

While awareness research has been creating systems to support peripheral awareness of remote communication partners in one way or another for two decades, these systems have almost exclusively been using visual cues and speech audio [350]. In this work, we present subliminal copresence systems as a new class of systems whose key properties are calmness (following the definition of calm technology by Weiser and Brown [566]) and a preservation of richness in the transmission that we call *directness* (cf. Definition 11).

In the following chapters, we

- define the new class of awareness systems we call SCS and discuss its design considerations,
- showcase – through implementation and case studies – two systems derived from our design considerations,
- develop questionnaire-based measures specifically targeted at SCS and apply them to our systems, and
- propose a workflow and implementation to enable symbolic presentation and manipulation of subsymbolic automatic filters to manage the diversity of data flows from awareness systems and thereby enhance their calmness.

2 Relevant Concepts

There are a lot of technological, psychological and sociological concepts that are used in awareness systems and (tele-)presence research and that are in turn relevant for the class of systems called subliminal copresence systems (SCS) we propose. Awareness and presence are terms that are accompanied by a huge cloud of related terms and concepts;[1] frustratingly, one of the strongest common links between these terms is that their definitions, understandings and interpretations vary – sometimes greatly – from researcher to researcher, even within the same field [11, 38, 39, 321, 350, 360, 440]. For presence-related terms, Lee [321] attributes this at least partly to the great number of disciplines that are investigating presence, naming business, communication, computer science, education, industrial engineering, psychology, and sociology. Reis et al. [440] similarly attribute the remarkable influx of new constructs in the relationship literature to the great number of fine distinctions between concepts that are in principle related and the examination of different aspects of the same phenomena. The consequential diversification of the literature itself has led to definitions and theories that are similar but nevertheless hard to relate to each other. Aron and Mashek [11] make out such a problem of a fragmentation of definitions and interpretations even for core concepts and terminology in their

[1] These include telepresence, copresence, telecopresence, social presence, virtual presence, mediated presence, remote presence, objective and subjective presence, physical presence, self presence, immediacy, intimacy, mutual awareness, connected presence, connectedness, social connectedness, interpersonal connectedness, togetherness, engagement, and involvement.

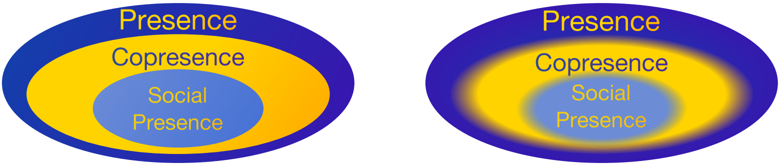

(a) (b)

Figure 2.1. Illustration of presence, copresence and social presence as used in this work. These distinctions can be seen as classes (a) or as points on a continuum (b).

field (in their case *intimacy* and *closeness*).

We will not be able to solve this conundrum here but we will nonetheless try to present some of the more relevant and most commonly found terms. We will also at least partly alleviate the problem by focusing on those aspects that are most relevant to SCS and by liberally brushing over some of the finer details and distinctions that might arguably not be as important to this work.

2.1 Definitions

We will define presence/telepresence, social presence and copresence separately in the following, with remote presence being used synonymously with telepresence. Take Figure 2.1 as a first overview that is going to be explained in more detail after the definitions. Figure 2.4 on page 35 is a much more comprehensive overview of all of the concepts presented in the following which also includes awareness, connectedness, immediacy, closeness, and intimacy.

Telepresence

Telepresence is a term invented by Pat Gunkel and publicised by

Minsky [380] in 1980 while the idea itself goes back at least to works Raymond Goertz did in the mid-1940s [493]. Minsky described the concept that users of teleoperation systems were to have the impression of actually being at the remote site when operating their particular manipulator. Motivations ranged from hazardous environments over the global availability of rare expertise (such as in remote surgery) to the sheer comfort of telecommuting. Later authors kept the basic definition but often dropped the association with telemanipulation[2] and extended the term to virtual environments [371, 401, 439, 445, 499], even though Sheridan [492] tried to establish the term **virtual presence** for these cases.

Another term used synonymously with the latter definition of telepresence is **mediated presence**, used in communication research to emphasise that physical presence should not be a subject matter of presence research [38, 321]. Lee however argues in his much regarded review article that even physical presence involves some form of mediation and that there is no fundamental difference between natural perception and virtual experiences [25, 337, 343]. This means that it is advisable to pay close attention to this distinction when studying the presence literature because the views that presence – as opposed to telepresence – refers *only* to the direct and unmediated presence in an environment [517] or that presence works as superordinate term including natural and technologically mediated presence (telepresence) [474] can also be found.

In this work, telepresence is used interchangeably with remote presence. Both terms are qualifiers for presence to emphasise the lack of co-location of the interactants.

Presence

Presence, as just indicated, is a term on whose exact definition,

[2] While Schloerb [474] argues that some form of teleoperation is a prerequisite of what he calls *subjective telepresence*, i. e. the feeling of being at a remote place as opposed to *objective telepresence* which is defined by the mere fact of a "causal interaction" with a remote place.

judging from the literature, there does not appear to be a consensus within presence research. It is interpreted in widely varying ways, to the extent of even being charged metaphysically [321]. The closest to a consensus definition was issued by the International Society for Presence Research [262] in 2000. They define presence in the very first sentence of their "Explication Statement" to be "a shortened version of the term 'telepresence'" but qualify this statement a paragraph later, saying that natural perception of the world constitutes "a broader conception of the term 'presence' – i. e., not a shortened version of 'telepresence'". They still narrow the understanding of presence and all its aspects to that part of perception (of a real or simulated environment) that is mediated by technology and in fact it is essential for presence that "part or all of a person's perception fails to accurately acknowledge the role of technology in her/his perception".

The International Society for Presence Research [262] lists a number of dimensions that constitute presence. These are a) immersion, b) perceptual realism, c) social realism (believability of social interactions encountered), d) social presence (convincing illusion of a two-way communication where it is only one-way[3]), e) copresence (see below), and f) the somewhat abstract "medium as social actor" which describes the believability of an agent or a similar technological entity.

In any case, the subtle differentiation of what constitutes presence and especially whether it is different from telepresence is not needed within the scope of this work. Therefore, we content ourselves with the term's meaning in standard English:

Definition 1 (From the Oxford Dictionary of English): Presence is "the state or fact of existing, occurring, or being present" [410].

The multi-volume Oxford English Dictionary slightly expands on this definition, in that there, presence is the "fact or condition of being

[3] Note that this is of social presence is utterly different from what we discuss below.

present; the state of being with or in the same place as a person or thing; attendance, company, society, or association" [412].

The definition from Merriam-Webster's dictionary reads quite similar since there, presence is "the fact or condition of being present" [188]. These definitions are neutral to any kind of mediation. Whether or not presence is equivalent to telepresence or a subordinate to it is not relevant with regard to our understanding of copresence and social presence. This is illustrated in Figure 2.1 on page 16: presence includes both social presence and copresence (and indeed every other form of "-presence").

Copresence

Copresence has also seen its fair share of definitions. The aforementioned International Society for Presence Research [262] define it very broadly and say: "'Co-presence' [occurs] when part or all of a person's perception fails to accurately acknowledge the role of technology in her/his perception that the person or people with whom s/he is engaged in two-way communication is/are in the same physical location and environment when in fact they are in a different physical location". Rogers et al. [452] define it completely differently, as the possibility for *co-located* people to share an interface. While the other differentiations are more or less nuances, this is a completely different conception of copresence and it has to be watched out for when reading the literature. Sometimes copresence is also used synonymously with social presence [38, 548].

Regarding copresence as a form of telepresence however, Biocca et al. [38] conclude: "Of all the work defining co-presence, Goffman's is by far the most subtle, elaborated, and developed even though [it] dates back to [the] early 1960s."

As briefly mentioned in Chapter 1, the term copresence was introduced to the field of sociology by Goffman [182] in 1963 and describes the ability to perceive others and a sense of another's ability to perceive oneself. It is described by Goffman as a very basic form of interaction and he opposes copresence to what he calls *focused interaction* or *full scale coparticipation*.

The subtlety Biocca et al. [38] mention probably stems from the fact that Goffman discusses many interesting examples, including different modalities and conditions of voluntarily or involuntarily asymmetric presence perception. We will not repeat all these considerations here but to retain said subtlety, we will quote Goffman in some length in the following. Goffman [182] leads up to the term by writing:

> First, sight begins to take on an added and special role. Each individual can *see* that he is being experienced in some way, and he will guide at least some of his conduct according to the perceived identity and initial response of his audience. [...] Further, he can be seen to be seeing this, and can see that he has been seen seeing this.

Goffman's concentration on the visual modality can also be seen in the following quote, where he differentiates between focused and unfocused interaction, saying that

> *unfocused interaction*, that is, the kind of communication that occurs when one gleans information about another person present by glancing at him, if only momentarily, as he passes into and then out of one's view. Unfocused interaction has to do largely with the management of sheer and mere copresence. The second step deals with focused interaction, the kind of interaction that occurs when persons gather close together and openly cooperate to sustain a single focus of attention, typically by taking turns at talking.

This is a distinction that we follow closely with the concepts of copresence and social presence (see below).

However, Goffman also generalises copresence to other modalities; he writes the following after having a passage on the lack of a good word to describe perception regardless of the senses involved.

> The full conditions of *copresence*, however, are found in less variable circumstances: persons must sense that they are close enough to be perceived in whatever they are doing, including their experiencing of others, and close enough to be perceived in this sensing of being perceived.

Later in the same work, he talks about the sometimes fuzzy border between simple copresence and more involved forms of interaction that – as already hinted at – we will be calling *social presence*:

> I have suggested that a face engagement is a sufficiently clear-cut unit that an individual typically must either be entirely within it or entirely outside it. This is nicely borne out by the trouble caused when a person attempts to be half-in and half-out. None the less, there are communication arrangements that seem to lie halfway between mere copresence and full scale coparticipation, one of which should be mentioned here. When two persons walk silently together down the street or doze next to each other at the beach, they may be treated by others as "being together," and are likely to have the right to break rather abruptly into spoken or gestured communication, although they can hardly be said to sustain continuously a mutual activity. This sense of being together constitutes a kind of lapsed verbal encounter, functioning more as a means of excluding nonmembers than as a support for sustained focused interaction among the participants.

We recognise this lack of clear borders between different kinds of presence in Figure 2.1(b) on page 16. Such verbosity as used by Goffman in discussing copresence can be somewhat unwieldy though, so we will define copresence in the context of this work in a much more compact form. We chose to adapt the definition by Youngblut [585],[4] because to our minds it covers the essence of copresence and we can take advantage of the work Youngblut did on how to measure this concept.

Definition 2: Copresence is the subjective experience of being together with others who are situated at different physical locations.

This is obviously a definition that equates copresence and telecopresence. If one wanted to distinguish the two, however, it would work equally well if the latter part about physical co-location were omitted.

Zhao [591] suggested a taxonomy of copresence (shown in Table 2.1 on the facing page) but their understanding of copresence seems to be extremely general while at the same time focused on human-computer interaction (HCI).

Social Presence

Social presence is no less diverse in how it is defined than the previous concepts were but Biocca et al. [38] made out certain clusters. They summarise social presence definitions by providing a list of words that are synonyms under each definition: "co-presence", "co-location", "apparent existence, feedback, or interactivity of the other", "sense of being together", "perceived access to another intelligence", "salience of the other", "mutual understanding", and "interdependent multichannel exchange of behaviours". Biocca et al. continue by pointing out some limitations of the previous definitions and conceptualisations.

[4]The original definition in the work of Youngblut [585] is: "Copresence is the subjective experience of being together with others in a computer-generated environment, even when participants are physically situated in different sites".

Mode of Presence \ Interactant Location	Physical Proximity	Remote Location
in Person	Corporeal Copresence (face-to-face)	Corporeal Telecopresence (face-to-device)
in Simulation	Virtual Copresence (physical simulation: instrumental robots, communicative robots)	Virtual Telecopresence (digital simulation: intrumental agents, communicative agents)

Table 2.1. Basic categories of copresence after Zhao [591]. The table header names are slightly adapted for brevity.

Following closely the early definition of Short et al. [494],[5] we define social presence as follows:

Definition 3: Social presence is the degree to which a person is feeling involved in a particular interaction with another person or agent and the consequent perceived intensity of the interpersonal relationship.

We concur with Short et al. that social presence is a subjective property of the medium and that it depends on objective qualities of the medium and on the user. It is therefore not an objective measure.

While achieving social presence is the goal of most mediated communication techniques, it is not always desirable. There are obvious privacy concerns but beyond that, social presence uses cognitive resources, even when the individual is not actively communicating

[5] Short et al. [494] define social presence to be "the degree of salience of the other person in the interaction and the consequent salience of the interpersonal relationships" (cf. Section 1.1). We want to avoid the term "salience" because it has different meanings in different disciplines and replace it by the more universally accessible notion of involvedness.

but occupied with other tasks and moreover even when the link is only unidirectional [53].

Awareness

Awareness, again, is a word that has been associated with a wide variety of meanings. Schmidt [477] criticises that the "very word 'awareness' is one of those highly elastic English words that can be used to mean a host of different things. Depending on the context it may mean anything from consciousness or knowledge to attention or sentience, and from sensitivity or apperception to acquaintance or recollection". He even goes so far as to say that in the context of CSCW, "the term 'awareness' is being used in increasingly contradictory ways. [...] In fact, it is hardly a concept any longer". Rettie [443] also refers to this passage but tries to order different definitions regardless. She quotes an early and – according to Nova [400] the best-known – definition by Dourish, Bly and Bellotti [124, 125] which is that awareness is "an understanding of the activities of others, which provides a context for your own activity". Rettie then highlights work by Christiansen and Maglaughlin [81] who identified a total of 41 different awareness phrases whom they categorised into four types: workplace awareness, availability awareness (where "availability" means the availability of people and objects), group awareness, and contextual awareness. Rettie [443] concludes that a central element of awareness is being synchronous or near-synchronous which is a surprisingly fuzzy criterion. She however points out that the concept of *connectedness* is more well-defined and should therefore – counterintuitively – be preferred over awareness when evaluating awareness systems. This is a view that IJsselsteijn et al. [257] seem to share since a main focus in the questionnaire they developed to evaluate awareness systems is connectedness. Romero et al. [457] call their concept of connectedness a consequence of their research on peripheral awareness of close family members. Dey and de Guzman [114] relate awareness and connectedness by saying that awareness is "often a means to an end, with researchers actually being interested in how to support increased feelings of connectedness between people."

Markopoulos et al. [351] regard the critique by Schmidt – that awareness is used in contradictory ways – as met by the subsequent concretisation of the meaning of awareness in the respective fields that had each overloaded the term. They redefine awareness themselves in the context of awareness systems. While acknowledging earlier definitions as well as aspects such as *situation awareness* as studied in the field of human factors,[6] they shift their concept of awareness to suit awareness systems, focusing on "the social interactions between individuals and groups":

> Individuals may seek awareness for their own sake, as a means for understanding their own self, reflecting on relations with others or simply as a means for engaging with their social network. At work they may seek awareness of the social context, giving rise to informal and serendipitous interactions [125] and knowledge sharing. At leisure they may seek the formation and strengthening of social ties [352] or provide affective support to each other [457].
>
> Contrary to the assumptions underlying the notion of situation awareness, accuracy and completeness of awareness is not the golden standard to strive for. Full knowledge of activities of others is usually not at all desirable, with people preferring to control the flow of information from others [414] or preferring to cooperatively agree on setting limits to this awareness in order to facilitate social processes and allow equivocation and politeness [8].
>
> [The citations in the quotation were adapted to match this document's referencing scheme.]

This notion of incompleteness of awareness is (literally) best illustrated by the focus and nimbus model of awareness by Rodden

[6] One of the definitions for situation awareness they highlight is the one by Endsley and Garland [132] who define it to be "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future".

[450], shown in Figure 2.2 on the next page. Rittenbruch and McEwan [447] describe this model as "regarded as influential" but "not widely adopted". We regard this model as useful and will define awareness in a way that is compatible with it while still maintaining applicability to Rettie's [443] notion of including objects as well as people as subjects of awareness. We do this because it facilitates some comparisons between mediated communication and ambient displays.

Definition 4: Awareness is a person's perception of cues that are relevant for a certain activity.

Based on this we can define **peripheral awareness** or **peripheral attention** as follows:

Definition 5: Peripheral awareness is awareness that happens with minimal or completely without distraction.

Gross et al. [197] further divide awareness into five categories they call group awareness, social awareness, task-specific awareness, situation awareness and objective self-awareness. See Table 3.2 on page 54 for a short overview and their speculations on the best technological means to influence these types of awareness.

Awareness systems are defined by Markopoulos et al. [351] as "systems intended to help people construct and maintain awareness of each others' activities, context or status, even when participants are not co-located".

This term has also suffered from being applied very generally. Markopoulos et al. [351] and Rittenbruch and McEwan [447] even describe e-mail and instant messaging as awareness systems.

Markopoulos [348] describes the goal of awareness systems as very similar to that of SCS:

> A common denominator for research in [the field of awareness systems] is the ambition to go beyond means for rich

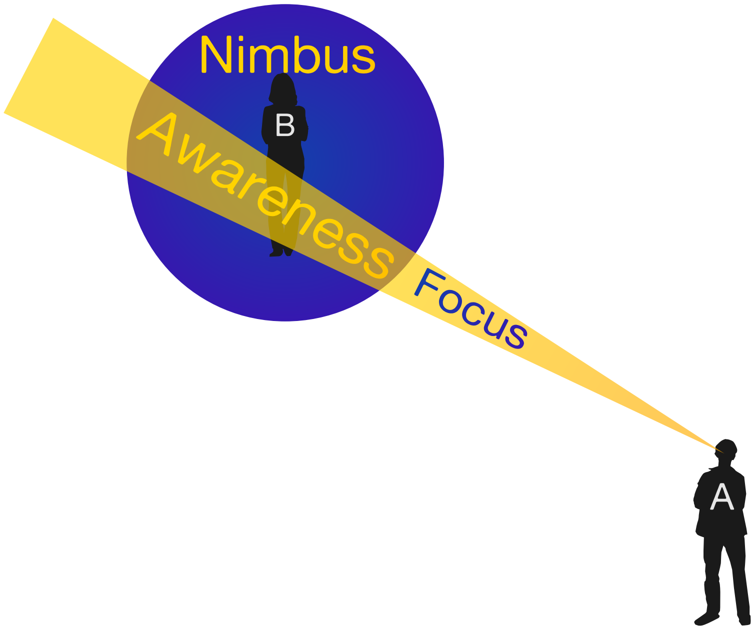

Figure 2.2. Illustration and varnished Venn diagram of the set-theoretical focus and nimbus model of awareness by Rodden [450], based on the more specialised COMIC model by Benford and Fahlén [33].
The nimbus $N_B(t)$ is the information that is in principle perceivable of a person B at time t and the focus $F_A(t)$ the information perceived by person A at time t, restricted by their attentional capabilities (adapted from McEwan and Greenberg [366]). The awareness is then $W_{A \to B}(t) = F_A(t) \cap N_B(t)$. The focus is explicitly *not* restricted to the visual modality, even if this abstracted illustration might suggest otherwise.

> and efficient information exchange and onto supporting sustained and almost effortless communication between individuals or groups. Such sustained communication, it is hoped, will enable users over time to build up and maintain an understanding of the activities of each other. Awareness of others can directly address social and affective needs, or, less ambitiously, act as a trigger and frame of reference for communication through other media [457] that are better suited for addressing these needs.
>
> [The citation in the quotation was adapted to match this document's referencing scheme and punctuation was silently corrected.]

This means that – even though the motivations for awareness systems and SCS are very similar – not only are the types of systems built so far different (cf. Chapter 1) but so is the immediate goal: awareness systems try to create an awareness of the other person in whatever way while subliminal copresence systems specifically try to create a sense of copresence. Abstraction is a perfectly adequate device in awareness systems [422, 423] whereas interviews with the participants of our study described in Chapter 6 indicate that SCS rely on a feeling of immediacy.

Nevertheless, most definitions of awareness systems implicitly include SCS as a sub-category (cf. Section 2.3).

Connectedness
Social Connectedness

Connectedness, also sometimes called social connectedness, is among the more recent concepts described here. Although it has seen an increased usage in recent years, it has not been studied as extensively as most other concepts and there is also no great number of definitions to choose from. Markopoulos [348] defined it broadly and somewhat indirectly as "an assessment made by individuals regarding their emotional distance from others". The most cited definition is the one by van Baren et al. [549] who define connectedness as "a positive emotional appraisal which is characterized by a feeling of staying in touch within ongoing social relationships".[7] Romero et al. [457] say that connectedness is characterised by a persistence beyond the duration of the actual social interaction. IJsselsteijn et al. [257] describe social connectedness as an alternative measure to social presence since they also postulate that awareness systems do not promote social presence because they lack the informational content or "media richness" as they put it. Social connectedness is proposed as a measure that can capture the emotional closeness that can nonetheless be promoted by awareness systems.

The Cambridge Business English Dictionary [67] defines connectedness as "the state of being connected and having a close relationship with other things or people".

In social psychology, the seeking of connectedness is one of the three basic motivational principles [500].[8] This is defined as people's "attempt to create and maintain feelings of mutual support, liking, and acceptance from those they care about and value".

Based on these definitions an attempt of a consensus definition can be made by unifying most of the properties mentioned:

Definition 6: Connectedness is the mutual feeling of support, liking, acceptance, caring, valuing, staying in touch and a lack of emotional distance that can persist beyond the duration of the actual experience of presence.

The actual relative and absolute importance of each of these properties would have to be determined by future research in this area.

Immediacy

Immediacy is a concept that is related to closeness (see below). It is defined as follows.

[7] Van Baren et al. [549] attribute this definition to IJsselsteijn et al. [258].
[8] The other two are the *striving for mastery* and the *valuing of "me and mine"*.

Definition 7: Immediacy is the psychological distance speakers put between themselves and their listeners [258, 494].

The role of the concept in language has been explored extensively by Wiener and Mehrabian [575] and Short et al. [494] summarise the extensions of Wiener and Mehrabian's work by other authors.

Closeness

Closeness is defined in the influential work by Aron et al. [12] in a somewhat behaviourist manner as the inclusion of someone else's resources, perspectives and identities in oneself (conjointly called *inclusion of other in the self* or IOS).

Aron and Mashek [11] structure the definitions of closeness they perceived in three groups: closeness as intensity and frequency of interactions, closeness as mutual influence, and closeness as cognitive connectedness (this includes IOS).

Closeness seems a very similar concept to immediacy, especially since the definition of immediacy above can be restated in a semantically equivalent way as the "amount of psychological closeness" between two people. What can generally be seen in the literature is that authors who use immediacy as a well-defined term tend not to use closeness or only use it in a loose understanding. If one was to differentiate the two concepts, one could argue that closeness is a quality of a relationship between two people that works on a longer time scale than immediacy. For this work, however, the two terms can be seen as synonymous.

Intimacy

We will not use the concept of intimacy in any formal sense in this work but since it is a frequently used concept in related work, it is helpful to understand how it has been used and how it relates to the concepts we already discussed. Intimacy is a quality of a relationship that is created, for example, by self-disclosure or joint activities [141]. At this point it should not come as a surprise that there is no consensus in the literature on an exact definition [288, 426, 431]. However, Prager and Roberts [433] consolidate and build upon previous work [333, 432, 441] when they define intimacy using the following three conditions: self-revealing behaviour (possibly including nonverbal and physical interaction), positive involvement with the other person (i. e. a positive regard of and undivided attention towards one another) and shared understandings. Prager and Roberts describe the latter point in a way similar to the description of inclusion of another's perspectives in the self by Aron et al. [12] and the criterion might indeed be more comprehensively and more concisely be described as *closeness*.

As they did for closeness, Aron and Mashek [11] again make out several broad classes into which all definitions of intimacy fall: intimacy as moments of felt connection and mutuality, intimacy as self-disclosure and vulnerability, intimacy as mutual responsiveness and felt concern for the other's welfare, and intimacy as physical touch and sexuality.

One can also often find an understanding of intimacy that is based on the influential early work of Argyle and Dean [10]. They argue in their Intimacy Equilibrium Model that interaction partners seek to keep the amount of intimacy during a conversation at a certain level that, amongst other things, of course depends on their relationship. They do so using verbal and non-verbal means. If the intimacy level deviates from the optimum in either direction, this is perceived as unpleasant. This understanding, although often mixed up with the previous one, seems to work on a different time scale and the intimacy of a relationship under the former set of definitions can maybe be compared to the optimum *amount* of intimacy under the Intimacy Equilibrium Model.

Subliminal Copresence Systems

With the above terms discussed and – where necessary – defined, we have not only laid the groundwork for further discussion but we also have all parts in place to formally define subliminal copresence systems.

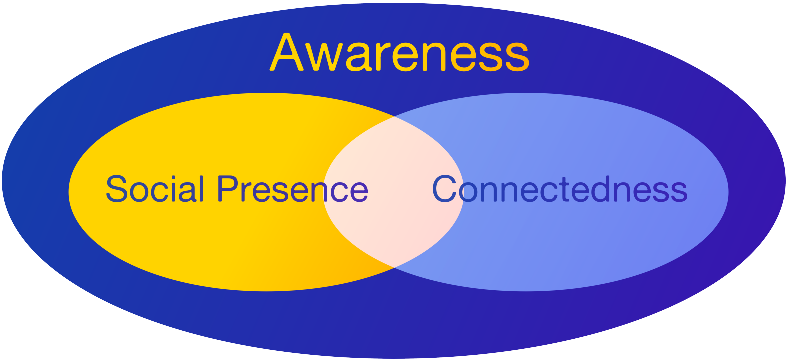

Figure 2.3. Relationships between awareness, social presence and connectedness as an Euler diagram after Rettie [443].

Definition 8: Subliminal copresence systems are systems that connect two or more people in a way that aims to increase copresence and connectedness while minimising any concomitant increases in social presence and distractions, even during prolonged use.

The exact implications for the design of such systems are not trivial and will be discussed in detail in Chapter 3.

2.2 Relationships Between Concepts

In the following we will try to draw parallels and connections between the various terms presented above and put them in relation to each other.

With the definitions used in this work, both social presence and copresence are a kind of presence and copresence is a more basic concept than social presence. Social presence and connectedness are both subsets of awareness and are not disjoint [443]. These relations are shown in Figures 2.1 and 2.3. Since awareness in the wider sense also encompasses the awareness of objects as well as people [443], awareness can in turn be declared a superset of presence. More formally, it can be said that

$$\text{awareness} \supsetneq \text{presence} \supsetneq \text{copresence} \supsetneq \text{social presence} \tag{2.1}$$

and

$$\begin{aligned} \text{awareness} &\supsetneq \text{social presence} \\ \text{awareness} &\supsetneq \text{connectedness} \\ \text{connectedness} \cap \text{social presence} &\neq \emptyset. \end{aligned} \tag{2.2}$$

Nowak and Biocca [401] give an overview of the usage of the three presence terms that are important in the context of this work. They distinguish *telepresence*, *social presence* and *copresence*. Telepresence is defined as the sense of immersion in the simulation of a remote place. For social presence, Nowak and Biocca highlight the definition of Short et al. [494] who define it to be "the degree of salience of the other person in the interaction and the consequent salience of the interpersonal relationships". The usage of the term "salience" points to a conscious aspect; this is why we could declare that our goal with SCS is to be able to achieve a maximum of copresence *without* simultaneously and inevitably increasing social presence (cf. Definition 8). However, considering the large number of varying definitions of each of the presence terms, this concise statement is in danger of being misunderstood when taken out of context.

Rettie [443] compares *social presence*, *awareness* and *connectedness*, calling connectedness "a more fundamental concept" than the other two in the context of awareness systems. She especially remarks that with awareness systems "there may be virtually no social presence, i. e. little salience of the other person, and no access to their intelligence, intentions or sensory impressions, however, at the same time there may be a feeling of psychological involvement and experience of connectedness."

Short et al. [494] explain the difference between *immediacy* through the use of a certain medium (so called *technological immediacy*) and *social presence* by using the example of someone who uses a telephone to call his colleague in the adjacent office instead of walking over. This creates a psychological distance and thus non-immediacy between the two but does not affect the social presence which is solely a function of the medium. Similarly, when the physical distance between two people makes a telephone call an option that in itself does not decrease immediacy, an interlocutor might still speak in a manner that puts distance between the two, still without influencing social presence.

As mentioned earlier, *closeness* and *immediacy* are very similar concepts and based on the definitions, *connectedness* and *immediacy* also have a large overlap. Connectedness is actually one of the classes of definitions Aron and Mashek [11] identified for closeness and they say this includes their own influential definition called IOS [13]. It is therefore probably safe to say that closeness, immediacy and connectedness are very related concepts. In this work, for the sake of simplicity, we will predominantly use the term connectedness because it is the most prevalent in the awareness systems literature.

From the definitions above also follows that intimacy is a superset of closeness:

$$\text{closeness} \subsetneq \text{intimacy} \tag{2.3}$$

Figure 2.4 on the next page is extending the Euler diagrams[9] of Figures 2.1 and 2.3 and combining them with the conclusion drawn in the last paragraphs in an attempt to unify all concepts presented in Section 2.1. Contrary to Equations (2.1) to (2.3) and the relations shown in Figures 2.1 and 2.3, this unified Euler diagram is not completely backed by the literature though and details might therefore be subject to discussion and change.

[9] Euler diagrams are a generalised form of the better known Venn diagrams. Venn diagrams show all 2^n intersection between the sets depicted while an Euler diagram shows only the relationships that actually represent the reality it is meant to illustrate [21, 451]. Venn diagrams must use shaded areas to mark those subsets that are empty because they cannot omit them.

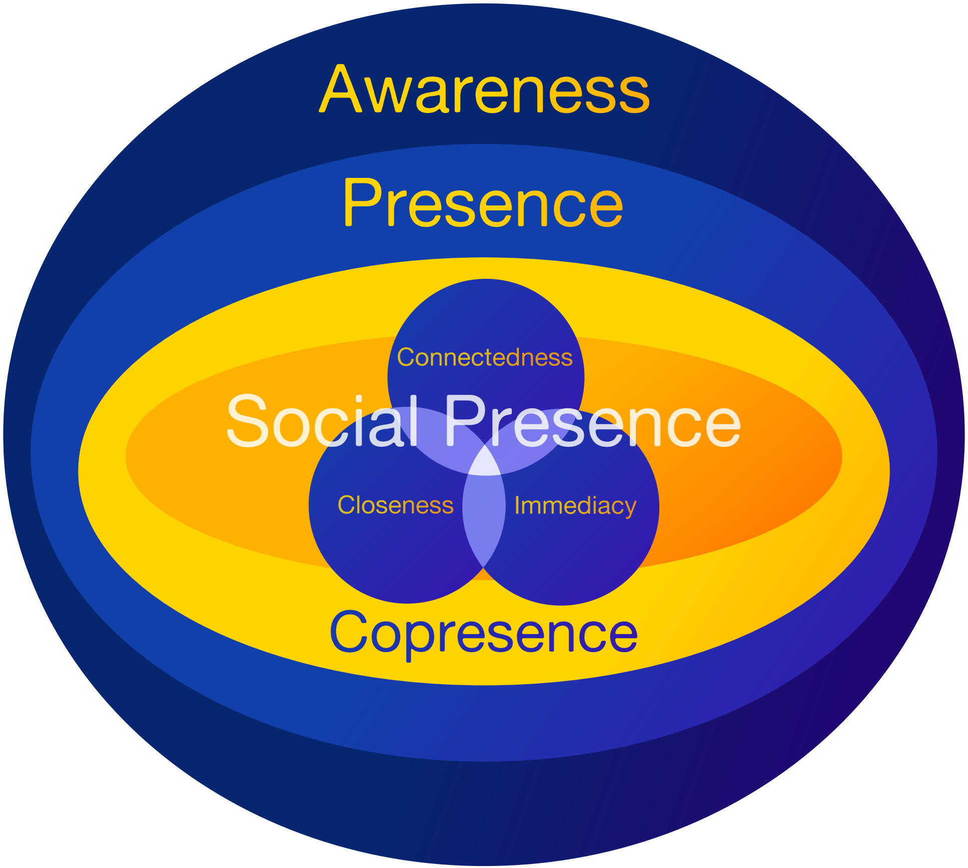

Figure 2.4. Proposition of a unified Euler diagram of all the concepts and terms presented here. The sizes of the sets and their intersection do not bear any information themselves.

The definitions above and also Figure 2.1 on page 16 might suggest that presence is a prerequisite for copresence and indeed this has been hypothesised [69, 455, 498, 537]. Yet of all the studies reviewed by Youngblut [585][10] only Slater et al. [498] have actually found a correlation between presence and copresence that was independent of the medium while Axelsson et al. [16, 17] found a correlation between presence and copresence that disappeared when a less immersive display was used. If confirmed, this could be good news for everyone trying to create copresence without wanting to recreate a maximally immersive experience such as shared virtual environments or highly realistic videoconferencing.

Similarly, Thie and van Wijk [530] were the only ones who reported on the relationship between presence and social presence according to Youngblut [585] although they did indeed find a significant relationship between the two. This is coherent with the association of immersion with social presence but a single study is of course not enough to draw any reliable conclusions.

Youngblut [585] conclude that current evidence is not sufficient to either assume nor dismiss any general relationship between presence, copresence and social presence.

IJsselsteijn et al. [258] argue that supporting immediacy and intimacy is especially important to promote social presence.

2.3 Awareness Systems

Driven by the findings of Kraut et al. [306] on the importance of informal communication and physical proximity, media spaces set out to mitigate the drawbacks of being physically remote and enhance collaboration over a distance and especially facilitate the opportunistic communications that often initiate collaborations in the first place. Although they were successful in enabling brief unplanned communication, they fell

[10] These are Axelsson et al. [16, 17], Basdogan et al. [28], Bystrom and Barfield [65, 66], Casanueva [69], Romano et al. [455], Schroeder et al. [482], Slater et al. [497, 498], Steed et al. [514], Whitelock et al. [572] and Wideström et al. [574].

Box 2: Speculations on Serendipity
Many authors emphasise the importance of serendipitous encounters – the benefits of just bumping into someone – yet they struggle to replicate these benefits with mediated communication. We would like to take a little detour to speculate on possible reasons for this.
We would argue that several factors might be responsible for the effectiveness of serendipitous interactions: these events' relative rarity, the (almost automatic) interruptibility of the persons involved and the availability of social cues to ensure this availability, and the implicit provision of context and memory aid. When you serendipitously meet someone in a corridor, you usually are not occupied with something else, you can use social cues to see if the other person is (e. g. when they appear particularly hurried or preoccupied), and the very fact of bumping into someone reminds you of points you wanted to address in relation to that person. These encounters are also relatively rare, so they automatically encourage interactions because they cannot be postponed without a certain insecurity as to when the next opportunity arises.
Whittaker et al. [573] tried to emulate serendipitous interactions by randomly connecting two people with each other. This violates the principles of interruptibility, though, and social cues are not available before the interaction starts. Similarly, media spaces, depending on what types of rooms are connected, do not provide interruptibility and artificial scarcity of events. In fact, what might work best conceptually is virtually connecting two corridors although technically this would be much harder to implement given screen size constraints and the difficulty of where to place the cameras.

short when it came to emulating the kind of opportunities that randomly "bumping into" someone create [573], a mode of interaction called *serendipitous interaction* [132, 246, 353] (cf. Box 2 above).

Awareness systems were then designed to surpass media spaces with regards to calmness and peripheral awareness and to tackle the lack of "informal awareness" between people [366].

SCS are even below awareness systems in the sense that they try to transmit emotion more than content. While working on an emotional level is also the aim of some awareness systems, they usually work on a level of content transmission that is still much nearer to face-to-face communication than SCS are. However, given that there were no real examples of SCS prior to this work, it might well be that the terms are going to merge in the future. Either way, we deem it safe to consider SCS one type of awareness systems for now. The only occasion where someone provided a definition of awareness systems that excluded SCS was Markopoulos [348] who said that a defining property of awareness systems is social translucence [135]: awareness of what the interaction partner is aware of. He calls systems lacking this social translucence *monitoring systems*. While it might be desirable for a particular SCS to have this property, we do not require it in general.

While Section 2.1 took a look at definitions of awareness systems and tried to pinpoint the difference between them and SCS, Section 3.2 lists some awareness systems built or conceived so far.

2.4 Measures of Presence and Connectedness

To evaluate and assess the efficacy of SCS and awareness systems, it is crucial to have reliable ways to measure the concepts we discussed in the previous sections. This is a particular challenge since many of the scales and methods developed over the years rely on spoken or at least deliberate communication. We will therefore take a look at a variety of scales used by other authors to measure connectedness and different forms of presence in particular, and then develop our own questionnaire, using as much of this prior work as possible to ensure a certain amount of consistency and comparability but also to profit from the considerable work and number of development cycles that went into some of these questionnaires. However, our own modifications, additions and recombinations are based on purely theoretical considerations which we will be put to a first practical test in Chapter 6.

Van Baren and IJsselsteijn [548], Youngblut and Huie [586], Youngblut [585], Schuemie et al. [483] and the *Measures Statement and Compendium* by the International Society for Presence Research [261] provide the most comprehensive overviews of presence measurements to date. Although they are all mostly concerned with virtual environments or conventional media usage, they are still a good basis to develop presence measures for SCS because of the large amount of research that has been done, often with many studies and a large number of participants for each measure.

Van Baren and IJsselsteijn [548] give a comprehensive overview of all kinds of subjective and objective presence measures, including a total of 24 questionnaires, qualitative measures such as the autoconfrontation method, ethnographic observations, focus group explorations, free-format self-reports and interaction analysis, psychophysical measures, and subjective corroborative measures of mental processes related to physical presence on the subjective side and psychophysiological measures (cardiovascular, skin conductance and skin temperature measurements, eye tracking, pupil response, and facial electromyography), electroencephalogram (EEG), functional magnetic resonance imaging (fMRI), behavioural measures (facial expressions and postural, reflex and social responses), and task performance measures on the objective side. Most of the objective measures are focusing on physical presence.

Although van Baren and IJsselsteijn describe the ability of each measure to discriminate between different levels of presence (called *sensitivity*), the exact definition of these different levels varies a lot and is often not well defined or not easily generalisable. The paper only discriminates between what is called "physical presence" and "social presence", with copresence being treated as a kind of social presence. Physical presence is defined as "the sense of being physically located somewhere" and is therefore a concept that is very much related to virtual reality environments.

A further criterion for tests alongside sensitivity is called *reliability* and describes the stability and consistency of a test over time. The most common metric for questionnaires is the internal consistency of scales which is almost always measured using either the average inter-item correlation or – even more commonly – Cronbach's α which is defined as follows [112, 311]:

$$\alpha(S) = \frac{k}{k-1}\left(1 - \frac{1}{\sigma_S^2}\sum_{i=1}^{k}\sigma_{S_i}^2\right) \tag{2.4}$$

with $\alpha(S)$ denoting Cronbach's α for a scale $S = \{S_1, \ldots, S_k\}$, $k = |S|$ being the number of items in this scale, σ_S^2 the variance between participants of the total test scores of all k items[11] and $\sigma_{S_i}^2$ the variance[12] of the ith item S_i. Every mentioning of an α value in the following will refer to this Cronbach's α unless otherwise noted. The parameter S is almost never written but was included here for the sake of clarity.

We would like to highlight a few of the many questionnaires presented by van Baren and IJsselsteijn [548] that were especially interesting for SCS in terms of sensitivity and the kinds of presence measured.

The Networked Minds Questionnaire [39, 40, 41] measures social presence defined as "the moment-by-moment awareness of the co-presence of another sentient being accompanied by a sense of engagement with the other. It is an outcome of cognitive simulations of the other's cognitive, emotional, and behavioural dispositions" [41].[13] It also identifies different factors called copresence (divided in isolation/inclusion and mutual awareness), psychological involvement (consisting of mutual attention, empathy, and mutual understanding), and behavioural engagement (behavioural interaction, mutual assistance, and dependent action) [548]. The standardised internal consistency ranged between $\alpha = .69$ and $\alpha = .87$ with an average of $\alpha = .77$

[11] This could be denoted as $\sigma_S^2 := s_n^2(X),\ X = \left\{\sum_{i=1}^{k} S_{i_p} : p = 1, \ldots, n\right\}$ if the response of participant p to item S_i is called S_{i_p} and $s_n^2(X)$ is the sample variance of X.

[12] Using the above notation this would be $\sigma_{S_i}^2 := s_n^2(Y),\ Y = \left\{S_{i_p} : p = 1, \ldots, n\right\}$.

[13] Despite the mixing up of social presence and copresence in this definition, the subsequent additions make it still conform broadly with our own definition of social presence, merely at the same time claiming that copresence is a prerequisite of social presence and the amount of copresence is positively correlated with social presence.

The questionnaire by Nowak and Biocca [401] measures telepresence, copresence and social presence by combining previous questionnaires by Short et al. [494], Burgoon and Hale [63] and Lombard et al. [338]. The internal consistency scores of the different scales ranged from $\alpha = .78$ to $\alpha = .90$ with an average of $\alpha = .85$. Some items of the questionnaire were eliminated after a first round of consistency measurements.

Basdogan et al. [28] also measured "the sense of being together" (and did so in the context of touch which is interesting for our FEELABUZZ system detailed in Chapter 4) but did not re-use any existing scales nor provided any measure of internal consistency for their questions.

The survey paper by Youngblut [585] also focuses on virtual environments but is still more useful for SCS since it divides presence into place presence (roughly what van Baren and IJsselsteijn [548] call "physical presence"), copresence and social presence.

Youngblut makes out a distinct lack in the number of copresence questionnaires (and a total lack of other methods to measure copresence) but lists nine of them[14] and highlights the one by Casanueva [69] as the most comprehensive and as the only one to report a reliability score (Cronbach's α of .79). He draws a lot on Slater et al. [498] in his questionnaire.

In the table provided by Youngblut [585], Schroeder et al. [482] also stand out due to the number of studies conducted; they measure presence and copresence among other things.

Concerning the measurement of social presence, Youngblut [585] observes that all questionnaires are derived from Short et al. [494] (who present some direct questions as well as a set of items using the semantic differential technique developed by Osgood et al. [409]) with the exception of Bailenson et al. [20]. For the latter – while having an internal consistency value of $\alpha = .83$ – concerns about the face validity are raised, i. e. whether the questionnaire actually measures social

[14]described in these works: Axelsson et al. [16, 17], Basdogan et al. [28], Bystrom and Barfield [65, 66], Casanueva [69], Romano et al. [455], Schroeder et al. [482], Slater et al. [497, 498], Steed et al. [514], Whitelock et al. [572] and Wideström et al. [574]

presence. We share this concern on the basis of questions regarding awareness and being in the same room which resemble very much questions that are otherwise used to measure copresence.

Regarding the fundamental choice between self-reports and objective measures, Biocca and Harms [39] argue that self-reports are more meaningful than methods that do not have access to introspection. Other authors disagree, however, and think of objective measures as a reasonable or even necessary addition to questionnaires [256, 548, 585].

2.4.1 Questionnaire Design for Subliminal Copresence Systems

For the studies conducted for our systems FEELABUZZ (cf. Chapter 4) and upstairs (cf. Chapter 5) that we present in Chapter 6, we designed a custom questionnaire based on several of the questionnaires mentioned above [13, 39, 69, 257, 401, 494, 498, 550, 584]. The studies were conducted with native German speakers so the questionnaire was written in German, too. We used or adapted existing and validated translations where available. In particular, we used the adaptation of the Networked Minds Questionnaire by Biocca and Harms [39] by Blicker-Dielmann [43], the translations done by Doll et al. [118] and partially by Asendorpf et al. [15] of the attachment styles and relationship questionnaire (RQ) by Bartholomew and Horowitz [26] which was in turn based on Hazan and Shaver [217], and the German social desirability scale by Stöber [518].

Other questionnaires used and translated by ourselves were the Affective Benefits and Costs of Communication Technologies Questionnaire (Adult ABCCT) by Yarosh [584] which uses many sources but is mainly based on the very similarly named Affective Benefits and Costs of Communication Questionnaire (ABC-Q) by IJsselsteijn et al. [257] and the social connectedness questionnaire derived from it [550]. The ABC-Q tests for affective benefits, including “a feeling of having company, a stronger group attraction, a feeling of staying in touch, of keeping up-to-date with other people’s lives, and a sense of sharing, belonging, and intimacy", while the four cost scales are called *obligations*, *expectations*, *threats to privacy* and *process effort* [257, 258]. IJsselsteijn et al. [257] identify *emotional involvement*, *thinking about each other*, and *staying in touch* as measures of connectedness and address them with appropriate scales.

Both the Networked Minds Questionnaire and the ABC-Q (after its first extension) and its derivatives make use of reciprocity in their questions, i. e. extending self-assessment questions with analogue questions about the interaction partner.

The ABC-Q contains an effort scale which is, in principle, very interesting with its approach of differentiating between effort that the interaction partner appreciates (*personal effort*) and that which the other person is ignorant of (*process effort*). We felt, though, that the manner of asking for these kinds of effort does not translate well to SCS because they are supposed to primarily use input that is not explicit but the secondary result of a primary task (cf. Section 3.1.2). This is especially the case for the personal effort. Two items from the process effort subscales could be adapted into our new cognitive load subscale that we added to the costs part of the ABCCT.

Paas et al. [413] showed that self-reports of perceived mental effort are a valid measure of cognitive load and DeLeeuw and Mayer [108] in principle support this, although they hold a somewhat more complicated view: they provided experimental support for the older idea that there are different kinds of cognitive load for which different measures are adequate. Lewis [329] states that preference to subjective measures should be given preference when user satisfaction is the primary goal while objective measures are more important where increased performance is paramount. Rating scales to assess cognitive load are usually built after Paas et al. [413] and take the form of a difficulty rating of any given task [551] although in HCI measures such as ease-of-use and likability have been used [6, 329].

When reusing scales and questions, social presence was the most problematic measure. Maybe unsurprisingly, most of the scales and questions to measure social presence rely on the existence of a full-fledged dialogue or other kind of social interaction that SCS cannot provide. Therefore we had to leave out most items to the point that the face validity of the social presence measure can no longer be undoubtedly assumed.

The specific difficulties when trying to apply existing scales from awareness systems and presence and connectedness research to SCS can be seen as another indication that SCS were not a class of systems that most past researchers had in mind, despite the surface-level similarities between awareness system and SCS definitions.

2.5 Long-Distance Relationships

We believe a kind of relationship that is especially suitable to examine the effects of subliminal copresence systems are long-distance relationships (LDRs) and – as a special subclass – so called long-distance romantic relationships (LDRRs).[15] This is because of their almost total dependence on mediated communication as opposed to proximal romantic relationships[16] (PRRs) and the intensity and long-term nature of the relationship compared to short-term lab experiments or studies with less intense kinds of interpersonal relationships.

Conversely, we believe SCS might be of special interest to LDRs and LDRRs as well, for the reasons already mentioned but also because continuous telepresence provides means to decouple variables that are otherwise tied to physical presence and therefore allows researchers to test hypotheses that are otherwise difficult to examine. What this means in detail will be further explained in the following sections.

LDRs and especially LDRRs have become a topic of intense research in the last two decades (cf. the preface to Stafford [507]) which is why

[15] LDRRs are sometimes further subclassed in the literature. As college students appear to be a frequent object of studies (and maybe a particularly bad one according to Guldner and Swensen [201]), long-distance dating relationships (LDDRs) is a term one will encounter quite often. LDDRs are contrasted to adult romantic relationships or adult romantic LDRs, with further subtyping into long-distance marriages and dual career, dual residence (DCDR) types of relationships.

[16] i. e. romantic relationships that involve the availability of physical presence most of the time

a look at the results of this research can inform the development and exploration of SCS. Although we are going to focus on LDRRs in this work (because of the large body of research to build on), we think SCS might be of equal relevance for any strong emotional bond. As Fehr [141] summarises, while romantic partners often have the "deepest, closest, most intimate" relationships, this is not always the case and the bonds people form with friends and family can also be very significant.

2.5.1 Overview of Research on Long-Distance Romantic Relationships

One particularly interesting aspect of LDRRs that may defy first intuition is that very few differences were found between LDRRs and PRRs[17] in terms of various relationship satisfaction measures [449, 466, 467]. Relationship stability was repeatedly found to be even higher in LDRRs than in PRRs [442, 508, 515] although it has to be said that elevated levels of depression, detachment and resentment are also well documented in LDRRs [200, 557].

Mashek and Sherman [361] write generally about the desire for less closeness in intimate relationships and Le et al. [318] found the experience of missing a partner in an LDRR to be only weakly correlated with loneliness but strongly with commitment and attachment. Knox et al. [292] received mixed reports after five months of separation, with the 80 % of surviving relationships[18] being equally often described as having improved as they were as having worsened from the separation. College students in LDRRs were also found to be better rested and perform better than those in PRRs [507] and working LDRR couples were more satisfied with their work [62]. People generally seem to focus more on school and career when separated and more on their relationship when temporarily reunited and seem to appreciate this duality [173, 507].

[17] Geographically close dating relationships (GCDRs) are the equivalent counterpart to LDDRs.

[18] As compared to 100 % of LDRRs and 74 % of PRRs still being intact after six months in a study by Reske and Stafford [442].

A certain amount of such temporary reunions has also been shown to be positively correlated with trust among the partners; too much face-to-face contact in LDRRs, however, seems to decrease satisfaction and the chance of relationship survival [510].

Indeed, permanent reunion (i. e. an LDRR becoming a PRR again) ended a third of relationships within three months in a study by Stafford et al. [509]. Reasons reported by the participants were the loss of autonomy, including increased control and monitoring (one participant was quoted: "[now that we live in the same city] I have to check everything I do with her."), time rescheduling difficulties, increased awareness of negative character traits (although the opposite happened as well), presumably because of the previous maintenance of an idealised image of the partner and the facilitated perpetuation of a positive self-presentation [508, 509],[19] and – surprising to us – increased jealousy. As much as 85 % of participants reported to have missed at least one aspect of the LDRR after their reunion. Stafford et al. grouped these aspects in five categories they called closeness via distance, quality time, anticipation/novelty, autonomy, and time management ease. The first describes the fact that the spatial distance made some couples feel closer emotionally. The next two categories sum up comments that indicated an increased incentive to make good use of the face-to-face time that was given to the couples and a heightened relishing and anticipation of these moments. The last two categories describe the increased difficulties to find time for oneself and one's own interests, friends and family, and for each other.

On the positive aspects of PRRs, Stafford et al. found primarily the increased time spent together, quoting one participant: "We finally got to do all the 'little' things we'd been wanting to do for so long; we get

[19]Although it has to be said that idealisation of the partner is normal in a relationship [559] and is not necessarily a bad thing either [227, 378, 464]. Stafford [507] argues, however, that the increase in idealisation in an LDRR is bad even though it promotes satisfaction and longevity of the relationship – or precisely because of this. She argues that the delusional views the partners uphold of each other and the lack of information on the partner's bad traits could simply defer the development and eventual termination of a relationship and should in this case not be considered a good thing [98].

to hold each other, wake up next to each other, eat together, etc."

As Roberts and Pistole [449] point out, intimacy and commitment were not found to be any different between LDRRs and PRRs [201, 552]. However, the higher relationship stability of LDRRs that was found at least for college students [507] could in turn lead to a higher relational satisfaction [484]. Whether satisfaction itself was higher in LDRRs or PRRs varied from study to study [201, 507, 552] but the more recent work by Roberts and Pistole [449] looked at this question much more scrutinously. The increased freedom and autonomy in LDRRs that counterbalances the lack of physical closeness had previously been one of the main reasons brought up to explain the satisfaction with LDRRs [467]. Roberts and Pistole [449] presented a more differentiated model based on character traits.

Existing studies on the effect of different media on LDRRs have naturally focused on telephone, letters and written communication over the Internet (mainly e-mail and instant messaging). Dainton and Aylor [103] found that the telephone was positively correlated with relational commitment and satisfaction, and the use of written communication over the Internet had a positive association with trust. Stafford and Reske [510] found letters to be more highly correlated with feelings of satisfaction and love than telephone or face-to-face contact; the Internet was not an important factor at the time of that particular study (1990). A similar effect was found by Guldner [200] who reported the chance for long-distance couples of staying together to be higher when they exchanged letters.

Gerstel and Gross [173] found varying results for the use of telephone by married commuters whose reports ranged from an increase in connectedness to an increase in loneliness, whereas letters are seen more uniformly positively [511]. Stafford [507] attributes this to the non-real-time nature of letters which allows to choose tone and content more carefully. On the other hand, this does not seem to translate as easily to comparable forms of communication on the Internet [459], even though the opposite had been suggested [428].

Jacobs and Hicks [267] found that soldiers during deployment reported the possibility of regular communications with their partners at home to be important with regards to a feeling of intimacy. Interestingly, it was found to be the case that the availability of communication possibilities made the separation less difficult, regardless of whether these possibilities were used or not [454].

Much research has been done and is still ongoing on how to measure relationship quality and even on what constitutes a healthy relationship [507]. The overview of Roberts and Pistole [449] on methods to measure relationship quality uses the Relationship Questionnaire (RQ) by Bartholomew and Horowitz [26], the Relationship Closeness Inventory (RCI) by Berscheid et al. [34], the Inclusion of the Other in the Self scale (IOS) by Aron et al. [12, 13] which also measures closeness, and the Dyadic Satisfaction subscale of the Dyadic Adjustment Scale (DAS) by Spanier [506].

2.5.2 Potential Effects of SCS on Long-Distance Relationships

The results of the research presented above suggests some influence SCS might have on LDRRs on various levels. On the one hand, they could support them by increasing presence (presumably copresence to be precise), either overall or in-between other forms of communication. They could also simply provide one of the "creative ways of communicating while at a distance" that Westefeld and Liddell [571] suggested to develop in order to help maintain an LDRR or even open up new maintenance strategies usually related to physical proximity that were impossible to study for LDRRs before. A mitigation of the increase in depression, detachment and resentment found in some LDRRs would of course be ideal.

On the other hand, long-term communication as provided by SCS could have a negative effect on the formation and maintenance of the positive illusions and idealisation that Stafford [507] proposed as the main mechanism behind the deleterious effects of face-to-face interaction and permanent reunion on relationships. This possibility exists because the maintenance of illusive self-presentation should become more difficult over prolonged periods of time as opposed to the very short-lived episodes of high-level communication. It is of course not clear to what extent the systems presented in this work can have a detrimental effect on positive self-presentation at all. Plausible deniability might be reduced under certain circumstances though (upstairs might reveal actual household activity to name just one example).

It would also be interesting to see if any effect of SCS on LDRRs correlates with the chance of a given relationship to survive permanent reunion. This hypothesised effect of SCS might go so far as to strengthen relationships that would survive a reunion while having a detrimental effect on those that do not (making the measurements of relationship quality changes even more challenging).

3 A Design Framework for Subliminal Copresence Systems

3.1 Awareness Systems Taxonomy

In Chapters 1 and 2, we introduced the concept of subliminal copresence systems (SCS), looked at important concepts and measures used in related fields and defined what we mean by SCS through the goals these systems are supposed to achieve (Definition 8). In this chapter, we will take a look at the design decisions and derived goals and concepts that follow from these goals, possible channels SCS can exploit, and similar systems built so far that might inspire future SCS. While we will take a closer look at systems that are similar to those we already implemented and are going to present in Chapters 4 and 5, we intend this chapter not only to bolster and motivate our own decisions but to be a helpful reference for anybody intending to explore this new field with their own systems and ideas.

While we do not want to deny the possibility that SCS might be useful in a work-related setting, conditions and goals are very different in work compared to home environments [169, 392, 447], so the following design considerations do not focus on collaborative or task-centred applications. Specialised awareness-related design frameworks for computer-supported cooperative work (CSCW) exist [203, 429]. Greenberg et al. [191] give a good overview of some particularities of ambient displays for the home environment and a breakdown of the three different social groups that one encounters there (called *home* *inhabitants*, *intimate socials* and *extended socials*) and their characteristics and communicative needs.

3.1.1 Related Taxonomies and Design Guidelines

As a basis for our taxonomy of SCS, we will take a look at taxonomies in related fields. Maybe one of the more obvious of those after the discussions of Chapter 2 is the taxonomy of copresence proposed by Zhao [591]. It proposes four basic categories for copresence systems that are shown in Table 2.1 on page 23. They are of limited use for this work, however, since all discussed systems here would fall in the *corporeal telecopresence* category. More relevant are the properties Zhao calls "interface parameters". These are *embodiment*, *immediacy*, *scale* and *mobility*. We adopted immediacy and scale below, while embodiment is implicit for corporeal telecopresence and mobility is discussed under *group configuration*.

Mobility and immediacy are concepts that were already brought up under the same names in the telecommunications taxonomy by Cleevely and Cawdell [82] which we do not further discuss here because it focuses on technical aspects while at the same time predating the World Wide Web. Tsui et al. [543] use only two axes, *interactivity* and *personalness*.

In the related field of ubiquitous computing, Scholtz and Consolvo [480] reviewed different taxonomies that had varying focus areas [32, 158, 269] and aggregated them into a taxonomy of their own. Of particular note is the approach by Friedman et al. [158], who define the ethical dimensions human welfare, ownership and property, freedom from bias, privacy, universal usability, trust, autonomy, informed consent, accountability, identity, calmness, and environmental sustainability [480]. We have a comparatively long discussion of calmness below and also mention privacy. Universal usability (accessibility) is of particular import for unusual modalities and hardware. While, in principle, new systems can be good or bad news depending on particular disabilities of an individual, it is important to keep in mind that there has been considerable development time and market pressure for standard computer

Desirable aspects		
Satisfying	Helpful	Fun
Enjoyable	Motivating	Provocative
Engaging	Challenging	Surprising
Pleasurable	Enhancing sociability	Rewarding
Exciting	Supporting creativity	Emotionally fulfilling
Entertaining	Cognitively stimulating	
Undesirable aspects		
Boring	Unpleasant	Frustrating
Patronising	Making one feel guilty	Making one feel stupid
Annoying	Cutesy	Childish
Gimmicky		

Table 3.1. Goals as to what a user experience should and should not be like according to Rogers et al. [452]. The authors themselves state that this list is not exhaustive. It also seems subject to cultural and individual differences in the perception of certain aspects as desirable or undesirable.[a]

[a]See for example the discussion of cuteness by Kinsella [286].

usage which newly developed systems tend to lack.

The framework developed by Scholtz and Consolvo [480] called *ubiquitous computing evaluation areas* (UEAs) is hierarchical in nature and quite extensive. For example, *trust* is a UEA containing the metrics *privacy*, *awareness* and *control*. The UEAs range from technical properties of a specific implementation such as application robustness and aesthetics, over social properties such as cost and status gain to more conceptional properties such as attentional overhead. On the other hand, the detailedness of their framework reduces the transferability out of the area of HCI on to awareness systems and SCS.

Type of Awareness	Recommended Means	Expected Benefit
Group awareness	Provision of information for task completion	Establishment of transactive memory system facilitating collaboration and increasing group satisfaction (particularism)
	Provision of simultaneous source of information	Symmetry of awareness
Social awareness	Visualisation through diagrams or symbols (2D, 3D)	Provision of information about one's own appearance enabling feedback on the effect of the appearance on other group members
	Context-sensitive feedback mechanism	Reflection on communication rules through feeding back on one's own or other experience from another group member Establishment of sense for community (informal awareness)
	Mutual sensing of presence	Respect of other's spheres or interaction spaces (peripheral awareness)
	Motivation for and control of engagement	Each member is heard and can be brought into the group
Task-specific awareness	Visualisation: bar charts, flow diagrams	Mutual transparency of working process leading to increased understanding of group members
Situation awareness	Extrapolation tools	Provision of explicit projections into the future of collaborative activities leading to increased understanding of individual behavior of group members and allowing sound revision of individual expectations
	Feedback provision at an informal level	Improving the sense of community
	Presentation of group member's location	Improving the spatial feeling of community (location awareness)
Objective self-awareness	Mirroring video camera	Increased awareness of oneself leading to accurate support of self-reflection

Table 3.2. Kinds of awareness by Gross et al. [197].

In 1994, Grudin [198] proposed his "eight challenges for groupware developers". Among these, *disparity of work and benefit* and *disruption of social processes* seem especially relevant to SCS.

Gross et al. [197] classified awareness-related CSCW systems from a social sciences point of view. They tried to consolidate awareness research in CSCW and results from small group research and found five types of awareness that presented common themes: group awareness, social awareness, task-specific awareness, situation awareness and objective self-awareness (cf. Table 3.2 on page 54). They suggested that social awareness – which includes a "mutual sensing of presence" – and situation awareness provide especially fertile and versatile opportunities to support them technologically.

Markopoulos [348] developed a design framework for awareness systems in which they motivate awareness systems using the basic *social needs* from social psychology, citing Kenrick et al. [285]. They discuss how awareness systems can help people to address these needs, which Markopoulos presented as

– to affiliate with others,
– to obtain social information and present oneself,
– to attract and retain mates, and
– to protect ourselves and those we value.

While the works of Markopoulos and the others who take a broader perspective on the topic are a good reminder that a deeper sociological or psychological understanding of the needs one wants to address can often be helpful and we would therefore recommend reading them to anyone interested in such matters, we do not think that a design framework should necessarily try to list all needs that might be addressed by such a class of systems or differentiate between *basic* social needs and those that are specific to a certain culture or other social group.

The work by Markopoulos [348] also includes an excellent discussion of design dimensions and challenges for awareness systems. He says himself that he focuses on important design decisions instead of a classification of awareness systems for which he referred to Gross et al. [197]. He defines his design space for awareness systems along the following seven dimensions.

- *Awareness of place versus awareness of people* covers the context of communication and raises questions we discuss with our dimensions of group configuration, scale and privacy.
- *Precision* is the information content of a channel which can be deliberately reduced through degradation (e. g. blurring or distorting) or abstraction to enhance privacy or – especially in the case of abstraction – reduce cognitive load.
- *Accuracy* is the reliability of the information provided by a channel which should be balanced against the often opposing goal of precision. It can again be used deliberately to provide privacy and plausible deniability but too much user control over the reliability of a channel's data can severely endanger the trust users put in the system.
- The *notification level* was introduced as *interruption dimension* by McCrickard et al. [365] and further developed by other authors [362, 430]. This concept describes a classification of different levels at which a system can grab the attention of a user when information changes, ranging from active polling by the user to interruptions that cannot be ignored. Since SCS generally transmit information continuously, these categories are less relevant to them than other aspects of calmness but we will come back to these distinctions in Chapter 7.
- *Input automation* means the distinction between explicit versus implicit input. We discuss explicit versus implicit communication in Section 3.1.3. Markopoulos points out that implicit input can be error-prone when it involves automatic pattern recognition and does not convey intentionality. Directness as discussed below can help with both of these concerns.
- The *private versus shared nature of awareness displays* considerations focus more on the factual than on the intentional side of input and output devices than the earlier aspect of place versus people but touches upon similar areas and finally
- the *level of user control* over information flows is discussed and classified in four gradations, which are 1) switching all streams of a system on and off as a whole or 2) individually (publish/subscribe),

3) filtering information streams or 4) providing even more flexible control. Markopoulos raises the point that more flexibility also means a much greater overhead for the users. Again, this is a topic that we will further discuss in Chapter 7.

From these design dimensions, Markopoulos derives seven challenges for designers of awareness systems. These are:

1. minimising the procedural effort (which refers to the overhead of operating the *system* as opposed to the effort of communication itself[1]),
2. supporting transitions to other media,
3. designing agency in awareness systems (meaning that a sense of human agency and intentionality should be preserved instead of too much automation),
4. upholding reciprocity and an equity of costs and benefits of the communication between the participants,
5. balancing accountability and autonomy (under this topic, Markopoulos discusses the need to provide plausible deniability and possibilities to deceive a system but balance them with the necessary trust in the system),
6. seamful design [71] (designing with useful and usable transition between media in mind, instead of trying to cover them up), and
7. data proportionality [251, 252, 349], i. e. being reasonable with regard to the privacy intrusions introduced.

In the overarching field of interaction design, Rogers et al. [452] define a range of "user experience goals" in the form of desirable and undesirable aspects. These are listed in Table 3.1 on page 53. They have a complementary list of items they call "design principles" which are 1) the *visibility* of functions, 2) *feedback* to user actions, 3) reasonable *constraints* as a means to unclutter an interface and to avoid bad input, 4) *consistency* of the interface, and 5) usage of *affordances* (as popularised by Norman [398, 399]).

[1] Markopoulos actually thinks of the personal effort of communication as a positive thing as he says that is generally appreciated by the recipient. Clearly we pursue a different strategy with SCS, though, where any effort should be minimised and not only the procedural effort.

McEwan and Greenberg [366] define eleven design principles for their Community Bar system that was intended to facilitate and spur informal, spontaneous and serendipitous interaction between colleagues. For this system, McEwan and Greenberg drew on the Locales Framework by Greenberg et al. [190] (see Box 3 on the next page) and the focus/nimbus model of awareness by Rodden [450] shown in Figure 2.2 on page 27 in Section 2.1 as a theoretical basis. The design principles that came out of this research were:

1. Information should be permanently visible at the periphery and it should be kept up-to-date. It should not interfere with other tasks.
2. Interaction in small groups should be supported.
3. Switching from the peripheral awareness of people to interacting with them should be easy and quick ("lightweight").
4. There should be rich information sources and channels to provide many different cues.
5. The "informal awareness and interaction tools" should support the different social contexts (or locales, see Box 3 on the next page) between which people switch constantly and
6. they should provide ways to organise and relate these locales to one another.
7. The tools should allow customisation and
8. allow to manage the interactions as they evolve over time.
9. The users should be able to change their focus (cf. Figure 2.2 on page 27) and
10. control their nimbus (ibid.).
11. A change in the focus/nimbus combination should be represented in the information display and interaction affordances.

Of course, goal 1 is at the centre of our own work but it is important to note that much further elaboration is needed on how to pursue this goal. In fact, the things we can process subconsciously are very limited [334], they vary from person to person [330] and the mere expectation that we might be interrupted by something from the periphery already has the potential to decrease someone's ability to concentrate [266].

Many of the other items are either very generic and apply to HCI in

Box 3: The Locales Framework
The Locales Framework was developed by Greenberg et al. [190] and forms a quite comprehensive framework of awareness and context. Rittenbruch and McEwan [447] describe it as "revolutionary" and "unique in its coverage". They also acknowledge its limited impact in the field of CSCW, though, and try to explain this by its complexity and its "descriptive rather than prescriptive nature".
The Locales Framework consists of social worlds that describe different social contexts that people organise themselves in. This is related to social roles [181] and in fact *aspects* in Diaspora [248] or the better known *circles* in Google+ [536] could even be described as an implementation of these social worlds. There is one important difference, though, which is that social worlds are fuzzy so that people are not either in a certain social world or not but can rather be closer to the centre where they are more involved or they can be at the periphery when they skip group activities and just do not participate as much in that particular social world.
Social worlds are connected to *sites* and *means* in an *n*-to-*m* fashion by *locales*. Sites are places – either real or virtual – where the social worlds perform their activities and means are the tools (or artifacts in a wider sense) used to do so. All components are modelled as being dynamical and their changes over time are called *trajectories*.

general or they are rather specific to the Community Bar system. What McEwan and Greenberg call richness in item 4 resonates with SCS though and we explore this a bit further under the keywords of *signal resemblance* and *directness*. We discuss the lightweight switching to conversation under *complementarity*.

We do not see focus and nimbus so tightly coupled when it comes to design aspects and subsume focus issues under *calmness* while control over one's nimbus is discussed under the keyword of *privacy*.

For ambient displays, de Guzman et al. [105] mark out the three design objectives *aesthetic appeal*, *ease of interpretation* and *minimal distraction*.

Important Aspects of Subliminal Copresence Systems	
Calmness	Complementarity
Directness	Signal Resemblance
Privacy	Plausible Deniability
Information Capacity	Relevance
Affectiveness	Scale
Immediacy	Group Configuration
Trustworthiness	

Table 3.3. Tabular overview of all the concepts defined in this chapter.

Pousman and Stasko [430] assign values to a number of systems for the four dimensions *information capacity*, *notification level*, *representational fidelity* and *aesthetic emphasis* and mark out four categories of systems. They call these four ambient display design archetypes *symbolic sculptural displays*, *multiple-information consolidators*, *information monitor displays* and *high-throughput textual displays*. The high-throughput textual displays (for which a web-based feed aggregator is named as an example) and the multiple-information consolidators (with the Mac OS X Dashboard[2] being an example here) might not meet our definition of calmness as described below, however.

3.1.2 Definitions and Terminology

We will in the following present, define and discuss concepts that we deem of particular import to the area of subliminal copresence systems. See Table 3.3 for an overview.

Calmness

Weiser and Brown [565, 566] coined the term *calm technology* back

[2]The Dashboard is an easily accessible desktop overlay that displays a number of user-configurable apps.

in 1995 in the context of the then-dawning area of ubiquitous and pervasive computing. They define calmness through what they call the *periphery*, which is a term to denote the things people can process subconsciously and effortlessly, without drawing our attention when not needed.[3] According to Weiser and Brown, it is also a property of calm technology to ease the transition between peripheral processing and back.

The idea is very much entrenched in the area of pervasive computing, even when often called differently. Saha and Mukherjee [465] for example speak of applications that should “disappear into the environment” and of “invisibility” as a desired property, as do Scholtz and Consolvo [480]. Scholtz and Consolvo also notably name customisability an aspect of invisibility.

Awareness systems share a lot of aims with pervasive computing and ambient displays in particular. It is interesting to note, however, that the awareness systems literature only rarely refers back to Weiser.[4] It is instead mostly understood that the term *awareness* or *peripheral awareness*, depending on the definition, implies calmness (cf. Section 2.1).

Therefore no definition of calmness that is specific to communication systems could be found and we describe calmness in the context of this work to be more specifically relevant for SCS.

Definition 9: A system is **calm** to the extent that it does not intrude into the daily lives of users through unwanted distraction or

[3]Weiser [564] attributes the term “periphery” to his coauthor in Weiser and Brown [566] and equates it with the terms “compiling”, “tacit dimension”, “visual invariants”, “horizon”, and “ready-to-hand” by other authors.

[4]It is in this context interesting to note that Weiser and Brown [566] themselves mentioned video conferencing as an example of a technology that was *more* calm than the traditional telephone conference. This was asserted on the basis that an additional video signal provides more cues that can be processed peripherally, namely nonverbal communication cues. Clearly though, this is not the concept of “calm” communication that this thesis aims to explore or awareness systems research for that matter.

the introduction of obtrusive hardware, cognitive or temporal burden or unfamiliar concepts.

This means that a calm system should be unobtrusive where its output is concerned and the input should happen implicitly, during actions users perform independently of the system.[5] Where possible, the output of a system output should be processable subconsciously by the users without requiring a long familiarisation period. An already familiar metaphor such as a simulated or transmitted real-world phenomenon is therefore preferable. From the definition also follows that the number of artifacts added to the environment should be minimised. Using already existing artifacts not only decreases the burden of additional visible hardware and takes away the need to *learn* to interact with new objects, it also takes away the need to interact with more objects than before in the first place and lets the system make use of already existing interaction patterns which means the communication is implicit rather than explicit and one cannot simply forget to use a new communication device or refrain from using it because it is too cumbersome to integrate it into one's routines. Hindus et al. [235] found with their Presence Light, however, that intentional presence indicators can be more desirable than implicit ones (cf. Section 3.2).

In light of the fact that no one communication system is ideal for every communication requirement (cf. *complementarity* below), changing communication media and systems should be fluent and effortless.

Aesthetics can also form an important part of calmness. Although often seen as a goal in itself, aesthetics also are functionally important for systems that are continually present in the vicinity of the users and whose signals are also meant to be perceived incessantly [430]. This means that those hardware parts of a system that are visible should look nice in that particular environment and that whatever

[5] Even though this in no way rules out the possibility to use the system for intentional communication as well (cf. Section 3.1.3).

output modalities are used, the stimuli should be pleasing to the respective senses.

The doctrine of calmness follows from the intent to create *subliminal* systems, i. e. those that do not draw attention to themselves and their mode of communication except under very special circumstances that indicate that special attention to the interaction partner is indeed required.

Complementarity
This describes the notion that calmness is not always desirable in systems. Weiser and Brown also expressed this in their original work on calm technology [566] and mention the example of a "calm videogame" which they say would be pointless. For awareness systems, we believe something similar to be true. IJsselsteijn et al. [258] put it like this: "awareness systems are not seen as replacing existing communication means, but rather as enriching them, strengthening existing social bonds and enabling new kinds of interactions."

What is more, we are convinced that using a subliminal communication channel exclusively for a prolonged period of time without complementing it at least from time to time by a channel that is less calm but richer in information and that is potentially language-dominated (most likely telephone calls, video conferencing or real-world encounters; but maybe even IM or microblogging [293]), incurs significant drawbacks. This is the result of an extensive discussion we[6] had at Aalto University's Media Lab in Helsinki in 2011 and that is backed by the proven role especially the telephone plays in the maintenance of long-distance relationships (LDRs) [103]. To substantiate, consider our application called *upstairs* (cf. Chapter 5) that records footsteps and other structure-borne sound from the floor of one room and transmits them to the ceiling of another room to make it sound as if both rooms were atop one another. As we had already conjectured earlier, we believe that without an intrinsic interest in

[6] my colleagues Till Bovermann, René Tünnermann, and myself

the other person, this would just be perceived as annoying rather than desirable.

Definition 10: A high **complementarity** in an SCS means that the system is able to make good use of context provided by high-level channels or that conversely the system is able to provide context for such higher-level channels.

As a shorthand, we will call the channels that are less calm, more information-dense, and that tend to be short-lived and often language-dominated **high-level channels** and the channels provided by SCS conversely **low-level channels**. Because they are more dense in information, the former most likely are more apt at providing the *context* that is crucial for any communication system [130] and that we deem important for the lower level channels and other key concepts such as privacy. Note that the terms high-level and low-level relate to the concept of richness in media richness theory [100, 101]. However, media richness tried to pack together even more different aspects of media in one measure (richness meant the ability to convey natural language instead of just numeric information, a greater ability to convey nonverbal channels, the ability to personalise a message, and less lag [110]) and after it seemed not to live up to its promises [110], Dennis et al. [109, 110] proposed a more fine-grained alternative. However, we think that while such a coarse term like richness or high-level channels does not do justice to the complex interrelationships between the various properties of awareness systems and media in general mentioned above, it can still be useful to use such broader terms to convey certain ideas as long as one is always weary of the shortcomings of these terms.

As already mentioned under *calmness*, from the notion of complementarity also follows that escalating to higher-level channels should be easy and dropping in and out of such additional streams should be quick and seamless.[7]

[7]Something that was shown to be true at least in certain contexts [306, 573] and

Directness
Signal Resemblance

These two related terms describe the idea that we are used to very rich interactions and the equally rich stimuli these interactions produce, often for more than one of our senses at the same time. By refraining from automatic interpretation of the signals produced by various activities, the richness of the signal is not lost and users can use their naturally acquired repertoire of interpretation abilities as well as *context* [123]. This concept is also tightly linked to *calmness*, as stimuli that people are used to are more easily interpreted subconsciously (in the *periphery*) and might be perceived as less alien. It should be mentioned, however, that Hudson and Smith [246] argued that a higher information content also means a greater interruption and thus less calmness, a design principle that was also used in the AROMA project [422, 423] which was based on abstraction and symbolic representation. We would count information content only among one of the factors at play here, however, along with the familiarity of the stimulus facilitating its cognitive processing and the potential annoyance an information-poor and thus highly repetitive stimulus might bring with it.[8]

Definition 11: Directness is the extent to which a communication system preserves *all* features the signal produced by the sender originally had – not only the features the system designer had in mind – and the extent to which the receiver can interpret these features using the same capabilities that could interpret the original signal. **Signal resemblance** additionally describes the extent to which an output signal feels familiar with regard to the original signal.

This distinction between directness and signal resemblance might seem subtle but take a visual link as an example. Depending on

which Weiser and Brown [566] also believed to be generally true.

[8]Privacy and plausible deniability (see below) surely benefit from abstraction and the stripping away of information, though.

the transmitted images, an inverted or colour-shifted signal might be as easy or even easier to interpret than the original signal and thus the transformation does not affect directness (unless the image contents can possibly change to something where the transformation does interfere, such as faces) but it does affect signal resemblance (cf. Figure 3.1 on the facing page). Both directness and signal resemblance are anthropocentric features, so inverting an audio signal for example would affect neither of the two.

An obviously related aspect is the *modality* of the input and the output (cf. Section 3.1.4). Both should ideally match or at least be related in some way to facilitate the processing in specialised areas of the brain (directness) and reduce the alienness of the stimulus (signal resemblance).

Uninterpreted subsymbolic signal transmission also allows users to become creative with the way they use a channel that is given to them. We distinguish two basic ways such a channel can be used, called *explicit communication* and *implicit context communication*. These terms are explained in Section 3.1.3.

It has to be said, though, that there are actions or signals that are symbolic in nature and therefore signal resemblance in these cases deviates a lot from the above understanding of the term. From the premise that symbols tend to be interpreted consciously, it follows that it is especially challenging to keep such a display in line with the notion of calmness. Warnock et al. [563] found little difference, however, when comparing the interruption potential of notifications using different modalities, including symbolic and subsymbolic ones.

Privacy

Privacy is of particular importance for systems that transmit continuously from private places, both from an ethical point of view as well as most likely in terms of user acceptance.

There are different aspects to privacy and following from that different threats to it.

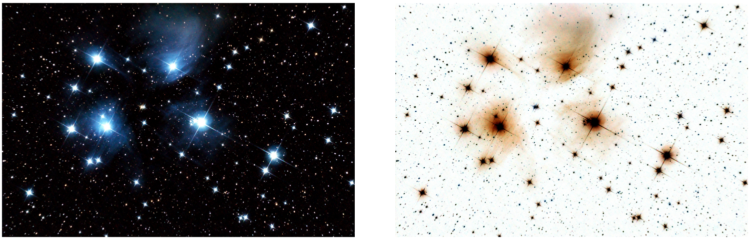

(a) The Messier 45 star cluster in a normal visible light photograph and the same image inverted. ©Photographed by Fred Locklear, published under the Creative Commons *by-nc-sa* licence on Flickr.

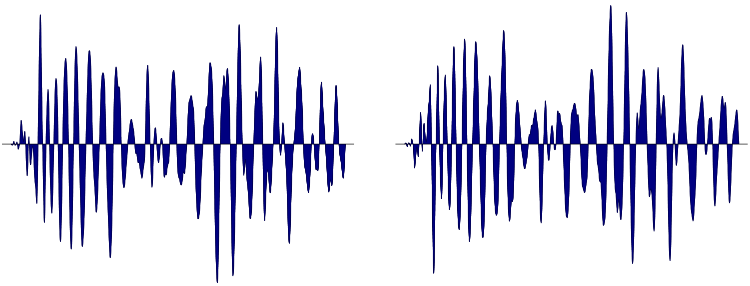

(b) An audio signal and its inverted form. While arithmetically the difference is very large, they both sound exactly the same. ©Recorded by Herbert Boland, published under the Creative Commons attribution licence on Freesound.

Figure 3.1. Illustration how the same transformation does not affect directness regardless of modality but does affect signal resemblance in the visual domain.

Gilbert [177] gives a good overview of the privacy discussion and mentions several aspects to privacy: "privacy as confidentiality: we might want to keep certain information about ourselves, or certain things that we do, secret from everyone else or selected others", "privacy as self-determination: we might consider some of our behaviour private in that it is 'up to us' and no business of others (where those 'others' may range from the state to our employers)", "privacy as freedom to be 'left alone', to go about our business without being checked on", and "privacy as control of personal data: we might desire the right to control information about us – where it is recorded, who sees it". A right to anonymity is also mentioned.

Solove [503] created a complete taxonomy of privacy that features four main ways in which privacy can be compromised. These are called information collection, information processing, information dissemination, and invasion. Each of these four aspects is broken down into several subaspects. A summary of this structure can be found in Solove [504].

Gaver et al. [170] highlight four aspects of privacy they see as important for media spaces: control over surveillance, knowledge of surveillance, knowing the intention of a privacy intrusion, and the ability to avoid disturbance.

Among the theoretical models of awareness, the focus/nimbus model of awareness is especially suitable to model privacy by representing the amount of shared information as *nimbus*, the scope of queried information as *focus* and the overlap between the two as awareness (cf. Figure 2.2 on page 27).

Patil and Kobsa [419] discuss privacy with special regard to awareness systems. They distinguish between normative, social and technical privacy as different concepts that come with different threats and solutions, some of which they list but do not further discuss.

For the scope of this work, however, we believe two broad aspects are a sufficient approximation of the complex and politically charged problem that is privacy. The two aspects are taken from the Merriam-Webster dictionary [189], where they are expressed as follows.

Definition 12 (From Merriam-Webster's Dictionary):
Privacy is "the quality or state of being apart from company or observation" and the "freedom from unauthorized intrusion".

Since the first aspect is in essence about the *feeling* of not being observed [29, 140, 276], reliable knowledge of observation is an important part of this aspect of privacy (while Gaver et al. [170] called this a separate dimension).

Directness and *context-dependence* have the potential to alleviate the second aspect to privacy when they mean that the transmitted data itself is harder to interpret automatically and without knowing the necessary context (something also noted by some participants of our FEELABUZZ study; cf. Section 6.5.2.2). Still, when designing a system it is worth considering what exactly is transmitted, how it can be misused, and how the misuse can be prevented or how the usefulness of the data to a third party can be diminished.

Regarding the first aspect of privacy, it is then to consider how users feel about the data sent away, even in the case that this data only reaches the desired receiver. We generally believe it to be important that a user of an SCS must be able to monitor their own signal to learn what they are sending off. This can simply be seen as an implementation of the design principle of *feedback* after Rogers et al. [452]. Closer to the actual case, we argue that while with unmediated interactions, there is usually no "remote end", with SCS there is. Although traditional telecommunication systems have a spatial divide between cause and effect of an action, the individual connection is short-lived. Furthermore, in the real world (and also with usual communication systems), although we might not always have the possibility to monitor what others perceive from us, we have a vastly larger number of samples of interactions to mostly give us a good notion about what other interactants perceive. There can be misjudgements and false assumptions – a wall might be thinner than expected or people where we thought none to be – but in general most people do not feel insecure about such matters. On the other hand, an environment equipped with sensors whose characteristic we have no way of knowing and whose data is processed by a mapping we might not fully comprehend poses a different kind of challenge for our self-assessment. We postulate that without providing a possibility for monitoring directly what the effect of our actions on the remote end are, it is much more difficult for people to feel comfortable with the privacy aspect of such a constantly running system.[9] McEwan and Greenberg [366] implemented this quite naturally as part of their nimbus control.

Of course a user must be able to switch off a system at any time. This does not free a system designer from considering privacy aspects though, as the act of switching off a system that is otherwise constantly running also tells the remote end something. A possible remedy would be the option to send a dummy signal but this might endanger the *trust* in the system itself which is why the introduction of such a possibility has to be considered carefully.

Privacy has also been much discussed in the context of media spaces where it is seen as a good that has to be traded off against functionality while the most commonly suggested solutions are a symmetry of communication and the ability to switch the system off [48, 170, 447]. Relying on social norms is also mentioned by Gaver et al. [170] as a way to cope with the issue although a potential difficulty to scale up this approach is admitted. To inform users of privacy intrusions, Gaver et al. [170] found auditory icons [55, 305] with a nonpercussive envelope to be accepted best. Theoretical work has been done by Palen and Dourish [414] who based their model on the social psychologist Altman [4, 5]. Palen and Dourish [414] describe their work closely following Altman as follows:

[9] However, as we will see in Chapter 6, users actually did not care about privacy that much, nor did they use the self-monitoring feature we built into the system ever again after the initial setup phase.

> As a *dialectic* process, privacy regulation is conditioned by our own expectations and experiences, and by those of others with whom we interact. As a *dynamic* process, privacy is understood to be under continuous negotiation and management, with the *boundary* that distinguishes privacy and publicity refined according to circumstance.

The work the most praised by Rittenbruch and McEwan [447] is that by Boyle and Greenberg [51] though. Boyle and Greenberg combine environmental psychology and CSCW research to build up a vocabulary to describe privacy and its constituents, especially in the context of video media spaces. Their work is very detailed but they characterise three main ways individuals interact with their environment that are aspects of privacy. They call these solitude, confidentiality and autonomy and describe them as (in this order) the control over social interactions (e. g. whether or not to communicate or be the focus of attention), the control over information access (e. g. what personal information others have access to) and the control over how the individuals themselves behave and express their identity. Any technological influence on one of these interactions of individuals with their environments constitutes a possible privacy violation that a designer ought to be aware about and ought to balance against the potential gain.

Technical aspects that are *loosely* tied to privacy but that should at least be briefly mentioned are authentication (identity verification), authorisation (permission and access control) and encryption. These might or might not play a part during the conceptualisation of a system but they almost certainly should play a part during the implementation phase of a production-ready system. More of such technical aspects that might help with privacy protection are listed by Patil and Kobsa [419].

Data economy and data proportionality [251, 252, 349] are important aspects of privacy. Data that is never produced is a more effective privacy measure than any security measures that follow. So if a system does not need a camera, it should not have one. If it does but the video features are extracted in-camera, the raw video stream does not have a chance to leave the system even in case of a compromised main system.

One should not forget to consider physical properties of systems as well. A flap in front of a camera is more immediate, transparent and trustworthy than any software system can ever be.

Plausible Deniability

Another possibility to mitigate the effect of privacy intrusions is the transmission of data that allows for plausible deniability. Plausible deniability is a term that was first used in the context of the CIA during the Kennedy era [68] but has come to be used in many different fields and with varying meanings and connotations [259, 347, 594]. The Oxford Essential Dictionary of the U. S. Military defines plausible deniability "with respect to clandestine operations" as "the state of being capable of being denied by those in authority" [411]. Zurunkel [594] describes the essential aspect of plausible deniability in a political context to be the "condition in which a subject can safely and [believably] deny knowledge of any particular truth that may exist" while Mao and Paterson [347] summarise the understanding of the term in the context of network protocols as the situation when "an adjudicator might reasonably conclude that it is plausible that the entity denying involvement in a protocol run is telling the truth". For Bradner [52], plausible deniability is a social affordance in mediated communication and he seems to use the term in the same sense we do but he does not define it. Following this informal usage within the context of privacy and HCI [e. g. 8, 319, 350, 393], we define the term where communication systems are concerned as follows:

Definition 13: Plausible deniability is the extent to which an individual at the transmitting end of a communication system generally can deny an interpretation of this system's signal by the receiver.[10]

This means that an ideal system in this regard will let the receiver come to the right conclusion about the source of the transmitted signal as often as possible but will at the same time let the person who was the cause of that signal provide plausible – while not necessarily true – alternative explanations for the signal.

It is clear that these two goals are in conflict with each other though, as well as with the perceived soundness, expressiveness and trustworthiness of a system. We nonetheless believe it to be an important aspect of privacy to be able to have some degree of interpretive influence on the effects of one's actions – which goes as far as being able to believably lie if necessary.

Bradner [53] found plausible deniability to be an important factor in the readiness of users to adopt new technologies and one that was ignored in the turn-of-the-century media theories.

For LDRRs in particular, preservation of a positive self-presentation through a high plausible deniability might be an important factor in prolonging relationships – for good or for bad (cf. Section 2.5 and Stafford [507]).

It is not clear whether a system that produces symbolic output from a subsymbolic signal and thus one that provides a ready-made interpretation of a signal performs better or worse regarding plausible deniability than a system that follows our notion of *directness*. This would most likely depend on the perceived reliability of the automatic classification but other factors such as memory plasticity might play a role [35, 167, 301, 372, 581].[11]

[10] As this definition makes plausible deniability an *extent*, a better term would arguably be something along the lines of *denial plausibility* but for the sake of established terminology, we will stick with this semantic inconvenience.

[11] That is to say the more an individual trusts their own sensory and episodic memory and the less willing they are to adapt their recollection of events, the less they are willing to accept an alternative interpretation of a direct signal.

Information Capacity

Pousman and Stasko [430] distinguish between the *information capacity* as the number of information sources an ambient display can show and the *representational fidelity* which is related to feature degradation and feature abstraction by Matthews et al. [362] and encompasses five categories from indexical to symbolic. Markopoulos [348] distinguishes between the precision and the accuracy of a channel. Accuracy describes the reliability of information whereas precision can mean actual numerical precision, granularity or the amount of information delivered and can be artificially diminished through degradation or abstraction as described by Matthews et al. [362].

An orthogonal, more direct approach to information capacity would be to measure the channel bandwidth or throughput, for example as the bit rate of the modalities involved [377].

However, we believe that in the context of SCS, the information a channel can transmit should always be seen as the actual meaning a user perceives as a function of the cognitive load the necessary processing to extract this meaning invariably induces. Here, a symbolic representation of something does not need to be more difficult or time-intensive to process than the thing itself [520]. Furthermore, information can be amplified by context information the communication participants already possess and it should be differentiated from the affective function a channel might have in addition to the actual information. Information that has no relevance to the recipient should not be considered. We therefore define:

Definition 14: The **relative information capacity** is the amount of relevant factual information a user can extract from the output of a system (called the *absolute* information capacity) in relation to the cognitive load the user incurs through this extraction.

While maximising the relative information capacity is clearly an important aspect of SCS, it is very difficult to measure with the means we developed in Section 2.4. Therefore, we will in this work mostly talk about cognitive load and about various measures of *effectiveness* for SCS, such as copresence. Relative information capacity, on the other hand, can be seen as a measure of *efficiency* for SCS.

Relevance

Demanding the relevance of the transmitted information sounds like a trivial design principle for a communication system and it is indeed quite fundamental as we are going to explain but still it is worth considering *how* a designer can maximise the probability that the transmitted information is relevant. The transmitted signals should be produced while performing activities that are frequent or important for the goal of the connection (such as a feeling of connectedness). The channel should provide the feeling of participating in the interaction partner's life, which is also linked to the aspect of *affectiveness*. Everyday activities are therefore an important source of signals that can be transmitted via SCS which also further strengthens the argument made for *calmness*.

Relevance is indeed fundamental as it is one of the four so-called Gricean Maxims that underlie any reasonable conversation according to Grice [194]. These are the Maxims of Quality, Quantity, Relevance and Manner.[12] We believe the Maxim of Relevance also applies to SCS even though they are in general not meant for carrying out conversations. The other principles, which could be restated as truthfulness, brevity and clarity in this order, are harder to apply as such. Truthfulness and clarity contradict our own principle of plausible deniability but could be related to *directness* and a system's *trustworthiness*. Brevity goes against the premise of systems that run indefinitely and is therefore clearly not applicable here, as one would likewise not think of applying the Maxim of Quantity to body language either.

[12] Together these maxims comprise the so-called cooperative principle that was formulated by Grice as follows: "Make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged." [194]

Affectiveness

Affectiveness can be defined as follows:

Definition 15: Affectiveness is the capability of systems to work on an emotional level.

This might or might not be a desirable property of a system as a high affectiveness might reduce calmness. On the other hand it can be a design goal for an SCS to increase the emotional content of a connection since there is in general less symbolical content than in other communication systems. It was also one of our motivations for SCS in the first place to heighten the feeling of connectedness, so again, working on an emotional level might be of special importance.

Modality can be an important aspect in this regard [415], as can be the aesthetic design [397]. We do not think informational capacity to necessarily be a factor though. Imagine for example how a letter from someone dear to us or a piece of clothes a beloved person once wore or a place at which he or she used to sit can make us feel just because that person touched it [533]. Often, the less sensory stimuli there are, the more room is left for imagination. Instead of transmitting a certain reality with as much sensory bandwidth as possible, it might therefore often be more effective and more compatible with the other goals to give users something to build upon and depend on their minds to add in all the details.

Scale

Scale, as defined by Zhao [591], “refers to the number of people enabled by a given interface to interact with one another”. For Zhao all copresence is at least bidirectional although we do not see a reason to exclude unidirectional modes of operation out of principle. For higher scales, the classic distinctions of one-to-many and many-to-many can be made.

We say of a system that can act on a wide range of scales to have a high **scalability**.

Scale and scalability also have practical relevance where the adoption of a system is concerned. The more participants are needed to derive a real benefit from a system, the more difficult it becomes to acquire this critical mass in the first place [130, 198, 354].

Immediacy

Zhao called immediacy "the speed at which messages travel back and forth between copresent individuals"[13] [591] and Cleevely and Cawdell [82] described it as a binary property that can take the values "real-time" and "delayed". Ehrlich [130] distinguishes the requirement for speed from the differentiation between synchronous and asynchronous media. The definition by Zhao however does not specify whether the speed is limited by technological reasons (latency) or behavioural reasons. Were it only for infrastructure and software usability, there had to be no reason why email should be much slower than instant messaging. While there are small differences where the GUI and the application layer of the network architecture are concerned, the real difference between the two are user habits and consequently user expectations. So while there are still media that are inherently real-time and synchronous (telephone would be the classic example) and others allow varying degrees of delays to respond, these delays are more cultural and conceptual in nature and we think the technological distinction is better described by **persistence**.

The interviews with participants of our upstairs study (cf. Chapter 6) indicate that latency at least with this channel is not felt to be a problem but the feeling that the copresent person immediately caused the signal was said to be important. It is therefore likely that asynchronous media, persistent ones or those with extreme lag have more difficulties in creating a sense of copresence than volatile real-time channels. Work in telepresence systems has also shown that the amount of latency that is tolerable varies hugely from application to application [90].

[13]This is not to be confused with the concept of immediacy from Definition 7.

The evaluation cycle done by Hindus et al. [235] led away from synchronous to asynchronous channels. We do not believe this to be the case invariably, though, and the first results with our upstairs system seem to corroborate this belief (cf. Chapter 6).

Group Configuration

Rittenbruch and McEwan [447] distinguish private setups from semi-public (i. e. public areas of the workplace), co-located, partially distributed (i. e. a co-located group that is extended by a few remote individuals) and mixed presence ones (i. e. setups where each participant has co-located and remote communication partners).

McEwan and Greenberg [366] emphasise the importance of the possibility to communicate within groups in an ad-hoc manner.

From the peer discussions on upstairs (Section 5.3), we differentiated between the requirements arising from different kinds of places the system is used in. More precisely, we differentiated between public spaces (like city squares or entrance areas of museums), working environments (like workrooms, the area around specific desks, or recreation areas such as lunch rooms), and private spaces (either the whole area of apartments or parts like hallways that are regularly frequented).

Hindus et al. [235] say that the point that "homes are a distinct domain [...] cannot be emphasized enough."

Mobile systems are special with regards to group configuration. While they are usually the most personalised of all kinds of systems, they are also used in a range of situations, locations and with a varying number of bystanders of different closeness.

Trustworthiness

By "trust" we in this case do not mean the trust between the communication partners but the trust the users put into the system itself. Dimensions of trust according to Schneider [478] comprise correctness, security, reliability, safety, and survivability. Mostly, system designers try to maximise all of these aspects as much as possible given the circumstances. Markopoulos [348] remarks that trust can also be deliberately traded off against other aspects such as privacy and plausible deniability by reducing what he calls the *accuracy* of the system.

Our experiences with the FEELABUZZ system showed that trust in the form of reliability is an important aspect for the calmness and general effectiveness of a system and can be especially challenging to achieve for mobile systems (cf. Chapter 6).

3.1.3 Consequences of Directness

The information conveyed by SCS that follow the directness paradigm as explicated above can be split into two parts that we call *explicit communication* and *implicit context communication*.

3.1.3.1 Implicit Context Communication

The original idea of *directness* was the transmission of implicit context in a *calm* way. These can either originate from the users or from the environment, as already proposed by Murray-Smith et al. [390].

Using the example of our FEELABUZZ system (cf. Chapter 4), the time-series data in Figure 4.2 on page 103 show that different kinds of activities by the users of – in this case – mobile phones lead to the very different acceleration profiles that are displayed in the plots.

Likewise, sitting in a driving vehicle will lead to an acceleration pattern that is notably different from those caused by human movements.

Note that none of this has to be detected by pattern recognition software. There are no predefined classes. Instead, we expect the interpretation of many movement patterns to come quite naturally (cf. Section 4.3) and involve all the rich context information and world knowledge humans have. Additionally, we expect the sophistication of the interpretations to be able to fluently increase with the user experience. As there are rarely class boundaries in the real world that are not artificially drawn, we want to enable the perception of transitions between different types of movement in all their ambiguity and fuzziness in a near-analogue fashion without the need to make clear distinctions. While regression models could do so as well, the subsequent mapping back to artificial vibrotactile stimuli in a way that allows intuitive access as well as in-depth learning of subtle features that are possibly unknown to the system designer would be a major challenge. Relying on the human's long-evolved ability to interpret rich real-world data streams to us seems to be a more promising way in terms of effectiveness and a much more interesting way in terms of unintended uses and exploration by future users.

3.1.3.2 Explicit Communication

Providing users with a new communication channel that is not pre-interpreted enables them to develop their own unintended uses. While the main idea is to have a channel that is processed subconsciously by the users, the idea of calmness does not exclude the notion of using the same channel for intentional, symbolic and conscious communication. For the systems upstairs and FEELABUZZ that we present in Chapters 4 and 5, the possibilities for readily understood signals are limited though. Apart from *knocking* to do simple things such as requesting attention, synchronising or timing pre-decided behaviour, or giving short binary feedback, few intentional tactile or audible communication events will be understood by the naïve user. Although there are sophisticated means of communication through such narrow channels, most notably Morse code, we expect those to be employed only by experts and not to become widespread. Instead, we rely on people's ability to develop their own adapted communication strategies using a mixture of implicit and explicit negotiation. Quite complex and effective communication systems can emerge via such mechanisms [37, 165, 183, 218, 283, 488] and language itself is thought to have evolved this way [80, 139, 210].

3.1.4 Modality

3.1.4.1 Input Modalities

Regarding the goal of calmness, nonverbal communication cues seem the most likely candidates for SCS. Knapp and Hall [290] compiled an exhaustive list of such nonverbal communication channels of which an aggregation follows to give a broad overview of cues that can be transmitted in principle by SCS. Their taxonomy includes body motions and reactions such as hand and head gestures, including touching, movements of other body parts like legs, feet, shoulders or the torso, facial expressions, in particular gaze direction and length, posture and vegetative reactions like pupil dilation, rubescence or perspiration, prosodic features such as voice pitch, loudness and tempo, rhythm and syllable length and voice timbre, other nonverbal vocal cues such as laughing, crying and moaning to name a few and parts of the lexicon usually not considered to be words (like "uh-huh" or "um" in English), interpersonal distance (proxemics), and finally even the presentation of objects such as cloths, make-up, accessories and furnishing.

Other cues are conceivable that do not serve a purpose of communication Knapp and Hall [290] might have had in mind but still might have emotional or informational value such as someone's whereabouts, movement and object usage.

Of course, which of these features are practical to use for SCS is ultimately limited by the availability of adequate sensors and the processing capabilities, both of which are changing rapidly and constantly.

In 1976, four years after the first edition of Knapp and Hall's book, Short et al. [494] had put together an overview of nonverbal communication with a special focus towards telecommunication, although for them at the time this predominantly meant comparing telephone to early video conferencing and consequently they mainly discussed prior research on the role of nonverbal signals in face-to-face communication (such as Birdwhistell [42] and Argyle [9]).

Summarising Argyle [9], they name six functions of nonverbal channels for face-to-face communication: 1) feedback on mutual attention (using gaze direction, head gestures, hand gestures, and non-word utterances; posture might be added [316]), 2) turn-taking control (using gaze direction and head gestures, although more cues could be added [120, 344, 405]), 3) feedback on content, 4) illustrations and emphasis (cues are likely to include gaze direction, hand gestures and facial expressions), 5) nonverbal parts of the lexicon such as head shakes or symbolic hand gestures, and 6) interpersonal attitude (using all possible nonverbal cues, including proximity, hand gestures, facial expressions, and gaze direction). Their list seems to focus on the social functions of face-to-face communication and therefore does not highlight more task-oriented aspects such as joint attention.

Short et al. continue by describing the effects of certain groups of nonverbal channels in some detail with regard to telephony and videotelephony. They also warn, however, not to "attempt to predict the effects of media by extrapolation from the know qualities of the medium and the known functions of non-verbal cues". For this they name four main arguments: 1) the redundancy between several channels and cues, combined with 2) the adaptability of people to the media they use, 3) a change of people's expectations and of social norms depending on the media, and 4) the general lack of the studies underlying the knowledge on the roles and functions of nonverbal cues to account for interpersonal complexity that goes beyond their usually simplistic setup, especially where some kind of performance measure is used – some important aspects of communication are hard to measure in this way.

It is important to note that Short et al. use these shortcomings to introduce their concept and subsequent measure of social presence. We, however, add copresence as a layer below social presence that Short et al. did not have in mind because they did not consider SCS.

3.1.4.2 Output Modalities

Arroyo et al. [14] compared the disruptiveness of five different modalities: heat, smell, sound, vibration and light.[14] They found no statistically

[14] The concrete stimuli were a heat lamp directed at the user's hand, a scent dispenser delivering the smell Elmer's glue and soy sauce, a phone ring sound file, a vibrating chair,

significant difference in their small sample size but concluded from the large variations in the data that individual differences and experiences play a large role in their perception of the disruptiveness of different stimuli.

Bodnar et al. [46] compared smell, colour and sound[15] in a setup similar to Arroyo et al. [14] and with the same number of participants. They, however, found various differences between the modalities in terms of disruptiveness and effectiveness (error rate when responding or not responding to the stimuli depending on its kind, as well as self-reports). The olfactory stimuli were found to be the least disruptive but also the least effective.

Kaye [280] provide a good theoretical backing for the use of olfactory cues in HCI, describing amongst other things the fundamental issues with the huge dimensionality of scents compared to other stimuli, some dispersion methods, cultural issues and come to the conclusion that olfactory displays are best suited for data that change slowly and have a medium amount of persistence.

Sallnäs [468] looked specifically at the effect of tactile feedback on presence and social presence in collaborative virtual environments. She found a significant increase in perceived social presence when verbal communication was not possible but found no such effect when voice communication was enabled.

Warnock et al. [563] found that the modality of interruptions had no effect on the decrease in primary task performance caused by the interruptions. They compared a total of eight types of stimuli: written words, spoken words, pictograms, colour, earcons, auditory icons, vibrations, and smell.[16]

To summarise, there seems to be more research needed in this area and it might well be that the individual, the interrupted task, the social context, and the exact way of delivery of a stimulus have as much of an influence as its modality.

and indirect light via spotlights and a screen.

[15] They chose cloves and mint as scents, an on-screen coloured rectangle as the visual stimulus, and two different bell-like sounds as auditory cues.

[16] For a list of the 24 stimuli used in total, please refer to Table 1 in Warnock et al. [563].

3.2 Related Systems

3.2.1 Calm Information Displays

There is a large body of work on ambient displays, also called peripheral displays, notification systems, or ambient information systems. While these are not SCS or even real predecessors, they share the core goal of finding calm ways to convey information, although it is not data about another human being (or if it is, we will look at these systems in Section 3.2.2). Such ambient information displays exist in a huge variety of shapes and forms, from simple lamps [111] over informative installations [172, 223] to dynamic paintings [192, 206, 438, 513].

Ishii and Ullmer [264], Ishii et al. [265] and Dahley et al. [102] did seminal work in tangible interaction that inspired many later systems [128] but they did not focus on tangible computer-mediated communication between people as did the systems described in Section 3.2.3. Dahley et al. [102] in particular developed some early ambient displays, however. One of them was an array of pinwheels whose rotation was controlled by data-driven airflow, the other was a lamp that shone through a water surface that was disturbed by falling drops to convey information (Figure 3.2(o) on page 90).

Eggen and van Mensvoort [128] proposed a number of related systems that primarily display information in a calm and integrated way, e. g. as fountain height (Figure 3.2(f) on page 89) or bird song.

Gandhi et al. [166] developed and marketed several ambient displays, including the Ambient Orb, a lamp that can display various data by changing its colour (Figure 3.2(e) on page 89). In their more recent focus on awareness of energy consumption, they developed a very similar device with an embedded display called Energy Orb (Figure 3.2(h) on page 89), possibly indicating that a mere colour change was not perceived as a rich enough data channel.

Energy awareness is also the goal and motivation of the power-aware cord by Gustafsson and Gyllenswärd [202] that uses the power cord itself to display glow patterns that are a function of the current that flows through it (Figure 3.2(c) on page 89).

However, Consolvo et al. [89] found in their study with an ambient picture frame that the emittance of light was not a desirable property for their ambient display; at least not when it does not adapt itself to low-light conditions because it then drew to much attention and therefore “stopped being ambient.”

3.2.2 Calm Communication

There have been quite a few systems exploring alternative ways to communicate and more specifically, to communicate in a less obtrusive way.

Chang et al. [73] developed a picture frame where a user could touch one and the other would light up (Figure 3.2(l) on page 90).

A common way to let people know of each other in a non-obtrusive way is to use status information, mostly IM status because it is readily available.

Dey and de Guzman [114] suggested a whole range of devices displaying IM status. Quoting their suggestions and descriptions in full, these were:

PresenceStool a “plastic stool that changes colors to reflect a loved one’s online presence” (Figure 3.2(k) on page 90)

PictureFrame a “picture frame holding a picture of a loved one, that displays a physical icon to current presence”

StatusGrid a “photo wall in a grid pattern, with colored lights for each photo of a person indicating presence”

StuffedAnimal a “stuffed animal or pillow that changes color or temperature and vibrates to reflect presence”

ScentsOfPresence a “small toy that releases different scents providing a reminder of a loved one, when that person’s presence status changes”

AugmentedMirror a “mirror augmented with colored lights and physical icons to indicate [a] loved one’s presence”

RotatingLights a “floor lamp with multiple bulbs that flash in different patterns to reflect a loved one’s status”

InformationAppliance “augmented displays of information such as thermometers or clocks that use color to reflect presence”

Chime a “wind chime that chimes at different frequencies to indicate presence status”

PhotoDisplay a “copy of a personal photograph where the background imagery is replaced with a translucent background that changes colors to reflect changes in presence status”

Based on preferences expressed in focus group studies, they chose to implement the picture frame and the mirror as prototypes. In the subsequent study they found that these prototypes did increase the awareness of the remote person’s online status and the self-reported feeling of connectedness.

De Guzman et al. [105] implemented four prototypes displaying IM status. These were *Expanding Ball*, a plastic toy Hoberman sphere [237] (cf. Figure 3.2(n) on page 90) that opens and closes according to the status of the remote user, *Spinner*, a blue box with wheels that would spin a picture of a contact when they change their status, *IMFrame*, a digital picture frame displaying someone’s online status, and *Chime*, a wind chime with colour LEDs, displaying status and status change with light and sound. Initial results from focus groups were mixed.

Kuwabara et al. [312] transmit status information via an ambient display in form of a plant with an integrated fibre optic lamp.

Bhandari and Bardzell [36] suggested a series of systems transmitting voluntary but non-verbal information. These were the MatchUs Board, a synchronised touch-enabled drawing program (Figure 3.2(q) on page 90), the Together Aquarium, a synchronised aquarium simulator, CanCan, a button on a watch or similarly worn object that signals a short “talk to you later” to the partner, and Audible Gifts, an artifact that enables voice messages, that are artificially restricted in length to increase their perceived value.

Go et al. [180] suggest a family of devices they call Familyware, intended to provide “support for small and intimate communities”, and more specifically friends and family. They propose three exemplary systems, two that connect a toy (a teddy bear and a rattle) to a digital photo frame, changing the picture when the toys are used, and one more symmetrical concept, designed for adult users that consists of two paired necklaces, each holding a stone with temperature readings and the ability to control its own temperature so that one user holding their necklace in their hand will trigger the stone of the paired necklace to grow warm (Figure 3.2(d) on page 89).

Mynatt et al. [392, 461] also developed a system based on a digital photo frame, intended for family watching after an elderly relative who is living alone. In their case, they use activity data to change the border of the frame to a number of icons that each represent a day's worth of activity (Figure 3.2(g) on page 89).

Hindus et al. [235] developed several concepts and prototypes in their Casablanca project, namely CommuteBoard and MessageBoard, both of which are shared electronic sketch pads, ScanBoard, the same idea using scanning instead of pad input, RoomLink, a shared audio space using a permanent audio connection, and inTouch, a pair of tokens where touching one lets another vibrate, light up or become warmer (Figure 3.3(d) on page 97).

PresenceLight was another concept presented by Hindus et al., a light whose brightness is determined by the activity at a remote place. The focus group feedback actually led to the renunciation of activity sensing because it was perceived as too intrusive and too imprecise and to the development of the Intentional Presence Lamp where presence had to be actively indicated. A very similar concept was implemented by Deschamps-Sonsino [111] (Figure 3.2(j) on page 90).

Light is also used in some concepts by Tollmar and Persson [533], including a branching lamp where each branch represents one person and is touch-sensitive, and a light sculpture where each individual lamp's brightness is determined by remote movement and the amount of energy and water used.

McEwan and Greenberg built a system they dubbed Community Bar [366, 367, 456] (cf. Figure 3.2(p) on page 90). It is a software component that runs on the users' desktop and shows a small window that serves as a monitor bar of one's contacts and integrates different channels like video and chat. Escalation to audio links or better video connections is done very easily. Users can also use a slider to change their nimbus and focus (cf. Figure 2.2 on page 27 and Rodden [450]).

Gibbs et al. [176] presented a design sketch for an example of what they called *phatic technology*. They use the following definition:

> Phatic technologies are technologies specifically designed to sustain social interactions, rather than convey information. Phatic technologies are not concerned with the utility of the interaction, the usefulness of the information, nor the ease-of-use of the device – although each of these may be important for end-user experience. Phatic technologies should be judged by the degree to which they contribute to a feeling of ongoing connectedness. Technologies that support phatic exchanges are similar to devices that support peripheral awareness [however] phatic interactions are not limited to being peripheral.

What Gibbs et al. [176] actually propose is a palm-mounted device that will indicate the fact that a remote person is composing a text message to the wearer through visual and haptic cues. So it uses implicit input – the sender would write the message anyway – but none that is continuous. Assessing the obtrusiveness of the output is difficult without studies or even a prototype device but it appears similar in nature to notifications in general, although checking them would be easier with the palm-mounted design than with a phone in your pocket, more akin to wrist-top computing that is gaining popularity [145].

There are also a number of commercially available or announced products that enable novel ways of communication.

Nabaztag, with the successor now called Karotz [204] (Figure 3.2(b) on the facing page), were amongst the early commercially available ambient displays. They were little rabbit-shaped devices that could signal to the user through sound, lights and movable ears. The Nabaztag could function as an ambient display for arbitrary information, including information about interpersonal communication such as e-mail, SMS and phone calls. They also introduced the novel communication channel of mirroring the movable ears between devices so one user could see when the remote communication partner moved an ear.

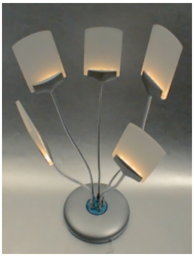

(a) 6th sense by Tollmar and Persson [533]

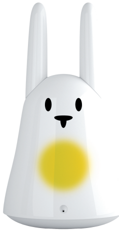

(b) Karotz by Haladjian and Mével [204]

(c) Power-Aware Cord by Gustafsson and Gyllenswärd [202]

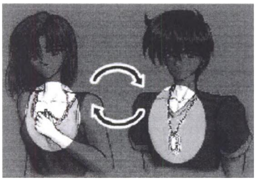

(d) Familyware necklaces by Go et al. [180]

(e) Ambient Orb by Gandhi et al. [166]

(f) Data Fountain by Eggen and van Mensvoort [128]

(g) Family Portrait by Rowan and Mynatt [461]

(h) Energy Orb by Gandhi et al. [166]

(i) Tree-lamp by Tollmar and Persson [533]

Figure 3.2. Gallery of related systems, part 1. Images reproduced without express permission in accordance with German legislation (§ 51 UrhG).

(j) Good Night Lamp by Deschamps-Sonsino [111]

(k) PresenceStool by Dey and de Guzman [114]

(l) LumiTouch by Chang et al. [73]

(m) SynchroMate by Gibbs et al. [176]

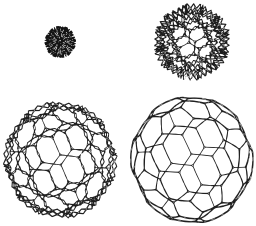

(n) Hoberman sphere in different stages of expansion (Hoberman [237]), used by de Guzman et al. [105].

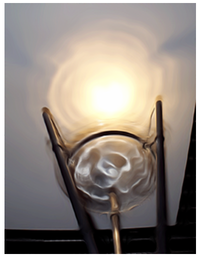

(o) Water Lamp by Dahley et al. [102]

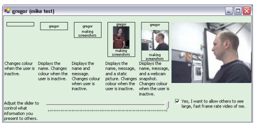

(p) Community Bar by McEwan and Greenberg [366]

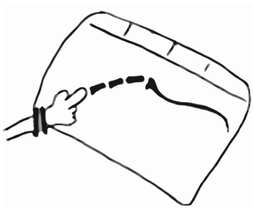

(q) MatchUs Board design by Bhandari and Bardzell [36] (vectorised from low resolution original)

(r) Artificial communication channel concept by Mertes [375]

Figure 3.2. Gallery of related systems, part 2. Images reproduced without express permission in accordance with German legislation (§ 51 UrhG).

Kowalski et al. [303] proposed a system they called "Cubble", explicitly designed to support long-distance romantic relationships (LDRRs; cf. Section 2.5). They use a hand-holding metaphor by having both users touch their respective devices. These can interchangeably be a base station or a smartphone (Figure 3.3(b) on page 97). The device will display colours and also give tactile feedback.

A project called Pillow Talk [384] is being prepared for a crowdfunding campaign at the time of this writing. It is a pair of pillows where one lights up when the other is being laid on. Additionally, the other person's heartbeat is recorded using a wristband and replayed acoustically.

Kaye [281] explored the use of minimal communication in a desktop context. Users could click a circle which would colour in this circle at the remote end. This colour would fade over time. They found that users would ascribe a lot of meaning to this minimal channel, depending on their knowledge of the other person and the specific context. Kaye and Goulding [282] had previously designed three more physical concepts for LDRR couples; *How do I love thee* are connected abacuses, *Hand Holding* are two design variants using warmth to convey remote hand-holding, and *Love Egg*, an egg that you can speak into to leave a message which triggers the remote egg to roll around in the dish it lies in.

Sound was also explored as an output modality. Isaacs et al. [263] used earcons [370] to convey activities, deliberate predefined short messages and the identity of the sender. Mynatt et al. [391] tried different approaches, namely a soundscape enriched with auditory icons [55], musical earcons, voices and a hybrid approach to convey information. Tünnermann and Bovermann [49, 544] developed a technique called *auditory augmentation* by which data modulate sounds that are already present in the environment, a concept which Tünnermann et al. [545] later expanded to their broader concept of *blended sonification* that also includes our upstairs system.

While communication systems are almost exclusively understood as *remote* communication systems, attempts to calmly improve face-to-face communication have been made. Mertes et al. [116, 375, 376] called this *artificial communication channels* and proposed or built several systems that use augmented reality technology to overlay someone's surroundings with another person's attentional focus while both were sharing the same physical space (cf. Figure 3.2(r) on page 90).

3.2.3 Tactile Communication

Because one of our SCS is FEELABUZZ, a tactile way to communicate, we will at this point take a closer look at other systems that experimented with tactile or haptic communication.

Brave and Dahley [54] already proposed a system for haptic communication in 1997 called inTouch. It consisted of two paired devices, made up of three rollers that transmit all movement to the paired device where it is reproduced (Figure 3.3(d) on page 97).

Other early work in this area was done by Strong and Gaver [519]. They developed three prototypes called Feather, Scent and Shaker. While the first two are not haptic and basically unidirectional,[17] Shaker is a more symmetrical design (Figure 3.3(a) on page 97), consisting of a device that physically translates movement into a current and a receiver that translates this current back into movement. Each partner has one such pair of devices. Strong and Gaver mention themselves that it would be better to have both functions in one device but the analogue nature of their prototype made this impractical.

Smith and MacLean [501, 502] tested a prototype where users handled knobs that were each connected to a motor that could measure how the user turned the knob but could also act as a force-feedback device to mirror the actions performed on the other end (Figure 3.3(e) on page 97). Users were separated by a wall. Smith and MacLean wanted to find out how effectively certain emotions could be transmitted through such a tactile link. They found that participants were much better than chance when presented with a known set of four emotions but still a long way from 100 % (cf. Table 3.4 on page 94 for exact results).

[17] Feather and Scent use a picture frame as the sender and as the receiver a tall transparent cone with an air current supporting a feather in Feather and a fragrance oil burner in Scent. Handling the picture frame will cause the feather to be blown up or fragrance oil to be heated.

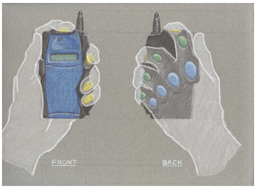

(s) ComTouch concept by Chang et al. [72]

(t) Prototype with optional glove by Eichhorn et al. [131]

(u) Hug Shirt by Rosella and Genz [458]

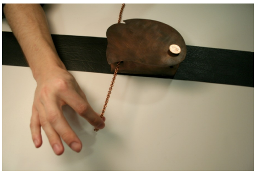

(v) Tug n' Talk by Adcock et al. [2]

(w) Breathing phone prototype by Hemmert et al. [226]

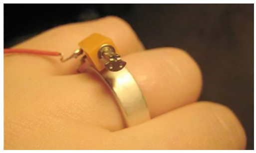

(x) United-pulse by Werner et al. [570]

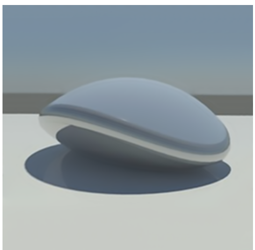

(y) The feelybean by Kontaris et al. [300]

Figure 3.2. Gallery of haptic communication systems, part 1. Images reproduced without express permission in accordance with German legislation (§ 51 UrhG).

Perceived Emotion	Sent Emotion Angry	Delighted	Relaxed	Unhappy
Angry	.62	.09	.03	.06
Delighted	.23	.49	.15	.15
Relaxed	.07	.25	.57	.31
Unhappy	.08	.16	.26	.48

Table 3.4. Data from Smith and MacLean [502] on the ability of their tactile link to convey certain emotions. The confusion matrix shows the portion of trials in which the emotion in the rows was perceived when the emotion in the columns was sent. The cell colouring is detailed in Box 4 on page 149.

Chang et al. [72] developed concept and prototype for vibrotactile communication that is meant to accompany spoken language. In their ComTouch system, each finger has one button to measure pressure, one point where the same finger gets feedback on the resulting vibrotactile signal and a second point where the vibrotactile signal of the partner can be felt. The ergonomic design Chang et al. envisioned is shown in Figure 3.2(s) on page 93. The evaluation of their one-finger prototype focused on how vibration can augment or replace spoken communication.

Sommerer and Mignonneau [505] developed two paired hand-held devices that transmitted heartbeat and breath as an art project. The heartbeat was conveyed visually and as a tangible pulsing, the actuator for the breath was a small ventilator (Figure 3.3(c) on page 97). Such involuntary cues such as the heartbeat [570] (Figure 3.2(x) on page 93) or breathing [224, 225, 226] (Figure 3.2(w) on page 93) have been used in other systems as well.

Oakley and O'Modhrain [403] built a prototype they called Contact IM which focused on the ability to convey haptic communication in an asynchronous way. The tactile information itself was the movement of throwing a ball over a virtual tennis court, recorded and replayed either by force-feedback joysticks or a Phantom device.[18]

Kaerlein [275] built a small anthropomorphic robot to be used as a tactile communication device (Figure 3.3(g) on page 97).

Eichhorn et al. [131] developed prototypes of hand-held devices, featuring a knob that would act as both input and output. The devices were covered by fabric to make them more comfortable (Figure 3.2(t) on page 93).

O'Brien and Mueller [404] designed and implemented a series of devices that were to imitate different styles of holding hands but do so remotely. An early design probe intended for data collection can be seen in Figure 3.3(h) on page 97 but these were found to be too conspicuous and so the participants of an early study rejected them for social reasons. Later models were therefore built to be more discreet. Figure 3.3(i) on page 97 shows one of three mock-ups that were used to investigate what style of holding hands people preferred. One result from O'Brien and Mueller that is interesting with regard to our FEELABUZZ system is that participants generally objected using a dedicated device for the tactile communication and one participant mentioned "that she wanted something she could relate to personally, and not 'everyone has the same yellow ball.' "

Fogg et al. [151] developed a few prototypes they called HandJive. They consisted of two balls that could be moved along one axis which would then translate to a movement of the corresponding ball on the paired device but along an orthogonal axis so that users would not struggle for control of the ball position. Fogg et al. also tried to establish a simple form of symbolic communication they called "Tactilese" (cf. Figure 3.3(f) on page 97) but unfortunately did not pursue this much further.

Adcock et al. [2] developed a belt buckle that simulates tugging someone's shirt (Figure 3.2(v) on page 93).

Goodman and Misilim [184] presented a pair of beds where an array of pressure sensors and heating pads would indicate the position of one partner as warmth in the remote bed. Dunne and Raby [127]

[18] http://geomagic.com/en/products/phantom-omni/overview

did something similar with public benches which people were uncomfortable with, demonstrating that mediated intimacy can be perceived in similar ways to physically close intimacy [168].

Rosella and Genz [458] developed a shirt that transmits hugs as part of their fashion collection, a concept similar to one of the prototypes developed by Vetere et al. [554].

Gorilovsky et al. [186, 187] at the time of this writing are taking preorders for a wristband that transmits to touch and tapping through a capacitive sensor and an accelerometer and makes a paired wristband vibrate.

Mobile phones with their built-in vibrotactile actuators have also been used for tactile communication. Crossan et al. [96] implemented a multimodal contact list that includes vibrotactile feedback. The contact list presents context information similar to the earlier unimodal work by Siewiorek et al. [495] but also provides a tangible *mood vector* and then allows users to contact the people on the list via bidirectional vibrotactile messages. Hoggan et al. [239] transmit vibrotactile signals called *pressages* during phone calls when the side of the phone is pressed while Park et al. [417] use touch input for the same purpose in their system called CheekTouch. Kontaris et al. [300] pursue a similar approach but use special hardware to be used during Skype sessions (Figure 3.2(y) on page 93).

There is also a considerable amount of work that focuses either on the recognition of haptic gestures or on mapping non-haptic cues to haptic feedback [58, 106, 133, 460]. Murray-Smith et al. [390] detect footsteps through accelerometer data and turn it into vibrotactile feedback on a remote device to help gait alignment during mobile phone conversations.

Finally, there are also a number of haptic and otherwise unconventional remote communication devices that are quite forthright about targeting a more specific use case for couples [57, 225, 382, 576].

(a) Shaker by Strong and Gaver [519]

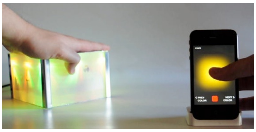

(b) Cubble by Kowalski et al. [303]

(c) Mobile Feelings by Sommerer and Mignonneau [505]

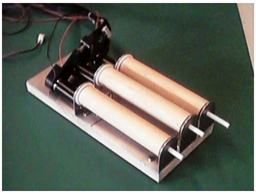

(d) inTouch by Brave and Dahley [54]

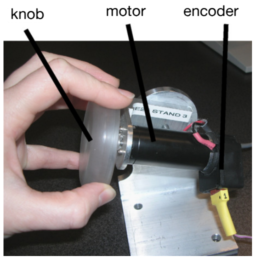

(e) Experimental setup from Smith and MacLean [502]

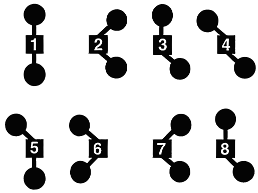

(f) Representation of a movement routine on HandJive by Fogg et al. [151]

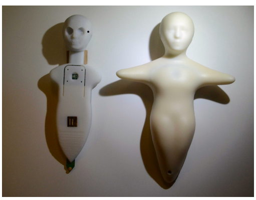

(g) Elfoid by Kaerlein [275]

(h) Earlier hand-holding prototype by O'Brien and Mueller [404]

(i) Later hand-holding prototype by O'Brien and Mueller [404]

Figure 3.3. Gallery of haptic communication systems, part 2. Images reproduced without express permission in accordance with German legislation (§ 51 UrhG).

4 FEELABUZZ – A Vibrotactile Smartphone-Based Link

The idea implemented in our system called FEELABUZZ is to use smartphones as a widely available, in fact almost ubiquitous platform to transmit and convey movement. We do so by mapping the accelerometer of one phone to the vibration motor of another phone as illustrated in Figure 4.1 on the following page. This principle was implemented on Openmoko,[1] WebOS[2] and Android.[3]

The choice of modality was made because touch is arguably the most immediate, the most affective and intimate, and – when it comes to media – one of the most overlooked modalities used for human communication. It can convey emotions and feelings on a direct and primordial level [131, 221, 383, 554] and plays a crucial role in human social interaction [87, 290]. It is also one of the modalities people being apart miss the most, not only for couples but also with friends [406]. These circumstances support our goal of *affectiveness* while the choice of unmodified smartphones as a platform goes along with the notion of *calmness* (cf. Section 3.1.2).

We gave an overview of previous work on tactile communication in Section 3.2.3. What can be seen by looking at these approaches is a

[1] http://www.openmoko.org/
[2] https://developer.palm.com/
[3] http://www.android.com/

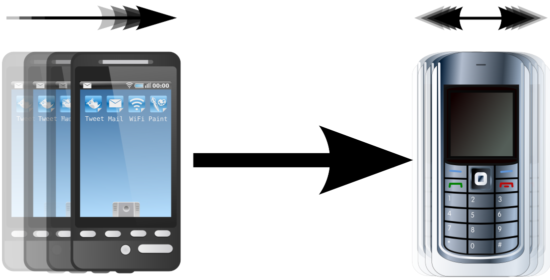

Figure 4.1. Schematic illustration of FEELABUZZ. The change of movement of one mobile phone as detected by its built-in accelerometers is mapped to the vibration motor of another phone.

focus on specialised devices and on an event-based nature of communication, often lacking richness and therefore directness. With FEELABUZZ, we did not want to use additional hardware and we wanted to provide a rich, continuous communication stream. Additional hardware is cumbersome and intrusive, and we believe that the lack of directness comes at the cost of feelings of immediacy and connectedness.

A system with somewhat similar goals is Shoogle [579] which uses techniques from model-based sonification [231] to generate rich vibrotactile interaction and exploration possiblities for data on mobile devices.

The effect on connectedness by any comparable systems has not been investigated so far but Sallnäs et al. [469] looked for an increase in perceived social presence using haptic force feedback devices. They could not find a significant effect but their setup was very different from the SCS presented here (i. e. very explicit communication, force feedback instead of vibrotactile feedback, no mobile devices).

The choice of modalities for FEELABUZZ may seem obvious because both acceleration and vibration are available on about every modern smartphone. The accelerometers are normally used for nothing else than to rotate the screen according to the direction of gravity, although in recent times, motion controls for games have become another important use for these sensors. Vibration, on the other hand, has been available since way before the smartphone era, serving as a silent alternative to ring tones. However, the choice of vibration as an output modality not merely stems from the prevalence of this unobtrusive stimulus on the chosen platform and its availability and discreetness when carrying the phone in one's pocket but also from the fact that movement is naturally transformed into vibration and similar tactile feedback in the real world. Everyday examples include footsteps on the floor, multiple persons using one stair rail, someone stirring next to you on a sofa or even the feeling to one's own hand when stroking something. In fact, feeling vibrations is vital to the ability to discern fine surface textures by touch [274].

A maybe more subtle advantage of mobile phones other than their virtual omnipresence and the availability of the sensible combination of sensor and actuator just discussed is that they are usually permanently with their users wherever they go and are thus a kind of constant companion that is able to record movements without any setup time or danger of forgetting an extra device. Not having to buy and to carry around an extra piece of hardware is also an important property. Using phones also makes it easier to integrate the new haptic channel with existing auditory, visual and maybe textual channels in the future, thereby extending the phone's capabilities as a communication device. As we have our phones with us or nearby most of the time, they are well suited not only for *explicit communication* but also for *implicit context communication* (e. g. walking or riding the bus). These concepts are detailed in Section 3.1.3.

Additionally, when considering to call someone, it would be helpful to know the callee's current context. Parts of the context can often be guessed or are already known to the caller. This is for example the case when someone calls someone else to make sure they are already on their way to an appointment, or when somebody calls their partner to see whether he or she is still on their way home from work or has already arrived. At the time of this writing, there is no easy way of acquiring that kind of information *without* actually calling someone and thereby risking to interrupt that person, although attempts in that direction have been made [96, 107, 495].

As modern mobile phones gather a lot of data that can be used as context information, attempts have been made to transfer this data in one way or another [96, 435, 495]. This can be quite an intrusion into one's privacy though. We believe that movement data is relatively harmless here, as opposed to for example location data, an assessment that our preliminary results from Chapter 6 seem to support.

Related work by Haverinen [216] provides additional motivation for the explicit communication aspect of FEELABUZZ. Haverinen chose to use explicitly triggered smartphone vibrations called *pressages* to augment phone calls. In her evaluation, she states the following:

> When directly asked whether the pressage feature could be useful, both couples were quite doubtful. On the other hand, the participants mentioned some interesting ideas how the vibro-tactile messages could be used. The first couple suggested that pressages might be more interesting and usable if they were not mixed with the voice.

For FEELABUZZ, we opted for a direct transmission of acceleration data in accordance with the directness paradigm (cf. Definition 11). With the first two prototypes, this transmission was permanently active.

With the Android GUI, a new interaction mode was created where User A can touch the entry of User B on the contact list upon which the activity of User B is rendered as vibration of User A's phone. This is supposed to further strengthen the metaphor of touching another person and feeling their movements.

We generally hope that the tactile nature of the link – in addition to providing useful context information – serves to create an increased sense of connectedness through the basal nature of touch, thereby providing the affectiveness called for in Section 3.1.2. First results from our user studies were very promising with regard to the creation of an emotional connection (cf. Figure 4.16 on page 126 and Section 6.4.2).

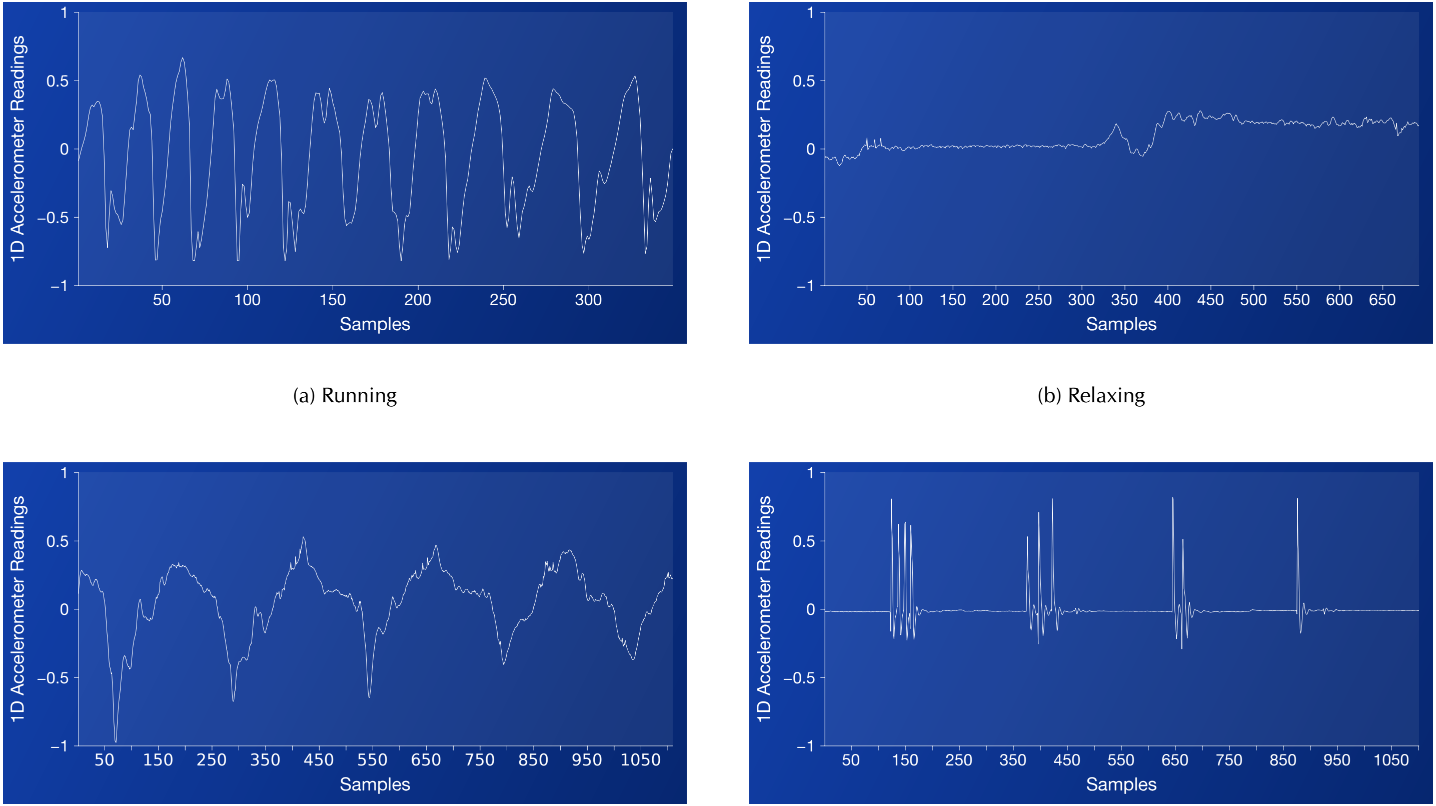

(a) Running

(b) Relaxing

(c) Swinging

(d) Hand Tapping on the Phone

Figure 4.2. Raw 1-DoF accelerometer data of different activities recorded with 100 Hz. The ordinate is deliberately kept unit-less and serves only to show that the maximum of the scale is equal for all plots.

4.1 Goals and Challenges

Albeit the vibrations today's mobile phones can make are a poor substitute for the actual touch of another person, we believe that the knowledge that it is the very movement the other person is doing in that very moment that makes a user's phone vibrate in a certain way can give them a real feeling of presence and intimacy. As we argued in Section 3.1.2, a large sensory bandwidth is not needed to achieve the goal of affectiveness; the only thing needed is something for the users' imagination to build upon.

Still, how much is there to really build on? Figure 4.2 on page 103 shows accelerometer data for different activities. It is not necessary to be an expert to distinguish these four sample activities. We are optimistic, though, that people *will* become experts in a sense, in that they will learn to pick up even the comparatively subtle cues that separate the way of movement of close persons from everyone else's way. Provided that a strong social tie is a profound enough motivation to train their sense of touch to achieve this, we think the actual sensitivity of touch is often underestimated [179].

Figures 4.2(a)–(c) on page 103 show acceleration data of different activities that would serve to transmit *implicit context* as defined per Section 3.1.3. Figure 4.2(d), on the other hand, shows a tapping with the palm on a mobile phone resting in one's pocket, representing an activity that would fall under *explicit communication*. We believe that people might become quite creative once given such a straightforward tool as FEELABUZZ and could for example develop their own signals to quickly inform someone about things without even taking their phones out of their pockets (cf. Section 3.1.3.2). In this example, it is evident that the sensor had been tapped on first four times, then three times, then twice and finally once.

The challenge now is to find a mapping from acceleration data to vibration output that makes it as easy for the users to discern all of these kinds of patterns with their skin as it is to tell them apart visually in the graphs and to deduce the underlying activity in an intuitive way, relying as much on pre-existing world knowledge on part of the user

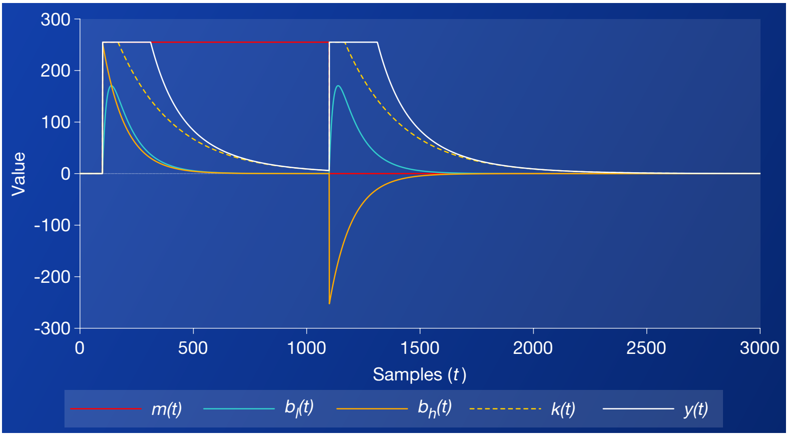

Figure 4.3. Step response to the rectangular signal $m(t)$. There is a linear interpolation between samples for the sake of a smoother rendering.

as much as possible. The mapping also needs to take the inertia and limited fidelity of the mobile phone actuators into account.

4.2 Implementation

As mentioned above, we implemented FEELABUZZ on three different platforms. The first was the open platform Neo FreeRunner (GTA2) [386] running Openmoko SHR [385]. The complete openness of this platform allowed for very easy prototyping. This first implementation was followed by an adaptation for the WebOS platform Palm Pre [486] and then a reimplementation for Android [185] on HTC Desire smartphones [332].

When choosing a mapping from acceleration to vibration strength, we followed the notion of *signal resemblance* as closely as possible (cf. Section 3.1.2). Heikkinen et al. [221] additionally provided some insight into general expectations of users towards haptic communication which include varying intensities of touch and poking and knocking as metaphors.

4.2.1 Signal Processing and Vibrotactile Mapping

Let $\boldsymbol{s}(t) = \left(s_1(t), s_2(t), \ldots, s_S(t)\right)^{\top}$ be the readings of the S accelerometers at time step t, where $s_i(t) \in [0, s_{max}]$, and $y(t) \in [0, y_{max}]$ be the output in form of the vibration strength sent to the vibration actuator, i. e. the target speed of the vibration motor.[4] The mapping from $\boldsymbol{s}(t)$ to $y(t)$ consists of several steps. First, we compute the magnitude of the vector of sensor values.

$$m(t) = \rho \|\boldsymbol{s}(t)\| = \rho \sqrt{\sum_{i=1}^{S} s_i(t)^2} \tag{4.1}$$

with ρ being a normalisation factor:

$$\rho = \frac{y_{max}}{\sqrt{S s_{max}^2}} \tag{4.2}$$

Now an RC high-pass filter [291] is applied to the sensor values with the decay constant α_h

$$b_h(t) = \alpha_h \Big(b_h(t-1) + \big(m(t) - m(t-1) \big) \Big) \tag{4.3}$$

which gets rid of the gravitational acceleration and other constant or long-term acceleration influences[5] without losing as much inertia as a simple derivation would.

[4]For the specific values that S, s_{max}, y_{max} and other constants take, please refer to Table 4.1 on the facing page.

[5]When using the maximum sample rate of 30 Hz on the Palm Pre it is possible to shake the phone so hard that the accelerometers will register a constant acceleration. The 100 Hz of the Neo FreeRunner were enough to circumvent this phenomenon. To prevent the high-pass filter from eliminating the constant maximum acceleration on platforms that cannot read from the sensors fast enough, it turned out to be painfully inelegant yet obscenely effective to artificially set the sensor value of one time step to 0 when a threshold number of successive near-maximum acceleration frames is exceeded.

$\alpha_h = 0.99$	$\beta_a = 3$
$\alpha_l = 0.05$	$\beta_{b_h} = 2$
$\alpha_e = 0.4$	$\beta_{b_l} = 3$
$\eta = 1$	$S = 3$ or 6
$y_{max} = 255$	$s_{max} = 2268$

(a) Neo FreeRunner

$\alpha_h = 0.967$	$\beta_a = 5$
$\alpha_l = 0.157$	$\beta_{b_h} = 2$
$\alpha_e = 0.4$	$\beta_{b_l} = 2$
$\eta = 2.5$	$S = 3$
$y_{max} = 100$	$s_{max} = 2$

(b) Palm Pre

Table 4.1. Constants used on the (a) Neo FreeRunner (sensor sampling rate 100 Hz) and the (b) Palm Pre (sampling rate 30 Hz). The plots in this chapter all use the Neo FreeRunner constants because its higher sampling rate is less prone to artifacts.

Subsequently, an exponential smoothing is applied with smoothing factor α_l:

$$b_l(t) = \alpha_l |b_h(t)| + (1 - \alpha_l) b_l(t-1) \tag{4.4}$$

This is important in order to give more inertia to the system in a controlled way so that a lot of activity from the sender will add up to give an increasingly strong signal on the receiving end (cf. Figure 4.4 on the next page). This turned out to be what best matched our intuitive a-priori expectations of how the system *should* behave. The effect of this is that one can "charge" the system by prolonged high-energy shaking, feeling a bit like spinning an unbalanced flywheel faster and faster.

If we were to output $b_l(t)$ as vibration, it would feel quite unresponsive and slow, however. The low-pass filter has the drawback of levelling out all of the more impulse-like parts of the signal which are a salient feature and also quite important for signalling. To preserve these impulse components, we add them back in with a simple kind of spike detection. This also has the benefit of making the system more responsive to quick accelerations as the then-detected spike will kick-start the acceleration motor.

For this, we compute the moving average over the last n time steps,

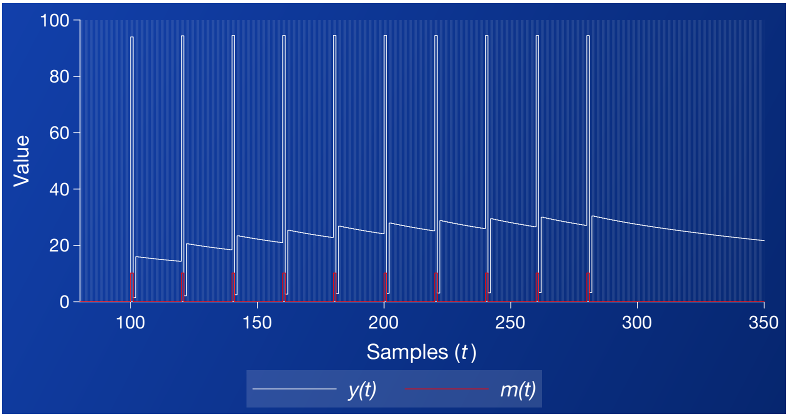

Figure 4.4. Filter response $y(t)$ to a burst of delta pulses $m(t)$. No interpolation has been applied in order to preserve the form of the Dirac deltas; the background pattern marks samples alternatingly to give a better impression of the discrete nature of the t axis.

defined for any function $x(t)$ as

$$MA_n(x, t) = \frac{1}{n} \sum_{i=0}^{n-1} x(t - i) \tag{4.5}$$

and check if the high-pass-filtered signal $b_h(t)$ exceeds the threshold of β_a times the moving average. In less rigorous terms this means that a spike is a sample that is significantly higher than its local neighbourhood. If this is the case, we perform an exponential mapping of the spike signal and add it back to the low-pass-filtered signal with the adjusting

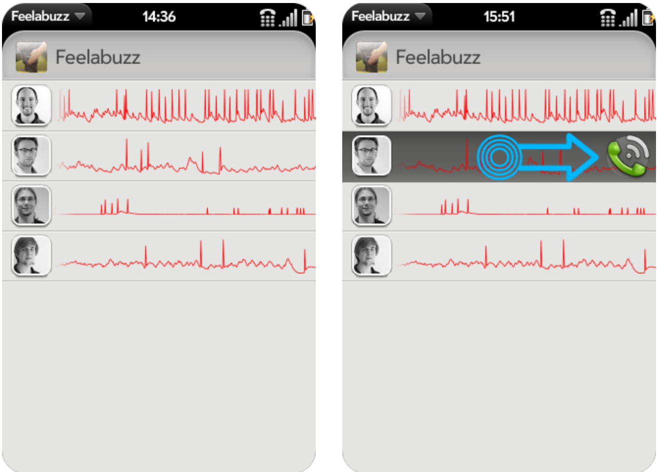

(a) (b)

Figure 4.5. Screenshots of the WebOS dialer application. The interaction is illustrated in (b): the activity of a remote contact is rendered to the vibration motor when this contact is touched (blue circle). Additionally, a call icon is displayed. By moving the finger and releasing it on the icon, a call to the contact is initiated. Pictures by Sebastian Hammerl

coefficients β_{b_h} and β_{b_l}:

$$k(t) = \begin{cases} y_{max} \left(\frac{\beta_{b_h} b_h(t)}{y_{max}} \right)^{\alpha_e} & \text{if } b_h(t) > \beta_a MA_n(b_h, t), \\ 0 & \text{else.} \end{cases} \tag{4.6}$$

$$y(t) = \min\Big(\eta\big(k(t) + \beta_{b_l} b_l(t)\big), y_{max}\Big) \tag{4.7}$$

We used a moving average window of $n = 5$. The normalisation constant η is necessary on some platforms to linearly correct for sensor or actuator sensitivities that are too low. For the Palm Pre we found a value of $\eta = 2.5$ to work well. For the Neo FreeRunner, no such constant is needed. Finally, the output is cropped to y_{max}. For the Neo FreeRunner we noticed that vibration strengths of $y(t) < 30$ are not noticeable so we set them to 0 to not unnecessarily strain the battery (this cutoff is not included in Equation 4.7 nor in Figures 4.3 to 4.8).

In the following, we plot several filter responses to artificial input functions to illustrate various properties of the above transformation, including intermediary steps where this helps to understand what is happening. Figure 4.3 on page 105 shows a step response, with the intermediary steps $b_h(t)$, $b_l(t)$ and $k(t)$. The high- and low-pass characteristics of $b_h(t)$ and $b_l(t)$ can clearly be seen. The output signal $y(t)$ subsequently shows an immediate response as well as a strong inertia that can be configured independently from $b_h(t)$. Functions $y(t)$ and $k(t)$ are clipped to y_{max} which obviouslydistorts their shape. In Figure 4.4 on page 108, a burst of delta pulses increasingly excites the system and this excitation takes a comparatively long time to wear off. At the same time, the pulses themselves are perfectly preserved and amplified. They are also clearly high-pass filtered as made apparent by the downward spikes that help them stand out in noisier signals.

4.2.2 User Interfaces

Our first prototype platform, the Neo FreeRunner, had no graphical user interface whatsoever. The transmitting and receiving IP addresses were hard-coded into the Python code that ran on both phones.

The Palm Pre implementation had the user interface that can be seen in Figure 4.5 on page 109, although it was mostly static and the contact list could not be changed at runtime. It served to demonstrate the concept and the interaction indicated in Figure 4.5(b) on page 109. The application shows a short activity level history from all available contacts. The contact overview serves as a visual representation of the users' accelerations. Thereby a user can quickly see which contacts show interesting activity patterns at any time. The visual overview also serves as a *chat room* in which all users can simultaneously see and react to each other's activity patterns. It would also be easy to discern similar acceleration patterns arising from joint activities (e. g. sitting in

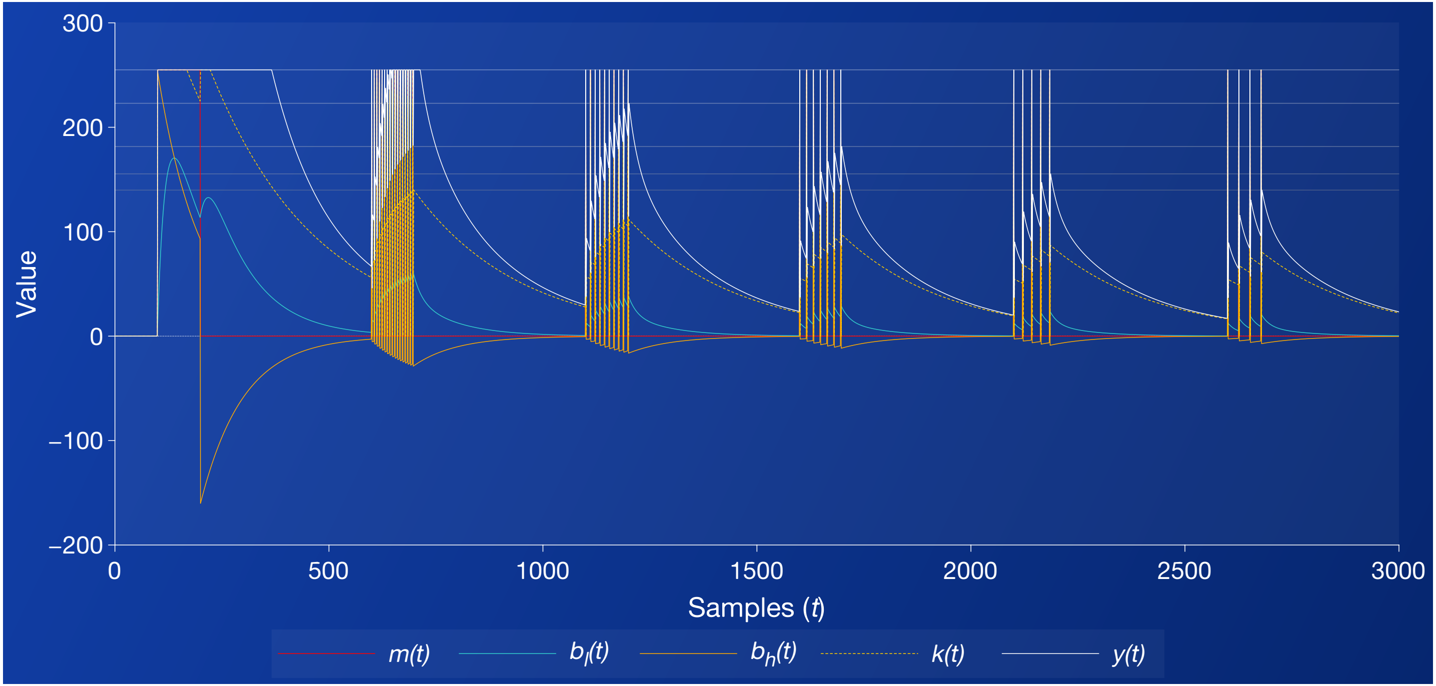

Figure 4.6. Filter response to bursts of Dirac deltas with varying frequency. The first (degenerated) burst gets the maximum input value $m(t) = 255$ at every sample; the delay between delta pulses increases by 5 samples for every subsequent burst. The time between beginnings of bursts is 500 samples. The maximum length of a burst is 100 samples but bursts can be shorter when 100 is not a multiple of the pulse interval. The horizontal lines show the non-spiking maximum of the output signal $y(t)$ for each burst. It can be seen that the output value depends on the number and frequency of previous delta pulses. Linear interpolation has been applied to reduce line overlap and because – at most resolutions – the spikes will appear more realistic like this.

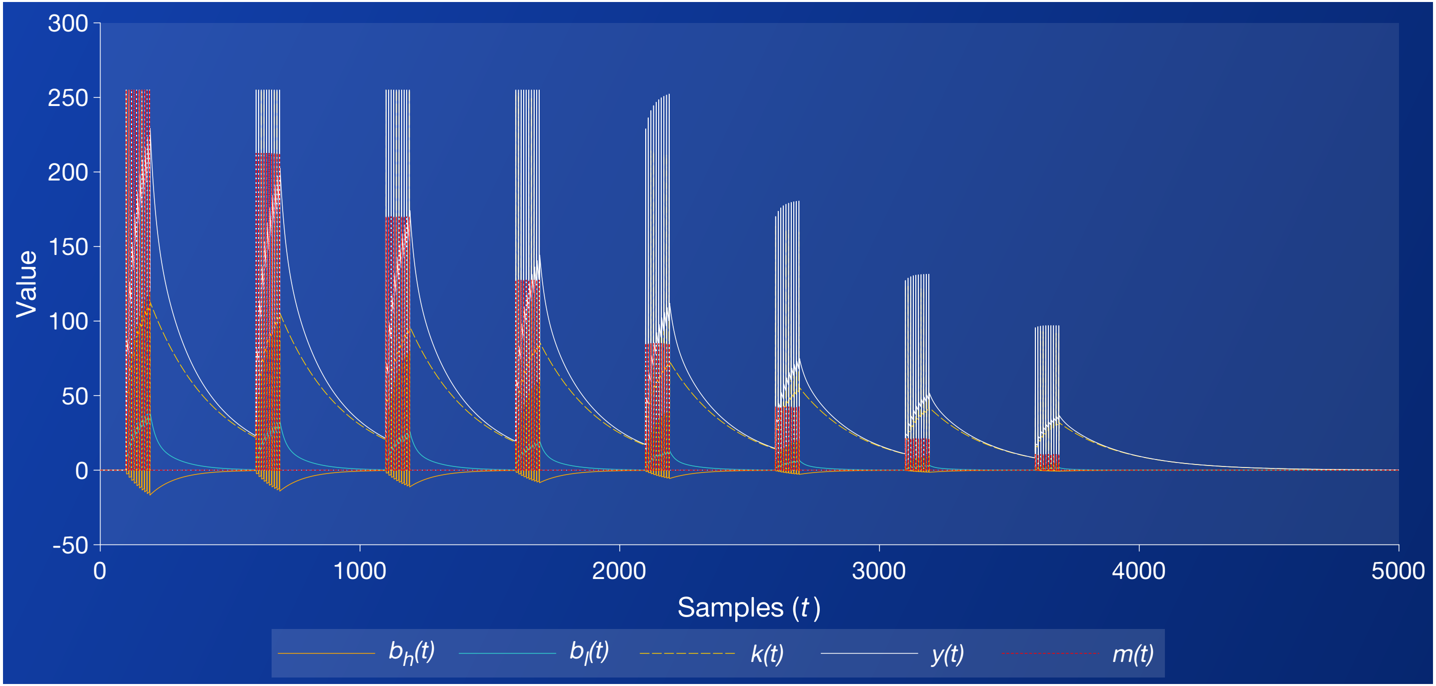

Figure 4.7. Filter response to bursts of Dirac deltas with varying intensity. All bursts consist of 10 Dirac deltas with a frequency of 1 pulse every 11 samples. With each burst, the amplitude of the pulses in the input signal $m(t)$ is decreased by one sixth of the initial value and is then halved for the last two bursts. This plot shows the dependence of the low-pass activation and to a lesser degree the spiking height in $y(t)$ on the input intensity. Please note that contrary to other plots, $m(t)$ was moved to the foreground to see it more clearly. Linear interpolation has been applied to reduce line overlap and because at most resolutions, the spikes will appear more realistic like this.

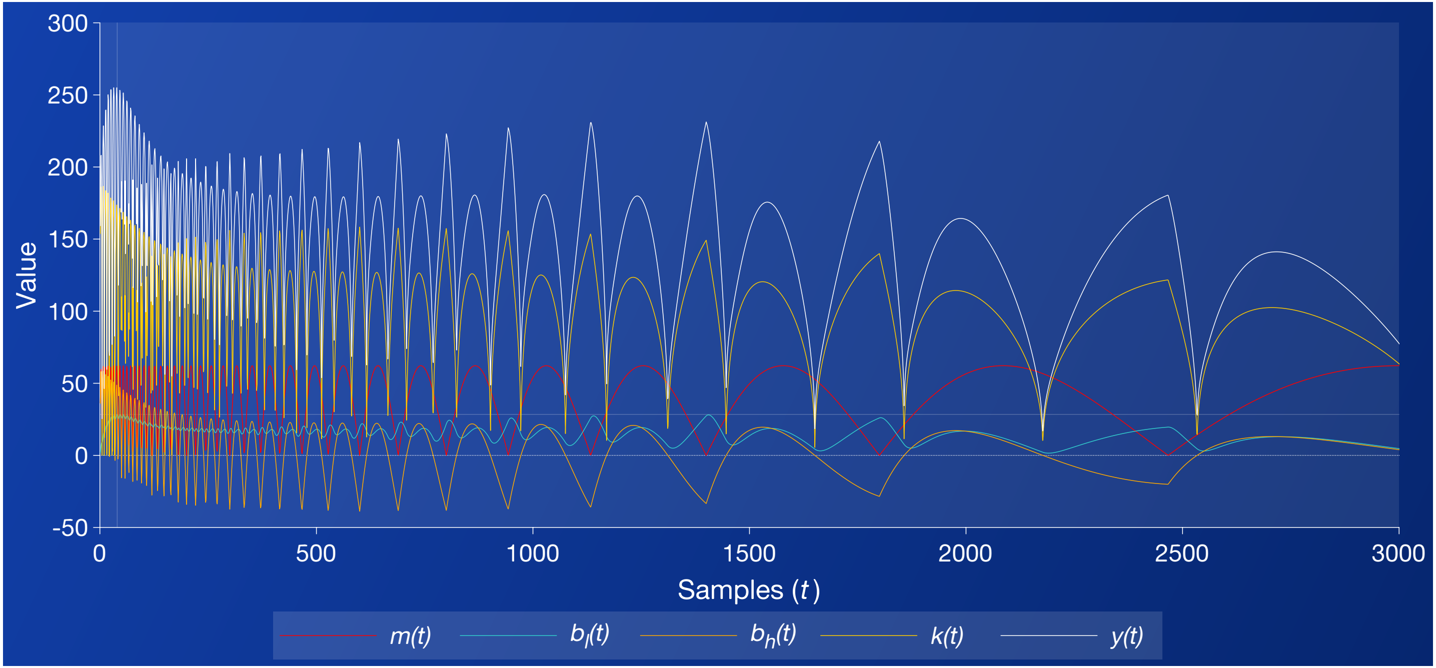

Figure 4.8. Filter response to a rectified sine wave of continuously decreasing frequency. Note especially how the input signal $m(t) = a\left|\sin\left(2t\pi\left(10 + \frac{t}{20}\right)^{-1}\right)\right|$, where a is a gain factor, initially puts enough energy into the system to make the amplitude of the output $y(t)$ quickly build up. This is driven by the low-pass system $b_l(t)$ whose curve shows the same hump as $y(t)$. When the frequency then drops below a threshold (in this case at about a wavelength of 12 samples) the low-pass system loses its energy. Also notice how each each cycle of the input signal causes two cycles of the output signal (one per edge), with the abrupt change in direction because of the rectification being imparted as a peak to the output as well.

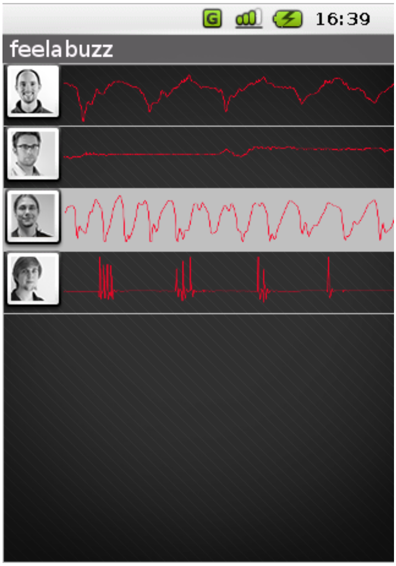

Figure 4.9. Initial concept of a FEELABUZZ Android app.

the same accelerating vehicle or doing sport together). An additional audification [119] of the current history upon selecting it might also be useful for an eyes-free interaction. Preliminary tests indicate that such acceleration data can be quite characteristic when played back as a short sound snippet, too. The history can also function as an activity footprint, although surely this raises privacy concerns which need to be addressed.

Some of these ideas have been realised in the Android implementation. Its basic design is the same but it features a full user management. The users can also use the contact list to configure their privacy settings. They can for example decide to which contacts they want to transmit their activity. The transmission of a random baseline activity that would not reveal that activity transmission has been disabled has been discussed but it has been decided that this would undermine the

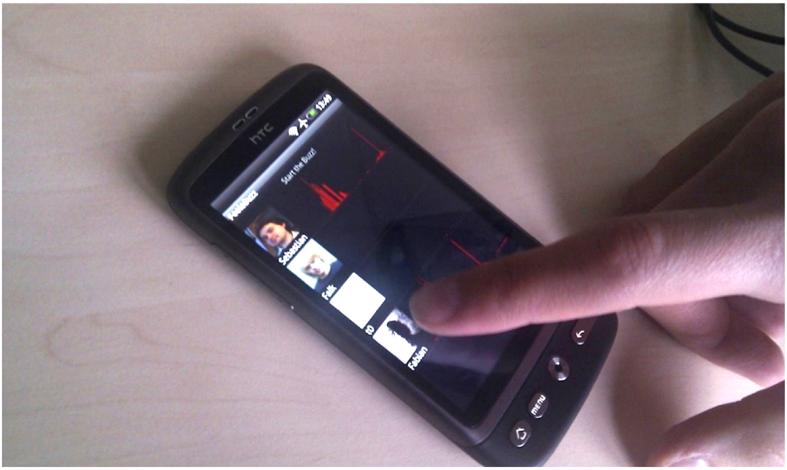

Figure 4.10. Finished Android app in use, showing the contact list screen. Photo courtesy of Falk Olaf Altheide, Sebastian Kahl and Fabian Vöhl.

trustworthiness of the system too much.

For the study described in Section 6.3.2, we went back to a much simpler GUI and a hard-coded phone-to-phone link. This later app just showed a static image of a mobile phone to indicate that the phone was still running. The screen was maximally dimmed to conserve battery life. Showing anything at all turned out to be necessary because we were forced to turn off all interactivity of the phone, including touch screen and power button so that the app could not be accidentally switched off while the phone was in the pocket. Having a phone with a black screen that does not react to any user input at all is essentially indistinguishable from a phone that is switched off or out of battery. This irritated even us during our test runs so we introduced the static display to give our participants at least a rudimentary status indicator.

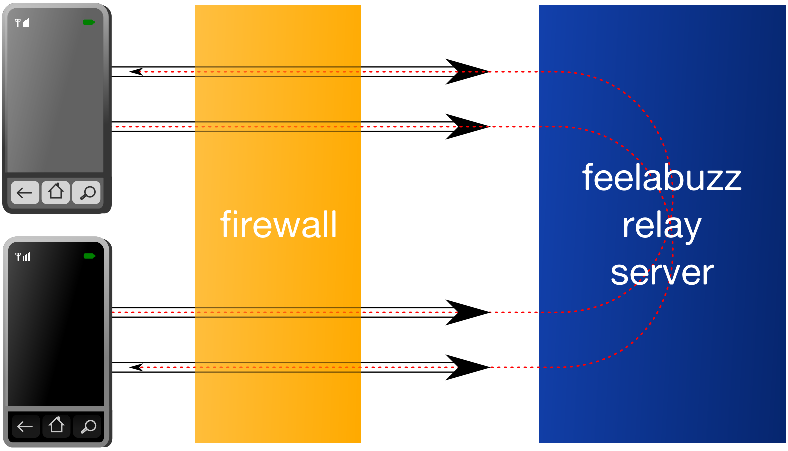

Figure 4.11. Network connection schema for the second iteration of the WebOS implementation. Each phone opens two connections to the server, using one for sending and one for receiving. The firewall prohibits connections from phone to phone and from server to phone, ruling out easier variants.

4.2.3 Technology

Both the Openmoko and the WebOS implementations are written in Python and are restricted to one-to-one connections. Both initially relied on Open Sound Control (OSC) [582] over UDP where each sensor value was transmitted as soon as it was recorded and processed. This was well-suited for demonstrations within one room and network as these implementations are very low-latency but also prone to packet loss and latency jitter. The software coped with the undetectability of a connection loss by cutting the vibration strength in half for every 500 ms without any new packets.

A change in university wireless network policy made this direct point to point UDP connection no longer an option and we reimplemented the network module for the WebOS application. We experimented for a while with a Point-to-Point Protocol (PPTP) [208] VPN to circumvent these new firewall restrictions but had to give this up eventually due to a buggy PPTP implementation in WebOS.

We then opted for a stateful connection relayed by a server to break through our university network's firewall, using only outbound TCP connections because inbound connections were also not allowed (cf. Figure 4.11 on page 116).[6] Because of the increased overhead of TCP, sending values individually was no longer practical and so we introduced buffering. The buffering brings with it some necessary design decisions, especially in unreliable networks, of how to keep latency low, handle disconnects and irregularities in the update rate, while at the same time discarding as few samples as possible.

Figure 4.13 on page 119 shows a recording of the actual number of data points remaining in the input buffer on the receiving end. When setting the accelerometer update frequency to "as quickly as possible", we measured an average sensor update frequency of 38 Hz so this is what is assumed on the receiving end. We did not want to send a timestamp with each data point for reasons of bandwidth and implementation complexity. However, the update rate on the receiving end is dynamically adapted depending on the amount of data received in order to keep the buffer within a certain fill level corridor. In Figure 4.12 on the following page it can be seen that if more than one second worth of data is available, the wait time between updates is linearly decreased until it reaches zero at two seconds worth of data. If more than three seconds of data are available, the oldest data points are silently discarded. The latter case happens especially right after a connection is established or re-established. The relay server had time to buffer a lot of data in these cases and discarding most of it is necessary in order for the time delay not to become too large.

Figure 4.13 on page 119 shows that these strategies seem to be effective in keeping the buffer length in a certain corridor, despite a slight jitter in packet reception times. Subjectively, the variations in

[6] Note that TCP connections allow bidirectional data exchange once they are established.

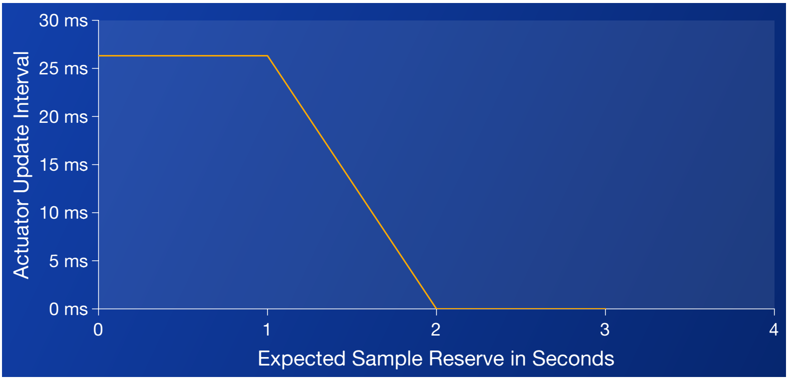

Figure 4.12. Time slept between updates by the thread that updates the vibration strength from the pool of received values, depending on the number of samples left in the queue (here shown in seconds as determined by the nominal update rate, not number of samples). An update rate of 38 Hz is used here but the qualitative behaviour is always the same: up to a buffer of one second, the sleep time is determined by the nominal update rate, ignoring processing time. With a growing queue, the sleep time drops linearly until samples are used up as quickly as possible once the buffer reaches two seconds. Beyond a buffer of three seconds, samples are silently dropped from the queue. See Figure 4.13 on the facing page for the effect of the behaviour described above.

update frequency were not noticeable to us.

Employing these strategies to circumvent the university firewall came in handy when we decided to use 3G UMTS in the user study described in Section 6.3.2 instead of WLAN because of its much greater reach and issues with reconnects due to WLAN roaming that caused frequent interruptions in connectivity between the devices [171].

The Android [185] application was a complete reimplementation in Java, developed as an undergraduate project supervised by my colleague René Tünnermann and myself. Unfortunately, the tactile rendering did not reach the same maturity as with the other prototypes and could

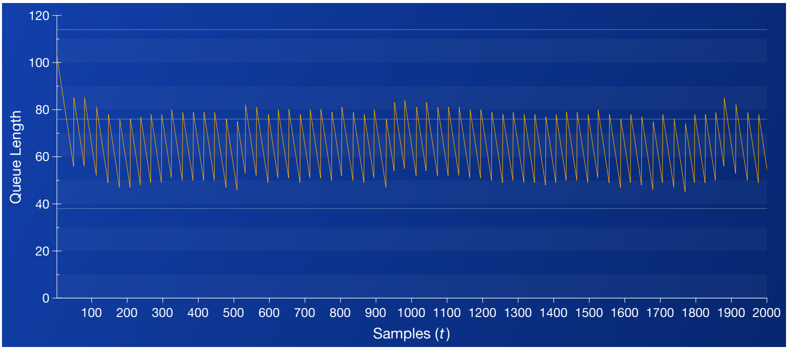

Figure 4.13. The effect of the delay variation shown and described in Figure 4.12 on page 118. Roughly every 38 samples, a new set of 38 samples is received. Initially, a much larger amount is received, most of which is discarded right away. Note how the buffer length is kept in a narrow corridor and irregularities in reception are handled well. Also note the linear look of the segments, even though the behaviour of the first derivative is non-linear due to the constant regulation of the delay.

therefore not be used for studies, even though the greater availability of the hardware would have made this an appealing option.

The Android implementation uses XMPP [254] over TCP as the transport layer. Also, sensor values are bundled before they are sent away, which increases latency but also reduces network load, which is especially important regarding the move from the stateless UDP to TCP with its acknowledgement packets.

XMPP already provides many of the features we need such as contact management, session management and dynamically connecting phones from arbitrary networks. The XMPP extension XEP60 [427] is used to provides a publish-subscribe pattern and server-sided buffering of recent activity which is needed when freshly tuning in to a contact or rendering the overview visualisations from scratch.

Regarding the vibration output, Android has the special difficulty that its API does not provide a call to control the strength of the vibration motor. The only call that can be used is `android.os.Vibrator.vibrate(pattern, repeat)` where `pattern` is an array describing the pattern of the vibration in the following way: the first value is the number of milliseconds for which to activate the vibration motor; the second value is the number of milliseconds of inactivity following that. The third value then describes a period of activity again and so forth. This makes it possible to emulate a command to control the vibration strength using a crude pulse-width modulation. Although the switching frequency for this is very low with only 100 Hz, the inherent inertia of the vibration motor will low-pass filter such a pattern into the semblance of a variable-strength vibration feedback.

4.3 Evaluation

We did a preliminary evaluation whose primary goal was to determine how well basic activity patterns can be discerned and recognised with the given hardware and mapping. However, we also used the occasion to get a first impression of the success of the affectiveness and calmness aspects that can be expected from our system.

For a study that looks at actual copresence through longer-term use, see Section 6.3.2.

4.3.1 Method

Ten participants took part in the study, 5 male and 5 female. The participants went through the study in pairs who were known to each other. The participants' phones were bidirectionally transmitting the acceleration data. As the first step of each trial, the general idea and basic properties of the acceleration-vibration mapping were explained to the participants. Each participant was then given the opportunity to familiarise him- or herself with both phones at the same time to get a better first impression of the mapping. When they both felt familiar with the system, they split up the phones so that both participants had one of them. They were again asked to explore the system until feeling familiar with it. They were then explained the following procedure.

The two participants were separated so that they could no longer see or hear each other. One of them was asked to perform one of three activities while wearing the telephone in their pocket: resting, walking or running. The other participant was instructed to guess, on the basis of the vibration they felt, which of these activities was being performed, holding the telephone in their hand. We asked the participants to guess as often as necessary. Auditory communication between the separated participants was never direct but mediated through the experimenters who were using a telephone line. This step was repeated ten times before the roles were switched between the two participants. The schedule of activities each participant had to perform was randomly generated in advance and different for each participant. The study was conducted using the Neo FreeRunner prototype.

Finally, the participants were asked to fill in a questionnaire. The questionnaire we used is based on the Computer System Usability Questionnaire (CSUQ) by Lewis [329]. We removed or adapted questions that did not make sense in our scenario and ended up with 12 multiple-choice questions using a 7-point Likert scale (cf. Figure B.5 on page 444). We also added six free-response questions.

4.3.2 Results

The results of the activity classification can be seen in Table 4.2 on the next page as a confusion matrix. All four misclassifications occurred between the classes “running” and “walking” and only when a participant was first confronted with one of these activities (i. e. it was the first activity to be performed or it was preceded by “resting”).

The most favourably answered items were “It was simple to use this system.” and “It was easy to learn to use this system.”, both of which were “strongly agree”d upon by all participants (average 1.0). The items that scored worst were “I believe I would use this system on a regular basis.” with an average of 3.7 (cf. Figure 4.16 on page 126)

		actual activity		
		resting	walking	running
guess	resting	**35**	0	0
	walking	0	**29**	2
	running	0	2	**32**

Table 4.2. Class confusion matrix of the participants' activity recognition. The overall misclassification rate was 4%, all examples of which occurred when the classes *walking* or *running* first occurred for a given participant.

and "This system has all the functions and capabilities I expect it to have." with an average of 3.714.

Figure 4.16 on page 126 shows the responses to four of the questions as histograms. It includes a question "This system created a sense of connectedness" which we added to the questionnaire to get a first estimate regarding the affectiveness of the system and "I believe I would use this system on a regular basis" which was intended to hint at the unobtrusiveness of FEELABUZZ.

We regard the responses to the affectiveness question as obviously very promising and those to the calmness question as somewhat lacking although not utterly hopeless.

It must be emphasised, though, that this study did not include any reliable measure of either the affectiveness or the calmness aspect of FEELABUZZ. The study described in Section 6.3.2 was conducted to make more reliable claims regarding these questions.

4.3.3 Outlook

One of the main obstacles towards an application of FEELABUZZ beyond the lab is the restricted battery run time. We are going to take a brief look at possible angles for optimisation.

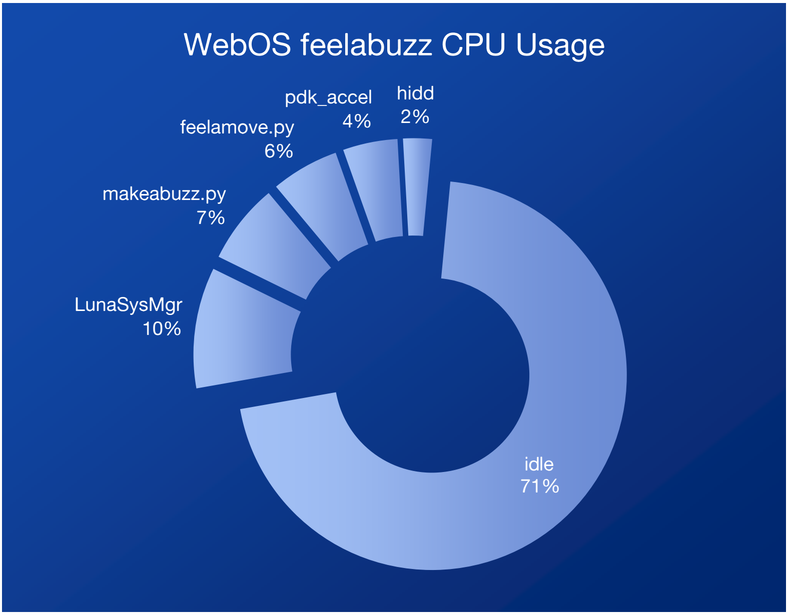

Figure 4.14. CPU Usage for FEELABUZZ, averaged over three 30 s measurements. Values are only approximations because of slight `top`[a] measurement inconsistencies.

[a]http://en.wikipedia.org/w/index.php?title=Top_(software)&oldid=557693140

On the Palm Pre platform, we measured an average ARMv7 Processor usage[7] of about 26 % while FEELABUZZ was running.[8] As can be seen

[7]The CPU frequency could scale from 125 MHz to 1 GHz but remained at 125 MHz while running FEELABUZZ once we had adapted the CPU governor setting to *conservative* and the upscaling threshold to 80 %.

[8]The attentive reader will find that this percentage and the idle percentage from Figure 4.14 do not add up to 100. This is due to a mismatch between the measured sum of CPU usage of all processes and the sum of user and system CPU usage over the same time interval. These measurement inaccuracies caused the overall CPU usage to overshoot 100 by a couple of points. Because of this comparatively small effect, we chose not to investigate

in Figure 4.14 on page 123, the `LunaSysMgr`, which is the WebOS process that does the main app execution, is the process that uses the most CPU. It is followed by the two Python processes that are responsible for processing and sending of the sensor data (`feelamove.py`) and the reception and actuator control (`makeabuzz.py`) which take up about an equal share of the CPU. This latter fact is somewhat surprising since all the processing logic described in Section 4.2.1 is done in `feelamove.py` without this showing in the CPU usage. So even though the signal processing could surely be optimised on many levels, it remains to be seen if this significantly reduces the processing time needed. Obviously reimplementing both Python scripts in a faster language would almost certainly reduce the CPU usage.

While `pdk_accel` – the process that reads out the accelerometer and which is already written in C – causes about 69 wakeups per second and thus many more than the sampling rate would make one expect, these are only about 10 % of the total wakeups and therefore the optimisation potential is probably limited at this point.

The vast majority of wakeups are caused by generic general purpose timers and the OMAP[9] interrupt. It is possible that switching off the screen could not only safe energy in itself but also reduce the load caused by the `LunaSysMgr` or the OMAP device but, as stated above, we found that completely switching off the screen, combined with the inability to interact with the phone, was highly confusing. We therefore maximally dim the screen and only show a static image but the screen never switches off. Restoring interactivity with the device while making sure that the application cannot be switched off accidentally would therefore have a high potential for prolonging battery life.

Reducing the number of wakeups might be an important step towards a less power-hungry app. We suspect this because the overall CPU usage is not unreasonably high, even though it might surely be improved. However, Figure 4.15 on the facing page shows that the CPU only resides in its deepest sleep state a quarter of the time and

the reasons for this any further.

[9] http://en.wikipedia.org/w/index.php?title=OMAP&oldid=558389536

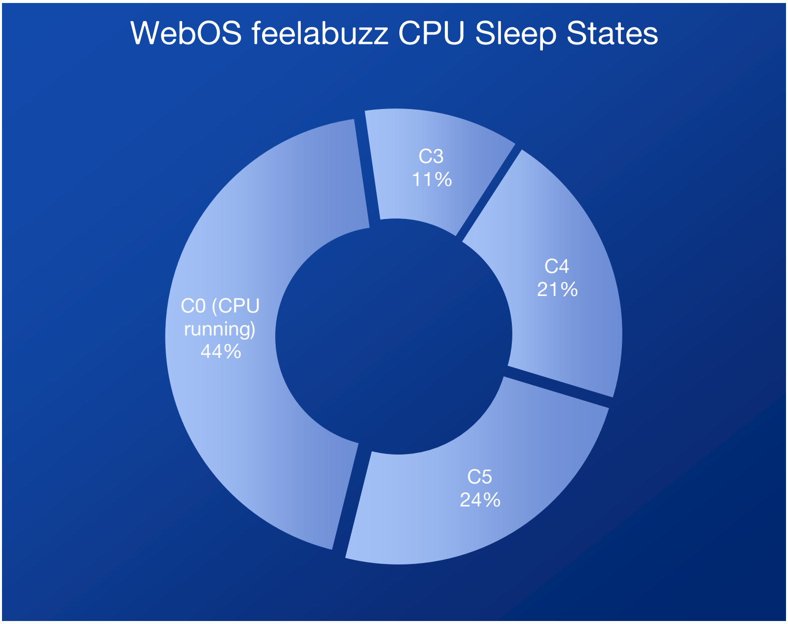

Figure 4.15. CPU sleep states for FEELABUZZ, measured over the course of a 300 s interval. C5 is the deepest sleep state. C1 and C2 were never used.

is actually running in C0 almost half of the time.

Reducing the power drain by display and CPU probably will only go so far, though. Real-time data processing, continuous sensor readings and wireless activity, and using the vibration motor a significant fraction of the time are all inherently power-hungry tasks on current-day hardware. For the future, however, one can with some confidence expect an increase in the prevalence of sensor coprocessors such as the Apple M7 [561]. Then, the CPU could be put to sleep for most of the time and together with an optimised network protocol that only actually transmits data when the vibration motor in the remote device would be active, a significantly extended battery life could be expected.

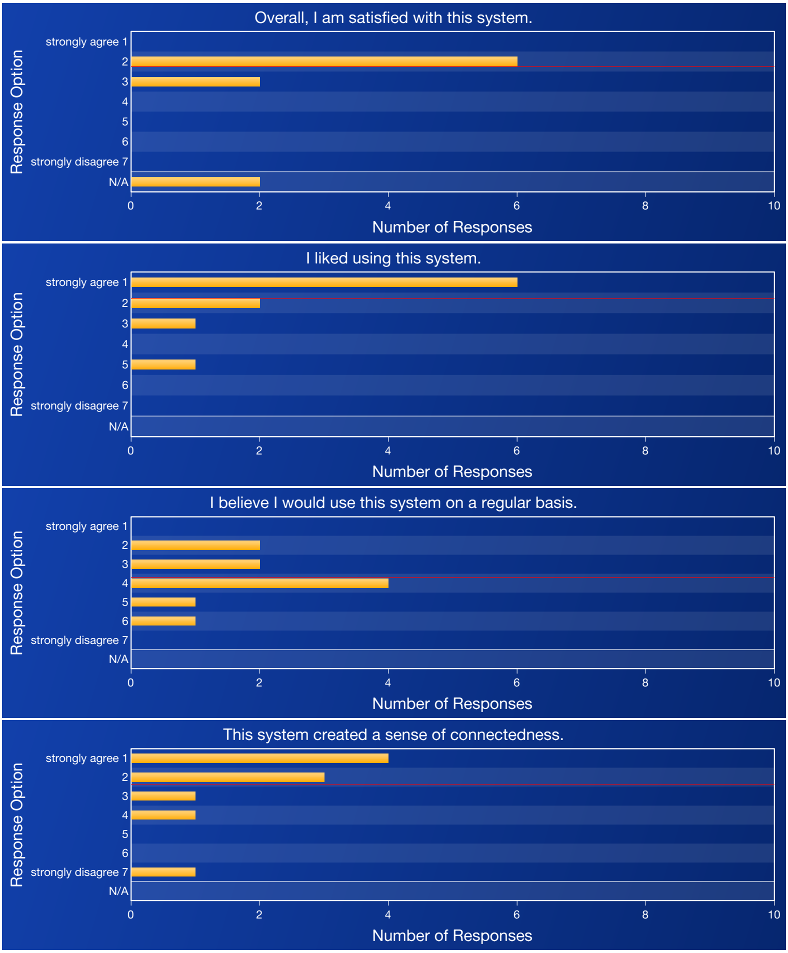

Figure 4.16. Responses to four of the questions in the questionnaire. The average values (also indicated by the red horizontal line in each plot) for these questions from top to bottom are 2.25, 1.8, 3.7 and 2.4.

5 upstairs – Acoustic Adjacency Simulation of Remote Rooms

5.1 Introduction

The idea for upstairs is directly related to the scenario presented in Section 1.1 of two people sharing a space and the notion of an *adjustable wall* (cf. Section 1.3) following from this scenario. What upstairs does actually *is* to simulate a wall – or more precisely a ceiling – that connects two rooms as if they were adjacent. There has been extensive research in the field of media spaces (cf. Section 2.3) where the goal was to connect two remote spaces as if they were one single space [45, 143, 270]. Unobtrusiveness was paid little attention to, though. Maybe for this reason, there was also little notion of not only connecting spaces with technical means but also *separating* them at the same time.

As illustrated roughly in Figure 5.1 on the next page and more detailedly in Figure 5.12 on page 140, upstairs connects two rooms by virtually stacking them mutually atop each other. It will appear to Person A that Person B lives upstairs and vice versa.[1]

Anecdotal evidence of people getting a good grasp of the daily routine of their upstairs neighbours even without other forms of communication was part of the inspiration for upstairs. We still believe, however, that complementarity as per Section 3.1.2 is crucial to make

[1] Note that the sound is always coming from above – and *not* from above for one user and from below for the other as would be the case with the real physical analog of two actual rooms one atop the other.

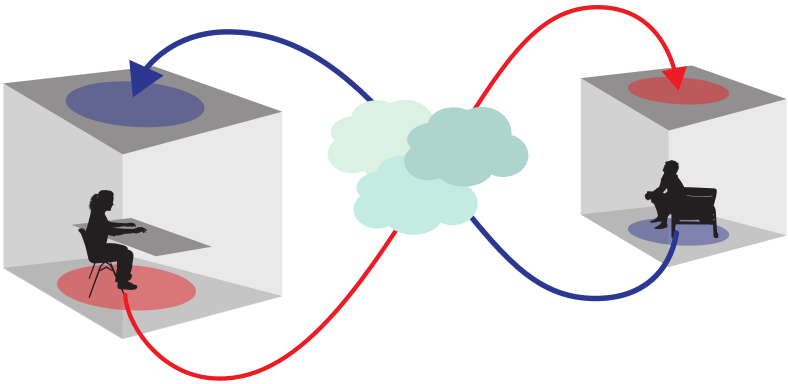

Figure 5.1. Schema of two rooms being connected via upstairs. The setup for each individual room can be seen in Figure 5.5 on page 133. Picture by Till Bovermann

real use of a system that simulates the noises diffused through walls because we do not think that without context and background knowledge there can be any long-term interest in the other person.

In fact, during initiatory discussions with colleagues and people from other backgrounds, the concern was frequently raised that people were already fairly annoyed by the upstairs neighbours they actually had and wondered why anyone would want such a system. We believe, though, that it is indeed the lack of interest and the lack of context information about the actual neighbours that are responsible for the annoyance in the first place.[2] It is a well-known phenomenon that noise perception depends largely on the personal feelings towards the source of the noise [271, 396, 421, 527]. If the upstairs neighbours were people they cared about *and* they knew enough about to deduce useful information about their activities from the sparse information

[2]Especially since the circumstance that neighbourhood relationships are prone to conflict has been true for a very long time as Cockayne [84] showed at least for Britain – probably without loss of generality though.

that the signal holds per se, then they would appreciate the auditory link. The positive first reactions we received for our example SCS in Section 6.5.1 and especially the cognitive load that was significantly less than what people anticipated beforehand could be interpreted as supporting this hypothesis.

5.2 Hardware Setup

Transmitting footsteps and other contact sounds is a four-step process: the sound must be 1) captured, 2) filtered, 3) transferred over the network and 4) eventually played back (cf. Figures 5.5 and 5.4 on page 133).[3]

We experimented with different kinds of floors and flooring materials. Hard concrete or stone floor (cf. Figure 5.6 on page 134), soft vinyl flooring (cf. Figure 5.2(a) on page 131) and the very flexible and elastic floor of a shipping container office transmitted footsteps very poorly: footsteps could only be heard up to 40 cm away from the microphone for concrete and even less for the elastic floor. Carpeting, even thin one, even worsened the situation considerably. We found wooden floor to work best; the only real loss of energy seemed to occur at plank boundaries (cf. Figures 5.7 on page 135 and 5.2(b) on page 131). Laminate flooring also worked quite satisfactory (cf. Figure 5.6 on page 134). All in all, we observed that the perceived loudness of the footsteps in the room itself is not a measure for the amount of acoustic energy that is transmitted to and by the floor.

Common means to insulate against noise transmission to adjacent rooms have little effect on the walking noise within the room[4] [174] or on the effectiveness of our sound recordings with contact microphones. In fact, when we used a sheet of felt between the laminate and the concrete floor (cf. Figure 5.6 on page 134), even soft movements without shoes became audible and sounded surprisingly natural.

[3] For additional photos and a detailed technical description see http://tai-studio.org/?page_id=808

[4] In German, the former is called *Trittschall*, the latter *Gehschall*. The structure-borne sound upstairs uses does not directly correspond to either, though.

The capturing of the contact sound was done with capacitive contact microphones (depicted in Figure 5.2(a) on the facing page). We used AKG C411 microphones[5] because of their qualities in both capturing performance – especially a low signal-to-noise ratio – and physical robustness. Their capacitive recording principle required phantom powered microphone input channels which in our earlier test setups we provided with RME[6] or MOTU[7] interfaces that used IEEE 1394 to connect to the computer that did the further processing. Due to hardware availability, we used Yamaha MG 102 C mixers[8] (cf. Figure 5.10 on page 138) and an analogue sound input to the built-in audio adapters of the laptop computers that did the sound processing for the later setups (cf. Figure 5.11 on page 139) and for the user study (cf. Figure 5.9 on page 137).

We experimented with variations in the number of microphones for stereo or multichannel sound but found that in all our more conventional setups no spatial resolution was achieved. Only the setup shown in Figure 5.7 on page 135 was able to deliver convincing spatial audio. There, the two microphones were spaced over 3 m apart in an even larger room, allowing movement beyond the stereo base.[9]

To create the illusion of sound travelling through floor and ceiling, we used a band-pass filter[10] with cut-off frequencies at about 50 Hz and above 300 Hz.[11] The processing was done using the SuperCollider

[5] http://www.akg.com/pro/p/c411group

[6] http://www.rme-audio.de/en_index.php

[7] http://www.motu.com/audio

[8] http://uk.yamaha.com/en/products/music-production/mixers/mg_series_c_models/mg102c/

[9] Recordings from this stereo setup can be found e. g. at https://soundcloud.com/lfsaw/2011-walkingonwood02?in=lfsaw/sets/test-recordings-for-the-shared and https://soundcloud.com/lfsaw/2011-2?in=lfsaw/sets/test-recordings-for-the-shared. Even in this generous hallway setup, a relatively sharp transition from one stereo channel to the other can be perceived.

[10] more specifically the combination of SuperCollider's second-order Butterworth low-pass and high-pass filters

[11] These numbers were determined empirically. We had first tried a low-pass filter only but this did not sound satisfactory. Measuring the actual filter characteristics of different

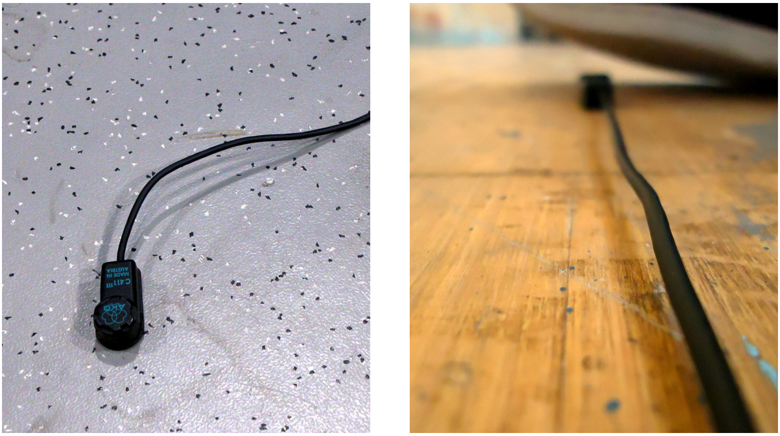

(a) Contact microphone on vinyl floor.

(b) Contact microphone on wooden floor. Photo by René Tünnermann

Figure 5.2. Contact microphone on different types of floor.

programming language.[12]

For the network transmission, we used the open source streaming software Darkice[13] and an Icecast2[14] server, configured for minimum latency (e. g. approximately 1 s in contrast to the usual 3–5 s). Since latency was only of moderate importance to us because it is important but not as crucial as e. g. in bidirectional speech transmission, we did not go to any greater lengths to optimise it (for example by using a low-latency audio codec – we used MP3 instead for maximum

ceilings could be a future step towards an even more convincing sound.

[12] http://www.audiosynth.com/

[13] http://code.google.com/p/darkice/

[14] http://www.icecast.org/

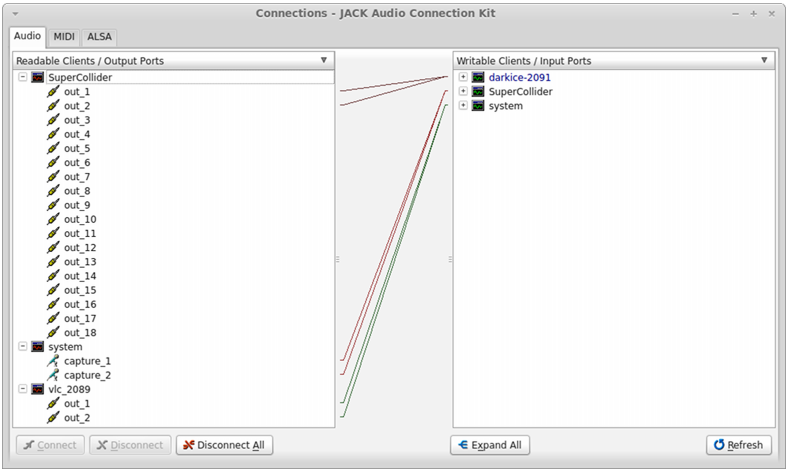

Figure 5.3. Screenshot of the Jack connections on the laptops used during the study.

compatibility). The Jack[15] patching to connect the different software components is illustrated in Figure 5.8 on page 136 and directly shown in Figure 5.3.

We played back the resulting HTTP audio stream on the receiving end using a software audio player, again with a cache size configured for minimum workable latency (although this minimum obviously depends on the network quality). As the type of loudspeakers plays an important role for the perceived sound due to the two factors frequency response and radiating properties, we adapted filter parameters accordingly. This turned out to be a crucial factor in the system that needs to be tweaked depending on the specific installation location. By facing the speakers upwards, the sound was distributed towards the ceiling so that the first reflections are heard most prominently. The resulting large emitting

[15]http://jackaudio.org/

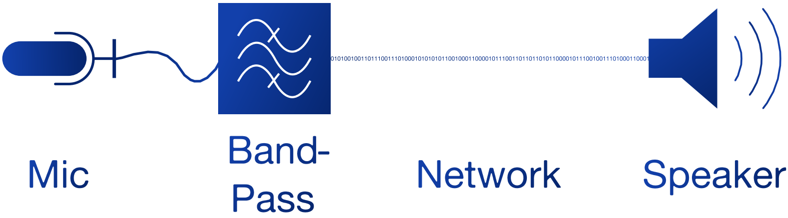

Figure 5.4. Essential steps of the upstairs data flow.

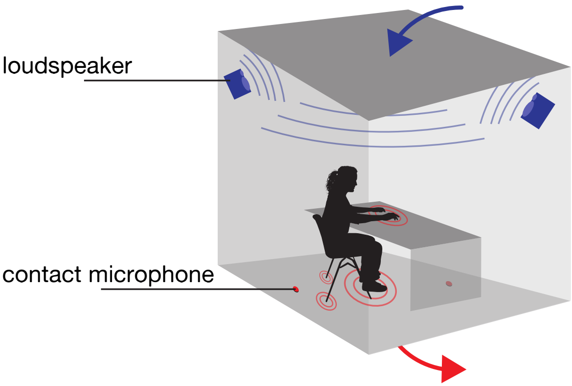

Figure 5.5. Schema of the upstairs setup in a single room. Two rooms are then connected as shown in Figure 5.1 on page 128. Picture by Till Bovermann

angle helped to increase the illusion of the whole ceiling giving off the sound. The wider the emitting angle of the speakers to begin with, though, the better the overall result that was achievable.

Figure 5.6. Construction of one of the upstairs test setups. The partial felt padding can be seen. Photo by René Tünnermann

Figure 5.7. Stereo upstairs hallway test setup from above.

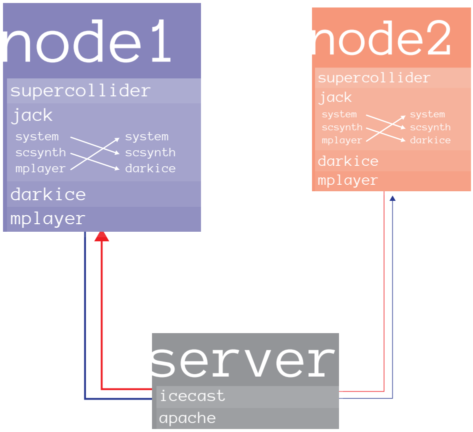

Figure 5.8. Illustration of the connections between the software components used. Picture by Till Bovermann

5.3 First Impressions

To spark some fruitful discussion beforehand, we confronted people from various backgrounds with the system, ranging from interaction design over interactive media art to photography and sound design. They were asked for their general impression of the system, their opinions towards it and how they think it could be improved.

In these discussions, one person reported that, although he rarely meets his upstairs neighbours, his knowledge about the daily routine of them is more detailedly than that of his friends. Another person suspected that, when having this system installed in his home, he would be more self-conscious about coming home late.

A question that arose from the discussion was on how to differentiate between the system-induced auditory illusion and the sounds actually originating from the neighbours living above the user. Also, people mentioned that many environments already suffer from a lot of

Figure 5.9. Two setups of upstairs by participants of our study (cf. Chapter 6). To get the speaker as close to the ceiling as possible, either a 165–350 cm speaker stand or available furniture was used. Photographed by the participants.

environmental noise and asked how we intended to deal with this.

Along the same lines, people emphasised the importance to add self-monitoring possibilities and to let users adjust the system in terms of overall amplification and filter parameters. Another suggestion that came up during the discussion was to augment only certain spots of places instead of covering the whole area.

These inputs and our own long-term test of the setup in our own offices over a couple of months were quite valuable for the user study we then conducted (cf. Chapter 6).

Figure 5.10. Mixer as used during the user study. The multitude of knobs and sockets had the potential to overwhelm less tech-savvy users. On the other hand, this way users did have the possibility to do some adjustments to filter settings if necessary.

Figure 5.11. Office setup of upstairs.

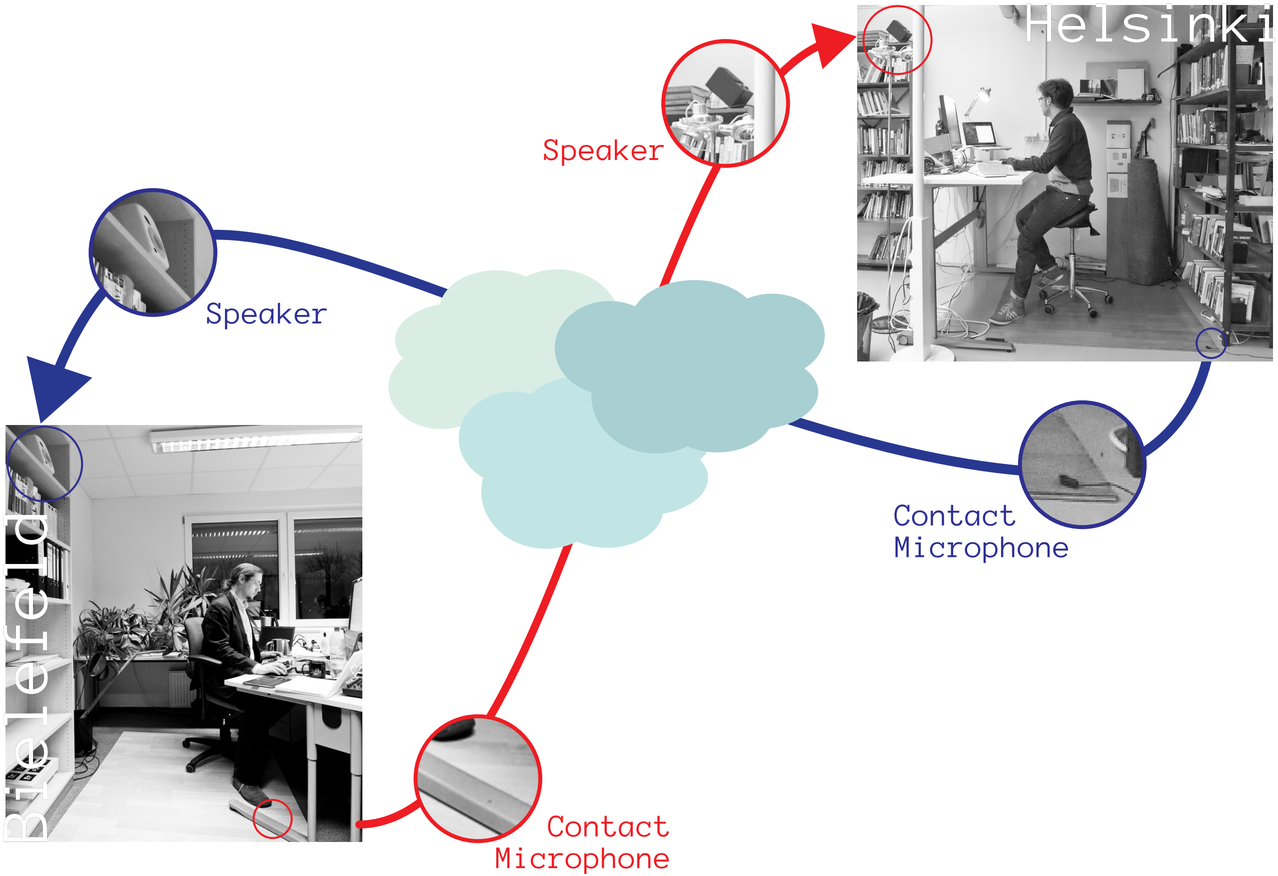

Figure 5.12. Upstairs setup connecting two offices in Helsinki and Bielefeld. Diagram and Helsinki photo by Till Bovermann

6 Exploring the Effects of Subliminal Copresence Systems

The systems upstairs and FEELABUZZ presented in the two preceding chapters are examples of subliminal copresence systems (SCS). In Chapters 1, 2 and 3, we postulated a number of effects SCS might have. We conducted a study where participants used FEELABUZZ for only a couple of minutes while performing mandated activities (cf. Section 4.3). That study was thus not designed to investigate how the system works on an emotional level, possibly affecting connectedness, presence or even relationship quality. It was also not designed to verify our important design goal of SCS having a low cognitive load so that they can be running over extended periods of time. For these reasons, we wanted to conduct a study whose participants had to use a given SCS for days instead of minutes. In the following, we are going to present the studies we conducted – including our expectations, method and results – and finish with a discussion of these results, even though the iterative nature of the exploratory part of our analysis blurs the line between method, results and discussion somewhat.

6.1 Expectations

We expect both our SCS to induce a high amount of copresence while maintaining a lower level of social presence. Accordingly, cognitive load should be low and people should find it convenient to have the system running continuously. Ideally, the effects of the SCS should occur without the participants even immediately noticing them or at least without them being distracted from other activities.

If telepresence is a prerequisite for copresence as some authors have said [69, 455, 498, 537] (cf. Chapter 2), one might expect telepresence to measure at least as high as copresence, even though nobody ever undertook any means to actually calibrate the different presence scales in relation to each other so this and other similar statements are only approximate and speculative.

As detailed in Section 1.2 and Section 2.5, SCS could have the potential to increase *or* to decrease relationship quality. However, given that transmitting movements of a partner seems like an insignificant thing compared to the noticeable alteration of the quality of a relationship, it is hard to be optimistic about any measurable and definitive change in relationship quality over the two-week study period. In any case, we expect relationship quality to be a factor that is particularly prone to variations that we cannot control for, so any fluctuations that we might see are hard to link back to the system itself, especially given the small sample size of the studies (cf. Section 6.3).

If connectedness is a better measure for the effectiveness of awareness systems than awareness itself as has been suggested (cf. Chapter 2) and upstairs and FEELABUZZ are effective awareness systems, then they should rank high on the Connectedness scale.

Among all the positively connotated scales present on the questionnaires we will present in the following section, we especially hope to find a low Cognitive Load and a low Threat to Privacy since this would correspond with our postulates from Section 3.1.2.

We would also like to get an indication towards the validity of the relations between the various concepts presented in Chapter 2 postulated by ourselves or other authors. This means especially that we would like to confirm or reject at least part of the overall picture shown by Figure 2.4 on page 35.

Generally however, it should be noted that both studies have very small sample sizes compared with the complexity of the questions raised and the tentativeness of the results can therefore not be overemphasised.

The upstairs study will be mainly exploratory and hypotheses-generating in nature, while the FEELABUZZ study will be confirmatory in as far as it is able to be.

6.2 Hypotheses

Putting the expectations from the previous section more formally, these are the hypotheses we want to test:

(H1) Copresence ratings are better than Social Presence ratings.

(H2) Copresence rates no better than Telepresence.

(H3) Relationship Quality before system use differs from afterwards.

(H4) Social Connectedness is better after system use than before.

(H5) Cognitive Load is low.

(H6) The perceived Threat to Privacy is low.

(H7) The measure of Copresence is better than the scale centre.

(H3) and (H4) can not be tested for FEELABUZZ since the questionnaire before system use did not comprise the necessary questions.

See Table 6.3 on page 150 for an overview, including the results from Sections 6.4.1.1 and 6.4.2.1.

6.3 Method

The two small-scale studies we conducted in order to get a first indication regarding the effect of upstairs on couples in long-distance relationships and to collect some user experience from people using our SCS upstairs and FEELABUZZ over a longer period of time were constructed similarly. The FEELABUZZ study was in essence a shorter version of the upstairs evaluation – by this allowing more people to take part.

Sample Statistic	Population Parameter	Description
$\bar{x}$	μ	Mean
SD	σ	Standard Deviation
Md	–	Median

Table 6.1. Frequently used symbols. Since there are no symbols for these concepts that are used consistently across disciplines, we define them in a way that should minimise the risk of confusion.

6.3.1 upstairs

We evaluated upstairs by giving two connected setups to couples living in a long-distance romantic relationship (LDRR) who installed the systems in their homes (cf. Figure 5.9 on page 137). There were three couples, making for $n = 6$ participants in total. Each couple was supposed to use the system for two weeks but in one of the couples, one partner moved to a new flat without the necessary wooden or laminate floor before this two-week period could be finished.

Participants were asked to complete two basic types of questionnaires before, during and after the study: a detailed one once a week (called Q_β, Q_γ and $Q_{\beta\gamma}$) and a simpler one every second day (called Q_α and Q'_α; all the questionnaires are shown in full in Appendix A.1). Before the start of the study, a baseline period took place in which participants were asked to answer the applicable questions from the simple questionnaire (Q_α).

The baseline period took place for all participants simultaneously but since there was only one set of hardware, the time between baseline and actual study period varied.

See Table 6.2 on the facing page and Figure 6.1 on page 146 for an overview of the study structure and for a detailed description; Section 2.4.1 describes the design of the questionnaires.

The participants were all heterosexual couples and two of the three were married. The participants were between 29 and 35 years old ($\bar{x} = 31.3\,\mathrm{a}$, $SD = 2.1\,\mathrm{a}$, $Md = 31\,\mathrm{a}$). The length of their relationships

Icon	Names	Number of Items	Measures/Influences
	Q_α, Q'_α (alternate days)	14, 30	inclusion of other in the self [13], copresence (aloneness, mutual awareness and closeness) [39, 401], social presence [494], affective costs and benefits [584], technical issues, free answers
	Q_β (before system use)	93	social desirability [518], relationship closeness [34], relationship quality [26], connectedness [550], affective costs and benefits expectations [584]
	Q_γ (half-time)	89	presence [498], copresence [39, 401], social presence [494], connectedness [550], affective costs and benefits [584]
	$Q_{\beta\gamma}$ (final)	128	presence [69, 498], copresence [39, 401], social presence [494], connectedness [550], affective costs and benefits [584], relationship closeness [34], relationship quality [26], system satisfaction

Table 6.2. Legend for Figure 6.1 on the next page. The icons are licensed under the GNU Lesser General Public License since being derived from the Oxygen icon set.

ranged between 5.3 and 12.8 years ($\bar{x} = 9.2$ a, $SD = 3.4$ a, $Md = 9.5$ a) while the part of the relationship that they themselves considered to be an LDRR[1] varied between 41 days and 3.5 years ($\bar{x} = 1.7$ a,

[1] Stafford [507] proposed to prefer such subjective self-reports over objective but

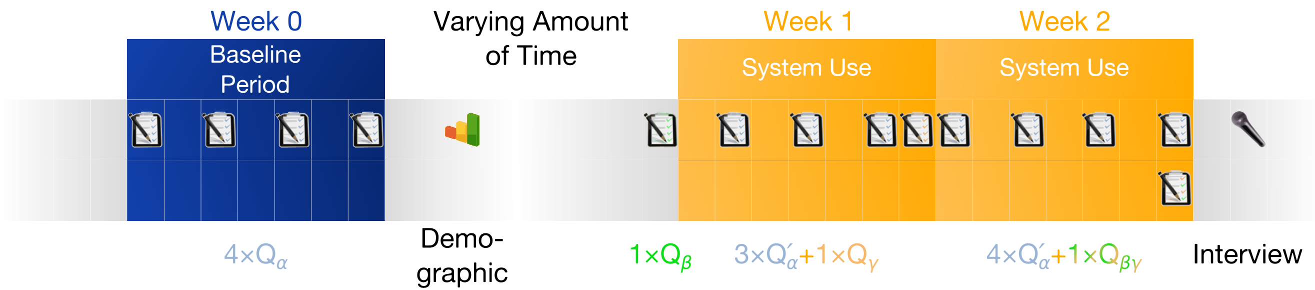

Figure 6.1. Schedule of the phases of the upstairs study and the questionnaires that were used. All participants had to complete the 14-item questionnaire Q_α every other day for one week as a baseline measurement. After a variable period of time, participants got the setup and before using it for the first time had to fill in another baseline questionnaire Q_β with 93 items which also included a social desirability scale [518]. They then used the system for two weeks, filling in Q'_α every other day which was an extended 30-item version of Q_α, a 89-item evaluation questionnaire Q_γ after a week and finally $Q_{\beta\gamma}$ which added one more item to Q_γ and combined it with Q_β (omitting the affective costs and benefits *expectations* from Q_β in favour of the actual affective costs and benefits from Q_γ) for a total of 128 items. A small 8-item questionnaire asking for some demographic data had to be completed after the baseline period. The questionnaire colour coding used in the icons is explained in Table 6.2 on page 145.

SD = 1.4 a, *Md* = 1.8 a). These self-reports were not exactly the same between partners but reasonably close (with an average standard deviation of 26 days and an average coefficient of variation[2] of .04). Regarding the separation-reunion cycle [507] at the time of the study, all participants said to meet their respective partner two to three times a month. Participants were also asked for their experience with computers and how much they would consider themselves to have a disposition for technology. All participants rated themselves within the two response options signifying the most experience and disposition for technology on a 7-point scale.[3]

Additionally to completing the $Q_{\beta\gamma}$ questionnaire (cf. Figure 6.1 on page 146) at the end of the usage period, an unstructured interview was held with each participant to also get subjective feedback, reflecting the exploratory stage of this research.

The questionnaires were anonymised by using code names,[4] and randomly permuting the data rows and stripping the timestamps from them before analysis.

The last couple were also given the opportunity to use the system longer than we asked them to do. This was part of the evaluation without the participants knowing this. We wanted to find out if the participants used upstairs only because we asked them to or if they liked the system so much that they used it for as long as they could.

---

arbitrary criteria in order to determine whether a relationship is long distant. As for objective measures, the participants lived between about 250 km and 760 km driving distance apart.

[2]The coefficient of variation is the standard deviation normalised by the mean, i. e. $CV = \frac{\sigma}{\mu}$.

[3]We actually required the non-local participants to have at least some technical background because setting up the prototype was not trivial (cf. Figure 5.9 on page 137 and Figure 5.10 on page 138).

[4]To avoid impersonal and easy to forget numbers, the participants were given a pool of fictional movie couples from which they could choose without the experimenter knowing the association between actual and fictional couple. Every proposed name was unique. This had the advantage that the gender was implicitly coded in the token, as well as the information which two partners belonged together. The questionnaire was programmed to recognise these names and adapt the gender-specific wordings accordingly.

We deemed this a more reliable method to assess the motivation of the participants than asking them whether they liked using the system or if they would use it if they had the opportunity to do so (which we also did). Unfortunately, the schedule did not allow for the other couples to be given this opportunity unsuspectingly.

6.3.2 FEELABUZZ

The studies for FEELABUZZ only went for two days per pair of participants and being in an LDRR was not a prerequisite for participating. Consequently, all questions assessing the relationship quality and similar measures were removed from the questionnaires (except for a single IOS item on the pre-study questionnaire) because any change here due to the system was deemed virtually impossible. There also were only two questionnaires: one very short questionnaire before using the system and a longer one after having used it (cf. Appendix A.2).

Two questions were additionally removed because they did not apply to FEELABUZZ, "I was disappointed when my partner wasn't there when I tried to contact him/her using the system" from the Unmet Expectations scale and "I had the impression that my partner actually lived upstairs" from the Telepresence scale.[5]

The study was completed by 7 couples, where one participant failed to complete the post-study questionnaire in a timely manner, leaving us with $n = 13$ participants.[6] There were 2 more couples that aborted the study because one of the partners found the vibrating phone too distracting. From those participants who completed the study and both questionnaires, 7 were female and 6 male. Their relationship had lasted between 5 months and 9 years ($\bar{x} = 3.74$ a, $SD = 3.15$ a, $Md = 2.5$ a).

[5] The latter question actually accidentally remained on the questionnaire and was ignored in the subsequent evaluation. It might be a troubling testament to the impression the long final questionnaire left with the participants that actually not a single one of them complained about this nonsensical item.

[6] We decided that we could still use the data from the partner that did fill in both questionnaires because we do not use any metric that associates partners with one another in this study.

Box 4: Remarks on the Analysis

The following statistical analysis was done using SPSS[a] version 20 and the high-precision spreadsheet [3] Gnumeric[b] in versions 1.10.17, 1.12.1 and 1.12.6.[c]

Tabular data is almost always augmented by an automatic table cell colouring. The cell value is mapped to a luminance–saturation scale in accordance with the findings of Rogowitz and Treinish [453]. The scale, direction and extreme values used are usually apparent from the kind of data displayed. Usually, this will be either a scale ranging from white to yellow for data in $[0, 1]$ (or $[-\infty, 1]$ for Cronbach's α, with the cell colour being white for any value ≤ 0) or vice versa (e. g. for p-values), or a scale from blue over white to yellow for data in $[-1, 1]$ or sometimes $[1, 7]$ for results on the 7-point scale that was used almost exclusively on the questionnaires (in this case, the neutral white point is at 4). The scales yellow–white–yellow and blue–white–blue also sometimes occur when the magnitude is the only measure for the value of an entry but direction is too interesting to discard the sign altogether (e. g. for correlations or factor loadings). When there is no natural or mathematical boundary for the content, the extremes used for the colour scales are given in the caption.

Unless otherwise noted, the questionnaires used 7-point scales ranging from "1 – I agree strongly" to "7 – I disagree strongly". Additionally, items were reversed in such a way that a small scale result represents a positive result for the system while a large value is a negative result. This means the common expectation that a high scale result means a high expression of its face value is only met for cost scales but not for benefit scales. See Table B.1 on page 432 for a full reference of which items were reversed.

We used an alpha level of .05 for all statistical tests.

[a] http://www-01.ibm.com/software/analytics/spss/

[b] http://projects.gnome.org/gnumeric/

[c] Almiron et al. [3] advise not to use spreadsheets for nonlinear regression and Monte Carlo experiments but found Gnumeric to deliver exact results for the kinds of calculations we used it for.

Hypothesis	Summary	upstairs	FEELABUZZ
(H1)	Copresence < Social Presence	✓	✓
(H2)	Copresence ≥ Telepresence	X	X
(H3)	Relationship quality changes	?	?
(H4)	Connectedness increases	?	?
(H5)	Cognitive Load < 4	✓	X
(H6)	Perceived Threat to Privacy < 4	✓	✓
(H7)	Copresence < 4	✓	✓

Table 6.3. Overview of upstairs hypotheses. ✓ means that the hypothesis can be confirmed,[a] X signifies that the relationship in the data is opposite to the one in the hypothesis, and ? means that no statistically significant statement on the hypothesis is possible.

[a]meaning that the null hypothesis assuming no change or difference was rejected with $p < .05$

All couples were heterosexual and unmarried. The participants' age ranged from 20 to 29 years ($\bar{x} = 25$ a, $SD = 2$ a, $Md = 25$ a). Self-assessed experience with computers was $\bar{x} = 5.85$, $SD = .99$ and disposition for technology was $\bar{x} = 6.31$, $SD = .85$ when given on a scale from 1 as the minimum and 7 as the maximum experience or disposition respectively.

Because the code name scheme for anonymisation needs coordination between all of the participants without the experimenter orchestrating this, it was impractical for this larger study. We therefore used conventional numerical tokens for the participants. Forgetting the token was less of a potential issue because of the much shorter time span.

6.4 Results

In the following we will, for each system separately, first test the hypotheses from Section 6.2 and then go into a more exploratory analysis of the results, seeing if there are any patterns that we did not expect. For FEELABUZZ, this part will be less extensive since we will focus on verifying the patterns that appeared in the upstairs data.

We will also take an especially close look at the Cognitive Load scale because it is the one scale that is not entirely or for the most part adapted from scales other authors already tested and analysed. We therefore want to analyse it in more detail than the other scales. It is also a relatively large scale and could profit from some reduction in the number of items.

A summary of the results of the hypothesis tests can be found in Table 6.3 on page 150.

6.4.1 upstairs

6.4.1.1 Hypothesis Testing

The Overall[7] Copresence ($\bar{x} = 3.26$, $SD = .59$) is significantly higher than the Overall Social Presence ($\bar{x} = 4.04$, $SD = .99$); paired one-tailed $t(5) = -3.32$, $p = .011$. Cohen's d indicates a large effect size ($d = -.96$) [85]. This confirms (H1) and the non-parametric Wilcoxon test corroborates this result, $Z = -1.99$, $p = .023$. Low values on these scales always signify a high presence (cf. Box 4 on page 149).

The Overall Copresence is significantly higher than Telepresence ($\bar{x} = 4.03$, $SD = 1.43$); paired one-tailed $t(5) = -2.52$, $p = .027$ with a large effect size ($d = -.84$), rejecting (H2) ($Z = -1.99$, $p = .023$).

Relationship Quality (H3) could not be compared before and after system use (cf. Section 6.4.1.2 below).

The Overall Social Connectedness before and after system use does not differ significantly overall or on any subscale (cf. Table 6.7 on page 162), so (H4) can not be confirmed.

The Cognitive Load self-reports ($\bar{x} = 3.38$, $SD = .33$, low values mean a low cognitive load) are significantly lower than the scale centre

[7]See Table 6.4 on page 155 for an overview of which scales are fused to form combined scales based on their common face value. See Figure 6.6 on page 160 for the internal consistencies of these combined scales.

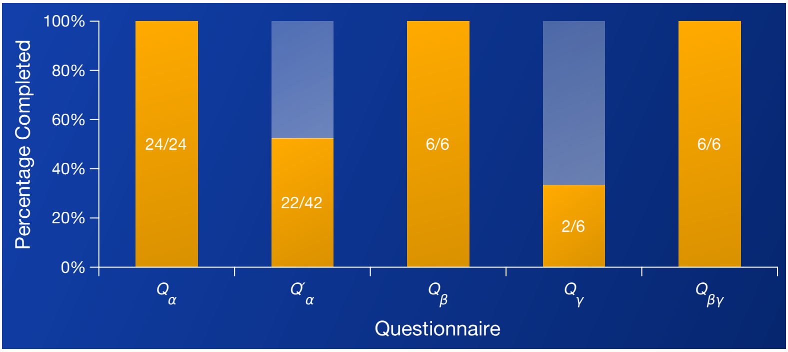

Figure 6.2. Response rates for the upstairs questionnaires. 100 % would have been the number of questionnaires completed by three couples running the study for the full three weeks each and filling in all questionnaires. This is the second number in the columns of the plot. The first number is the number of questionnaires actually completed. Questionnaires are either counted or not counted for this plot; there are no partially completed or counted questionnaires.

of $\mu = 4$, paired one-tailed $t(5) = -2.68$, $p = .022$ with a large effect size ($d = -1.69$), confirming (H5) ($Z = -2.00$, $p = .023$).

The perceived Threat to Privacy ($\bar{x} = 2.13$, $SD = 1.35$, low values mean a low perceived threat) is also significantly lower than the scale centre, paired one-tailed $t(5) = -3.93$, $p = .006$ with a large effect size ($d = -2.48$), confirming (H6) ($Z = -2.20$, $p = .014$).

Overall Copresence is also significantly higher than the scale centre, paired one-tailed $t(5) = -2.36$, $p = .032$ with a large effect size ($d = -1.49$), confirming (H7) ($Z = -1.78$, $p = .038$).

6.4.1.2 Exploratory and In-Depth Analysis

The response rates for some of the questionnaires were quite low (in part due to the couple that aborted the study prematurely and therefore could not complete the later questionnaires) as is shown in Figure 6.2

on page 152. The overall rate of completed questionnaires compared to three full runs was 71 % with the lowest rate for a single questionnaire being 33 % for Q_γ. Only one couple completed all the questionnaires.

Many items could therefore not be evaluated. Q_α, Q'_α, and Q_γ could not be used at all and planned comparisons between Q_γ and $Q_{\beta\gamma}$ could also not be performed.

Figures 6.3 and 6.7 (as well as Figure B.1 on page 433) show confidence intervals for questions from the final questionnaire $Q_{\beta\gamma}$ and Figures 6.4(a) and 6.5(a) also indicate confidence intervals, here from $Q_{\beta\gamma}$ in comparison to Q_β. These intervals represent the range in which the true mean of the population would fall in 95 % of cases, assuming an underlying normal distribution [121]. This can be nothing more than an approximation of course, since we cannot assume a normally distributed process given the limited and discrete choice of answers a 7-point scale provides. A Shapiro-Wilk normality test [489, 490] rejected 33.9 % of the 127 7-point items of the final questionnaire and the usually less powerful [516] Lilliefors test [331] (a modified Kolmogorov-Smirnov test) rejected 29.9 %.[8]

The significance tests we did were performed on the scales instead of the individual items though, for which only 10.7 % failed the Shapiro-Wilk test and 17.9 %[9] the Lilliefors test (5 and 3 in absolute numbers).[10]

On many occasions, we provide the more intuitively understood confidence intervals in addition to (or sometimes – as for example in plots – instead of) the usual measures like mean and standard deviation. The margin of error for the confidence intervals was computed as follows:

[8] Although the numbers seem similar, only 22.1 % of items were rejected by both tests compared to 41.7 % of items that were rejected by at least one of the tests. See Figure B.2 on page 434 for a more detailed picture of these results.

[9] not counting derived scales

[10] These were *Shared Understandings* and the Relationship Closeness Inventory on $Q_{\beta\gamma}$ to fail both tests, *Knowing Each Others' Experiences* on $Q_{\beta\gamma}$ to fail Shapiro-Wilk narrowly ($W(6) = .79$, $p = .048$), and *Emotional Expressiveness* and *Engagement & Playfulness* on $Q_{\beta\gamma}$ and *Feelings of Closeness* on Q_β to fail only the Kolmogorov-Smirnov test.

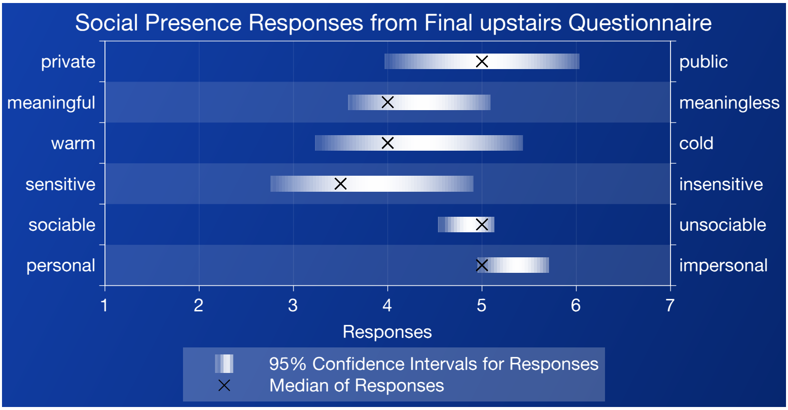

Figure 6.3. Responses from the final upstairs questionnaire ($Q_{\beta\gamma}$). These are social presence items from Short et al. [494] using the semantic differential technique [409]. The bars indicate the 95 % confidence intervals. The crosses additionally mark the median of the responses.

$$\text{Margin of Error} = \Phi^{-1}\left(\frac{\gamma+1}{2}\right)\frac{\sigma}{\sqrt{n}} \tag{6.1}$$

where Φ^{-1} is the inverse of the normal cumulative distribution function, $\gamma = 0.95$ is the confidence level, n the sample size and σ^2 the sample variance or the variance of the n mean responses per participant when computing the confidence of a whole scale instead of individual items.

Scale means and medians were computed according to Table B.1 on page 432 except for the comparisons between Q_β and $Q_{\beta\gamma}$ (Figure 6.4 on page 157 and Figure 6.5 on page 159) where the items that were not used when creating Q_β were not included in the $Q_{\beta\gamma}$ scales either. For the social presence scale that uses the semantic differential technique [409, 494] (Figure 6.3), all items were used except for the private—public dichotomy which was taken from the Privacy scale in Short et al. [494]. The resulting means and medians can be found in Figure 6.6 on page 160 and Table B.3 on page 436 for $Q_{\beta\gamma}$ and in Figure 6.4

Combined Scale	Substituent Scales
Overall Copresence	Isolation/Aloneness, Mutual Awareness, Perceived Other's Copresence
Networked Minds Copresence	Isolation/Aloneness, Mutual Awareness
Overall Social Connectedness	Relationship Salience, Shared Understandings, Knowing Each Others' Experiences, Feelings of Closeness
Emotion Transmission	Perceived Emotional Contagion/Empathy, Emotional Expressiveness
Presence in Absence and Staying in Touch	Presence in Absence, Staying in Touch
Overall Social Presence	Social Presence (Direct Questions), Social Presence (Semantic Differential Technique)

Table 6.4. Composition of the combined scales used in Figure 6.6 on page 160, Table 6.6 on page 158 and Table B.3 on page 436.

on page 157 and Figure 6.5 on page 159 for the comparisons between Q_β and $Q_{\beta\gamma}$. Some derived scales are developed in Section 6.5; their composition can be found in Table 6.4 and Table 6.15 on page 176.

Cronbach's α (cf. Equation (2.4) on page 40) was computed for all scales as a measure of their internal consistency (Table 6.6 on page 158 and Figure 6.6 on page 160). Different sources assess various values for α differently but one commonly found threshold for a good internal consistency is $\alpha = .7$ [289], so we highlighted the area under the interval $[.7, 1]$ in our plots.

Social Connectedness after van Bel et al. [550] was compared before

	Secure	Dismissing	Preoccupied	Fearful
Secure	1.00			
Dismissing	.52	1.00		
Preoccupied	.62	.28	1.00	
Fearful	.26	.11	.08	1.00

Table 6.5. Correlations between attachment types after Bartholomew and Horowitz [26] but much unlike the findings of Asendorpf et al. [15].

and after the two-week system-use period (Figure 6.4 on the next page). The expectations of the participants towards the system from Q_β were compared to the corresponding items asking for the actual experiences from $Q_{\beta\gamma}$ (Figure 6.5 on page 159). For both sets of comparisons, statistical significance was tested with paired *t*-tests and, because of the questionable normality and the small sample size [76, 137], we also performed Wilcoxon signed rank tests as a non-parametric alternative.

All results can be found in Table 6.7 on page 162. The only statistically significant difference between expectations and experiences is found for the Cognitive Load scale. There, both tests revealed such a statistically significant difference. The *t*-test shows a higher expected Cognitive Load[11] ($\bar{x} = 4.17$, $SD = .50$) than the actually self-reported Cognitive Load after using the system ($\bar{x} = 2.94$, $SD = 1.37$), paired two-tailed $t(5) = 4.82$, $p = .005$ with a large effect size ($d = -1.83$). The Wilcoxon signed rank test shows the same better than expected cognitive load $z = -2.21$, $p = .030$. The median response[12] for the expectations is higher ($Md = 5$) than for the self-reports ($Md = 2$) with a large effect size ($r = -.64$).

[11] For the Cognitive Load scale, items were reverse coded in such a way that a higher scale value means a higher cognitive load.

[12] The median scale response combines the responses in reverted form where indicated in Table B.1 on page 432.

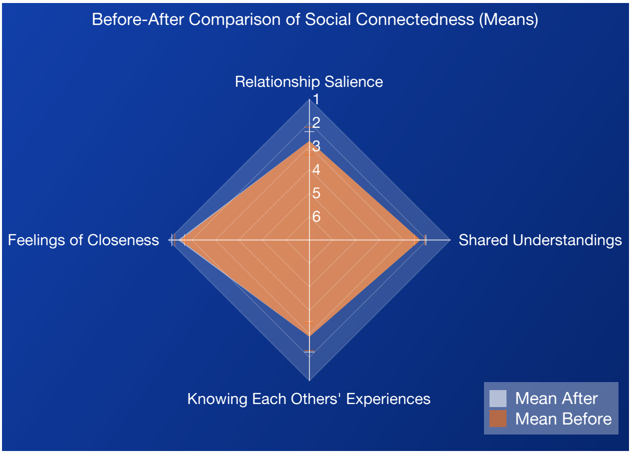

(a) Means

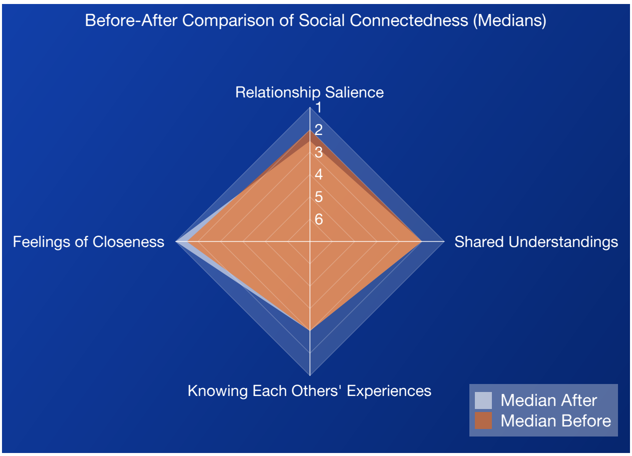

(b) Medians

Figure 6.4. Comparison of four social connectedness scales based on van Bel et al. [550] before and after using upstairs. The mean plot additionally has confidence interval indicators in the corresponding colour.

Scale	Cronbach's α upstairs	Cronbach's α FEELABUZZ
Relationship Salience	.34	.80
Shared Understandings	.00	.84
Knowing Each Others' Experiences	.76	.37
Feelings of Closeness	−.94	−.15
Overall Social Connectedness	.71	.86
Relationship Questionnaire	.76	—
Relationship Closeness Inventory	.52	—
Isolation/Aloneness	−.64	.00
Mutual Awareness	.88	.34
Perceived Emotional Contagion/Empathy	.80	.38
Behavioural Interdependence	.17	.69
Perceived Other's Copresence	.78	.74
Telepresence	.76	.59
Social Presence (Direct Questions)	.82	.59
Social Presence (Semantic Differential Technique)	.81	.86
Overall Social Presence	.88	.77
Emotional Expressiveness	.86	.87
Staying in Touch	.48	.79
Engagement & Playfulness	−1.33	.70
Presence in Absence	.70	.76
Opportunity for Social Support	.89	.84
Feeling Obligated	.60	.22
Unmet Expectations	.13	.48
Threat to Privacy	.65	.80
Cognitive Load	.50	.40
Vexatiousness	.81	.75
Ability to Ignore	.65	.37
Overall Copresence	.89	.70
Networked Minds Copresence	.85	.24
Emotion Transmission	.91	.52
Presence in Absence and Staying in Touch	.81	.87

Table 6.6. Cronbach's α values from $Q_{\beta\gamma}$ and the final FEELABUZZ questionnaire.

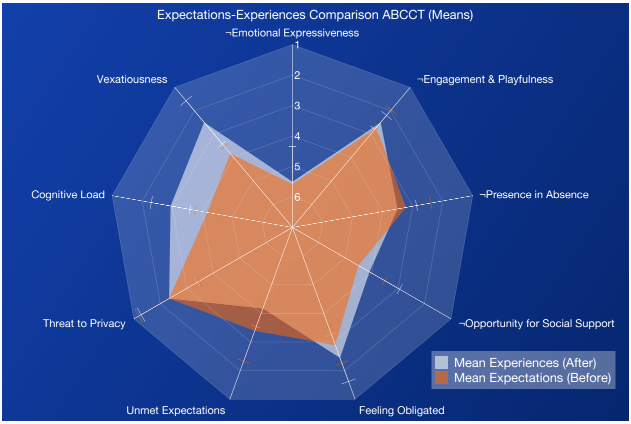

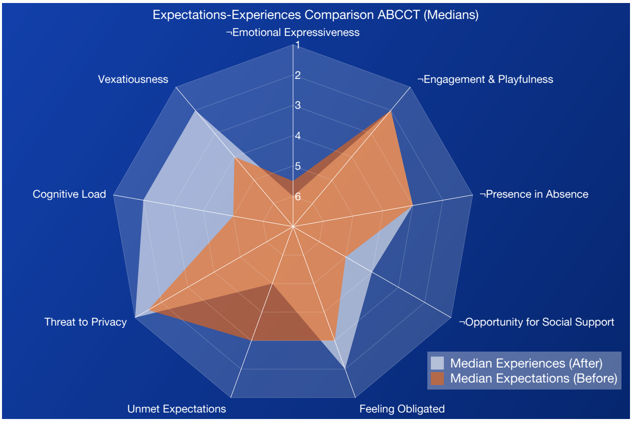

Figure 6.5. Comparison of the expectations regarding the affective costs and benefits of upstairs (from Q_β) and the actual self-reports at the end of the study (from $Q_{\beta\gamma}$). The scales are mainly based on the ABCCT [584]. The larger the area the better (which means that the cost scales – from "Feeling Obligated" to "Vexatiousness" – are reversed while the benefits scales preserve 1 to be the strongest agreement). The markings in the means plot represent the confidence intervals, based on the variance of the mean scale response per participant. Tests for statistical significance can be found in Table 6.7 on page 162. $Q_{\beta\gamma}$ scales exclude items missing from Q_β so values can differ from Figure 6.6 on the following page.

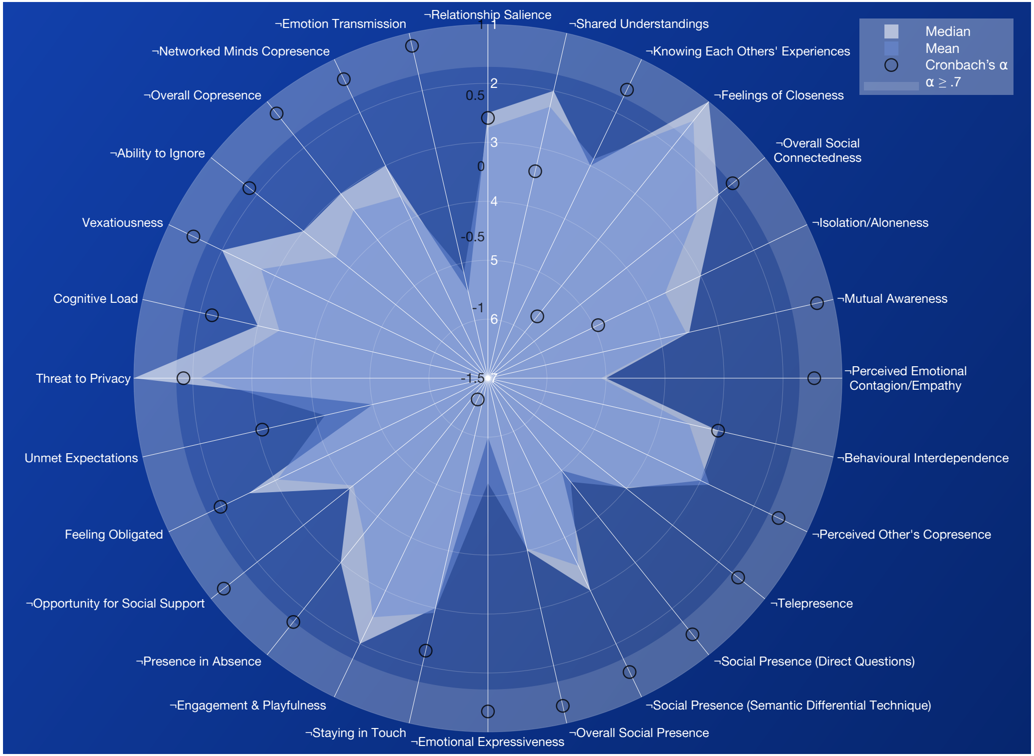

Figure 6.6. Scales from the final upstairs questionnaire ($Q_{\beta\gamma}$). Networked Minds Copresence, Emotion Transmission and all scales that start with "Overall" are combined scales. See Table 6.4 on page 155 for details. All items and scales were reversed when necessary to ensure a larger area (smaller response value) to mean a better result. To keep scale name semantics consistent during this process, benefit scales are prefixed by ¬ so a low value always means a high expression of the property described by the scale name. See Table B.3 on page 436 for the exact results.

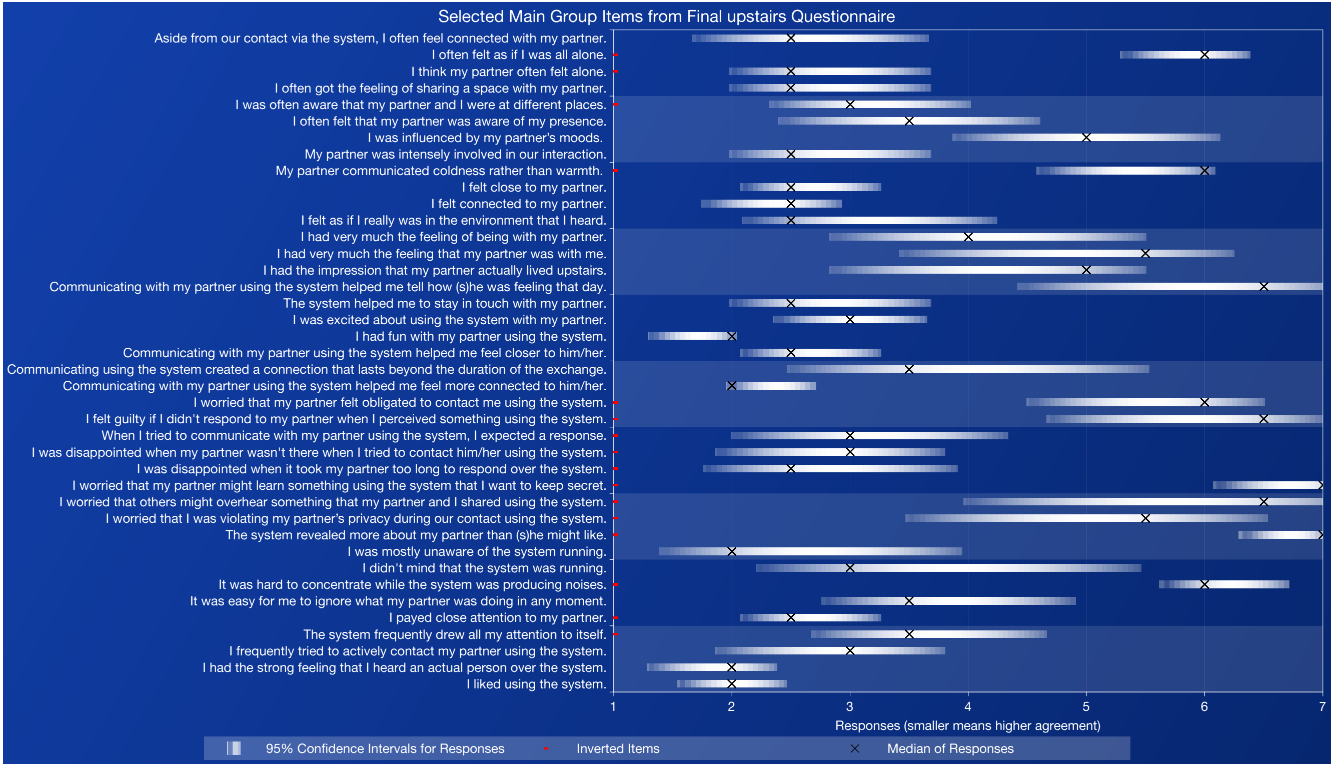

Figure 6.7. Confidence intervals and medians of selected items from the final upstairs questionnaire $Q_{\beta\gamma}$ (cf. Figure 6.3 on page 154 for an explanation of the plot).[a] Items that were reversed in the scale computations are marked in with red spots. No items were reversed for this plot though, so they range from "1 – I agree strongly" to "7 – I disagree strongly" in all cases. For a full list and pre-reversed items see Figure B.1 on page 433.

[a]Bear in mind that ~ 2 of the 39 items in this plot can be expected to have a population mean outside the confidence intervals shown.

Scale	v/df	$t(v)$	p-Value	Cohen's d	Z	Wilcoxon's T	r
Relationship Salience	5	.000	1.000	.000	.000	1.000	.000
Shared Understandings	5	.000	1.000	.000	.000	1.000	.000
Knowing Each Others' Experiences	5	.000	1.000	.000	−.105	.938	−.030
Feelings of Closeness	5	1.000	.363	−.263	−1.000	1.000	−.289
Overall Social Connectedness	5	.236	.823	−.037	−.526	.656	−.152
Relationship Closeness Inventory	5	.542	.611	−.136	−.315	.813	−.091
Emotional Expressiveness	5	.159	.880	−.077	.000	1.000	.000
Engagement & Playfulness	5	.432	.684	−.308	−.647	.625	−.187
Presence in Absence	5	−1.222	.276	.303	−1.051	.344	−.303
Opportunity for Social Support	5	.689	.521	−.311	−.271	.875	−.078
Feeling Obligated	5	.830	.444	−.485	−.944	.438	−.273
Unmet Expectations	5	−1.750	.141	.714	−1.625	.188	−.469
Threat to Privacy	5	.000	1.000	.000	−.137	1.000	−.587
Cognitive Load	5	4.824	.005	−1.827	−2.214	.031	−.639
Vexatiousness	5	2.794	.038	.784	−1.841	.125	−.531

Table 6.7. Wilcoxon signed rank test results [577] and paired two-tailed t-test results [521] of the social connectedness scales (and Overall Social Connectedness) before and after upstairs use (Figure 6.4 on page 157), the RCI Strength subscale before and after system use (Section 2.5) [34] and expectations towards the system versus the actual experiences (Figure 6.5 on page 159). For the cell colour of the effect sizes d and r, the threshold value for "large effect" from Cohen [85] was used as the maximum intensity (i. e. $|r| \geq .5$ and $|d| \geq .8$). The significance values p and T are coloured from 0 to 1.[a] The column v/df shows the degrees of freedom for the t-tests.

[a] Yellow and blue are only used to increase the comprehensibility of the table by visually differentiating between significance and effect size. They do not indicate a difference in judgement as in Table 6.9 on page 165.

Scale	Pearson's r
Relationship Salience	.05
Shared Understandings	.47
Knowing Each Others' Experiences	.27
Feelings of Closeness	.13
Overall Social Connectedness	.37
Relationship Questionnaire	.58
Relationship Closeness Inventory	.49
Emotional Expressiveness	.45
Engagement & Playfulness	.23
Presence in Absence	.23
Opportunity for Social Support	.45
Feeling Obligated	.50
Unmet Expectations	.65
Threat to Privacy	.67
Cognitive Load	.16
Vexatiousness	−.03

Table 6.8. Pearson correlation of the Social Connectedness scales (and Overall Social Connectedness) before and after system use (Figure 6.4 on page 157), the RCI Strength subscale [34] and Relationship Questionnaire [26] before and after upstairs use (Section 2.5) and expectations towards the system versus the actual experiences (Figure 6.5 on page 159). For the cell colour, the threshold value for "large effect" from Cohen [85] was used as maximum intensity ($|r| \geq .5$).

The Vexatiousness scale which is derived from the Cognitive Load scale (cf. Section 6.4.1.3) achieved similar results with expectations being higher ($\bar{x} = 3.83$, $SD = .41$, $Md = 4$) than the post-hoc reports ($\bar{x} = 2.50$, $SD = 1.26$, $Md = 2$). This result is statistically significant according to the paired two-tailed t-test $t(5) = 2.79$, $p = .038$ but not according to the Wilcoxon signed rank test $z = -1.84$, $p = .125$, even though the effect size only indicates a medium effect for the difference of means ($d = .61$) but a large effect for the difference of medians ($r = -.53$) [85].

The Relationship Questionnaire [15, 26, 118] looks at different attachment styles. Here, Asendorpf et al. [15] had previously found a strong negative correlation between the secure and the fearful attachment style, as had been suggested by Bartholomew and Horowitz [26]. Contrary to the orthogonal model by Bartholomew and Horowitz though, the supposedly uncorrelated axis dismissing—preoccupied was found to have both poles negatively correlated with the secure type. Table 6.5 on page 156 shows that our data matches neither model. Because of these inconsistencies and the small sample size, we decided not to regard this scale in any further analysis.

We also calculated the correlations between all the scales on $Q_{\beta\gamma}$ to find any striking patterns (Figure 6.21 on page 216) and since the large number of 325 potentially interesting correlations[13] bore the risk of actually masking relevant patterns, we additionally used subsets of this correlation matrix, partially augmented by single items that seemed relevant during the exploratory process described in Section 6.5. The resulting plots are shown in Figure 6.17 on page 197, Figure 6.18 on page 198 and Figure 6.19 on page 199. Since the correlations between the $Q_{\beta\gamma}$ scales and the Social Desirability scale are more about the confirmation that there does not seem to be a social desirability bias instead of finding patterns and because there is only one relevant column from the correlation matrix, we chose a tabular format for these results (Table B.5 on page 438). Similarly, the correlations between expectations and experiences – as one measure of effect size – form a single column which is presented in tabular form in Table 6.8 on page 163.

[13] the lower triangular part of the correlation matrix for 26 scales without the diagonal elements: $\frac{26^2-26}{2} = 325$

Item	Item Number	Item-Total Correlation	Cronbach's α without Item
I was mostly unaware of the system running.	2	.83	.23
It was easy for me to push the system to the back of my mind.	4	.81	.22
I had the feeling that the noises produced by the system disturbed my concentration.	7	.70	.32
The system frequently drew all my attention to itself.	11	.52	.39
It was easy for me to ignore what my partner was doing in any moment.	9	.49	.39
It was hard to concentrate while the system was producing noises.	6	.23	.48
I payed close attention to my partner.	10	.10	.50
The system distracted me from important tasks.	1	.01	.52
I often only realised in hindsight that I had heard something through the system.	12	−.10	.54
I didn't mind that the system was running.	3	−.28	.66
Using the system was effortless.	8	−.30	.58
I recognised important events with the system, even when I was not actively monitoring it.	5	−.32	.58

Table 6.9. Cognitive load items sorted by their corrected item-total correlation.

Item	Component 1	Component 2	Component 3
I had the feeling that the noises produced by the system disturbed my concentration.	.93	−.30	.13
I was mostly unaware of the system running.	.88	.34	−.08
The system frequently drew all my attention to itself.	.76	.09	−.42
The system distracted me from important tasks.	.58	−.11	.07
I recognised important events with the system, even when I was not actively monitoring it.	−.26	−.83	.37
It was hard to concentrate while the system was producing noises.	.46	−.81	.27
I often only realised in hindsight that I had heard something through the system.	−.20	.74	−.32
I payed close attention to my partner.	.39	.72	−.45
It was easy for me to ignore what my partner was doing in any moment.	.42	.47	.74
Using the system was effortless.	.14	−.61	−.67
I didn't mind that the system was running.	−.46	.51	.62
It was easy for me to push the system to the back of my mind.	.59	.21	.62

Table 6.10. Component matrix for the first three principle components of the Cognitive Load scale on $Q_{\beta\gamma}$.

Item	Component 1	Component 2	Component 3
I had the feeling that the noises produced by the system disturbed my concentration.	.95	−.27	.04
I was mostly unaware of the system running.	.86	.38	.14
The system frequently drew all my attention to itself.	.77	.35	−.29
The system distracted me from important tasks.	.58	−.10	.05
I recognised important events with the system, even when I was not actively monitoring it.	−.21	−.91	−.06
I payed close attention to my partner.	.35	.88	−.05
It was hard to concentrate while the system was producing noises.	.51	−.80	−.09
I often only realised in hindsight that I had heard something through the system.	−.24	.78	.03
It was easy for me to ignore what my partner was doing in any moment.	.36	.01	.90
Using the system was effortless.	.20	−.14	−.86
I didn't mind that the system was running.	−.51	.06	.76
It was easy for me to push the system to the back of my mind.	.55	−.12	.68

Table 6.11. Pattern matrix for the first three principle components of the Cognitive Load scale on $Q_{\beta\gamma}$.

Item	Component 1	Component 2	Component 3
I had the feeling that the noises produced by the system disturbed my concentration.	.95	−.28	.01
I was mostly unaware of the system running.	.85	.38	.17
The system frequently drew all my attention to itself.	.76	.31	−.26
The system distracted me from important tasks.	.58	−.10	.03
I recognised important events with the system, even when I was not actively monitoring it.	−.19	−.91	−.15
I payed close attention to my partner.	.34	.86	.04
It was hard to concentrate while the system was producing noises.	.53	−.82	−.18
I often only realised in hindsight that I had heard something through the system.	−.25	.79	.11
It was easy for me to ignore what my partner was doing in any moment.	.35	.10	.90
Using the system was effortless.	.21	−.23	−.88
I didn't mind that the system was running.	−.52	.15	.77
It was easy for me to push the system to the back of my mind.	.55	−.07	.67

Table 6.12. Structure matrix for the first three principle components of the Cognitive Load scale on $Q_{\beta\gamma}$.

Component	1	2	3
1	1.00	−.02	−.01
2	−.02	1.00	.10
3	−.01	.10	1.00

Table 6.13. Component correlation matrix for the first three principle components of the Cognitive Load scale on $Q_{\beta\gamma}$.

6.4.1.3 Cognitive Load Scale

While most scales were at least in large parts taken or adapted from previous works by other authors and thus already tested to some degree with a much larger sample size, the Cognitive Load scale presented in this work is new. We will therefore investigate it in more detail and more scrutinously.

The scale as a whole has a poor internal consistency with an α of .50. Although the very small sample size calls the attempt of a factor analysis into question [88, 525], we will nonetheless do so with a constant awareness of the tentativeness of any results and we will try to use the FEELABUZZ study to test the plausibility of our results.

The item correlation matrix (Table B.2 on page 435) was subjected to a principal component analysis (PCA) [243]. The scree plot in Figure 6.8 on page 171 shows the resulting components, sorted by their eigenvalue [70]. However, the usual subjective criterion of an abrupt change in steepness of the graph is rather inconclusive here since there is no sharp bend at any component with a non-zero eigenvalue. More objective criteria often applied include the Kaiser criterion which simply discards all factors for which the eigenvalue $\lambda_i < 1$. This would yield four factors, which is a lot for such a small sample. The Kaiser criterion also has been criticised in the past as being too simplistic (e. g. Kaiser [279], as cited in Ford et al. [152]). Horn [241] proposed a test that is regarded as more adaptive which generates random data and compares their eigenvalues with the data at hand. Normally, only eigenvalues that

Item	Code	Factor Loading	Item-Total Correlation	Cronbach's α without Item
I had very much the feeling of being with my partner.	TP3	.97	.89	.87
I think my partner often got the feeling of sharing a space with me.	MA4	.96	.83	.88
I often got the feeling of sharing a space with my partner.	MA1	.93	.78	.88
I often felt that my partner was aware of my presence.	MA5	.90	.86	.88
I felt as if I really was in the environment that I heard.	TP2	.89	.78	.88
I had very much the feeling that my partner was with me.	TP4	.87	.93	.87
It was as if we were both in the same environment.	MA6	.86	.93	.87
I was often aware that my partner and I were at different places.	MA2	.49	.25	.91
I hardly noticed my partner.	MA3	.45	.44	.90
I had the impression that my partner actually lived upstairs.	TP5	.09	.16	.92
The experience with the system was very intense.	TP1	−.01	.11	.91

Table 6.14. Factor loadings for the first principal component of the combination of Mutual Awareness and Telepresence scale on $Q_{\beta\gamma}$, along with the corrected item-total correlation for this scale and the combined α without the item.

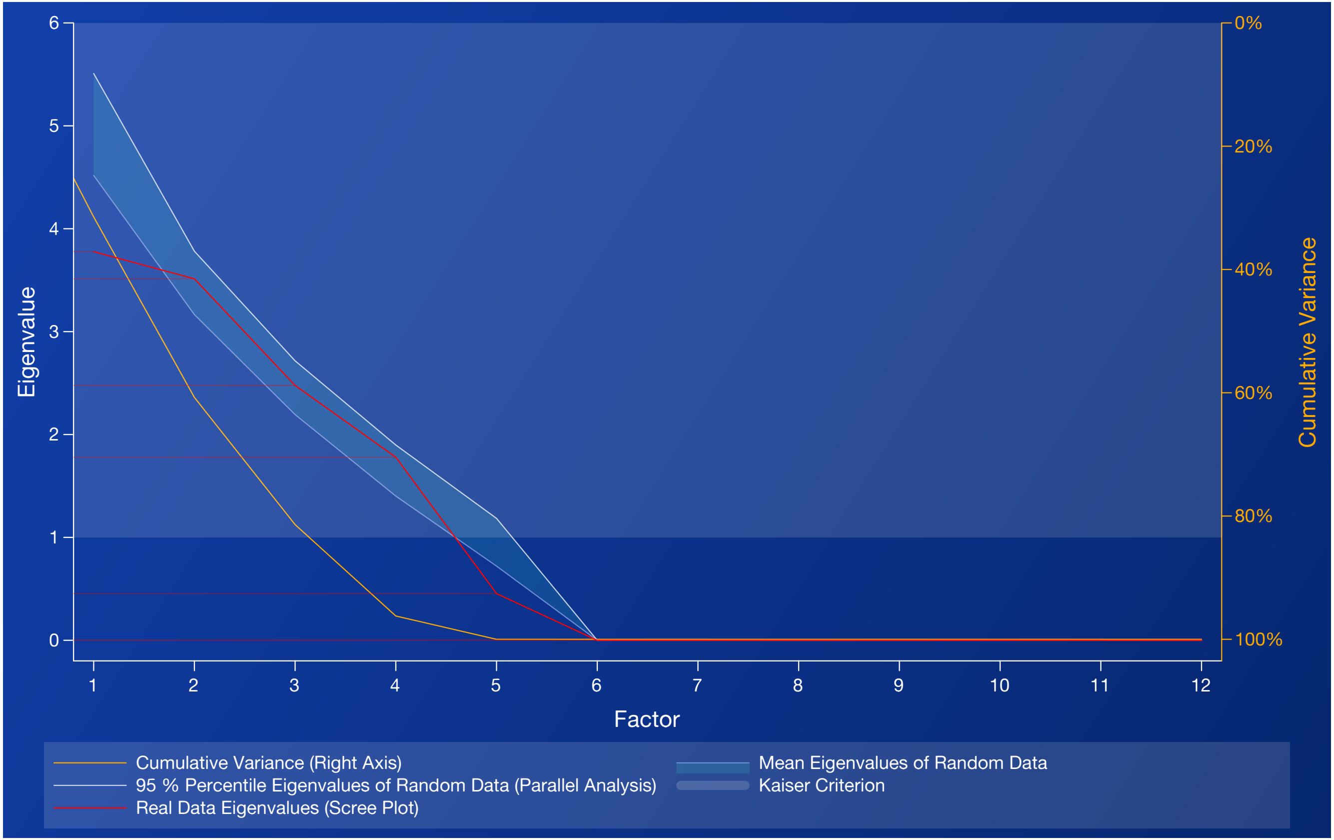

Figure 6.8. Scree plot for the Cognitive Load scale on $Q_{\beta\gamma}$. The Kaiser criterion (discard factors with eigenvalue < 1) is indicated as a lighter background and the average .95 quantile and means of the eigenvalues of 10 000 sets of randomly generated data are indicated as additional lines (parallel analysis after Horn [241]).

are larger than the respective 95th percentiles are considered but in our case not a single factor meets this criterion. Using the means of the random data eigenvalues yields the same factors as the Kaiser criterion.

In summary, these tests either argue for four or five components or for none at all which is most likely an artifact of the small sample size and the fact that we have fewer samples than factors. Looking at the cumulative variance explained (Figure 6.8 on page 171), three factors explain 81.4 % of the variance, so we chose to retain this still manageable number of factors for further investigation (Table 6.10 on page 166).

Although the component matrix (Table 6.10 on page 166) already seems to show a relatively clear structure in the factor loadings, a factor rotation was performed to obtain factors that are easier to interpret. Because we do not expect the Cognitive Load factors to be independent, we chose an oblique (non-orthogonal) rotation [136, 463]. A direct oblimin transform [209] could be performed with a relatively high parameter $\delta = .4$ [253] and still result in only weakly correlated components (cf. Table 6.13 on page 169).[14]

Both pattern matrix[15] (Table 6.11 on page 167) and structure matrix[16] (Table 6.12 on page 168) show the same very clear attributions of items to the three factors.[17] All items load highly on one factor (with the possible exception of "The system distracted me from important tasks" with a loading of .58) and much less on the other two (here one might call "It was easy for me to push the system to the back of my mind" an exception which does not load much less on Component 1 than on Component 3). However, three items stand out by having a high negative loading whereas the other factors in each scale load

[14] Russell [463] suggests to perform the orthogonal varimax rotation [278] first to see if the assumption of correlated factors is correct. We did this and found that the average of the off-diagonal elements of the factor correlation matrix (as shown for the direct oblimin transform in Table 6.13 on page 169) is .20 for the varimax rotation as compared to .04 for the direct oblimin rotation.

[15] The pattern matrix contains the rotated factor loadings (i. e. the rotated principal component coefficients) [142].

[16] The structure matrix contains the correlations between variables and factors [142].

[17] Expressed as item numbers, the three scales are 1, 2, 7, 11; 5, 6, 10, 12; and 3, 4, 8, 9.

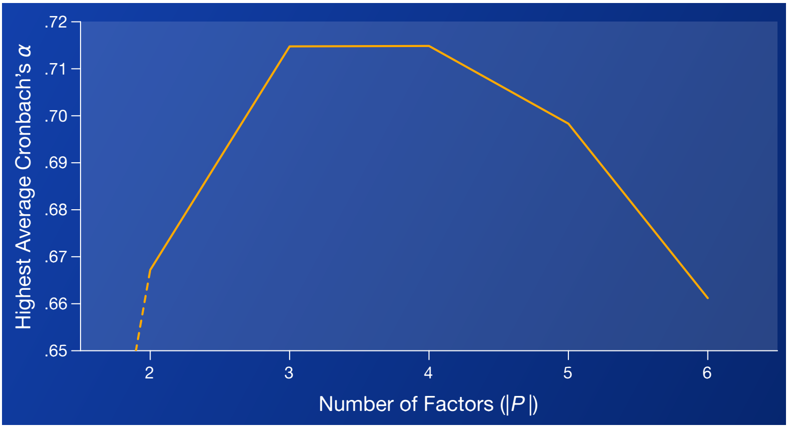

Figure 6.9. Highest achievable $\bar{\alpha}$ for all possible subscales of the Cognitive Load scale with $2 \leq |P| \leq 6$ (cf. Equations 6.3 and 6.4 on the next page).

positively. Reverting the coding of these items (or all other items for that matter) does not seem appropriate from a semantic point of view. Generally, the interpretation of the three factors found through the PCA does not seem obvious.

As an approach from a different angle, we performed a global optimisation of Cronbach's α to directly get the subscales with a maximally high internal consistency. This is to say that we took all partitions of the set of items where the smallest cardinality of any element of a partition was 2 (to exclude single-item "scales") and computed Cronbach's α for each element in a given partition.

More formally, P is a valid partition of the set of all items in a scale S if and only if the following three conditions are true:

$$\begin{gathered} \bigcup P = S \\ \forall A, B \in P \; (A \cap B = \varnothing), \qquad A \neq B \\ \forall A \in P \; (|A| \geq 2) \end{gathered} \tag{6.3}$$

Then, the average α for a given partition was used as a metric for the performance of this partition.

Box 5: Optimisation Time Complexity
The naïve approach we pursued to compute the maximum $\bar{\alpha}$ by generating all possible partitions P limits severely the size of the scale that can feasibly be optimised. The time complexity of this algorithm is $O(B_n)$ with B_n being the nth Bell number and n being the number of items on the scale that is to be optimised [134]. Bell numbers can be recursively defined as

$$B_n = \sum_{k=0}^{n-1} \binom{n-1}{k} B_k \tag{6.2}$$

with $B_0 = 1$ (cf. Weisstein [567]). As Figure 6.10 on the facing page illustrates, this means that while our optimisation only took about 12 minutes, a slightly larger scale would already have been impossible to optimise using this approach. Optimising the RCI Strength subscale from the $Q_{\beta\gamma}$ questionnaire for example would have taken 300 times the age of the universe on the same machine. Clearly, one of the more sophisticated techniques from the arsenal of computational optimisation is due if this method is to be applied more generally.

Figure 6.9 on page 173 shows a comparison of the maximum achievable average α for different numbers of factors (i. e. different cardinalities of P). If we denote the average α of a partition P as $\bar{\alpha}(P)$, then using Equation (2.4) on page 40 and P as defined above, we can write the maximum $\bar{\alpha}$ as follows:

$$\max_P \bar{\alpha}(P) = \max_P \left(\frac{1}{|P|} \sum_{A \in P} \alpha(A) \right) \tag{6.4}$$

As can be seen in the plot, three and four factors yield the highest value of $\bar{\alpha} = .71$. We went for the smaller number of factors, again with the small sample size in mind. These empirically found three scales consist of item numbers 1, 2, 7, 8, 10, 11; 3, 4, 9, 12; and 5 and 6,

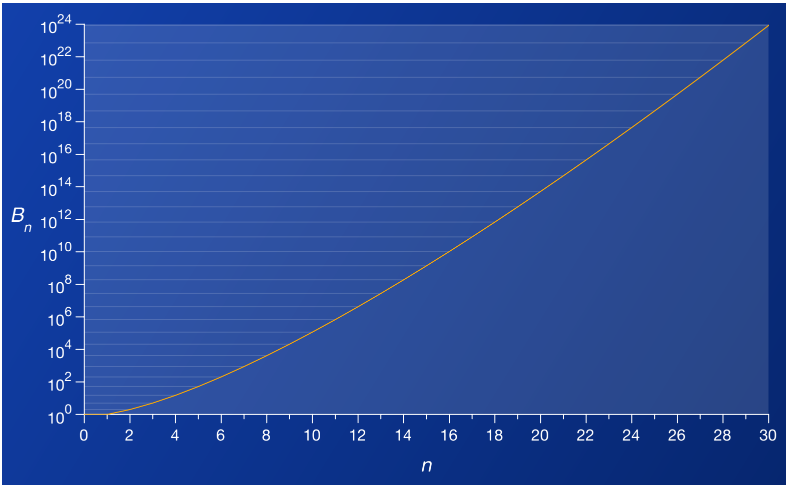

Figure 6.10. Log-lin plot of Bell numbers B_0 to B_{30}, showing the overexponential growth of the Bell numbers that determine the time complexity of the optimisation we performed.

with α values of .75, .65, and .74 respectively.[18] When looking at the associated questions, it is striking that the first scale mainly consists of items that ask for how annoying the system is while the second scale is mainly concerned with how well the user was able to ignore the system. This is a subtle difference but these two constructs are not the same thing and it is interesting that the α optimisation drew this line.

Given that these new subscales have a much better internal consistency than the original scale and that they can be interpreted semantically, we will define them as new scales called *Vexatiousness* and *Ability to Ignore*. The third scale contains miscellaneous items and we decided to move two items that ended up in the Vexatiousness scale without sounding as if they belonged to this construct to that

[18]Cf. Table 6.9 on page 165 for the item numbering or Table B.1 on page 432 for all items in the right order.

Vexatiousness	Ability to Ignore	Other/Miscellaneous
The system distracted me from important tasks.	It was easy for me to ignore what my partner was doing in any moment.	I recognised important events with the system, even when I was not actively monitoring it.
I was mostly unaware of the system running.	I often only realised in hindsight that I had heard something through the system.	It was hard to concentrate while the system was producing noises.
I had the feeling that the noises produced by the system disturbed my concentration.	I didn't mind that the system was running.	Using the system was effortless.
The system frequently drew all my attention to itself.	It was easy for me to push the system to the back of my mind.	I payed close attention to my partner.

Table 6.15. Empirically found Cognitive Load subscales and suggested names for these scales.

Miscellaneous scale, which actually improved the internal consistency for the Vexatiousness scale ($\alpha = .81$) while it naturally decreased the α for Miscellaneous, actually quite drastically to $\alpha = -.07$. This is acceptable, though, because this scale is not meant to be regarded any further – in fact, these items are good candidates for omittance in future uses of the Cognitive Load scales.

It is a disadvantage of the α optimisation method that it does not account for items that do not fit well with any of the scales because

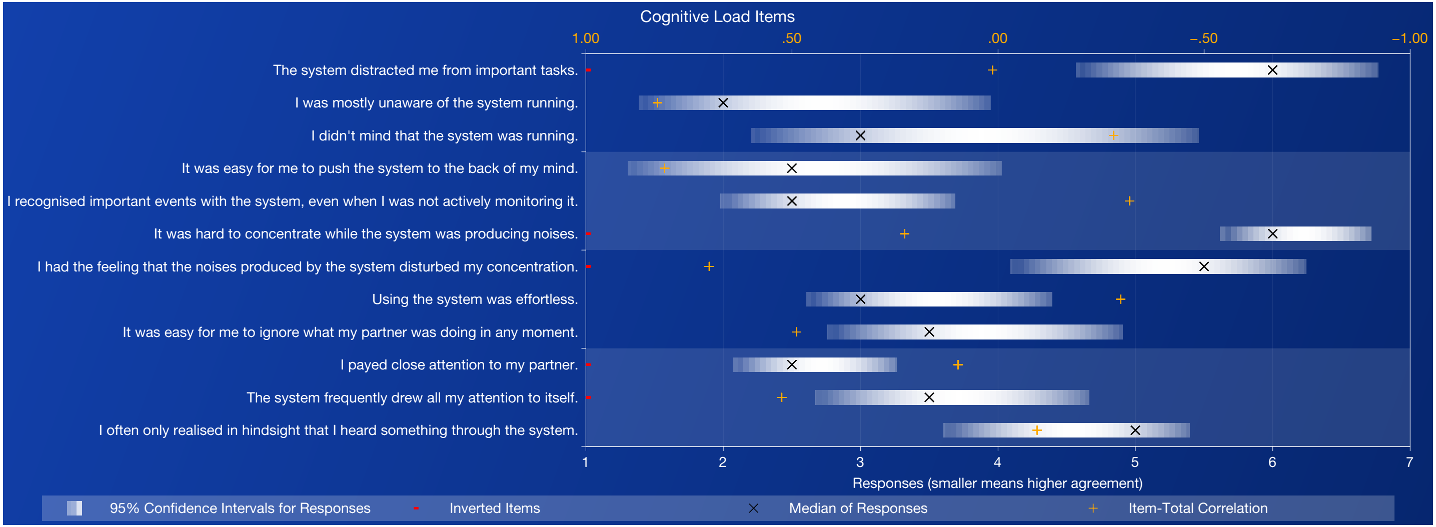

Figure 6.11. Responses from the final upstairs questionnaire ($Q_{\beta\gamma}$) to the items of the Cognitive Load scale. Values are shown as given on the questionnaire, i. e. lower values mean stronger agreement with a particular statement. Items that were recoded for the scale computation are marked red. Item-total correlation was also given a measure of how strongly a particular item measures what the scale measures.

it has to find a global optimum containing all items.[19] This means that such scales have to be tweaked manually, introducing the risk of researcher bias.

In item number notation, the new scales are 3, 4, 9, 12; 1, 2, 7, 11; and 5, 6, 8, 10. Table 6.15 on page 176 lists the items in these scales.

Looking back at the results from the PCA (Table 6.11 on page 167), the newly defined Vexatiousness scale matches the first factor perfectly. The Miscellaneous scale consists of the three items with high negative loadings and "I payed close attention to my partner."

6.4.2 FEELABUZZ

With FEELABUZZ, we chose not to perform non-parametric tests in addition to the *t*-tests, mainly because of a lack of a non-parametric alternative to the Welch test [568] that is not sensitive to differences in variance [137, 138]. The FEELABUZZ sample size is only slightly larger but even with the very small sample size of the upstairs study, the differences between the *t*-tests and the Wilcoxon tests were not dramatic in most cases (cf. Table 6.7 on page 162), so we felt confident to rely on the usual *t*-tests only and drop the analysis that does not require normality.[20]

6.4.2.1 Hypothesis Testing

As seen in Table 6.3 on page 150, the results for the hypotheses from Section 6.2 are identical, with the notable exception of (H5) that stated that the Cognitive Load would be better than the scale centre. In fact,

[19] We tried to further adapt the method by ignoring one of the scales but this lead to all but a few items being dumped into this ignored scale. We then weighted the scales by their size which solved the problem but introduced the artifact of choosing a specific cost function. Small differences here lead to vastly different results which is why we decided not to pursue this further and stick with the method that introduces the least amount of choice by the researcher and at the same time yielded results that are comprehensible and largely semantically coherent.

[20] The relation between sample size and normality is given by the central limit theorem [539].

Cognitive Load ($\bar{x}$ = 4.40, *SD* = .64) was significantly higher than 4 (two-tailed $t(12.00) = 2.27$, $p = .043$) with a large effect size ($d = .93$).

As with upstairs, there was more perceived Copresence ($\bar{x}$ = 3.47, *SD* = .75) than Social Presence ($\bar{x}$ = 4.36, *SD* = .99); one-tailed $t(12.00) = 3.98$, $p < .001$ with a large effect size ($d = 1.05$).

The result for (H2) was equally negative as it was for upstairs. Copresence ($\bar{x}$ = 3.47, *SD* = .75) was rated better than Telepresence ($\bar{x}$ = 4.15, *SD* = 1.21); two-tailed $t(12.00) = -3.06$, $p = .010$ with a medium effect size ($d = .71$).

The Threat to Privacy was again perceived as lower than the scale centre ($\bar{x}$ = 1.52, *SD* = .78); one-tailed $t(12.00) = -11.50$, $p < .001$ with a large effect size ($d = 4.69$) and likewise, Copresence was rated as better than the scale centre, $t(12.00) = -2.52$, $p = .013$ with a large effect size ($d = 1.03$).

6.4.2.2 Scale Analysis

As Figure 6.13 on page 185 shows, overall, the scale results between upstairs and FEELABUZZ were quite similar, with a strong correlation between scale means,[21] $r(19) = .76$, $p < .001$. There were some notable differences, though. Table 6.18 on page 183 shows how strongly scales differed and with which *p*-value. In particular, the following scales showed a significant difference between upstairs and FEELABUZZ when performing an unpaired two-tailed *t*-test [568]. Mutual Awareness was significantly worse for FEELABUZZ than for upstairs, $t(97.44) = -2.91$, $p = .005$, with a small effect size ($d = .26$). Behavioural Interdependence was also worse in FEELABUZZ $t(26.77) = -3.10$, $p = .005$, with a medium effect size ($d = .55$), as well as Staying in Touch, $t(35.55) = -2.95$, $p = .006$, with a small effect size ($d = .42$), Presence in Absence, $t(46.46) = -2.76$, $p = .008$, with a small effect size ($d = .37$), Opportunity for Social Support, $t(45.61) = -3.86$, $p < .001$, with a medium effect size ($d = .57$), and Cognitive Load, $t(164.49) = -3.60$, $p < .001$, with a small effect size ($d = .26$) and

[21] excluding combined scales and subscales

Kind of Noise	Number of Users Affected
Static (*Rauschen*)	4
Interruptions	3
Blips/Artifacts	3
Phantom Steps	0
Laptop Fan	1
Speech	1
Mobile Phone Interference	1

Table 6.16. Kind and frequency of noise reports for upstairs. Only reports in the questionnaires are counted (either from the multiple choice or the free answers – the latter are all entries below "Phantom Steps"). Cf. Table 6.17 on the facing page for intentional sounds that were transmitted. Speech was heard by all participants but only put down once on a questionnaire. A buzzing noise was mentioned in one of the interviews but not put on a questionnaire.

its derivative scale Vexatiousness, $t(57.22) = -2.51$, $p = .015$ with a small effect size ($d = .29$). Unmet Expectations was the only scale that was actually significantly better for FEELABUZZ than for upstairs, $t(15.24) = 2.87$, $p = .012$, with a medium effect size ($d = .71$).

The difference in this last scale seems to stem primarily from one item, "I was disappointed when it took my partner too long to respond over the system," unpaired two-tailed $t(10.96) = 3.78$, $p = .003$, with a large effect size ($d = 1.89$). This disappointment was rated as negative by our item reversion scheme (cf. Box 4 on page 149) but this was a difficult call since these kinds of expectations also point towards an emotional engagement.

The internal consistencies of the scales were also mostly similar, although the correlation was less pronounced than for the scale means,[22]

[22] All scales shown in Figure 6.13 on page 185 were included.

Kind of Sound	Number of Mentions
Footsteps	6
Knocking	6
Door Sounds	1
Chair Movement	3
Other Impact Sounds with Floor	2
Rustling	1
Impact Sounds with Furniture	1
Window Blind	1
Speech (unintelligible)	6
Speech (intelligible)	0
Sneezing / Other Non-Structure-Borne Sounds	2

Table 6.17. Types of sounds transmitted by upstairs as reported by the participants in the interviews. All mentions were spontaneous, therefore the true numbers may be higher if the participants forgot to mention a certain type of sound. Multiple mentions by one participant were only counted once. Noise not transmitted but created by the system itself is not covered here but in Table 6.16 on page 180.

$r(26) = .36$, $p = .057$. Overall, more scales showed a value of $\alpha > .7$, possibly due to the larger sample size.[23]

Against the overall trend, the combined presence scales show lower internal consistencies (cf. Table 6.6 on page 158) but except for the Networked Minds Copresence, they are all still good. The latter shows a marked drop to $\alpha = .24$ and no single item is to blame.[24] Telepresence also shows a drop to $\alpha = .59$ but here, the reduction in size from five

[23]The median α is almost equal, though (.70 for upstairs versus .69 for FEELABUZZ).

[24]The greatest improvement is achieved by dropping "I often felt that my partner was aware of my presence" but that also only raises the α to .42.

to four items provides a possible reason [161].

Figure 6.13 on page 185 and Table 6.18 on the next page compare the results of the final questionnaires of both systems.

The strong correlation between Social Connectedness and Presence in Absence that showed up for upstairs could not be found in the FEELABUZZ data, $r(11) = .29$, $p = .336$.

The question asking directly for a feeling of connectedness, which looked good for upstairs, is significantly worse for FEELABUZZ ($CI_{.95} = [3.31, 5.46]$, $Md = 4$), unpaired two-tailed $t(14.67) = -3.75$, $p = .002$.

The Overall Social Connectedness did not look significantly different from the upstairs results ($\bar{x} = 2.34$, $Md = 2$). Since there are no two sampling points for Social Connectedness, we can not test our hypothesis from Section 6.5.1.1.

The presence scales did not show any significant difference between the two systems and they look virtually identical (cf. Figure 6.15 on page 191). Correlations between the different presence scales look very similar to Figure 6.18 on page 198 in that no particularly weak correlations stand out (not shown). Indeed, 4 of the 10 correlations now reach a significance level of $p < .05$.

6.4.2.3 Cognitive Load

Figure 6.16 on page 193 shows that the Cognitive Load items were almost unanimously worse for FEELABUZZ than for upstairs, with an average difference of scale medians of 1.5 and .91 for the means.

With the FEELABUZZ study, we now have enough participants to test the suitability of the data for factor analysis, even though it is clear that there are still by far too few samples to call the analysis anything but tentative. Bartlett's Test of Sphericity [27] reached statistical significance ($p = .008$) while the Kaiser-Meyer-Olkin value was only .358 and therefore below the critical value of .5 [99, 279].[25]

The number of factors within the scale remains unclear, with the Scree plot suggesting 5 factors, the Kaiser criterion 4, and the parallel

[25] For a detailed look at the items responsible for this low value, cf. Table B.6 on page 439.

Scale	ν/df	t(ν)	p-Value	Cohen's d
Relationship Salience	54.548	1.275	.208	.217
Shared Understandings	54.972	−.018	.986	−.003
Knowing Each Others' Experiences	57.479	.915	.364	.154
Feelings of Closeness	32.160	−.838	.408	−.259
Isolation/Aloneness	19.435	1.892	.073	.429
Mutual Awareness	97.444	−2.906	.005	−.262
Perceived Emotional Contagion/Empathy	21.363	−1.989	.060	−.536
Behavioural Interdependence	26.774	−3.098	.005	−.550
Perceived Other's Copresence	107.122	.888	.376	.092
Telepresence	49.040	−.138	.891	−.020
Social Presence (Direct Questions)	45.002	−1.583	.120	−.214
Social Presence (SDT)	87.033	.619	.537	.081
Emotional Expressiveness	28.704	−1.051	.302	−.215
Staying in Touch	35.553	−2.952	.006	−.419
Engagement & Playfulness	48.745	−1.194	.238	−.189
Presence in Absence	46.464	−2.757	.008	−.365
Opportunity for Social Support	45.611	−3.864	< .001	−.575
Feeling Obligated	45.864	1.467	.149	.201
Unmet Expectations	15.245	2.869	.012	.709
Threat to Privacy	38.190	1.672	.103	.313
Cognitive Load	164.494	−3.599	< .001	−.260
Vexatiousness	57.221	−2.508	.015	−.292
Ability to Ignore	51.382	−.766	.447	−.096
Overall Social Connectedness	186.276	1.061	.290	.106
Overall Copresence	213.566	−1.635	.104	−.098

Table 6.18. Results of two-tailed *t*-tests for samples with unequal variances [568] between the final questionnaires of upstairs and FEELABUZZ. For the cell colour of the effect size *d*, the threshold value for "large effect" from Cohen [85] was used as the maximum intensity (i. e. $|d| \geq .8$). The *p*-values are coloured from 0 to 1.[a] The column ν/df shows the degrees of freedom for the *t*-tests.

[a] Yellow and blue are only used to visually discriminate between significance and effect size. They do not indicate a difference in judgement as in Table 6.9 on page 165.

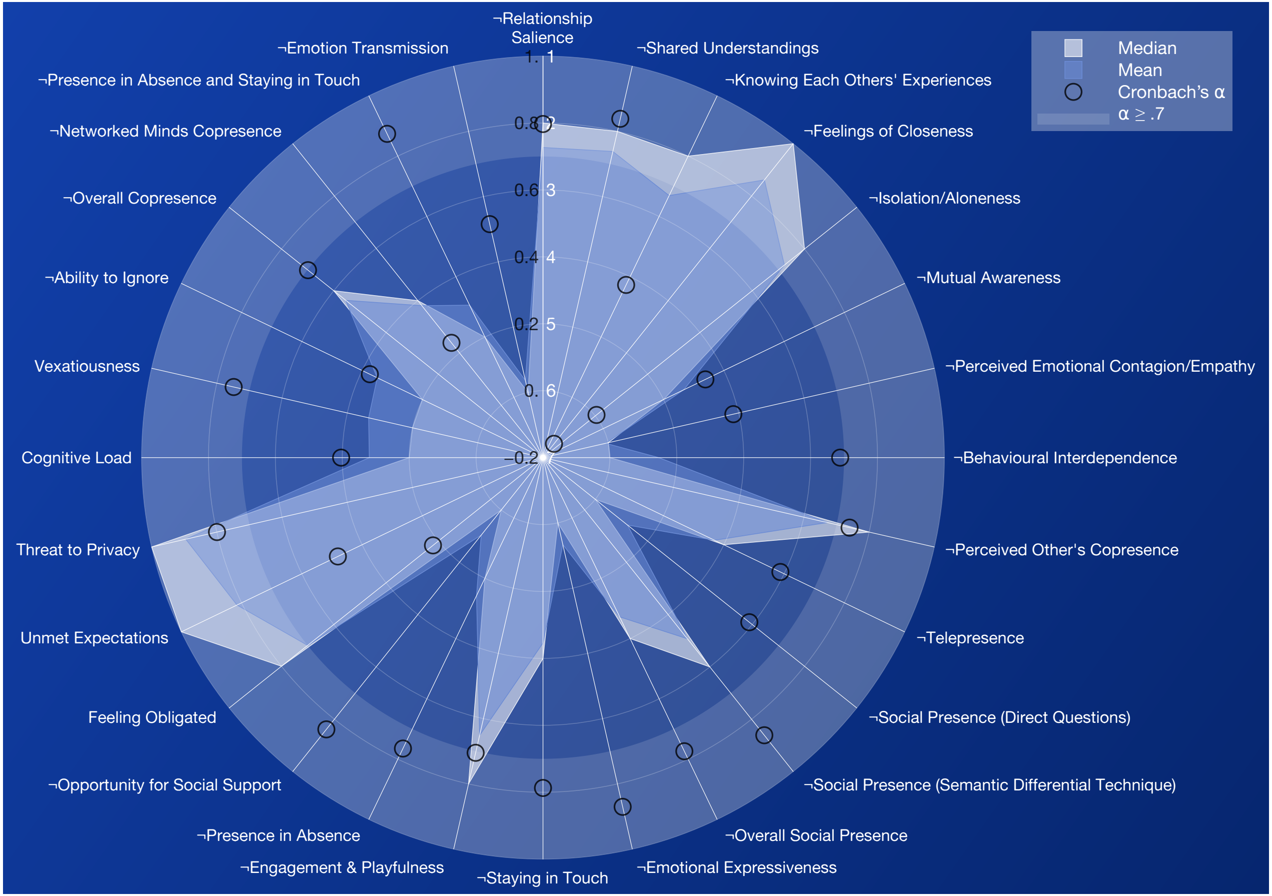

Figure 6.12. Scale results from the final FEELABUZZ questionnaire. Cf. Figure 6.6 on page 160 for an explanation of this chart.

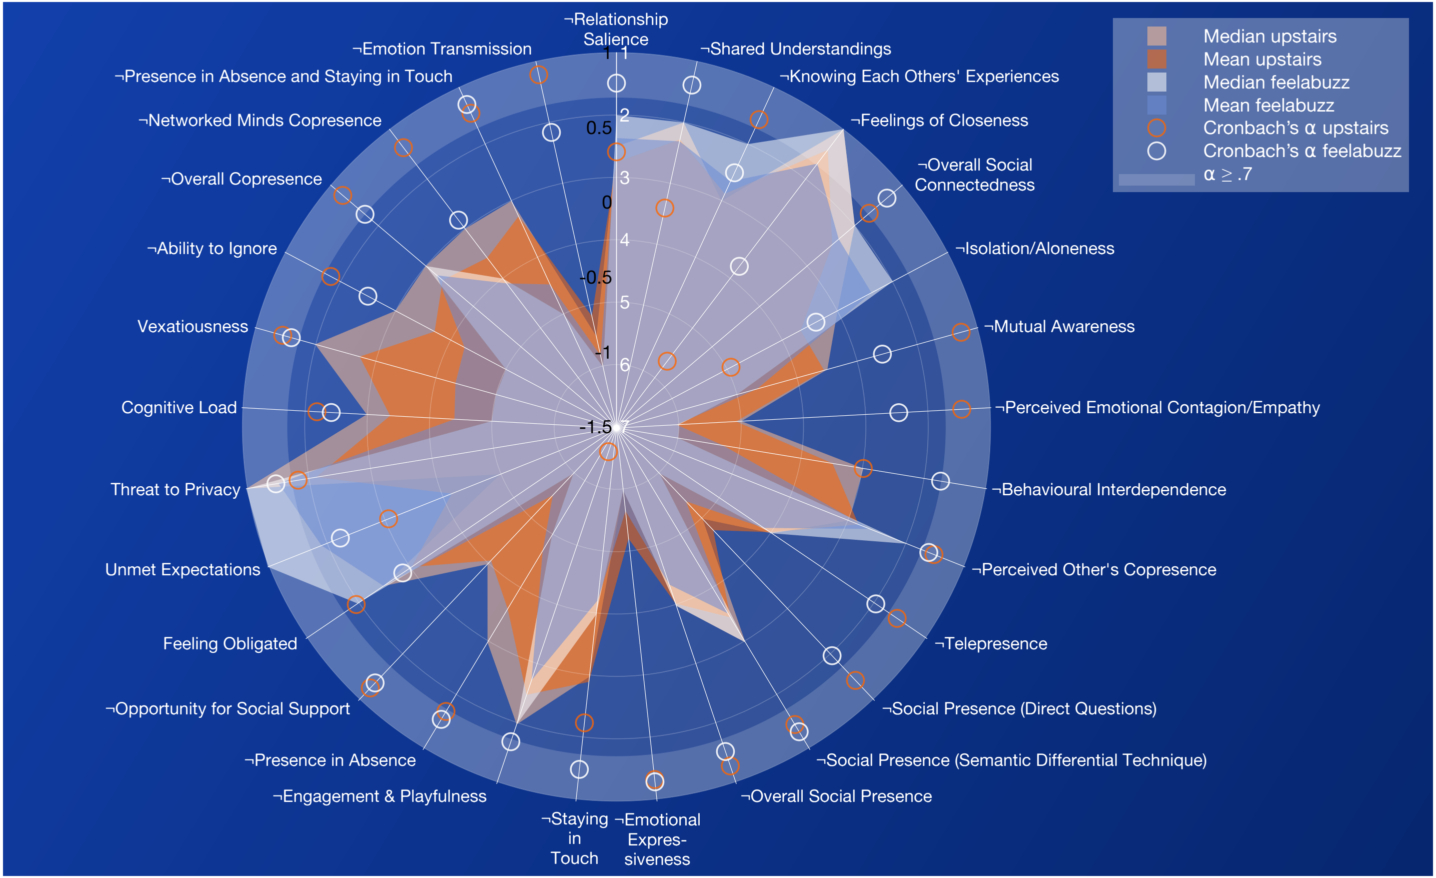

Figure 6.13. Comparison of responses to the final questionnaires of upstairs and FEELABUZZ, combining data from Figure 6.12 on page 184 and Figure 6.6 on page 160 (see there for an explanation of this chart).

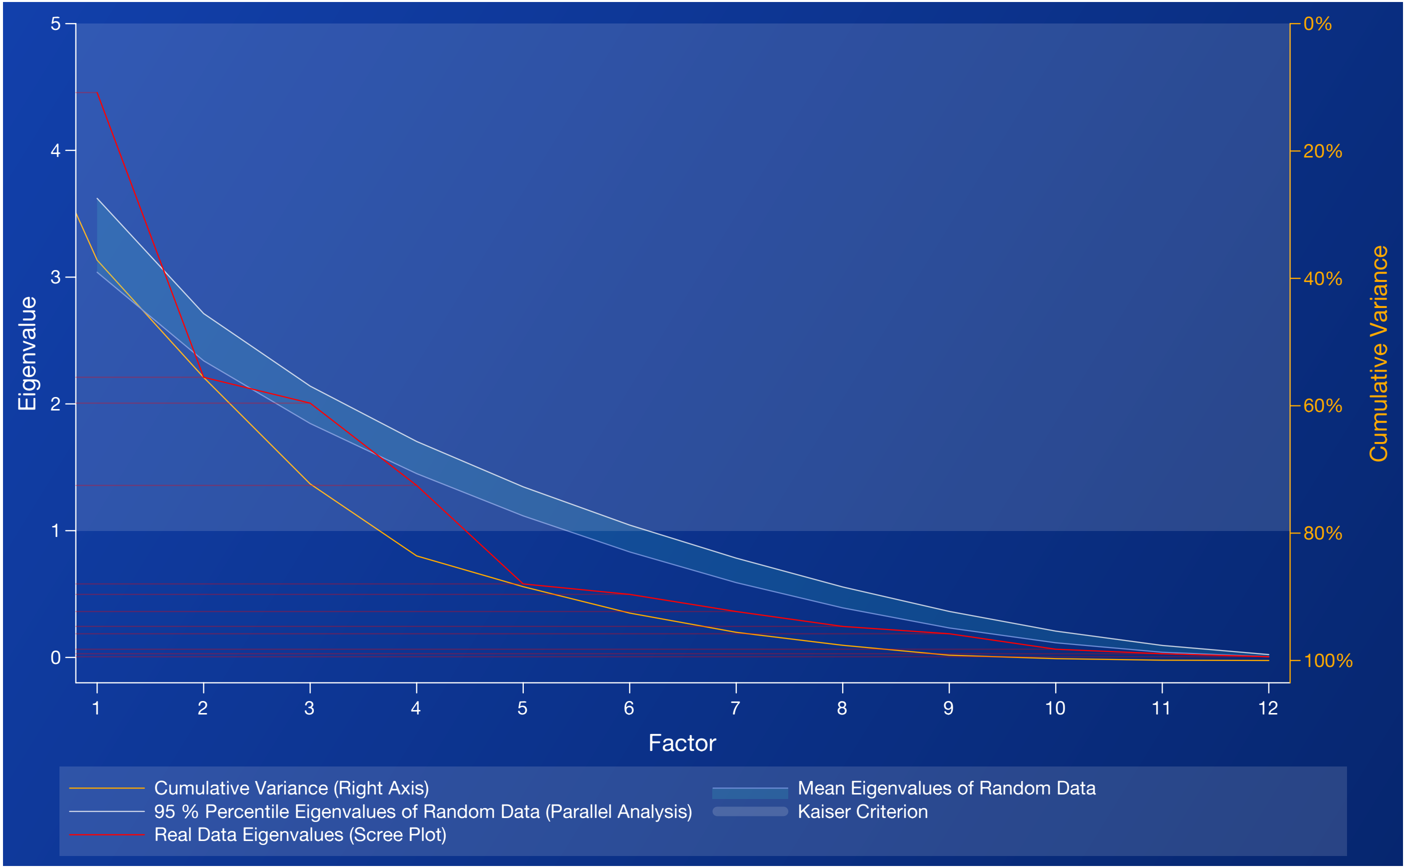

Figure 6.14. Scree plot for the Cognitive Load scale on the final FEELABUZZ questionnaire. See Figure 6.8 on page 171 for a more in-depth description.

analysis either 2 or 4 (cf. Figure 6.14 on page 186).[26] The maximum $\bar{\alpha} = .67$ was again achieved with 3 factors.

These difficulties and discrepancies point clearly to the need of a more solid analysis based on a significantly larger sample size. However, checking if the subscales we suggested based on the upstairs data remain at least plausible, Table 6.19 on the following page shows the component matrix for 3 factors on the Cognitive Load scale. The items in the first column are not coloured after the structure the matrix itself suggests but after the subscales from the upstairs study in order to provide a better comparison between the previously suggested subscales and the performance of these same scales in the second study. There are several items that could go either way according to their loadings but even when placing them in scales in favour of the previously established structure, things are not looking too rosy for those scales. The Vexatiousness scale still has the firmest footing as its good α already implied (cf. Figure 6.13 on page 185) but the FEELABUZZ data strongly suggests to also add "It was hard to concentrate while the system was producing vibrations"[27] to this scale. Note that "Using the system was effortless" loads strongly but negatively so removing it from the scale entirely remains a plausible choice (see also Table B.7 on page 442). "I was mostly unaware of the system running" and even more so "I didn't mind that the system was running" do not load strongly on the scales that they were put on after the upstairs study.

We performed a direct oblimin transform again but because of the higher component correlations, we chose a smaller parameter $\delta = 0$. While the rotation makes a simple structure more apparent as intended, this structure is farther from our original subscales than the unrotated factors (Table 6.20 on page 189 and Table 6.21 on page 190).

[26] Interestingly, 4 or 5 were actually the number of suggested factors for the upstairs study as well, even if we chose to ignore this fact.

[27] Note the slightly adapted wording.

Item	Component 1	Component 2	Component 3
The system distracted me from important tasks.	.86	−.34	−.03
I had the feeling that the vibrations produced by the system disturbed my concentration.	.84	−.20	−.27
The system frequently drew all my attention to itself.	.84	.08	.42
It was hard to concentrate while the system was producing vibrations.	.82	−.10	−.40
Using the system was effortless.	−.81	.22	.00
It was easy for me to ignore what my partner was doing in any moment.	.56	.43	.47
I recognised important events with the system, even when I was not actively monitoring it.	−.51	.44	.11
It was easy for me to push the system to the back of my mind.	.25	.87	.24
I often only realised in hindsight that I had felt something through the system.	.38	.74	−.35
I was mostly unaware of the system running.	.29	.45	.00
I payed close attention to my partner.	−.17	−.29	.81
I didn't mind that the system was running.	−.31	.15	−.72

Table 6.19. Component matrix for the first three principal components of the Cognitive Load scale on the FEELABUZZ final questionnaire. The questions are colour-coded to indicate the original subscale from the upstairs study. Light blue is Vexatiousness and orange Ability to Ignore.

Item	Component 1	Component 2	Component 3
The system distracted me from important tasks.	.91	−.03	.12
I had the feeling that the vibrations produced by the system disturbed my concentration.	.90	.06	−.15
It was hard to concentrate while the system was producing vibrations.	.86	.11	−.29
Using the system was effortless.	−.81	−.06	−.12
I recognised important events with the system, even when I was not actively monitoring it.	−.68	.26	−.03
The system frequently drew all my attention to itself.	.58	.43	.47
It was easy for me to push the system to the back of my mind.	−.26	.95	.08
I often only realised in hindsight that I had felt something through the system.	.07	.77	−.46
It was easy for me to ignore what my partner was doing in any moment.	.16	.69	.42
I was mostly unaware of the system running.	.04	.52	−.06
I payed close attention to my partner.	−.22	−.18	.84
I didn't mind that the system was running.	−.15	−.09	−.77

Table 6.20. Pattern matrix for the first three principal components of the Cognitive Load scale on the FEELABUZZ final questionnaire. The colour-coding is identical to Table 6.19 on page 188

Item	Component 1	Component 2	Component 3
The system distracted me from important tasks.	.91	.14	.19
I had the feeling that the vibrations produced by the system disturbed my concentration.	.89	.22	−.08
It was hard to concentrate while the system was producing vibrations.	.86	.27	−.23
Using the system was effortless.	−.83	−.21	−.18
The system frequently drew all my attention to itself.	.70	.54	.52
I recognised important events with the system, even when I was not actively monitoring it.	−.63	.14	−.08
It was easy for me to push the system to the back of my mind.	−.08	.90	.06
I often only realised in hindsight that I had felt something through the system.	.18	.78	−.45
It was easy for me to ignore what my partner was doing in any moment.	.32	.72	.44
I was mostly unaware of the system running.	.14	.53	−.06
I payed close attention to my partner.	−.20	−.22	.82
I didn't mind that the system was running.	−.22	−.12	−.78

Table 6.21. Structure matrix for the first three principal components of the Cognitive Load scale on the final FEELABUZZ questionnaire. The colour-coding is identical to Table 6.19 on page 188.

Component	1	2	3
1	1.00	.19	.07
2	.19	1.00	.00
3	.07	.00	1.00

Table 6.22. Component correlation matrix for the first three principal components of the Cognitive Load scale on the final FEELABUZZ questionnaire.

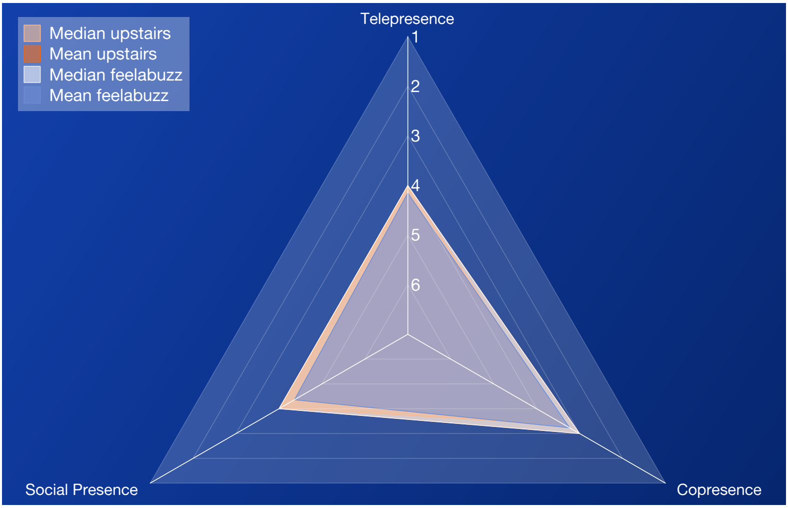

Figure 6.15. Comparison of presence scales from Figure 6.20 on page 201 between upstairs and FEELABUZZ.

Kind of Movement	Number of Mentions
Walking/Bicycle Ride	11
Deliberate Shaking	3
(Uneasy) Sitting	5
Climbing Stairs	3
Getting Up	1

Table 6.23. Kinds of movements participants reported to have felt through FEELABUZZ.

Kind of Glitch	Number of Mentions
Baseline Vibration	1
Gaps/Discontinuities	12
Vibration Pulses	1
Distortions	3
Phantom Steps	0
Vibrational Echo	1

Table 6.24. Kinds of glitches participants reported in the vibration patterns of FEELABUZZ.

Figure 6.16. Comparison of all twelve Cognitive Load items between upstairs and FEELABUZZ (cf. Table B.2 on page 435 for the exact wording for each item by abbreviation).

6.5 Discussion

6.5.1 upstairs

As can be seen in Table 6.2 on page 145 and Figure 6.1 on page 146, the total number of items over the course of the study was quite high and while the frequency of questionnaires was not raised as a concern by the participants themselves, the length and redundancy of the longer questionnaires was indeed mentioned during the interviews by most of the participants, occasionally in a quite emphatic way. We knew in advance that the high number of necessarily vague hypotheses and the comparable lack of experience as to which scales were reasonable and would give useful results were to make this study very arduous for the participants. Even where we used existing scales, there was little information in the literature regarding the influence of individual items on the internal consistency of the scales used, their item-total correlations or factor loadings and thus we had no valid basis to significantly reduce the size of the scales we used.[28]

The Social Desirability scale [518] results (cf. Table B.5 on page 438) give no indication towards a notable social desirability bias since the correlations between that scale and any other scale are not statistically significant, with the exception of a strong negative correlation with the Feelings of Closeness scale ($r(4) = -.92$, $p = .010$). A negative correlation in this case means that the stronger the social desirability of participants, the *more* feelings of closeness they will express.[29] We can offer no explanation as to why this one scale should be so much more related to a tendency of participants to answer questions in a way that they think others would approve of. We calculated the correlations for

[28] While we would love to provide such basis ourselves, we do not think the small sample sizes of our studies is a solid enough foundation to base such decisions on.

[29] Low values on the Feelings of Closeness scale mean a high agreement, while the Social Desirability scale consists of a number of yes/no questions that results in a higher value meaning a larger social desirability.

26 scales, so the probability of a chance occurrence at this significance level is about 24 %.[30]

6.5.1.1 Connectedness

Overall, the responses on the Social Connectedness scales based on van Bel et al. [550] were positive ($\bar{x} = 2.47$, $Md = 2$). Comparing the two measurements before and after system use (cf. (H4) in Section 6.2), the scales have exactly the same mean except for a slight difference in Feelings of Closeness (overall $d = -.04$) and almost exactly the same median ($r = -.15$) between before and after using the system as is visualised in Figure 6.4(a) on page 157 and can also be easily seen by looking at the first four rows of Table 6.7 on page 162. This does not mean, however, that there was a huge consistency of the averaged items between the two points in time on a per-participant basis (cf. Table 6.8 on page 163), even though the overall correlation on a per-participant and per-item basis was relatively high ($r(112) = .55$, $p < .001$). As it is not obvious which of the two correlations is more meaningful, it is not clear whether or not there was a notable change in each individual's social connectedness but it is unmistakable that there was no difference in the average amount of social connectedness.

However, when the participants were asked directly whether the system made them "feel more connected" to their respective partner, they on average did have this impression (confidence interval $CI_{.95} = [1.96, 2.71]$, $Md = 2$). A possible explanation for this discrepancy might be a different understanding of "feeling connected" by the participants compared to the definition of social connectedness that underlies the Social Connectedness scales.[31] A second question – as the last one coming from the Presence in Absence scale – might hint in this direction, too: "Communicating using the system created a connection that lasts beyond the duration of the exchange" was answered with a median

[30] $\sum_{n=0}^{N-1} p(1-p)^n = 1-(1-p)^N$ for significance level p and N scales.

[31] Although it is noteworthy that participants already felt quite often connected with their partners without using the system ($CI_{.95} = [1.67, 3.66]$, $Md = 2.5$; Figure 6.7 on page 161).

of only 3.5 and a confidence interval of $CI_{.95} = [2.47, 5.53]$ (again cf. Figure 6.7 on page 161) and the Presence in Absence scale itself shows similar results ($\bar{x} = 3.63$, $Md = 3$).

Social Connectedness, however, is often characterised by long-term effects (cf. Definition 6) and the questionnaire by van Bel et al. [550] reflects this. In fact, the Presence in Absence scale has a strong correlation with the Social Connectedness scale ($r(4) = .92$, $p = .009$, cf. Figure 6.17 on the next page) and many Social Connectedness questions, especially from the Shared Understandings and Feelings of Closeness scales, read a lot like relationship quality questions (cf. Table B.1 on page 432), even though there is no strong correlation between them and relationship strength as measured by the Relationship Closeness Inventory (cf. Figure 6.17 on the next page[32]). If this result should hold with a more extensive study and other systems, this would suggest that social connectedness as understood by van Bel et al. [550] is not a good measure for SCS and it would therefore further mean that subliminal copresence systems are not a subclass of awareness systems, at least as understood by van Bel et al.[33] In this case, the defining difference would be that awareness systems work beyond the duration of the connection while SCS do not (or to a lesser degree).

6.5.1.2 Presence

Copresence All Copresence scales yielded similar and moderately good results (Isolation/Aloneness: $\bar{x} = 3.67$, $Md = 3$; Mutual Awareness: $\bar{x} = 3.56$, $Md = 3.5$; Perceived Other's Copresence: $\bar{x} = 2.83$, $Md = 3$) that were significantly better than the scale centre when put together (cf. (H7) in Section 6.2). Combining the Perceived Other's Copresence [401] and the two Copresence scales from the Networked Minds Questionnaire [39] into one scale reveals that these scales seem

[32] Shared Understandings $r(4) = .08$, $p = .885$; Feelings of Closeness $r(4) = .08$, $p = .886$; Relationship Salience $r(4) = -.18$, $p = .728$; Knowing Each Others' Experiences $r(4) = -.20$, $p = .711$

[33] An alternative conclusion would be that upstairs is not an effective awareness system but this should be reflected on some other level such as cognitive load or user satisfaction.

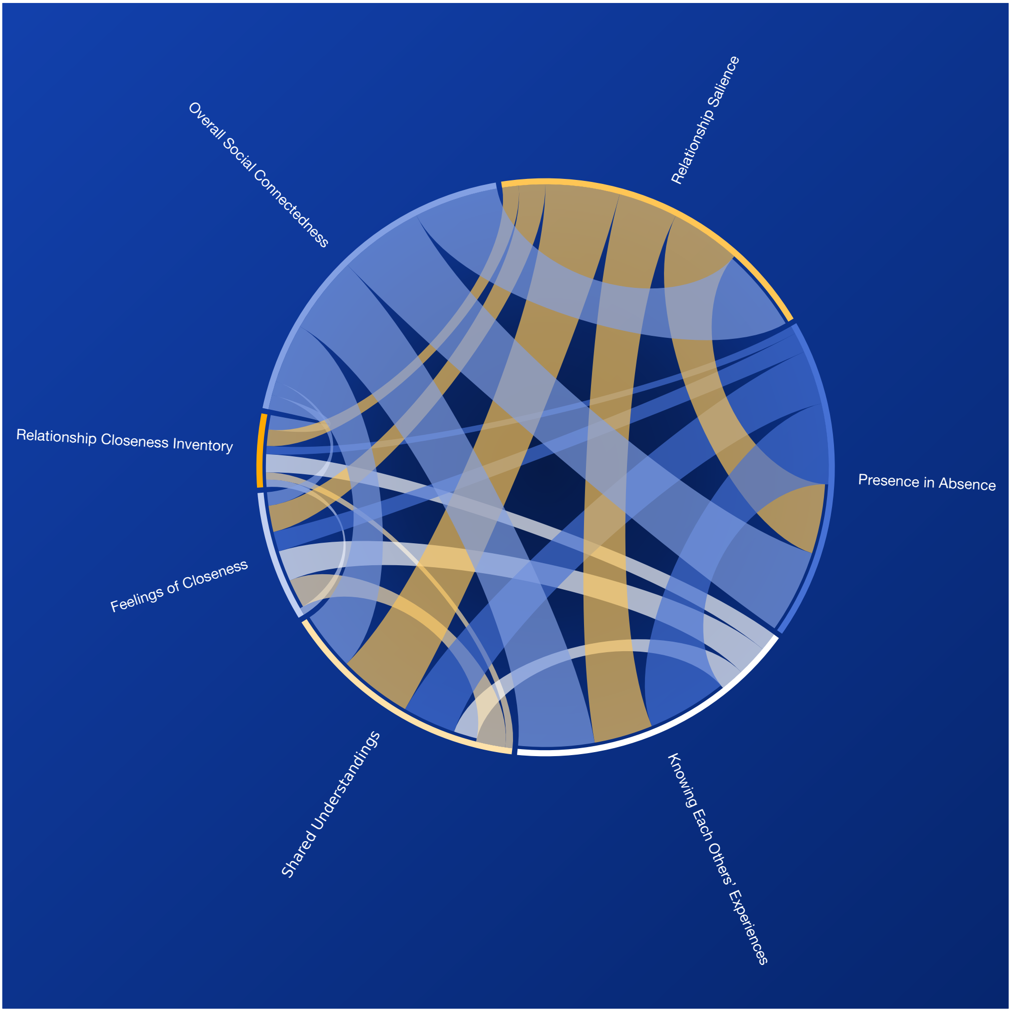

Figure 6.17. Correlations between Social Connectedness scales, Overall Social Connectedness, and the Relationship Closeness Inventory. Thicker connections mean higher correlation. Matrix diagonal elements ($r = 1$) are not represented. Only the following correlations reached statistical significance: RS–SU, RS–SC, KE–SC, KE–PiA, PiA–SC (see Figure 6.21 on page 216 for a table of abbreviations). Image created using Circos [310].

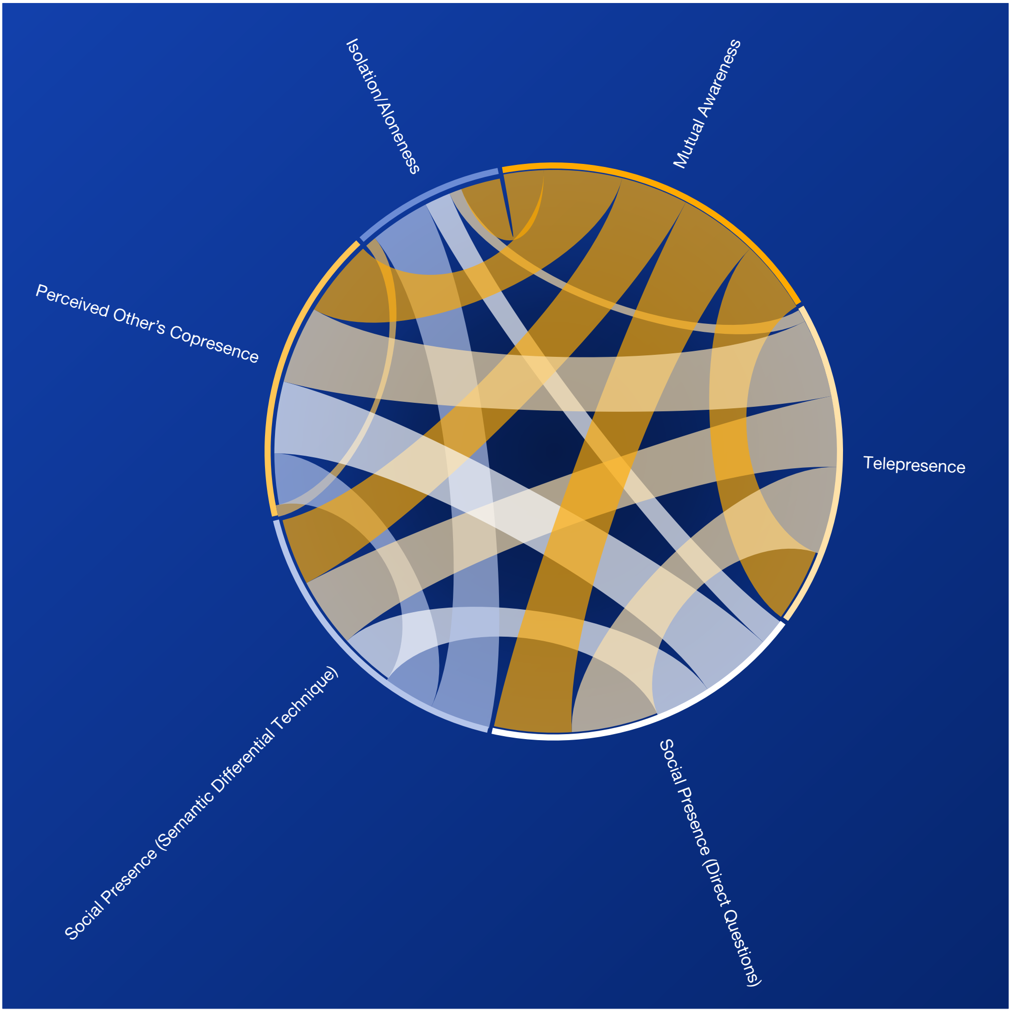

Figure 6.18. Correlations between different Presence scales. Only the correlation between Social Presence (Direct Questions) and Telepresence reached statistical significance ($r(4) = .90$, $p = .015$). Thicker connections mean higher correlation. Matrix diagonal elements ($r = 1$) are not represented. No combined scales were included.

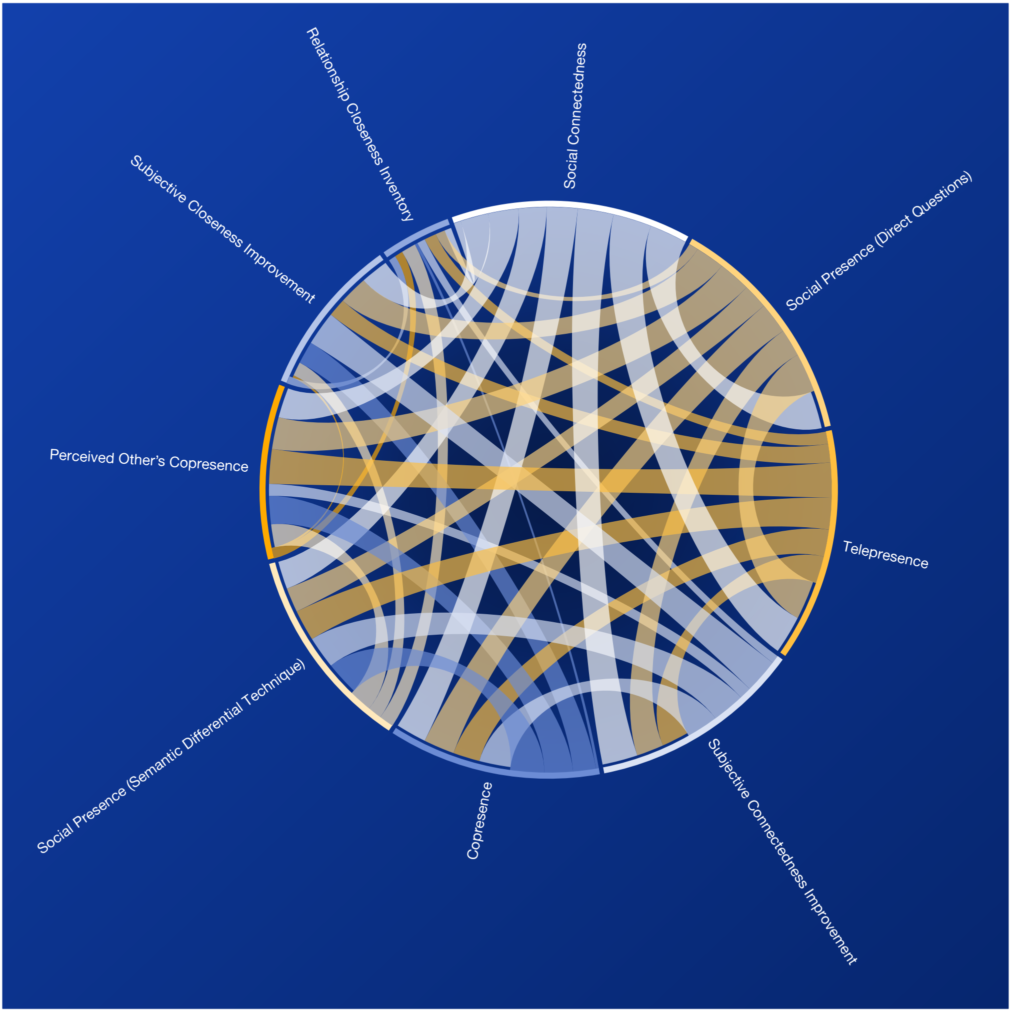

Figure 6.19. Correlations between concepts from Figure 2.4 on page 35. "Subjective Closeness/ Connectedness Improvement" means the item "Communicating with my partner using the system helped me feel closer/more connected to him/her" from Presence in Absence. Significant correlations were SP1–TP, SC–TP, SC–SP1 and between the Social Connectedness scale and the above item.

to measure the same construct (combined $\alpha = .89$) which is not surprising since – other than the name might imply – the Perceived Other's Copresence scale is *not* mainly about gauging the partner's state of copresence feelings. The two Networked Minds scales[34] also work nicely together, ($\bar{x} = 3.58$, $Md = 3$, $\alpha = .85$).

The participants indicated that upstairs managed to create a common space between the two remote locations ("I often got the feeling of sharing a space with my partner." $CI_{.95} = [1.98, 3.69]$, $Md = 2.5$), even though participants did not forget that the two locations were actually separate ("I was often aware that my partner and I were at different places." $CI_{.95} = [2.31, 4.02]$, $Md = 3$; cf. Figure 6.7 on page 161 for both items).

Social Presence The social presence indicators are relatively low as we had hypothesised, at least on the scale using direct questions (cf. Table B.1 on page 432, $\bar{x} = 4.75$, $Md = 5$, $SD = 5.03$). On the scale using the semantic differential technique [409], however, this is not the case ($\bar{x} = 3.47$, $Md = 3$, $SD = 3.73$; cf. Figure 6.6 on page 160). Both scales feature a similarly good Cronbach's $\alpha > .8$ and their combined α is .88 whereas their correlation is only $r(4) = .57$, $p = .237$. However, if the different means (paired two-tailed $t(5) = 2.94$, $p = .032$) and the mediocre correlation should mean that the two scales measure two different constructs, the item-total correlations and a PCA do not support this view (Table B.4 on page 437).

The system failed to transmit the moods from one participant to the other as shown by the combined Emotion Transmission scale ($\bar{x} = 5.17$, $Md = 5.5$) and in particular by items such as "I was influenced by my partner's moods" ($CI_{.95} = [3.87, 6.13]$, $Md = 5$). This point was further confirmed by all participants in the interviews. Participants felt that there simply was not enough information about the emotional state of the other in the signal and could not imagine there to be such information with the exception of extreme rage.

[34] Isolation/Aloneness and Mutual Awareness; referring to the Networked Minds Questionnaire [39, 40, 41]

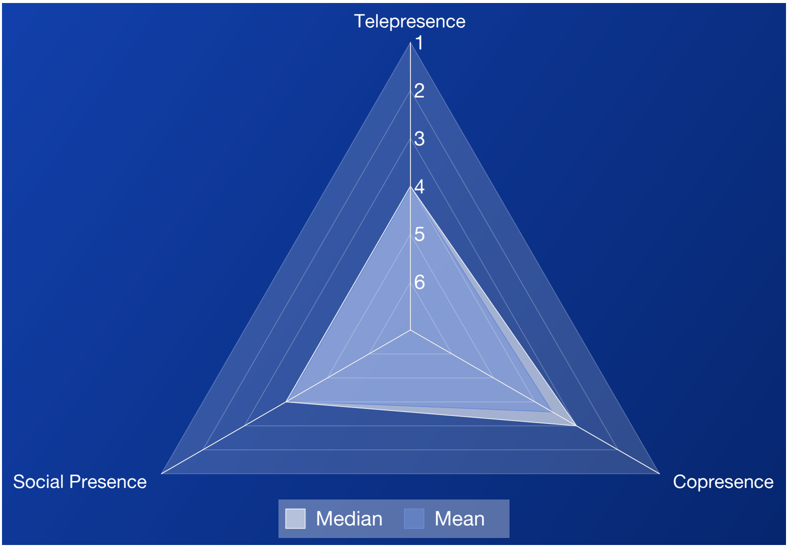

Figure 6.20. Excerpt from Figure 6.6 on page 160 showing only the Telepresence, Overall Copresence and Overall Social Presence Scales of $Q_{\beta\gamma}$.

Without having been familiarised with any of the terms or definitions of Chapter 2, one participant made a distinction between what she called “closeness” and “presence” during her interview. The way she used the terms she said that closeness was something upstairs could not convey but only something like a phone call or a Skype connection could, while a sense of presence was indeed conveyed by upstairs.

Telepresence The Telepresence scale ($\bar{x} = 4.03$, $Md = 4$) seems semantically related to the Mutual Awareness scale ($\bar{x} = 3.56$, $Md = 3.5$). The correlation between the two is strong even though only marginally significant $r(4) = .75$, $p = .083$ (cf. Figure 6.18 on page 198) and they have a common $\alpha = .90$. As shown by Table 6.14 on page 170, there are two items in each scale that do not fit with the rest but otherwise the two scales are virtually indistinguishable. The Perceived Other's Copresence scale shows no such obvious semantic similarities (but a similar correlation $r(4) = .81$, $p = .053$, common $\alpha = .86$) but this still poses the interesting question if there really is a clear difference between copresence and telepresence with SCS. It would stand to reason that for systems that try to create the presence of a remote partner, the presence *in* a common space *with* the remote partner and the presence *of* the remote partner might be one and the same thing. Since, as opposed to virtual environments, with SCS there is no sense of space other than what is transmitted from the remote partner.

This would also be in accordance with the findings of Axelsson et al. [16, 17], whose correlation between telepresence and copresence went away when immersion decreased.

Relations Between Kinds of Presence To compare the three main presence types (cf. Figure 2.1 on page 16) as they present themselves with upstairs in a more simplified picture than in Figure 6.6 on page 160, we decided to only show the combined Social Presence and Copresence scales alongside the Telepresence scale. The result can be seen in Figure 6.20 on page 201. Still no clear picture emerges, however. Copresence is highest as we had hoped (cf. (H1) in Section 6.2) but there is no hierarchy since Social Presence and Telepresence are practically on the same level. Most notably, telepresence is not on a level with copresence (cf. (H2) in Section 6.2) as suggested by the assumption that telepresence is a prerequisite for copresence, as had been hypothesised by different authors [69, 455, 498, 537]. This is all the more surprising as we had just established that it is very difficult to draw a line between copresence and telepresence with SCS.

In Section 2.2 a couple of Euler diagrams were shown. The study presented can not determine their validity in either direction but we will nonetheless see what the corresponding correlations in the data look like. Figure 6.19 on page 199 shows all correlations between items and scales that can be said to represent the concepts shown in

Figure 2.4 on page 35 (using mostly combined scales to retain clarity as much as possible).

Of course correlation can in no way be equated to conceptual overlap but we would nonetheless be tempted to question whether concepts with a very low correlation really can have a large conceptual overlap. However, apart from the Strength subscale of the Relationship Closeness Inventory (RCI) and to a lesser degree the Perceived Other's Copresence, no exceptionally weak correlations show up. This means all proposed overlaps remain plausible. Concerning the existing weak correlations, the interpretation we propose is that the Strength subscale of the RCI, whose questions mainly revolve around the influence one's partner has on one's own life (cf. Figure A.19 on page 405), simply does not measure closeness as understood in Section 2.1 (as a more long-term form of immediacy). Another aspect that might help to explain the weak correlations is that the internal consistency of the RCI in our study is poor ($\alpha = .52$; the average inter-item correlation is .004) and therefore a poor *external* correlation might not be that surprising.

The IOS scale that virtually *is* one of the definitions for closeness (cf. Chapter 2) unfortunately could not be evaluated since it was only part of the highly incomplete Q_α/Q'_α questionnaires.

We only briefly want to emphasise that the relation between social presence and social connectedness that was proposed by Figure 2.3 on page 32 and shown in Rettie [443] in particular also remains plausible and the correlation between the Overall Social Connectedness and the direct questions Social Presence is actually among the few statistically significant correlations ($r(4) = .87$, $p = .026$).

6.5.1.3 Relationship Satisfaction

As just mentioned, we could unfortunately not reasonably evaluate the Inclusion of Other in the Self scale due to too many gaps in the relevant questionnaires. Other than with the RCI Strength subscale, we did also not feel comfortable to compare single points in time because of the considerable time that had passed for some couples between the baseline period and the trying out of the system, increasing the chance that random fluctuations of relationship satisfaction significantly affect the measurements that can occur due to various reasons we could not control for. This would have been less of an issue with pairs of complete time series to compare.

We therefore considered (H3) as not tested but as far as there is a result with the RCI, Table 6.7 on page 162 clearly shows no statistically significant difference between the two measurements and very small effect sizes.

As indicated in Section 6.1, it is no real surprise that upstairs apparently failed to significantly alter the relationship. Not only was it a difficult target for such a system in the first place, it also becomes more difficult with an already good relationship that has little room for improvement.

6.5.1.4 Cognitive Load

The internal consistency of our new scale was initially poor. Although the factor analysis was a far shot to begin with, given the ratio of samples to items, and it did not turn up easily interpretable results following the usual route, we managed to find a pattern in the data that actually makes sense and that we had not thought of ourselves. While at first glance it might seem the same thing if a system annoys the user or if the user is able to ignore it, one actually can see these two constructs as separate. Just because someone's attention is drawn to a system does not mean that they are annoyed or even distracted by this.

The reports from the interviews largely matched the data from the Vexatiousness and Ability to Ignore scales.

The perception of cognitive load measured significantly lower than the scale average, with the Vexatiousness scale being markedly lower than the original scale or Ability to Ignore (Cognitive Load: $\bar{x} = 3.38$, $Md = 3$; Vexatiousness: $\bar{x} = 2.74$, $Md = 2$; Ability to Ignore $\bar{x} = 3.71$, $Md = 3$; cf. Figure 6.6 on page 160 and (H5) in Section 6.2). So while participants did not completely ignore the sounds produced by the system, they did not distract or annoy them. This is best shown by the item "It was hard to concentrate while the system was producing noises" which with all participants disagreed ($CI_{.95} = [5.62, 6.72]$, $Md = 6$; cf. Figure 6.11 on page 177).

This notion was also strongly emphasised in the interviews. Most participants said that they were aware of the sounds but they did not bother them in any way. One couple said that they often did not consciously perceive the sounds and described them as being part of the background. One participant described how he often only became aware of the system's output when the sounds had stopped because his partner had gone to bed.

Many participants commented on the presence of unwanted noises such as blips and static and described these as the most annoying part of the system, even though none said to have been more than mildly annoyed. One participant heard a constant low buzzing noise which was the most annoying sound reported by any participant. See Table 6.16 on page 180 for a quantitative assessment of these issues. The intended sounds transmitted by the system were unanimously described as pleasant. It is therefore our impression after the interviews that the cognitive load would have been even better if the sound output had been free from digital or analogue artifacts.

With the benefit of hindsight, we can try to explain why four of the items on the Cognitive Load scale did not fare so well and ended up in the Miscellaneous scale. "I recognised important events with the system, even when I was not actively monitoring it" and "It was hard to concentrate while the system was producing noises" semantically fall in between Vexatiousness and the Ability to Ignore which might account for their bad match with the other items. "Using the system was effortless" is very vague, has no clear meaning and can be easily understood as a usability question which – together with the challenging setup of the system due to the mixer – might also explain why the item was answered so differently from the rest of the scale (Table 6.9 on page 165). Finally, "I payed close attention to my partner" also has a completely different angle than the rest of the scale. It was the only item directly taken from another questionnaire and given that the item was taken from the Perceived Attentional Engagement/Attentional Allocation scale from the Networked Minds Questionnaire, this is a hint that Cognitive Load as measured by our scale and Attentional Allocation as understood by Biocca et al. [41] are quite distinct concepts.

6.5.1.5 Privacy

Since privacy was a concern that we ourselves had, it is good to see that this was not perceived to be much of an issue, with the Threat to Privacy being rated significantly lower than the scale centre ($\bar{x} = 2.13$, $Md = 1$; cf. (H6) in Section 6.2).

The only concern mentioned during the interviews was that a third party might gain knowledge over when the participants are not at home. Other concerns were not mentioned and most participants explicitly said that they were not worried to inadvertently disclose anything to their partner over the system. This is also strongly expressed in the questionnaire item "I worried that my partner might learn something using the system that I want to keep secret", $CI_{.95} = [6.07, 7.26]$, $Md = 7$ (cf. Figure 6.7 on page 161. Some also mentioned, though, that this might be different with people that they were not as close to as their partner.

6.5.1.6 Emotional Value

All participants reported at least some moments in which upstairs created a feeling of closeness and attachment towards the partner. For most, this feeling was a constant sentiment throughout using upstairs while for some it was confined to such singular moments. Many participants reported feelings of sadness or sentimentality during the dismantling of their system or a feeling of loss afterwards.

Although upstairs managed to play this emotional role for many participants, as briefly mentioned in Section 6.5.1.2 the system apparently failed to transmit emotions or moods experienced by the remote partners themselves. This can be seen from the bad values on the highly correlated ($r(4) = .85$, $p = .033$, common $\alpha = .91$; Figure 6.21 on page 216) Emotional Expressiveness ($\bar{x} = 5.22$, $Md = 6$) and Perceived

Emotional Contagion/Empathy ($\bar{x} = 5.08$, $Md = 5$) scales (Figure 6.6 on page 160) and also from unanimous reports during the interviews. Few participants saw any potential for such transmissions, no matter what quality improvements would be made or how much time users had to get used to the signals.

The reports from the questionnaires about the expectations towards the partner regarding upstairs are a mixed bag; on the one hand there are items like "I was disappointed when my partner wasn't there when I tried to contact him/her using the system" ($CI_{.95} = [1.86, 3.80]$, $Md = 3$) and "I was disappointed when it took my partner too long to respond over the system" ($CI_{.95} = [1.76, 3.91]$, $Md = 2.5$) from the Unmet Expectations scale ($\bar{x} = 4.17$, $Md = 5$),[35] indicating that there is a certain emotional investment in the system. On the other hand, "I worried that my partner felt obligated to contact me using the system" ($CI_{.95} = [4.49, 6.51]$, $Md = 6$) and "I felt guilty if I didn't respond to my partner when I perceived something using the system" ($CI_{.95} = [4.66, 7.00]$, $Md = 6.5$) from Feeling Obligated ($\bar{x} = 3.04$, $Md = 2.5$) seem to say that the participants did not act on that emotional investment (cf. Figure 6.7 on page 161 for a visualisation of the responses to the individual items mentioned in this section and Figure 6.6 on page 160 for the scale results). Similarly, some participants reported in the interviews that they felt sad or disappointed when they used knocking as a direct way of communicating through the system and their partner did not answer.

Depending on the interpretation, this emotional investment need not be an entirely bad thing, even if it means that users are disappointed when their partner is not there or otherwise does not react. Seen within the greater context of technology dependence, the new channel becomes just another of those pieces of technology that we develop some level of dependence on, exactly *because* they are useful.

[35] As a sidenote, the Unmet Expectations scale gets a dramatically increased $\alpha = .66$ (new $\bar{x} = 5.17$, $Md = 5$) if the poorly worded "I worried that I was not meeting my partner's expectations for our contact using the system" is left out (original $\alpha = .13$).

6.5.1.7 Discussion of Miscellaneous Results

As an interesting sidenote, the participants were either very bad at assessing their own loneliness or their partner's or both. "I often felt as if I was all alone" was mostly rejected ($CI_{.95} = [5.28, 6.38]$, $Md = 6$) while "I think my partner often felt alone" was mostly agreed to ($CI_{.95} = [1.98, 3.69]$, $Md = 2.5$). It is difficult to draw conclusions from this without knowing the true amount of loneliness but this result emphasises that self- and friend reports can diverge a lot even for seemingly simple assessments.

Table 6.17 on page 181 summarises the types of noises participants perceived via upstairs. As shown in Table 6.16 on page 180 and mentioned in Section 6.5.1.4, unwanted noises produced by the system itself were mentioned by all participants with varying impressions on how annoying they were perceived to be. For most of the participants, though, solving the problems causing these acoustic artifacts was a prerequisite for a hypothetical prolonged use of the system.

Some less technical problems with the usefulness of upstairs that were mentioned during the interviews were the need to wear shoes,[36] having the wrong flooring,[37] and obviously the problem of not producing any sounds when not moving, especially for extended periods of time.

The participants mentioned several possible improvements and aspects to think of for the eventuality of developing upstairs into an actual product. These were cost, power consumption, sound localisation with multiple microphones and speakers,[38] and extending the system to multiple rooms.

[36] See Section 5.1 for a possible way to overcome this obstacle by padding the flooring with felt, enabling even sneaking on socks to be heard. This is quite intrusive of course so we did not require or even suggest this in our study.

[37] We required all participants to have wooden or laminate flooring. One participant had different flooring in other rooms and thus effectively went silent on leaving the room.

[38] One aspect that was brought up by a participant, is the fact that the loudness depended too much on the distance from the microphone. This would also be alleviated by a well-designed multi-channel setup since then, moving away from one microphone would mean moving closer to another.

All participants used explicit communication in addition to the implicit context communication (cf. Section 3.1.3). The participants used simple knocking patterns (e. g. three times in a row) and no more than two semantically distinct patterns were used (for greeting and parting).

One couple had thought about using Morse code but ended up not doing it, saying that there were better ways to communicate verbally than upstairs. Another couple developed the knocking into a ritual that they would perform each time one of the partners came home.

In the interviews, participants described the sounds transmitted by the system as pleasant. One participant said she needed two days to get used to this new sound source because she at first confused the sounds with real neighbours or animals in the walls. Most participants did not think that upstairs sounded like real upstairs neighbours but this was not seen as negative. Some participants even said that they found the quality of the transmission more pleasant than that of a real ceiling as it was less muffled and more crisp while still being dampened enough as to not be distracting. During the interviews, "fitting", "right" and "natural" were words used to describe how the footsteps sounded. One participant called the sounds "80 % authentic."

When asked to speculate on the effect of a system that simulated footsteps *within* the same living space as opposed to the spatially separated one upstairs provided, some participants felt that the "poltergeist" worries we raised in Section 1.2 were plausible but none came up with those worries by themselves.

Most participants said they did not have the impression that upstairs changed their communication behaviour with other media with the notable exception of letting them know when their partner was home so that it was worthwhile to call by phone.

All participants said that they would have used the system for longer if they had been given the chance. The last couple who actually had that chance continued using upstairs for several weeks and only stopped doing so when one of the partners had a favourable opportunity to personally bring the hardware back to us and used this opportunity.

One participant said that he would not want to use such a system indefinitely because he liked being alone from time to time and with such a system he would never feel really alone.

Being asked if they could imagine using a visual analogue to upstairs, participants dismissed the idea and said that not having to look somewhere to monitor the signal was crucial. The normal *modus operandi* of upstairs as a system that was constantly running was very much appreciated, too. One participant kept the system running even when leaving home for the weekend. It was also emphasised that not having to adjust any system parameters after the initial setup except for adjusting the volume from time to time was very important. Nobody reported to have used the possibility of self-monitoring after the initial setup phase.

The importance of directness (cf. Chapter 3) was also often highlighted. The knowledge that it was actually their partner whom they heard was deemed important and when confronted with ideas of more abstract, mediated or persistently direct systems (such as presence lamps [111, 235] or footstep recognition [390]), participants rejected them as much less appealing. The importance of the non-visual channel was also emphasised unanimously as a visual signal was felt to be distracting or less persistent and easily forgotten.

As many participants used upstairs in the room they also slept in, this was a recurring topic. One participant liked lying in bed in the morning, hearing that her partner had already gotten up and also found this a good motivation to get up herself. Two participants found it comforting to hear their partners still being awake when falling asleep themselves and one of them mentioned that she felt like not going to bed alone when hearing her partner. Another mentioned a sense of safety conveyed through the sounds of her partner. One participant simply found it useful to hear when his partner had already gone to bed so he could call her late without risking to wake her.

Finally, it is good to see that the system seemed to be enjoyable. Even though the Engagement & Playfulness scale itself has an outright abysmal internal consistency ($\alpha = -1.33$), its overall value is quite good ($\bar{x} = 2.50$, $Md = 2$) and there are encouraging items like "I was excited about using the system with my partner" ($CI_{.95} = [2.35, 3.65]$, $Md = 3$), "I had fun with my partner using the system" ($CI_{.95} = [1.29, 2.04]$, $Md =$

2), and “I liked using the system”[39] ($CI_{.95} = [1.54, 2.46]$, $Md = 2$).

In the interviews, using the system was described as “comforting”, “fun”, “enjoyable”, “entertaining”, “pleasant”, conveying a “warm feeling” and as an opportunity to be closer to the partner without having to actively do something.

6.5.2 FEELABUZZ

If the choice of touch as a modality has any effect on the affectiveness of the system as we had hypothesised in Chapter 4, this does not show on the Emotion Transmission scale ($\bar{x} = 5.82$, $Md = 6$) or one of its constituent scales in any way, all of which show even worse results for FEELABUZZ than for upstairs.

The scales that differ most strikingly between upstairs and FEELABUZZ are Opportunity for Social Support and Cognitive Load, both of which are worse in FEELABUZZ than in upstairs (cf. Figure 6.13 on page 185). The lack of social support might be linked to a similar problem as the emotion transmission has in that there is not enough information to provide such a support using this link. This does not explain, however, why FEELABUZZ performs so much worse in this regard.

The correlations between Telepresence and Mutual Awareness ($r(11) = .72$, $p = .006$; common $\alpha = .70$) and Perceived Other’s Copresence respectively ($r(11) = .69$, $p = .010$; common $\alpha = .82$) are slightly less strong but this time, both correlations are statistically significant, lending further credence to the hypothesis that there is no real difference between copresence and telepresence for SCS at all.

There is no longer a statistically significant correlation between Social Presence and Social Connectedness, $r(11) = .36$, $p = .228$.

Again, the relation between the two scales meant to assess social presence remains unclear. The correlation is even weaker than for upstairs, $r(11) = .34$, $p = .258$ and the means again differ significantly ($\bar{x} = 3.54$, $SD = 1.47$ for the SDT versus $\bar{x} = 5.38$, $SD = 1.96$ for the

[39] not part of the Engagement & Playfulness scale

direct questions; paired two-tailed $t(51) = -6.56$, $p < .001$) but the combined Cronbach's $\alpha = .77$ remains high.

6.5.2.1 Cognitive Load

Given how central the goal of calmness is for our understanding of SCS, the bad performance of FEELABUZZ in terms of cognitive load is maybe the most important result in comparison to the upstairs study. Not only are the results of the Cognitive Load scale clearly worse than for upstairs but two couples actually aborted the experiment because they found the vibrations too distracting. We are going to dedicate Chapter 7 to discussing the consequences and possible technical solutions to SCS that fail to fulfil this central goal.

Other lessons can be learnt from the second round of the factor analysis of the Cognitive Load scale.

One item whose bad factor loadings should be explicitly mentioned is "I payed close attention to my partner" because this underpins our hypothesis that Cognitive Load as measured by our scale and Attentional Allocation as understood by Biocca et al. [41] are distinct concepts (cf. Section 6.5.1.4).

In conclusion, it can be said that the scales we suggested definitely need more refinement with better data but seemed to be a step in the right direction; at least more so than the adventurous factor analysis with $n = 6$ on a 12-item scale might have given us reason to hope. The least it did was to identify some items that can probably be safely omitted from future uses of the Cognitive Load scale (cf. Table 6.11 on page 167 and Table 6.20 on page 189).

6.5.2.2 Free Answers

Of the 13 participants, 5 wanted to continue using FEELABUZZ as it was, while 5 more required improvements before being willing to use the system. The top suggestions for improvement were a longer battery life and reduced heat production, the ability to switch the system off, an increased stability and reliability, and the integration of the functionality with the user's regular phone. Three participants did not want to use FEELABUZZ no matter what improvements were to be made to it. The overall impression of the free answers on the FEELABUZZ questionnaires compared to the upstairs questionnaires is that users were much more critical towards FEELABUZZ and conversely more satisfied with upstairs.

Table 6.23 on page 192 shows how often different kinds of movements were mentioned when participants were asked to write down the kinds of movements they could make out through FEELABUZZ. However, at least the 3 participants who mention deliberate shaking of the phone are an underestimation because 8 of the 13 participants reported to have used deliberate signalling (i. e. explicit communication; cf. Section 3.1.3.2) and asked for the kind of signalling, all of them said that they shook the phone.

The sensitivity or definition of the vibrations themselves were different for different people. While some described walking as a constant vibration, others felt it as a rhythmic pattern.

Among the technical issues that affected the perception of vibration patterns, disconnects due to poor reception were by far the most common, as shown by the responses to the corresponding item on the questionnaire (cf. Table 6.24 on page 192) and by free answers. These answers also show that this was the main issue that makes the system impractical for many of the participants, along with the short battery life. Distortions on the other hand were not a real problem, showing that the techniques described in Section 4.2.3 seem to work.

Our hypothesis from Section 3.1.2 on complementarity, stating that context plays an important part for these kinds of channels and that contextualisation allows the recognition of otherwise unrecognisable data is supported by the following two remarks:

> [I couldn't recognise] whole activities through the vibration pattern alone. But yesterday, I knew he went grocery shopping. I then felt how he was on his way, possibly waited in line, and went back home again.

> Walking was easy to recognise. I knew that she was out on

> her bike one day but couldn't have recognised that without already knowing it. Once I could deduct from the time of day and the sequence "rhythmic vibrations—10 min pause (tram)—rhythmic vibrations" that she was on her way home.

6.6 Conclusion

In spite of the limited sample sizes of the studies, we feel confident enough to say that it showed that permanent, synchronous, low-bandwidth channels can be applied successfully in a personal and home environment, a result that is complementary to the design process by Hindus et al. [235] who moved away from synchronous to asynchronous channels (cf. Section 3.1.2). They also found implicit presence information not to be well-received by their focus group. With upstairs we showed that people can accept implicit presence signals. From the interview, the free answers and the comparison with the Presence Light system presented by Hindus et al., we suggest that such systems should be rich, specific (Presence Light reacted to pets as well as to humans), unobtrusive, using the right modality and work as flawlessly as possible which is a particular challenge for mobile systems. Hindus et al. also do not focus on LDRRs but on general family members living elsewhere and on a work–home connection.

We have some evidence that copresence and telepresence are no distinct concepts in the context of SCS, while copresence and social presence behave the way we had hoped, with copresence being higher than social presence and them not showing a high correlation.

Users of upstairs were extremely positive about its effect and all of them wished to continue using it. For FEELABUZZ, feedback was more mixed but still three quarters of the users said that they would want to continue using the system, although half of those would only do so if technical improvements were made. With the systems we built so far – while they can have an emotional impact – the actual transmission of emotion is low.

Connectedness did not appear to be a good measure for the effectiveness of either of our SCS, calling the classification of SCS as awareness systems into question. It is too early, though, to call this question already. It might just be that the Social Connectedness scales by van Bel et al. [550] are not able to measure the subclass of awareness systems that we call subliminal copresence systems.

Privacy was not a great concern for the users of our system, meaning either that our measures to ensure data economy were effective or simply that users did not care much for their data.

Maybe most importantly, the cognitive load of SCS can be kept low but this varies strongly with the actual system and FEELABUZZ performed much worse in this regard. Also, the users of upstairs experienced a lower cognitive load than they themselves expected. When it comes to measuring cognitive load, we proposed a new way to do so that is applicable to SCS and does so with a good internal consistency, even though further refinement of our scale is needed.

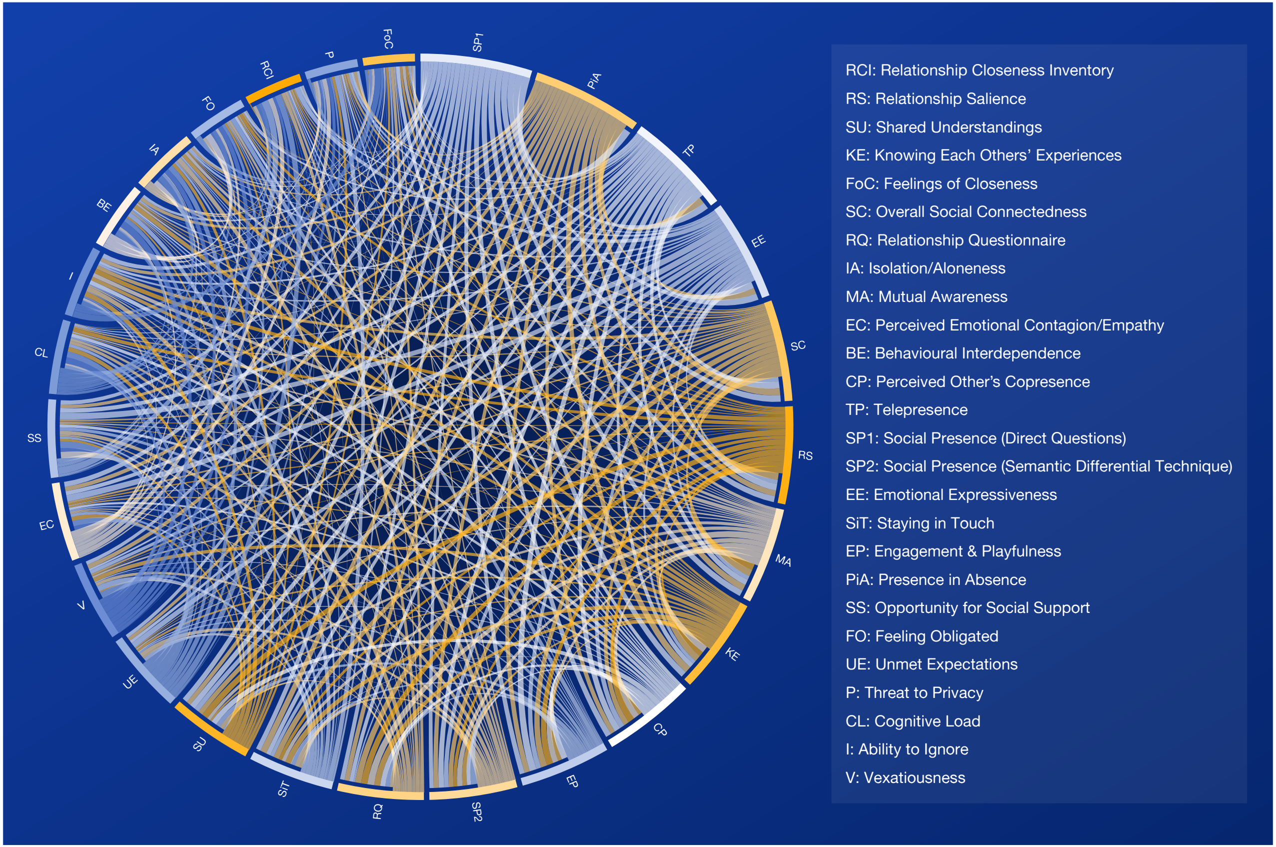

Figure 6.21. Correlations between all $Q_{\beta\gamma}$ scales. Thicker connections mean higher correlation. Matrix diagonal elements ($r = 1$) are not represented. The 37 correlations reaching statistical significance are not specially marked but $|r(4)| \geq .81 \Leftrightarrow p < .05$. Image created using Circos [310].

7 Umber – Symbolic End-User Interaction with a Subsymbolic Black-Box Predictor

7.1 Motivation

The results from Chapter 6 show that not every subliminal copresence systems (SCS) reaches its goal of being a piece of calm technology. While it is of course advisable and even indispensable to design for calmness in the first place, we believe that – depending on design constraints, such as the constraint of using unmodified smartphone hardware in the case of FEELABUZZ – it will not always be possible to achieve a level of calmness that is acceptable in each situation or required by every user. However, forcing the users to switch such systems on and off or even tuning certain parameters is in itself a burden that might not be offset by the usefulness or perceived usefulness of such background systems. Additionally, as indicated in Chapter 3, SCS inherently and deliberately have are subliminal and discreet and each system has a high likelihood to have a number of blind spots regarding the types of user behaviour it can monitor. Therefore, instead of relying on only one or two systems to already work well, we expect it to be a good strategy to bundle a larger number of SCS together in an environment to achieve a more pronounced sense of copresence.[1] This in turn means, however, that

[1] One project heading in this direction was the RemoteHome installation by Schneidler et al. [479].

the need to control so many individual devices in order to minimise vexatiousness and privacy intrusion is a liability that seems especially inconsolable with that primary aim of calmness.

Markopoulos [348] writes on the topic of *seamful design* [71] the following:

> Unless in very simple cases, one cannot assume that awareness information presented is reliable. Erroneous information or system behaviour can arise because of malfunction of context sensing technology or a problem in the network, etc. While one can hope that technology improvements can reduce these problems, another source of erroneous information is harder to prevent. In cases where context sensing is involved for capturing awareness information, it is typical that heuristics are applied relating what information is sensed by technical means and what inferences need to be drawn.

Investigating a similar line of thought, Romero and Greenberg [456] found that providing fine-grained awareness controls (or rather *focus* controls as the authors use the focus/nimbus model shown in Figure 2.2 on page 27) did not have a great success with the users of their system because the distraction management did in itself distract and interrupt too much. We believe this to be a general problem of complex SCS and other ambient displays, as well as a problem for scenarios where a large number of simple SCS with continuous, remotely controlled or proactive behaviours are present. While in principle it is good and important to design configurability into these systems, in practice it is probably not feasible to fine-tune such calmness parameters all the time without taking more time and effort than can ever be saved by the better adjusted balance between calmness on the one hand and bringing important information to the user's attention on the other. We therefore think that this is an area where automated system can help.

There is an additional obstacle, though. Although machine learning techniques exist to adequately assess a user's status within a given domain [77, 107, 149, 312, 470, 593], these trained classifiers are usually black box models. The information they contain is hidden, mostly these days in the form of support vectors or the weights of a neural network. We believe that this can pose a problem when introducing such systems into people's daily lives and letting these systems decide about such important matters as interruptibility and privacy. We therefore believe it to be important for these systems to be comprehensible by the users. There are techniques to approximate the behaviour of a subsymbolic black box classifier by a set of symbolic rules that are then more likely to be understood by humans. These techniques are called *rule extraction* (cf. Section 7.7.2).

Once we have this symbolic abstraction of a classifier, this opens another very interesting prospect: using a technique called *knowledge insertion* (cf. Section 7.7.3), users could be given the ability not only to understand the black box model but also to manipulate it by manipulating the symbolic approximation. Such systems are called *hybrid transformation architectures* (cf. Section 7.7). This would then form a closed human-in-the-loop workflow between the symbolic and the subsymbolic domain, enabling the user to directly change or even prevent undesired behaviour by the system, while the system can benefit from the considerable world knowledge the users can provide (cf. Figure 7.24 on page 261 and Figure 7.23 on page 260). After all, it is their private lives that the system is trying to get a handle on so they can safely be called domain experts in this regard.

We believe that the lessons learnt from such systems might find a wider application than SCS themselves. After all, people today are indulged in an ever increasing flood of information we need or wish to sort through. Who did not already wish for a secretary to presort the important mail from the unimportant, the interesting news from the noise and hold back calls until a later time when we can better appreciate them – who but the lucky few who can afford one?

The idea to build software to do this job is of course not new and fiction is ripe with agents performing their job perfectly [64, 487, 578]. However, fiction as well as daily life can also teach us that agents, human and non-human alike, that are always around are prone to cause conflict.

The first iterations of such automatic assistants are already in use, sometimes without us noticing them as such: most of us use spam filters on a regular basis and most of these filters use some kind of low-level automatisms combined with high-level user control. Proactive intelligent personal assistants such as Google Now[2] are already context-aware to some degree. Yet the level of sophistication is still quite low. Assistants need to adapt to a certain user and handle context in a more complex way than can be provided by easily understood rules.

With the advent of pervasive computing and ambient intelligence and with the ever increasing amount of information streams that need personalised filtering, systems using machine learning are going to appear increasingly in everyday applications. This means that they are going to affect users who were never in direct contact with such adaptive systems before. While this is no problem for systems that "just work" such as web searches, fingerprint scanners or face recognition in photo applications, it might pose a problem for systems that try to tackle problems that have a less comfortable ratio of problem complexity to available training data. Even the long-standing application and prime example of machine learning for end-users, the Bayes mail filtering, suffers somewhat from its unpredictability and attempts have been made to increase the comprehensibility of the filter decisions [86]. Context-aware applications on mobile platforms or in smart rooms, tailored to each user's individual behaviour patterns, including custom locations, daily schedules, preferences, personality traits and biometric features such as gait, present a much more difficult class of challenges. Finding a one-size-fits-all kind of solution might not in general be possible. This is especially true since the system cannot ask for an arbitrary amount of feedback from the user for its supervised learning. Otherwise, the assistance of the system would be more of an impediment than it can actually contribute to e. g. reduce interruptions of the user or save time through context-aware services. A system, however sophisticated, is useless when users switch off the adaptivity feature because it does not seem to help and they simply do not understand why it does the things

[2] http://www.google.com/landing/now/

it does or how they can make it do the things they want it to do.

As initially marked out, the best-performing machine learning algorithms are subsymbolic in nature, adjusting weights, probabilities or high-dimensional vectors in order to duplicate complex real-world patterns. It is very unlikely that end-users can be expected to interpret, let alone purposefully change these parameters. Even experts in these models are hardly able to do so. Symbolic, rule-based systems on the other hand, albeit still being artificial and often hard to grasp, are much more similar to natural language and actually find application in tasks such as simple mail filtering or smart playlists (cf. Figure 7.2 on page 224).

We therefore suggest to increase the accessibility and transparency of black-box algorithms with comprehensible and modifiable symbolic representations (a process called *white-boxing*; cf. Section 7.7.1.1). This addresses the foreboding acceptance problems we see for black-box systems when they start governing important aspects of people's personal lives. We believe that this can be achieved by giving users a sense and a means of control beyond (but still including) the passive learning-by-example capabilities such systems usually have.

With the advent of intelligent personal assistants on mobile devices, the still somewhat cryptic rule sets of early systems might even eventually inform these assistants' language output, giving them the ability to explain themselves and their behaviour.

7.2 Overview

In this chapter, we will present Umber, which is our conceptual framework within which we developed two prototypes for such systems as the one we described above. More precisely, Umber consists of two prototype systems that can learn user contexts on different platforms. The first one is called Lumber and runs on a laptop computer, using the built-in sensors to assess the context and taking the user's instant messenger status to label these contexts automatically from the fixed set of statuses the instant messenger provides. In the second iteration we moved to a mobile platform. This prototype is called Cucumber and

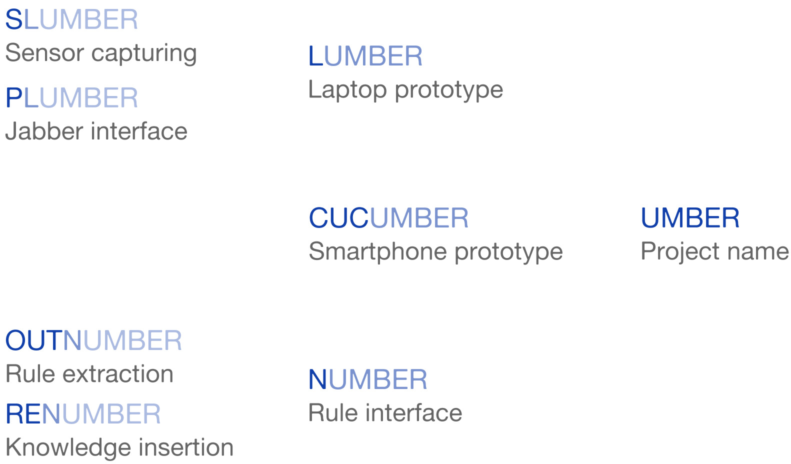

Figure 7.1. Overview of Umber software modules, their basic functions and the overarching naming scheme.

runs on smartphones, again using the sensors built into the devices. On this platform, no implicitly provided contexts are available, so contexts have to be labelled explicitly. Any number of user-defined labels can be assigned to any given time interval. Finally, the rule extraction, editing and knowledge insertion takes place in a module called Number. See Figure 7.1 for an overview of all software modules that constitute Umber.

7.3 Comprehensibility

As our goal is to design systems that make complex and potentially complicated relations between sensor measurements on the one hand and states and machine behaviours on the other hand approachable by end users without special training in understanding the system, *comprehensibility* or *understandability* is a core concept that we will briefly present in the following.

7.3.1 Comprehensibility of Rule Sets

Tickle et al. [531] define the comprehensibility of a set of rules using the number of rules and the number of antecedents per rule. Rule comprehensibility in practical terms (i. e. as a cost function for optimisation) is almost always such a simple measure that is inversely proportional to the number of conditions per rule and the total number of rules [157, 249, 296, 297, 534]. Neumann [394], on the other hand, points out that there is no measure of actual comprehensibility that is independent of the individual user.

Regarding rule comprehensibility at the design stage, it is assumed that different rule types have a different comprehensibility (cf. Section 7.7.2.2 for an overview of different rule types, especially propositional if-then rules versus oblique rules) but there is little work that compares the actual ability of users to understand different types of rules for a fixed accuracy, how comprehensibility decreases with increasing rule size, or what effect different logically equivalent rearrangements of rules have [420]. Huysmans et al. [250] looked at some of these questions; in particular they compared how well and quickly rules were understood when presented as decision tables, decision trees and written rules. They also showed that oblique rules are indeed difficult to understand, even though for small rule sets users could handle them better than the authors had anticipated and they substantiated the long-standing assumption that more complex rule sets are less comprehensible.

What Huysmans et al. [250] also did was to define comprehensibility precisely enough to investigate it whereas before the term had been used in a less clarified and more equivocal way [420, 534]. Huysmans et al. challenged their study participants with a number of rule-related problem solving tasks and measured the accuracy and response time of the participants and asked them for their confidence. It was found that this subjective measure correlated well with the actual accuracy and users therefore seem to be able to assess their performance and the difficulty of the task. Decision tables performed best across the board, whereas textual representations were found to be perceived as

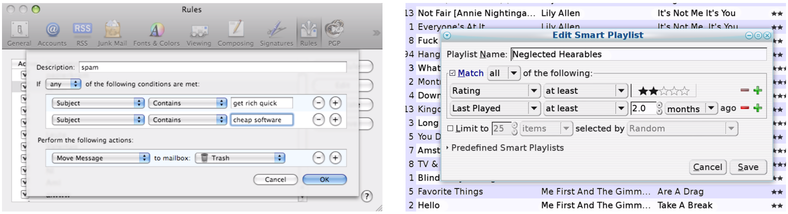

(a) Mail User Agent (b) Music Player

Figure 7.2. Examples for rule interfaces in common software, i. e. Apple Mail (a) and the Songbird (now renamed Nightingale) music player (b).

the least easy to use (apart from oblique rules)[3] but actually did not perform clearly worse than decision trees.

For our Number prototype we had already chosen the conventional representation of rules within the context of graphical user interfaces (cf. Figure 7.2) which is a hybrid of decision tree and textual representation. For a future iteration it will surely be interesting to evaluate decision tables, though, and see if they also perform well when it comes to manipulating rules instead of just interpreting them.

7.3.2 Comprehensibility of User Interfaces

Gajos et al. [164] found that user satisfaction with UIs is determined both by the *predictability* of a UI's behaviour and its *accuracy*, that is the effectiveness with which the UI tries to offer context-dependent functionality in a dynamic way. While they see both aspects as opposing each other, adding explanation capability to a system is a way to have dynamic systems that are still comprehensible and thus predictable.

[3] Oblique rules could not be compared to the other three representations performance-wise due to their different expressive power.

7.4 Context Awareness

Context has long been in the focus of ubiquitous computing research, which lead to the subfield of so-called context-aware systems. Context awareness describes the idea that devices can sense their environment and adapt their behaviour accordingly. Mobile context-aware devices, however, at first primarily regarded location [476]. This quickly began to change, though, as devices became ever more mobile and sensors smaller, cheaper, less power-hungry and therefore more widespread in consumer devices [1, 30, 113, 122, 388, 475, 476, 562].

The most commonly quoted definition for context is the one by Howe [244], who defines context to be “[that] which surrounds, and gives meaning to, something else.” Brown et al. [61] distinguish between the user’s context and the application’s context and say that the context of the user has to be the more important defining aspect for context-aware applications since the context of the application includes the user input and therefore almost all software would have to be called context aware. This is a distinction Crowley et al. [97] also make. Of course, user input or information invisible to the user can still be part of the information a context-aware application uses. Other authors do not share this distinction in the first place [475].

In any case, what we essentially try to predict in the machine learning systems for which we will be using the rule interface is the user’s context.

Since this is the same kind of context we deem so important in communication and in contextualising presence cues, synergies between SCS and context recognition systems are conceivable, for example in sharing sensors or feature extraction, but even when it comes to directly transmitting detected status or similar context condensates as symbolic cues to a remote person [435, 473].

7.5 Lumber – Laptop Prototype

As an experimentation testbed and first proof-of-concept system we wished for a context-aware system that has a differentiated interruptibility status, is unobtrusive in that it integrates nicely with every-day computer usage and translates well to other applications such as mobile platforms. Therefore we chose the availability status of an instant messenger as classes that are to be learnt by the subsymbolic classifier. The idea there being that initially, users set their status normally and this is directly used as class labels for the supervised learning algorithm. At a later state, the classifier can then be used to set the labels itself. As sensors and data sources, only built-in capabilities of standard laptop computers were used.

7.5.1 User Status

7.5.1.1 Manual Status Changes

Many people these days share their status in a variety of ways: e. g. via microblogging, social networking sites, photo sharing, location sharing, and instant messaging status. Some of this information is not generated for a purpose different from sharing one's status or it is generated automatically, but oftentimes people are willing to put a surprising amount of effort into keeping their status up-to-date. Rittenbruch et al. [446] discuss this phenomenon in a more general sense when introducing the concept of *intentionally enriched awareness*. This is the idea that users that are observed by awareness systems should have an active role in providing information about their situational context.

Dey and de Guzman [114] found out that their IM users in their focus groups found status information important, even when they were not interested in communicating with someone. Nardi et al. [393] had come to similar conclusions using different methods.

While status information maintained by the users themselves can in theory be the most reliable form of status information to base context-dependent behaviour on, in practice not all users can or want to carry the burden of keeping their status up-to-date all the time and automatic systems become necessary. Another aspect might be that user status either takes the form of natural language which would have to be interpreted by software or status comprises a number of discrete steps that usually denote availability or interruptibility. While this might just be the information needed for some applications, for others this will be of less use than information about e. g. location, physical activity or the proximity of certain people.

7.5.1.2 Automatic Status Changes

Existing automatic status change systems are widespread but usually very primitive and take the form of "set to *away* after *n* minutes without mouse and keyboard activity and to *not available* after *m* minutes" (cf. Figure 7.3 on the following page for examples). Compared to other sensors that often would be available, these inactivity-based cues are unnecessarily crude and have an obvious inherent lag. On the other hand they are very reliable compared to other sensors and even compared to a stochastic learning algorithm that would only be given mouse and keyboard activity as input. In fact, they might be too reliable, compromising plausible deniability (cf. Definition 13 on page 73). If there is even a very small possibility of misbehaviour of a mechanism (as a stochastic algorithm always provides), plausible deniability is granted.

Predicting user status and interruptibility for desktop usage has been shown to be possible before [149]. In related fields this has also been demonstrated practically, for example in ambient intelligence [77, 312, 470] and on smartphone platforms [107, 593].

7.5.2 Architecture

The main user interaction for collecting data is done using a conventional instant messenger. As instant messaging platform, we chose to use XMPP (the *Extensible Messaging and Presence Protocol*), more commonly known as Jabber [254]. This decision was made because XMPP is an open protocol, it is XML-based and thus human-readable, it

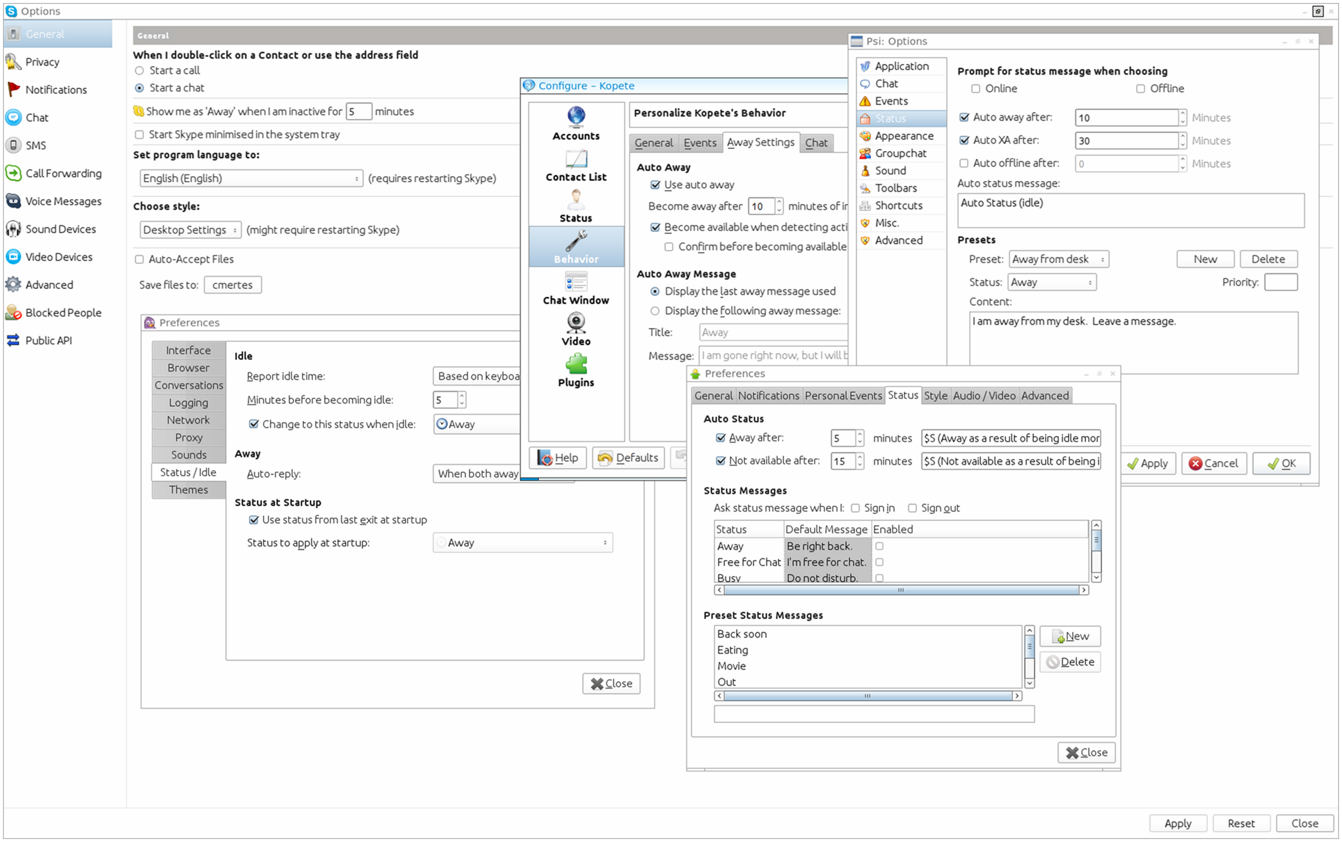

Figure 7.3. Screenshot of the very similar auto-away preferences of the Skype (http://www.skype.com/en/) voice-over-IP software and the Pidgin (http://www.pidgin.im/), Kopete (http://kopete.kde.org/), Gajim (http://gajim.org/?lang=en) and Psi (http://psi-im.org/) instant messengers (from left to right), all in their Linux version.

is in widespread use, a large number of clients exist[4] and also a large number of implementations for various languages.[5]

We also enabled an XMPP infrastructure for the inter-process communication (IPC) between individual modules of Lumber but for now we are still using an alternative, more robust means of IPC via the file system using comma-separated values (CSV) [255].

The user is informed about automatic status changes from the system using desktop notifications [220, 325] (Figure 7.15 on page 243). Negative feedback for a wrong status change can be given by just resetting the status via the Jabber client.

The GUI for controlling the data collection activity and the automatic setting of the status by the system can be seen in Figure 7.14 on page 241. This is also where the machine learning and rule extraction is initiated and the rule viewer and editor is called. The latter is shown in Figure 7.35 on page 293.

Figure 7.23 on page 260 gives an overview of how these components exchange information. As can be seen, there are two basic loops in the information flow. One where the data from Plumber (the status; cf. Section 7.5.2.2) and Slumber (the sensor features; cf. Section 7.5.2.1) are fed into the subsymbolic classifier after a simple preprocessing step that includes feature selection, normalisation, optional class reduction (cf. Figure 7.4 on the following page), string feature extraction (cf. Section 7.5.2.1) and conversion into Weka's Attribute-Relation File Format (ARFF) [444],[6] whereas the output of that classifier constitutes the status prediction for Plumber. Section 7.7.2.7 goes into some more detail on the rule extraction and knowledge insertion part.

7.5.2.1 Slumber – Lumber Sensor Framework

Table C.1 on page 454 lists all the features used by Lumber along with the description that is given to the users, too. There is a total of 21 video

[4] http://xmpp.org/xmpp-software/clients/

[5] http://xmpp.org/xmpp-software/libraries/

[6] There is also an alternative LibSVM output module to interface with LibSVM's tools more easily (cf. Section 7.5.3).

free/DnD	present/away	free/away/DnD	regression	
free	free	free	free	1
online	online	online	online	2
away	away	away	away	3
not available	not available	not available	not available	4
busy	busy	busy	busy	5

Figure 7.4. Possibilities for class reduction in the preprocessing step. Five classes might be too fine-grained for some applications. The first column shows a reduction to two classes based on interruptibility, the second are two classes as well but based on physical presence, the third collapses the gradual distinctions free/online and away/unavailable into one class each, leaving three classes and the fourth column represents an ordering that can be used for regression which might also lead to less severe classification errors.

features, 21 audio features, 4 time-related features, 2 related to mouse and keyboard activity and a variable number of window title features. These features will be described in more detail in the following. The relatively high number of features stems from the fact that the last five time steps $n-5$ to $n-1$ and their mean are included as features in time step n to reduce the susceptibility to noise where necessary and to enable a kind of short-term sensory memory. Sensor values and user status are polled every 5 s by default which results in a time span of 30 s being covered by these features.

If these repeated features are not counted, there are only 3 video features, 3 audio features and 1 mouse and keyboard activity feature. Earlier prototypes contained a much larger number of features, especially more audio features such as band energies for various frequencies. We decided to cut down on these because we felt that this large number of features would require too many training examples to work and make decision trees too varied and too complex. The decision to introduce the past 5 time steps in addition to the usual averaged features was made because we felt that the exact behaviour over time of these features might also be important and what these features provide in combination is basically a six-sample plot of the dynamic of a given feature.

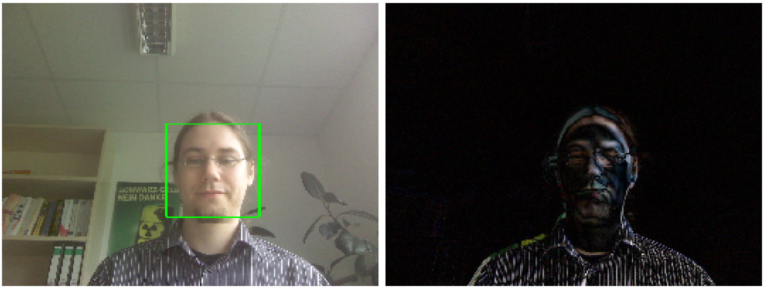

(a) (b)

Figure 7.5. Face detection to recognise user presence and a difference image as motion detection.

Visual Features Laptop computers virtually universally come with cameras built into them that constantly face the user. We use this to assess the presence of the user and the physical activity to some degree. To extract features from the camera, we use the iceWing image capturing and processing framework [340, 341]. We use it to detect the number of faces in the camera image using the Viola and Jones [555, 556] algorithm,[7] the overall brightness, and the squared average of the difference image, i. e. a measure for the overall movement in the picture.[8] For all three features, the last five time steps and their averages are also included.

Audio Features From the microphone data, only spectral features are used (illustrated in Figure 7.6 on the following page). The spectral centroid is extracted which is a measure that approximates the perceived

[7] The output of this module can be seen in Figure 7.5(a).
[8] Figure 7.5(b) shows the output of the difference image module.

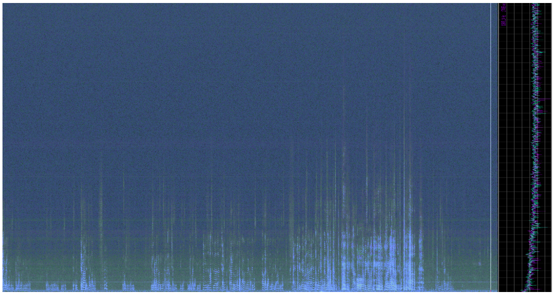

Figure 7.6. Audio energy in different bands is used to spectrally analyse the acoustic input.

brightness of a sound [193, 355], as well as the spectral spread [589], both as a measure of the signal's timbre. We also extract the overall energy in a given window.[9] As with the visual features, the recent past of the audio features was included as separate features in their own right.

Time Features Although time progresses linearly and is usually also linearly represented in data structures (e. g. Unix time [211]), there are repeating characteristics like the passing of night and day, weeks, months and years that are used in the representations of virtually all calendar systems [307]. Using such common clock and calendar notations to represent time features has the disadvantage that one-dimensional measures such as the time of day wrap around suddenly – in this case at midnight – and other regularities such as the weekday/weekend cycle and the seasons have to be modelled explicitly. We therefore find a circular, two-dimensional representation better suited – a simple Euclidean distance function can be used to get a measure that avoids any discontinuities and locally approximates the more appropriate angular distance, while maintaining a reasonable global behaviour (cf. Figure 7.7 on the next page).

[9] The effective window size used was 5 seconds, sampled with 44.1 kHz.

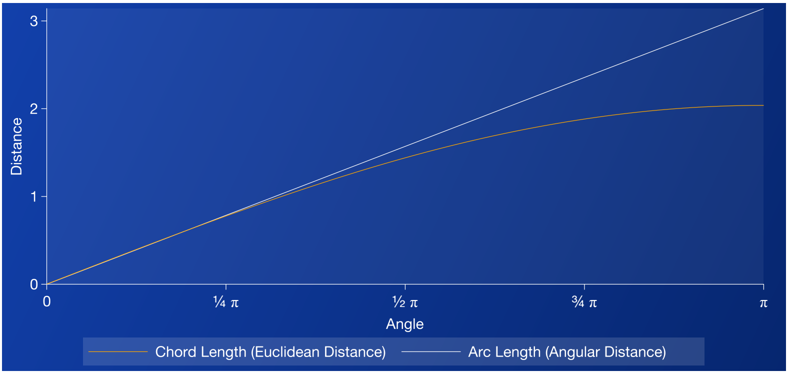

Figure 7.7. Comparison of angular and Euclidean distance for two points on the unit circle (the radius does not change the ratio of the two functions) with the angle between them ranging from 0° to 180°.

We decided to use three circles in parallel to represent time: one comprising one day to catch repeating characteristics that depend on the time of day, another one comprising one week for weekly features such as the weekend and one circle comprising a whole year to be able to learn season-specific behaviour where this occurs or other annual particularities. This last circle has been left out for our short-term evaluation, leading to 4 time features that were actually used.

These features are computed as follows. For the time of day $\boldsymbol{t}$, with $h \in \mathbb{R}$ being the hours since midnight:

$$t_1 = -\cos\left(\frac{h-6}{12}\pi\right) \tag{7.1}$$

$$t_2 = \sin\left(\frac{h-6}{12}\pi\right) \tag{7.2}$$

Midnight is at the bottom of the unit circle, midday at the top, 6 o'clock in the evening is to the right and 6 o'clock in the morning to the left (cf. Figure 7.8 on the following page).

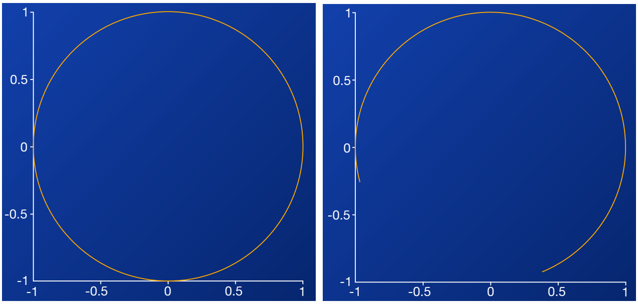

Figure 7.8. The complete unit circle on the left represents the time range from 0:00 to 24:00, the arc on the right the range from 5:00 to 22:30.

The day of the week $1 \leq \delta < 7$ is represented as $\boldsymbol{\theta}$ along the same lines with $l_{week} = 7$.

$$\theta_1 = -\cos\left(\delta \frac{2\pi}{l_{week}}\right) \tag{7.3}$$

$$\theta_2 = \sin\left(\delta \frac{2\pi}{l_{week}}\right) \tag{7.4}$$

Likewise, the date $\boldsymbol{\tau}$, with d being the day of the year and l_{year} being 365 or 366 depending on the length of the particular year in days:

$$\tau_1 = -\cos\left(d \frac{2\pi}{l_{year}}\right) \tag{7.5}$$

$$\tau_2 = \sin\left(d \frac{2\pi}{l_{year}}\right) \tag{7.6}$$

There is no normalisation regarding the position of the seasons on the unit circle in this equation as this is not perfectly possible with a simple rotation in the Gregorian calendar (cf. Figure 7.9 on the next page).

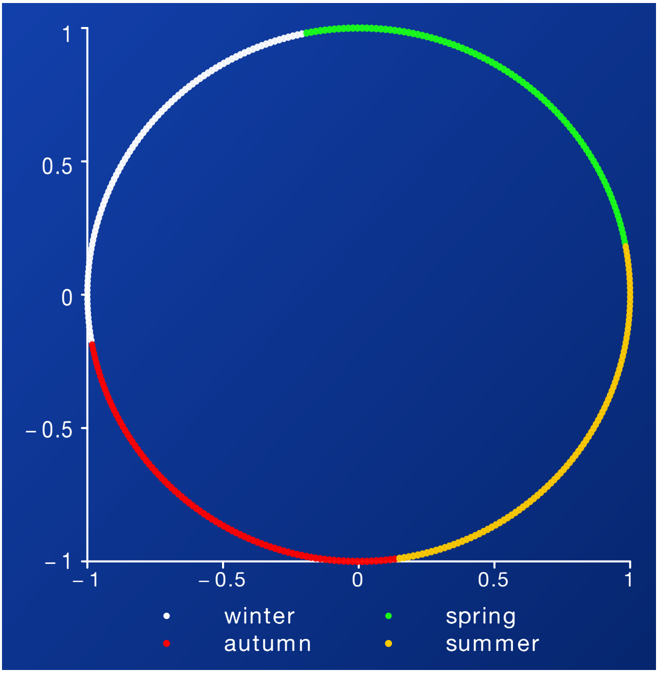

Figure 7.9. The position of the astronomical seasons day by day during a typical year on the Northern hemisphere using the *day of year* feature.

Activity Features To get a sense of the general activity of the user on a desktop system, the time elapsed since the last mouse or keyboard activity can be used. This information is provided by the X11 Screen Saver extension library (`libXss`) [160]. However, the library only provides a single value representing the combination of both activities. So any activity of either mouse or keyboard will reset the timer. Any possibility to intercept all mouse and keyboard activity manually to get separate values for both types of activities would have been very technically invasive[10] and was therefore dismissed.

The data this sensor delivers has a very characteristic, irregular sawtooth-like look. Cf. Figure 7.10(a) on the facing page for a plot of an exemplary dataset. As the time since the last activity can grow without bound, linear normalisation to a fixed interval turned out to be problematic since the most extreme values have a large effect on the overall variance whereas normalisation to unit variance can leave quite extreme values in the data. For future implementations, it might therefore be expedient to introduce some nonlinearity that limits the values this feature can take (cf. Figure 7.10(b) on the next page).

Window Properties As the main activities performed by users tend not to be identifiable by any means other than analysing the screen content itself (imagine telling what someone is doing on their computer just by looking at them but not at the screen), we decided to use the currently active window as an indicator for the current activity of a user. The X server stores metadata about windows in attributes and properties [472]. As per the Freedesktop window manager specification [553], the currently active window is stored in the `_NET_ACTIVE_WINDOW` property of the root window. From the active window thus determined, the three properties application title (`_OB_APP_TITLE`), application class (`_OB_APP_CLASS`) and window title (`_NET_WM_NAME`) are read.[11]

In a production system, a fixed number of generic string features (such as string kernels for SVMs or string distance metrics for kNN) might be used; however, this would come at the cost of highly reduced comprehensibility of the rules eventually produced. In any case, for the

[10] These possibilities include a screen-sized invisible window that stays always on top and passes along all mouse and keyboard activity and rewriting the X Window mouse and keyboard drivers.

[11] To get an appreciation of the kind of strings usually stored in these three properties, on any X window system a command line such as the following can be used while changing windows (tested in bash 4.2.25 and zsh 5.0.0 with X.Org 1.11.3):

```
watch 'xprop -id `xprop -root | grep "^_NET_ACTIVE_␣WINDOW"| cut -
f5 -d"␣"` | grep "_NET_WM_NAME\|_OB_APP_CLASS\|_OB_APP_TITLE"'
```

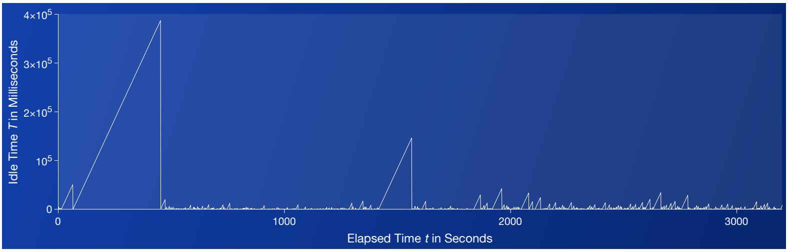

(a) Raw Data

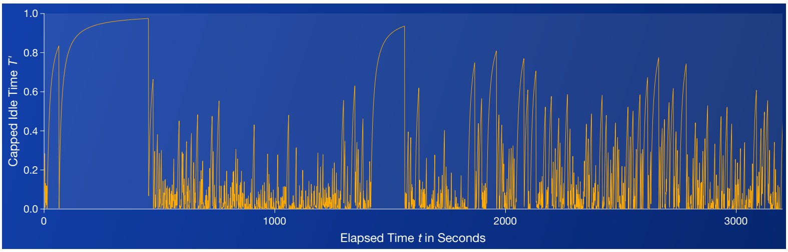

(b) Capped Data

Figure 7.10. Exemplary activity data for the mouse/keyboard idle sensor. (a) shows the raw data and (b) a suggested derived feature $T' = \frac{t}{\tau + t}$ for the same data that has the asymptote $T' = 1$. In this case, the time constant takes the value $\tau = 10\,\text{s}$.

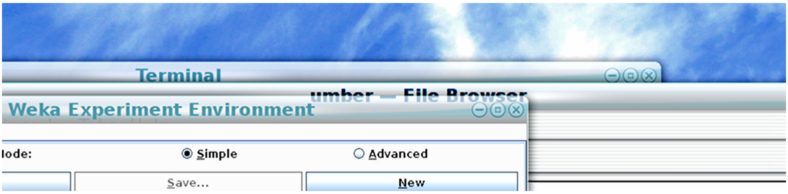

Figure 7.11. Several overlapping windows, the rightmost of which is active. The window title feature therefore would have the value "umber – File Browser"; application name and application class would in this case both be "Nautilus", designating the name of the Gnome Desktop Environment's standard file browser.

prototype system, we allowed ourselves a variable number of token features which implies the necessity to rebuild the whole model every time new data points are added. This was less of an issue given that no incremental learning was implemented anyway. It makes tree consistency even more difficult, though.

Token features means that the strings are split into tokens separated by spaces and punctuation. For example, the terminal window name `cmertes@Saffron:~/Documents/Diss` would be split into the tokens `cmertes`, `Saffron`, `~`, `Documents` and `Diss`. Depending on an internal switch, these tokens are then put into either one common list or three separate lists, one for each window property. These lists are then pruned by sorting the tokens by the frequency with which they occur and discarding the least frequent tokens. The percentage of discarded tokens can be freely chosen; we ended up discarding the bottom 90 % of tokens because of the relatively large number of rarely used tokens. Each token then becomes a binary feature where we store whether the token was present or not for every data point.

Unlike the other features mentioned, these string features are generated offline during postprocessing, immediately before the generation of the model. So the sensor log files only contain the raw values of the window properties.

An obvious drawback of this approach is that the frequency of a token does not say anything about its predictive power. While very rare tokens are generally not useful to train a model, we might still discard tokens that are important. A more advanced approach would be to generate all possible token features and then use some sort of feature selection algorithm [424]. This would significantly increase the time needed to generate the subsymbolic model though which might not be acceptable after all.

7.5.2.2 Plumber – Lumber Jabber Interface

Thanks to the XMPP remote control extension [538] we can integrate with any client supporting this extension without having to change any client code itself or writing client-specific plugins. More specifically, it is possible to remotely change the status of a client. Only commands originating from the same account are accepted and thus the same server-side authentication is used as for the account login in general. Reading out the status is part of the normal functioning of Jabber and is possible without special privileges (other than being mutual contacts). Jabber allows the specification of priorities for multiple resources using the same account so Plumber takes the role of a normal client, just with a low priority (Figure 7.12 on the following page). Unfortunately, Jabber does not have an *invisible* status so it is still possible to send messages to Plumber instead of the actual user if someone orders their client explicitly to do so, if a user goes offline without shutting down Plumber as well, or if a client does not handle priorities correctly. Plumber reacts to this with an automated reply, explaining what it is and that the message did not reach its destination.

User Interface The main user interface is the user's usual Jabber client itself. A minimal control GUI is needed nonetheless but the user should not have to bother with too many options. After all, the goal is to increase calmness. So initially, there were only four buttons and one option as can be seen in Figure 7.13 on the next page: 1) start Plumber (at minimum this means recording sensor data along with Jabber status; the *auto-set status* option controls whether or not the Jabber status is

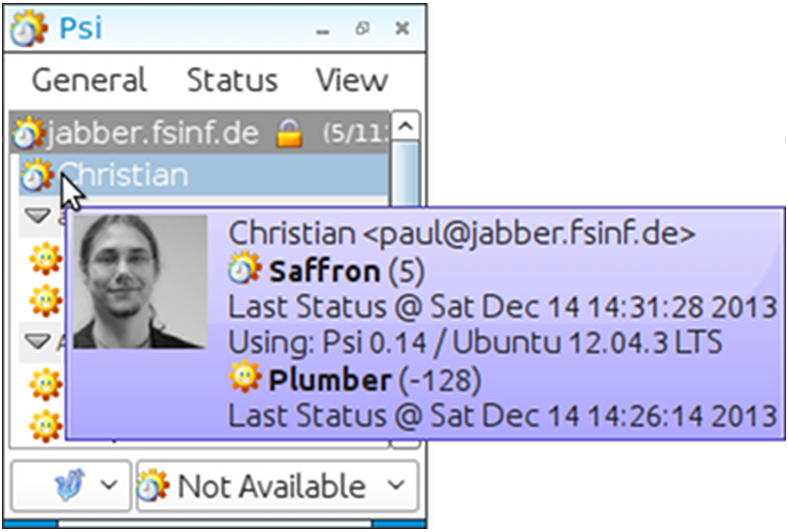

Figure 7.12. Screenshot of how Plumber reveals itself in a Jabber client (Psi in this case). The tooltip shows the two resources for the same account with their respective name, status and priority (here 5 for the regular user and −128 for Plumber).

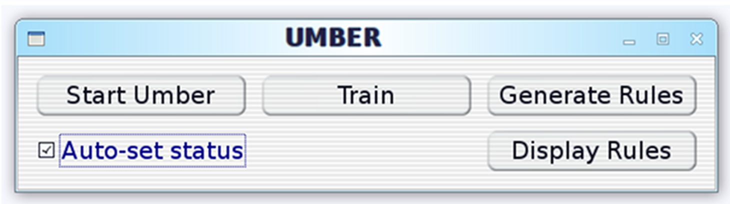

Figure 7.13. The main window of the Lumber GUI is deliberately kept simple. This is an earlier prototype which allows the user to start the recording, select whether or not to automatically set the Jabber status from the trained model and to generate and inspect rule sets.

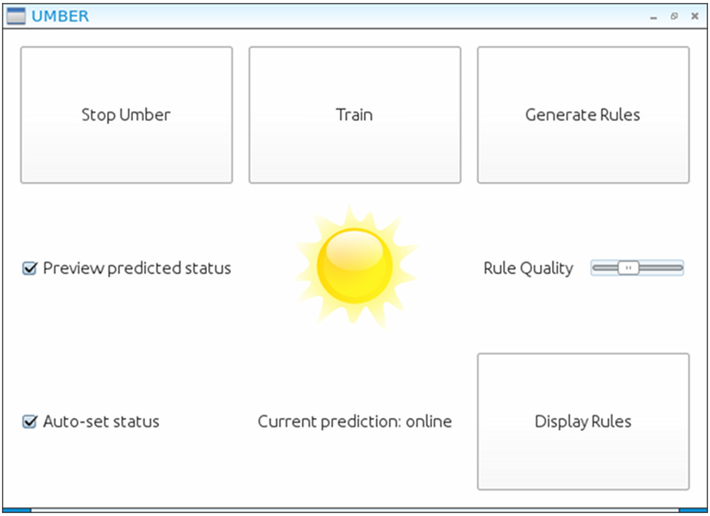

Figure 7.14. This later version of the Lumber control window adds an optional display of the current status prediction as a symbol (cf. Figure 7.16 on page 243) and as text and also adds a slider to control the trade-off of rule quality versus generation speed.

also changed according to the predictions),[12] 2) train the subsymbolic classifier from the recorded data, 3) generate rules from the subsymbolic classifier, and 4) display and optionally change the generated rules. Then this was slightly extended to the picture shown in Figure 7.14 which optionally shows the current prediction so that the users can decide whether or not the predictions are already good enough for them to activate the automatic status change or whether more data has to be collected for the classifier. Also, a slider has been added to control the trade-off between rule set quality and generation speed because

[12] The button that calls Plumber is actually labelled "Start Umber". This is because the subsystem names are not transparent to the user who does not need to know they even exist, let alone their names.

generating good rule sets with G-REX can be quite a time-consuming task and improving rule sets even a bit means investing much more time for additional runs of G-REX. Therefore users might want to trade some quality for large speedups. We chose not to expose the individual parameters of G-REX though because this would have introduced a huge complexity. Instead, three parameters are all changed at the same time by this slider. Specifically, the slider linearly varies the population size from 200 to 1000, the length factor between 0.01 and 0.005, and the batch size from 1 to 5. These specific values are not the result of a user study but stem from personal experience and recommendations Rikard König, the creator of G-REX, gave in private communication.

Because the Jabber client is the only way to influence the class labels of the training data, it is important to inform the user about automatic status changes immediately. As mentioned earlier, desktop notifications are used for this, as shown in Figure 7.15 on the next page. Additionally, most Jabber clients offer to continuously display the current status in the notification area.

7.5.3 Lumber Performance

To evaluate the basic generalisation capability of the classifier using the features we selected, we tested it using a preliminary single-user data set with a small sample size of 1531. The data set was recorded during several shorter sessions in different situations (such as in the office, at home and during the weekend) over the course of a week using the standard 5 s polling interval. We left out the time features since they cannot contain useful information with so few and short sessions and might lead to an undeservedly good result by overfitting the few long stretches of identical status which are covered by both training and test data. For classification we used a support vector machine with an RBF kernel from LibSVM [74], wrapped by the Weka toolkit [205]. We did a random ⅔-split into a training and a test data set and performed a hyperparameter grid optimisation on the training

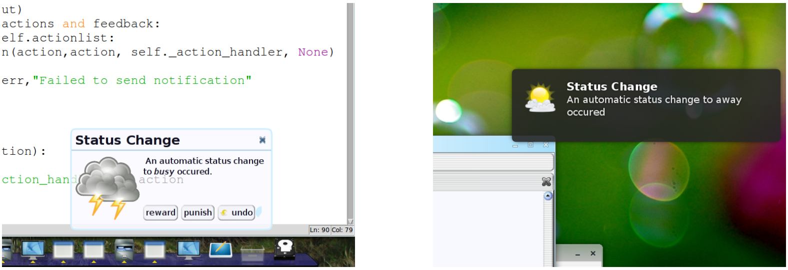

Figure 7.15. Desktop notifications can take different forms but are generally a good way to unobtrusively inform the user about an occurring status change. The icon symbolises the new status. The screenshot on the left shows an earlier concept that still includes the ability for the user to influence the classifier more directly than by changing the status. This idea was abandoned early in the implementation process because it was deemed unnecessary and overly complicated.

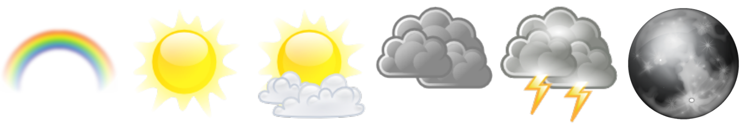

Figure 7.16. Weather-themed icons used for the different statuses. From left to right: free, online, away, not available, busy, offline. The moon and rainbow icons are used under the demo licence from icons-land.com. All other icons are distributed under the GNU Lesser General Public License as part of the Oxygen icon set .

data with a 5-fold cross validation using LibSVM's grid tool [245].[13] On the test set, the generalisation accuracy was 98.2 % with parameters $C = 512$ and $\gamma = \frac{1}{8192}$.

While this seems like a good result, it is not clear what the threshold for a system would be to be perceived as calm by most users. On a more technical note it has also to be said that the generalisation ability for completely new situations is drastically lower and we do not know at this point how much training data would be necessary to better generalise to new situations and how this would vary between users (cf. Section 7.6.3 for variability between users in the case of Cucumber).

7.6 Cucumber – Mobile Prototype

Because of the restricted context domains of the desktop and technical difficulties in rolling out Lumber on a larger scale to collect data from more users, we also implemented a mobile application to collect context data from standard smartphone hardware. The field of context recognition using mobile phone sensors has mostly focused on the implementation of recognition using a single sensor rather than on the integration of all available sensors although there are notable exceptions [117, 159, 238, 287, 302, 317, 448]. To our knowledge there is no architecture though that combines both machine learning and on-device context recognition. We therefore implemented the Cucumber system that does online, on-device subsymbolic context recognition.

Cucumber lets the users define their own contexts and initially polls them for their current context every 15 minutes but suppresses polling when the classification probability is high enough. Users define their context in the form of an arbitrary number of self-defined activities.

7.6.1 Implementation

Cucumber is a self-contained mobile phone application running on Android [185]. The implementation was done by Tobias Rodehutskors,

[13] https://github.com/cjlin1/libsvm/blob/master/tools/grid.py

Box 6: On-Device Processing
In Cucumber we made an effort to keep data on the mobile device. This has usability implications because users do not need to switch devices but more importantly, it has privacy implications because the data that never leaves a device cannot be misused. The obvious alternative in a productive system is to send all data to a central server to allow users access to their data in the cloud from any machine. We think, however, that this is not desirable given the sensitive nature of these data.

co-supervised by my colleague Sebastian Hammerl and myself.

The extracted features are stored in an SQLite database,[14] alongside the raw data (except for audio data for legal and privacy reasons), classifier predictions and manual labels from the users.

7.6.1.1 Features Extracted

Accelerometer Features Activity recognition using accelerometer data has in the past mostly been done by using multiple accelerometers worn at different points on the user's body [23, 146, 308, 346, 363, 418, 526] or with dedicated sensor devices instead of the consumer-grade hardware that is built into mobile phones [78, 79, 320, 322, 327, 342, 436, 522].

There were only a few attempts at activity recognition using off-the-shelf smartphone accelerometers. Miluzzo et al. [379] used the accelerometer in conjunction with other sensors for their CenceMe application but only extracted mean, standard deviation and number of peaks. Yang [583] used a selection of spectral and time-series features to recognise activities using only the acceleration data from a mobile phone. Brezmes et al. [59] put great emphasis on an arbitrary position and rotation of the mobile phone. They tried time- and frequency-domain features but dismissed both approaches, ultimately choosing to use a kNN on unprocessed time-series records. Kwapisz

[14] http://sqlite.org/

et al. [313] relied solely on a number of time-domain features while Wang et al. [560] use a wavelet transform to generate features from the accelerometer signals.

We chose to use the features proposed by Kwapisz et al. [313] because they worked quite well in a previous system we built [222]. These features are computed over a 10-second window and consist of the mean and standard deviation, the mean difference from the sample mean, a time between peaks algorithm and a 10-bin histogram – all computed per axis – and the average magnitude of the 3D acceleration vector.

Magnetometer Features Mean and standard deviation over the sample window are computed per axis for the magnetometer readings. Additionally, an absolute rotation of the device in space is computed, using the accelerometer to determine the device's rotation relative to earth's centre of mass. This effectively yields a 2D compass value and the mean azimuth.

Audio Features Instead of the rectangular windows of equal size in frequency space used in Section 7.5.2.1, for Cucumber we used the well established Mel-frequency cepstral coefficients (MFCCs) [60, 374, 558] that are closer to human perception. They are overlapping triangular windows on the cepstrum [47]. Additionally, the root mean square (RMS) and zero crossings were used as features in the time domain.

Time Features We used the same circular time features as were described in Section 7.5.2.1 (day, week and year) but we modelled the weekend–working week dichotomy explicitly to make learning it unnecessary.

Location From the GPS data, mean and variance of longitude, latitude and altitude is computed. Speed and direction in the x/y plane are also determined, as well as overall mean speed and true bearing, i. e. the movement's orientation towards the geographic north pole. The

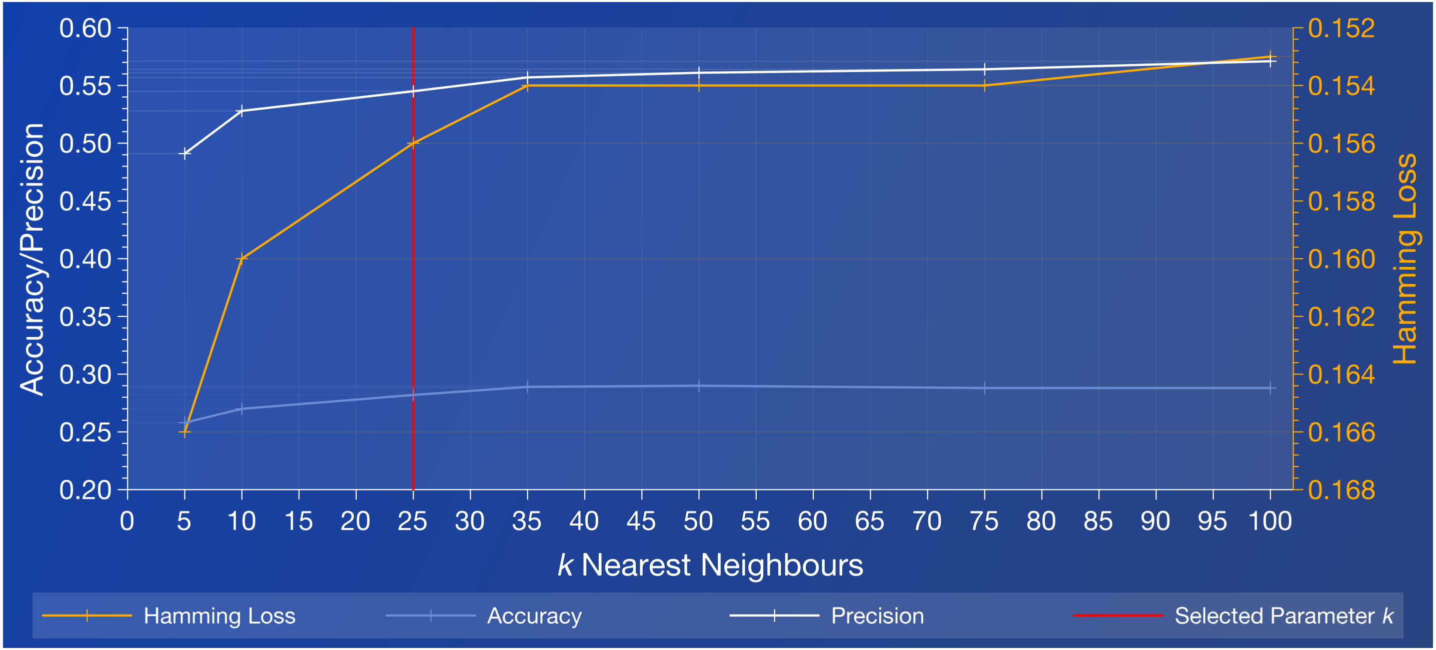

Figure 7.17. Parameter selection for the kNN classifier after Hammerl [207].

GPS accuracy is stored but not further processed. The same features are again computed for the Wifi-supported position.

Bluetooth and Wifi Additional to the location information, the Wifi scan raw data is also stored (i. e. SSIDs, BSSIDs, signal strengths and frequencies). Bluetooth scan data was stored for later inspection but not used in the classification because most Bluetooth devices these days are switched off by default to conserve energy and therefore the utility of these scans is very low.

Other Features Proximity and light sensor produce a scalar value[15] for which mean and standard deviation over each window are computed.

Additionally, the charging state of the battery is queried as a binary value (charging or not charging).

7.6.1.2 Machine Learning

Since users can choose more than one class for each data point, a multi-label learning algorithm is needed. For this, a variant of the Weka toolkit called Mulan was used [542]. In a previous study, Hammerl [207] had tried ML-kNN [588], RAkEL [541] and BPMLL [587] as alternative machine learning techniques for this type of data. He found that the classifier that worked best was ML-kNN, a multi-label version of the k nearest neighbours classifier [588]. Hammerl also tried different values for parameter k (cf. Figure 7.17 on page 247) and found that more neighbours meant a better classification but also found a saturation at about $k = 35$ neighbours. For Cucumber, $k = 25$ was chosen because it delivers already 97 % of the maximum measured accuracy and 95 % of the precision (cf. Equations 7.8 and 7.9 on page 253 for a definition of these multi-label measures) while needing much less samples and therefore less time that a user has to train a certain class. Using less

[15] Although on the platformed used, the proximity sensor only supported taking one of two values. Cf. http://developer.android.com/reference/android/hardware/SensorEvent.html#values

neighbours also cuts down on the computing the mobile device has to accomplish.

7.6.2 Usage

As said earlier, Cucumber initially prompts users every 15 minutes for their current activities (Figure 7.18(b) on the following page) unless set to incognito or silent mode. It is possible to select more than one activity. Users can define their own activities and thus do not have to select from a fixed set.

When prompting the user, a class c is preselected if the probability $P(c)$ that it represents a current activity as determined by the classifier is larger than a certain threshold θ. As the software learns to predict the user's activity, it will cease to prompt the user if the probabilities of the classifier for all the classes are either above or below a certain threshold (i. e. only show the prompt if there is at least one class for which the propability $0.2 < P(c) < 0.8$). That a probability value below a threshold can silence the prompt might seem counterintuitive but it is based on the rationale that if the probability that the current activity is correctly classified by a class c is $P(c)$ then $1 - P(c) = P(\neg c)$. In other words, if the probability for a class is very small, then this corresponds to a high probability that the current activity does not belong to that class.

Data recording is triggered every 3 minutes for 10 seconds unless in incognito mode.

7.6.3 User Study

In a study designed and conducted by my colleague Sebastian Hammerl, optimal values for the probability thresholds could be determined empirically. The study comprised $n = 8$ participants (6 male, 2 female; mean age 27 years) who used Cucumber for 11 days (although a few participants continued shortly beyond this time frame).

During the study, the level we call the *preselection confidence threshold* that determines at which classification probability to preselect a class was set to $\theta = 0.5$. This value was chosen after some loose

(a) Main Screen

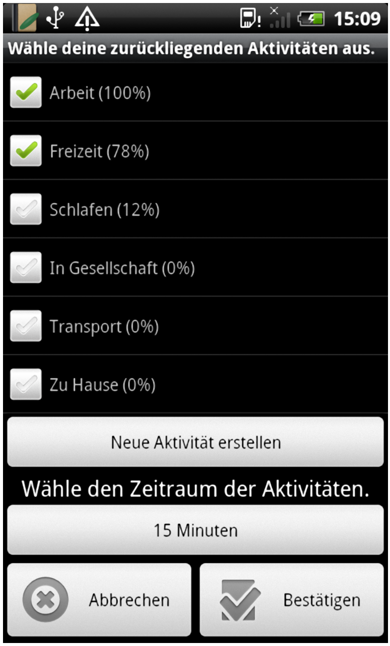

(b) Activity Selection

Figure 7.18. Two screenshots of the Cucumber application. The first is the main screen, the second an exemplary activity selection with two activities preselected due to their high classification probability. The activity selection is either triggered by a timer or explicitly by the user.

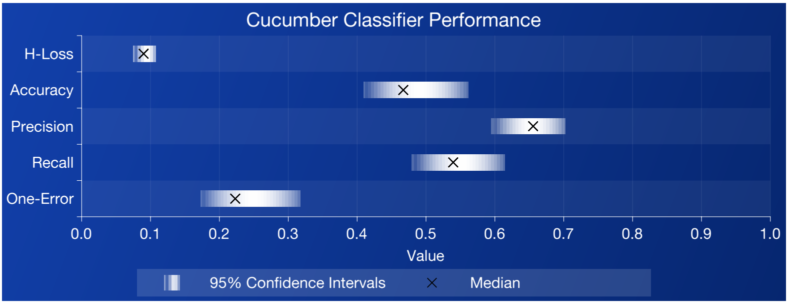

Figure 7.19. Classifier performance measures for the Cucumber study.

testing with a single data set but was later found to be quite close to the optimal value,[16] as determined by the following method.

Figure 7.20 on the next page shows the relative amount of necessary changes to the classes suggested by the classifier, i. e. the ratio

$$\begin{aligned} \Gamma &= \frac{|C_{\text{user}}| - |C_{\text{classifier}} \triangle C_{\text{user}}|}{C_{\text{user}}} \\ &= \frac{|C_{\text{classifier}} \cap C_{\text{user}}| - |C_{\text{classifier}} \setminus C_{\text{user}}|}{|C_{\text{user}}|} \end{aligned} \tag{7.7}$$

of necessary changes[17] to the suggested set of classes over the number of entries that would be necessary without any automatic preselection by the classifier. This ratio is plotted for varying preselection thresholds θ, showing an optimum very close to the value of $\theta = 0.5$ chosen for the study. The rationale behind this measure is that it corresponds to the number of "taps" or other necessary inputs to correctly set the state of the device and thus the classifier. If the ratio is negative, the user

[16] within about 98 % to be exact

[17] $\triangle$ denotes the symmetric difference $(A \setminus B) \cup (B \setminus A)$ of two sets A and B. C_{user} is the set of classes that the user actually selected and $C_{\text{classifier}}$ the set of classes that the classifier preselected.

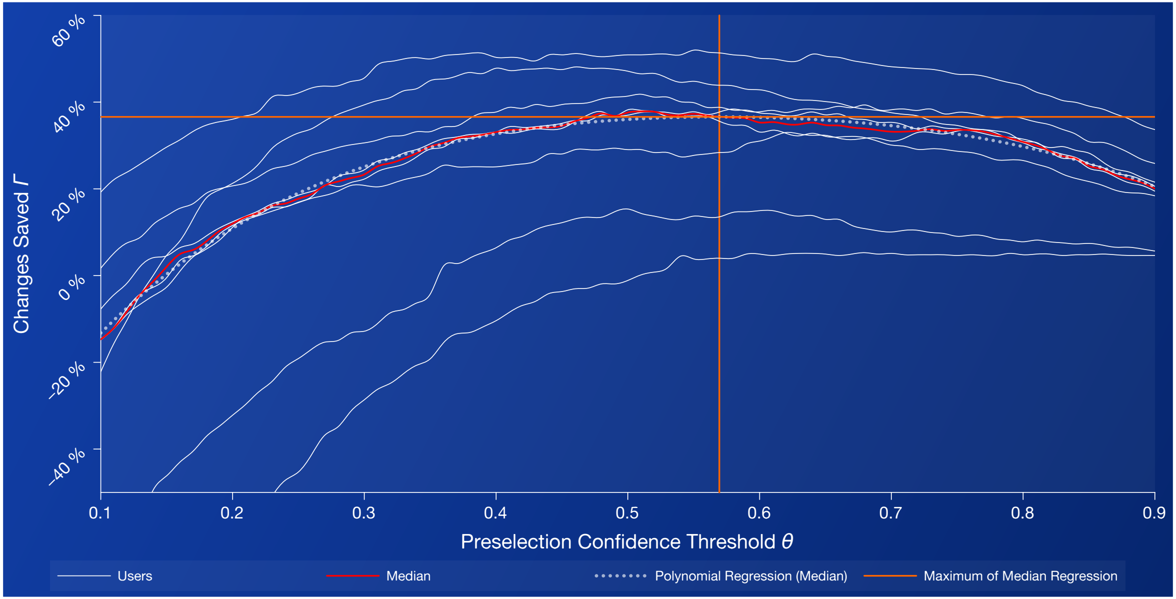

Figure 7.20. This figure shows the relative amount of necessary changes to the suggested set of current classes Γ, i. e. the edit distance[a] between the sorted lists of classes with each class counting as one symbol, divided by the edit distance between the empty string (i. e. list of classes) and the desired one. This ratio is shown for varying probability value thresholds θ for the preselection. The optimal preselection confidence threshold according to the median of all users' data is $\theta \approx 0.51$ (reaching $\Gamma \approx 37.9\,\%$). The polynomial regression of the median data[b] is of order 4. Its maximum $\Gamma \approx 36.6$ is at $\theta \approx 0.57$ (orange drop lines).
Note that the plot only shows the interesting part around the maxima of the graphs, before edge effects set in. All graphs will converge against 0 % for $\theta \to 1$ since fewer and fewer preselections will happen, mimicking the "no classifier" state in its extreme position. This means that the regression only works for values close to the maximum.

[a] modified Levenshtein distance [328] where substitutions count double [268]

[b] being $\Gamma = -400.667 * \theta^4 + 955.677 * \theta^3 - 970.742 * \theta^2 + 471.82 * \theta - 51.6359$

would have been better off without any preselection by the classifier while any value larger than 0 corresponds to saved time.

The plot shows that there is quite a broad plateau, making the choice of θ very robust. For the regression line, Γ stays within 10 % of its maximum for the whole interval $\theta \in [0.41, 0.74]$. What can also be seen, however, is that there is quite a bit of variability between users. The one user that could not profit from the system at all reported to have had a particularly irregular schedule at the time of the study. Others could slightly increase their saved changed by varying their personal value of θ.

Being faced with a multi-label problem makes the classifier evaluation more complicated [588]. A number of metrics have been proposed to deal with this [540], from which we use a selection of the more popular ones [175, 471, 535, 592].

There are more or less direct multi-label generalisations of the usual measures accuracy, precision and recall. They are defined as follows, with X_i being the set of correct labels for data point x_i and Y_i being the set of labels predicted by the classifier,[18] where $i \in \{1, \ldots, N\}$.

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^{N} \frac{|X_i \cap Y_i|}{|X_i \cup Y_i|} \tag{7.8}$$

$$\text{Precision} = \frac{1}{N} \sum_{i=1}^{N} \frac{|X_i \cap Y_i|}{|Y_i|} \tag{7.9}$$

$$\text{Recall} = \frac{1}{N} \sum_{i=1}^{N} \frac{|X_i \cap Y_i|}{|X_i|} \tag{7.10}$$

Two more measures we used were the Hamming loss (H-Loss) and the one error. The Hamming loss counts the number of misclassified labels which means that it is better, the smaller it is. Let L be the set

[18] Be careful not to confuse these with the y_i from Equations 7.13 and 7.14 on page 271.

of all labels, then

$$\text{Hamming Loss} = \frac{1}{N}\sum_{i=1}^{N}\frac{|X_i \triangle Y_i|}{|L|} \tag{7.11}$$

The one error is a kind of metric that in a way tries to evaluate what the classifier's classification error would have been if it was a single-label problem. It only looks at the top-ranked label and checks if it is in Y_i. Note that this is of course an easier problem than actual single-label classification. Let $r_1(x_i)$ be the function that returns the top-ranked label.

$$\text{One Error} = \frac{1}{N}\left|\{x_i \mid r_1(x_i) \notin Y_i\}\right| \tag{7.12}$$

The results for the study data can be seen in Figure 7.19 on page 251. It is also interesting to look at the temporal dynamic of these measures. In Figure 7.21 on the facing page it can be seen that the classifier reaches near-peak performance very quickly but then drops drastically when daily routines change during the weekend.

On average, users ended up with 13.9 classes ($SD = 3.72$), with the minimum number of classes being 10 and the maximum being 20. The average number of classes for each selection was 1.7.

7.6.3.1 Discussion

Overall, the data shows that it is possible to predict user status on mobile devices with sensors available on the device itself accurately enough to support applications where the users are interested in their status anyways (i. e. lifelogging). This is even possible when users choose the classes and number of classes they want and that are important for their lives, including multiple classes at the same time, instead of relying on a fixed set of classes that are easy to detect.

The study also shows that people did indeed use the multi-label feature and being in more than one relevant state at the same time is not the exception but the rule.

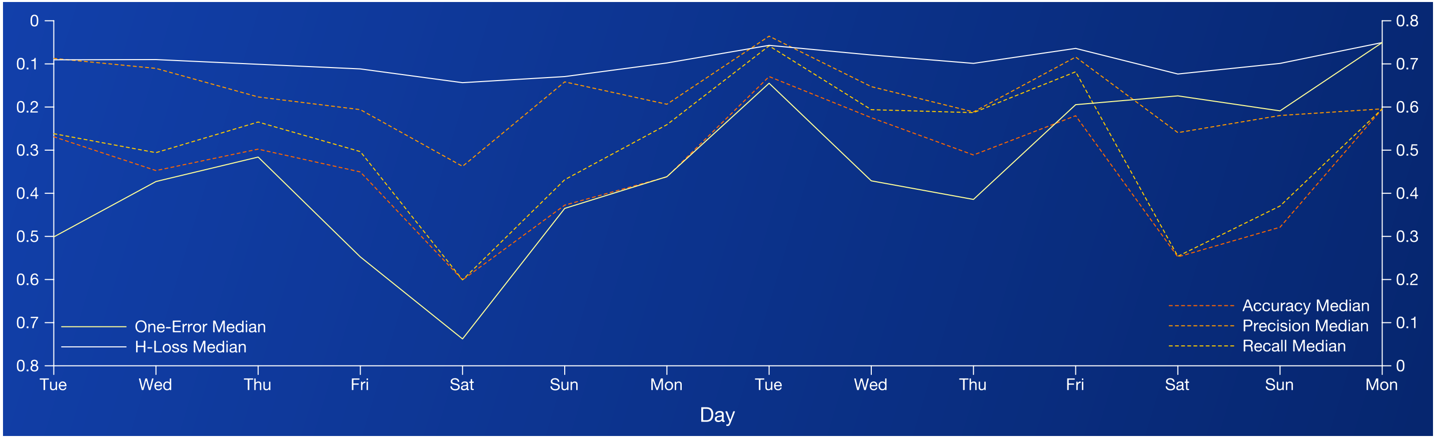

Figure 7.21. Change of classifier performance over time. The inverted left ordinate is used by Hamming loss and one error and the right by accuracy, precision and recall. All time series use the median of all the participants' results. The second weekend and the following Monday are based on significantly less data and should be interpreted with caution. It can be well seen that performance reaches very good values almost immediately but drops on the weekend, as would be expected with a change of daily routines.

Unsurprisingly, manual corrections were still necessary and what such a study therefore cannot show, is the willingness of users to accept the burden of such manual corrections in any given application that uses user status or in turn their willingness to accept the error rate, in this case about 10 % misclassified labels. We expect this to depend largely on the importance of the application to the users themselves and on the effectiveness for a particular user. That is, for a minor gain in comfort that depends on a very accurate prediction, this kind of system will not be accepted and neither will this be acceptable for people with very irregular schedules. On the other hand, a large personal improvement for a user with a fairly predictable schedule and an application that is robust against occasional misclassifications will very likely be accepted. The large middle ground between these two extremes is an area for future research.

The balance between automated classification and manual correction is also very much relevant to the area of hybrid transformation architectures that we are going to explore in the following.

7.7 Hybrid Transformation Architectures

A hybrid transformation architecture (HTA) is a system that translates between the world of the subsymbolic in which most commonly used machine learning techniques live and the world of symbolic rules that have largely fallen out of favour for intelligent systems[19] but that are arguably much closer to human language than subsymbolic systems are.

Wermter and Sun [569] coined the term *hybrid transformation architectures* in the context of neural networks but the principle has been applied to many different machine learning techniques [249]. Wermter and Sun argue that such hybrid systems are able to combine the advantages of symbolic and subsymbolic systems. Among the advantages the symbolic side brings to the table, they mention the ease of interpretation, the explicit control and fast initial coding. For neural

[19] as for example expressed by the derogatory nickname "good old-fashioned AI" given to the symbolic approach to intelligent systems [178]

Box 7: Wermter and Sun Taxonomy
Wermter and Sun [569] used their taxonomy to describe systems that combine neural networks and symbolic representations but their classification scheme is oblivious to the kind of subsymbolic system used. The names of the categories, however, are not; beginning with their term for the whole class of systems: **hybrid neural architectures**.
They differentiate these in three subclasses: **unified neural architectures** that are purely subsymbolic systems but allow the symbolic interpretation of its components (the archetypal example is the grandmother cell [195]), **hybrid transformation architectures** that translate between symbolic and subsymbolic representations (see main text), and **hybrid modular architectures** that combine symbolic and subsymbolic architectures in various ways to solve a specific task.
The work by Wermter and Sun [569] was not the first to provide a taxonomy of hybrid symbolic–subsymbolic systems (cf. Wermter and Sun [569] itself) but they integrated earlier work, often by the authors themselves. The alternatives names used to describe this class of systems are similarly variations on the same theme, e. g. hybrid connectionist-symbolic models [523, 524], hybrid intelligent systems [373], or neurosymbolic integration [234].

systems, they mention better learning, robust fault-tolerant processing, generalisation and what they call "gradual analog plausibility", an at that point unexplained term that we take as alluding to the inherent fuzziness of subsymbolic systems.

Wermter and Sun [569] presented a taxonomy of what they called *hybrid neural architectures*, combining and simplifying earlier classification schemes (cf. Box 7 above). Hybrid transformation architectures are one of the three basic classes of hybrid neural architectures in this taxonomy. However, they used the term for any system that performs *either* rule extraction (i. e. the process of extracting symbolic information from subsymbolic systems) *or* knowledge insertion (i. e. the opposite process of modifying a subsymbolic system according to symbolic data). In an important deviation from this taxonomy, we use the term hybrid transformation architecture only for systems that combine *both* these properties (cf. Figure 7.22 on the facing page) since rule extraction and knowledge insertion are well-established terms already and are sufficient to describe any system that combines symbolic and subsymbolic data in these ways while for the combination of the two, no term yet exists so we used the one by Wermter and Sun to fill that gap but excluded the already labelled subsets from the definition. The term *knowledge refinement* has been used to describe this combination [83] but the term implies the exclusion of using knowledge insertion to initialise a subsymbolic model.

7.7.1 Possible Uses of Hybrid Transformation Architectures in Human-Computer Interaction

7.7.1.1 White-Boxing

Practically all of the most effective machine learning algorithms to date produce a model that is essentially a black box [249], such as a collection of neuronal weights or support vectors. The process of turning an incomprehensible black box into its comprehensible counterpart[20] is sometimes called *white-boxing* [247]. Huysmans et al. [249] argue that such a term (they themselves use "white box extraction") is generally more appropriate in this context than rule extraction, since the extracted symbolic model does not necessarily consist of rules (cf. Section 7.7.2). Either way, as we established the importance of comprehensible systems as a design goal in Section 7.3, white-boxing is an important motivation behind Umber.

Many traditional reasons to use rule extraction techniques also fall in this category, especially in the areas of medicine and credit scoring which are fields that have been large drivers behind the development of rule extraction techniques [18, 249, 357, 359, 425]. It has been argued

[20] The opposite of a black box is mostly called a white box even though the more rarely used "glass box" stays arguably truer to the original metaphor.

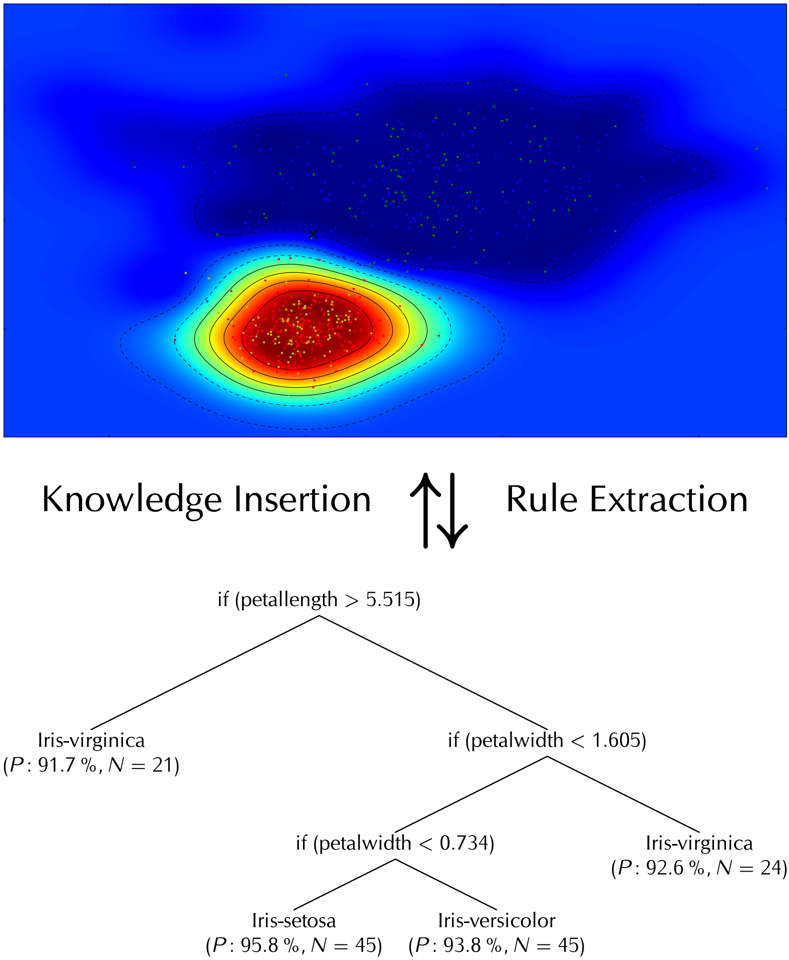

Figure 7.22. Hybrid transformation architectures translate from subsymbolic to symbolic systems and back. The decision tree was generated from the Iris data set [144] using G-REX; the SVM visualisation uses simple artificial two-Gaussian data.

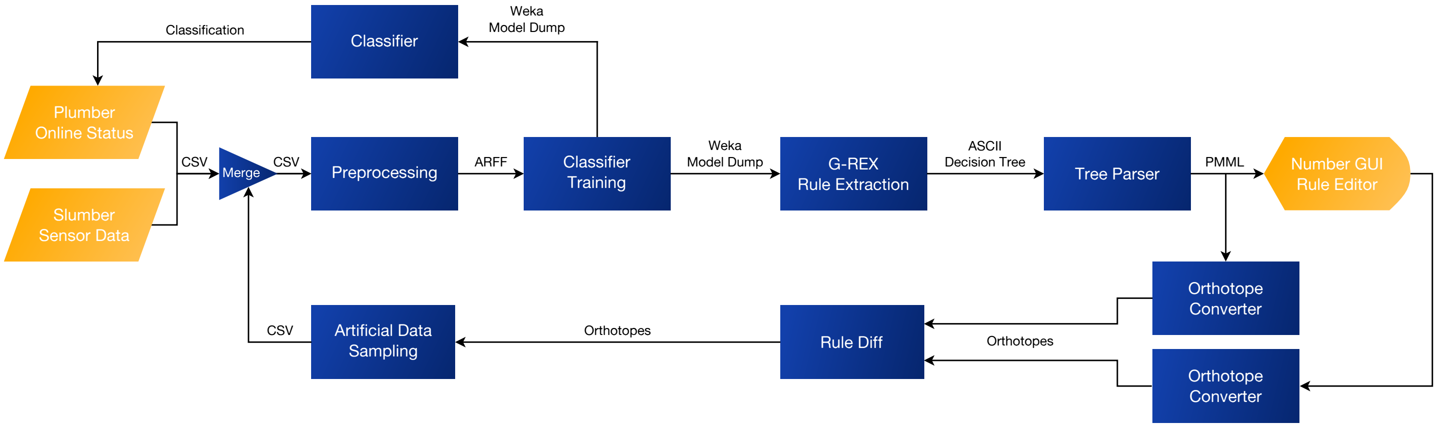

Figure 7.23. Lumber information flow pipeline with box shape semantics loosely based on flowchart semantics. Yellow boxes represent points at which information between user and software is exchanged while blue boxes are largely automatic, even though the user might initiate and parametrise them.
The box labels denote functions or module names and the arrow captions indicate the format of the data transfer. Boxes can be separate software components, scripts or just functionally distinct submodules within a larger software component.
The real-time flow from Slumber to the classifier was omitted from this diagram for better readability.

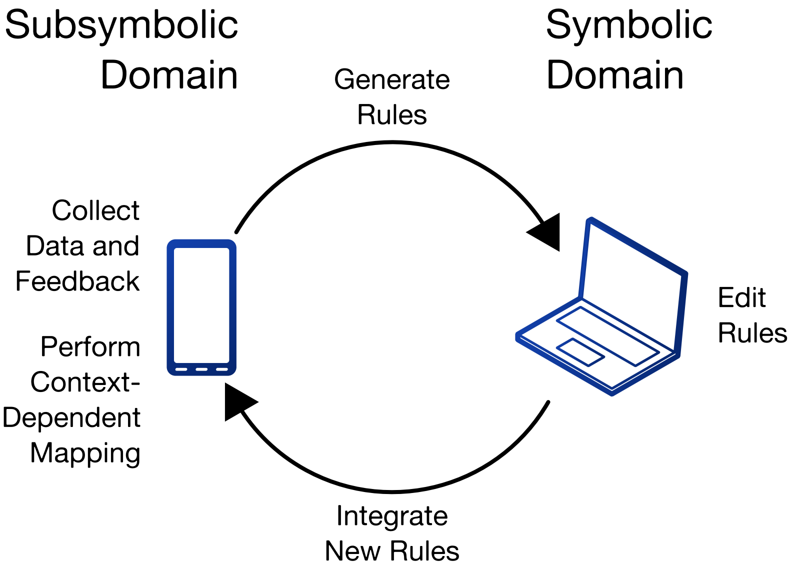

Figure 7.24. Illustration of a symbolic–subsymbolic workflow for mobile phones (Cucumber).

that many medical practitioners do not want to trust the diagnosis of a patient to a system that would not explain how it came to its conclusion [212].

In credit scoring, the Equal Credit Opportunity Act in the US [547] prohibits to “discriminate against any applicant [. . .] on the basis of race, color, religion, national origin, sex or marital status, or age”. Therefore creditors are obliged to provide the “specific reasons” of any denial of credit. This is where rule extraction techniques come in handy because they can do this without having to train a rule-based credit scoring model to begin with but rather generate one later from a subsymbolic model.

7.7.1.2 Explainability and the Potential of Verbalising Probabilities

The above applications of rule extraction in medicine and credit scoring show a tight link between comprehensibility and the ability of a system to explain and motivate its actions. The latter can be expected to be especially important in robots and other agents that interact with humans using natural language. Neither a string of numbers, nor a generic answer along the lines of "Just trust me on this." would likely satisfy most human users who ask for the reasons for a recommendation the machine makes [336, 416].

Rule-based systems are not as easily understood as natural language is but they are probably largely superior to probability-based representations since humans are known to be notoriously bad at intuitively dealing with probabilities [229, 324]. HTAs might form a bridge between machine learning and human understanding by enabling input and output of rules to subsymbolic systems. They could also directly feed into agents that support speech output of symbolic systems however. These systems would then have the advantage of being able to respond to enquiries such as "Why did you just do this?", an ability called *explainability* [91].

7.7.1.3 User Control

The previous applications of HTAs emphasise the uses on the output side, i. e. the rule extraction part of these systems. However, HTAs are not only supposed to increase user acceptance by making it transparent what the underlying system does but they also allow the user to influence what the system does. This could potentially give the user a greater sense of agency compared to a black-box system that reacts to user behaviour through learning but otherwise does not provide users with any means to directly influence the system's behaviour. Of course the goal would not only be to make users *feel* better about the system but to effectively increase its performance. This ties in with the following aspect of using domain knowledge.

7.7.1.4 Using Domain Knowledge

HTAs also make it possible to directly tap into the users' domain knowledge. This especially applies for ubiquitous computing systems where that *domain knowledge* is the users' knowledge about their daily lives. To put this less stiltedly, it can be expected nobody knows better what the users' priorities and preferences are than the users themselves. Therefore providing users with a way of formalising these preferences directly might prove much more efficient and much less obtrusive and time-consuming than learning them by example. After all, generating the training data required might be desirable or even not be practical at all.

Imagine a context-aware phone that receives a call from work-related number at 9 PM. Some users might not want business calls to be put through during their leisure time, others might tolerate them during the late afternoon but not after say 8 PM whilst others still might want to receive business calls regardless of the time. Having to put through a great number of unsolicited calls that the user then marks as *incorrectly put through* in order to generate enough data to train the subsymbolic model would be inconvenient to say the least. On the other hand, one call might be enough to remind the user to define a rule that forbids calls after a certain time.

7.7.1.5 Handling Rare Classes

Significant imbalances in the number of training examples for different classes are often in itself a problem for classifiers [364]. The anticipation of such problems could even lead users to not provide training data in the first place, especially if the class distribution is not only skewed but there is no training data for one of the classes in the first place. Take for example a user training their e-mail filter. This user gets frequent phishing mails that pretend to be from Paypal while this company has not send the user any legitimate mail yet. The user does have an account with Paypal, however, so when marking the phishing mails as spam, the user becomes (rightfully) concerned that the classifier might latch onto the company name as a keyword to classify spam, causing any legitimate mail that might one day be sent from Paypal to be classified as spam as well. This might cause the user to delete the phishing mails but not mark them as spam, depriving the classifier of potentially valuable training data.

Being able to inspect and change learnt models can consequently encourage users to provide training data that they might otherwise have discarded because they were afraid of "spoiling" the classifier. In the example provided, the user can look out for a rule like "if mail contains 'paypal' then classify as spam" to appear and delete the rule when this happens, forcing the classifier to look for more subtle indicators of spam in these mails.

7.7.2 Rule Extraction

While there have been various definitions for rule extraction, we will define it as follows:

Definition 16: Rule extraction is the generation of an approximated representation of a predictive model that improves either the ability to correctly predict new data over the underlying model, or has a superior comprehensibility, or both.

Huysmans et al. [249] present a good survey on rule extraction algorithms, including a detailed description of the most common taxonomy for rule extraction algorithms, the ADT taxonomy. It is named after Andrews, Diederich and Tickle [7, 531] but has since been extended and used variedly by other authors. We will briefly describe it in one of its extended forms in Section 7.7.2.1. A more recent but much less detailed look at the field as a whole can be found in Diederich et al. [115].

In the field of data mining, rule extraction techniques are often equated with the term "knowledge discovery in data(-bases)" (KDD) [163, 335, 420]. In this field though, the systems are mainly targeted at highly skilled experts who are willing to put a lot of effort into gaining insight into hidden structures in data. This is quite the opposite of the casual non-expert users who mainly want the technology they use not to bother them.

As Huysmans et al. [249] point out, it is actually misleading to call the process of extracting comprehensible information from incomprehensible models "rule extraction" as there are other forms of decision support, i. e. presentations that are more comprehensible than support vectors or neuronal weights. Huysmans et al. [249] name finite state machines, graphical models and decision trees, although we consider the latter to be just a form of presentation of a given rule set, likewise the decision tables that Huysmans et al. [250] found to perform quite well in terms of comprehensibility. Huysmans et al. [249] propose *white box extraction* as an appropriate term (cf. Section 7.7.1.1). *Knowledge extraction* might be another one, mirroring the common name of the contrary process *knowledge insertion* [369]. However, as "rule extraction" is so widespread and we do in fact only extract rules, we will only use this term throughout the thesis.

The obvious alternative to rule extraction is using a machine learning algorithm that directly produces sets of rules such as C4.5 [434] and CART [56]. However, at least propositional if-then rules (see Section 7.7.2.2 for a brief overview of rule types) partition the feature space using hyperrectangles, so they have little chance of reaching the generalisation power of subsymbolic systems (cf. Figure 7.26 on page 268). Moreover, the more fine-grained an area is sampled using rules, the less comprehensible it will be for humans. The same is generally true for a more powerful type of rules such as oblique rules [250]. See Section 7.7.2.2 for details.

7.7.2.1 Taxonomies of Rule-Extraction Algorithms

There have been a number of additions to and modifications of the commonly used ADT taxonomy [93, 219, 249, 309]. One generalisation of the ADT taxonomy is presented by Huysmans et al. [249]. In it, there are three main dimensions to place any given rule extraction algorithm. These are (although differently named in the original report) 1) predictive model type, 2) translucency regarding the underlying model and 3) rule type and quality.[21]

The type of the predictive model is either regression or classification. The second dimension is concerned with the dependence of the rule extraction algorithm on exactly what model is used underneath. Some algorithms work on internals like support vectors or individual neurons while others are completely oblivious to how the black box works. This aspect is a bit more confusing when it comes to the exact definition in the literature. While Huysmans et al. [249] call this dimension "dependent versus independent", the widespread ADT taxonomy called this distinction "pedagogical versus decompositional" (the latter term coined by Craven and Shavlik [94]) and also introduced a hybrid form that was called "eclectic" [7].[22]

Decompositional algorithms look at individual components of the black box model whereas pedagogical algorithms directly map the input of the model to its output. While this seems similar or even the same as saying that the algorithms are dependent on versus independent from the type of the underlying black box model, according to Huysmans et al. [249] the definition of pedagogical systems does not forbid a look inside the black box. This lead them to coin new terms with definitions that they think to be better applicable where not only artificial neural networks are considered as subsymbolic models (as has traditionally been the focus in the field), even though the authors of the ADT taxonomy themselves applied them to SVMs as well [24]. However, these terms and definitions failed to get any widespread use in the field so we will also stick to the original terms and ignore the special cases where their original definition might fail to overlap with the Huysmans et al. taxonomy.

[21] The ADT taxonomy has five dimensions that are called 1) translucency, 2) rule quality, 3) expressive power, 4) portability and 5) algorithmic complexity.

[22] The exact definition of an eclectic algorithm is to "utilise knowledge about the internal architecture and/or weight vectors in the trained artificial network to complement a symbolic learning algorithm".

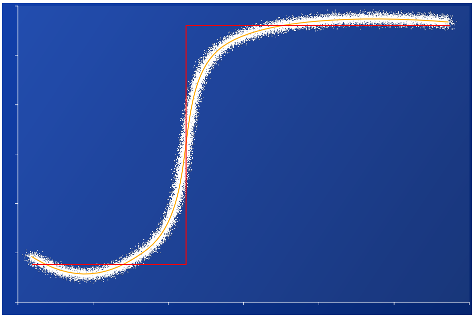

(a) Propositional with Focus on Comprehensibility

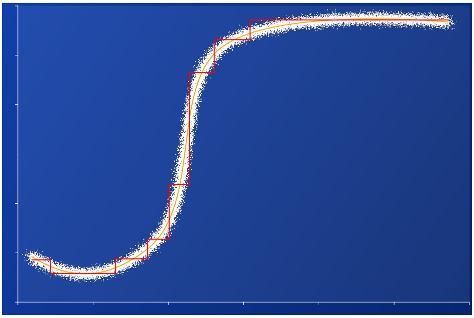

(b) Propositional with Focus on Accuracy

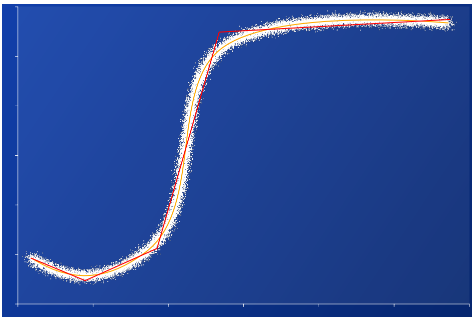

(c) Oblique Rules

Figure 7.25. Illustration of the approximation capabilities of propositional (a) vs. oblique rules (c) and of the trade-off between comprehensibility (a) and accuracy/fidelity (b). The white point clouds represent the data, the orange lines in their centre a perfect fit and the red straight line segments the approximation by the according rule set.

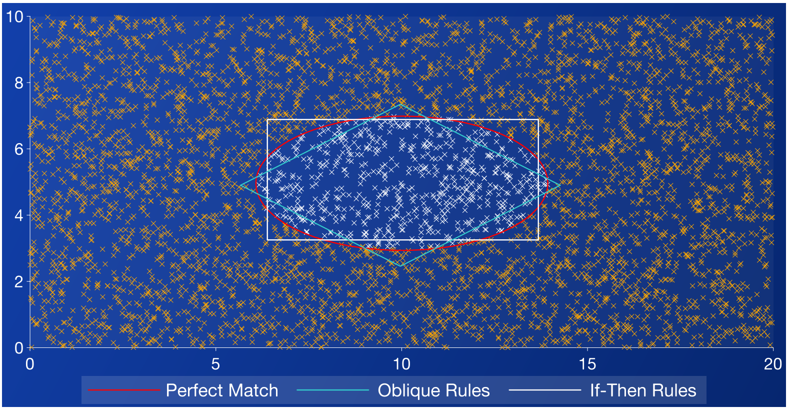

Figure 7.26. Illustration of the approximation capabilities of different kinds of rules in 2D.

Because the requirements vary so greatly for SCS (e. g. multi-label, incremental learning, speed requirements, mobile requirements) and we saw in Section 7.6.1.2 how kNN won over SVMs that are considered by some to be the best off-the-shelf supervised learning algorithm [395], it is very valuable for us to use a pedagogical rule extraction algorithm so we can keep the machine learning part modular and simply change it when the requirements change for a given application or because we learnt something through our prototypes.

7.7.2.2 Rule Types

When talking about rules, mostly simple, so-called *propositional if-then rules* are meant. There are other kinds of rules however, the most important of which we will briefly describe in the following. The different kinds of rules differ not only in their visual representation but also

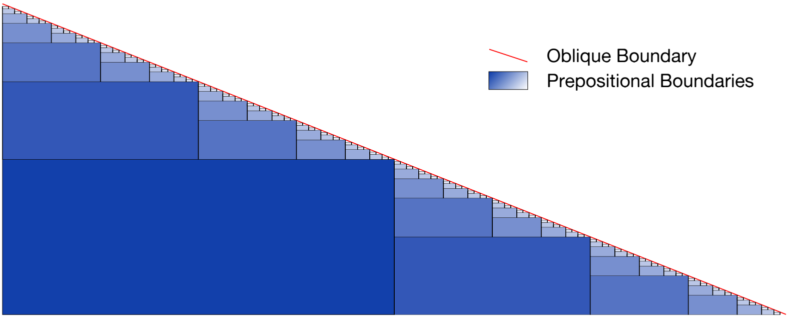

Figure 7.27. Illustration that – in general – oblique rules can only be approximated but never be perfectly represented using prepositional rules. The number of prepositional rules also increases exponentially with the desired accuracy. Adapted from Figure 2 in Fung et al. [162].

in their descriptive power, i. e. in how easily they can approximate arbitrary shapes, and in their comprehensibility.

The most common type, the propositional if-then rules (if (feature < upper boundary) ∧ (feature > lower boundary) then class), describe paraxial hyperrectangles in the feature space, although these hyperrectangles (also called *orthotopes*) do not need to be bounded from all directions, i. e. some edges can be at infinity.

Mathematically, oblique rules represent only a slight adaptation of this basic type of rules. They allow the addition of features and their multiplication with a factor to the effect of "if factor 1 × feature 1 + factor 2 × feature 2 > boundary then class". This means that boundaries in the feature space are no longer required to be parallel to the coordinate axes but can form arbitrary linear functions. On the other hand, this mathematically simple adaptation is thought to significantly reduce the comprehensibility of the rules [249] and such an effect has indeed been shown by Huysmans et al. [250].

TREPAN (cf. Section 7.7.2.5) produces rules called *m-of-n rules*. These were first proposed by Towell and Shavlik [534] and describe a type of rules where *m* conditions from a set of *n* conditions have to be true for the whole expression to be true. Such rules are a compact form of expressing propositional rules in disjunctive normal form [249] and are therefore equivalent to propositional rules with respect to their logical expressiveness. It has yet to be shown, however, if they are as easily understood as propositional rules as well.

An interesting type of rules regarding comprehensibility are *fuzzy if-then rules* [126]. These resemble propositional if-then rules but instead of a comparison statement they take linguistic variables which might result in something like "if feature 1 is large" instead of "if feature 1 $> x$". The problem of course is to derive meaningful linguistic variables where an area within the feature space is described in a way that is meaningful for the symbolic predictive system as well as for the human user and, most importantly, where these meanings match each other. Also, for hybrid transformation architectures, the problem of the creation of linguistic variables for the rule set manipulation and subsequent knowledge insertion process arises, as well as the re-use of created variables during the next iteration of rule extraction. Typically, fuzzy rules do not describe orthotopes in the feature space but – as their name implies – fuzzy areas realised by windowing functions. As will be seen below, we chose not to use fuzzy rules in the system built but we did put emphasis on the condition that the rule extraction algorithm can in principle generate fuzzy rules as well for the sake of future potential improvements.

7.7.2.3 Rule Quality

The four most important [115] criteria for rule quality Andrews et al. [7] describe are accuracy, fidelity, consistency and comprehensibility. While comprehensibility was covered to a larger extent in Section 7.3, we will briefly describe the other three in the following.

Accuracy Accuracy is the ability of the extracted, rule-based model to predict and generalise the original data, i. e. to accurately predict a different but equally structured dataset as the training data.

This is the same as the well-known measure commonly used to estimate the quality of predictive systems. Its formal definition for data points x_i with $i \in \{1, \ldots, N\}$ and the correctness of the corresponding classification $y_i \in \{0, 1\}$ where 1 is a correct classification and 0 an incorrect one (with "sym" denoting the classification of the symbolic and "subs" in Equation (7.14) that of the underlying subsymbolic system) is:

$$Accuracy_{\text{sym}} = \frac{1}{N} \sum_{i=1}^{N} y_{\text{sym}_i} \tag{7.13}$$

Fidelity Fidelity is the ability of the extracted rules to match the predictions of the underlying model. Other than accuracy, this is a measure that only makes sense for two-step processes such as the extraction of rules from trained models. More formally again, with $\oplus$ denoting the *exclusive or* operator:

$$Fidelity_{\text{sym}} = 1 - \frac{1}{N} \sum_{i=1}^{N} y_{\text{subs}_i} \oplus y_{\text{sym}_i} \tag{7.14}$$

Consistency There are slightly varying definitions for consistency. Generally, consistency is about the similarity between the result of different runs of a rule extraction algorithm for the same or similar input. Andrews et al. [7] define it to be the consistency of the *prediction* of different rule sets generated from the same subsymbolic model, whereas Johansson et al. [273] define it as the similarity of these different rule sets themselves. This gives rise to the difficulty of defining similarity between rule sets though. Huysmans et al. [249] propose a formal definition for this but in our case it is sufficient to define similarity to be the average size of the change set computed by our orthotope difference module in Outnumber between one run of the rule extraction and another without any change in the subsymbolic model (cf. Section 7.7.2.7).

To contrast accuracy and fidelity: if the trained subsymbolic system makes a prediction error for new data but the extracted rules do not, this increases accuracy but decreases fidelity. While sometimes an increased accuracy over the underlying model and thus a better generalisation capability is actually a goal of rule extraction techniques [249, 357], in our case we want to give the user an impression of how the subsymbolic system actually behaves. Therefore fidelity is our prime goal, even though it must be balanced with comprehensibility.

7.7.2.4 Umber Rule Extraction Requirements

For choosing a rule extraction algorithm we first defined a 7-point wish-list of items we would ideally like to see in a rule extraction algorithm. However, we will see that it is not possible to accommodate all of these requirements in the same algorithm. These requirements are:

- a high comprehensibility of the extracted rules,
- prioritisation of fidelity (model prediction) over accuracy (data prediction),
- independence from the learning algorithm used (i. e. being *pedagogical*),
- the option to perform regression as well as classification (even though we did not actually implement regression within the scope of this work),
- consistency, meaning a certain degree of stability of the generated rules (i. e. small changes in the model lead to small changes in the rule set), but
- on the other hand we would like to let the users generate a forest of different decision trees from which the one with the best comprehensibility or fidelity can be selected. And finally we
- would like to be able to generate different types of rules (e. g. propositional if-then, *m*-of-*n*, fuzzy rules) because some types might be easier to understand than others.

Of course, a good overall performance and availability on relevant platforms are implicit requirements.

The tabular overviews of rule extraction algorithms found in Huysmans et al. [249] and Martens et al. [356] were very helpful in narrowing down the large amount of algorithms to the following two, following the above items.

The two methods that best fit our requirements and that were found to perform very well overall [358] were TREPAN [95] – which works somewhat similar to other decision tree algorithms such as C4.5 but prominently uses a very sophisticated way to estimate missing data points from the underlying trained model – and G-REX [272] which uses evolutionary programming to find its rule sets. Neither one perfectly fits all requirements though. TREPAN is largely deterministic and thus does not generate alternatives by itself. G-REX on the other hand does so quite naturally and can even be encouraged to do so [298]. Because of the stochastic algorithm, it conversely has an increased instability of rules though, beyond the normal informational instability [295]. There is a feature that is only partially implemented to date,[23] allowing the usage of existing trees to seed the initial population. This would greatly increase the needed consistency between runs while maintaining the variability when this is desired.

Regarding most other criteria, both techniques perform equally well. G-REX, however, also has the advantage of being able to produce different types of rules. We therefore chose G-REX over TREPAN but we will briefly present both candidates.

The decision trees G-REX generates can be tuned to be more accurate or less complex as desired by changing the parameters of the benefit function that selects the best trees in each generation.

7.7.2.5 TREPAN

Developed by Craven [92], TREPAN might be the best-known rule extraction algorithm. Despite its age, it is still relevant today due to its good performance [19, 295, 358] which is mainly attributed to its data generation feature [356]. In the following, we will briefly explain how it works.

[23] Source: private communication with the author of G-REX.

TREPAN is similar to traditional decision tree learning algorithms [95] like C4.5 [434] or ID2-of-3 [389]. This means it recursively splits a training data set into subsets according to which new rule optimally splits the remaining data into different classes, thereby approximating the class boundaries. Like ID2-of-3, it produces *m*-of-*n* rules, thereby using more than one feature for each split. Unlike all conventional algorithms, however, it makes use of a so-called *oracle*, i. e. a predictive model that is used to generate new data points for nodes further down the tree where the remaining data sets are becoming too small. This means that a guaranteed minimum number of instances is used for each node. The artificial instances are generated using each feature's marginal distribution to sample the input space. These marginal distributions are either determined via frequency counting for discrete data or by using a probability density estimation after Silverman [496] for continuous features.

The oracle is also used to relabel all training data (which, amongst other things, increases the fidelity) and to determine the probability a given node only covers one class, serving as a stopping criterion for the algorithm.

Craven [92] and Huysmans et al. [249] describe TREPAN in more detail.

7.7.2.6 G-REX

G-REX is a rule extraction mechanism by König et al. [272, 295, 297] that uses genetic programming.[24] Genetic programming is a type of biology-inspired optimisation strategy proposed by Koza [304] that – unlike similar approaches in evolutionary computation such as evolutionary programming [150], evolution strategies [437, 485] or genetic algorithms [240] – uses tree-based representations, even though other data structures have since been proposed [22, 129, 315]. As such it generates one or more populations of decision trees and each population is then evolved by crossing over and mutating individuals that

[24] G-REX correspondingly stands for G*enetic* R*ule* Ex*traction*.

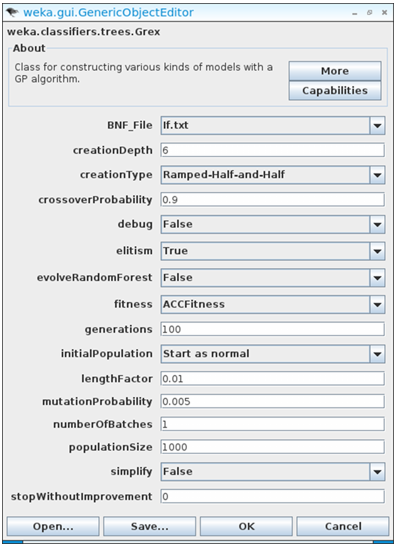

Figure 7.28. G-REX configuration window in Weka.

perform best according to a fitness function. Different fitness functions can be used but usually it will contain a punishment term for rule size and a term for rule quality such as fidelity, accuracy or area under curve (AUC) [273, 299].

7.7.2.7 Outnumber – Umber Rule Extraction

When the user presses the "Generate Rules" button in the Lumber control window (cf. Figure 7.14 on page 241), the software first checks whether a current model exists, i. e. whether building one was manually triggered from the "Train" button. If this is not the case, this is done implicitly.

G-REX runs as a module within the Weka framework[25] [205]. The

[25]The module itself is called `weka.classifiers.meta.RuleExtract` whereas the

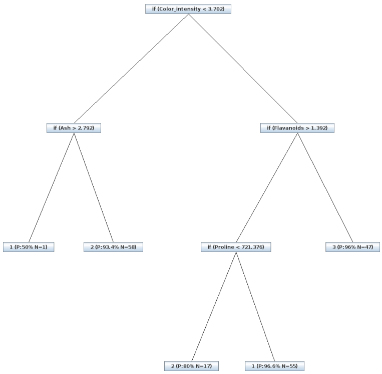

(a) G-REX GUI Tree

```
if Color_intensity < 3.702
|T: if Ash > 2.792
|   |T: 1 {N=1 P=0.500}
|   |F: 2 {N=58 P=0.934}
|F: if Flavanoids > 1.392
|   |T: if Proline < 721.376
|   |   |T: 2 {N=17 P=0.800}
|   |   |F: 1 {N=55 P=0.966}
|   |F: 3 {N=47 P=0.960}
```

(b) G-REX ASCII Tree

Figure 7.29. Comparison of the tree output visible in the G-REX GUI and the plain text output it produces that is used in the Lumber processing pipeline. The tree was generated from the wine data set [153] using G-REX.

only way G-REX can output the decision trees it generates in an easily parsable format is a custom human-readable plain text format (cf. Figure 7.29(b)). This is then parsed and one rule is generated for each leaf by tracing the tree up towards its root. The resulting rule set is then exported as PMML, the Predictive Model Markup Language [104, 199] which is an XML standard and allows interaction with third-party software.

This PMML decision tree is then fed into the rule editor (cf. Figure 7.35 on page 293) where the actual revisiting and manipulation of the rules by the user takes place. The resulting rule set is then converted to an orthotope data structure and compared to the set of orthotopes

configuration window shown in Figure 7.28 on page 275 is from `weka.classifiers.trees.Grex`. This is because the meta classifier needs two further modules as a parameter, one of them being the black box model, the other one the rule extraction algorithm.

corresponding to the rule set before the user manipulation by computing the pairwise difference between the two sets. Per orthotope, this means that

1. either two orthotopes are disjoint – in which case the output set is equal to the input set – or
2. the first lies completely within the other – in which case it is just removed – or
3. a set of new orthotopes is generated that constitutes the volume of the first orthotope without the second one.

The representation of rules as sets of orthotopes has the advantage of being more universal and versatile but poses problems for the special cases of orthotopes with volume $V = 0$ or $V = \infty$. The former would happen for "is equal to" rules and is solved by giving these orthotopes a fixed minimum dimension.[26] The infinite case describes orthotopes with one or more missing boundaries, i. e. a variable that is not restricted to an interval but only defined to be smaller or larger than a certain value. The problem for sampling from such rule sets is not specific to the orthotope representation, however, and is currently not properly addressed. One quick but not entirely adequate fix would be to include a step before the sampling that adds missing boundaries based on maximum values derived from existing data. Hardcoding these maximum boundaries per variable based on world knowledge and user testing would be another possibility. A more definitive solution would be to transform all features into a space that has inherent bounds, similar to the step we suggested for the idle time feature of Lumber (cf. Figure 7.10(b) on page 237).

The result of the difference operation is a new set of orthotopes to be added that can then be directly passed on to the sampling module. However, it might turn out to be a reasonable step to ignore orthotopes below a certain minimum extent in all directions because those might arise in great number from small fluctuations in the exact orthotope boundaries.

[26] Special treatment during the sampling step would be an alternative as this, unlike the orthotopes of infinite volume, clearly poses no fundamental problem.

The deletion of rules is currently not handled. Deletion of rules has some fundamental challenges, too.

7.7.3 Knowledge Insertion

Knowledge insertion – also sometimes called *knowledge synthesis* [368], *rule insertion* or *insertion of symbolic information* [491] – is the process of initialising or adapting a subsymbolic model using symbolic information.

While being pedagogical, i. e. independent from the subsymbolic model, is an established property of rule extraction techniques (cf. Section 7.7.2.1), a strong dependence on the model used is the norm for knowledge insertion methods [155, 156, 242, 294, 368, 369, 381, 402, 407, 408, 491, 535].

The only alternative to manipulating the model directly is to create virtual training data according to the symbolic information. However, these training data need not be positive examples for a particular class [590].

Creating such negative examples is one possibility to implement the deletion of rules. Logically, this is equivalent to the creation of positive data points in the complete feature space *outside* the orthotope that is to be deleted or suppressed, thereby effectively diluting the existing points. However, this approach is most likely to be infeasible in practice. Probably the most practical way to handle deletion is to introduce a module that removes a percentage of data points that fall within the orthotope boundaries that are to be removed from the stored data set, since these data points must exist in the first place or else the rule would not have been generated. If changing the underlying black box is also an option, artificial data points that carry negative weight could be used as well.

7.7.3.1 Challenges

While generating artificial training data from a given set of orthotopes might sound like a straightforward process, the devil is in the detail

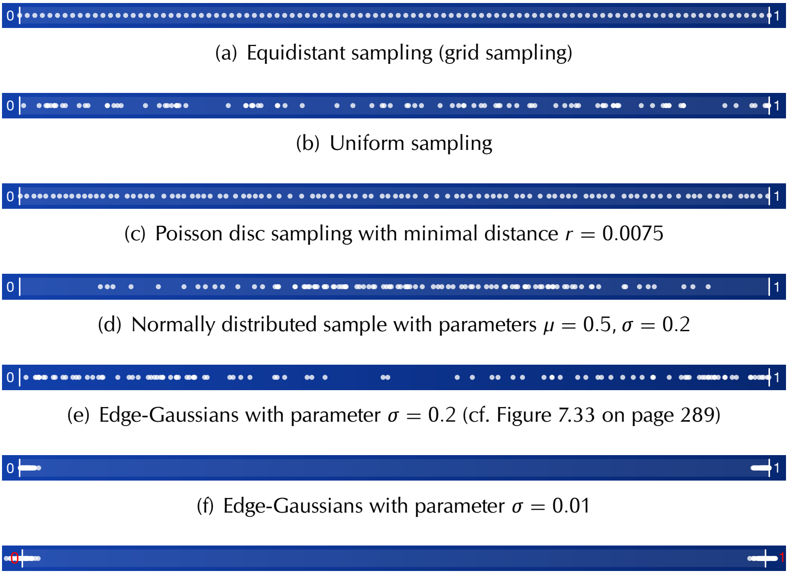

(a) Equidistant sampling (grid sampling)

(b) Uniform sampling

(c) Poisson disc sampling with minimal distance $r = 0.0075$

(d) Normally distributed sample with parameters $\mu = 0.5, \sigma = 0.2$

(e) Edge-Gaussians with parameter $\sigma = 0.2$ (cf. Figure 7.33 on page 289)

(f) Edge-Gaussians with parameter $\sigma = 0.01$

(g) Gaussian sampling with interval boundaries as means ($\sigma = 0.01$)

Figure 7.30. Different sampling strategies exemplified in 1D. All examples use 100 points and are meant to represent the interval [0, 1]. Of course higher dimensions are likely to behave quite differently due to the curse of dimensionality [31, 214].

here. The most important question that has to be answered in this regard is that of the probability distribution to use. We will explore this further in Section 7.7.3.3 but for now let us consider the different design goals that might demand different answers. So when a user defines an orthotope boundary, how exact is this expected to be represented in the subsymbolic model? How malleable by future “real” training examples is it supposed to be? Should it stand fixed for eternity because the user decreed it or was it meant just as a starting point and to be refined by actual experience? Is change in one direction preferred over change in the other direction? How bad would “holes” in the orthotope be in case the coverage over the whole volume is not sufficient?

These questions directly affect the choice of density of virtual training data compared to that of the real training data and it also affects the distribution of the relative density of that artificial data. It is also obvious that it is unlikely that there can ever be one perfect answer to these questions that suits all scenarios and all users equally. So one would either have to make a design decision suited to the application at hand or pass some of these decisions along to the user as a parameter.

7.7.3.2 Renumber – Umber Knowledge (Re-)Insertion

For the time being, Umber uses a simple multivariate Gaussian distribution with parameter $\mathbf{\Sigma}$ being configurable by the user.[27] This is little more than a placeholder distribution though and in the following sections we will explore further avenues that can be taken here.

7.7.3.3 Sampling Methods and Distributions

Generating data to add the information from new rules to a subsymbolic model can follow different strategies. Arguably the simplest strategy would be to equidistantly sample the interval representing the new rule (grid sampling; cf. Figure 7.30(a) on page 279). To avoid any aliasing effects and get more realistic data points, one could also produce uniformly distributed random numbers (cf. Figure 7.30(b) on page 279) or use more sophisticated methods like a Poisson disc sampling [314] (cf. Figure 7.30(c) on page 279). With high-dimensional orthotopes, these approaches will need a lot of data points, though, since the points will become more and more spread out. It is therefore worthwhile to think of ways to concentrate data points in a meaningful way. This leads to a number of different patterns and only tests with users in real-world situations can show which of these patterns is the most appropriate. In fact, as previously said, several of them might be similarly viable for different purposes and it could be that in the end it turns out to

[27] The covariance matrix $\mathbf{\Sigma}$ is required to be a diagonal matrix.

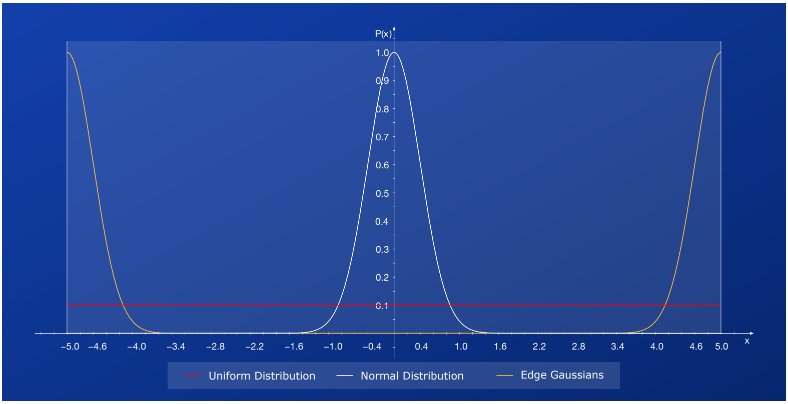

Figure 7.31. Examples for 1D probability density distributions for knowledge insertion. Exemplary orthotope edges are at $x = -5$ and $x = 5$. The Gaussians use standard deviation $\sigma = \frac{1}{\sqrt{2\pi}} \approx .4$ for unit $\max_x P(x), x \sim \mathcal{N}(\mu, \sigma^2)$.

be most practical to offer users the possibility to choose from different methods (probably described by their function such as *exact thresholds* or *treat values as rough estimates*).

Since many real data sources will be approximately normally distributed thanks to the central limit theorem [539], emulating this by using normally distributed random numbers also suggests itself. The Gaussian approach has the drawback, though, that rules can be expected to be usually specified by the user with rather meaningful interval borders, whereas the Gaussian samples these borders only very sparsely which will lead to a lot of noise regarding the exact borders.[28] Sampling only the class boundaries would be one obvious way out but this would make it very difficult for the system to refine the borders from further

[28] It should be noted that this problem can be expected to be much less pronounced for high-dimensional data as the number of samples in the tail of the distribution increases with the number of dimensions. It might still be advantageous, however, to have a sampling function that emphasises the edges inherently.

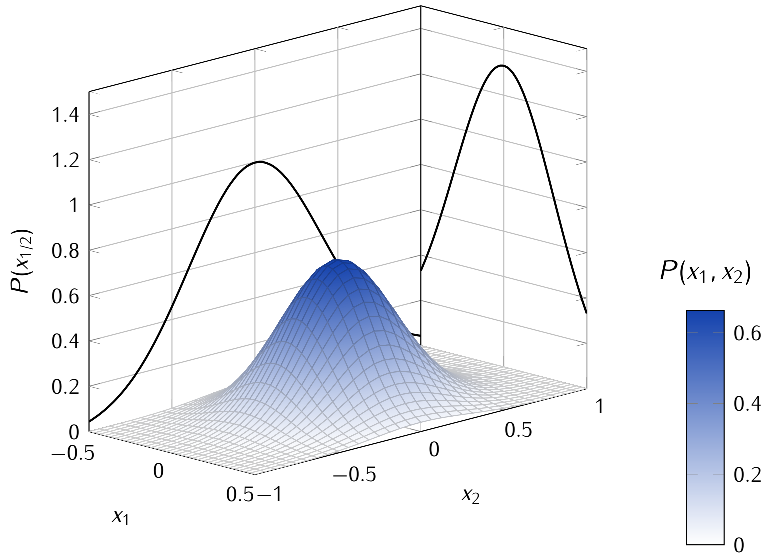

Figure 7.32. Bivariate normal distribution to sample two-dimensional rules and demonstrate multivariate distributions. Orthotope edges are $x_1 \in [-0.5, 0.5]$ and $x_2 \in [-1, 1]$ and standard deviations used are $\sigma_1 = 0.2$ and $\sigma_2 = 0.3$. In the background, the marginal distributions are shown, i. e. the univariate Gaussians with the same parameters.

experience and also there is the danger of the classifier omitting the entire volume of the orthotope.

Figure 7.31 on page 281 suggests one approach to solve this problem in 1D, by adding two Gaussians with their means exactly at the interval borders instead of at the centre[29] (we are going to call these *edge Gaussians*). The exaggeratedly small variance used in the figure

[29]Of course the actual mean of the whole distribution is the orthotope centre $\boldsymbol{\mu}$ while the local mean around a class boundary will not be the boundary itself. One might therefore extend the orthotope boundaries by one standard deviation to compensate this or just try to add random Gaussian noise after sampling the class boundaries; a technique

is probably not advisable since then there will be said danger of the classifier omitting the centre of the orthotope.[30]

For a one-dimensional interval, uniform and Gaussian probability density functions are illustrated in Figure 7.31 on page 281 and actual random samples of all of the different sampling strategies discussed above are shown in Figure 7.30 on page 279. While generalising uniform and Gaussian distribution to higher dimensions is trivial (although the consequences might be everything but), the edge Gaussians behave less benevolently.

7.7.3.4 Multivariate Edge-Gaussian

A naïve approach to construct a two-dimensional analogue to the edge Gaussians we just saw would be to just add four univariate Gaussians instead of the two used in Figure 7.31 on page 281. With the edges of the 2D orthotope given by $H = [x_{11}, x_{12}] \times [x_{21}, x_{22}]$, this yields

$$\begin{aligned} P(x_1, x_2) = a \Bigg(& \tfrac{1}{\sigma_1\sqrt{2\pi}} \exp\left(-\tfrac{(x_1-x_{11})^2}{2\sigma_1^2}\right) + \tfrac{1}{\sigma_1\sqrt{2\pi}} \exp\left(-\tfrac{(x_1-x_{12})^2}{2\sigma_1^2}\right) \\ & + \tfrac{1}{\sigma_2\sqrt{2\pi}} \exp\left(-\tfrac{(x_2-x_{21})^2}{2\sigma_2^2}\right) + \tfrac{1}{\sigma_2\sqrt{2\pi}} \exp\left(-\tfrac{(x_2-x_{22})^2}{2\sigma_2^2}\right)\Bigg) \\ = a' & \sum_{i=1}^{2} \left(\frac{1}{\sigma_i} \sum_{j=1}^{2} \exp\left(-\frac{(x_i - x_{ij})^2}{2\sigma_i^2}\right)\right) . \end{aligned} \tag{7.15}$$

With a and a' being factors to normalise to unit integral. When plotted for the same interval as Figure 7.32 on page 282, the following surface is the result (normalisation not applied):

that we will not further discuss here but that might be perfectly viable given an adaptive and sufficiently large variance.

[30] If the variance is high enough for the definite integral within the interval to significantly differ from 1, a renormalisation is due. This was not done in Figure 7.31 on page 281 since $1 - \int_{-5}^{+5} P(x)\,dx < 5 \cdot 10^{-36}$. The same normalisation factor can be used for the normal Gaussian and the edge Gaussians, though.

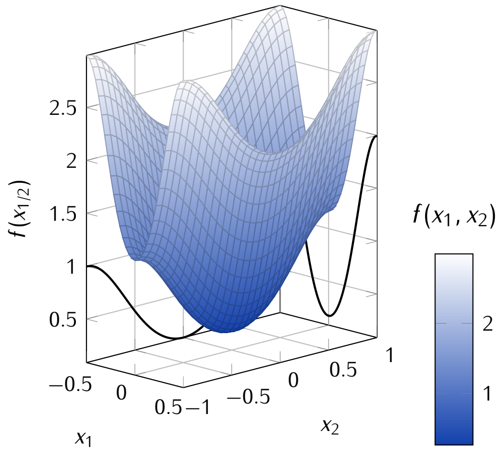

As can be seen, this approach leads to an extreme emphasis on the corners of the orthotope. While it cannot be ruled out that this might actually be useful, we want to try to find a distribution that does not show this behaviour.

The idea is to start with a vertically mirrored version of the normal distribution, i. e. $f(\boldsymbol{x}) = \max\big(P(\boldsymbol{x})\big) - P(\boldsymbol{x})$, where $P(\boldsymbol{x})$ is the probability density function of the multivariate normal distribution.

We will in the following quickly sketch how we got to our sampling distribution in two dimensions and then properly develop the function in k dimensions. We exemplify and test the result in two dimensions again in Appendix C.1, including a fully normalised version.

For the initial development of our distribution, we will use a unit variance, zero mean Gaussian in the interval $x_{1/2} \in [-1, 1]$ without paying attention to any integral normalisation. We renormalise to unit height instead and get

$$f_1(x_1, x_2) = 1 - e^{-(x_1^2 + x_2^2)} \tag{7.16}$$

which looks like this:

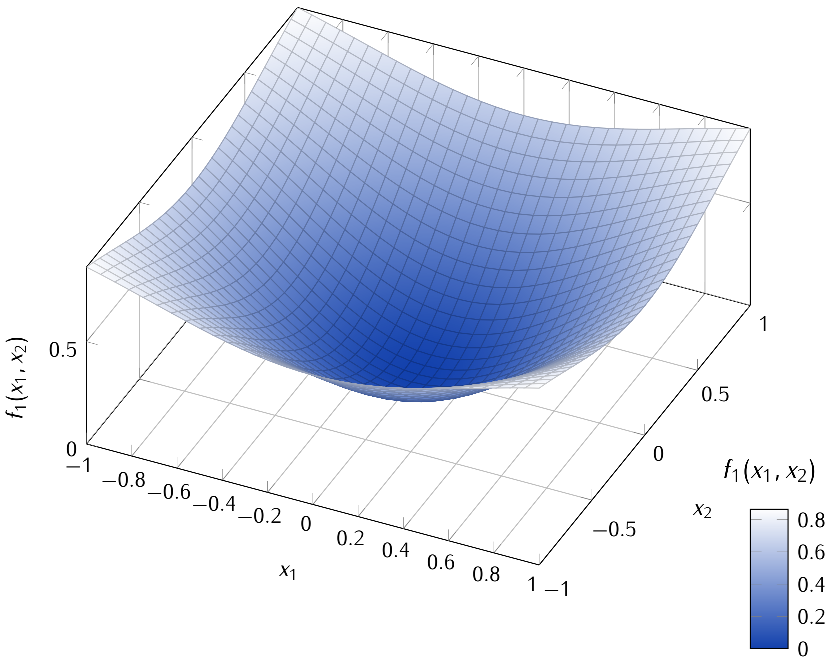

The corner preference is less pronounced but still present. We fix this by adding the difference to 1 at the orthotope edges to the function value, which is $f_1(1, x_2) = e^{-1-x_2^2}$ and $f_1(x_1, 1) = e^{-1-x_1^2}$ respectively. Due to the inherent symmetry, each term straightens both the $x_i = -1$ and the $x_i = +1$ edge. The effect of these correction terms on the shape of the function decreases with the distance to the respective edge. With both correction terms added, we get

$$f_2(x_1, x_2) = 1 - e^{-(x_1^2 + x_2^2)} + e^{-(1+x_2^2)} + e^{-(1+x_1^2)} \tag{7.17}$$

looking as follows:

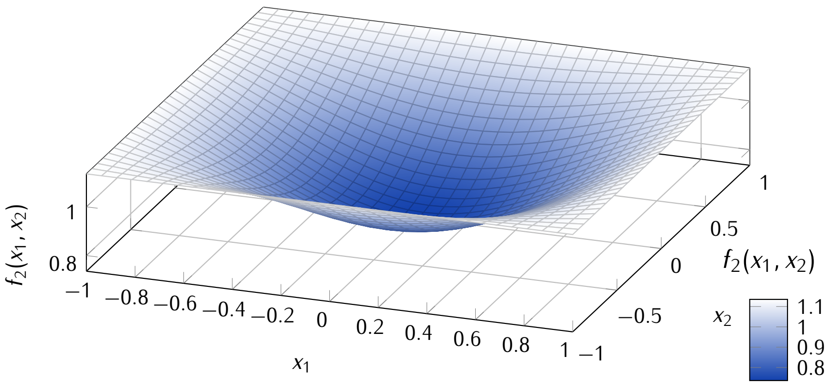

The shape of the function can best be seen in the contour plot.

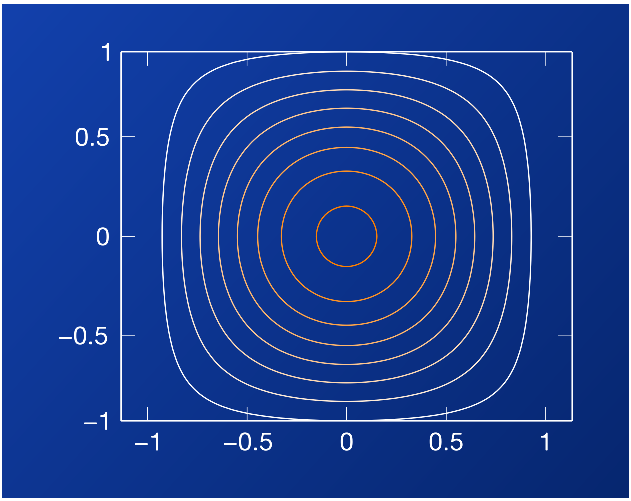

Generalisation to k Dimensions For the generalised k-dimensional case we can without loss of generality still assume the domain of $\boldsymbol{x} \in H$ to be a hypercube of double unit height around $\boldsymbol{\mu}$ – defined as a Cartesian product of intervals $H = [\mu_1 - 1, \mu_1 + 1] \times [\mu_2 - 1, \mu_2 + 1] \times \ldots \times [\mu_k - 1, \mu_k + 1]$ – and an identity covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{I}$. The reason is that in our case, the variance is a function of the edge length of the orthotope and therefore anisotropic scaling can be used on the hypercube in a second step, keeping the covariance matrix a diagonal matrix since the orthotope is paraxial.

We can therefore simplify the density distribution function of the multivariate normal distribution:

$$\begin{aligned} P(x_1, \ldots, x_k) &= \frac{1}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{\mu})^\mathsf{T} \boldsymbol{\Sigma}^{-1} (\boldsymbol{x} - \boldsymbol{\mu})\right) \\ &= \frac{1}{\sqrt{(2\pi)^k}} \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{\mu}\|^2}{2}\right) \end{aligned} \tag{7.18}$$

In fact, we will start with $\boldsymbol{\mu} = \mathbf{0}$ and drop the normalisation factor because it obviously ceases to work properly when restricting the domain to a finite orthotope. This leaves us with the following function, defined on $[-1, 1]^k$.

$$P'(\boldsymbol{x}) = \exp\left(-\frac{\|\boldsymbol{x}\|^2}{2}\right) \tag{7.19}$$

Using k correction terms, we will then straighten all $2k$ edges, each of which is an orthotope of dimension $k - 1$.

$$f_1(\boldsymbol{x}) = P'(\boldsymbol{x}) - \sum_{i=1}^{k} P'(x_1, \ldots, x_{i-1}, 1, x_{i+1}, \ldots, x_k) \tag{7.20}$$

To scale the hypercube to a general orthotope, we need a vector of scaling factors $\boldsymbol{s}, s_i > 0$. Given an orthotope $O = [x_{11}, x_{12}] \times [x_{21}, x_{22}] \times \ldots \times [x_{k1}, x_{k2}]$ we can get $\boldsymbol{\mu} = \left(\frac{x_{11}+x_{12}}{2}, \frac{x_{21}+x_{22}}{2}, \cdots, \frac{x_{k1}+x_{k2}}{2}\right)^\mathsf{T}$ and $\boldsymbol{s} = \left(\frac{|x_{11}-x_{12}|}{2}, \frac{|x_{21}-x_{22}|}{2}, \ldots, \frac{|x_{k1}-x_{k2}|}{2}\right)^\mathsf{T}$. As can be seen, the orthotope edge length has to be divided by 2 to get the scaling factors s_i. So the unscaled hypercube $[-1,1]^k$ we started with has an edge length of 2 but a scaling factor of 1. We can then write

$$f_2(\boldsymbol{x}) = f_1\left(\operatorname{diag}(\boldsymbol{s})^{-1}\boldsymbol{x}\right) \quad , \quad \boldsymbol{x} \in [-s_1, s_1] \times [-s_2, s_2] \times \ldots \times [-s_k, s_k] \tag{7.21}$$

where $\operatorname{diag}(\boldsymbol{s}) = \begin{pmatrix} s_1 & & \\ & \ddots & \\ & & s_k \end{pmatrix}$ We subsequently translate the origin-centred orthotope to get a function that is finally defined on the target orthotope O, i.e. $\boldsymbol{x} \in O$.

$$f_3(\boldsymbol{x}) = f_2(\boldsymbol{\mu} - \boldsymbol{x}) \tag{7.22}$$

We then mirror the function so that it has its minimum at $f_4(\boldsymbol{\mu}) = 0$

$$f_4(\boldsymbol{x}) = f_3(\boldsymbol{\mu}) - f_3(\boldsymbol{x}) \tag{7.23}$$

Finally, we can renormalise the integral within the orthotope to 1 to get a probability distribution, although this step can be optional (for example when using the common Metropolis–Hastings sampling [215, 580]) and it might even be advisable to leave it out because of the significantly increased complexity (cf. Equation (C.8) on page 450).

$$f_5(\boldsymbol{x}) = f_4(\boldsymbol{x}) \left(\int_{x_{11}}^{x_{12}} \cdots \int_{x_{k1}}^{x_{k2}} f_4(\boldsymbol{x})\, dx_k \ldots dx_1 \right)^{-1} \tag{7.24}$$

To see the expanded forms of these equations for the two-dimensional case and additional plots, see Appendix C.1.

Test Sampling Figure 7.33 on the next page shows a 1D histogram of a large number of samples drawn from the non-normalised Equation (7.23) using Metropolis–Hastings sampling. Some artifacts of this method are visible when comparing the shape of the histogram to the shape of the function itself but the qualitative behaviour is well approximated. Figure 7.34 on page 290 shows a 2D Metropolis–Hastings sampling. Here, a decreased concentration towards the corners is apparent compared to

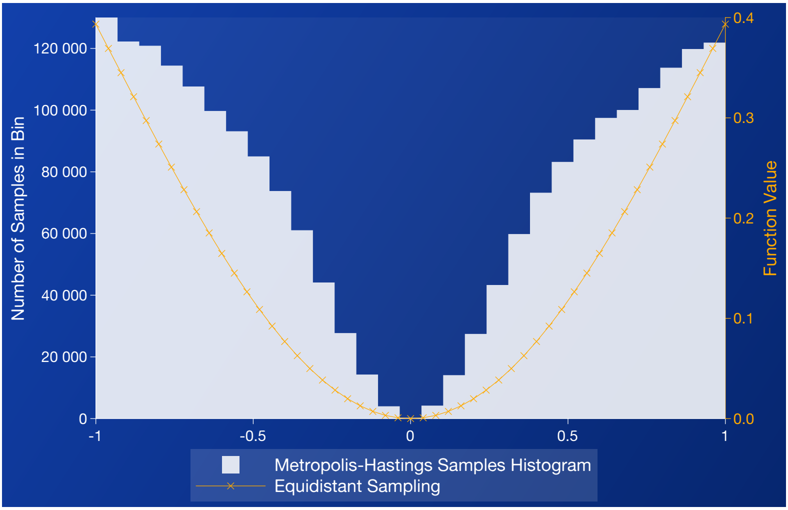

Figure 7.33. Histogram of 1D Metropolis–Hastings sampling of the edge sampling probability distribution from Equation (7.23) on page 288 (left axis; accepted samples $1.5 \cdot 10^5$, $\boldsymbol{x} \in [-1, 1]$), compared to its equidistantly sampled and linearly interpolated function values (right axis).

the centre of the edges. A modification of the density distribution function to compensate for this effect (maybe even something along the lines of Equation (7.16) on page 284) or using a different sampling method might be steps worth considering if this poses problems in practice.

It might be worth emphasising that the large number of samples used in both plots is solely to provide a more accessible and more pleasant visualisation. It is in no way indicative of the typical number of samples that would be used for training.

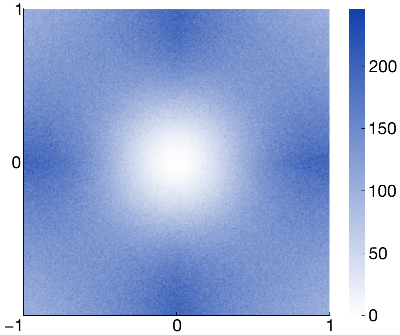

Figure 7.34. 500 × 500 bin 2D histogram of Metropolis–Hastings sampling of the edge sampling probability distribution from Equation (7.23) on page 288 or Equation (C.7) on page 448 respectively (accepted samples $3.1 \cdot 10^7$).

7.8 Conclusion

7.8.1 Summary

In this chapter, we showed avenues towards a fully context-aware system that filters information and controls SCS channels in order to minimise interruptions and maximise calmness, and we reported on the technological and conceptual progress we have made so far. This should allow the user to understand and influence its subsymbolically determined decisions in a symbolic way. *Context awareness* here describes the ability of the system to assess the situation the user is in, and modify its behaviour accordingly (see Section 7.4 for details).

The focus of this work was on data production, rule extraction and user interaction. This means that the evaluation of the quality and usability of the rule extraction and the full closure of the symbolic–subsymbolic loop fall outside the scope of this work and are to be left to future research.

We presented the current state of our prototype system Umber and its components. The data collection on laptop and smartphone platforms is working well, as is the feature extraction on both platforms, except for the string features that would need to be changed in a productive system and even in its current state, the parametrisation is not based on empirical data.

The merging and preprocessing stages of the data pipeline would require some improvements and unification to be ready for a fully closed loop while the feature selection needs to be more empirically grounded on both platforms. The classifiers work equally well on both platforms but only the Lumber classifier has a connection to the rule extraction whereas only the Cucumber classifier has been tuned using real user data and tested under realistic conditions. Both need to be exposed to longer term data to see how especially the memory requirements scale with the amount of data. Rule extraction and display of the extracted rules in the Number GUI work but the extracted rules from the limited data collection with Slumber lack any comprehensibility and intuitive sense. They also vary greatly between runs so that seeding G-REX with the last result is a necessity, a feature that has not yet been fully implemented. Alternatively, the rule extraction module had to be replaced by a different algorithm, the most likely candidate being TREPAN.

Rule set difference and orthotope conversion are implemented but not yet tested with real-world data due to the complications just mentioned. We expect some work to be found in ironing out the particularities of these steps, even though the scaffolding is in place. Still missing is a module that handles the deletion of rules as a special case, however (cf. Section 7.7.3).

Lastly, the knowledge insertion is only rudimentarily implemented by using multivariate Gaussians which we believe to be inadequate. We proposed an alternative that we view as more promising conceptually but this is as of now unproven. We also explained that we doubt that only one method of sampling would be viable in practise when it comes to user friendliness and comprehensibility. Instead, we believe it to be important to have multiple a distributions in the toolbox, including one that generates edge-focused samples which is why we derived such a method.

Another important step towards a comprehensible system would be to replace the raw features by meaningfully labelled linear combinations of features but this is not yet considered in the existing code and also bears some conceptual difficulties regarding circular dependencies and overlap between what constitutes classes and what situational labels that take the role of features (see below).

7.8.2 Outlook

The hypotheses we would like to test with Umber regarding increased usability (meaning effectiveness, efficiency, safety, utility, learnability, and memorability [452]) or – to put it in more modern terms – a better user experience[31] (which includes among many other things usefulness, desirability, credibility, accessibility, and user preferences [260, 452]) through hybrid transformation architectures are the following:

1. Non-expert users can understand the rules produced from a subsymbolic adaptive system (and can do so better than the original system's parameters itself).
2. Such inspection possibilities make them more comfortable with such systems.
3. They can manipulate such rule sets in a sensible way, thereby contributing their domain knowledge.
4. Such manipulation possibilities give them a sense of control over the system.
5. This further enhances the acceptance of adaptive systems.

[31] While user experience (in ISO 9241-210 [260] defined to be a "person's perceptions and responses resulting from the use and/or anticipated use of a product, system or service") does not seem to be a comparable term to usability (the "extent to which a system, product or service can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use"), it has become to stand for a more holistic view of HCI that includes the whole experience a user has with a system. The definition of usability has followed suit though and in ISO 9241-210, it is remarked that "[usability], when interpreted from the perspective of the users' personal goals, can include the kind of perceptual and emotional aspects typically associated with user experience. Usability criteria can be used to assess aspects of user experience".

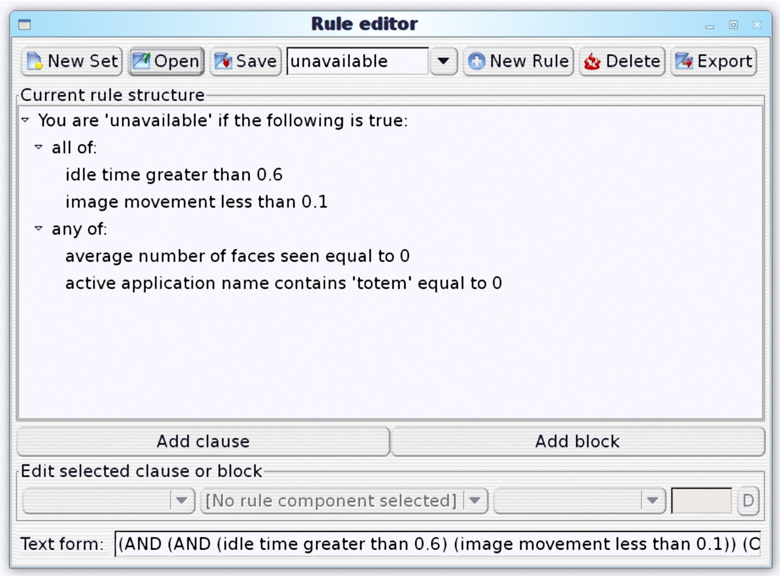

Figure 7.35. The current draft of the rule editor. It serves to present the generated rules to the users and allows them to edit them. The knowledge insertion back into the subsymbolic model is currently a work in progress, as are the usability of the GUI and the presentation of the features.

6. Hybrid systems are actually preferred by users over purely rule-based systems due to their non-absolute interpretation of rules, learning-from-example-capability and better generalisation abilities.

There are also some more technical hypotheses that need to be tested:

7. It is practically possible to predict complex context relations such as a user's interruptibility with available sensors (with or without the intervention of the user through rule interfaces). There are strong indications in works by other authors that this is indeed the case [77, 107, 147, 148, 149, 213, 236, 387, 470, 593] and our own results seem very promising in this regard as well.

8. A practical amount of fidelity for comprehensible rule sets is possible.
9. Inserted knowledge from rule changes is reflected in re-extracted rules.
10. A practical amount of rule stability is possible.

A further aspect that needs to be addressed is the preferred presentation and manipulation interface. This includes the kind of rules that are extracted from the model (fuzzy rules for example are often said to be more comprehensible) as well as the user interface itself.

Presentation alone can not compensate for a fundamental issue with this kind of generated rules though, which is the fact that the symbols used are low-level features meant to be interpreted by machines, not humans. For some features, a simple translation layer might be sufficient; for example the time features described in Section 7.5.2.1 might not be very legible in their raw format but can be easily translated into a human-readable format. Fuzzy rules can themselves be seen as another form of translation layer.

Such a translation is surely not an option for all kinds of features, though, and a more fundamental escape from this problem would be to use unsupervised learning to find structures in the data and then present recordings from typical points in these clusters to the user who can then label them. The hope is that this would mean that separate kinds of situations that are present in the data can be associated with their natural-language description and this can then be used as a more readily understood symbol that might represent not a single feature but an arbitrary linear combination of features, defined by the cluster centres.

Lastly, it is clear that generated rules can always be only a rough approximation of the underlying model or the incomprehensibility of the rules would not be far from the obscurity of the stochastic model itself. This also explains why rule-based classifiers are no real alternative to our approach: in order to decently model the necessary patterns, their complexity would exceed the comprehensible amount of rules by far.

8 Conclusion

8.1 Summary

Subliminal copresence systems (SCS) are systems designed to fill a gap left by telecommunication systems that simulate face-to-face communication. They do this by creating copresence without necessarily creating social presence at the same time. They are intended to be left running in the background over hours at a time and a necessary prerequisite for this is calmness while privacy is of particular concern. We looked at these and other design considerations for SCS and gave an overview of the literature in relevant fields, especially awareness research. In the process, we tried to consolidate concepts and taxonomies from many different areas as much as possible.

We implemented two exemplary SCS called FEELABUZZ and upstairs. While FEELABUZZ is a mobile system that makes use of readily available hardware by using off-the-shelf smartphones and uses vibrations to transmit movements, upstairs is a fixture that uses special contact microphones to make people living far apart virtual upstairs neighbours of each other.

Having had proven the fundamental capability of FEELABUZZ to recognisably transmit basic activity types, we conducted one small scale user study for each of the two systems with couples as users of the system. For these studies, we used a set of questionnaires that we developed specifically to investigate SCS while building on existing and established measures where possible. Results indicate that both systems achieved their goal of creating copresence and doing so to a larger extent than they created social presence.

To assess calmness, we developed a cognitive load scale that exhibited a satisfactory internal consistency. The overall cognitive load it measured was low for upstairs while FEELABUZZ showed signs of being less calm.

The upstairs data appeared to show that the new Cognitive Load scale consists of two factors that measure the ability of a user to ignore a system and how annoyed the user feels by a system (we called the two subscales *Ability to Ignore* and *Vexatiousness*). Upstairs scored very good on the Vexatiousness scale and moderately good on the Ability to Ignore scale. However, the data from the FEELABUZZ study shows that further research is needed to verify the existence and composition of these subscales.

Participants felt that privacy was no concern with either of the two systems. The role context can play in helping to interpret low-level signals was confirmed.

We also found indications that copresence and telepresence might not be distinct concepts where SCS are concerned and that the systems we built do not seem to transmit emotion well, regardless of modality.

However, while we measured the *transmission* of emotion over SCS (or lack thereof), we still do not possess of a good way to assess the emotional impact of these systems, other than the feeling of copresence itself. Given that we postulate this to be one of their key characteristics, finding such a measure would be an important extension of our research method. However, the participants of our studies did seem to get emotionally attached even to the system itself, since most of the upstairs users described feelings of sadness or loss when they had to stop using it. About 77 % of FEELABUZZ users and 100 % of upstairs users stated that they would have wished to continue using the system, albeit some made some technical improvements a prerequisite for this.

A further refinement of our method should take a look at whether or not the notion of connectedness is salvageable in the context of SCS, given the important role this concept plays in the awareness systems literature and how potentially relevant it might be for SCS as well, just looking at its definition instead of the existing scale and also looking at the fact that people described a feeling of increased connectedness for both systems.

Based on our experiences with those first two SCS, we saw a need for automatic filtering and began building a system called Umber that not only learns to automatically detect the situational context of users but allows them to investigate the results of this machine learning as symbolic rules. We extended this idea to constitute a full human-in-the-loop system that would allow users not only to investigate machine learning systems but also to influence them using the same rule interface. To this end, we built an interactive rule interface and did some precursor work on how to close the loop using orthotope representations of arbitrary rule sets and suggested a sampling method that extends the set of commonly used distribution functions, believing that no single distribution of artificial data points can fit all situations perfectly.

Our status recognition systems even without help from a symbolic branch reached a classification rate of 98 % for the laptop prototype and reduced the number of necessary inputs to classify the current situation by 38 % in the mobile and multi-label case.

8.2 Outlook

While some interesting preliminary results came from the two studies we performed with FEELABUZZ and upstairs, the most obvious circumstance preventing further and more profound conclusions to be drawn was the small number of participants. This is especially true given that some of the most interesting and most ambitious aspects of the upstairs study could not be pursued due to the high number of missing questionnaires. So the obvious next step has to be a follow-up study that is larger in scale, probably at least by an order of magnitude.

Such a large number would be especially important if the effect of SCS on relationship satisfaction is to be analysed, given how many hard to control for factors influence relationship satisfaction. The studies we conducted did serve to suggest that one would probably have to take care to ensure a wider variety of relationship qualities to begin with. It would also be interesting to follow up on relationships after permanent reunion. Our considerations in Section 2.5 indicate that such an effort might be worthwhile, given that SCS could bring something to LDRRs that normally only PRRs have, making this an interesting avenue of research even from a humanities point of view. As mentioned above, we find it impossible to predict whether SCS would stabilise or destabilise a relationship and how such an effect would relate to the kinds of relationships that survive a permanent reunion compared to those that do not.[1] The negative effects of the copresence created by SCS through a lack of solitude as voiced by at least one of the participants would also be an interesting aspect to investigate in the future.

Another vast area of possible future research is determining the exact relationship between the design guidelines and aspects of SCS we postulated in Chapter 3 and the measures such as copresence and closeness defined in Chapter 2. While we substantiated our guidelines with the findings of other authors and the results from our studies wherever possible, much of it are still merely opinions and convictions, be they shared among many researchers or held by only a few.

In particular, some of the research questions that would be interesting to investigate in the future are the following.

- Does implicit context communication as it was described in Section 3.1.3.1 help to stay in touch emotionally over a distance?
- How to sculpt information cues such that they can easily be processed on a subconscious level?
- Can implicit context communication help to simulate the effect of serendipitous interaction (cf. Section 2.3 and Box 2 on page 37) in vocational relationships or does implicit communication support explicit communication in any other way?
- Is the following hypothesis true: the less well you know a person or a group of persons (i. e. the less connected you feel to someone emotionally), the more you want to use artifical implicit communication channels with them in public spaces rather

[1] Of course, finding that SCS have no effect on relationship satisfaction or stability whatsoever is also always an option.

than in private places; i. e. what is the relation between group configuration and connectedness and intimacy?

– Does explicit-only communication work better than implicit-only?
– Does using only implicit or only explicit communication for a longer period of time make the other type of communication wither or otherwise negatively affects it?
– What is it about usually explicitly used communication channels that prevents them from being implicitly used?[2]

8.2.1 Other Possible Subliminal Copresence Systems

Sonification [230] is a promising candidate for future SCS where input and output modality are not the same. Study participants highlighted upstairs' use of sound as one of its major advantages with regard to subconscious monitoring and a low cognitive load in general.

Earcons and auditory icons have already been proposed by some authors in the context of awareness systems [263, 277, 563]. One of the techniques that we find promising in this regard is model-based sonification [232, 233] because it allows a large amount of data to be displayed in a way that is still processable by the human auditory system, potentially making it possible to find a mapping that follows the notion of *directness* to enable richness (cf. Definition 11). Blended sonification as a relatively new sonification paradigm seems especially well suited to implement calm sonification by blending in with the environment [545].

We had hypothesised that there are synergistic effects when a number of SCS are used simultaneously. Therefore developing and testing a wider variety of different SCS and looking at how they play together when used simultaneously is also an interesting prospect.

Ideas for such systems include seating furniture that warms up or glows when its remote counterpart is used, similar to the connected beds and benches by Goodman and Misilim [184] and Dunne and

[2] We put this personal impression of ours to the upstairs study participants in the interviews and they unanimously agreed that letting, say, a Skype connection open is not a desirable communication channel, although most had tried to do so. Nobody could pinpoint the exact reason for this, although some hypotheses were brought forward.

Raby [127] respectively. Further acoustic links that use ambient noise instead of structure-borne sounds could be small localised connected spaces such as kitchen cabinets, sinks or stoves.

Of course, technology is a crucial driver of which systems are possible, practical and affordable. With the advent of consumer-grade and unobtrusive augmented reality hardware (cf. Figure 8.1 on the facing page), this might soon be a readily available channel to use for calm communication as well, for example exploiting attention focus in various ways, such as a heat map of a moving average of the remote partner's attention focus (comparable to the concept by Mertes [375] but remotely instead of only in physically close interactions), although something similar could be achieved with projections on the output side and location on the input side. Colour could be added as an additional dimension, for example to indicate certain activity types.

Affordable printable displays could be used to better integrate visual media spaces into a home environment, creating things such as virtual frosted glass walls between remote rooms.

More speculatively, a future set of sensors and actuators might be able to reproduce the smell of someone's cooking by using a set of odorants typically found in food ingredients.

Finally, an interesting deviation from SCS could be to share context but *not* any cues created by the users. For example a screen could simulate a window that shows brightness, light colour, weather[3] or even the complete outside view of a remote partner's location. It would be interesting to see how such shared context facilitates interaction between time zones or climates, or whether it also leads to a sense or closeness even though nothing is transmitted that can be in any way influenced by the other person. Such a system could also contribute to a better understanding of SCS by investigating in what way the amount of information transmitted and displayed (i. e. only brightness or colour versus a full video transmission) contributes to cognitive load and how little information is sufficient to maintain any potential closeness benefits.

[3]Again, multimodality could be introduced by controlling the temperature or even humidity in the direct vicinity of the virtual window.

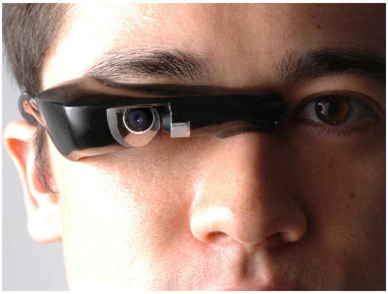

(a) EyeTap video see-through (Mann et al. [345])

(b) Shimadzu DataGlass 3/A (http://www.shimadzu.com/)

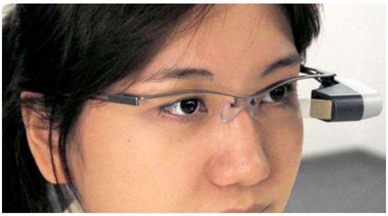

(c) Brother prototype (http://www.brother.com/)

(d) Lumus prototype (http://www.lumus-optical.com/)

(e) Google Glass (http://www.google.com/glass/)

Figure 8.1. Unobtrusive augmented reality devices. Most are either prototypes or saw a very small-scale release but were designed with a mass market in mind.

8.2.2 Umber Improvements

Evidently, the first goal for any further development of Umber needs to be to close the hybrid transformation loop and see what lessons can be learnt from this alone. The main stumbling blocks we expect are increasing the consistency of the generated rules and the subsequent orthotope differences, finding an end-user ready interface to manage different knowledge insertion and sampling strategies, and the translation of low-level features to comprehensible symbols.

Right now, Umber uses symbols that stem directly from the sensory features used during the machine learning stage. These features and their ranges are not necessarily well understood by human users. A better approach would be to detect features that are both meaningful when used in rules as well as semantic units.

One approach might be to detect clusters or similar structures in the data and present some contextual recordings to the user to label these structures. The idea behind this is that the clusters represent certain situations that are meaningful to the user and this could then be used as a complex combination or transformation of basic input features. Besides better comprehensibility, this also could well be combined with the benefit of much better approximation capability of the orthotopes represented by the rule set. When there is a transformation from feature space to rule space and back, orthotopes could – given the right kernel or similar transfer function – have arbitrarily good expressive power.

This in turn raises the question what the relation between a kernel that maximises comprehensibility and one that maximises expressive power or compactness of the rules is. There might be structures that yield a very small rule set but are nearly impossible to understand because the symbol used have no meaning to a human.

After solving these issues and thus then actually possessing a system as we envisioned it at the beginning of Chapter 7, investigating its effectiveness would naturally be the next step. In particular, we see the following research questions as the most interesting ones.

- Is a hybrid architecture actually superior to a purely rule-based one when it comes to comprehensibility and predictive power?

- Do users appreciate and use the white-boxing and manipulation capabilities provided by a hybrid transformation architecture?
- Do people feel better about the behaviour of the system because they could *in principle* manipulate it, even if they do not actually do it and even in cases where it does not actually achieve a better classification rate?
- How do people react to different rule types and kinds of presentation, especially to fuzzy rules and decision tables? How do they react to regression versus classification?

Bibliography

[1] Gregory Abowd, Anind K. Dey, Peter Brown, Nigel Davies, Mark Smith and Pete Steggles. 'Towards a Better Understanding of Context and Context-Awareness'. In: *Handheld and Ubiquitous Computing*. Edited by Hans-Werner Gellersen. Volume 1707. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 1999, pages 304–307. ISBN: 9783540665502. DOI: 10.1007/3-540-48157-5_29 (cited on page 225).

[2] Matthew Adcock, Drew Harry, Matthew Boch, Raul-David V. Poblano and Vanessa Harden. 'Tug n' Talk: A Belt Buckle for Tangible Tugging Communication'. In: alt.chi 2007. (San Jose, CA, USA, 28 Apr.–3 May 2007). ACM. URL: http://www.media.mit.edu/~dharry/tugandtalk/tugandtalk_FINAL.pdf (cited on pages 93, 95).

[3] Marcelo G. Almiron, Bruno Lopes, Alyson L. C. Oliveira, Antonio C. Medeiros and Alejandro C. Frery. 'On the Numerical Accuracy of Spreadsheets'. In: *Journal of Statistical Software* 34.4 (12 Apr. 2010), pages 1–29. ISSN: 1548-7660. URL: http://www.jstatsoft.org/v34/i04 (cited on page 149).

[4] Irwin Altman. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. Monterey, CA, USA: Brooks/-Cole Publishing Company, 1975. ISBN: 9780818501685 (cited on page 70).

[5] Irwin Altman. 'Privacy Regulation: Culturally Universal or Culturally Specific?' In: *Journal of Social Issues* 33.3 (1977), pages 66–84. ISSN: 1540-4560. DOI: 10.1111/j.1540-4560.1977.tb01883.x (cited on page 70).

[6] James L. Alty. ‘Can We Measure Usability’. In: *Proceedings of Advanced Information Systems*. 1992, pages 95–106 (cited on page 43).

[7] Robert Andrews, Joachim Diederich and Alan B. Tickle. ‘A Survey and Critique of Techniques for Extracting Rules from Trained Artificial Neural Networks’. In: *Knowledge-Based Systems* 8.6 (1995), pages 373–389 (cited on pages 264, 266, 270 sq.).

[8] Paul M. Aoki and Allison Woodruff. ‘Making Space for Stories: Ambiguity in the Design of Personal Communication Systems’. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’05. Portland, OR, USA: ACM, 2005, pages 181–190. ISBN: 1581139985. DOI: 10.1145/1054972.1054998 (cited on pages 25, 72).

[9] Michael Argyle. *Social interaction*. Methuen (London), 1969. ISBN: 0416139000 (cited on page 81).

[10] Michael Argyle and Janet Dean. ‘Eye-Contact, Distance and Affiliation’. In: *Sociometry* 28.3 (Sept. 1965), pages 289–304. DOI: 10.2307/2786027 (cited on page 31).

[11] Arthur P. Aron and Debra J. Mashek. ‘Conclusion’. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 23, pages 415–428. ISBN: 9780805842845 (cited on pages 15, 30 sq., 34).

[12] Arthur P. Aron, Debra J. Mashek and Elaine N. Aron. ‘Closeness as Including Other in the Self’. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 3, pages 27–41. ISBN: 9780805842845 (cited on pages 30 sq., 48).

[13] Arthur Aron, Elaine N. Aron and Danny Smollan. ‘Inclusion of Other in the Self Scale and the Structure of Interpersonal Closeness’. In: *Journal of Personality and Social Psychology* 63 (1992), pages 596–612 (cited on pages 34, 42, 48, 145).

[14] Ernesto Arroyo, Ted Selker and Alexandre Stouffs. 'Interruptions as Multimodal Outputs: Which are the Less Disruptive?' In: *Proceedings of the 4th International Conference on Multimodal Interfaces*. ICMI '02. Washington, DC, USA: IEEE, 2002, pages 479–482. ISBN: 0769518346. DOI: 10.1109/ICMI.2002.1167043 (cited on pages 82 sq.).

[15] Jens B. Asendorpf, Rainer Banse, Susanne Wilpers and Franz J. Neyer. 'Beziehungsspezifische Bindungsskalen für Erwachsene und ihre Validierung durch Netzwerk- und Tagebuchverfahren: Relationship-Specific Attachment Scales for Adults and Their Validation with Network and Diary Procedures'. German. In: *Diagnostica* 43.4 (1997), pages 289–313. ISSN: 0012-1924 (cited on pages 42, 156, 164).

[16] Ann-Sofie Axelsson, Åsa Abelin, Ilona Heldal, Alexander Nilsson, Ralph Schroeder and Josef Wideström. 'Collaboration and Communication in Multi-User Virtual Environments: A Comparison of Desktop and Immersive Virtual Reality Systems for Molecular Visualization'. In: *Proceedings of the 6th UKVRSIG Conference*. (Salford University, Salford, UK). 1999, pages 107–117 (cited on pages 36, 41, 202).

[17] Ann-Sofie Axelsson, Åsa Abelin, Ilona Heldal, Ralph Schroeder and Josef Wideström. 'Cubes in the Cube: A Comparison of a Puzzle-Solving Task in a Virtual and a Real Environment'. In: *CyberPsychology & Behavior* 4.2 (5 June 2001), pages 279–286. DOI: 10.1089/109493101300117956 (cited on pages 36, 41, 202).

[18] Bart Baesens, Rudy Setiono, Christophe Mues and Jan Vanthienen. 'Using Neural Network Rule Extraction and Decision Tables for Credit-Risk Evaluation'. In: *Management Science* 49.3 (2003), pages 312–329. DOI: 10.1287/mnsc.49.3.312.12739 (cited on page 258).

[19] Bart Baesens, Tony van Gestel, Stijn Viaene, Maria Stepanova, Johan A. K. Suykens and Jan Vanthienen. 'Benchmarking State-of-the-Art Classification Algorithms for Credit Scoring'. In: *The Journal of the Operational Research Society* 54.6 (2003), pages 627–635. ISSN: 01605682 (cited on page 273).

[20] Jeremy N. Bailenson, Jim Blascovich, Andrew C. Beall and Jack M. Loomis. 'Equilibrium Theory Revisited: Mutual Gaze and Personal Space in Virtual Environments'. In: *Presence: Teleoperators and Virtual Environments* 10.6 (Dec. 2001), pages 583–598. ISSN: 1054-7460. DOI: 10.1162/105474601753272844 (cited on page 41).

[21] Bill Bailey, Chris Burks et al. *Euler Diagram*. Wikipedia. 27 Aug. 2012. URL: http://en.wikipedia.org/w/index.php?title=Euler_diagram&oldid=509366490 (visited on 2012-11-19) (cited on page 34).

[22] Wolfgang Banzhaf, Peter Nordin, Robert E. Keller and Frank D. Francone. *Genetic Programming – An Introduction: On the Automatic Evolution of Computer Programs and Its Applications*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1998. ISBN: 155860510X (cited on page 274).

[23] Ling Bao and Stephen Intille. 'Activity Recognition from User-Annotated Acceleration Data'. In: *Pervasive Computing*. Edited by Alois Ferscha and Friedemann Mattern. Volume 3001. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2004, pages 1–17. ISBN: 9783540218357. DOI: 10.1007/978-3-540-24646-6_1 (cited on page 245).

[24] Nahla Barakat and Joachim Diederich. 'Eclectic Rule-Extraction from Support Vector Machines'. In: *International Journal of Computational Intelligence* 2.1 (2008), pages 59–62 (cited on page 266).

[25] Robert A. Baron and Michael J. Kalsher. *Psychology*. 5th edition. Allyn and Bacon, 2001. ISBN: 9780205314027 (cited on page 17).

[26] Kim Bartholomew and Leonard M. Horowitz. 'Attachment Styles Among Young Adults: A Test of a Four-Category Model'. In: *Journal of Personality and Social Psychology* 61 (1991), pages 226–244 (cited on pages 42, 48, 145, 156, 163 sq., 398, 404).

[27] Maurice Stevenson Bartlett. 'Properties of Sufficiency and Statistical Tests'. In: *Proceedings of the Royal Society of London: Series A, Mathematical and Physical Sciences* 160.901 (18 May 1937), pages 268–282. ISSN: 00804630 (cited on page 182).

[28] Cagatay Basdogan, Chih-Hao Ho, Mandayam A. Srinivasan and Mel Slater. 'An Experimental Study on the Role of Touch in Shared Virtual Environments'. In: *ACM Transactions on Compututer-Human Interaction* 7.4 (Dec. 2000), pages 443–460. ISSN: 1073-0516. DOI: 10.1145/365058.365082 (cited on pages 36, 41).

[29] Melissa Bateson, Daniel Nettle and Gilbert Roberts. 'Cues of being Watched Enhance Cooperation in a Real-World Setting'. In: *Biology Letters* 2.3 (2006), pages 412–414. ISSN: 1744-957X. DOI: 10.1098/rsbl.2006.0509. URL: http://rsbl.royalsocietypublishing.org/content/2/3/412.abstract (cited on page 69).

[30] Paolo Bellavista, Antonio Corradi, Mario Fanelli and Luca Foschini. 'A Survey of Context Data Distribution for Mobile Ubiquitous Systems'. In: *ACM Computing Surveys*. CSUR 44.4 (Sept. 2012), 24:1–24:45. ISSN: 0360-0300. DOI: 10.1145/2333112.2333119 (cited on page 225).

[31] Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957. ISBN: 9780691079516 (cited on page 279).

[32] Victoria Bellotti, Maribeth Back, W. Keith Edwards, Rebecca E. Grinter, Austin Henderson and Cristina Lopes. 'Making Sense of Sensing Systems: Five Questions for Designers and Researchers'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing our World, Changing Ourselves*. CHI '02. Minneapolis, MN, USA: ACM, 2002, pages 415–422. ISBN: 1581134533. DOI: 10.1145/503376.503450 (cited on page 52).

[33] Steve Benford and Lennart Fahlén. 'A Spatial Model of Interaction in Large Virtual Environments'. In: *Proceedings of the Third European Conference on Computer-Supported Cooperative Work*. ECSCW '93. Milan, Italy: Kluwer Academic Publishers, 1993, pages 109–124. ISBN: 0792324471 (cited on page 27).

[34] Ellen Berscheid, Mark Snyder and Allen M. Omoto. 'The Relationship Closeness Inventory: Assessing the closeness of interpersonal relationships'. In: *Journal of Personality and Social Psychology* 57 (1989), pages 792–807 (cited on pages 48, 145, 162 sq., 399, 405).

[35] Andrew L. Betz, John J. Skowronski and Thomas M. Ostrom. 'Shared Realities: Social Influence and Stimulus Memory'. In: *Social Cognition* 14.2 (1996), pages 113–140 (cited on page 73).

[36] Shruti Bhandari and Shaowen Bardzell. 'Bridging Gaps: Affective Communication in Long Distance Relationships'. In: *CHI '08 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '08. Florence, Italy: ACM, 2008, pages 2763–2768. ISBN: 9781605580128. DOI: 10.1145/1358628.1358758 (cited on pages 6, 86, 90).

[37] Derek Bickerton. *Roots of Language*. Karoma, 1981. ISBN: 9780897200448 (cited on page 80).

[38] Frank Biocca, Judee Burgoon, Chad Harms and Matt Stoner. 'Criteria and Scope Conditions for a Theory and Measure of Social Presence'. In: *Proceedings of the 4th Annual International Workshop on Presence*. Presence 2001. (Temple University, Philadelphia, PA, USA). 21–23 May 2001. URL: http://www.temple.edu/ispr/prev_conferences/proceedings/2001/Biocca1.pdf (cited on pages 15, 17, 19 sq., 22).

[39] Frank Biocca and Chad Harms. 'Defining and Measuring Social Presence: Contribution to the Networked Minds Theory and Measure'. In: *Proceeding of the 5th International Workshop on Presence*. Edited by Feliz Gouveia Gouveia and Frank Biocca. 2002, pages 7–36. URL: http://www.temple.edu/ispr/prev_conferences/proceedings/2002/final%20papers/biocca%20and%20harms.pdf (cited on pages 1, 6, 15, 40, 42, 145, 196, 200, 436).

[40] Frank Biocca, Chad Harms and Judee Burgoon. 'Towards a More Robust Theory and Measure of Social Presence: Review and Suggested Criteria'. In: *Presence: Teleoperators and Virtual Environments* 12 (2004), pages 456–480 (cited on pages 40, 200, 421 sq., 436).

[41] Frank Biocca, Chad Harms and J. Gregg. 'The Networked Minds Measure of Social Presence: Pilot Test of the Factor Structure and Concurrent Validity'. In: *Proceedings of 4th International Workshop on Presence*. (Philadelphia, USA). 21–23 May 2001 (cited on pages 40, 200, 206, 212, 436).

[42] Ray L. Birdwhistell. *Kinesics and Context – Essays on Body Motion Communication*. University of Pennsylvania Press Philadelphia, 1970. ISBN: 0812276051 (cited on page 81).

[43] Alexander Blicker-Dielmann. *Gemeinsam einsam fernsehen: Eine Untersuchung zum Einfluss sozialer Hinweisreize auf die Filmrezeption*. German. Diplomica Verlag, 2011. 156 pages. ISBN: 9783842853492 (cited on page 42).

[44] Sara A. Bly, Steve R. Harrison and Susan Irwin. 'Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment'. In: *Communications of the ACM* 36.1 (Jan. 1993), pages 28–46. ISSN: 0001-0782. DOI: 10.1145/151233.151235 (cited on page 3).

[45] Sara A. Bly, Steve R. Harrison and Susan Irwin. 'Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment'. In: *Commun. ACM* 36.1 (Jan. 1993), pages 28–46. ISSN: 0001-0782. DOI: 10.1145/151233.151235 (cited on page 127).

[46] Adam Bodnar, Richard Corbett and Dmitry Nekrasovski. 'AROMA: Ambient Awareness through Olfaction in a Messaging Application'. In: *Proceedings of the 6th international conference on Multimodal interfaces*. ICMI '04. State College, PA, USA: ACM, 2004, pages 183–190. ISBN: 1581139950. DOI: 10.1145/1027933.1027965 (cited on page 83).

[47] Bruce P. Bogert, Michael J. R. Healy and John Wilder Tukey. 'The Quefrency Alanysis of Time Series for Echoes: Cepstrum, Pseudo Autocovariance, Cross-Cepstrum and Saphe Cracking'. In: *Proceedings of the Symposium on Time Series Analysis*. Edited by Murray Rosenblatt. New York, NY, USA: John Wiley, 1963. Chapter 15, pages 209–243 (cited on page 246).

[48] Alan Borning and Michael Travers. 'Two Approaches to Casual Interaction over Computer and Video Networks'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Reaching through Technology*. CHI '91. New Orleans, LA, USA: ACM, 1991, pages 13–19. ISBN: 0897913833. DOI: 10.1145/108844.108847 (cited on page 70).

[49] Till Bovermann, René Tünnermann and Thomas Hermann. 'Auditory Augmentation'. In: *International Journal on Ambient Computing and Intelligence (IJACI)* 2.2 (2010), pages 27–41. ISSN: 1941-6237. DOI: 10.4018/jaci.2010040102 (cited on page 91).

[50] Till Bovermann, René Tünnermann and Christian Leichsenring. *Upstairs*. July 2011. URL: http://tai-studio.org/index.php/category/project-shared-presence-2/ (cited on page 6).

[51] Michael Boyle and Saul Greenberg. 'The Language of Privacy: Learning from Video Media Space Analysis and Design'. In: *ACM Transactions on Computer-Human Interaction*. TOCHI 12.2 (June 2005), pages 328–370. ISSN: 1073-0516. DOI: 10.1145/1067860.1067868 (cited on page 71).

[52] Erin Bradner. 'Social Affordances: Understanding Technology Mediated Social Networks at Work'. In: *Annual Meeting of the Association of Internet Researchers*. (Minneapolis, MN, USA). 2000 (cited on page 72).

[53] Erin Bradner. 'Social Affordances of Computer-Mediated Communication Technology: Understanding Adoption'. In: *CHI '01 Extended Abstracts on Human Factors in Computing Systems*. CHI '01. Seattle, WA, USA: ACM, 2001, pages 67–68. ISBN: 1581133405. DOI: 10.1145/634067.634111 (cited on pages 24, 73).

[54] Scott Brave and Andrew Dahley. 'inTouch: A Medium for Haptic Interpersonal Communication'. In: *Extended Abstracts on Human Factors in Computing Systems*. CHI '97. Atlanta, GA, USA: ACM, 1997, pages 363–364. ISBN: 0897919262. DOI: 10.1145/1120212.1120435 (cited on pages 92, 97).

[55] Eoin Brazil and Mikael Fernström. 'Auditory Icons'. In: *The Sonification Handbook*. Edited by Thomas Hermann, Andy Hunt and John G. Neuhoff. Berlin, Germany: Logos Publishing House, 2011. Chapter 13, pages 325–338. ISBN: 9783832528195. URL: http://sonification.de/handbook/chapters/chapter13/ (cited on pages 70, 91).

[56] Leo Breiman, J. H. Friedman, R. A. Olshen and C. J. Stone. *Classification and Regression Trees*. The Wadsworth and Brooks-Cole Statistics-Probability Series. Chapman & Hall, 1984. ISBN: 9780412048418 (cited on page 265).

[57] Johanna Brewer, Joseph Kaye, Amanda Williams and Susan Wyche, editors. *Sexual Interactions: Why We Should Talk About Sex in HCI*. Workshop Proceedings. CHI '06. Montréal, Québec, Canada: ACM, 2006. ISBN: 1595932984. DOI: 10.1145/1125451.1125765. URL: http://www.ics.uci.edu/~johannab/sexual.interactions.2006/chi2006.sex.PAPERS.htm (cited on page 96).

[58] Stephen Brewster and Lorna M. Brown. 'Tactons: Structured Tactile Messages for Non-Visual Information Display'. In: *AUIC '04: Proceedings of the Fifth Conference on Australasian User Interface*. Dunedin, New Zealand: Australian Computer Society, Inc., 2004, pages 15–23 (cited on page 96).

[59] Tomas Brezmes, Juan-Luis Gorricho and Josep Cotrina. 'Activity Recognition from Accelerometer Data on a Mobile Phone'. In: *Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living*. Edited by Sigeru Omatu, Miguel Rocha, José Bravo, Florentino Fernández, Emilio Corchado, Andrés Bustillo and Juan Corchado. Volume 5518. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2009, pages 796–799. ISBN: 9783642024801. DOI: 10.1007/978-3-642-02481-8_120 (cited on page 245).

[60] J. S. Bridle and M. D. Brown. 'An Experimental Automatic Word Recognition System'. In: *JSRU Report* 1003 (1974), page 5 (cited on page 246).

[61] Peter J. Brown, John D. Bovey and Xian Chen. 'Context-Aware Applications: From the Laboratory to the Marketplace'. In: *Personal Communications, IEEE* 4.5 (Oct. 1997), pages 58–64. ISSN: 1070-9916. DOI: 10.1109/98.626984 (cited on page 225).

[62] Barbara B. Bunker, Josephine M. Zubek, Virginia J. Vanderslice and Robert W. Rice. 'Quality of Life in Dual-Career Families: Commuting versus Single-Residence Couples'. In: *Journal of Marriage and Family* 54.2 (May 1992), pages 399–407. ISSN: 00222445 (cited on page 45).

[63] Judee Burgoon and J. L. Hale. 'Validation and Measurement of the Fundamental Themes of Relational Communication'. In: *Communication Monographs* 54 (1987), pages 19–41 (cited on page 41).

[64] Bill Buxton. 'The Making of Knowledge Navigator'. In: *Sketching User Experience: Getting the Design Right and the Right Design*. Interactive Technologies. Morgan Kaufmann. ISBN: 9780080552903. URL: http://www.dubberly.com/articles/the-making-of-knowledge-navigator.html (cited on page 219).

[65] Karl-Erik Bystrom and Woodrow Barfield. 'Effects of Participant Movement Affordance on Presence and Performance in Virtual Environments'. In: *Virtual Reality: The Journal of the Virtual Reality Society* 2.2 (1996), pages 206–216 (cited on pages 36, 41).

[66] Karl-Erik Bystrom and Woodrow Barfield. 'Collaborative Task Performance for Learning Using a Virtual Environment'. In: *Presence: Teleoperators and Virtual Environments* 8.4 (Aug. 1999), pages 435–448. ISSN: 1054-7460. DOI: 10.1162/105474699566323 (cited on pages 36, 41).

[67] Cambridge Business English Dictionary. 'Connectedness'. In: *Cambridge Business English Dictionary*. Cambridge University Press, 2011. ISBN: 9780521122504. URL: http://dictionary.cambridge.org/dictionary/business-english/connectedness?q=connectedness (cited on page 29).

[68] Rodney P. Carlisle. *Complete Idiot's Guide to Spies and Espionage*. Complete Idiot's Guide to ... Alpha Books, 2003, page 213. ISBN: 9780028644189 (cited on page 72).

[69] Juan S. Casanueva. 'Presence and Co-Presence in Collaborative Virtual Environments'. Master's thesis. Department of Computer Science, Faculty of Science, University of Cape Town, Apr. 2001. URL: http://people.cs.uct.ac.za/~edwin/MyBib/2000-casanueva-thesis.pdf (cited on pages 36, 41 sq., 142, 145, 202, 432).

[70] Raymond B. Cattell. 'The Scree Test For The Number Of Factors'. In: *Multivariate Behavioral Research* 1.2 (1966), pages 245–276. DOI: 10.1207/s15327906mbr0102_10 (cited on page 169).

[71] Matthew Chalmers and Areti Galani. 'Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems'. In: *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*. DIS. Cambridge, MA, USA: ACM, 2004, pages 243–252. ISBN: 1581137877. DOI: 10.1145/1013115.1013149 (cited on pages 57, 218).

[72] Angela Chang, Sile O'Modhrain, Rob Jacob, Eric Gunther and Hiroshi Ishii. 'Design of a Vibrotactile Communication Device: ComTouch'. In: *Proceedings of the 4th Conference on Designing Interactive Systems*. DIS '02. London, UK: ACM, 2002, pages 312–320. ISBN: 1581135157. DOI: 10.1145/778712.778755 (cited on pages 6, 93 sq.).

[73] Angela Chang, Ben Resner, Brad Koerner, XingChen Wang and Hiroshi Ishii. 'LumiTouch: An Emotional Communication Device'. In: *CHI '01 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '01. Seattle, WA, USA: ACM, 2001, pages 313–314. ISBN: 1581133405. DOI: 10.1145/634067.634252 (cited on pages 85, 90).

[74] Chih-Chung Chang and Chih-Jen Lin. *LibSVM: A Library for Support Vector Machines*. 2001. URL: http://www.csie.ntu.edu.tw/~cjlin/libsvm (cited on page 242).

[75] Wei-Chao Chen, Herman Towles, Lars Nyland, Greg Welch and Henry Fuchs. 'Toward a Compelling Sensation of Telepresence: Demonstrating a Portal to a Distant (Static) Office'. In: *Proceedings of the Conference on Visualization '00*. VIS. IEEE Computer Society Press, 2000, pages 327–333 (cited on page 1).

[76] Michael R. Chernick and Robert H. Friis. *Introductory Biostatistics for the Health Sciences: Modern Applications Including Bootstrap*. Wiley, 2003. ISBN: 9780471458654 (cited on page 156).

[77] Belkacem Chikhaoui, Shengrui Wang and Helene Pigot. 'A Frequent Pattern Mining Approach for ADLs Recognition in Smart Environments'. In: *Advanced Information Networking and Applications, International Conference on* (2011), pages 248–255. ISSN: 1550-445X. DOI: 10.1109/AINA.2011.13 (cited on pages 219, 227, 293).

[78] Yongwon Cho, Yunyoung Nam, Yoo-Joo Choi and We-Duke Cho. 'SmartBuckle: Human Activity Recognition Using a 3-Axis Accelerometer and a Wearable Camera'. In: *Proceedings of the 2nd International Workshop on Systems and Networking Support for Health Care and Assisted Living Environments*. HealthNet '08. Breckenridge, CO, USA: ACM, 2008, 7:1–7:3. ISBN: 9781605581996. DOI: 10.1145/1515747.1515757 (cited on page 245).

[79] Tanzeem Choudhury, Gaetano Borriello, Sunny Consolvo, Dirk Haehnel, Beverly Harrison, Bruce Hemingway, Jeffrey Hightower, Predrag Klasnja, Karl Koscher, Anthony LaMarca, James A. Landay, Louis LeGrand et al. 'The Mobile Sensing Platform: An Embedded Activity Recognition System'. In: *IEEE Pervasive Computing: Mobile and Ubiquitous Systems* 7.2 (Apr. 2008), pages 32–41. ISSN: 1536-1268. DOI: 10.1109/MPRV.2008.39 (cited on page 245).

[80] Morten H. Christiansen and Simon Kirby. *Language Evolution*. Oxford University Press, 2003. 414 pages. ISBN: 9780191581663 (cited on page 80).

[81] Nina Christiansen and Kelly Maglaughlin. 'Crossing from Physical Workspace to Virtual Workspace: Be AWARE!' In: *Proceedings of HCII*. International Conference on Human-Computer Interaction. New Jersey: Lawrence Erlbaum Associates, 2003, pages 1128–1132 (cited on page 24).

[82] David Cleevely and Richard Cawdell. 'A Telecommunications Taxonomy'. In: *Telecommunications Policy* 10.2 (1986), pages 107–119. ISSN: 0308-5961. DOI: 10.1016/0308-5961(86)90018-2 (cited on pages 52, 77).

[83] Ian Cloete and Jacek M. Zurada. *Knowledge-Based Neurocomputing*. MIT Press, 2000. ISBN: 9780262032742 (cited on page 258).

[84] Emily Cockayne. *Cheek by Jowl: A History of Neighbours*. Bodley Head, 5 Apr. 2012. ISBN: 9781847921345 (cited on page 128).

[85] Jack Cohen. *Statistical Power Analysis for the Behavioral Sciences*. 2nd edition. Lawrence Erlbaum Associates, 1988. 567 pages. ISBN: 9780805802832 (cited on pages 151, 162 sqq., 183).

[86] William W. Cohen. 'Learning Rules that Classify E-Mail'. In: *AAAI Spring Symposium on Machine Learning in Information Access*. Volume 18. 1996, page 25 (cited on page 220).

[87] Gary James Collier and Donna DiCarlo. *Emotional Expression*. Hillsdale, NJ, England: Lawrence Erlbaum Associates, 1985 (cited on page 99).

[88] Andrew L. Comrey and Howard B. Lee. *A First Course in Factor Analysis*. 2nd edition. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, 1992. 430 pages. ISBN: 0805810625 (cited on page 169).

[89] Sunny Consolvo, Peter Roessler and BrettE. Shelton. 'The CareNet Display: Lessons Learned from an In Home Evaluation of an Ambient Display'. In: *UbiComp 2004: Ubiquitous Computing*. Edited by Nigel Davies, ElizabethD. Mynatt and Itiro Siio. Volume 3205. Lecture Notes in Computer Science. Springer, 2004, pages 1–17. ISBN: 9783540229551. DOI: 10.1007/978-3-540-30119-6_1 (cited on page 85).

[90] Jeremy R. Cooperstock. 'Multimodal Telepresence Systems'. In: *Signal Processing Magazine* 28.1 (Jan. 2011), pages 77–86. ISSN: 1053-5888. DOI: 10.1109/MSP.2010.939040 (cited on page 77).

[91] Mark G. Core, H. Chad Lane, Michael van Lent, Dave Gomboc, Steve Solomon and Milton Rosenberg. 'Building Explainable Artificial Intelligence Systems'. In: *Proceedings of the 18th Conference on Innovative Applications of Artificial Intelligence*. Volume 2. IAAI'06. Boston, MA, USA: AAAI Press, 2006, pages 1766–1773. ISBN: 9781577352815 (cited on page 262).

[92] Mark W. Craven. 'Extracting Comprehensible Models from Trained Neural Networks'. PhD thesis. University of Wisconsin-Madison, 1996. ISBN: 0591144956 (cited on pages 273 sq.).

[93] Mark W. Craven and Jude Shavlik. *Rule Extraction: Where do We Go from Here*. Working Paper 99-1. University of Wisconsin Machine Learning Research Group, 1999 (cited on page 265).

[94] Mark W. Craven and Jude W. Shavlik. 'Using Sampling and Queries to Extract Rules from Trained Neural Networks'. In: *In Proceedings of the Eleventh International Conference on Machine Learning*. Morgan Kaufmann, 1994, pages 37–45 (cited on page 266).

[95] Mark W. Craven and Jude W. Shavlik. 'Extracting Tree-Structured Representations of Trained Networks'. In: *Proceedings of the 1995 Conference on Neural Information Processing Systems*. Edited by David S. Touretzky, Michael C. Mozer and Michael E. Hasselmo. Volume 8. Advances in Neural Information Processing Systems. MIT Press, 1996, pages 24–30. ISBN: 9780262201070 (cited on pages 273 sq.).

[96] Andrew Crossan, Grégoire Lefebvre, Sophie Zijp-Rouzier and Roderick Murray-Smith. 'A Multimodal Contact List to Enhance Remote Communication'. In: *Mobile Social Signal Processing*. Edited by Roderick Murray-Smith. Volume 8045. Lecture Notes in Computer Science. Springer, 2014. Chapter 9, pages 84–100. ISBN: 9783642543241. DOI: 10.1007/978-3-642-54325-8_9 (cited on pages 96, 102).

[97] James Crowley, Joëlle Coutaz, Gaeten Rey and Patrick Reignier. 'Perceptual Components for Context Aware Computing'. In: *UbiComp 2002: Ubiquitous Computing*. Edited by Gaetano Borriello and Lars Holmquist. Volume 2498. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2002, pages 117–134. ISBN: 9783540442677. DOI: 10.1007/3-540-45809-3_9 (cited on page 225).

[98] Lisa Takeuchi Cullen. 'Till Work Do Us Part'. In: *Time Magazine* (27 Sept. 2007). URL: http://www.time.com/time/magazine/article/0,9171,1666269,00.html (visited on 2012-08-22) (cited on page 46).

[99] Edward E. Cureton and Ralph B. D'Agostino. *Factor Analysis: An Applied Approach*. Lawrence Erlbaum Associates, 1993. ISBN: 9780805815467 (cited on page 182).

[100] Richard L. Daft and Robert H. Lengel. 'Organizational Information Requirements, Media Richness and Structural Design'. In: *Management Science* 32.5 (1986), pages 554–571. DOI: 10.1287/mnsc.32.5.554 (cited on page 64).

[101] Richard L. Daft, Robert H. Lengel and Linda Klebe Trevino. 'Message Equivocality, Media Selection and Manager Performance: Implications for Information Systems'. In: *MIS Quarterly* 11.3 (Sept. 1987), pages 355–366. ISSN: 0276-7783. DOI: 10.2307/248682 (cited on page 64).

[102] Andrew Dahley, Craig Wisneski and Hiroshi Ishii. 'Water Lamp and Pinwheels: Ambient Projection of Digital Information into Architectural Space'. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. (Los Angeles, CA, USA). CHI '98. New York, NY, USA: ACM, 1998, pages 269–270. ISBN: 1581130287. DOI: 10.1145/286498.286750 (cited on pages 84, 90).

[103] Marianne Dainton and Brooks Aylor. 'Patterns of Communication Channel Use in the Maintenance of Long-Distance Relationships'. In: *Communication Research Reports* 19.2 (2002), pages 118–129. DOI: 10.1080/08824090209384839 (cited on pages 47, 63).

[104] Data Mining Group. *PMML 4.1 – RuleSet*. 2011. URL: http://www.dmg.org/v4-1/RuleSet.html (visited on 2013-12-03) (cited on page 276).

[105] Edward S. de Guzman, Margaret Yau, Anthony Gagliano, Austin Park and Anind K. Dey. 'Exploring the Design and Use of Peripheral Displays of Awareness Information'. In: *CHI '04 Extended Abstracts on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: ACM, 2004, pages 1247–1250. ISBN: 1581137036. DOI: 10.1145/985921.986035 (cited on pages 59, 86, 90).

[106] Mathew Deepa. *vSmileys: Imaging Emotions through Vibration Patterns*. Tampere, Finland, 2005 (cited on page 96).

[107] Amnon Dekel and Dan Nacht. 'Minimizing Mobile Phone Disruption Via Smart Profile Management'. In: *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*. (Bonn, Germany). Edited by Reinhard Oppermann, Markus Eisenhauer, Matthias Jarke and Volker Wulf. Mobile HCI. ACM. ACM, Sept. 2009. ISBN: 9781605582818. URL: http://webhotel2.tut.fi/emmi/forum/sites/webhotel2.tut.fi.emmi.forum/files/2008/entries/smartprofile/Smart_Profile.pdf (cited on pages 102, 219, 227, 293).

[108] Krista E. DeLeeuw and Richard E. Mayer. 'A Comparison of Three Measures of Cognitive Load: Evidence for Separable Measures of Intrinsic, Extraneous, and Germane Load'. In: *Journal of Educational Psychology* 100.1 (2008), pages 223–234. ISSN: 0022-0663. DOI: 10.1037/0022-0663.100.1.223 (cited on page 43).

[109] Alan R. Dennis, Robert M. Fuller and Joseph S. Valacich. 'Media, Tasks, and Communication Processes: A Theory of Media Synchronicity'. In: *MIS Quarterly* 32.3 (2008-09), pages 575–600. ISSN: 0276-7783 (cited on page 64).

[110] Alan R. Dennis and Joseph S. Valacich. ‘Rethinking Media Richness: Towards a Theory of Media Synchronicity’. In: *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences*. Volume 1. HICSS ’99. Washington, DC, USA: IEEE Computer Society, 1999, pages 1–10. ISBN: 0769500013. DOI: 10.1109/HICSS.1999.772701 (cited on pages 11, 64).

[111] Alexandra Deschamps-Sonsino. *Good Night Lamp*. Designswarm Industries Limited. 2007–. URL: http://web.archive.org/web/20110301170603/http://goodnightlamp.com/ (visited on 2012-06-06) (cited on pages 7, 84, 87, 90, 210).

[112] Robert F. DeVellis. *Scale development: Theory and applications*. Volume 26. Applied social research methods series. Sage Publications, Inc., 1991. ISBN: 080393775X (cited on page 40).

[113] Anind K. Dey. ‘Understanding and Using Context’. In: *Personal and Ubiquitous Computing* 5 (1 2001). 10.1007/s007790170019, pages 4–7. ISSN: 1617-4909 (cited on page 225).

[114] Anind K. Dey and Edward S. de Guzman. ‘From Awareness to Connectedness: The Design and Deployment of Presence Displays’. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. (Montréal, Québec, Canada). CHI ’06. New York, NY, USA: ACM, 2006, pages 899–908. ISBN: 1595933727. DOI: 10.1145/1124772.1124905 (cited on pages 24, 85, 90, 226).

[115] Joachim Diederich, Alan B. Tickle and Shlomo Geva. ‘Quo Vadis? Reliable and Practical Rule Extraction from Neural Networks’. In: *Advances in Machine Learning I* 262 (2010), pages 479–490. DOI: 10.1007/978-3-642-05177-7_24 (cited on pages 264, 270).

[116] Angelika Dierker, Christian Mertes, Thomas Hermann, Marc Hanheide and Gerhard Sagerer. ‘Mediated Attention with Multimodal Augmented Reality’. In: *Proceedings of the 2009 International Conference on Multimodal Interfaces and the 2009 Workshop on Machine Learning for Multimodal Interaction*. (Cambridge, MA, USA). ICMI-MLMI ’09. New York, NY, USA: ACM, Nov. 2009, pages 245–252. ISBN: 9781605587721. DOI: 10.1145/1647314.1647368 (cited on page 91).

[117] Trinh Minh Tri Do, Jan Blom and Daniel Gatica-Perez. 'Smartphone Usage in the Wild: A Large-scale Analysis of Applications and Context'. In: *Proceedings of the 13th International Conference on Multimodal Interfaces*. ICMI '11. Alicante, Spain: ACM, 2011, pages 353–360. ISBN: 9781450306416. DOI: 10.1145/2070481.2070550 (cited on page 244).

[118] Jörg Doll, Michael Mentz and Erich H. Witte. 'Zur Theorie der vier Bindungsstile: Messprobleme und Korrelate dreier integrierter Verhaltenssysteme'. (On the Theory of Four Attachment Styles: Measurement Issues and Correlates of Three Integrated Behavioral Systems). In: *Zeitschrift für Sozialpsychologie* 26 (1995), pages 148–159 (cited on pages 42, 164, 398, 404).

[119] Florian Dombois and Gerhard Eckel. 'Audification'. In: *The Sonification Handbook*. Edited by Thomas Hermann, Andy Hunt and John G. Neuhoff. Berlin, Germany: Logos Publishing House, 2011. Chapter 12, pages 301–324. ISBN: 9783832528195. URL: http://sonification.de/handbook/chapters/chapter12/ (cited on page 114).

[120] Hans-Peter Dommel and Jose J. Garcia-Luna-Aceves. 'Network Support for Turn-Taking in Multimedia Collaboration'. In: (San Jose, CA, USA). Edited by Martin Freeman, Paul Jardetzky and Harrick M. Vin. Volume 3020. 1997, pages 304–315. DOI: 10.1117/12.264303 (cited on page 82).

[121] Frederick Dorey. 'Confidence Intervals: What is the Real Result in the Target Population?' In: *Clinical Orthopaedics and Related Research*. Statistics in Brief 468 (11 2010), pages 3137–3138. ISSN: 0009-921X. DOI: 10.1007/s11999-010-1407-4 (cited on page 153).

[122] Paul Dourish. 'What We Talk About When We Talk About Context'. In: *Personal Ubiquitous Computing* 8.1 (Feb. 2004), pages 19–30. ISSN: 1617-4909. DOI: 10.1007/s00779-003-0253-8 (cited on page 225).

[123] Paul Dourish. *Where the Action is: The Foundations of Embodied Interaction*. MIT Press, 2004 (cited on page 65).

[124] Paul Dourish and Victoria Bellotti. 'Awareness and Goordination in Shared Workspaces'. In: *Proceedings of the ACM conference on Computer-Supported Cooperative Work*. CSCW '92. Toronto, Ontario, Canada: ACM, 1992, pages 107–114. ISBN: 0897915429. DOI: 10.1145/143457.143468 (cited on page 24).

[125] Paul Dourish and Sara Bly. 'Portholes: Supporting Awareness in a Distributed Work Group'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '92. Monterey, CA, USA: ACM, 1992, pages 541–547. ISBN: 0897915135. DOI: 10.1145/142750.142982 (cited on pages 24 sq.).

[126] Didier Dubois and Henri Prade. 'What are Fuzzy Rules and How to Use Them'. In: *Fuzzy Sets and Systems* 84.2 (1996), pages 169–185. ISSN: 0165-0114. DOI: 10.1016/0165-0114(96)00066-8 (cited on page 270).

[127] Anthony Dunne and Fiona Raby. 'Fields and Thresholds'. In: *Architects in Cyberspace*. Edited by Martin Pearce and Neil Spiller. Volume 1. Architectural Design Profile 118. Academy Editions, 1995, pages 60–65 (cited on pages 95, 299).

[128] Berry Eggen and Koert van Mensvoort. 'Making Sense of What is Going on Around: Designing Environmental Awareness Information Displays'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 4, pages 99–124. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 84, 89).

[129] Jeroen Eggermont. *Genetic Programming*. Technical report. Leiden University Medical Center, 2007. URL: http://www.win.tue.nl/ipa/archive/falldays2007/HandoutEggermont.pdf (cited on page 274).

[130] Susan F. Ehrlich. 'Strategies for Encouraging Successful Adoption of Office Communication Systems'. In: *ACM Transactions on Information Systems*. TOIS 5.4 (Oct. 1987), pages 340–357. ISSN: 1046-8188. DOI: 10.1145/42196.42198 (cited on pages 64, 77).

[131] Elisabeth Eichhorn, Reto Wettach and Eva Hornecker. 'A Stroking Device for Spatially Separated Couples'. In: *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services*. MobileHCI '08. Amsterdam, The Netherlands: ACM, 2008, pages 303–306. ISBN: 9781595939524. DOI: 10.1145/1409240.1409274 (cited on pages 6, 93, 95, 99).

[132] Mica R. Endsley and Daniel J. Garland. *Situation Awareness: Analysis and Measurement*. London: Taylor & Francis, 2000. 408 pages. ISBN: 9780805821345 (cited on pages 25, 37).

[133] Mario J. Enriquez and Karon E. MacLean. 'The Hapticon Editor: A Tool in Support of Haptic Communication Research'. In: *International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems* (2003), page 356. DOI: 10.1109/HAPTIC.2003.1191310 (cited on page 96).

[134] M. C. Er. 'A Fast Algorithm for Generating Set Partitions'. In: *The Computer Journal* 31.3 (1988), pages 283–284. DOI: 10.1093/comjnl/31.3.283. URL: http://comjnl.oxfordjournals.org/content/31/3/283.abstract (cited on page 174).

[135] Thomas Erickson, Christine Halverson, Wendy A. Kellogg, Mark Laff and Tracee Wolf. 'Social Translucence: Designing Social Infrastructures That Make Collective Activity Visible'. In: *Communications of the ACM: Supporting Community and Building Social Capital* 45.4 (Apr. 2002), pages 40–44. ISSN: 0001-0782. DOI: 10.1145/505248.505270 (cited on pages 2, 38).

[136] Leandre R. Fabrigar, Duane T. Wegener, Robert C. MacCallum and Erin J. Strahan. 'Evaluating the Use of Exploratory Factor Analysis in Psychological Research'. In: *Psychological Methods* 4.3 (1999), pages 272–299. ISSN: 1082-989X. DOI: 10.1037/1082-989X.4.3.272 (cited on page 172).

[137] Morten W. Fagerland. '*t*-Tests, Non-Parametric Tests, and Large Studies: A Paradox of Statistical Practice?' In: *BMC Medical Research Methodology* 12.1 (2012), pages 1–7. DOI: 10.1186/1471-2288-12-78. URL: http://www.biomedcentral.com/1471-2288/12/78 (cited on pages 156, 178).

[138] Morten W. Fagerland and Leiv Sandvik. 'The Wilcoxon-Mann-Whitney Test under Scrutiny'. In: *Statistics in Medicine* 28.10 (2009), pages 1487–1497. ISSN: 1097-0258. DOI: 10.1002/sim.3561 (cited on page 178).

[139] Nicolas Fay, Michael Arbib and Simon C. Garrod. 'How to Bootstrap a Human Communication System'. In: *Cognitive Science* (2013). ISSN: 1551-6709. DOI: 10.1111/cogs.12048 (cited on page 80).

[140] Judgement of the Federal Constitutional Court of Germany (BVerG) from 15 Dec. 1983. Az. 1 BvR 209/83, 1 BvR 440/83, 1 BvR 420/83, 1 BvR 362/83, 1 BvR 269/83, 1 BvR 484/83 (Volkszählungsurteil). URL: http://openjur.de/u/268440.html (visited on 2012-10-11) (cited on page 69).

[141] Beverly Fehr. 'A Prototype Model of Intimacy IInteraction in Same-Sex Friendships'. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 2, pages 9–26. ISBN: 9780805842845 (cited on pages 30, 45).

[142] Andy Field. *Discovering Statistics Using SPSS*. 3rd edition. Introducing Statistical Methods Series. SAGE Publications, 17 Aug. 2007. 856 pages. ISBN: 9780857020963 (cited on page 172).

[143] Robert S. Fish, Robert E. Kraut and Barbara L. Chalfonte. 'The VideoWindow System in Informal Communication'. In: *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*. CSCW '90. Los Angeles, CA, USA: ACM, 1990, pages 1–11. ISBN: 0897914023. DOI: 10.1145/99332.99335 (cited on page 127).

[144] Ronald Aylmer Fisher. *UCI Machine Learning Repository: Iris Data Set*. Edited by Kevin Bache and Moshe Lichman. 1 July 1988. URL: http://archive.ics.uci.edu/ml/datasets/Iris (cited on page 259).

[145] Patrick Fletcher, Jon Mentor and Ari Brockman, editors. *Top 5 Smart Watches 2014*. Smart Watch News. 2014. URL: https://web.archive.org/web/20140108180513/http://www.smartwatchnews.org/top-5-smart-watches/ (cited on page 88).

[146] Friedrich Foerster and Jochen Fahrenberg. 'Motion Pattern and Posture: Correctly Assessed by Calibrated Accelerometers'. In: *Behavior Research Methods* 32 (3 2000), pages 450–457. ISSN: 1554-351X. DOI: 10.3758/BF03200815 (cited on page 245).

[147] James Anthony Fogarty. 'Constructing and Evaluating Sensor-Based Statistical Models of Human Interruptibility'. PhD thesis. Pittsburgh, PA, USA, 2006. ISBN: 9780542976346 (cited on page 293).

[148] James Fogarty and Scott E. Hudson. 'Toolkit Support for Developing and Deploying Sensor-Based Statistical Models of Human Situations'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '07. San Jose, CA, USA: ACM, 2007, pages 135–144. ISBN: 9781595935939. DOI: 10.1145/1240624.1240645 (cited on page 293).

[149] James Fogarty, Scott E. Hudson, Christopher G. Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny C. Lee and Jie Yang. 'Predicting Human Interruptibility with Sensors'. In: *Transactions on Computer-Human Interaction*. TOCHI 12 (1 Mar. 2005), pages 119–146. ISSN: 1073-0516. DOI: 10.1145/1057237.1057243 (cited on pages 219, 227, 293).

[150] Lawrence Jerome Fogel, Alvin J. Owens and Michael John Walsh. *Artificial Intelligence Through Simulated Evolution*. New York, NY, USA: John Wiley & Sons, 1966 (cited on page 274).

[151] B. J. Fogg, Lawrence D. Cutler, Perry Arnold and Chris Eisbach. 'HandJive: A Device for Interpersonal Haptic Entertainment'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '98. Los Angeles, CA, USA: ACM Press/Addison-Wesley, 1998, pages 57–64. ISBN: 0201309874. DOI: 10.1145/274644.274653 (cited on pages 95, 97).

[152] J. Kevin Ford, Robert C. MacCallum and Marianne Tait. 'The Application of Exploratory Factor Analysis in Applied Psychology: A Critical Review and Analysis'. In: *Personnel Psychology* 39.2 (1986), pages 291–314. ISSN: 0031-5826. DOI: 10.1111/j.1744-6570.1986.tb00583.x (cited on page 169).

[153] D. Michele Forina, Riccardo Leardi, C. Armanino and S. Lanteri. *UCI Machine Learning Repository: Wine Data Set*. Edited by Kevin Bache and Moshe Lichman. 1 July 1991. URL: http://archive.ics.uci.edu/ml/datasets/Wine (cited on page 276).

[154] Ben Foster. *How Many Users on Facebook?* 26 Oct. 2013. URL: https://web.archive.org/web/20131026074853/http://www.benphoster.com/facebook-user-growth-chart-2004-2010 (cited on pages i, 8 sq.).

[155] Paolo Frasconi, Marco Gori and Giovanni Soda. *Injecting Nondeterministic Finite State Automata into Recurrent Neural Networks*. Technical report. Florence, Italy: Università di Firenze, 1993 (cited on page 278).

[156] Paolo Frasconi, Marco Gori and Giovanni Soda. 'Recurrent Neural Networks and Prior Knowledge for Sequence Processing: A Constrained Nondeterministic Approach'. In: *Knowledge-Based Systems* 8.6 (1995), pages 313–332. ISSN: 0950-7051. DOI: 10.1016/0950-7051(96)81916-2 (cited on page 278).

[157] Alex A. Freitas. 'A Survey of Evolutionary Algorithms for Data Mining and Knowledge Discovery'. In: *Advances in Evolutionary Computing*. Edited by Ashish Ghosh and Shigeyoshi Tsutsui. Natural Computing Series. Springer, 2003, pages 819–845. ISBN: 9783642623868. DOI: 10.1007/978-3-642-18965-4_33 (cited on page 223).

[158] Batya Friedman, Peter H. Kahn Jr. and Alan Borning. *Value Sensitive Design: Theory and Methods*. Technical report. University of Washington, 1 Dec. 2002, pages 02–12 (cited on page 52).

[159] Jon Froehlich, Mike Y. Chen, Sunny Consolvo, Beverly Harrison and James A. Landay. 'MyExperience: A System for *In situ* Tracing and Capturing of User Feedback on Mobile Phones'. In: *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services*. MobiSys '07. San Juan, Puerto Rico: ACM, 2007, pages 57–70. ISBN: 9781595936141. DOI: 10.1145/1247660.1247670 (cited on page 244).

[160] Jim Fulton and Keith Packard. Version 1.0. *X11 Screen Saver Extension*. 1992. URL: http://www.x.org/releases/X11R7.7/doc/scrnsaverproto/saver.html (cited on page 235).

[161] Neal Fultz. *SPSS FAQ: What Does Cronbach's alpha Mean?* UCLA Statistical Consulting Group. 30 Sept. 2013. URL: http://www.ats.ucla.edu/stat/spss/faq/alpha.html (cited on page 182).

[162] Glenn Fung, Sathyakama Sandilya and R. Bharat Rao. 'Rule Extraction from Linear Support Vector Machines'. In: *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*. KDD '05. Chicago, IL, USA: ACM, 2005, pages 32–40. ISBN: 159593135X. DOI: 10.1145/1081870.1081878 (cited on page 269).

[163] Mohamed Medhat Gaber. *Scientific Data Mining and Knowledge Discovery – Principles and Foundations*. Heidelberg, Germany: Springer, 2010. ISBN: 9783642027871 (cited on page 264).

[164] Krzysztof Z. Gajos, Kartherine Everitt, Desney S. Tan, Mary Czerwinski and Daniel S. Weld. 'Predictability and Accuracy in Adaptive User Interfaces'. In: *Proceedings of the 2008 CHI*. Apr. 2008 (cited on page 224).

[165] Bruno Galantucci. 'An Experimental Study of the Emergence of Human Communication Systems'. In: *Cognitive Science: A Multidisciplinary Journal* 29.5 (2005), pages 737–767 (cited on page 80).

[166] Pritesh Gandhi, Lou Casavant, Ted Vann and Hiroshi Ishii. *Ambient Devices*. Ambient Devices. 2001–. URL: http://www.ambientdevices.com/ (cited on pages 84, 89).

[167] Maryanne Garry1 and Devon L. L. Polaschek. 'Imagination and Memory'. In: *Current Directions in Psychological Science* 9.1 (2000), pages 6–10 (cited on page 73).

[168] Bill Gaver. 'Provocative Awareness'. In: *Computer Supported Cooperative Work*. CSCW 11.3-4 (2002), pages 475–493. ISSN: 0925-9724. DOI: 10.1023/A:1021277326673 (cited on page 96).

[169] Bill Gaver and Heather Martin. 'Alternatives: Exploring Information Appliances Through Conceptual Design Proposals'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '00. The Hague, The Netherlands: ACM, 2000, pages 209–216. ISBN: 1581132166. DOI: 10.1145/332040.332433 (cited on page 51).

[170] William Gaver, Thomas Moran, Allan MacLean, Lennart Lövstrand, Paul Dourish, Kathleen Carter and William Buxton. 'Realizing a Video Environment: EuroPARC's RAVE System'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '92. Monterey, CA, USA: ACM, 1992, pages 27–35. ISBN: 0897915135. DOI: 10.1145/142750.142754 (cited on pages 68 sqq.).

[171] Jim Geier. *How Wi-Fi Roaming Really Works*. Wireless-Nets. 2013. URL: http://www.wireless-nets.com/resources/tutorials/how_roaming_works.html (cited on page 118).

[172] Hans-W. Gellersen, Albrecht Schmidt and Michael Beigl. 'Ambient Media for Peripheral Information Display'. In: *Personal Technologies* 3.4 (1999), pages 199–208. ISSN: 0949-2054. DOI: 10.1007/BF01540553 (cited on page 84).

[173] Naomi Gerstel and Harriet Gross. *Commuter Marriage: A Study of Work and Family*. Perspectives On Marriage And The Family. Guilford Press, 5 Oct. 1984. 228 pages. ISBN: 9780898620764 (cited on pages 45, 47).

[174] Gesellschaft für Akustikforschung Dresden mbH. *Walking Noise*. 2012. URL: http://www.akustikforschung.de/en/leistungen/bau-und-raumakustik/bauakustik/gehschall/ (cited on page 129).

[175] Nadia Ghamrawi and Andrew McCallum. 'Collective Multi-Label Classification'. In: *Proceedings of the 14th ACM International Conference on Information and Knowledge Management*. CIKM '05. Bremen, Germany: ACM, 2005, pages 195–200. ISBN: 1595931406. DOI: 10.1145/1099554.1099591 (cited on page 253).

[176] Martin R. Gibbs, Frank Vetere, Marcus Bunyan and Steve Howard. 'SynchroMate: A Phatic Technology for Mediating Intimacy'. In: *Proceedings of the 2005 Conference on Designing for User eXperience*. DUX '05. San Francisco, CA, USA: AIGA: American Institute of Graphic Arts, 2005. ISBN: 159593250X (cited on pages 88, 90).

[177] Nigel Gilbert, editor. *Dilemmas of Privacy and Surveillance: Challenges of Technological Change*. The Royal Academy of Engineering, Mar. 2007. ISBN: 1903496322. URL: http://www.raeng.org.uk/news/publications/list/reports/dilemmas_of_privacy_and_surveillance_report.pdf (cited on page 68).

[178] Charles Gillingham et al. *GOFAI*. Wikipedia. 17 May 2012. URL: http://en.wikipedia.org/w/index.php?title=GOFAI&oldid=493055185 (visited on 2013-10-11) (cited on page 256).

[179] Evelyn Glennie. *Evelyn Glennie Shows How to Listen*. TED. Feb. 2003. URL: http://www.ted.com/talks/evelyn_glennie_shows_how_to_listen.html (visited on 2012-06-06) (cited on page 104).

[180] Ken Go, John Carroll and Atsumi Imamiya. 'Familyware'. In: *Home Informatics and Telematics*. Edited by Andy Sloane and Felix Rijn. Volume 45. IFIP – The International Federation for Information Processing. Springer, 2000, pages 125–140. ISBN: 9781475754155. DOI: 10.1007/978-0-387-35511-5_9 (cited on pages 86, 89).

[181] Erving Goffman. *The Presentation of Self in Everyday Life*. Doubleday Anchor Books. Doubleday, 1959. ISBN: 9780385094023 (cited on page 59).

[182] Erving Goffman. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Volume 91194. The Free Press, 1963. URL: http://solomon.soth.alexanderstreet.com/cgi-bin/asp/philo/soth/documentidx.pl?sourceid=S10019969 (cited on pages i, 5, 19 sqq.).

[183] Susan Goldin-Meadow and Carolyn Mylander. 'Spontaneous Sign Systems Created by Deaf Children in Two Cultures'. In: *Nature* 391.6664 (15 Jan. 1998), pages 279–281 (cited on page 80).

[184] Elizabeth Goodman and Marion Misilim. 'Sensing Beds'. In: UbiComp 2003. (Seattle, WA, USA). 2003. URL: http://www.confectious.net/ITP/netobj/sensing/how.html (cited on pages 95, 299).

[185] Google Inc. *Android*. URL: http://www.android.com (cited on pages 105, 118, 244).

[186] Dmitry Gorilovsky, Olga Perova, Aleksey Dolgushin, Mike Sannikov, Sergey Zonovyev, Artem Zamlyanukhin, Andrew Gadjala and Andrey Korytsev. *TapTap: A Touch Communication Wristband*. Woodenshark. 16 Oct. 2013. URL: http://www.kickstarter.com/projects/woodenshark/taptap-wristband-a-new-way-to-say-i-love-you (cited on page 96).

[187] Dmitry Gorilovsky, Olga Perova, Aleksey Dolgushin, Mike Sannikov, Sergey Zonovyev, Artem Zamlyanukhin, Andrew Gadjala, Andrey Korytsev and Anton Nekhaenko. *TapTap.me*. Woodenshark. 17 Oct. 2013. URL: http://taptap.me/ (cited on page 96).

[188] Philip Babcock Gove et al., editors. *Merriam-Webster Online*. Merriam-Webster, Incorporated. 2012. URL: http://www.merriam-webster.com/dictionary/presence (visited on 2012-06-28) (cited on page 19).

[189] Philip Babcock Gove et al., editors. *Merriam-Webster Online*. Merriam-Webster, Incorporated. 2012. URL: http://www.merriam-webster.com/dictionary/privacy (visited on 2012-06-14) (cited on page 69).

[190] Saul Greenberg, Geraldine Fitzpatrick, Carl Gutwin and Simon Kaplan. 'Adapting the Locales Framework for Heuristic Evaluation of Groupware'. In: *Australasian Journal of Information Systems* 7.2 (2000), pages 102–108. ISSN: 1326-2238 (cited on pages 58 sq.).

[191] Saul Greenberg, Carman Neustaedter and Kathryn Elliot. 'Awareness in the Home: The Nuances of Relationships, Domestic Coordination and Communication'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 3, pages 73–96. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 6, 51).

[192] Saul Greenberg and Michael Rounding. 'The Notification Collage: Posting Information to Public and Personal Displays'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '01. Seattle, WA, USA: ACM, 2001, pages 514–521. ISBN: 1581133278. DOI: 10.1145/365024.365339 (cited on page 84).

[193] John M. Grey and John W. Gordon. 'Perceptual Effects of Spectral Modifications on Musical Timbres'. In: *The Journal of the Acoustical Society of America* 63 (5 1978), pages 1493–1500. ISSN: 0001-4966. DOI: 10.1121/1.381843 (cited on page 232).

[194] H. Paul Grice. 'Logic and Conversation'. In: *Syntax and Semantics*. Edited by Peter Cole and Jerry L. Morgan. Volume 3. New York, NY, USA: Academic Press, 1975 (cited on page 75).

[195] Charles G. Gross. 'Genealogy of the "Grandmother Cell"'. In: *The Neuroscientist* 8.5 (2002), pages 512–518. DOI: 10.1177/107385802237175 (cited on page 257).

[196] Elisheva F. Gross. 'Adolescent Internet Use: What We Expect, What Teens Report'. In: *Journal of Applied Developmental Psychology* 25.6 (2004), pages 633–649. ISSN: 0193-3973. DOI: 10.1016/j.appdev.2004.09.005 (cited on pages i, 8).

[197] Tom Gross, Chris Stary and Alex Totter. 'User-Centered Awareness in Computer-Supported Cooperative Work-Systems: Structured Embedding of Findings from Social Sciences'. In: *International Journal of Human-Computer Interaction* 18.3 (2005), pages 323–360. DOI: 10.1207/s15327590ijhc1803_5 (cited on pages 26, 54 sq.).

[198] Jonathan Grudin. 'Groupware and Social Dynamics: Eight Challenges for Developers'. In: *Communications of the ACM* 37.1 (Jan. 1994), pages 92–105. ISSN: 0001-0782. DOI: 10.1145/175222.175230 (cited on pages 55, 77).

[199] Alex Guazzelli, Michael Zeller, Wen-Ching Lin and Graham Williams. 'PMML: An Open Standard for Sharing Models'. In: *The R Journal* 1.1 (2009), pages 60–65. URL: http://journal.r-project.org/2009-1/RJournal_2009-1_Guazzelli+et+al.pdf (cited on page 276).

[200] Gregory T. Guldner. 'Long-Distance Romantic Relationships: Prevalence and Separation-Related Symptoms in College Students'. In: *Journal of College Student Development* 37 (1996), pages 103–106 (cited on pages 45, 47).

[201] Gregory T. Guldner and Clifford H. Swensen. 'Time Spent Together and Relationship Quality: Long-Distance Relationships as a Test Case'. In: 12.2 (1995), pages 313–320. DOI: 10.1177/0265407595122010 (cited on pages 44, 47).

[202] Anton Gustafsson and Magnus Gyllenswärd. 'The Power-Aware Cord: Energy Awareness Through Ambient Information Display'. In: *CHI '05 Extended Abstracts on Human Factors in Computing Systems*. CHI '05. Portland, OR, USA: ACM, 2005, pages 1423–1426. ISBN: 1595930027. DOI: 10.1145/1056808.1056932. URL: http://soda.swedish-ict.se/3932/1/CHI_2005.pdf (cited on pages 84, 89).

[203] Carl Gutwin and Saul Greenberg. 'A Descriptive Framework of Workspace Awareness for Real-Time Groupware'. In: *Computer Supported Cooperative Work (CSCW)* 11.3-4 (2002), pages 411–446. ISSN: 0925-9724. DOI: 10.1023/A:1021271517844 (cited on page 51).

[204] Rafi Haladjian and Olivier Mével. *Nabaztag/Karotz*. Apr. 2012–. URL: http://store.karotz.com/ (cited on pages 6, 88 sq.).

[205] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann and Ian H. Witten. 'The WEKA Data Mining Software: An Update'. In: *SIGKDD Explorations* 11.1 (2009), pages 10–18. URL: http://www.kdd.org/explorations/issues/11-1-2009-07/p2V11n1.pdf (cited on pages 242, 275).

[206] Lars Hallnäs and Johan Redström. 'Slow Technology – Designing for Reflection'. In: *Personal Ubiquitous Computing* 5.3 (Jan. 2001), pages 201–212. ISSN: 1617-4909. DOI: 10.1007/PL00000019 (cited on page 84).

[207] Sebastian Hammerl. 'Development of a Smartphone Based System to Support Personal Journaling of Daily Activities'. PhD thesis. Bielefeld University, 16 Apr. 2013 (cited on pages 247 sq.).

[208] Kory Hamzeh, Gurdeep Singh Pall, William Verthein, Jeff Taarud, W. Andrew Little and Glen Zorn. RFC 2637. *Point-to-Point Tunneling Protocol (PPTP)*. July 1999. URL: http://www.ietf.org/rfc/rfc2637.txt (cited on page 117).

[209] Harry H. Harman. *Modern Factor Analysis*. 3rd edition. University of Chicago Press, 1976. 487 pages. ISBN: 9780226316529 (cited on page 172).

[210] 'Origins and Evolution of Language and Speech'. In: *Annals of the New York Academy of Sciences* 280 (1976). Edited by Steven R. Harnad, Horst D. Steklis and Jane Lancaster. ISSN: 0077-8923 (cited on page 80).

[211] Guy Harris, David R. Tribble, Harris Macon et al. *Unix Time*. Wikipedia. 2 Sept. 2013. URL: https://en.wikipedia.org/w/index.php?title=Unix_time&oldid=571153236 (visited on 2013-09-04) (cited on page 232).

[212] Anna Hart and Jeremy Wyatt. 'Evaluating Black-Boxes as Medical Decision Aids: Issues Arising from a Study of Neural Networks'. In: *Informatics for Health and Social Care* 15.3 (1990), pages 229–236. DOI: 10.3109/14639239009025270 (cited on page 261).

[213] Satoshi Hashimoto, Takahiro Tanaka, Kazuaki Aoki and Kinya Fujita. 'Estimation of Interruptibility during Office Work Based on PC Activity and Conversation'. In: *Human Interface and the Management of Information: Information and Interaction for Learning, Culture, Collaboration and Business*. Edited by Sakae Yamamoto. Volume 8018. Lecture Notes in Computer Science. Springer, 2013, pages 297–306. ISBN: 9783642392252. DOI: 10.1007/978-3-642-39226-9_33 (cited on page 293).

[214] Trevor Hastie, Robert Tibshirani and Jerome Friedman. 'Local Methods in High Dimensions'. In: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd edition. Springer Series in Statistics. Springer, 26 Aug. 2009. Chapter 2, pages 22–26. ISBN: 9780387848587 (cited on page 279).

[215] W. Keith Hastings. 'Monte Carlo Sampling Methods Using Markov Chains and Their Applications'. In: *Biometrika* 57.1 (1970), pages 97–109. DOI: 10.1093/biomet/57.1.97 (cited on page 288).

[216] Laura Haverinen. 'Augmented Communication with Haptic I/O in Mobile Devices'. Master's thesis. University of Helsinki, Faculty of Science, Department of Computer Science, 4 Apr. 2012. URL: http://urn.fi/URN:NBN:fi-fe201204103207 (cited on page 102).

[217] Cindy Hazan and Phillip Shaver. 'Romantic Love Conceptualized as an Attachment Process'. In: *Journal of Personality and Social Psychology* 52.3 (1987), pages 511–524. ISSN: 1939-1315. DOI: 10.1037/0022-3514.52.3.511 (cited on page 42).

[218] Patrick G. T. Healey, Nik Swoboda, Ichiro Umata and James King. 'Graphical Language Games: Interactional Constraints on Representational Form'. In: *Cognitive Science* 31.2 (Mar.–Apr. 2007), pages 285–309. DOI: 10.1080/15326900701221363. URL: http://www.isrl.uiuc.edu/~amag/langev/paper/healey07graphicalLanguageGames.html (cited on page 80).

[219] Michael J. Healy. 'A Topological Semantics for Rule Extraction with Neural Networks'. In: *Connection Science* 11.1 (1999), pages 91–113. DOI: 10.1080/095400999116377 (cited on page 265).

[220] Mike Hearn and Christian Hammond. *Desktop Notification Specification*. Galago Project. Jan. 2006. URL: http://www.galago-project.org/specs/notification/ (cited on page 229).

[221] Jani Heikkinen, Thomas Olsson and Kaisa Väänänen-Vainio-Mattila. 'Expectations for User Experience in Haptic Communication with Mobile Devices'. In: *MobileHCI '09: Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Bonn, Germany: ACM, 2009, pages 1–10. ISBN: 9781605582818. DOI: 10.1145/1613858.1613895 (cited on pages 99, 105).

[222] Dimitri Heil. 'Klassifikation von Bewegungsmustern auf einem Smartphone'. German. Bachelor's thesis. Bielefeld University, 20 Dec. 2010. 35 pages. Unpublished (cited on page 246).

[223] Jeremy M. Heiner, Scott E. Hudson and Kenichiro Tanaka. 'The Information Percolator: Ambient Information Display in a Decorative Object'. In: *Proceedings of the 12th Annual ACM Symposium on User Interface Software and Technology*. UIST '99. Asheville, NC, USA: ACM, 1999, pages 141–148. ISBN: 1581130759. DOI: 10.1145/320719.322595 (cited on page 84).

[224] Fabian Hemmert. 'Ambient Life: Permanent Tactile Life-Like Actuation as a Status Display in Mobile Phones'. In: *Adjunct Proceedings of the 21st annual ACM symposium on User Interface Software and Technology*. (Monterey, CA, USA, 20–22 Oct. 2008). UIST AP '08. New York, NY, USA: ACM, 2008 (cited on page 94).

[225] Fabian Hemmert, Ulrike Gollner, Matthias Löwe, Anne Wohlauf and Gesche Joost. 'Intimate Mobiles: Near-Body Telepresence through Tightness, Wetness and Airflow in Mobile Phones'. In: *Erotic HCI'10: Proceedings of the First International Workshop on Erotic Life in HCI*. 2010 (cited on pages 7, 94, 96).

[226] Fabian Hemmert, Matthias Löwe, Anne Wohlauf and Gesche Joost. 'Animate Mobiles: Proxemically Reactive Posture Actuation As a Means of Relational Interaction with Mobile Phones'. In: *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*. TEI '13. Barcelona, Spain: ACM, 2013, pages 267–270. ISBN: 9781450318983. DOI: 10.1145/2460625.2460669 (cited on pages 93 sq.).

[227] Clyde Hendrick and Susan S. Hendrick. 'Lovers Wear Rose Colored Glasses'. In: *Journal of Social and Personal Relationships* 5.2 (1988), pages 161–183. DOI: 10.1177/026540758800500203 (cited on page 46).

[228] René Henn, Renate Hichert, Linda Sydow, Rainer Warnecke, Sabrina Werscheid and Sara Schlappa, editors. *Jahresbericht Bundesnetzagentur 2011*. Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen. 2011. URL: http://www.bundesnetzagentur.de/SharedDocs/Downloads/DE/BNetzA/Presse/Berichte/2012/Jahresbericht2011pdf.pdf?__blob=publicationFile (cited on page 8).

[229] Walter T. Herbranson and Julia Schroeder. 'Are Birds Smarter than Mathematicians? Pigeons Perform Optimally on a Version of the Monty Hall Dilemma'. In: *Journal of Comparative Psychology* 124.1 (2010), pages 1–13. ISSN: 1939-2087. DOI: 10.1037/a0017703 (cited on page 262).

[230] Thomas Hermann, Andy Hunt and John G. Neuhoff, editors. *The Sonification Handbook*. 1st edition. Berlin, Germany: Logos Publishing House, Nov. 2011. 586 pages. ISBN: 9783832528195. URL: http://sonification.de/handbook/ (cited on page 299).

[231] Thomas Hermann and Helge Ritter. 'Listen to Your Data: Model-Based Sonification for Data Analysis'. In: *Advances in Intelligent Computing and Multimedia Systems*. Edited by George E. Lasker, Mahbubur Rahman Syed and Orlando R. Baiocchi. Volume 76. Baden-Baden, Germany: International Institute for Advanced Studies in System Research and Cybernetics, pages 189–194. ISBN: 9780921836803. URL: http://www.techfak.uni-bielefeld.de/ags/ni/publications/media/HermannRitter1999-LTY.pdf (cited on page 100).

[232] Thomas Hermann and Helge Ritter. 'Listen to your Data: Model-Based Sonification for Data Analysis'. In: edited by George E. Lasker. Advances in Intelligent Computing and Multimedia Systems. International Institute for Advanced Studies in System Research and Cybernetics, 1999 (cited on page 299).

[233] Thomas Hermann and Helge Ritter. 'Model-Based Sonification Revisited: Authors' Comments on Hermann and Ritter, ICAD 2002'. In: *Transactions on Applied Perception*. TAP 2.4 (Oct. 2005), pages 559–563. ISSN: 1544-3558. DOI: 10.1145/1101530.1101557 (cited on page 299).

[234] Melanie Hilario. 'An Overview of Strategies for Neurosymbolic Integration'. In: *Proceedings of the Workshop on Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*. Lawrence Erlbaum Associates, 1997, pages 1–6 (cited on page 257).

[235] Debby Hindus, Scott D. Mainwaring, Nicole Leduc, Anna Elizabeth Hagström and Oliver Bayley. 'Casablanca: Designing Social Communication Devices for the Home'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. (Seattle, WA, USA). CHI '01. New York, NY, USA: ACM, 2001, pages 325–332. ISBN: 1581133278. DOI: 10.1145/365024.383749 (cited on pages 62, 78, 87, 210, 214).

[236] Joyce Ho and Stephen S. Intille. 'Using Context-Aware Computing to Reduce the Perceived Burden of Interruptions from Mobile Devices'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '05. Portland, OR, USA: ACM, 2005, pages 909–918. ISBN: 1581139985. DOI: 10.1145/1054972.1055100 (cited on page 293).

[237] Charles Hoberman. 'Radial Expansion/Retraction Truss Structures'. US Patent 5,024,031A. 18 June 1991 (cited on pages 86, 90).

[238] Thomas Hofer, Wieland Schwinger, Mario Pichler, Gerhard Leonhartsberger, Josef Altmann and Werner Retschitzegger. 'Context-Awareness on Mobile Devices – The Hydrogen Approach'. In: *Proceedings of the 36th Annual Hawaii International Conference on System Sciences*. HICSS '03. Washington, DC, USA: IEEE Computer Society, 2003, pages 292.1–. ISBN: 0769518745. DOI: 10.1109/HICSS.2003.1174831 (cited on page 244).

[239] Eve Hoggan, Craig Stewart, Laura Haverinen, Giulio Jacucci and Vuokko Lantz. 'Pressages: Augmenting Phone Calls with Non-Verbal Messages'. In: *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*. UIST '12. Cambridge, MA, USA: ACM, 2012, pages 555–562. ISBN: 9781450315807. DOI: 10.1145/2380116.2380185 (cited on page 96).

[240] John Henry Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Oxford, England: University of Michigan Press, 1975 (cited on page 274).

[241] John L. Horn. 'A Rationale and Test for the Number of Factors in Factor Analysis'. In: *Psychometrika* 30 (1965), pages 179–185 (cited on pages 169, 171).

[242] Bill G. Horne and Don R. Hush. 'On the Node Complexity of Neural Networks'. In: *Neural Networks* 7.9 (Nov. 1994), pages 1413–1426. ISSN: 0893-6080. DOI: 10.1016/0893-6080(94)90089-2 (cited on page 278).

[243] Harold Hotelling. 'Analysis of a Complex of Statistical Variables into Principal Components'. In: *Journal of Educational Psychology* 24.6 (1933), pages 417–441. ISSN: 0022-0663(Print). DOI: 10.1037/h0071325 (cited on page 169).

[244] Denis Howe. *Context*. Free On-Line Dictionary of Computing. 2010. URL: http://foldoc.org/context (visited on 2012-11-29) (cited on page 225).

[245] Chih-Wei Hsu, Chih-Chung Chang and Chih-Jen Lin. *A Practical Guide to Support Vector Classification*. Technical report. Taipei, Taiwan: National Taiwan University, 15 Apr. 2010 (cited on page 244).

[246] Scott E. Hudson and Ian Smith. 'Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems'. In: *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*. CSCW '96. Boston, MA, USA: ACM, 1996, pages 248–257. ISBN: 0897917650. DOI: 10.1145/240080.240295 (cited on pages 37, 65).

[247] Jina Huh, Mark W. Newman and Mark S. Ackerman. 'Supporting Collaborative Help for Individualized Use'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pages 3141–3150. ISBN: 9781450302289. DOI: 10.1145/1978942.1979408 (cited on page 258).

[248] Adam Hunt, Tom Morton et al. *Diaspora (Social Network)*. Wikipedia. 5 Apr. 2014. URL: https://en.wikipedia.org/w/index.php?title=Diaspora_(social_network)&oldid=602864499 (cited on page 59).

[249] Johan Huysmans, Bart Baesens and Jan Vanthienen. *Using Rule Extraction to Improve the Comprehensibility of Predictive Models*. Technical report. Katholieke Universiteit Leuven, Faculty of Business and Economics, 2006, pages 1–55. URL: http://www.econ.kuleuven.be/fetew/pdf_publicaties/KBI_0612.pdf (cited on pages 223, 256, 258, 264 sqq., 269 sqq.).

[250] Johan Huysmans, Karel Dejaeger, Christophe Mues, Jan Vanthienen and Bart Baesens. 'An Empirical Evaluation of the Comprehensibility of Decision Table, Tree and Rule Based Predictive Models'. In: *Decision Support Systems* 51.1 (2011), pages 141–154. ISSN: 0167-9236. DOI: 10.1016/j.dss.2010.12.003 (cited on pages 223, 265, 269).

[251] Giovanni Iachello. 'Privacy and Proportionality'. PhD thesis. Atlanta, GA, USA: Georgia Institute of Technology, 3 Apr. 2006. ISBN: 0542606631. URL: https://smartech.gatech.edu/handle/1853/10487 (cited on pages 57, 71).

[252] Giovanni Iachello and Gregory D. Abowd. 'Privacy and Proportionality: Adapting Legal Evaluation Techniques to Inform Design in Ubiquitous Computing'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI. Portland, OR, USA: ACM, 2005, pages 91–100. ISBN: 1581139985. DOI: 10.1145/1054972.1054986 (cited on pages 57, 71).

[253] IBM. *Delta Parameter in Oblimin Rotation in SPSS Factor*. Reference Number 1475052. IBM. 9 Mar. 2011. URL: http://www-01.ibm.com/support/docview.wss?uid=swg21475052 (visited on 2013-03-01) (cited on page 172).

[254] Peter Saint-Andre. RFC 6120. *Extensible Messaging and Presence Protocol (XMPP): Core*. Mar. 2011. URL: http://www.ietf.org/rfc/rfc6120.txt (cited on pages 119, 227).

[255] Yakov Shafranovich. RFC 4180. *Common Format and MIME Type for Comma-Separated Values (CSV) Files*. Oct. 2005. URL: http://www.ietf.org/rfc/rfc4180.txt (cited on page 229).

[256] Wijnand A. IJsselsteijn, Huib de Ridder, Jonathan Freeman and Steve E. Avons. 'Presence: Concept, Determinants and Measurement'. In: *Human Vision and Electronic Imaging V*. (San Jose CA, USA, 24–27 Jan. 2000). Volume 3959. Proceedings of SPIE. Bellingham, WA, USA: Society of Photo-Optical Instrumentation Engineers, 2 June 2000, pages 520–529. ISBN: 9780819435774. DOI: 10.1117/12.387188 (cited on page 42).

[257] Wijnand A. IJsselsteijn, Joy K. van Baren, Panos Markopoulos, Natalia Romero and Boris E. R. de Ruyter. 'Measuring Affective Benefits and Costs of Mediated Awareness: Development and Validation of the ABC-Questionnaire'. In: *Awareness Systems*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. Human-Computer Interaction Series. Springer, 2009. Chapter 20, pages 473–488. ISBN: 9781848824775 (cited on pages 24, 29, 42 sq., 425, 436).

[258] Wijnand A. IJsselsteijn, Joy K. van Baren and Froukje van Lanen. 'Staying in Touch: Social Presence and Connectedness through Synchronous and Asynchronous Communication Media'. In: *Proceedings of the International Conference on Human-Computer Interaction*. HCII 2003. New Jersey, USA: Lawrence Erlbaum Associates, 2003, pages 924–928 (cited on pages 6, 29 sq., 36, 43, 63).

[259] Ikip, Craig Burrows, Briaboru et al. *Plausible deniability*. Wikipedia. 29 May 2012. URL: http://en.wikipedia.org/w/index.php?title=Plausible_deniability&oldid=494880880#Other_examples (visited on 2012-06-06) (cited on page 72).

[260] International Organization for Standardization. ISO 9241-210:2010. *Ergonomics of Human-System Interaction – Part 210: Human-Centred Design Processes for Interactive Systems* (International Standard). 3 Mar. 2010. URL: http://www.iso.org/iso/catalogue_detail.htm?csnumber=52075 (cited on page 292).

[261] International Society for Presence Research. *Measures Statement and Compendium*. May 2000. URL: http://ispr.info/about-presence-2/tools-to-measure-presence/ispr-measures-compendium/ (visited on 2012-08-09) (cited on page 39).

[262] International Society for Presence Research. *The Concept of Presence: Explication Statement*. 29 Apr. 2000. URL: http://ispr.info/ (visited on 2012-06-28) (cited on pages 18 sq.).

[263] Ellen Isaacs, Alan Walendowski and Dipti Ranganthan. 'Hubbub: A Sound-Enhanced Mobile Instant Messenger That Supports Awareness and Opportunistic Interactions'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '02. Minneapolis, MN, USA: ACM, 2002, pages 179–186. ISBN: 1581134533. DOI: 10.1145/503376.503409 (cited on pages 91, 299).

[264] Hiroshi Ishii and Brygg Ullmer. 'Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms'. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. (Atlanta, GA, USA). CHI '97. New York, NY, USA: ACM, 1997, pages 234–241. ISBN: 0897918029. DOI: 10.1145/258549.258715 (cited on page 84).

[265] Hiroshi Ishii, Craig Wisneski, Scott Brave, Andrew Dahley, Matt Gorbet, Brygg Ullmer and Paul Yarin. 'ambientROOM: Integrating Ambient Media with Architectural Space'. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. (Los Angeles, CA, USA). CHI '98. New York, NY, USA: ACM, 1998, pages 173–174. ISBN: 1581130287. DOI: 10.1145/286498.286652 (cited on page 84).

[266] Pico Iyer. *on Connection*. Edited by Moses Znaimer. Ideacity. 2012. URL: http://www.ideacityonline.com/talks/pico-iyer-talk-at-ideacity-2012/ (visited on 2013-02-16) (cited on page 58).

[267] E. W. Jacobs and Mary W. Hicks. 'Periodic Family Separation: The Importance of Beliefs in Determining Outcomes'. In: *Military Family* 7.2 (1987), pages 3–5 (cited on page 47).

[268] Nico Jacobs. 'Relational Sequence Learning and User Modelling'. PhD thesis. Katholieke Universiteit Leuven, 15 Oct. 2004. URL: https://lirias.kuleuven.be/bitstream/123456789/245905/1/Jacobs.pdf (cited on page 252).

[269] Anthony Jameson. 'Adaptive Interfaces and Agents'. In: *Human-Computer Interaction: Design Issues, Solutions, and Applications* (2009), page 105 (cited on page 52).

[270] Gavin Jancke, Gina Danielle Venolia, Jonathan Grudin, J. J. Cadiz and Anoop Gupta. 'Linking Public Spaces: Technical and Social Issues'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '01. Seattle, WA, USA: ACM, 2001, pages 530–537. ISBN: 1581133278. DOI: 10.1145/365024.365352 (cited on page 127).

[271] Soames R. F. Job. 'Community Response to Noise: A Review of Factors Influencing the Relationship between Noise Exposure and Reaction'. In: *Journal of the Acoustical Society of America* 83.3 (1988), pages 991–1001. ISSN: 0001-4966. DOI: 10.1121/1.396524 (cited on page 128).

[272] Ulf Johansson, Rikard König and Lars Niklasson. 'Rule Extraction from Trained Neural Networks Using Genetic Programming'. In: *13th International Conference on Artificial Neural Networks*. 2003, pages 13–16 (cited on pages 273 sq.).

[273] Ulf Johansson, Rikard König and Lars Niklasson. 'Automatically Balancing Accuracy and Comprehensibility in Predictive Modeling'. In: *8th International Conference on Information Fusion*. Volume 2. IEEE. 2005, pages 1554–1560. DOI: 10.1109/ICIF.2005.1592040 (cited on pages 271, 275).

[274] Kenneth O. Johnson and Steven S. Hsiao. 'Neural Mechanisms of Tactual Form and Texture Perception.' In: *Annual Review of Neuroscience* 15.1 (1992), pages 227–250. ISSN: 0147-006X (cited on page 101).

[275] Timo Kaerlein. 'Presence in a Pocket: Phantasms of Immediacy in Japanese Mobile Telepresence Robotics'. In: *communication +1* 1.1 (2012). Article 6 (cited on pages 95, 97).

[276] Kai-Hendrik, Fgb et al. *Überwachungsdruck*. German. Wikipedia. 7 July 2011. URL: http://de.wikipedia.org/w/index.php?title=%C3%83%C2%9Cberwachungsdruck&oldid=90972321 (visited on 2012-10-20) (cited on page 69).

[277] Anssi Kainulainen, Markku Turunen and Jaakko Hakulinen. 'Awareness Information with Speech and Sound'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 10, pages 231–256. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on page 299).

[278] Henry F. Kaiser. 'The Varimax Criterion for Analytic Rotation in Factor Analysis'. In: *Psychometrika* 23 (3 1958), pages 187–200. ISSN: 0033-3123. DOI: 10.1007/BF02289233 (cited on page 172).

[279] Henry F. Kaiser. 'A Second Generation Little Jiffy'. In: *Psychometrika* 35 (4 1970), pages 401–415. ISSN: 0033-3123 (cited on pages 169, 182).

[280] Joseph Kaye. 'Making Scents: Aromatic Output for HCI'. In: *Interactions* 11.1 (Jan. 2004), pages 48–61. ISSN: 1072-5520. DOI: 10.1145/962342.964333 (cited on page 83).

[281] Joseph Kaye. 'I Just Clicked to Say I Love You: Rich Evaluations of Minimal Communication'. In: *CHI'06 Extended Abstracts on Human Factors in Computing Systems*. ACM. 2006, pages 363–368 (cited on pages 6, 91).

[282] Joseph Kaye and Liz Goulding. 'Intimate Objects'. In: *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*. DIS '04. Cambridge, MA, USA: ACM, 2004, pages 341–344. ISBN: 1581137877. DOI: 10.1145/1013115.1013175 (cited on page 91).

[283] Judy Kegl. 'The Nicaraguan Sign Language Project: An Overview'. In: *Signpost* 7.1 (1994), pages 24–31 (cited on page 80).

[284] Maurice G. Kendall. 'A New Measure of Rank Correlation'. In: *Biometrika* 30 (1938), pages 81–93 (cited on page 442).

[285] Douglas T. Kenrick, Steven L. Neuberg and Robert B. Cialdini. *Social Psychology: Unraveling the Mystery*. Allyn & Bacon, 2004. 648 pages. ISBN: 9780205165216 (cited on page 55).

[286] Sharon Kinsella. 'Cuties in Japan'. In: *Women, Media, and Consumption in Japan*. Edited by Lise Skov and Brian Moeran. ConsumAsiaN Book Series. University of Hawai'i Press, 1995. Chapter 6, pages 220–254. ISBN: 9780824817763 (cited on page 53).

[287] Niko Kiukkonen, Jan Blom, Olivier Dousse, Daniel Gatica-Perez and Juha K. Laurila. 'Towards Rich Mobile Phone Datasets: Lausanne Data Collection Campaign'. In: *Proceedings of the International Conference on Pervasive Services*. (Berlin, Germany). ICPS. July 2010. URL: http://publications.idiap.ch/downloads/papers/2010/Kiukkonen_ICPS,BERLIN,2010_2010.pdf (cited on page 244).

[288] Jesper Kjeldskov, Martin R. Gibbs, Franks Vetere, Steve Howard, Sonja Pedell, Karen Mecoles and Marcus Bunyan. 'Using Cultural Probes to Explore Mediated Intimacy'. In: *Australasian Journal of Information Systems* 11.2 (Dec. 2004), pages 102–115. ISSN: 1449-8618. URL: http://dl.acs.org.au/index.php/ajis/article/view/128/107 (cited on page 30).

[289] Paul Kline. *The Handbook of Psychological Testing*. 2nd edition. Routledge, 1999. 752 pages. ISBN: 9781317798040 (cited on page 155).

[290] Mark L. Knapp and Judith A. Hall. *Nonverbal Communication in Human Interaction*. 5th edition. Wadsworth/Thomson Learning, 2001. ISBN: 9780155063723 (cited on pages 81, 99).

[291] Michael Knowles, Ted P. Pavlic, Dick Lyon et al. *High-Pass Filter*. Wikipedia. 20 Mar. 2013. URL: http://en.wikipedia.org/w/index.php?title=High-pass_filter&oldid=545652081#Discrete-time_realization (visited on 2013-04-23) (cited on page 106).

[292] David Knox, Marty E. Zusman, Vivian Daniels and Angel Brantley. 'Absence Makes the Heart Grow Fonder? Long Distance Dating Relationships among College Students'. In: *College Student Journal* 36.3 (1 Sept. 2002), pages 364–367. ISSN: 0146-3934 (cited on page 45).

[293] Felix Köbler, Christoph Riedl, Céline Vetter, Jan Marco Leimeister and Helmut Krcmar. 'Social Connectedness on Facebook: An Explorative Study on Status Message Usage'. In: *Proceedings of the 16th Americas Conference on Information Systems*. (Lima, Peru, 12–15 Aug. 2010). AMCIS. 2010 (cited on page 63).

[294] Eyal Kolman and Michael Margaliot. *Knowledge-Based Neurocomputing: A Fuzzy Logic Approach*. Volume 234. Studies in Fuzziness and Soft Computing. Berlin, Germany: Springer, 2009. ISBN: 9783540880769 (cited on page 278).

[295] Rikard König. 'Predictive Techniques and Methods for Decision Support in Situations with Poor Data Quality'. Licentiate thesis. University of Borås, 8 May 2009. URL: http://hdl.handle.net/2320/5134 (cited on pages 273 sq.).

[296] Rikard König, Ulf Johansson, Tuve Lofstrom and Lars Niklasson. 'Improving GP Classification Performance by Injection of Decision Trees'. In: *IEEE Congress on Evolutionary Computation*. CEC. 2010, pages 1–8. DOI: 10.1109/CEC.2010.5585988 (cited on page 223).

[297] Rikard König, Ulf Johansson and Lars Niklasson. 'G-REX: A Versatile Framework for Evolutionary Data Mining'. In: *Data Mining Workshops, 2008, IEEE International Conference on*. IEEE. 2008, pages 971–974 (cited on pages 223, 274).

[298] Rikard König, Ulf Johansson and Lars Niklasson. 'Finding the Tree in the Forest'. In: *2010 IADIS International Conference Applied Computing*. Volume 135–142. IADIS Press, 2010 (cited on page 273).

[299] Rikard König, Ulf Johanssson and Lars Niklasson. 'Using Genetic Programming to Increase Rule Quality'. In: *Proceedings of the Twenty-First International FLAIRS Conference*. 2008, pages 288–293 (cited on page 275).

[300] Dimitrios Kontaris, Daniel Harrison, Evgenia-Eleni Patsoule, Susan Zhuang and Annabel Slade. 'Feelybean: Communicating Touch over Distance'. In: *CHI '12 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '12. Austin, TX, USA: ACM, 2012, pages 1273–1278. ISBN: 9781450310161. DOI: 10.1145/2212776.2212439 (cited on pages 93, 96).

[301] Asher Koriat, Morris Goldsmith and Ainat Pansky. 'Toward a Psychology of Memory Accuracy'. In: *Annual Review of Psychology* 51.1 (2000), pages 481–537. DOI: 10.1146/annurev.psych.51.1.481 (cited on page 73).

[302] Panu Korpipää, Jani Mäntyjärvi, Juha Kela, Heikki Keränen and Esko-Juhani Malm. 'Managing Context Information in Mobile Devices'. In: *IEEE Pervasive Computing* 2.3 (July–Sept. 2003), pages 42–51. ISSN: 1536-1268. DOI: 10.1109/MPRV.2003.1228526 (cited on page 244).

[303] Robert Kowalski, Sebastian Loehmann and Doris Hausen. 'Cubble: A Multi-Device Hybrid Approach Supporting Communication in Long-Distance Relationships'. In: *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction*. TEI '13. Barcelona, Spain: ACM, 2013, pages 201–204. ISBN: 9781450318983. DOI: 10.1145/2460625.2460656 (cited on pages 91, 97).

[304] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Complex Adaptive Systems. MIT Press, 1992. 819 pages. ISBN: 9780262111706 (cited on page 274).

[305] *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Sante Fe Institute Studies in the Sciences of Complexity. Perseus Books, 1994. ISBN: 9780201626049. 672 (cited on page 70).

[306] Robert Kraut, Carmen Egido and Jolene Galegher. 'Patterns of Contact and Communication in Scientific Research Collaboration'. In: *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*. CSCW '88. Portland, OR, USA: ACM, 1988, pages 1–12. ISBN: 0897912829. DOI: 10.1145/62266.62267 (cited on pages 36, 64).

[307] Joe Kress, Karl Palmen et al. *Calendar*. – Calendar Systems. Wikipedia. 3 Sept. 2013. URL: http://en.wikipedia.org/w/index.php?title=Calendar&oldid=571420017#Calendar_systems (visited on 2013-09-04) (cited on page 232).

[308] Narayanan C. Krishnan and Sethuraman Panchanathan. 'Analysis of Low Resolution Accelerometer Data for Continuous Human Activity Recognition'. In: *International Conference on Acoustics, Speech and Signal Processing*. (31 Mar.–4 Apr. 2008). ICASSP. IEEE. 2008, pages 3337–3340. DOI: 10.1109/ICASSP.2008.4518365 (cited on page 245).

[309] R. Krishnan, G. Sivakumar and P. Bhattacharya. 'A Search Technique for Rule Extraction from Trained Neural Networks'. In: *Pattern Recognition Letters* 20.3 (1999), pages 273–280. ISSN: 0167-8655. DOI: 10.1016/S0167-8655(98)00145-7 (cited on page 265).

[310] Martin Krzywinski, Jacqueline Schein, İnanç Birol, Joseph Connors, Randy Gascoyne, Doug Horsman, Steven J. Jones and Marco A. Marra. 'Circos: An Information Aesthetic for Comparative Genomics'. In: *Genome Research* 19.9 (Sept. 2009), pages 1639–1645. URL: http://genome.cshlp.org/content/19/9/1639.short (cited on pages 197, 216).

[311] Frederik Kunz, Torsten Stollen, Sigbert et al. *Cronbachs Alpha*. German. Wikipedia. 21 July 2012. URL: http://de.wikipedia.org/w/index.php?title=Cronbachs_Alpha&oldid=105846969 (visited on 2012-08-24) (cited on page 40).

[312] Kazuhiro Kuwabara, Takumi Watanabe, Takeshi Ohguro, Yoshihiro Itoh and Yuji Maeda. 'Connectedness Oriented Communication: Fostering a Sense of Connectedness to Augment Social Relationships'. In: *Proceedings of the Symposium on Applications and the Internet*. SAINT 2002. 2002, pages 186–193. DOI: 10.1109/SAINT.2002.994476 (cited on pages 5, 86, 219, 227).

[313] Jennifer R. Kwapisz, Gary M. Weiss and Samuel A. Moore. 'Activity Recognition Using Cell Phone Accelerometers'. In: *SIGKDD Explorations Newsletter* 12.2 (Mar. 2011), pages 74–82. ISSN: 1931-0145. DOI: 10.1145/1964897.1964918 (cited on pages 245 sq.).

[314] Ares Lagae and Philip Dutré. 'A Comparison of Methods for Generating Poisson Disk Distributions'. In: *Computer Graphics Forum* 27.1 (2008), pages 114–129. ISSN: 1467-8659. DOI: 10.1111/j.1467-8659.2007.01100.x (cited on page 280).

[315] William B. Langdon. *Genetic Programming and Data Structures: Genetic Programming + Data Structures = Automatic Programming!* Genetic Programming Series. Kluwer Academic Publishers, 1998. 278 pages. ISBN: 9780792381358 (cited on page 274).

[316] Stephen R. H. Langton, Roger J. Watt and Vicki Bruce. 'Do the Eyes Have It? Cues to the Direction of Social Attention'. In: *Trends in Cognitive Sciences* 4.2 (2000), pages 50–59. ISSN: 1364-6613. DOI: 10.1016/S1364-6613(99)01436-9 (cited on page 82).

[317] Juha K. Laurila, Daniel Gatica-Perez, Imad Aad, Jan Blom, Olivier Bornet, Trinh-Minh-Tri Do, Olivier Dousse, Julien Eberle and Markus Miettinen. 'The Mobile Data Challenge: Big Data for Mobile Computing Research'. In: *Mobile Data Challenge by Nokia Workshop*. (Newcastle, UK). 2012 (cited on page 244).

[318] Benjamin Le, Timothy J. Loving, Gary W. Lewandowski, Emily G. Feinberg, Katherine C. Johnson, Remy Fiorentino and Jennifer Ing. 'Missing a romantic partner: A prototype analysis'. In: *Personal Relationships* 15.4 (3 Dec. 2008), pages 511–532. ISSN: 1475-6811. DOI: 10.1111/j.1475-6811.2008.00213.x (cited on page 45).

[319] Scott Lederer, Jason I. Hong, Anind K. Dey and James A. Landay. 'Personal Privacy Through Understanding and Action: Five Pitfalls for Designers'. In: *Personal Ubiquitous Computing* 8.6 (Nov. 2004), pages 440–454. ISSN: 1617-4909. DOI: 10.1007/s00779-004-0304-9 (cited on page 72).

[320] Mi-hee Lee, Jungchae Kim, Kwangsoo Kim, Inho Lee, Sun Ha Jee and Sun Kook Yoo. 'Physical Activity Recognition Using a Single Tri-Axis Accelerometer'. In: *Proceedings of The World Congress on Engineering and Computer Science 2009*. International Conference on Computational Biology 2009. (San Francisco, CA, USA, 20–22 Oct. 2009). Edited by S. I. Ao, Craig Douglas, W. S. Grundfest and Jon Burgstone. Volume 1. 2 volumes. WCECS'09. International Association of Engineers. Hong Kong, China: Newswood Limited, 2009, pages 14–17. ISBN: 9789881701268. URL: http://www.iaeng.org/publication/WCECS2009/WCECS2009_pp14-17.pdf (cited on page 245).

[321] Kwan Min Lee. 'Presence, Explicated'. In: *Communication Theory* 14.1 (2004), pages 27–50. ISSN: 1468-2885. DOI: 10.1111/j.1468-2885.2004.tb00302.x (cited on pages 15, 17 sq.).

[322] Seon-Woo Lee and Kenji Mase. 'Activity and Location Recognition Using Wearable Sensors'. In: *IEEE Pervasive Computing* 1.3 (July 2002), pages 24–32. ISSN: 1536-1268. DOI: 10.1109/MPRV.2002.1037719 (cited on page 245).

[323] Sook Jung Lee. 'Online Communication and Adolescent Social Ties: Who Benefits More from Internet Use?' In: *Journal of Computer-Mediated Communication* 14.3 (2009), pages 509–531. ISSN: 1083-6101. DOI: 10.1111/j.1083-6101.2009.01451.x (cited on pages i, 8).

[324] Jonah Lehrer. *How We Decide*. Houghton Mifflin Harcourt, 2009. ISBN: 9780618620111 (cited on page 262).

[325] Christian Leichsenring et al. *Pop-up notification*. Wikipedia. 5 Sept. 2013. URL: http://en.wikipedia.org/w/index.php?title=Pop-up_notification&oldid=571645016 (visited on 2013-09-21) (cited on page 229).

[326] Christian Leichsenring, René Tünnermann and Thomas Hermann. 'feelabuzz – Direct Tactile Communication with Mobile Phones'. In: *International Journal of Mobile Human Computer Interaction* 3.1 (2011). to appear (cited on page 6).

[327] Jonathan Lester, Tanzeem Choudhury and Gaetano Borriello. 'A Practical Approach to Recognizing Physical Activities'. In: *Pervasive Computing*. Edited by Kenneth Fishkin, Bernt Schiele, Paddy Nixon and Aaron Quigley. Volume 3968. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2006, pages 1–16. ISBN: 9783540338949. DOI: 10.1007/11748625_1 (cited on page 245).

[328] Vladimir Iosifovich Levenshtein. 'Binary Codes Capable of Correcting Deletions, Insertions, and Reversals'. In: *Soviet Physics Doklady* (1966), pages 707–710 (cited on page 252).

[329] James R. Lewis. 'IBM Computer Usability Satisfaction Questionnaires: Psychometric Evaluation and Instructions for Use'. In: *International Journal of Human-Computer Interaction* 7.1 (Jan. 1995), pages 57–78. ISSN: 1044-7318. DOI: 10.1080/10447319509526110 (cited on pages 43, 121).

[330] Matthew D. Lieberman and Robert Rosenthal. 'Why Introverts Can't Always Tell Who Likes Them: Multitasking and Nonverbal Decoding'. In: *Journal of Personality and Social Psychology* 80.2 (2001), pages 294–310. ISSN: 1939-1315. DOI: 10.1037/0022-3514.80.2.294 (cited on page 58).

[331] Hubert W. Lilliefors. 'On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown'. In: *Journal of the American Statistical Association* 62.318 (1967), pages 399–402. DOI: 10.2307/2283970 (cited on pages 153, 434).

[332] D. Lindstrom, Lester, Phily D. et al. *HTC Desire*. Wikipedia. 25 July 2012. URL: http://en.wikipedia.org/w/index.php?title=HTC_Desire&oldid=504030882 (visited on 2012-08-29) (cited on page 105).

[333] Tonya Lippert and Karen J. Prager. 'Daily Experiences of Intimacy: A Study of Couples'. In: *Personal Relationships* 8.3 (2001), pages 283–298. ISSN: 1350-4126. DOI: 10.1111/j.1475-6811.2001.tb00041.x (cited on page 31).

[334] John Lisman and Eliezer J. Sternberg. 'Habit and Nonhabit Systems for Unconscious and Conscious Behavior: Implications for Multitasking'. In: *Journal of Cognitive Neuroscience* 25.2 (Nov. 2012), pages 273–283. ISSN: 0898-929X. DOI: 10.1162/jocn_a_00319 (cited on page 58).

[335] Huan Liu and Hiroshi Motoda. *Feature Selection for Knowledge Discovery and Data Mining*. Volume 454. Springer, 1998 (cited on page 264).

[336] Meghann Lomas, Robert Chevalier, Ernest Vincent Cross II, Robert Christopher Garrett, John Hoare and Michael Kopack. 'Explaining Robot Actions'. In: *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*. HRI '12. Boston, MA, USA: ACM, 2012, pages 187–188. ISBN: 9781450310635. DOI: 10.1145/2157689.2157748 (cited on page 262).

[337] Matthew Lombard. *Resources for the Study of Presence: Presence Explication*. July 2000. URL: http://web.archive.org/web/20011215044926/http://nimbus.temple.edu/~mlombard/Presence/explicat.htm (cited on page 17).

[338] Matthew Lombard, Theresa B. Ditton, Daliza Crane, Bill Davis, Gisela Gil-Egui, Karl Horvath and Jessica Rossman. 'Measuring Presence: A Literature-Based Approach to the Development of a Standardized Paper-and-Pencil Instrument'. In: *Proceedings of the Third International Workshop on Presence*. 2000 (cited on page 41).

[339] Matthew Lombard, Melissa Selverian, Karl Horvath, Sung Bok Park and Ha Sung Hwang, editors. *Telepresence in Popular Culture: A Study of Portrayals of Presence*. 16 Sept. 2006. URL: http://web.archive.org/web/20060916001823/http://lombardresearch.temple.edu/p4/index.htm (visited on 2012-08-30) (cited on page 3).

[340] Frank Lömker. *iceWing – A Graphical Plugin Shell*. Bielefeld University. 3 July 2013. URL: http://icewing.sourceforge.net/ (visited on 2013-07-13) (cited on page 231).

[341] Frank Lömker, Sebastian Wrede, Marc Hanheide and Jannik Fritsch. 'Building Modular Vision Systems with a Graphical Plugin Environment'. In: *Proceedings of the Fourth IEEE International Conference on Computer Vision Systems*. ICVS '06. Washington, DC, USA: IEEE Computer Society, 2006, pages 2–9. ISBN: 0769525067. DOI: 10.1109/ICVS.2006.18 (cited on page 231).

[342] Xi Long, Bin Yin and Ronald M. Aarts. 'Single-Accelerometer-Based Daily Physical Activity Classification'. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. (Minneapolis, MN, USA, 3–6 Sept. 2009). EMBC'09. 2009, pages 6107–6110. DOI: 10.1109/IEMBS.2009.5334925 (cited on page 245).

[343] Jack M. Loomis. 'Distal Attribution and Presence'. In: *Presence: Teleoperators and Virtual Environments* 1.1 (Jan. 1992), pages 113–119. ISSN: 1054-7460 (cited on page 17).

[344] Mark Maat and Dirk Heylen. 'Turn Management or Impression Management?' In: *Intelligent Virtual Agents*. Edited by Zsófia Ruttkay, Michael Kipp, Anton Nijholt and HannesHögni Vilhjálmsson. Volume 5773. Lecture Notes in Computer Science. Springer, 2009, pages 467–473. ISBN: 9783642043796. DOI: 10.1007/978-3-642-04380-2_51 (cited on page 82).

[345] Steve Mann, James Fung, Anurag Sehgal Interaction, Daniel Chen et al. 'Designing EyeTap Digital Eyeglasses for Continuous Lifelong Capture and Sharing of Personal Experiences'. In: *Proceedings of Alt.Chi*. 2005 (cited on page 301).

[346] Andrea Mannini and Angelo Maria Sabatini. 'Machine Learning Methods for Classifying Human Physical Activity from On-Body Accelerometers'. In: *Sensors* 10.2 (2010), pages 1154–1175. ISSN: 1424-8220. DOI: 10.3390/s100201154. URL: http://www.mdpi.com/1424-8220/10/2/1154 (cited on page 245).

[347] Wenbo Mao and Kenneth G. Paterson. *On the Plausible Deniability Feature of Internet Protocols*. Technical report. Information Security Group – Royal Holloway University of London, 23 Jan. 2003. URL: http://www.isg.rhul.ac.uk/~kp/IKE.ps (cited on page 72).

[348] Panos Markopoulos. 'A Design Framework for Awareness Systems'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 2, pages 49–72. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 26, 28, 38, 55 sqq., 74, 79, 218).

[349] Panos Markopoulos, Bert Bongers, Erik Van Alphen, Jasper Dekker, Wouter Van Dijk, Sebastiaan Messemaker, Joep Van Poppel, Bram Van der Vlist, Dirk Volman and Gilles Van Wanrooij. 'The PhotoMirror Appliance: Affective Awareness in the Hallway'. In: *Personal and Ubiquitous Computing* 10.2-3 (Jan. 2006), pages 128–135. ISSN: 1617-4909. DOI: 10.1007/s00779-005-0007-x (cited on pages 57, 71).

[350] Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay, editors. *Awareness Systems: Advances in Theory, Methodology and Design*. 1st edition. Human-Computer Interaction Series. Springer, 2009. 492 pages. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 3, 13, 15, 72).

[351] Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 'Preface'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. 1st edition. Human-Computer Interaction Series. Springer, 2009. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 5, 25 sq.).

[352] Panos Markopoulos, Wijnand A. IJsselsteijn, Claire Huijnen and Boris E. R. de Ruyter. 'Sharing Experiences through Awareness Systems in the Home'. In: *Interacting with Computers* 17.5 (2005), pages 506–521. ISSN: 0953-5438. DOI: 10.1016/j.intcom.2005.03.004 (cited on page 25).

[353] Panos Markopoulos, Natalia Romero, Joy K. van Baren, Wijnand A. IJsselsteijn, Boris de Ruyter and Babak Farshchian. 'Keeping in Touch with the Family: Home and Away with the ASTRA Awareness System'. In: *Extended Abstracts on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: ACM, 2004, pages 1351–1354. ISBN: 1581137036. DOI: 10.1145/985921.986062 (cited on page 37).

[354] M. Lynne Markus. 'Toward a "Critical Mass" Theory of Interactive Media: Universal Access, Interdependence and Diffusion'. In: *Communication Research* 14.5 (1987), pages 491–511. DOI: 10.1177/009365087014005003 (cited on page 77).

[355] Jeremy Marozeau, Alain de Cheveigné, Stephen McAdams and Suzanne Winsberg. 'The Dependency of Timbre on Fundamental Frequency'. In: *Journal of the Acoustical Society of America* 114 (5 25 June 2003), pages 2946–2957. ISSN: 0001-4966. DOI: 10.1121/1.1618239 (cited on page 232).

[356] David Martens, Bart Baesens and Tony van Gestel. 'Decompositional Rule Extraction from Support Vector Machines'. In: *IEEE Transactions on Knowledge and Data Engineering* 21.2 (Feb. 2009), pages 178–191 (cited on page 273).

[357] David Martens, Bart Baesens, Tony van Gestel and Jan Vanthienen. 'Adding Comprehensibility to Support Vector Machines Using Rule Extraction Techniques'. In: *Credit Scoring and Credit Control* IX (2005) (cited on pages 258, 272).

[358] David Martens, Bart Baesens, Tony van Gestel and Jan Vanthienen. 'Comprehensible Credit Scoring Models Using Rule Extraction from Support Vector Machines'. In: *European Journal of Operational Research* 183.3 (2007), pages 1466–1476. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2006.04.051 (cited on page 273).

[359] David Martens, Johan Huysmans, Rudy Setiono, Jan Vanthienen and Bart Baesens. 'Rule Extraction from Support Vector Machines: An Overview of Issues and Application in Credit Scoring'. In: *Rule Extraction from Support Vector Machines*. Edited by Joachim Diederich. Volume 80. Studies in Computational Intelligence. Springer Berlin Heidelberg, 2008, pages 33–63. ISBN: 9783540753896. DOI: 10.1007/978-3-540-75390-2_2 (cited on page 258).

[360] Debra J. Mashek and Arthur P. Aron, editors. *Handbook of Closeness and Intimacy*. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. 464 pages. ISBN: 9780805842845 (cited on page 15).

[361] Debra J. Mashek and Michelle D. Sherman. 'Desiring Less Closeness with Intimate Others'. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 19, pages 343–356. ISBN: 9780805842845 (cited on page 45).

[362] Tara Matthews, Anind K. Dey, Jennifer Mankoff, Scott Carter and Tye Rattenbury. 'A Toolkit for Managing User Attention in Peripheral Displays'. In: *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*. UIST. Santa Fe, NM, USA: ACM, 2004, pages 247–256. ISBN: 1581139578. DOI: 10.1145/1029632.1029676 (cited on pages 56, 74).

[363] Uwe Maurer, Asim Smailagic, Daniel P. Siewiorek and Michael Deisher. 'Activity Recognition and Monitoring Using Multiple Sensors on Different Body Positions'. In: *International Workshop on Wearable and Implantable Body Sensor Networks*. (Cambridge, MA, USA, 3–5 Apr. 2006). BSN'06. 2006, pages 113–116. ISBN: 0769525474. DOI: 10.1109/BSN.2006.6 (cited on page 245).

[364] Kate McCarthy, Bibi Zabar and Gary Weiss. 'Does Cost-Sensitive Learning Beat Sampling for Classifying Rare Classes?' In: *Proceedings of the 1st International Workshop on Utility-Based Data Mining*. UBDM '05. Chicago, IL, USA: ACM, 2005, pages 69–77. ISBN: 1595932089. DOI: 10.1145/1089827.1089836 (cited on page 263).

[365] D. Scott McCrickard, C. M. Chewar, Jacob P. Somervell and Ali Ndiwalana. 'A Model for Notification Systems Evaluation: Assessing User Goals for Multitasking Activity'. In: *ACM Transactions on Computer-Human Interaction*. TOCHI 10.4 (Dec. 2003), pages 312–338. ISSN: 1073-0516. DOI: 10.1145/966930.966933 (cited on page 56).

[366] Gregor McEwan and Saul Greenberg. 'Supporting Social Worlds with the Community Bar'. In: *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*. GROUP '05. Sanibel Island, FL, USA: ACM, 2005, pages 21–30. ISBN: 1595932232. DOI: 10.1145/1099203.1099207 (cited on pages 27, 38, 58 sq., 70, 78, 87, 90).

[367] Gregor McEwan, Saul Greenberg, Michael Rounding and Michael Boyle. 'Groupware Plug-ins: A Case Study of Extending Collaboration Functionality through Media Items'. In: *Proceedings of 2nd International Conference on Collaboration Technologies*. 23 Feb. 2006 (cited on page 87).

[368] Ken McGarry and Stefan Wermter. 'Training Without Data: Knowledge Insertion into RBF Neural Networks'. In: *Proceedings of the 19th International Joint Conference on Artificial Intelligence*. IJCAI'05. Edinburgh, Scotland: Morgan Kaufmann Publishers, 2005, pages 792–797 (cited on page 278).

[369] Kenneth J. McGarry and John MacIntyre. 'Knowledge Extraction and Insertion from Radial Basis Function Networks'. In: *IEE Colloquium on Applied Statistical Pattern Recognition*. (Brimingham, UK, 20 Apr. 1999). 1999, pages 15/1–15/6. DOI: 10.1049/ic:19990372 (cited on pages 265, 278).

[370] David McGookin and Stephen Brewster. 'Earcons'. In: *The Sonification Handbook*. Edited by Thomas Hermann, Andy Hunt and John G. Neuhoff. Berlin, Germany: Logos Publishing House, 2011. Chapter 14, pages 339–361. ISBN: 9783832528195. URL: http://sonification.de/handbook/chapters/chapter14/ (cited on page 91).

[371] Hilary McLellan. 'Virtual Realities'. In: *Handbook of Research for Educational Communications and Technology* (1996), pages 457–487 (cited on page 17).

[372] Michelle Meade and Henry Roediger. 'Explorations in the Social Contagion of Memory'. In: *Memory & Cognition* 30 (7 2002), pages 995–1009. ISSN: 0090-502X. DOI: 10.3758/BF03194318 (cited on page 73).

[373] Larry R. Medsker. *Hybrid Intelligent Systems*. Norwell, MA, USA: Kluwer Academic Publishers, 1995. ISBN: 0792395883 (cited on page 257).

[374] Paul Mermelstein. 'Distance Measures for Speech Recognition, Psychological and Instrumental'. In: *Pattern Recognition and Artificial Intelligence: Proceedings of the Joint Workshop on Pattern Recognition and Artificial Intelligence*. Edited by Chi Hau Chen.

Academic Press Rapid Manuscript Reproduction. New York, NY, USA: Academic Press, 1976, pages 374–388. ISBN: 9780121709501 (cited on page 246).

[375] Christian Mertes. 'Multimodal Augmented Reality to Enhance Human Communication'. Master's thesis. Bielefeld University, Aug. 2008. URL: http://bieson.ub.uni-bielefeld.de/volltexte/2009/1414/ (cited on pages 90 sq., 300).

[376] Christian Mertes, Angelika Dierker, Thomas Hermann, Marc Hanheide and Gerhard Sagerer. 'Enhancing Human Cooperation with Multimodal Augmented Reality'. In: *Proceedings of the 13th International Conference on Human-Computer Interaction*. Heidelberg, Germany: Springer, July 2009, pages 447–451 (cited on page 91).

[377] George A. Miller. 'The Magical Number Seven Plus or Minus Two: Some Limits on Our Capacity for Processing Information'. In: *The Psychological Review* 63.2 (Mar. 1956), pages 81–97. ISSN: 0033-295X. URL: http://cogprints.org/730/1/miller.html (cited on page 74).

[378] Paul J. E. Miller, Sylvia Niehuis and Ted L. Huston. 'Positive Illusions in Marital Relationships: A 13-Year Longitudinal Study'. In: *Personality and Social Psychology Bulletin* 32.12 (2006), pages 1579–1594. DOI: 10.1177/0146167206292691 (cited on page 46).

[379] Emiliano Miluzzo, Nicholas D. Lane, Kristóf Fodor, Ronald Peterson, Hong Lu, Mirco Musolesi, Shane B. Eisenman, Xiao Zheng and Andrew T. Campbell. 'Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application'. In: *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*. SenSys '08. Raleigh, NC, USA: ACM, 2008, pages 337–350. ISBN: 9781595939906. DOI: 10.1145/1460412.1460445 (cited on page 245).

[380] Marvin Minsky. 'Telepresence'. In: *Omni* 2.9 (1980), pages 45–52 (cited on page 17).

[381] Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1967. ISBN: 0131655639 (cited on page 278).

[382] Ben Moir et al. *Fundaware*. Durex. 17 Apr. 2013. URL: https://www.youtube.com/watch?v=T6vul95hwOY (cited on page 96).

[383] Ashley Montagu. *Touching: The Human Significance of the Skin*. Harper & Row, 1971. ISBN: 9780060802509 (cited on page 99).

[384] Joanna Montgomery. *Pillow Talk*. Little Riot. 2013. URL: http://littleriot.com/pillowtalk/ (cited on page 91).

[385] Sean Moss-Pultz and Timothy Chen. *Openmoko*. URL: http://www.openmoko.com/ (cited on page 105).

[386] Sean Moss-Pultz, Harald Welte, Michael Lauer, Werner Almesberg, Teenie Hung, William Lai, Tim Lee, Tina Kao and Harry Tsai. *Freerunner*. URL: http://wiki.openmoko.org/index.php?title=Neo_FreeRunner&oldid=90150 (cited on page 105).

[387] Martin Mühlenbrock, Oliver Brdiczka, Dave Snowdon and Jean-Luc Meunier. 'Learning to Detect User Activity and Availability from a Variety of Sensor Data'. In: *Proceedings of the Second IEEE International Conference on Pervasive Computing and Communications*. PERCOM '04. Washington, DC, USA: IEEE, 2004, pages 13–. ISBN: 0769520901 (cited on page 293).

[388] Max Mühlhäuser and Melanie Hartmann. 'Interacting with Context'. In: *Quality of Context*. Edited by Kurt Rothermel, Dieter Fritsch, Wolfgang Blochinger and Frank Dürr. Volume 5786. Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2009, pages 1–14. ISBN: 9783642045585. DOI: 10.1007/978-3-642-04559-2_1 (cited on page 225).

[389] Patrick M. Murphy and Michael J. Pazzani. 'ID2-of-3: Constructive Induction of *m*-of-*n* Concepts for Discriminators in Decision Trees'. In: *Proceedings of the Eighth International Workshop on Machine Learning*. 1991, pages 183–187 (cited on page 274).

[390] Roderick Murray-Smith, Andrew Ramsay, Simon Garrod, Melissa Jackson and Bojan Musizza. 'Gait Alignment in Mobile Phone Conversations'. In: *MobileHCI '07: Proceedings of the 9th International Conference on Human-Computer Interaction with Mobile Devices and Services*. Singapore: ACM, 2007, pages 214–221. ISBN: 9781595938626. DOI: 10.1145/1377999.1378009 (cited on pages 79, 96, 210).

[391] Elizabeth D. Mynatt, Maribeth Back, Roy Want, Michael Baer and Jason B. Ellis. 'Designing Audio Aura'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '98. Los Angeles, CA, USA: ACM/Addison-Wesley, 1998, pages 566–573. ISBN: 0201309874. DOI: 10.1145/274644.274720 (cited on page 91).

[392] Elizabeth D. Mynatt, Jim Rowan, Sarah Craighill and Annie Jacobs. 'Digital Family Portraits: Supporting Peace of Mind for Extended Family Members'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. (Seattle, WA, USA). CHI '01. New York, NY, USA: ACM, 2001, pages 333–340. ISBN: 1581133278. DOI: 10.1145/365024.365126 (cited on pages 51, 87).

[393] Bonnie A. Nardi, Steve Whittaker and Erin Bradner. 'Interaction and Outeraction: Instant Messaging in Action'. In: *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*. CSCW '00. Philadelphia, PA, USA: ACM, 2000, pages 79–88. ISBN: 1581132220. DOI: 10.1145/358916.358975 (cited on pages 72, 226).

[394] J. Neumann. *Classification and Evaluation of Algorithms for Rule Extraction from Artificial Neural Networks*. PhD summer project, University of Edinburgh. 1998 (cited on page 223).

[395] Andrew Ng. *Machine Learning: Lecture 6*. OpenCourseWare. Standford University. 2007. URL: http://see.stanford.edu/materials/aimlcs229/transcripts/MachineLearning-Lecture06.html (cited on page 268).

[396] Mary Ellen Nivison and Inger M. Endresen. 'An Analysis of Relationships among Environmental Noise, Annoyance and Sensitivity to Noise, and the Consequences for Health and Sleep'. In: *Journal of Behavioral Medicine* 16 (3 1993), pages 257–276. ISSN: 0160-7715. DOI: 10.1007/BF00844759 (cited on page 128).

[397] Don Norman. 'Emotion & Design: Attractive Things Work Better'. In: *Interactions* 9.4 (July 2002), pages 36–42. ISSN: 1072-5520. DOI: 10.1145/543434.543435 (cited on page 76).

[398] Donald A. Norman. *The Psychology Of Everyday Things*. New York, NY: Basic Books, 13 June 1988. ISBN: 0465067093 (cited on page 57).

[399] Donald A. Norman. 'Affordance, conventions, and design'. In: *interactions* 6.3 (May 1999), pages 38–43. ISSN: 1072-5520. DOI: 10.1145/301153.301168 (cited on page 57).

[400] Nicolas Nova. 'The Impact of Awareness Tools on Mutual Modelling in a Collaborative Game'. Master's thesis. University of Geneva, Oct. 2002. URL: http://tecfa.unige.ch/~nova/msc_nova.pdf (cited on page 24).

[401] Kristine L. Nowak and Frank Biocca. 'The Effect of the Agency and Anthropomorphism on Users' Sense of Telepresence, Copresence, and Social Presence in Virtual Environments'. In: *Presence: Teleoperators & Virtual Environments* 12.5 (2003), pages 481–494 (cited on pages 5, 17, 33, 41 sq., 145, 196, 423, 436).

[402] Haydemar Núñez-Castro, Cecilio Angulo-Bahón and Andreu Català-Mallofré. 'Hybrid Architecture Based on Support Vector Machines'. In: *Computational Methods in Neural Modeling*. Edited by José Mira and José R. Álvarez. Volume 2686. Lecture Notes in Computer Science. Springer, 2003, pages 646–653. ISBN: 9783540402107. DOI: 10.1007/3-540-44868-3_82 (cited on page 278).

[403] Ian Oakley and Sile O'Modhrain. 'Contact IM: Exploring Asynchronous Touch over Distance'. In: *Proceedings of the ACM conference on Computer-Supported Cooperative Work*. CSCW. 2002, pages 16–20 (cited on page 94).

[404] Shannon O'Brien and Florian Mueller. 'Holding Hands Over a Distance: Technology Probes in an Intimate, Mobile Context'. In: *Proceedings of the 18th Australia Conference on Computer-Human Interaction*. OZCHI '06. Sydney, Australia: ACM, 2006, pages 293–296. ISBN: 1595935452. DOI: 10.1145/1228175.1228226 (cited on pages 6, 95, 97).

[405] Daniel C. O'Connell, Sabine Kowal and Erika Kaltenbacher. 'Turn-Taking: A Critical Analysis of the Research Tradition'. In: *Journal of Psycholinguistic Research* 19.6 (1990), pages 345–373. ISSN: 0090-6905. DOI: 10.1007/BF01068884 (cited on page 82).

[406] Johanna Renny Octavia, Elise van den Hoven and Hans De Mondt. ‘Overcoming the Distance between Friends’. In: *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI … But Not as We Know It*. Volume 2. BCS-HCI ’07. University of Lancaster, United Kingdom: British Computer Society, 2007, pages 79–82. ISBN: 9781902505954 (cited on page 99).

[407] Christian W. Omlin and C. Lee Giles. ‘Stable Encoding of Large Finite-State Automata in Recurrent Neural Networks with Sigmoid Discriminants’. In: *Neural Computation* 8.4 (May 1996), pages 675–696. ISSN: 0899-7667. DOI: 10.1162/neco.1996.8.4.675 (cited on page 278).

[408] Christian W. Omlin, C. Lee Giles, Bill G. Horne, Laurens R. Leerink and Tsungnan Lin. ‘Training Recurrent Neural Networks with Temporal Input Encodings’. In: *IEEE International Conference on Neural Networks: IEEE World Congress on Computational Intelligence*. (Orlando, FL, USA). Volume 2. 1994, pages 1267–1272. DOI: 10.1109/ICNN.1994.374366 (cited on page 278).

[409] Charles Egerton Osgood, George J. Suci and Percy H. Tannenbaum. *The Measurement of Meaning*. Illini Books. Urbana, IL: University of Illinois Press, 1957. 342 pages (cited on pages 41, 154, 200).

[410] Oxford Dictionaries. *Presence*. Oxford University Press. Apr. 2010. URL: http://oxforddictionaries.com/definition/presence (cited on page 18).

[411] Oxford University Press. ‘plausible deniability’. In: *The Oxford Essential Dictionary of the U.S. Military*. Berkley, 2001 (cited on page 72).

[412] Oxford University Press. *Presence*. Oxford University Press. 2014. URL: http://www.oed.com/view/Entry/150669? (cited on page 19).

[413] Fred G. W. C. Paas, Jeroen J. G. van Merriënboer and Jos J. Adam. ‘Measurement of Cognitive Load in Instructional Research’. In: *Perceptual and Motor Skills* 79.1 (1994), pages 419–430. ISSN: 0031-5125 (cited on page 43).

[414] Leysia Palen and Paul Dourish. 'Unpacking "Privacy" for a Networked World'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '03. Ft. Lauderdale, FL, USA: ACM, 2003, pages 129–136. ISBN: 1581136307. DOI: 10.1145/642611.642635 (cited on pages 10, 25, 70).

[415] Jaak Panksepp. *Affective Neuroscience: The Foundations of Human and Animal Emotions*. Volume 4. USA: Oxford University Press, 30 Sept. 2004. 482 pages. ISBN: 9780195178050 (cited on page 76).

[416] Cécile L. Paris. 'Generation and Explanation: Building an Explanation Facility for the Explainable Expert Systems Framework'. In: *Natural Language Generation in Artificial Intelligence and Computational Linguistics*. Edited by Cécile L. Paris, William R. Swartout and William C. Mann. Volume 119. The Kluwer International Series in Engineering and Computer Science. Springer, 1991, pages 49–82. ISBN: 9781441951250. DOI: 10.1007/978-1-4757-5945-7_2 (cited on page 262).

[417] Young-Woo Park, Chang-Young Lim and Tek-Jin Nam. 'CheekTouch: An Affective Interaction Technique While Speaking on the Mobile Phone'. In: *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '10. Atlanta, GA, USA: ACM, 2010, pages 3241–3246. ISBN: 9781605589305. DOI: 10.1145/1753846.1753965 (cited on page 96).

[418] Juha Pärkkä, Miikka Ermes, Panu Korpipää, Jani Mäntyjärvi, Johannes Peltola and Ilkka Korhonen. 'Activity Classification Using Realistic Data from Wearable Sensors'. In: *Transactions on Information Technology in Biomedicine* 10.1 (Jan. 2006), pages 119–128. ISSN: 1089-7771. DOI: 10.1109/TITB.2005.856863 (cited on page 245).

[419] Sameer Patil and Alfred Kobsa. 'Privacy Considerations in Awareness Systems: Designing with Privacy in Mind'. In: *Awareness Systems*. Edited by Panos Markopoulos, Wendy Mackay and Boris E. R. de Ruyter. Human–Computer Interaction Series. 2009, pages 187–206. ISBN: 9781848824775 (cited on pages 68, 71).

[420] Michael J. Pazzani. 'Knowledge discovery from Data?' In: *IEEE Intelligent Systems and their Applications* 15.2 (2000), pages 10–12 (cited on pages 223, 264).

[421] Eja Pedersen and Kerstin Persson Waye. 'Perception and Annoyance due to Wind Turbine Noise: A Dose-Response Relationship'. In: *Acoustical Society of America Journal* 116 (Dec. 2004), pages 3460–3470. DOI: 10.1121/1.1815091 (cited on page 128).

[422] Elin Rønby Pedersen. 'People Presence or Room Activity: Supporting Peripheral Awareness Over Distance'. In: *Conference Summary on Human Factors in Computing Systems*. CHI '98. Los Angeles, CA, USA: ACM, 1998, pages 283–284. ISBN: 1581130287. DOI: 10.1145/286498.286763 (cited on pages 28, 65).

[423] Elin Rønby Pedersen and Tomas Sokoler. 'AROMA: Abstract Representation of Presence Supporting Mutual Awareness'. In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. CHI '97. Atlanta, GA, USA: ACM, 1997, pages 51–58. ISBN: 0897918029. DOI: 10.1145/258549.258584 (cited on pages 28, 65).

[424] Hanchuan Peng, Adam Vepstas Pocock, Linas, Jörg Kurt Wegner, Michael Hardy et al. *Feature Selection*. Wikipedia. 19 Aug. 2013. URL: https://en.wikipedia.org/w/index.php?title=Feature_selection&oldid=569190180 (visited on 2013-09-04) (cited on page 239).

[425] Will Penny and David Frost. 'Neural Networks in Clinical Medicine'. In: *Medical Decision Making* 16.4 (1996), pages 386–398. DOI: 10.1177/0272989X9601600409 (cited on page 258).

[426] Daniel Perlman and Beverley Fehr. 'The Development of Intimate Relationships'. In: *Intimate Relationships: Development, Dynamics, and Deterioration*. Edited by D. Perlman S. Duck. Thousand Oaks, CA, USA: Sage Publications, Inc, 1987, pages 13–42 (cited on page 31).

[427] Ralph Meijer Peter Millard Peter Saint-Andre. *XEP-0060: Publish-Subscribe*. 12 July 2010. URL: http://xmpp.org/extensions/xep-0060.html (cited on page 119).

[428] Simon H. Pincus, Robert House, Joseph Christenson and Lawrence E. Adler. 'The Emotional Cycle of Deployment: A Military Family Perspective'. In: *US Army Medical Department Journal* 4.5 (2001), page 6 (cited on page 47).

[429] David Pinelle and Carl Gutwin. 'A Groupware Design Framework for Loosely Coupled Workgroups: Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work'. In: *ECSCW 2005*. Edited by Hans Gellersen, Kjeld Schmidt, Michel Beaudouin-Lafon and Wendy Mackay. Springer Netherlands, 2005. Chapter 4, pages 65–82. ISBN: 9781402040221. DOI: 10.1007/1-4020-4023-7_4 (cited on page 51).

[430] Zachary Pousman and John Stasko. 'A Taxonomy of Ambient Information Systems: Four Patterns of Design'. In: *Proceedings of the Working Conference on Advanced Visual Interfaces*. AVI. Venezia, Italy: ACM, 2006, pages 67–74. ISBN: 1595933530. DOI: 10.1145/1133265.1133277 (cited on pages 56, 60, 62, 74).

[431] Karen J. Prager. *The Psychology of Intimacy*. New York, NY, USA: Guilford Press, 1995 (cited on page 31).

[432] Karen J. Prager and Duane Buhrmester. 'Intimacy and Need Fulfillment in Couple Relationships'. In: *Journal of Social and Personal Relationships* 15.4 (1998), pages 435–469. ISSN: 0265-4075. DOI: 10.1177/0265407598154001 (cited on page 31).

[433] Karen J. Prager and Linda J. Roberts. 'Deep Intimate Connection: Self and Intimacy in Couple Relationships'. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 4, pages 43–60. ISBN: 9780805842845 (cited on page 31).

[434] J. Ross Quinlan. *C4.5: Programs for Machine Learning*. San Francisco, CA, USA: Morgan Kaufmann Publishers, 1993. ISBN: 1558602380 (cited on pages 265, 274).

[435] Mika Raento, Antti Oulasvirta, Renaud Petit and Hannu Toivonen. 'ContextPhone: A Prototyping Platform for Context-Aware Mobile Applications'. In: *Pervasive Computing, IEEE* 4.2 (Jan. 2005), pages 51–59. ISSN: 1536-1268. DOI: 10.1109/MPRV.2005.29 (cited on pages 102, 225).

[436] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore and Michael L. Littman. 'Activity Recognition from Accelerometer Data'. In: *Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence*. Volume 3. IAAI'05. Pittsburgh, PA, USA: AAAI Press, 2005, pages 1541–1546. ISBN: 157735236X (cited on page 245).

[437] Ingo Rechenberg. 'Evolutionsstrategien'. German. In: *Simulationsmethoden in der Medizin und Biologie*. Edited by Berthold Schneider and Ulrich Ranft. Volume 8. Medizinische Informatik und Statistik. Springer Berlin Heidelberg, 1978, pages 83–114. ISBN: 9783540090502. DOI: 10.1007/978-3-642-81283-5_8 (cited on page 274).

[438] Johan Redström, Tobias Skog and Lars Hallnäs. 'Informative Art: Using Amplified Artworks As Information Displays'. In: *Proceedings of DARE 2000 on Designing Augmented Reality Environments*. DARE '00. Elsinore, Denmark: ACM, 2000, pages 103–114. DOI: 10.1145/354666.354677 (cited on page 84).

[439] Byron Reeves. 'Being there: Television as Symbolic versus Natural Experience'. Unpublished manuscript. Stanford, CA, USA, 1991 (cited on page 17).

[440] Harry T. Reis, Margaret S. Clark and John G. Holmes. 'Perceived Partner Responsiveness as an Organizing Construct in the Study of Intimacy and Closeness'. In: *Handbook of Closeness and Intimacy*. Edited by Debra J. Mashek and Arthur P. Aron. Mahwah, NJ, USA: Lawrence Erlbaum Associates Publishers, 11 Feb. 2004. Chapter 12, pages 201–225. ISBN: 9780805842845 (cited on page 15).

[441] Harry T. Reis and Phillip Shaver. 'Intimacy as an Interpersonal Process'. In: *Handbook of Personal Relationships: Theory, Research and Interventions*. Edited by S. Duck, D. F. Hay, S. E. Hobfoll, W. Ickes and B. M. Montgomery. Oxford, England: John Wiley & Sons, 1988, pages 367–389 (cited on page 31).

[442] James R. Reske and Laura Stafford. 'Idealization and Communication in Long-Distance and Geographically Close Premarital Relationships'. In: *Annual Conference of the International Communication Association*. San Francisco, CA, USA, May 1989 (cited on page 45).

[443] Ruth Rettie. 'Connectedness, Awareness and Social Presence'. Social Presence and the Experience of Connectedness. In: *Proceedings of Presence 2003*. The 6th Annual International Workshop on Presence. (Aalborg, Denmark, 6–8 Oct. 2003). 2003. URL: http://www.kingston.ac.ukhttp://eprints.kingston.ac.uk/2106/1/Rettie.pdf/~ku03468/includes/docs/Connectedness,%20Awareness%20and% http://eprints.kingston.ac.uk/2106/1/Rettie.pdf (cited on pages 24, 26, 32 sq., 203).

[444] Peter Reutemann et al. *ARFF (stable version)*. 26 Feb. 2012. URL: http://weka.wikispaces.com/page/view/ARFF+(stable+version)/305204182 (cited on page 229).

[445] Howard Rheingold. *Virtual Reality*. New York, NY, USA: Simon & Schuster Inc., 1991. ISBN: 0671778978 (cited on page 17).

[446] Markus Rittenbruch, Tim Mansfield and Stephen Viller. 'Design and Evaluation of Intentionally Enriched Awareness'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 16, pages 367–395. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on page 226).

[447] Markus Rittenbruch and Gregor McEwan. 'An Historical Reflection of Awareness in Collaboration'. In: *Awareness Systems: Advances in Theory, Methodology and Design*. Edited by Panos Markopoulos, Boris E. R. de Ruyter and Wendy Mackay. 1st edition. Human-Computer Interaction Series. Springer, 2009. Chapter 1, pages 3–48. ISBN: 9781848824768. DOI: 10.1007/978-1-84882-477-5 (cited on pages 26, 51, 59, 70 sq., 78).

[448] Oriana Riva and Cristiano di Flora. 'Contory: A Smart Phone Middleware Supporting Multiple Context Provisioning Strategies'. In: *Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops*. ICDCSW '06. July 2006, page 68. DOI: 10.1109/ICDCSW.2006.33 (cited on page 244).

[449] Amber Roberts and M. Carole Pistole. 'Long-Distance and Proximal Romantic Relationship Satisfaction: Attachment and Closeness Predictors'. In: *Journal of College Counseling* 12.1 (2009), pages 5–

17. ISSN: 2161-1882. DOI: 10.1002/j.2161-1882.2009.tb00036.x (cited on pages 45, 47 sq.).

[450] Tom Rodden. 'Populating the Application: A Model of Awareness for Cooperative Applications'. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. CSCW '96. Boston, MA, USA: ACM, 1996, pages 87–96. ISBN: 0897917650. DOI: 10.1145/240080.240200 (cited on pages 25, 27, 58, 87).

[451] Peter Rodgers, editor. *Euler Diagrams Overview*. 22–23 Sept. 2004. URL: http://www.cs.kent.ac.uk/events/conf/2004/euler/eulerdiagrams.html (visited on 2012-11-19). Euler Diagrams 2004 (cited on page 34).

[452] Yvonne Rogers, Helen Sharp and Jenny Preece. *Interaction Design: Beyond Human-Computer Interaction*. 3rd edition. John Wiley & Sons Ltd, 2011. ISBN: 9780470665763 (cited on pages 19, 53, 57, 69, 292).

[453] Bernice E. Rogowitz and Lloyd A. Treinish. *Why Should Engineers and Scientists be Worried about Color?* IBM Thomas J. Watson Research Center. 2000. URL: http://www.research.ibm.com/people/l/lloydt/color/color.HTM (visited on 2013-02-07) (cited on page 149).

[454] David E. Rohall, Mady Wechsler Segal and David R. Segal. 'Examining the Importance of Organizational Supports on Family Adjustment to Army Life in a Period of Increasing Separation'. In: *Journal of Political and Military Sociology* 27.1 (1999), pages 49–65 (cited on page 48).

[455] Daniela M. Romano, Paul Brna and John A. Self. 'Collaborative Decision-Making and Presence in Shared Dynamic Virtual Environments'. In: *Proceedings of the 1998 Presence in Shared Virtual Evironments Workshop*. 10–11 June 1998 (cited on pages 36, 41, 142, 202).

[456] Gregor Romero Natalia A. iand McEwan and Saul Greenberg. 'A Field Study of Community Bar: (Mis)-Matches Between Theory and Practice'. In: *Proceedings of the 2007 International ACM Conference on Supporting Group Work*. GROUP '07. Sanibel Island, FL, USA: ACM, 2007, pages 89–98. ISBN: 9781595938459. DOI: 10.1145/1316624.1316638 (cited on pages 87, 218).

[457] Natalia A. Romero, Panos Markopoulos, Joy K. van Baren, Boris E. R. de Ruyter, Wijnand A. IJsselsteijn and Babak Farshchian. 'Connecting the Family with Awareness Systems'. In: *Personal and Ubiquitous Computing* 11.4 (Apr. 2007), pages 299–312. ISSN: 1617-4909. DOI: 10.1007/s00779-006-0089-0 (cited on pages 3, 24 sq., 28 sq.).

[458] Francesca Rosella and Ryan Genz. *HugShirt*. CuteCircuit. 2006. URL: http://cutecircuit.com/collections/the-hug-shirt/ (cited on pages 93, 96).

[459] Barbara J. Ross. 'The Emotional Impact of E-Mail on Deployment'. In: *Proceedings of the Naval Institute* 127 (2001), pages 85–86 (cited on page 47).

[460] A. F. Rovers and H. A. van Essen. 'HIM: A Framework for Haptic Instant Messaging'. In: *CHI '04 Extended Abstracts on Human Factors in Computing Systems*. Vienna, Austria: ACM, 2004, pages 1313–1316. ISBN: 1581137036. DOI: 10.1145/985921.986052 (cited on page 96).

[461] Jim Rowan and Elizabeth D. Mynatt. 'Digital Family Portrait Field Trial: Support for Aging in Place'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. (Portland, OR, USA). CHI '05. New York, NY, USA: ACM, 2005, pages 521–530. ISBN: 1581139985. DOI: 10.1145/1054972.1055044 (cited on pages 87, 89).

[462] Lillian B. Rubin. 'On Men and Friendship'. In: *Psychoanalytic Review* 73.2 (1986), pages 165–181 (cited on page 1).

[463] Daniel W. Russell. 'In Search of Underlying Dimensions: The Use (and Abuse) of Factor Analysis'. In: *Personality and Social Psychology Bulletin* 28.12 (2002), pages 1629–1646. DOI: 10.1177/014616702237645. URL: http://psp.sagepub.com/content/28/12/1629.abstract (cited on page 172).

[464] Ronald M. Sabatelli. 'Exploring Relationship Satisfaction: A Social Exchange Perspective on the Interdependence between Theory, Research, and Practice'. In: *Family Relations* 37.2 (Apr. 1988), pages 217–222. ISSN: 01976664 (cited on page 46).

[465] Debashis Saha and Amitava Mukherjee. 'Pervasive Computing: A Paradigm for the 21st Century'. In: *Computer* 36.3 (Mar. 2003), pages 25–31. ISSN: 0018-9162. DOI: 10.1109/MC.2003.1185214 (cited on page 61).

[466] Erin M. Sahlstein. 'Relating at a Distance: Being Together and being Apart in Long-Distance Relationships'. PhD thesis. University of Iowa, 2000 (cited on pages 8, 45).

[467] Erin M. Sahlstein. 'Relating at a Distance: Negotiating being Together and being Apart in Long-Distance Relationships'. In: *Journal of Social and Personal Relationships* 21.5 (2004), page 689 (cited on pages 8, 45, 47).

[468] Eva-Lotta Sallnäs. 'The Effect of Modality on Social Presence, Presence and Performance in Collaborative Virtual Environments'. PhD thesis. Stockholm, Sweden: Royal Institute of Technology (KTH), 2004 (cited on page 83).

[469] Eva-Lotta Sallnäs, Kirsten Rassmus-Gröhn and Calle Sjöström. 'Supporting Presence in Collaborative Environments by Haptic Force Feedback'. In: *ACM Transactions on Computer-Human Interaction* 7.4 (Dec. 2000), pages 461–476. ISSN: 1073-0516. DOI: 10.1145/365058.365086 (cited on page 100).

[470] Dairazalia Sanchez, Monica Tentori and Jesus Favela. 'Hidden Markov Models for Activity Recognition in Ambient Intelligence Environments'. In: *Mexican International Conference on Computer Science* (2007), pages 33–40. DOI: h10.1109/ENC.2007.31 (cited on pages 219, 227, 293).

[471] Robert E. Schapire and Yoram Singer. 'BoosTexter: A Boosting-Based System for Text Categorization'. In: *Machine Learning* 39.2-3 (2000), pages 135–168. ISSN: 0885-6125. DOI: 10.1023/A:1007649029923 (cited on page 253).

[472] Robert W. Scheifler. 1.0. *X Window System Protocol*. 2004. URL: http://www.x.org/releases/X11R7.6/doc/xproto/x11protocol.html (cited on page 236).

[473] Bill N. Schilit, David M. Hilbert and Jonathan Trevor. 'Context-Aware Communication'. In: *IEEE Wireless Communications* 9.5 (Oct. 2002), pages 46–54. ISSN: 1536-1284. DOI: 10.1109/MWC.2002.1043853 (cited on page 225).

[474] David W. Schloerb. 'A Quantitative Measure of Telepresence'. In: *Presence: Teleoperators and Virtual Environments* 4.1 (1995), pages 64–80 (cited on page 17).

[475] Albrecht Schmidt. 'Ubiquitous Computing – Computing in Context: A Thesis on Context-Awareness, Context Aware Computing, and Ubiquitous Computing'. PhD thesis. Lancaster, UK: Computing Department Lancaster University, Nov. 2002. URL: http://www.comp.lancs.ac.uk/~albrecht/phd/ (cited on page 225).

[476] Albrecht Schmidt, Michael Beigl and Hans-Werner Gellersen. 'There is More to Context than Location'. In: *Computers & Graphics* 23.6 (1999), pages 893–901. ISSN: 0097-8493. DOI: 10.1016/S0097-8493(99)00120-X (cited on page 225).

[477] Kjeld Schmidt. 'The Problem with 'Awareness': Introductory Remarks on 'Awareness in CSCW''. In: *Computer Supported Cooperative Work* 11 (3 2002), pages 285–298. ISSN: 0925-9724. DOI: 10.1023/A:1021272909573 (cited on pages 24 sq.).

[478] Fred B. Schneider. 'Preface'. In: *Trust in Cyberspace*. Edited by Fred B. Schneider. Washington, DC, USA: National Academies Press, 1999, pages vii–xi (cited on page 78).

[479] Tobi Schneidler, Magnus Jonsson and Fredrik Petersson. *RemoteHome*. Oct. 2002. URL: http://web.archive.org/web/20090703112912/http://www.tobi.net/remotehome/remotehome.htm (cited on page 217).

[480] Jean Scholtz and Sunny Consolvo. 'Toward a Framework for Evaluating Ubiquitous Computing Applications'. In: *Pervasive Computing* 3.2 (Apr.–June 2004), pages 82–88. ISSN: 1536-1268. DOI: 10.1109/MPRV.2004.1316826 (cited on pages 52 sq., 61).

[481] Ralph Schroeder. 'Being There Together and the Future of Connected Presence'. In: *Presence: Teleoperators and Virtual Environments* 15.4 (Aug. 2006), pages 438–454. ISSN: 1054-7460. DOI: 10.1162/pres.15.4.438 (cited on page 1).

[482] Ralph Schroeder, Anthony Steed, Ann-Sofie Axelsson, Ilona Heldal, Åsa Abelin, Wideström, Alexander Nilsson and Mel Slater. 'Collaborating in Networked Immersive Spaces: As Good as Being There Together?' In: *Computer & Graphics* 25.5 (2001), pages 781–788. ISSN: 0097-8493. DOI: 10.1016/S0097-8493(01)00120-0 (cited on pages 36, 41).

[483] Martijn J. Schuemie, Peter van der Straaten, Merel Krijn and Charles A. P. G. van der Mast. 'Research on Presence in Virtual Reality: A Survey'. In: *CyberPsychology & Behavior* 4.2 (2001), pages 183–201. ISSN: 1094-9313. DOI: 10.1089/109493101300117884 (cited on page 39).

[484] Andrew I. Schwebel, Ryan L. Dunn, Barry F. Moss and Maureena A. Renner. 'Factors Associated with Relationship Stability in Geographically Separated Couples'. In: *Journal of College Student Development* 33.3 (1992), pages 222–30 (cited on page 47).

[485] Hans-Paul Schwefel. *Evolution and Optimum Seeking*. The Sixth Generation. New York, NY, USA: John Wiley & Sons, 1993. ISBN: 0471571482 (cited on page 274).

[486] Cameron Scott et al. *Palm Pre*. Wikipedia. 8 Aug. 2012. URL: http://en.wikipedia.org/w/index.php?title=Palm_Pre&oldid=506395247 (visited on 2012-08-29) (cited on page 105).

[487] Edo Segal. *I was Just Dead: A Tale of Augmented Reality*. 1993. URL: http://futurity.s3.amazonaws.com/iwasjustdead/IWasJustDeadeBook.pdf (cited on page 219).

[488] Ann Senghas, Sotaro Kita and Asli Özyürek. 'Children Creating Core Properties of Language: Evidence from an Emerging Sign Language in Nicaragua'. In: *Science* 305.5691 (17 Sept. 2004), page 1779 (cited on page 80).

[489] Samuel S. Shapiro and R. S. Francia. 'An Approximate Analysis of Variance Test for Normality'. In: *Journal of the American Statistical Association* 67.337 (1972), pages 215–216. DOI: 10.1080/01621459.1972.10481232 (cited on pages 153, 434).

[490] Samuel S. Shapiro and Martin B. Wilk. 'An Analysis of Variance Test for Normality (Complete Samples)'. In: *Biometrika* 52.3-4 (1965), pages 591–611. DOI: 10.1093/biomet/52.3-4.591 (cited on pages 153, 434).

[491] Jude W. Shavlik. 'Combining Symbolic and Neural Learning'. In: *Machine Learning* 14.3 (1994), pages 321–331. ISSN: 0885-6125. DOI: 10.1007/BF00993982 (cited on page 278).

[492] Thomas B. Sheridan. 'Musings on Telepresence and Virtual Presence'. In: *Presence: Teleoperators and Virtual Environments* 1.1 (Jan. 1992), pages 120–126. ISSN: 1054-7460 (cited on page 17).

[493] Thomas B. Sheridan. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, 12 Aug. 1992. 415 pages. ISBN: 9780262193160 (cited on page 17).

[494] John Short, Ederyn Williams and Bruce Christie. *The Social Psychology of Telecommunications*. John Wiley and Sons Ltd., Sept. 1976. 206 pages. ISBN: 9780471015819 (cited on pages 1, 4 sq., 23, 30, 33 sq., 41 sq., 81 sq., 145, 154, 424).

[495] Daniel Siewiorek, Asim Smailagic, Junichi Furukawa, Andreas Krause, Neema Moraveji, Kathryn Reiger, Jeremy Shaffer and Fei Lung Wong. 'SenSay: A Context-Aware Mobile Phone'. In: *Wearable Computers, 2003. Proceedings. Seventh IEEE International Symposium on*. IEEE. 2003, pages 248–249. ISBN: 0769520340 (cited on pages 96, 102).

[496] Bernard W. Silverman. 'Density Estimation for Statistics and Data Analysis'. In: *Monographs on Statistics and Applied Probability*. Volume 26. London, UK: Chapman and Hall, 1986 (cited on page 274).

[497] Mel Slater, David-Paul Pertaub and Anthony Steed. 'Public Speaking in Virtual Reality: Facing An Audience of Avatars'. In: *Computer Graphics and Applications, IEEE* 19.2 (Mar.–Apr. 1999), pages 6–9. ISSN: 0272-1716. DOI: 10.1109/38.749116 (cited on pages 36, 41).

[498] Mel Slater, Amela Sadagic, Martin Usoh and Ralph Schroeder. 'Small-Group Behavior in a Virtual and Real Environment: A Comparative Study'. In: *Presence: Teleoperators & Virtual Environments* 9.1 (2000), pages 37–51 (cited on pages 36, 41 sq., 142, 145, 202).

[499] Mel Slater and Martin Usoh. 'Representations Systems, Perceptual Position, and Presence in Immersive Virtual Environments'. In: *Presence: Teleoperators and Virtual Environments* 2.3 (1993), pages 221–233 (cited on page 17).

[500] Eliot R. Smith and Diane M. MacKie. *Social Psychology*. 3rd edition. Psychology Press, Jan. 2007. 657 pages. ISBN: 9781841694085. URL: http://psypress.co.uk/smithandmackie/ (cited on page 29).

[501] Jocelyn Darlene Smith. 'Communicating Emotion through a Haptic Link: A Study of the Influence of Metaphor, Personal Space and Relationship'. Master's thesis. Vancouver, Canada: University of British Columbia, 2005. URL: https://circle.ubc.ca/handle/2429/18176 (cited on page 92).

[502] Jocelyn Darlene Smith and Karon MacLean. 'Communicating Emotion through a Haptic Link: Design Space and Methodology'. In: *International Journal of Human-Computer Studies* 65.4 (2007), pages 376–387. ISSN: 1071-5819. DOI: 10.1016/j.ijhcs.2006.11.006 (cited on pages 92, 94, 97).

[503] Daniel J. Solove. 'A Taxonomy of Privacy'. In: *University of Pennsylvania Law Review* 154.3 (Jan. 2006). GWU Law School Public Law Research Paper No. 129, pages 477–560 (cited on page 68).

[504] Daniel J. Solove. ''I've Got Nothing to Hide' and Other Misunderstandings of Privacy'. In: *San Diego Law Review* 44 (2007). GWU Law School Public Law Research Paper No. 289 (cited on page 68).

[505] Christa Sommerer and Laurent Mignonneau. 'Mobile Feelings: Wireless Communication of Heartbeat and Breath for Mobile Art'. In: *The Mobile Audience: Media Art and Mobile Technologies*. Edited by Martin Rieser. Volume 5. Architecture, Technology, Culture. Rodopi, 2011. ISBN: 9789042031289 (cited on pages 94, 97).

[506] Graham B. Spanier. 'Measuring Dyadic Adjustment: New Scales for Assessing the Quality of Marriage and Similar Dyads'. In: *Journal of Marriage & the Family* 38 (1976), pages 15–28 (cited on page 48).

[507] Laura Stafford. *Maintaining Long-Distance and Cross-Residential Relationships*. LEA's Communication Series. Taylor & Francis, 2004. ISBN: 9780805851649 (cited on pages i, 8, 44 sqq., 73, 145, 147).

[508] Laura Stafford and Andy J. Merolla. 'Idealization, Reunions, and Stability in Long-Distance Dating Relationships'. In: *Journal of Social and Personal Relationships* 24.1 (2007), pages 37–54. DOI: 10.1177/0265407507072578 (cited on pages 45 sq.).

[509] Laura Stafford, Andy J. Merolla and Janessa D. Castle. 'When Long-Distance Dating Partners Become Geographically Close'. In: *Journal of Social and Personal Relationships* 23.6 (2006), pages 901–919. DOI: 10.1177/0265407506070472 (cited on page 46).

[510] Laura Stafford and James R. Reske. 'Idealization and Communication in Long-Distance Premarital Relationships'. In: *Family Relations* 39.3 (1990), pages 274–279. ISSN: 01976664 (cited on pages 46 sq.).

[511] Laura Stafford and S. C. Yost. 'The Role of Communication in Naval Couples' Marital Satisfaction'. In: *Annual Conference of the National Communication Association*. (Chicago, USA). Nov. 1990 (cited on page 47).

[512] Pink Stars. *Companionable Silence*. 14 Oct. 2010. URL: http://www.urbandictionary.com/define.php?term=companionable%20silence&defid=5281802 (visited on 2012-09-09). up/downvote ratio 7.5 (cited on page 1).

[513] John Stasko, Todd Miller, Zachary Pousman, Christopher Plaue and Osman Ullah. 'Personalized Peripheral Information Awareness Through Information Art'. In: *UbiComp 2004: Ubiquitous Computing*. Edited by Nigel Davies, Elizabeth D. Mynatt and Itiro Siio. Volume 3205. Lecture Notes in Computer Science. Springer, 2004, pages 18–35. ISBN: 9783540229551. DOI: 10.1007/978-3-540-30119-6_2 (cited on page 84).

[514] Anthony Steed, Mel Slater, Amela Sadagic, Adrian Bullock and Jolanda Tromp. 'Leadership and Collaboration in Shared Virtual Environments'. In: *Proceedings of the IEEE Virtual Reality*. VR '99. Washington, DC, USA: IEEE Computer Society, 1999, pages 112–115. ISBN: 0769500935 (cited on pages 36, 41).

[515] Timothy Stephen. 'Communication and Interdependence in Geographically Separated Relationships'. In: *Human Communication Research* 13.2 (1986), pages 191–210. ISSN: 1468-2958. DOI: 10.1111/j.1468-2958.1986.tb00102.x (cited on page 45).

[516] Michael A. Stephens. 'EDF Statistics for Goodness of Fit and Some Comparisons'. In: *Journal of the American Statistical Association* 69.347 (1974), pages 730–737. DOI: 10.1080/01621459.1974.10480196 (cited on page 153).

[517] Jonathan Steuer. 'Defining Virtual Reality: Dimensions Determining Telepresence'. In: *Journal of Communication* 42.4 (1992), pages 73–93. ISSN: 1460-2466. DOI: 10.1111/j.1460-2466.1992.tb00812.x (cited on page 17).

[518] Joachim Stöber. 'Die Soziale-Erwünschtheits-Skala-17 (SES-17): Entwicklung und erste Befunde zu Reliabilität und Validität'. German. In: *Diagnostica* 45.4 (Oct. 1999), pages 173–177. DOI: 10.1026/0012-1924.45.4.173 (cited on pages 42, 145 sq., 194, 397, 438).

[519] Rob Strong and Bill Gaver. 'Feather, Scent and Shaker: Supporting Simple Intimacy'. In: *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. CSCW '96. 1996, pages 29–30 (cited on pages 92, 97).

[520] J. Ridley Stroop. 'Studies of Interference in Serial Verbal Reactions'. In: *Journal of Experimental Psychology* 18.6 (1935), pages 643–662. ISSN: 0022-1015. DOI: 10.1037/h0054651 (cited on page 74).

[521] Student. 'The Probable Error of a Mean'. In: *Biometrika* 6.1 (Mar. 1908), pages 1–25 (cited on page 162).

[522] Amarnag Subramanya, Alvin Raj, Jeff A. Bilmes and Dieter Fox. 'Recognizing Activities and Spatial Context Using Wearable Sensors'. In: *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*. 2006. arXiv: 1206.6869 [cs.AI] (cited on page 245).

[523] Ron Sun. 'Hybrid Connectionist-Symbolic Models: A Report from the IJCAI'95 Workshop on Connectionist-Symbolic Integration'. In: *Artificial Intelligence Magazine* (1996) (cited on page 257).

[524] Ron Sun and Frédéric Alexandre, editors. *Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, 1997. ISBN: 0805823492 (cited on page 257).

[525] Barbara G. Tabachnick and Linda S. Fidell. *Using Multivariate Statistics*. 4th edition. Boston, MA, USA: Allyn and Bacon, 2001. 966 pages. ISBN: 9780321056771 (cited on page 169).

[526] Emmanuel Munguia Tapia, Stephen S. Intille, William Haskell, Kent Larson, Julie Wright, Abby King and Robert Friedman. 'Real-Time Recognition of Physical Activities and Their Intensities Using Wireless Accelerometers and a Heart Rate Monitor'. In: *Proceedings of the 11th IEEE International Symposium on Wearable Computers*. ISWC '07. Washington, DC, USA: IEEE Computer Society, 2007, pages 1–4. ISBN: 9781424414529. DOI: 10.1109/ISWC.2007.4373774 (cited on page 245).

[527] S. Martin Taylor. 'A Path Model of Aircraft Noise Annoyance'. In: *Journal of Sound Vibration* 96 (Sept. 1984), pages 243–260. DOI: 10.1016/0022-460X(84)90582-0 (cited on page 128).

[528] Kimberly Tee, Saul Greenberg and Carl Gutwin. 'Providing Artifact Awareness to a Distributed Group Through Screen Sharing'. In: *Proceedings of the 20th Anniversary Conference on Computer Supported Cooperative Work*. (Banff, Alberta, Canada). CSCW '06. New York, NY, USA: ACM, 2006, pages 99–108. ISBN: 1595932496. DOI: 10.1145/1180875.1180891 (cited on page 5).

[529] The American Heritage Dictionary of the English Language. *Companionable*. Houghton Mifflin Company. 2009. URL: http://www.thefreedictionary.com/companionable (visited on 2012-09-09) (cited on page 1).

[530] Stefan Thie and Jacoliene van Wijk. 'A General Theory on Presence'. In: *Proceedings of the First International Workshop on Presence*. (Ipswich, UK). 13 May 1998. URL: http://www0.cs.ucl.ac.uk/staff/m.slater/BTWorkshop/KPN/ (visited on 2012-08-29) (cited on page 36).

[531] Alan B. Tickle, Robert Andrews, Mostefa Golea and Joachim Diederich. 'The Truth Will Come to Light: Directions and Challenges in Extracting the Knowledge Embedded Within Trained Artificial Neural Networks'. In: *Neural Networks, IEEE Transactions on* 9.6 (1998), pages 1057–1068 (cited on pages 223, 264).

[532] Assaf Toledo. 'Measuring Absolute Adjectives'. Master's thesis. The Hebrew University of Jerusalem, June 2011.

[533] Konrad Tollmar and Joakim Persson. 'Understanding Remote Presence'. In: *Proceedings of the Second Nordic Conference on Human-Computer Interaction*. (Aarhus, Denmark). NordiCHI '02. New York, NY, USA: ACM, 2002, pages 41–50. ISBN: 1581136161. DOI: 10.1145/572020.572027 (cited on pages 76, 87, 89).

[534] Geoffrey G. Towell and Jude W. Shavlik. 'Extracting Refined Rules from Knowledge-Based Neural Networks'. In: *Machine Learning* 13.1 (1993), pages 71–101. ISSN: 0885-6125. DOI: 10.1023/A:1022683529158 (cited on pages 223, 269).

[535] Geoffrey G. Towell and Jude W. Shavlik. 'Knowledge-Based Artificial Neural Networks'. In: *Artificial Intelligence* 70.1-2 (Oct. 1994), pages 119–165. ISSN: 0004-3702. DOI: 10.1016/0004-3702(94)90105-8 (cited on pages 253, 278).

[536] Leonard H. Tower Jr., Diego Moya, Marek Kosniowski et al. *Google+ : Circles*. Wikipedia. 14 Apr. 2014. URL: https://en.wikipedia.org/w/index.php?title=Google+&oldid=604125852#Circles (cited on page 59).

[537] Jolanda Tromp, Adrian Bullock, Anthony Steed, Amela Sadagic, Mel Slater and Emmanuel Frécon. 'Small Group Behavior Experiments in the Coven Project'. In: *Computer Graphics and Applications* 18.6 (Nov. 1998), pages 53–63. ISSN: 0272-1716. DOI: 10.1109/38.734980 (cited on pages 36, 142, 202).

[538] Remko Tronçon and Peter Saint-Andre. XEP-0146. *Remote Controlling Clients*. 23 Mar. 2006. URL: http://xmpp.org/extensions/xep-0146.html (cited on page 239).

[539] Boris Semyonovich Tsirelson, Michael Hardy, Melcombe et al. *Central Limit Theorem*. Wikipedia. 6 May 2013. URL: http://en.wikipedia.org/w/index.php?title=Central_limit_theorem&oldid=553820285 (visited on 2013-05-16) (cited on pages 178, 281).

[540] Grigorios Tsoumakas, Ioannis Katakis and Ioannis Vlahavas. 'Mining Multi-Label Data'. In: *Data Mining and Knowledge Discovery Handbook*. Edited by Oded Maimon and Lior Rokach. Springer, 2010, pages 667–685. ISBN: 9780387098227. DOI: 10.1007/978-0-387-09823-4_34 (cited on page 253).

[541] Grigorios Tsoumakas, Ioannis Katakis and Ioannis Vlahavas. 'Random k-Labelsets for Multi-Label Classification'. In: *IEEE Transactions on Knowledge and Data Engineering* 23.7 (2011), pages 1079–1089 (cited on page 248).

[542] Grigorios Tsoumakas, Eleftherios Spyromitros-Xioufis, Jozef Vilcek and Ioannis Vlahavas. 'Mulan: A Java Library for Multi-Label Learning'. In: *Journal of Machine Learning Research* 12 (2011), pages 2411–2414 (cited on page 248).

[543] Katherine M. Tsui, Munjal Desai and Holly A. Yanco. 'Towards Measuring the Quality of Interaction: Communication through Telepresence Robots'. In: *Proceedings of the 2012 Performance Metrics for Intelligent Systems Workshop*. (College Park, MD, USA). New York, NY, USA: ACM, 20–22 Mar. 2012. URL: http://robotics.cs.uml.edu/fileadmin/content/publications/2012/31-tsui-etal-permis-measuring-communication.pdf (cited on pages 1, 52).

[544] René Tünnermann, Till Bovermann and Thomas Hermann. 'Auditory Augmentation at Your Fingertips'. In: *Sonic Interaction Design – Exhibition Catalogue*. Edited by Frauke Behrendt and Trond Lossius. Bergen Center for Electronic Arts, 2011, pages 14–17. ISBN: 9788299868006 (cited on page 91).

[545] René Tünnermann, Jan Hammerschmidt and Thomas Hermann. 'Blended Sonification: Sonification for Casual Information Interaction'. In: *Proceedings of the International Conference on Auditory Displays*. (Łódź, Poland). Volume 119–126. ICAD '13. 2013 (cited on pages 91, 299).

[546] René Tünnermann, Christian Leichsenring and Thomas Hermann. 'Direct Tactile Coupling of Mobile Phones with the FEELABUZZ System'. In: *Mobile Social Signal Processing*. Edited by Roderick Murray-Smith. Volume 8045. Lecture Notes in Computer Science. Springer, 2014. Chapter 8, pages 74–83. ISBN: 9783642543241. DOI: 10.1007/978-3-642-54325-8_9 (cited on page 6).

[547] United States Code. *Equal Credit Opportunity Act*. 1974. URL: http://www.law.cornell.edu/uscode/text/15/chapter-41/subchapter-IV. Title 15, chapter 41, subchapter IV (cited on page 261).

[548] Joy K. van Baren and Wijnand A. IJsselsteijn. *Measuring Presence: A Guide to Current Measurement Approaches*. Technical report IST-2001-39237. Version 1.0. Omnibus Presence Technology Assessment and Measurement Groups (OMNIPRES), 23 Mar. 2004. 86 pages. URL: http://www.temple.edu/ispr/p-r_PresenceMeasurement.pdf. European Commission – Information Society Technologies (cited on pages 19, 39 sqq.).

[549] Joy K. van Baren, Wijnand A. IJsselsteijn, Panos Markopoulos, Natalia A. Romero and Boris E. R. de Ruyter. 'Measuring Affective Benefits and Costs of Awareness Systems Supporting Intimate Social Networks'. In: *Proceedings of 3rd Workshop on Social Intelligence Design*. (Enschede, The Netherlands). Edited by T. Nishida. A. Nijholt. Volume 2. CTIT. 2004, pages 13–19. URL: http://alexandria.tue.nl/repository/articles/594140.pdf (cited on pages 28 sq.).

[550] Daniel T. van Bel, Karin C. H. J. Smolders, Wijnand A. IJsselsteijn and Yvonne A. W. de Kort. 'Social Connectedness: Concept and Measurement'. In: *Proceedings of the International Conference on Intelligent Environments*. Amsterdam, The Netherlands: IOS Press, 2009, pages 67–74. URL: http://repository.tue.nl/667830 (cited on pages 42, 145, 155, 157, 195 sq., 215, 400, 419 sqq.).

[551] Pascal W. M. van Gerven, Fred Paas, Jeroen J. G. van Merriënboer and Henk G. Schmidt. 'Modality and Variability as Factors in Training the Elderly'. In: *Applied Cognitive Psychology* 20.3 (2006), pages 311–320. ISSN: 1099-0720. DOI: 10.1002/acp.1247 (cited on page 43).

[552] K. Roger van Horn, Angela Arnone, Kelly Nesbitt, Laura Desilets, Tanya Sears, Michelle Giffin and Rebecca Brudi. 'Physical distance and interpersonal characteristics in college students' romantic relationships'. In: *Personal Relationships* 4.1 (1997), pages 25–34. ISSN: 1475-6811. DOI: 10.1111/j.1475-6811.1997.tb00128.x (cited on page 47).

[553] Sasha Vasko, Bradley T. Hughes, Dominik Vogt, Havoc Pennington, Jeff Raven, Jim Gettys, John Harper, Julian Adams, Matthias Ettrich, Micheal Rogers, Nathan Clemons, Tim Janik et al. 1.4.draft-2. *Extended Window Manager Hints*. 29 Sept. 2006. URL: http://standards.freedesktop.org/wm-spec/wm-spec-1.4.html (cited on page 236).

[554] Frank Vetere, Martin R. Gibbs, Jesper Kjeldskov, Steve Howard, Florian Mueller, Sonja Pedell, Karen Mecoles and Marcus Bunyan. 'Mediating Intimacy: Designing Technologies to Support Strong-Tie Relationships'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '05. Portland, OR, USA: ACM, 2005, pages 471–480. ISBN: 1581139985. DOI: 10.1145/1054972.1055038 (cited on pages 6, 96, 99).

[555] Paul Viola and Michael J. Jones. 'Rapid Object Detection using a Boosted Cascade of Simple Features'. In: *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on* 1 (2001), page 511. ISSN: 1063-6919. DOI: 10.1109/CVPR.2001.990517 (cited on page 231).

[556] Paul Viola and Michael J. Jones. 'Robust Real-Time Face Detection'. In: *International Journal of Computer Vision* 57.2 (1 May 2004), pages 137–154. ISSN: 0920-5691. DOI: 10.1023/B:VISI.0000013087.49260.fb (cited on page 231).

[557] Julia K. Vormbrock. 'Attachment Theory as Applied to Wartime and Job-Related Marital Separation'. In: *Psychological Bulletin* 114.1 (1993), page 122 (cited on page 45).

[558] Alex Waibel and Kai-Fu Lee. *Readings in Speech Recognition*. Morgan Kaufmann, 1990. ISBN: 9781558601246 (cited on page 246).

[559] Willard Waller. *The Family: A Dynamic Interpretation*. New York, NY, USA: The Cordon Company, 1938. 617 pages (cited on page 46).

[560] Jian-hua Wang, Jian-Jiun Ding, Yu Chen and Hsin-Hui Chen. 'Real Time Accelerometer-Based Gait Recognition Using Adaptive Windowed Wavelet Transforms'. In: *IEEE Asia Pacific Conference on Circuits and Systems*. (Kaohsiung, Taiwan, 2–5 Dec. 2012). APCCAS. 2012, pages 591–594. DOI: 10.1109/APCCAS.2012.6419104 (cited on page 246).

[561] Henrik Wannheden et al. *Apple M7*. Wikipedia. 18 Dec. 2013. URL: http://en.wikipedia.org/w/index.php?title=Apple_M7&oldid=586675407 (cited on page 125).

[562] Roy Want and Bill Schilit. 'Expanding the Horizons of Location-Aware Computing'. In: *IEEE Computer* 34.8 (2001), pages 31–34 (cited on page 225).

[563] David Warnock, Marilyn R. McGee-Lennon and Stephen Brewster. 'The Impact of Unwanted Multimodal Notifications'. In: *Proceedings of the 13th international conference on multimodal interfaces*. ICMI '11. Alicante, Spain: ACM, 2011, pages 177–184. ISBN: 9781450306416. DOI: 10.1145/2070481.2070510 (cited on pages 66, 83, 299).

[564] Mark Weiser. 'The Computer for the 21st Century'. In: *Scientific American* 265.3 (1991), pages 94–104 (cited on pages i, 61).

[565] Mark Weiser and John Seely Brown. *Designing Calm Technology*. Technical report. Xerox PARC, 21 Dec. 1995. URL: http://www.ubiq.com/hypertext/weiser/calmtech/calmtech.htm (cited on pages 60 sq.).

[566] Mark Weiser and John Seely Brown. *The Coming Age of Calm Technology*. Technical report. Xerox PARC, 5 Oct. 1996. URL: http://www.johnseelybrown.com/calmtech.pdf (cited on pages 13, 60 sq., 63, 65).

[567] Eric W. Weisstein, editor. *Bell Number*. Wolfram Research. 2013. URL: http://mathworld.wolfram.com/BellNumber.html (visited on 2013-03-27) (cited on page 174).

[568] B. L. Welch. 'The Generalization of "Student's" Problem When Several Different Population Variances are Involved'. In: *Biometrika* 34.1-2 (1947), pages 28–35. DOI: 10.1093/biomet/34.1-2.28 (cited on pages 178 sq., 183).

[569] Stefan Wermter and Ron Sun. 'An Overview of Hybrid Neural Systems'. In: *Hybrid Neural Systems*. Edited by Stefan Wermter and Ron Sun. Springer, 2000, pages 1–13. ISBN: 9783540673057 (cited on pages 256 sqq.).

[570] Julia Werner, Reto Wettach and Eva Hornecker. 'United-Pulse: Feeling Your Partner's Pulse'. In: *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services*. ACM. 2008, pages 535–538 (cited on pages 7, 93 sq.).

[571] John S. Westefeld and Debora Liddell. 'Coping with Long-Distance Relationships'. In: *Journal of College Student Personnel* 23.6 (Nov. 1982), pages 550–551 (cited on page 48).

[572] Denise Whitelock, Daniela M. Romano, Anne Jelfs and Paul Brna. 'Perfect Presence: What Does This Mean for the Design of Virtual Learning Environments?' In: *Education and Information Technologies* 5 (4 2000), pages 277–289. ISSN: 1360-2357. DOI: 10.1023/A:1012001523715 (cited on pages 36, 41).

[573] Steve Whittaker, David Frohlich and Owen Daly-Jones. 'Informal Workplace Communication: What is it Like and How Might We Support it?' In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating Interdependence*. CHI '94. Boston, MA, USA: ACM, 1994, pages 131–137. ISBN: 0897916506. DOI: 10.1145/191666.191726 (cited on pages 37, 64).

[574] Josef Wideström, Ann-Sofie Axelsson, Ralph Schroeder, Alexander Nilsson, Ilona Heldal and Åsa Abelin. 'The Collaborative Cube Puzzle: A Comparison of Virtual and Real Environments'. In: *Proceedings of the 3rd International Conference on Collaborative Virtual Environments*. CVE '00. San Francisco, CA, USA: ACM, 2000, pages 165–171. ISBN: 1581133030. DOI: 10.1145/351006.351035 (cited on pages 36, 41).

[575] Morten Wiener and Albert Mehrabian. *Language Within Language: Immediacy, a Channel in Verbal Communication*. The Century Psychology Series. New York, NY, USA: Appleton-Century-Crofts, 1968. 214 pages (cited on page 30).

[576] Justin Wilcox. *Closer ToGetHer: For Long Distance Loving*. Feb. 2012. URL: http://www.getclosertogether.com/ (visited on 2012-09-05) (cited on page 96).

[577] Frank Wilcoxon. 'Individual Comparisons by Ranking Methods'. In: *Biometrics Bulletin* 1.6 (Dec. 1945), pages 80–83 (cited on page 162).

[578] Tad Williams. *Otherland*. 4 volumes. 1996–2001 (cited on page 219).

[579] John Williamson, Roderick Murray-Smith and Stephen Hughes. 'Shoogle: Excitatory Multimodal Interaction on Mobile Devices'. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '07. San Jose, CA, USA: ACM, 2007, pages 121–124. ISBN: 9781595935939. DOI: 10.1145/1240624.1240642 (cited on page 100).

[580] Benjamin Wing, Ramin Charles Nakisa et al. *Metropolis–Hastings algorithm*. Wikipedia. 3 June 2013. URL: http://en.wikipedia.org/w/index.php?title=Metropolis%C3%A2%C2%80%C2%93Hastings_algorithm&oldid=558099618 (visited on 2013-06-16) (cited on page 288).

[581] Wendy Wood. 'Attitude Change: Persuasion and Social Influence'. In: *Annual Review of Psychology* 51.1 (2000), pages 539–570. DOI: 10.1146/annurev.psych.51.1.539 (cited on page 73).

[582] Matthew Wright and Adrian Freed. *Open Sound Control: A New Protocol for Communicating with Sound Synthesizers*. 1997 (cited on page 116).

[583] Jun Yang. 'Toward Physical Activity Diary: Motion Recognition Using Simple Acceleration Features with Mobile Phones'. In: *Proceedings of the 1st International Workshop on Interactive Multimedia for Consumer Electronics*. IMCE '09. Beijing, China: ACM, 2009, pages 1–10. ISBN: 9781605587585. DOI: 10.1145/1631040.1631042 (cited on page 245).

[584] Svetlana Yarosh. *Designing a Survey Instrument for Evaluating Communication Technologies*. Georgia Institute of Technology. 22 June 2011. URL: http://home.cc.gatech.edu/lana/31 (visited on 2012-09-04) (cited on pages 42, 145, 159, 425 sqq., 436).

[585] Christine Youngblut. *Experience of Presence in Virtual Environments*. Technical report. Alexandria, VA, USA: Institute for Defense Analyses, Sept. 2003. URL: http://handle.dtic.mil/100.2/ADA427495 (cited on pages 22, 36, 39, 41 sq.).

[586] Christine Youngblut and Odette Huie. 'The Relationship Between Presence and Performance in Virtual Environments: Results of a VERTS Study'. In: *Proceedings of the IEEE Virtual Reality 2003*. VR '03. Washington, DC, USA: IEEE Computer Society, 2003, pages 277–278. ISBN: 0769518826 (cited on page 39).

[587] Min-Ling Zhang and Zhi-Hua Zhou. 'Multi-Label Neural Networks with Applications to Functional Genomics and Text Categorization'. In: *IEEE Transactions on Knowledge and Data Engineering* 18 (2006), pages 1338–1351 (cited on page 248).

[588] Min-Ling Zhang and Zhi-Hua Zhou. 'ML-kNN: A Lazy Learning Approach to Multi-Label Learning'. In: *Pattern Recognition* 40.7 (July 2007), pages 2038–2048. ISSN: 0031-3203. DOI: 10.1016/j.patcog.2006.12.019 (cited on pages 248, 253).

[589] Xin Zhang and W. Ras Zbigniew. 'Analysis of Sound Features for Music Timbre Recognition'. In: *Proceedings of the 2007 International Conference on Multimedia and Ubiquitous Engineering*. MUE '07. Washington, DC, USA: IEEE Computer Society, 2007, pages 3–8. ISBN: 0769527779. DOI: 10.1109/MUE.2007.85 (cited on page 232).

[590] Qun Zhao and Jose C. Principe. 'Improving ATR Performance by Incorporating Virtual Negative Examples'. In: *International Joint Conference on Neural Networks*. Volume 5. IJCNN '99. 1999, pages 3198–3203. DOI: 10.1109/IJCNN.1999.836166 (cited on page 278).

[591] Shanyang Zhao. 'Toward a Taxonomy of Copresence'. In: *Presence: Teleoperators and Virtual Environments* 12.5 (Oct. 2003), pages 445–455. ISSN: 1054-7460. DOI: 10.1162/105474603322761261 (cited on pages 22 sq., 52, 76 sq.).

[592] Shenghuo Zhu, Xiang Ji, Wei Xu and Yihong Gong. 'Multi-Labelled Classification Using Maximum Entropy Method'. In: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '05. Salvador, Brazil: ACM, 2005, pages 274–281. ISBN: 1595930345. DOI: 10.1145/1076034.1076082 (cited on page 253).

[593] Sina Zulkernain, Praveen Madiraju and Sheikh Iqbal Ahamed. 'A Context Aware Interruption Management System for Mobile Devices'. In: *Mobile Wireless Middleware, Operating Systems, and Applications*. Springer, 2010, pages 221–234. ISBN: 9783642177583. DOI: 10.1007/978-3-642-17758-3_16 (cited on pages 219, 227, 293).

[594] Bob Zurunkel. *Plausible Deniability*. Urban Dictionary. 17 Oct. 2005. URL: http://web.archive.org/web/20051224150346/http://www.urbandictionary.com/define.php?term=plausible+deniability&defid=1484799 (visited on 2012-06-12). up/downvote ratio 7.0 (cited on page 72).

In the PDF version of this document, all titles in the bibliography above are hyperlinks. It should be noted, however, that all links pointing to Google Scholar were created automatically from publication title, first author and year of publication. This is meant as a convenience for the reader when searching for those references in this work that do not have a trustworthy and freely available web location that could also be trusted to remain available indefinitely. There is no guarantee for these automatically generated links to point to the right publication or even yield any results at all.

Acknowledgements

Putting only one name on the cover of this work, as isolated as the island that no man is, is as inappropriate as it is required. Therefore I am grateful for this opportunity to thank at least some of the many people who supported me during this time. While I wrote the words, without them I could never even have dreamt to accomplish any of this.

First of all, I wish to thank my advisors who enabled this work in the first place. Marc Hanheide always had a good sense for the big picture and for the way forward, nudging me there when necessary. Thomas Hermann was always not only a great source of expertise but also of unconventional ideas.

My heartfelt thanks also go to the people who were both colleagues and friends to me. The discussions with René Tünnerman, whose door was always open, even when it was not, were indispensable in both professional and personal matters and I hope nobody ever even does a back-of-the-envelope calculation of how many hours I actually made use of this generosity. Thank you also to Sebastian Hammerl for being a great office mate, never hesitating to help solve a problem or lend an ear. The same goes for Ulf Großekathöfer who had a rare combination of qualities by being fun, a competent machine learner, having a good taste in movies and using little rainbow squiggles as a science tool. I would also like to thank Till Bovermann amongst many things first of all for making our short stay in Helsinki an unforgettable holiday and an intensely productive time all at once.

Both the Ambient Intelligence and the Applied Informatics group have been providing a working atmosphere that few people had the pleasure to experience. Never did I feel anything but welcome and appreciated there. Unfortunately, I cannot address everyone who deserves it without bloating this already thick tome much further.

Friederike Eyssel was helpful far beyond what could be expected and taught me more about statistics than I ever thought there would be to learn.

I do not know how Franz Kummert managed to always keep an open door for me with the schedule he got but I am very thankful that he did.

I would like to thank the DFG and the CITEC for making this work possible and I would especially like to thank the people that made the CITEC more than an institution to me – Claudia Muhl especially, who always lifts the spirit of everyone around her. I also want to thank every one of the many graduate students for making this assortment of disciplines a close-nit group of trusty colleagues.

More thanks go out to Angelika Dierker, Marcel Martin, Sarah Ruthe and Steve Wolter for their ideas, suggestions, keen eyes and amicable support.

Thank you also to the student assistants, bachelor's, master's and diploma candidates, and student project members who I have had the pleasure to supervise and work with, although not all of them on the things covered in this thesis. These are Ralph Welsch, David Fleer, Matthias Heinrich, Nils Wöhler, Dimitri Heil, Sebastian Nitsch, Oliver Metz, Janik Jaskolski, Tobias Rodehutskors, Fabian Vöhl, Sebastian Kahl and Falk Olaf Altheide. Thanks also to Alexander Neumann and Bastian Kriesten for whom I had almost forgotten that they belong in this category, too, even though they have since wandered into the arguably and paradoxically even more important "undeservedly unmentioned" part of colleagues and friends above.

A group I will not name in person but to whom I am very grateful are the participants of the two studies, especially the participants of the upstairs study who tolerated that intrusion into their homes over weeks, without having known in advance that they would come to like it and who tolerated the endless questionnaires, without having known in advance quite how many redundant items they would have to answer.

I also want to thank all my family for their support and patience but most of all Andreas Hahn for being a great mentor to me, sharing his valuable experience and council and my mother for working through this difficult time with me and for being so understanding and supportive.

Last yet foremost, I want to express my deep gratitude to my wife Daniela, who supported me beyond measure in every possible way and when I initially said that all of these people here are proof that I could not have done this alone then it is her to whom this applies most of all.

Appendices

A Questionnaires

In the following, all the questionnaires used in the study described in Chapter 6 are shown. Since the study was done in German, the questionnaires are in German, too. The original questionnaires were dynamic in that the male or female form of pronouns and nouns was used after entering the code name that determined the gender of the participant (for the upstairs questionnaires) or after entering a value for the gender item (for FEELABUZZ). The dynamic gender feature had not been implemented in the very first questionnaire (Q_α) and two different static versions were sent to male and female participants separately (cf. Figure A.9 on page 395).

We did not re-typeset the online questionnaires or insert any page breaks for this overview. Therefore the print can be quite small for some pages. We would like to excuse this and kindly refer the reader to the electronic version of this document where all the following figures can be scaled up indefinitely.

Most scales were randomised but this randomisation was deactivated for this print version.[1] Questions marked with a red asterisk (*) were mandatory.

A few of the mandatory questions were only displayed when a previous question had been answered in a certain way (for example after checking an "other" item, a mandatory text field appeared asking

[1] For technical reasons, a copy of the original questionnaires was made for the non-randomised print versions. This is the reason for the appearance of the work *Druckversion* (print version) in the headers of the questionnaires shown here. This word was not found on the original questionnaires.

Druckversion Begleitender Fragebogen Upstairs ohne System

Dieser Fragebogen sollte zweitägig ausgefüllt werden, solange noch kein System vor Ort ist.

Hallo liebe Versuchsteilnehmerin!
Bitte fülle die folgende Seite mit Fragen aus und vergiss nicht, auf „Absenden" zu klicken.

Diese Umfrage enthält 9 Fragen.

Figure A.1. Q_α page 1: welcome screen

for specifics). In the static versions shown here, all of these conditional items are displayed without further indication but it will hopefully be apparent from the context that they had not to be filled in by each participant.

For more information on the questionnaire design, see Section 2.4.1; for information on the study design itself, refer to Chapter 6.

A.1 upstairs

Q_α was filled in every second day during the baseline period (i. e. without using the system). Q'_α was filled in every second day while using the system. Q_β was filled in right before the start of the system use period. Q_γ was filled in after one week of system use. $Q_{\beta\gamma}$ was filled in at the very end of every trial.

There was one short questionnaire that was used to collect some demographic data without compromising the identity of the participants. It could be filled in at any point during the trials period.

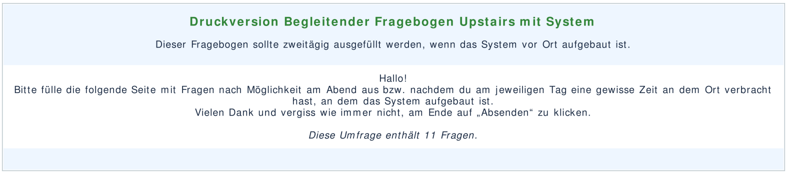
Druckversion Begleitender Fragebogen Upstairs mit System

Dieser Fragebogen sollte zweitägig ausgefüllt werden, wenn das System vor Ort aufgebaut ist.

Hallo!
Bitte fülle die folgende Seite mit Fragen nach Möglichkeit am Abend aus bzw. nachdem du am jeweiligen Tag eine gewisse Zeit an dem Ort verbracht hast, an dem das System aufgebaut ist.
Vielen Dank und vergiss wie immer nicht, am Ende auf „Absenden" zu klicken.

Diese Umfrage enthält 11 Fragen.

Figure A.2. Q'_α page 1: welcome screen

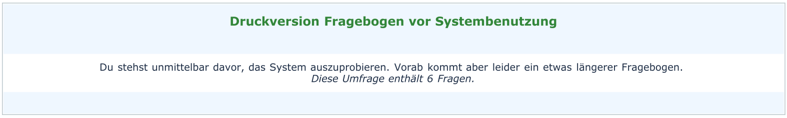
Druckversion Fragebogen vor Systembenutzung

Du stehst unmittelbar davor, das System auszuprobieren. Vorab kommt aber leider ein etwas längerer Fragebogen.
Diese Umfrage enthält 6 Fragen.

Figure A.3. Q_β page 1: welcome screen

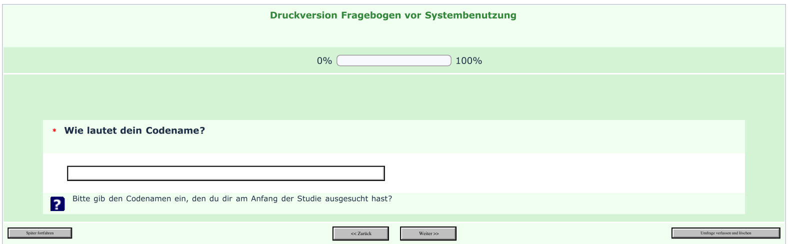
Druckversion Fragebogen vor Systembenutzung

0% 100%

* **Wie lautet dein Codename?**

? Bitte gib den Codenamen ein, den du dir am Anfang der Studie ausgesucht hast?

Später fortfahren | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.4. Q_β page 2: code name prompt

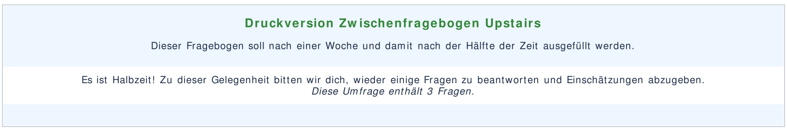
Druckversion Zwischenfragebogen Upstairs

Dieser Fragebogen soll nach einer Woche und damit nach der Hälfte der Zeit ausgefüllt werden.

Es ist Halbzeit! Zu dieser Gelegenheit bitten wir dich, wieder einige Fragen zu beantworten und Einschätzungen abzugeben.
Diese Umfrage enthält 3 Fragen.

Figure A.5. Q_γ page 1: welcome screen

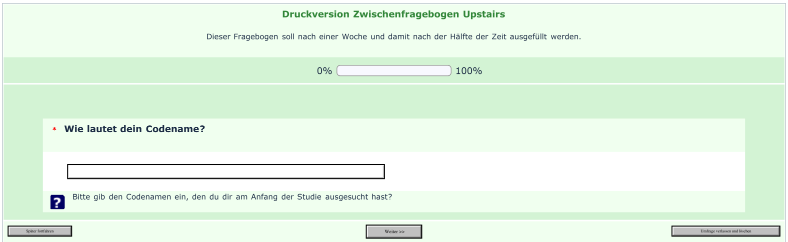

Figure A.6. Q_γ page 2: code name prompt

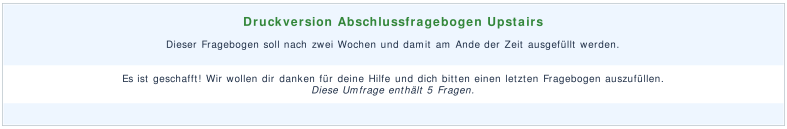

Figure A.7. $Q_{\beta\gamma}$ page 1: welcome screen

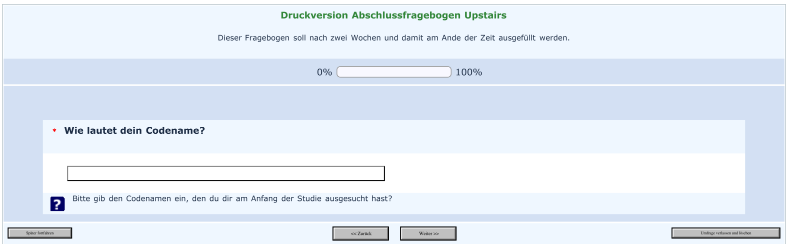

Figure A.8. $Q_{\beta\gamma}$ page 2: coden name prompt

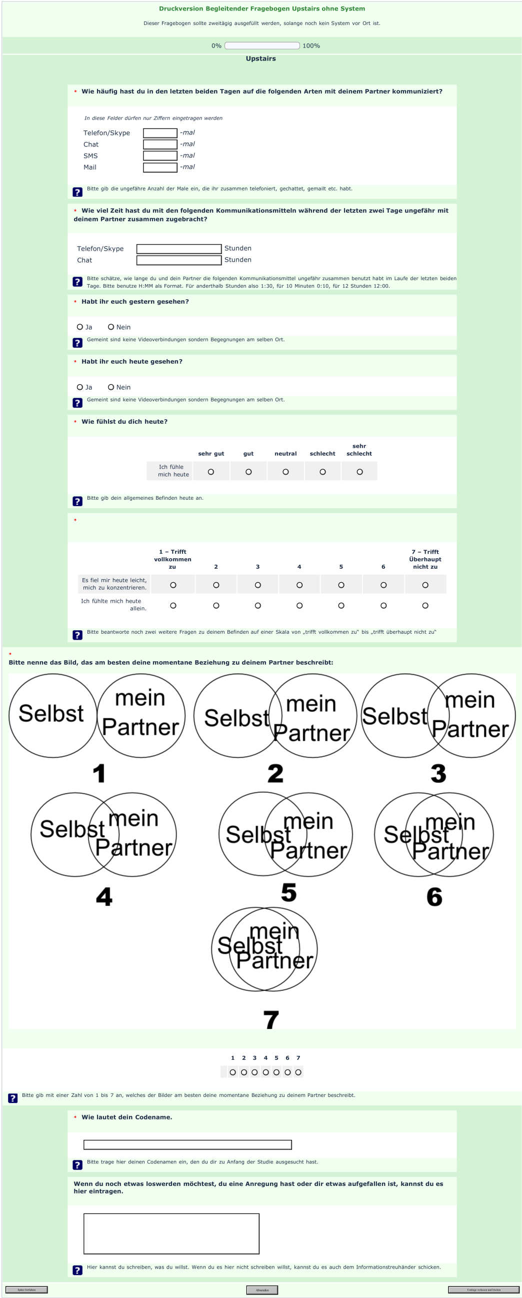

Druckversion Begleitender Fragebogen Upstairs ohne System

Dieser Fragebogen sollte zweitägig ausgefüllt werden, solange noch kein System vor Ort ist.

0% 100%

Upstairs

* **Wie häufig hast du in den letzten beiden Tagen auf die folgenden Arten mit deinem Partner kommuniziert?**

In diese Felder dürfen nur Ziffern eingetragen werden

Telefon/Skype *-mal*
Chat *-mal*
SMS *-mal*
Mail *-mal*

? Bitte gib die ungefähre Anzahl der Male ein, die ihr zusammen telefoniert, gechattet, gemailt etc. habt.

* **Wie viel Zeit hast du mit den folgenden Kommunikationsmitteln während der letzten zwei Tage ungefähr mit deinem Partner zusammen zugebracht?**

Telefon/Skype Stunden
Chat Stunden

? Bitte schätze, wie lange du und dein Partner die folgenden Kommunikationsmittel ungefähr zusammen benutzt habt im Laufe der letzten beiden Tage. Bitte benutze H:MM als Format. Für anderthalb Stunden also 1:30, für 10 Minuten 0:10, für 12 Stunden 12:00.

* **Habt ihr euch gestern gesehen?**

O Ja O Nein

? Gemeint sind keine Videoverbindungen sondern Begegnungen am selben Ort.

* **Habt ihr euch heute gesehen?**

O Ja O Nein

? Gemeint sind keine Videoverbindungen sondern Begegnungen am selben Ort.

* **Wie fühlst du dich heute?**

	sehr gut	gut	neutral	schlecht	sehr schlecht
Ich fühle mich heute	O	O	O	O	O

? Bitte gib dein allgemeines Befinden heute an.

*

	1 – Trifft vollkommen zu	2	3	4	5	6	7 – Trifft Überhaupt nicht zu
Es fiel mir heute leicht, mich zu konzentrieren.	O	O	O	O	O	O	O
Ich fühlte mich heute allein.	O	O	O	O	O	O	O

? Bitte beantworte noch zwei weitere Fragen zu deinem Befinden auf einer Skala von „trifft vollkommen zu" bis „trifft überhaupt nicht zu"

*
Bitte nenne das Bild, das am besten deine momentane Beziehung zu deinem Partner beschreibt:

1	2	3	4	5	6	7
O	O	O	O	O	O	O

? Bitte gib mit einer Zahl von 1 bis 7 an, welches der Bilder am besten deine momentane Beziehung zu deinem Partner beschreibt.

* **Wie lautet dein Codename.**

? Bitte trage hier deinen Codenamen ein, den du dir zu Anfang der Studie ausgesucht hast.

Wenn du noch etwas loswerden möchtest, du eine Anregung hast oder dir etwas aufgefallen ist, kannst du es hier eintragen.

? Hier kannst du schreiben, was du willst. Wenn du es hier nicht schreiben willst, kannst du es auch dem Informationstreuhänder schicken.

Später fortfahren Absenden Umfrage verlassen und löschen

Figure A.9. Q_α page 2: main group

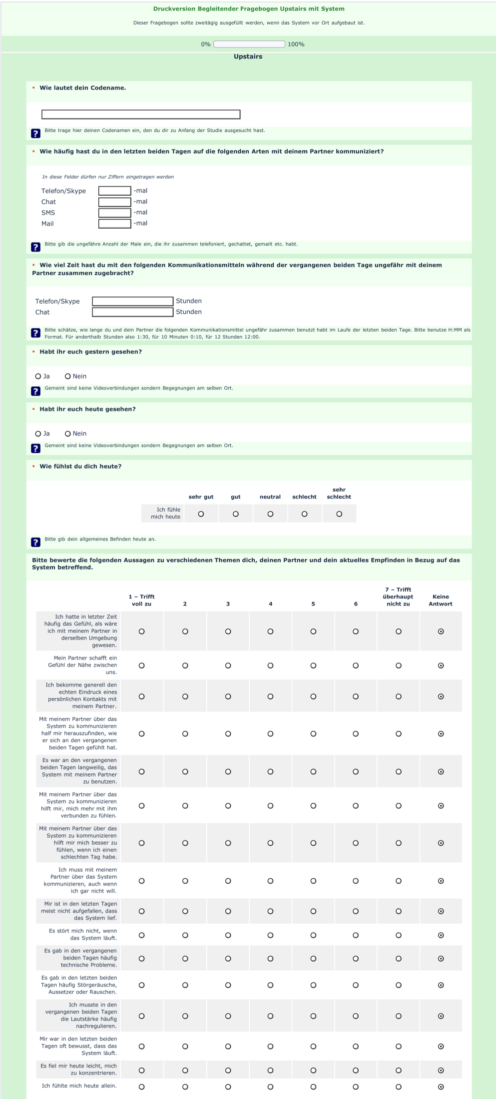

Druckversion Begleitender Fragebogen Upstairs mit System

Dieser Fragebogen sollte zweitägig ausgefüllt werden, wenn das System vor Ort aufgebaut ist.

0% 100%

Upstairs

* **Wie lautet dein Codename.**

? Bitte trage hier deinen Codenamen ein, den du dir zu Anfang der Studie ausgesucht hast.

* **Wie häufig hast du in den letzten beiden Tagen auf die folgenden Arten mit deinem Partner kommuniziert?**

In diese Felder dürfen nur Ziffern eingetragen werden

Telefon/Skype -mal
Chat -mal
SMS -mal
Mail -mal

? Bitte gib die ungefähre Anzahl der Male ein, die ihr zusammen telefoniert, gechattet, gemailt etc. habt.

* **Wie viel Zeit hast du mit den folgenden Kommunikationsmitteln während der vergangenen beiden Tage ungefähr mit deinem Partner zusammen zugebracht?**

Telefon/Skype Stunden
Chat Stunden

? Bitte schätze, wie lange du und dein Partner die folgenden Kommunikationsmittel ungefähr zusammen benutzt habt im Laufe der letzten beiden Tage. Bitte benutze H:MM als Format. Für anderthalb Stunden also 1:30, für 10 Minuten 0:10, für 12 Stunden 12:00.

* **Habt ihr euch gestern gesehen?**

○ Ja ○ Nein

? Gemeint sind keine Videoverbindungen sondern Begegnungen am selben Ort.

* **Habt ihr euch heute gesehen?**

○ Ja ○ Nein

? Gemeint sind keine Videoverbindungen sondern Begegnungen am selben Ort.

* **Wie fühlst du dich heute?**

	sehr gut	gut	neutral	schlecht	sehr schlecht
Ich fühle mich heute	○	○	○	○	○

? Bitte gib dein allgemeines Befinden heute an.

Bitte bewerte die folgenden Aussagen zu verschiedenen Themen dich, deinen Partner und dein aktuelles Empfinden in Bezug auf das System betreffend.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu	Keine Antwort
Ich hatte in letzter Zeit häufig das Gefühl, als wäre ich mit meinem Partner in derselben Umgebung gewesen.	○	○	○	○	○	○	○	⊙
Mein Partner schafft ein Gefühl der Nähe zwischen uns.	○	○	○	○	○	○	○	⊙
Ich bekomme generell den echten Eindruck eines persönlichen Kontakts mit meinem Partner.	○	○	○	○	○	○	○	⊙
Mit meinem Partner über das System zu kommunizieren half mir herauszufinden, wie er sich an den vergangenen beiden Tagen gefühlt hat.	○	○	○	○	○	○	○	⊙
Es war an den vergangenen beiden Tagen langweilig, das System mit meinem Partner zu benutzen.	○	○	○	○	○	○	○	⊙
Mit meinem Partner über das System zu kommunizieren hilft mir, mich mehr mit ihm verbunden zu fühlen.	○	○	○	○	○	○	○	⊙
Mit meinem Partner über das System zu kommunizieren hilft mir mich besser zu fühlen, wenn ich einen schlechten Tag habe.	○	○	○	○	○	○	○	⊙
Ich muss mit meinem Partner über das System kommunizieren, auch wenn ich gar nicht will.	○	○	○	○	○	○	○	⊙
Mir ist in den letzten Tagen meist nicht aufgefallen, dass das System lief.	○	○	○	○	○	○	○	⊙
Es stört mich nicht, wenn das System läuft.	○	○	○	○	○	○	○	⊙
Es gab in den vergangenen beiden Tagen häufig technische Probleme.	○	○	○	○	○	○	○	⊙
Es gab in den letzten beiden Tagen häufig Störgeräusche, Aussetzer oder Rauschen.	○	○	○	○	○	○	○	⊙
Ich musste in den vergangenen beiden Tagen die Lautstärke häufig nachregulieren.	○	○	○	○	○	○	○	⊙
Mir war in den letzten beiden Tagen oft bewusst, dass das System läuft.	○	○	○	○	○	○	○	⊙
Es fiel mir heute leicht, mich zu konzentrieren.	○	○	○	○	○	○	○	⊙
Ich fühlte mich heute allein.	○	○	○	○	○	○	○	⊙

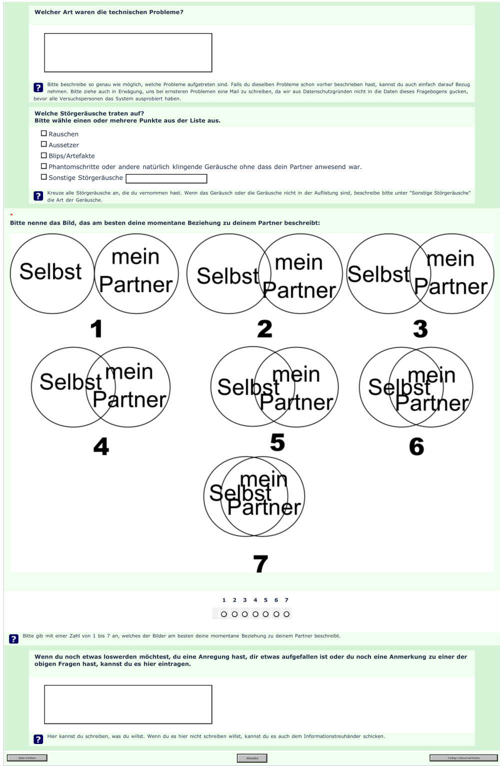

Welcher Art waren die technischen Probleme?

? Bitte beschreibe so genau wie möglich, welche Probleme aufgetreten sind. Falls du dieselben Probleme schon vorher beschrieben hast, kannst du auch einfach darauf Bezug nehmen. Bitte ziehe auch in Erwägung, uns bei ernsteren Problemen eine Mail zu schreiben, da wir aus Datenschutzgründen nicht in die Daten dieses Fragebogens gucken, bevor alle Versuchspersonen das System ausprobiert haben.

Welche Störgeräusche traten auf?
Bitte wähle einen oder mehrere Punkte aus der Liste aus.

☐ Rauschen
☐ Aussetzer
☐ Blips/Artefakte
☐ Phantomschritte oder andere natürlich klingende Geräusche ohne dass dein Partner anwesend war.
☐ Sonstige Störgeräusche

? Kreuze alle Störgeräusche an, die du vernommen hast. Wenn das Geräusch oder die Geräusche nicht in der Auflistung sind, beschreibe bitte unter "Sonstige Störgeräusche" die Art der Geräusche.

*
Bitte nenne das Bild, das am besten deine momentane Beziehung zu deinem Partner beschreibt:

1	2	3	4	5	6	7
○	○	○	○	○	○	○

? Bitte gib mit einer Zahl von 1 bis 7 an, welches der Bilder am besten deine momentane Beziehung zu deinem Partner beschreibt.

Wenn du noch etwas loswerden möchtest, du eine Anregung hast, dir etwas aufgefallen ist oder du noch eine Anmerkung zu einer der obigen Fragen hast, kannst du es hier eintragen.

? Hier kannst du schreiben, was du willst. Wenn du es hier nicht schreiben willst, kannst du es auch dem Informationstreuhänder schicken.

Später fortfahren | Absenden | Umfrage verlassen und löschen

Figure A.10. Q'_α page 2: main group

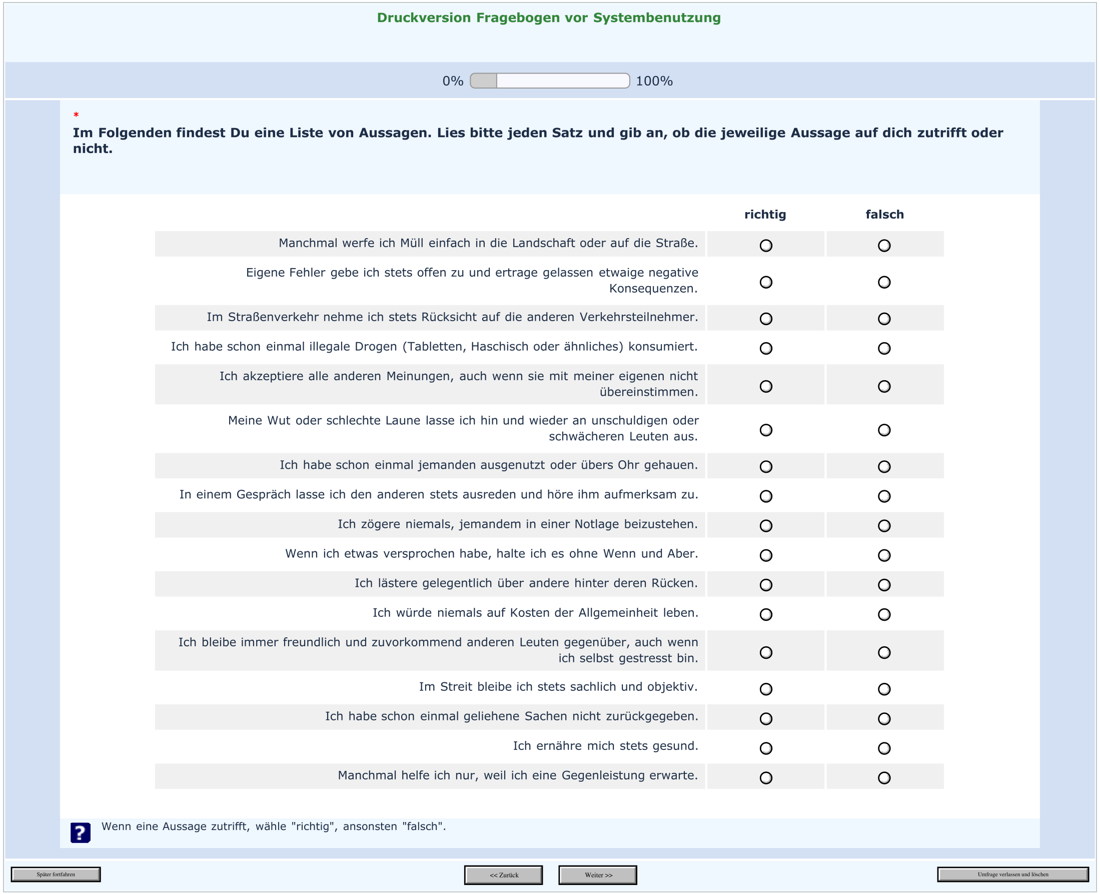
Druckversion Fragebogen vor Systembenutzung

0% 100%

*
Im Folgenden findest Du eine Liste von Aussagen. Lies bitte jeden Satz und gib an, ob die jeweilige Aussage auf dich zutrifft oder nicht.

	richtig	falsch
Manchmal werfe ich Müll einfach in die Landschaft oder auf die Straße.	○	○
Eigene Fehler gebe ich stets offen zu und ertrage gelassen etwaige negative Konsequenzen.	○	○
Im Straßenverkehr nehme ich stets Rücksicht auf die anderen Verkehrsteilnehmer.	○	○
Ich habe schon einmal illegale Drogen (Tabletten, Haschisch oder ähnliches) konsumiert.	○	○
Ich akzeptiere alle anderen Meinungen, auch wenn sie mit meiner eigenen nicht übereinstimmen.	○	○
Meine Wut oder schlechte Laune lasse ich hin und wieder an unschuldigen oder schwächeren Leuten aus.	○	○
Ich habe schon einmal jemanden ausgenutzt oder übers Ohr gehauen.	○	○
In einem Gespräch lasse ich den anderen stets ausreden und höre ihm aufmerksam zu.	○	○
Ich zögere niemals, jemandem in einer Notlage beizustehen.	○	○
Wenn ich etwas versprochen habe, halte ich es ohne Wenn und Aber.	○	○
Ich lästere gelegentlich über andere hinter deren Rücken.	○	○
Ich würde niemals auf Kosten der Allgemeinheit leben.	○	○
Ich bleibe immer freundlich und zuvorkommend anderen Leuten gegenüber, auch wenn ich selbst gestresst bin.	○	○
Im Streit bleibe ich stets sachlich und objektiv.	○	○
Ich habe schon einmal geliehene Sachen nicht zurückgegeben.	○	○
Ich ernähre mich stets gesund.	○	○
Manchmal helfe ich nur, weil ich eine Gegenleistung erwarte.	○	○

? Wenn eine Aussage zutrifft, wähle "richtig", ansonsten "falsch".

Später fortfahren | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.11. Q_β page 3: social desirability scale [518]

Druckversion Fragebogen vor Systembenutzung

0% 100%

Selbsteinschätzung und Einschätzung des anderen.

* **Im Folgenden findest du einige Aussagen sowohl über dich als auch über deinen Partner. Bitte lies sie dir sorgfältig durch und gib zu jeder Aussage an, inwieweit du findest, dass sie zutrifft.**

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich finde, dass es ziemlich leicht für mich ist, anderen gefühlsmäßig nahe zu sein. Es geht mir gut, wenn ich mich auf andere verlassen kann und wenn andere sich auf mich verlassen. Ich mache mir keine Gedanken darüber, dass ich allein sein könnte oder dass andere mich nicht akzeptieren könnten.	○	○	○	○	○	○	○
Ich finde, dass es ziemlich leicht für meinen Partner ist, anderen gefühlsmäßig nahe zu sein. Es geht ihm gut, wenn er sich auf andere verlassen kann und wenn andere sich auf ihn verlassen. Er macht sich keine Gedanken darüber, dass er allein sein könnte oder dass andere ihn nicht akzeptieren könnten.	○	○	○	○	○	○	○
Es geht mir auch ohne enge gefühlsmäßige Bindung gut. Es ist sehr wichtig für mich, mich unabhängig und selbständig zu fühlen und ich ziehe es vor, wenn ich nicht von anderen und andere nicht von mir abhängig sind.	○	○	○	○	○	○	○
Es geht meinem Partner auch ohne enge gefühlsmäßige Bindung gut. Es ist sehr wichtig für ihn, sich unabhängig und selbständig zu fühlen und er zieht es vor, wenn er nicht von anderen und andere nicht von ihm abhängig sind.	○	○	○	○	○	○	○
Ich möchte anderen gefühlsmäßig sehr nahe sein, aber ich merke oft, dass andere Widerstände dagegen errichten, mir so nahe zu sein, wie ich ihnen nahe sein möchte. Es geht mir sehr gut, wenn ich ohne enge Beziehung bin, aber ich denke manchmal, dass andere mich nicht so sehr schätzen wie ich sie.	○	○	○	○	○	○	○
Mein Partner möchte anderen gefühlsmäßig sehr nahe sein, aber er merkt oft, dass andere Widerstände dagegen errichten, ihm so nahe zu sein, wie er ihnen nahe sein möchte. Es geht ihm sehr gut, wenn er ohne enge Beziehung ist, aber er denkt manchmal, dass andere ihn nicht so sehr schätzen wie er sie.	○	○	○	○	○	○	○
Ich empfinde es manchmal als ziemlich unangenehm, anderen nahe zu sein. Ich möchte Beziehungen, in denen ich anderen nahe bin, aber ich finde es schwierig, ihnen vollständig zu vertrauen oder von ihnen abhängig zu sein. Ich fürchte manchmal, dass ich verletzt werde, wenn ich mir erlaube, anderen zu nahe zu kommen.	○	○	○	○	○	○	○
Mein Partner empfindet es manchmal als ziemlich unangenehm, anderen nahe zu sein. Er möchte Beziehungen, in denen er anderen nahe ist, aber er findet es schwierig, ihnen vollständig zu vertrauen oder von ihnen abhängig zu sein. Er fürchtet manchmal, dass er verletzt wird, wenn er sich erlaubt, anderen zu nahe zu kommen.	○	○	○	○	○	○	○

Figure A.12. Q_β page 4: relationship questionnaire [26, 118]

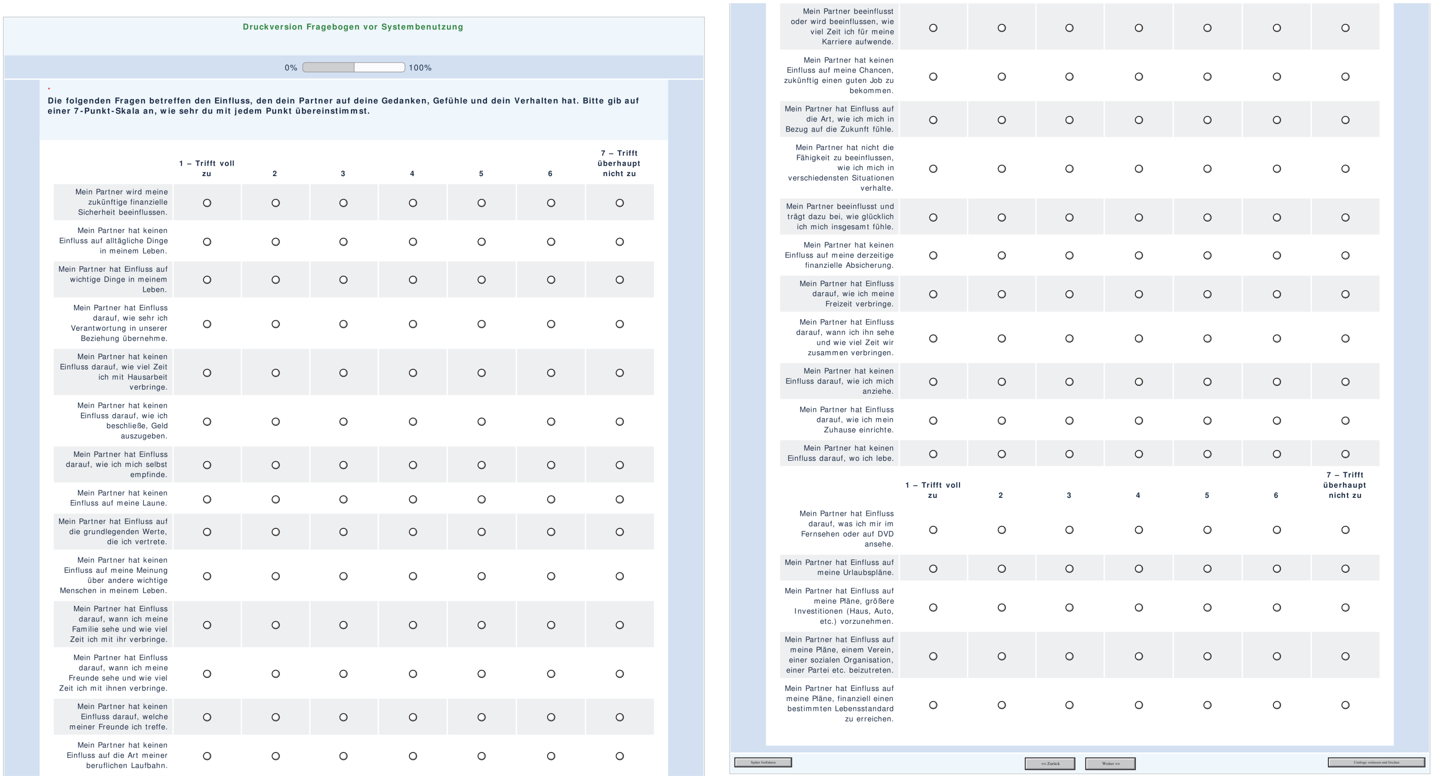

Druckversion Fragebogen vor Systembenutzung

0% 100%

Die folgenden Fragen betreffen den Einfluss, den dein Partner auf deine Gedanken, Gefühle und dein Verhalten hat. Bitte gib auf einer 7-Punkt-Skala an, wie sehr du mit jedem Punkt übereinstimmst.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner wird meine zukünftige finanzielle Sicherheit beeinflussen.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf alltägliche Dinge in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf wichtige Dinge in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie sehr ich Verantwortung in unserer Beziehung übernehme.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie viel Zeit ich mit Hausarbeit verbringe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie ich beschließe, Geld auszugeben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich mich selbst empfinde.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Laune.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf die grundlegenden Werte, die ich vertrete.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Meinung über andere wichtige Menschen in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich meine Familie sehe und wie viel Zeit ich mit ihr verbringe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich meine Freunde sehe und wie viel Zeit ich mit ihnen verbringe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, welche meiner Freunde ich treffe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf die Art meiner beruflichen Laufbahn.	○	○	○	○	○	○	○
Mein Partner beeinflusst oder wird beeinflussen, wie viel Zeit ich für meine Karriere aufwende.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Chancen, zukünftig einen guten Job zu bekommen.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf die Art, wie ich mich in Bezug auf die Zukunft fühle.	○	○	○	○	○	○	○
Mein Partner hat nicht die Fähigkeit zu beeinflussen, wie ich mich in verschiedensten Situationen verhalte.	○	○	○	○	○	○	○
Mein Partner beeinflusst und trägt dazu bei, wie glücklich ich mich insgesamt fühle.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine derzeitige finanzielle Absicherung.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich meine Freizeit verbringe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich ihn sehe und wie viel Zeit wir zusammen verbringen.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie ich mich anziehe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich mein Zuhause einrichte.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wo ich lebe.	○	○	○	○	○	○	○

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner hat Einfluss darauf, was ich mir im Fernsehen oder auf DVD ansehe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Urlaubspläne.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, größere Investitionen (Haus, Auto, etc.) vorzunehmen.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, einem Verein, einer sozialen Organisation, einer Partei etc. beizutreten.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, finanziell einen bestimmten Lebensstandard zu erreichen.	○	○	○	○	○	○	○

Figure A.13. Q_β page 5: Relationship Closeness Inventory, Strength subscale [34]

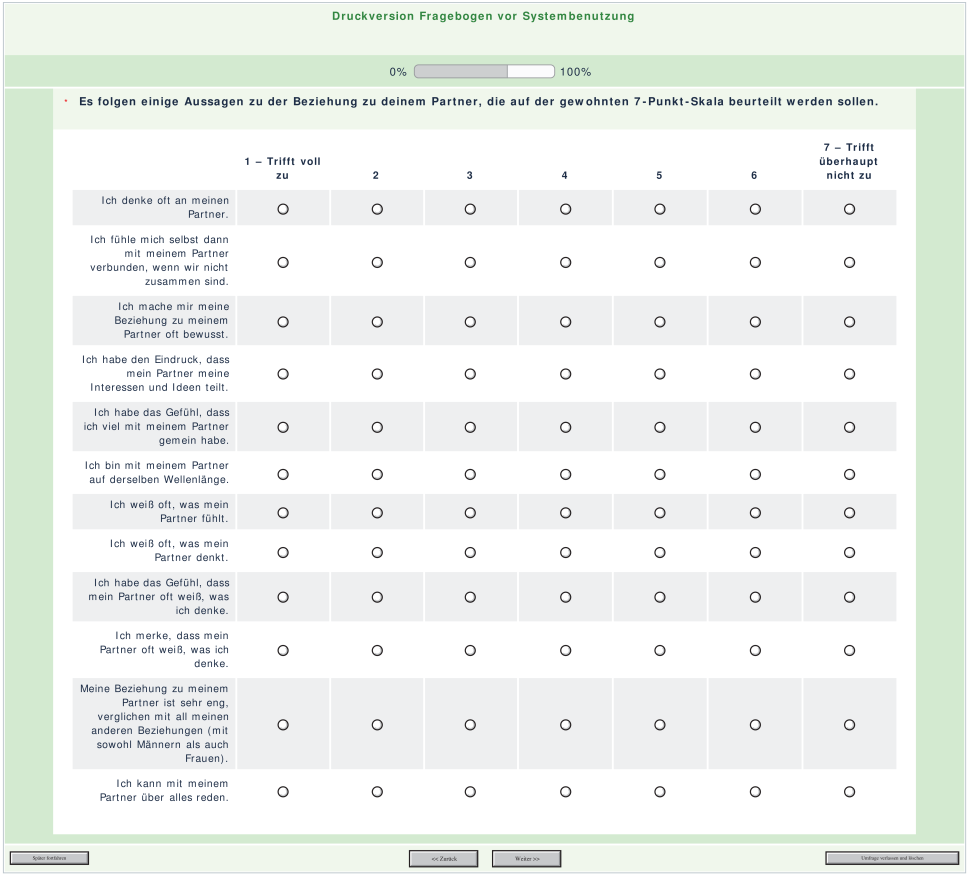

Druckversion Fragebogen vor Systembenutzung

0% 100%

* Es folgen einige Aussagen zu der Beziehung zu deinem Partner, die auf der gewohnten 7-Punkt-Skala beurteilt werden sollen.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich denke oft an meinen Partner.	○	○	○	○	○	○	○
Ich fühle mich selbst dann mit meinem Partner verbunden, wenn wir nicht zusammen sind.	○	○	○	○	○	○	○
Ich mache mir meine Beziehung zu meinem Partner oft bewusst.	○	○	○	○	○	○	○
Ich habe den Eindruck, dass mein Partner meine Interessen und Ideen teilt.	○	○	○	○	○	○	○
Ich habe das Gefühl, dass ich viel mit meinem Partner gemein habe.	○	○	○	○	○	○	○
Ich bin mit meinem Partner auf derselben Wellenlänge.	○	○	○	○	○	○	○
Ich weiß oft, was mein Partner fühlt.	○	○	○	○	○	○	○
Ich weiß oft, was mein Partner denkt.	○	○	○	○	○	○	○
Ich habe das Gefühl, dass mein Partner oft weiß, was ich denke.	○	○	○	○	○	○	○
Ich merke, dass mein Partner oft weiß, was ich denke.	○	○	○	○	○	○	○
Meine Beziehung zu meinem Partner ist sehr eng, verglichen mit all meinen anderen Beziehungen (mit sowohl Männern als auch Frauen).	○	○	○	○	○	○	○
Ich kann mit meinem Partner über alles reden.	○	○	○	○	○	○	○

Später fortfahren << Zurück Weiter >> Umfrage verlassen und löschen

Figure A.14. Q_β page 6: Connectedness scale based on van Bel et al. [550].

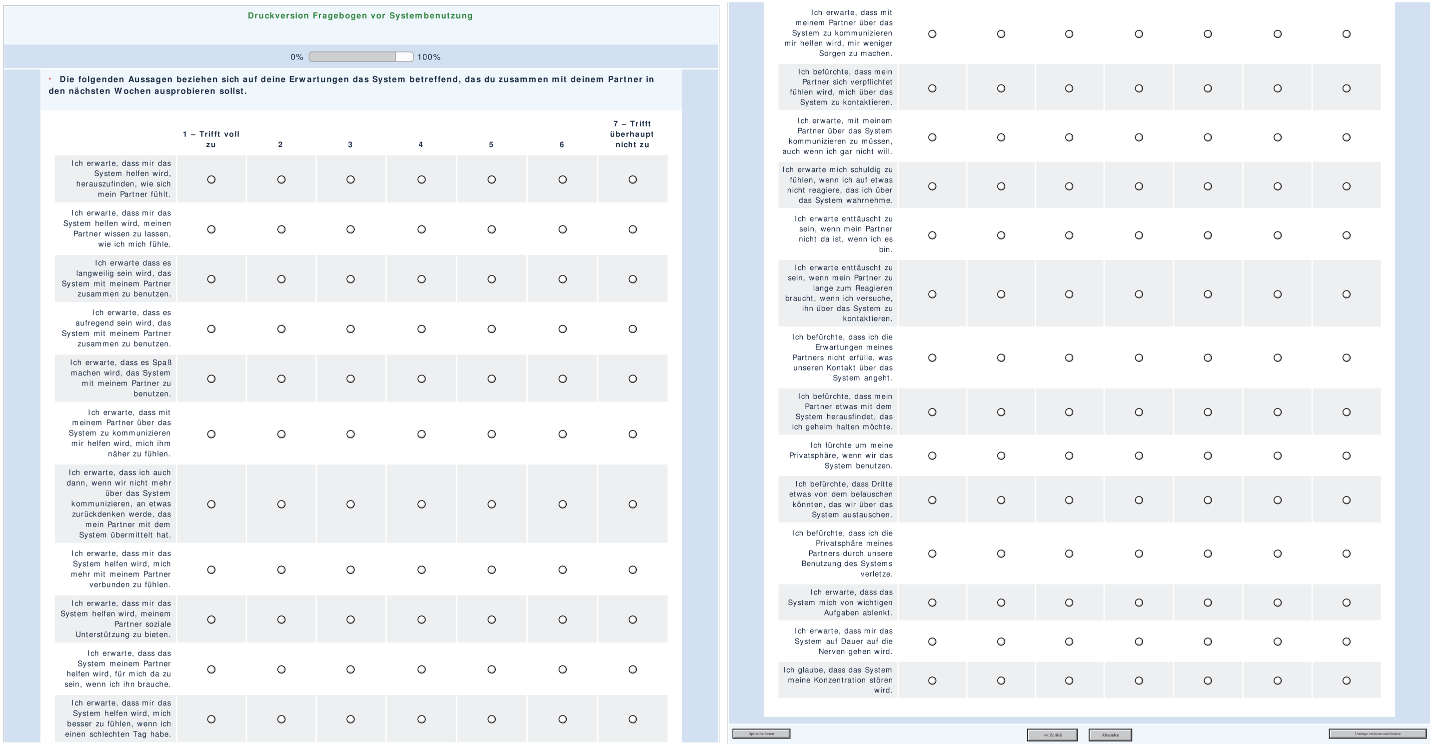
Druckversion Fragebogen vor Systembenutzung

0% 100%

* Die folgenden Aussagen beziehen sich auf deine Erwartungen das System betreffend, das du zusammen mit deinem Partner in den nächsten Wochen ausprobieren sollst.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich erwarte, dass mir das System helfen wird, herauszufinden, wie sich mein Partner fühlt.							
Ich erwarte, dass mir das System helfen wird, meinen Partner wissen zu lassen, wie ich mich fühle.							
Ich erwarte dass es langweilig sein wird, das System mit meinem Partner zusammen zu benutzen.							
Ich erwarte, dass es aufregend sein wird, das System mit meinem Partner zusammen zu benutzen.							
Ich erwarte, dass es Spaß machen wird, das System mit meinem Partner zu benutzen.							
Ich erwarte, dass mit meinem Partner über das System zu kommunizieren mir helfen wird, mich ihm näher zu fühlen.							
Ich erwarte, dass ich auch dann, wenn wir nicht mehr über das System kommunizieren, an etwas zurückdenken werde, das mein Partner mit dem System übermittelt hat.							
Ich erwarte, dass mir das System helfen wird, mich mehr mit meinem Partner verbunden zu fühlen.							
Ich erwarte, dass mir das System helfen wird, meinem Partner soziale Unterstützung zu bieten.							
Ich erwarte, dass das System meinem Partner helfen wird, für mich da zu sein, wenn ich ihn brauche.							
Ich erwarte, dass mir das System helfen wird, mich besser zu fühlen, wenn ich einen schlechten Tag habe.							
Ich erwarte, dass mit meinem Partner über das System zu kommunizieren mir helfen wird, mir weniger Sorgen zu machen.							
Ich befürchte, dass mein Partner sich verpflichtet fühlen wird, mich über das System zu kontaktieren.							
Ich erwarte, mit meinem Partner über das System kommunizieren zu müssen, auch wenn ich gar nicht will.							
Ich erwarte mich schuldig zu fühlen, wenn ich auf etwas nicht reagiere, das ich über das System wahrnehme.							
Ich erwarte enttäuscht zu sein, wenn mein Partner nicht da ist, wenn ich es bin.							
Ich erwarte enttäuscht zu sein, wenn mein Partner zu lange zum Reagieren braucht, wenn ich versuche, ihn über das System zu kontaktieren.							
Ich befürchte, dass ich die Erwartungen meines Partners nicht erfülle, was unseren Kontakt über das System angeht.							
Ich befürchte, dass mein Partner etwas mit dem System herausfindet, das ich geheim halten möchte.							
Ich fürchte um meine Privatsphäre, wenn wir das System benutzen.							
Ich befürchte, dass Dritte etwas von dem belauschen könnten, das wir über das System austauschen.							
Ich befürchte, dass ich die Privatsphäre meines Partners durch unsere Benutzung des Systems verletze.							
Ich erwarte, dass das System mich von wichtigen Aufgaben ablenkt.							
Ich erwarte, dass mir das System auf Dauer auf die Nerven gehen wird.							
Ich glaube, dass das System meine Konzentration stören wird.							

Figure A.15. Q_β page 7: affective costs and benefits expectations

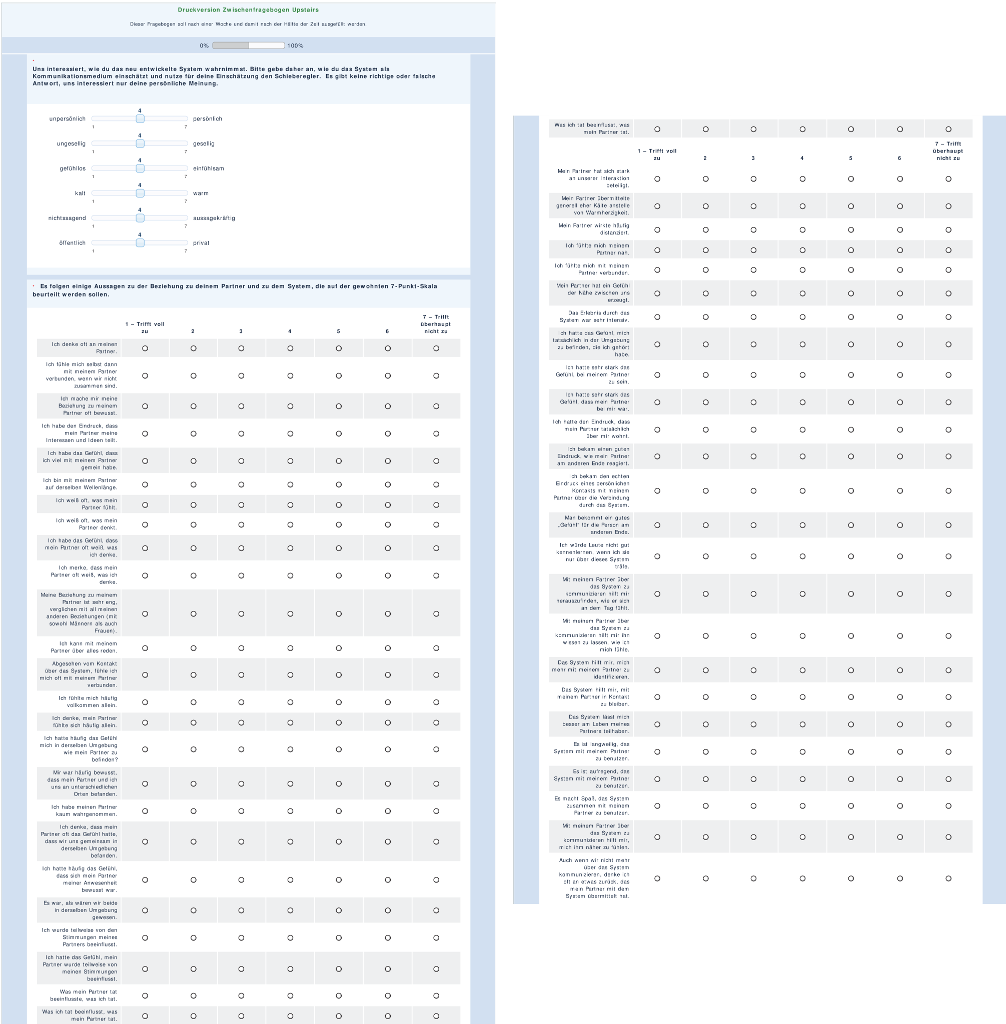

Druckversion Zwischenfragebogen Upstairs

Dieser Fragebogen soll nach einer Woche und damit nach der Hälfte der Zeit ausgefüllt werden.

0% 100%

Uns interessiert, wie du das neu entwickelte System wahrnimmst. Bitte gebe daher an, wie du das System als Kommunikationsmedium einschätzt und nutze für deine Einschätzung den Schieberegler. Es gibt keine richtige oder falsche Antwort, uns interessiert nur deine persönliche Meinung.

	1	4	7	
unpersönlich	1	4	7	persönlich
ungesellig	1	4	7	gesellig
gefühllos	1	4	7	einfühlsam
kalt	1	4	7	warm
nichtssagend	1	4	7	aussagekräftig
öffentlich	1	4	7	privat

Es folgen einige Aussagen zu der Beziehung zu deinem Partner und zu dem System, die auf der gewohnten 7-Punkt-Skala beurteilt werden sollen.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich denke oft an meinen Partner.	○	○	○	○	○	○	○
Ich fühle mich selbst dann mit meinem Partner verbunden, wenn wir nicht zusammen sind.	○	○	○	○	○	○	○
Ich mache mir meine Beziehung zu meinem Partner oft bewusst.	○	○	○	○	○	○	○
Ich habe den Eindruck, dass mein Partner meine Interessen und Ideen teilt.	○	○	○	○	○	○	○
Ich habe das Gefühl, dass ich viel mit meinem Partner gemein habe.	○	○	○	○	○	○	○
Ich bin mit meinem Partner auf derselben Wellenlänge.	○	○	○	○	○	○	○
Ich weiß oft, was mein Partner fühlt.	○	○	○	○	○	○	○
Ich weiß oft, was mein Partner denkt.	○	○	○	○	○	○	○
Ich habe das Gefühl, dass mein Partner oft weiß, was ich denke.	○	○	○	○	○	○	○
Ich merke, dass mein Partner oft weiß, was ich denke.	○	○	○	○	○	○	○
Meine Beziehung zu meinem Partner ist sehr eng, verglichen mit all meinen anderen Beziehungen (mit sowohl Männern als auch Frauen).	○	○	○	○	○	○	○
Ich kann mit meinem Partner über alles reden.	○	○	○	○	○	○	○
Abgesehen vom Kontakt über das System, fühle ich mich oft mit meinem Partner verbunden.	○	○	○	○	○	○	○
Ich fühlte mich häufig vollkommen allein.	○	○	○	○	○	○	○
Ich denke, mein Partner fühlte sich häufig allein.	○	○	○	○	○	○	○
Ich hatte häufig das Gefühl mich in derselben Umgebung wie mein Partner zu befinden?	○	○	○	○	○	○	○
Mir war häufig bewusst, dass mein Partner und ich uns an unterschiedlichen Orten befanden.	○	○	○	○	○	○	○
Ich habe meinen Partner kaum wahrgenommen.	○	○	○	○	○	○	○
Ich denke, dass mein Partner oft das Gefühl hatte, dass wir uns gemeinsam in derselben Umgebung befanden.	○	○	○	○	○	○	○
Ich hatte häufig das Gefühl, dass sich mein Partner meiner Anwesenheit bewusst war.	○	○	○	○	○	○	○
Es war, als wären wir beide in derselben Umgebung gewesen.	○	○	○	○	○	○	○
Ich wurde teilweise von den Stimmungen meines Partners beeinflusst.	○	○	○	○	○	○	○
Ich hatte das Gefühl, mein Partner wurde teilweise von meinen Stimmungen beeinflusst.	○	○	○	○	○	○	○
Was mein Partner tat beeinflusste, was ich tat.	○	○	○	○	○	○	○
Was ich tat beeinflusst, was mein Partner tat.	○	○	○	○	○	○	○

Was ich tat beeinflusst, was mein Partner tat.	○	○	○	○	○	○	○

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner hat sich stark an unserer Interaktion beteiligt.	○	○	○	○	○	○	○
Mein Partner übermittelte generell eher Kälte anstelle von Warmherzigkeit.	○	○	○	○	○	○	○
Mein Partner wirkte häufig distanziert.	○	○	○	○	○	○	○
Ich fühlte mich meinem Partner nah.	○	○	○	○	○	○	○
Ich fühlte mich mit meinem Partner verbunden.	○	○	○	○	○	○	○
Mein Partner hat ein Gefühl der Nähe zwischen uns erzeugt.	○	○	○	○	○	○	○
Das Erlebnis durch das System war sehr intensiv.	○	○	○	○	○	○	○
Ich hatte das Gefühl, mich tatsächlich in der Umgebung zu befinden, die ich gehört habe.	○	○	○	○	○	○	○
Ich hatte sehr stark das Gefühl, bei meinem Partner zu sein.	○	○	○	○	○	○	○
Ich hatte sehr stark das Gefühl, dass mein Partner bei mir war.	○	○	○	○	○	○	○
Ich hatte den Eindruck, dass mein Partner tatsächlich über mir wohnt.	○	○	○	○	○	○	○
Ich bekam einen guten Eindruck, wie mein Partner am anderen Ende reagiert.	○	○	○	○	○	○	○
Ich bekam den echten Eindruck eines persönlichen Kontakts mit meinem Partner über die Verbindung durch das System.	○	○	○	○	○	○	○
Man bekommt ein gutes „Gefühl" für die Person am anderen Ende.	○	○	○	○	○	○	○
Ich würde Leute nicht gut kennenlernen, wenn ich sie nur über dieses System träfe.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir herauszufinden, wie er sich an dem Tag fühlt.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir ihn wissen zu lassen, wie ich mich fühle.	○	○	○	○	○	○	○
Das System hilft mir, mich mehr mit meinem Partner zu identifizieren.	○	○	○	○	○	○	○
Das System hilft mir, mit meinem Partner in Kontakt zu bleiben.	○	○	○	○	○	○	○
Das System lässt mich besser am Leben meines Partners teilhaben.	○	○	○	○	○	○	○
Es ist langweilig, das System mit meinem Partner zu benutzen.	○	○	○	○	○	○	○
Es ist aufregend, das System mit meinem Partner zu benutzen.	○	○	○	○	○	○	○
Es macht Spaß, das System zusammen mit meinem Partner zu benutzen.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir, mich ihm näher zu fühlen.	○	○	○	○	○	○	○
Auch wenn wir nicht mehr über das System kommunizieren, denke ich oft an etwas zurück, das mein Partner mit dem System übermittelt hat.	○	○	○	○	○	○	○

Figure A.16. Q_γ page 3: main group parts 1 and 2

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Durch die Kommunikation mit meinem Partner über das System entsteht eine Verbindung, die über die Dauer des Austauschs hinaus wirkt.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir, mich mehr mit ihm verbunden zu fühlen.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir, ihm soziale Unterstützung zu bieten.	○	○	○	○	○	○	○
Das System hilft meinem Partner für mich da zu sein, wenn ich ihn brauche.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir mich besser zu fühlen, wenn ich einen schlechten Tag habe.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hilft mir, mir über etwas weniger Sorgen zu machen.	○	○	○	○	○	○	○
Ich mache mir Sorgen, dass mein Partner sich verpflichtet fühlt, mich über das System zu kontaktieren.	○	○	○	○	○	○	○
Ich muss mit meinem Partner über das System kommunizieren, auch wenn ich gar nicht will.	○	○	○	○	○	○	○
Ich fühle mich schuldig, wenn ich auf etwas nicht reagiere, das ich über das System wahrnehme.	○	○	○	○	○	○	○
Wenn ich versuche mit meinem Partner zu treten, erwarte ich eine Reaktion.	○	○	○	○	○	○	○
Ich bin enttäuscht, wenn mein Partner nicht da ist, wenn ich es bin.	○	○	○	○	○	○	○
Ich bin enttäuscht, wenn mein Partner zu lange zum Reagieren braucht, wenn ich versuche, ihn über das System zu kontaktieren.	○	○	○	○	○	○	○
Ich mache mir Sorgen, dass ich die Erwartungen meines Partners nicht erfülle, was unseren Kontakt über das System angeht.	○	○	○	○	○	○	○
Ich mache mir Sorgen, dass mein Partner etwas mit dem System herausfindet, das ich geheim halten möchte.	○	○	○	○	○	○	○
Ich machte mir Sorgen um meine Privatsphäre, während wir das System benutzt haben.	○	○	○	○	○	○	○
Ich mache mir Sorgen, dass Dritte etwas von dem belauschen können, das wir über das System austauschen.	○	○	○	○	○	○	○
Ich mache mir Sorgen, dass ich die Privatsphäre meines Partners durch unsere Benutzung des Systems verletze.	○	○	○	○	○	○	○
Das System verrät mir mehr über meinen Partner als ihm vielleicht lieb ist.	○	○	○	○	○	○	○
Das System lenkt mich von wichtigen Aufgaben ab.	○	○	○	○	○	○	○
Mir ist meist nicht aufgefallen, dass das System lief.	○	○	○	○	○	○	○
Es hat mich nicht gestört, dass das System lief.	○	○	○	○	○	○	○
Es fällt mir leicht, das System in den Hintergrund zu schieben.	○	○	○	○	○	○	○
Auch wenn ich nicht bewusst auf das System achte, bemerke ich darüber Ereignisse, die wichtig für mich sind.	○	○	○	○	○	○	○
Es war sehr schwierig, gleichzeitig etwas zu tun, das Aufmerksamkeit erforderte, während das System Geräusche produzierte.	○	○	○	○	○	○	○
Ich hatte das Gefühl, dass die Geräusche durch das System meine Konzentration gestört haben.	○	○	○	○	○	○	○

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Die Benutzung des Systems ist mühelos.	○	○	○	○	○	○	○
Ich kann einfach ignorieren, was mein Partner gerade tut.	○	○	○	○	○	○	○
Ich habe meinem Partner große Aufmerksamkeit geschenkt.	○	○	○	○	○	○	○
Das System hat häufig meine Aufmerksamkeit auf sich gezogen.	○	○	○	○	○	○	○
Ich hatte oft den Eindruck, erst im Nachhinein zu realisieren, dass ich Geräusche durch das System gehört habe.	○	○	○	○	○	○	○
Ich habe häufig aktiv versucht, mit meinem Partner in Kontakt zu treten.	○	○	○	○	○	○	○
Ich mag es, das System zu benutzen.	○	○	○	○	○	○	○

Später fortfahren | Absenden | Umfrage verlassen und löschen

Figure A.17. Q_{γ} page 3: main group part 3

Druckversion Abschlussfragebogen Upstairs

Dieser Fragebogen soll nach zwei Wochen und damit am Ande der Zeit ausgefüllt werden.

0% 100%

* Im Folgenden findest du einige Aussagen sowohl über dich als auch über deinen Partner. Bitte lies sie dir sorgfältig durch und gib zu jeder Aussage an, inwieweit du findest, dass sie zutrifft.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich finde, dass es ziemlich leicht für mich ist, anderen gefühlsmäßig nahe zu sein. Es geht mir gut, wenn ich mich auf andere verlassen kann und wenn andere sich auf mich verlassen. Ich mache mir keine Gedanken darüber, dass ich allein sein könnte oder dass andere mich nicht akzeptieren könnten.	○	○	○	○	○	○	○
Ich finde, dass es ziemlich leicht für meinen Partner ist, anderen gefühlsmäßig nahe zu sein. Es geht ihm gut, wenn er sich auf andere verlassen kann und wenn andere sich auf ihn verlassen. Er macht sich keine Gedanken darüber, dass er allein sein könnte oder dass andere ihn nicht akzeptieren könnten.	○	○	○	○	○	○	○
Es geht mir auch ohne enge gefühlsmäßige Bindung gut. Es ist sehr wichtig für mich, mich unabhängig und selbständig zu fühlen und ich ziehe es vor, wenn ich nicht von anderen und andere nicht von mir abhängig sind.	○	○	○	○	○	○	○
Es geht meinem Partner auch ohne enge gefühlsmäßige Bindung gut. Es ist sehr wichtig für ihn, sich unabhängig und selbständig zu fühlen und er zieht es vor, wenn er nicht von anderen und andere nicht von ihm abhängig sind.	○	○	○	○	○	○	○
Ich möchte anderen gefühlsmäßig sehr nahe sein, aber ich merke oft, dass andere Widerstände dagegen errichten, mir so nahe zu sein, wie ich ihnen nahe sein möchte. Es geht mir sehr gut, wenn ich ohne enge Beziehung bin, aber ich denke manchmal, dass andere mich nicht so sehr schätzen wie ich sie.	○	○	○	○	○	○	○
Mein Partner möchte anderen gefühlsmäßig sehr nahe sein, aber er merkt oft, dass andere Widerstände dagegen errichten, ihm so nahe zu sein, wie er ihnen nahe sein möchte. Es geht ihm sehr gut, wenn er ohne enge Beziehung ist, aber er denkt manchmal, dass andere ihn nicht so sehr schätzen wie er sie.	○	○	○	○	○	○	○
Ich empfinde es manchmal als ziemlich unangenehm, anderen nahe zu sein. Ich möchte Beziehungen, in denen ich anderen nahe bin, aber ich finde es schwierig, ihnen vollständig zu vertrauen oder von ihnen abhängig zu sein. Ich fürchte manchmal, dass ich verletzt werde, wenn ich mir erlaube, anderen zu nahe zu kommen.	○	○	○	○	○	○	○
Mein Partner empfindet es manchmal als ziemlich unangenehm, anderen nahe zu sein. Er möchte Beziehungen, in denen er anderen nahe ist, aber er findet es schwierig, ihnen vollständig zu vertrauen oder von ihnen abhängig zu sein. Er fürchtet manchmal, dass er verletzt wird, wenn er sich erlaubt, anderen zu nahe zu kommen.	○	○	○	○	○	○	○

Später fortführen | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.18. $Q_{\beta\gamma}$ page 3: relationship questionnaire [26, 118]

Druckversion Abschlussfragebogen Upstairs

Dieser Fragebogen soll nach zwei Wochen und damit am Ande der Zeit ausgefüllt werden.

0% 100%

*

Die folgenden Fragen betreffen den Einfluss, den dein Partner auf deine Gedanken, Gefühle und dein Verhalten hat. Bitte gib auf einer 7-Punkt-Skala an, wie sehr du mit jedem Punkt übereinstimmst.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner wird meine zukünftige finanzielle Sicherheit beeinflussen.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf alltägliche Dinge in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf wichtige Dinge in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie sehr ich Verantwortung in unserer Beziehung übernehme.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie viel Zeit ich mit Hausarbeit verbringe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie ich beschließe, Geld auszugeben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich mich selbst empfinde.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Laune.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf die grundlegenden Werte, die ich vertrete.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Meinung über andere wichtige Menschen in meinem Leben.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich meine Familie sehe und wie viel Zeit ich mit ihr verbringe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich meine Freunde sehe und wie viel Zeit ich mit ihnen verbringe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, welche meiner Freunde ich treffe.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf die Art meiner beruflichen Laufbahn.	○	○	○	○	○	○	○
Mein Partner beeinflusst oder wird beeinflussen, wie viel Zeit ich für meine Karriere aufwende.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine Chancen, zukünftig einen guten Job zu bekommen.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf die Art, wie ich mich in Bezug auf die Zukunft fühle.	○	○	○	○	○	○	○
Mein Partner hat nicht die Fähigkeit zu beeinflussen, wie ich mich in verschiedensten Situationen verhalte.	○	○	○	○	○	○	○
Mein Partner beeinflusst und trägt dazu bei, wie glücklich ich mich insgesamt fühle.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss auf meine derzeitige finanzielle Absicherung.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich meine Freizeit verbringe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wann ich ihn sehe und wie viel Zeit wir zusammen verbringen.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wie ich mich anziehe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss darauf, wie ich mein Zuhause einrichte.	○	○	○	○	○	○	○
Mein Partner hat keinen Einfluss darauf, wo ich lebe.	○	○	○	○	○	○	○
	1 – Trifft voll zu	**2**	**3**	**4**	**5**	**6**	**7 – Trifft überhaupt nicht zu**
Mein Partner hat Einfluss darauf, was ich mir im Fernsehen oder auf DVD ansehe.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Urlaubspläne.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, größere Investitionen (Haus, Auto, etc.) vorzunehmen.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, einem Verein, einer sozialen Organisation, einer Partei etc. beizutreten.	○	○	○	○	○	○	○
Mein Partner hat Einfluss auf meine Pläne, finanziell einen bestimmten Lebensstandard zu erreichen.	○	○	○	○	○	○	○

Später fortfahren | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.19. $Q_{\beta\gamma}$ page 4: Relationship closeness inventory strength subscale [34]

Druckversion Abschlussfragebogen Upstairs

Dieser Fragebogen soll nach zwei Wochen und damit am Ande der Zeit ausgefüllt werden.

0% 100%

Uns interessiert, wie du das neu entwickelte System wahrnimmst. Bitte gib daher an, wie du das System als Kommunikationsmedium einschätzt und nutze für deine Einschätzung den Schieberegler. Es gibt keine richtige oder falsche Antwort, uns interessiert nur deine persönliche Meinung.

	1	4	7	
unpersönlich	1	4	7	persönlich
ungesellig	1	4	7	gesellig
gefühllos	1	4	7	einfühlsam
kalt	1	4	7	warm
nichtssagend	1	4	7	aussagekräftig
öffentlich	1	4	7	privat

Es folgen einige Aussagen zu dem System und zu der Beziehung zu deinem Partner, die auf der gewohnten 7-Punkt-Skala beurteilt werden sollen.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich denke oft an meinen Partner.	O	O	O	O	O	O	O
Ich fühle mich selbst dann mit meinem Partner verbunden, wenn wir nicht zusammen sind.	O	O	O	O	O	O	O
Ich mache mir meine Beziehung zu meinem Partner oft bewusst.	O	O	O	O	O	O	O
Ich habe den Eindruck, dass mein Partner meine Interessen und Ideen teilt.	O	O	O	O	O	O	O
Ich habe das Gefühl, dass ich viel mit meinem Partner gemein habe.	O	O	O	O	O	O	O
Ich bin mit meinem Partner auf derselben Wellenlänge.	O	O	O	O	O	O	O
Ich weiß oft, was mein Partner fühlt.	O	O	O	O	O	O	O
Ich weiß oft, was mein Partner denkt.	O	O	O	O	O	O	O
Ich habe das Gefühl, dass mein Partner oft weiß, was ich denke.	O	O	O	O	O	O	O
Ich merke, dass mein Partner oft weiß, was ich denke.	O	O	O	O	O	O	O
Meine Beziehung zu meinem Partner ist sehr eng, verglichen mit all meinen anderen Beziehungen (mit sowohl Männern als auch Frauen).	O	O	O	O	O	O	O
Ich kann mit meinem Partner über alles reden.	O	O	O	O	O	O	O
Abgesehen vom Kontakt über das System, fühlte ich mich oft mit meinem Partner verbunden.	O	O	O	O	O	O	O
Ich fühlte mich häufig vollkommen allein.	O	O	O	O	O	O	O
Ich denke, mein Partner fühlte sich häufig allein.	O	O	O	O	O	O	O
Ich hatte häufig das Gefühl mich in derselben Umgebung wie mein Partner zu befinden.	O	O	O	O	O	O	O
Mir war häufig bewusst, dass mein Partner und ich uns an unterschiedlichen Orten befanden.	O	O	O	O	O	O	O
Ich habe meinen Partner kaum wahrgenommen.	O	O	O	O	O	O	O
Ich denke, dass mein Partner oft das Gefühl hatte, dass wir uns gemeinsam in derselben Umgebung befanden.	O	O	O	O	O	O	O
Ich hatte häufig das Gefühl, dass sich mein Partner meiner Anwesenheit bewusst war.	O	O	O	O	O	O	O
Es war, als wären wir beide in derselben Umgebung gewesen.	O	O	O	O	O	O	O
Ich wurde teilweise von den Stimmungen meines Partners beeinflusst.	O	O	O	O	O	O	O
Ich hatte das Gefühl, mein Partner wurde teilweise von meinen Stimmungen beeinflusst.	O	O	O	O	O	O	O
Was mein Partner tat beeinflusste, was ich tat.	O	O	O	O	O	O	O
Was ich tat beeinflusste, was mein Partner tat.	O	O	O	O	O	O	O

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner hat sich stark an unserer Interaktion beteiligt.	O	O	O	O	O	O	O
Mein Partner übermittelte generell eher Kälte anstelle von Warmherzigkeit.	O	O	O	O	O	O	O
Mein Partner wirkte häufig distanziert.	O	O	O	O	O	O	O
Ich fühlte mich meinem Partner nah.	O	O	O	O	O	O	O
Ich fühlte mich mit meinem Partner verbunden.	O	O	O	O	O	O	O
Mein Partner hat ein Gefühl der Nähe zwischen uns erzeugt.	O	O	O	O	O	O	O
Das Erlebnis durch das System war sehr intensiv.	O	O	O	O	O	O	O
Ich hatte das Gefühl, mich tatsächlich in der Umgebung zu befinden, die ich gehört habe.	O	O	O	O	O	O	O
Ich hatte sehr stark das Gefühl, bei meinem Partner zu sein.	O	O	O	O	O	O	O
Ich hatte sehr stark das Gefühl, dass mein Partner bei mir war.	O	O	O	O	O	O	O
Ich hatte den Eindruck, dass mein Partner tatsächlich über mir wohnt.	O	O	O	O	O	O	O
Ich bekam einen guten Eindruck, wie mein Partner am anderen Ende reagiert.	O	O	O	O	O	O	O
Ich bekam den echten Eindruck eines persönlichen Kontakts mit meinem Partner über die Verbindung durch das System.	O	O	O	O	O	O	O
Man bekommt ein gutes „Gefühl" für die Person am anderen Ende.	O	O	O	O	O	O	O
Ich würde Leute nicht gut kennenlernen, wenn ich sie nur über dieses System träfe.	O	O	O	O	O	O	O
Mit meinem Partner über das System zu kommunizieren half mir herauszufinden, wie er sich an dem Tag fühlte.	O	O	O	O	O	O	O
Mit meinem Partner über das System zu kommunizieren half mir ihn wissen zu lassen, wie ich mich fühlte.	O	O	O	O	O	O	O
Das System half mir, mich mehr mit meinem Partner zu identifizieren.	O	O	O	O	O	O	O
Das System half mir, mit meinem Partner in Kontakt zu bleiben.	O	O	O	O	O	O	O
Das System ließ mich besser am Leben meines Partners teilhaben.	O	O	O	O	O	O	O
Es war langweilig, das System mit meinem Partner zu benutzen.	O	O	O	O	O	O	O
Es war aufregend, das System mit meinem Partner zu benutzen.	O	O	O	O	O	O	O
Es machte Spaß, das System zusammen mit meinem Partner zu benutzen.	O	O	O	O	O	O	O
Mit meinem Partner über das System zu kommunizieren half mir, mich ihm näher zu fühlen.	O	O	O	O	O	O	O
Auch wenn wir nicht mehr über das System kommunizierten, dachte ich oft an etwas zurück, das mein Partner mit dem System übermittelt hatte.	O	O	O	O	O	O	O

Figure A.20. $Q_{\beta\gamma}$ page 5: main group parts 1 and 2

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Durch die Kommunikation mit meinem Partner über das System entstand eine Verbindung, die über die Dauer des Austauschs hinaus wirkte.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren half mir, mich mehr mit ihm verbunden zu fühlen.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren half mir, ihm soziale Unterstützung zu bieten.	○	○	○	○	○	○	○
Das System half meinem Partner für mich da zu sein, wenn ich ihn brauchte.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren half mir mich besser zu fühlen, wenn ich einen schlechten Tag hatte.	○	○	○	○	○	○	○
Mit meinem Partner über das System zu kommunizieren hat mir geholfen, mir über etwas weniger Sorgen zu machen.	○	○	○	○	○	○	○
Ich machte mir Sorgen, dass mein Partner sich verpflichtet fühlt, mich über das System zu kontaktieren.	○	○	○	○	○	○	○
Ich musste mit meinem Partner über das System kommunizieren, auch wenn ich gar nicht wollte.	○	○	○	○	○	○	○
Ich fühlte mich schuldig, wenn ich auf etwas nicht reagiert habe, das ich über das System wahrgenommen habe.	○	○	○	○	○	○	○
Wenn ich versucht habe mit meinem Partner in Kontakt zu treten, erwartete ich eine Reaktion.	○	○	○	○	○	○	○
Ich war enttäuscht, wenn mein Partner nicht da war, wenn ich es war.	○	○	○	○	○	○	○
Ich war enttäuscht, wenn mein Partner zu lange zum Reagieren gebraucht hat, wenn ich versucht habe, ihn über das System zu kontaktieren.	○	○	○	○	○	○	○
Ich machte mir Sorgen, dass ich die Erwartungen meines Partners nicht würde erfüllen können, was unseren Kontakt über das System angeht.	○	○	○	○	○	○	○
Ich machte mir Sorgen, dass mein Partner etwas mit dem System hätte herausfinden können, das ich geheim halten wollte.	○	○	○	○	○	○	○
Ich machte mir Sorgen um meine Privatsphäre, während wir das System benutzt haben.	○	○	○	○	○	○	○
Ich machte mir Sorgen, dass Dritte etwas von dem würden belauschen können, das wir über das System austauschten.	○	○	○	○	○	○	○
Ich machte mir Sorgen, dass ich die Privatsphäre meines Partners durch unsere Benutzung des Systems hätte verletzen können.	○	○	○	○	○	○	○
Das System verriet mir mehr über meinen Partner als ihm vielleicht lieb ist.	○	○	○	○	○	○	○
Das System lenkte mich von wichtigen Aufgaben ab.	○	○	○	○	○	○	○
Mir ist meist nicht aufgefallen, dass das System lief.	○	○	○	○	○	○	○
Es hat mich nicht gestört, dass das System lief.	○	○	○	○	○	○	○
Es fiel mir leicht, das System in den Hintergrund zu schieben.	○	○	○	○	○	○	○
Auch wenn ich nicht bewusst auf das System geachtet habe, bemerkte ich darüber Ereignisse, die wichtig für mich waren.	○	○	○	○	○	○	○
Es war sehr schwierig, gleichzeitig etwas zu tun, das Aufmerksamkeit erforderte, während das System Geräusche produzierte.	○	○	○	○	○	○	○
Ich hatte das Gefühl, dass die Geräusche durch das System meine Konzentration gestört haben.	○	○	○	○	○	○	○
	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Die Benutzung des Systems war mühelos.	○	○	○	○	○	○	○
Ich konnte einfach ignorieren, was mein Partner in einem bestimmten Moment tat.	○	○	○	○	○	○	○
Ich habe meinem Partner große Aufmerksamkeit geschenkt.	○	○	○	○	○	○	○
Das System hat häufig meine Aufmerksamkeit auf sich gezogen.	○	○	○	○	○	○	○
Ich hatte oft den Eindruck, erst im Nachhinein zu realisieren, dass ich Geräusche durch das System gehört habe.	○	○	○	○	○	○	○
Ich habe häufig aktiv versucht, mit meinem Partner in Kontakt zu treten.	○	○	○	○	○	○	○
Ich hatte stark das Gefühl, dass ich eine echte Person gehört habe und nicht Geräusche, die durch einen Computer erzeugt wurden.	○	○	○	○	○	○	○
Ich mochte es, das System zu benutzen.	○	○	○	○	○	○	○

Später fortfahren | << Zurück | Absenden | Umfrage verlassen und löschen

Figure A.21. $Q_{\beta\gamma}$ page 5: main group part 3

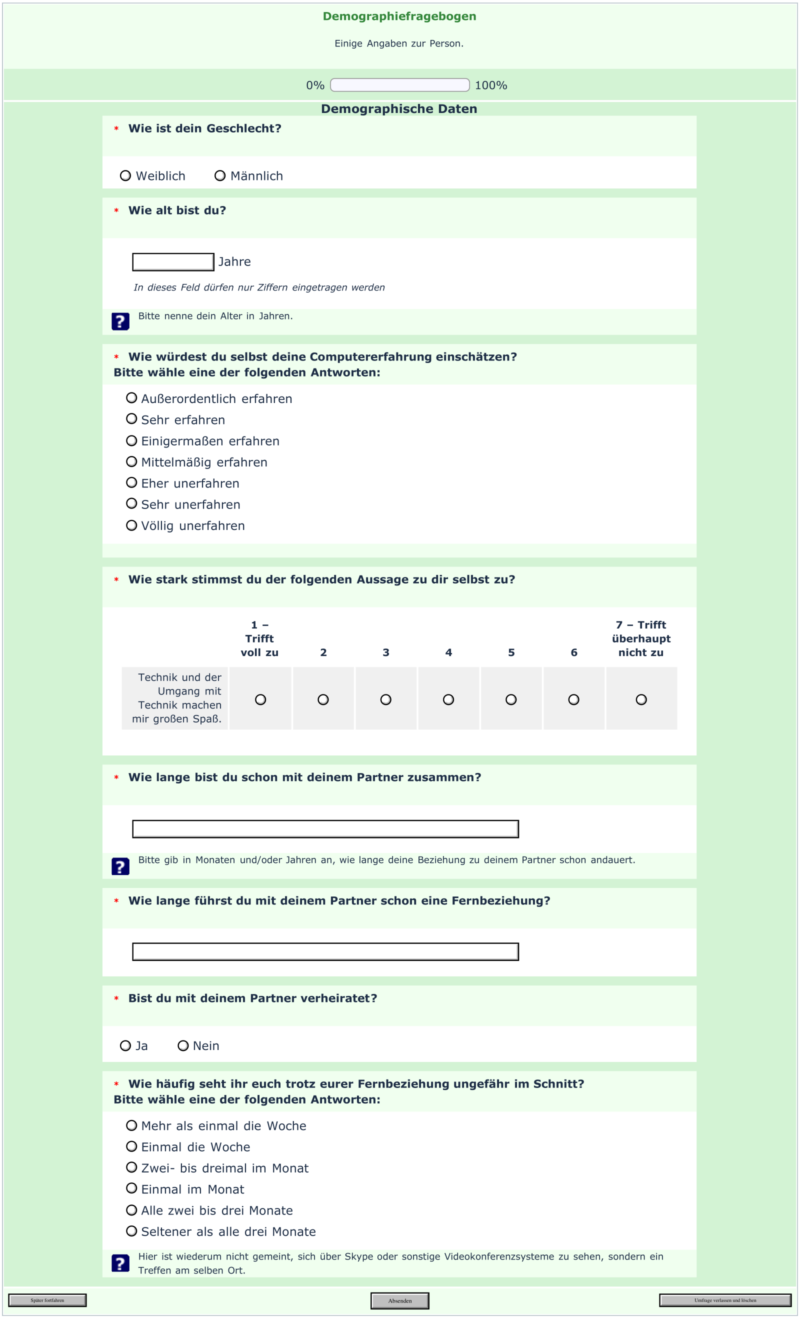

Demographiefragebogen

Einige Angaben zur Person.

0% 100%

Demographische Daten

* **Wie ist dein Geschlecht?**

- Weiblich
- Männlich

* **Wie alt bist du?**

Jahre

In dieses Feld dürfen nur Ziffern eingetragen werden

? Bitte nenne dein Alter in Jahren.

* **Wie würdest du selbst deine Computererfahrung einschätzen?**
Bitte wähle eine der folgenden Antworten:

- Außerordentlich erfahren
- Sehr erfahren
- Einigermaßen erfahren
- Mittelmäßig erfahren
- Eher unerfahren
- Sehr unerfahren
- Völlig unerfahren

* **Wie stark stimmst du der folgenden Aussage zu dir selbst zu?**

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Technik und der Umgang mit Technik machen mir großen Spaß.	O	O	O	O	O	O	O

* **Wie lange bist du schon mit deinem Partner zusammen?**

? Bitte gib in Monaten und/oder Jahren an, wie lange deine Beziehung zu deinem Partner schon andauert.

* **Wie lange führst du mit deinem Partner schon eine Fernbeziehung?**

* **Bist du mit deinem Partner verheiratet?**

- Ja
- Nein

* **Wie häufig seht ihr euch trotz eurer Fernbeziehung ungefähr im Schnitt?**
Bitte wähle eine der folgenden Antworten:

- Mehr als einmal die Woche
- Einmal die Woche
- Zwei- bis dreimal im Monat
- Einmal im Monat
- Alle zwei bis drei Monate
- Seltener als alle drei Monate

? Hier ist wiederum nicht gemeint, sich über Skype oder sonstige Videokonferenzsysteme zu sehen, sondern ein Treffen am selben Ort.

Später fortfahren | Absenden | Umfrage verlassen und löschen

Figure A.22. Demography questionnaire upstairs page 2: main group

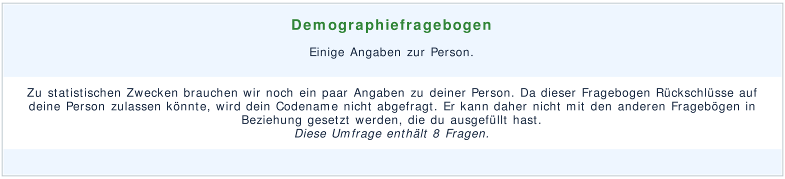

Demographiefragebogen

Einige Angaben zur Person.

Zu statistischen Zwecken brauchen wir noch ein paar Angaben zu deiner Person. Da dieser Fragebogen Rückschlüsse auf deine Person zulassen könnte, wird dein Codename nicht abgefragt. Er kann daher nicht mit den anderen Fragebögen in Beziehung gesetzt werden, die du ausgefüllt hast.
Diese Umfrage enthält 8 Fragen.

Figure A.23. Demography questionnaire upstairs page 1: welcome screen

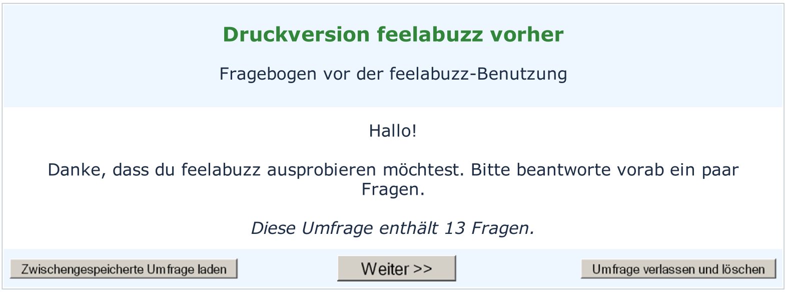

Druckversion feelabuzz vorher

Fragebogen vor der feelabuzz-Benutzung

Hallo!

Danke, dass du feelabuzz ausprobieren möchtest. Bitte beantworte vorab ein paar Fragen.

Diese Umfrage enthält 13 Fragen.

Zwischengespeicherte Umfrage laden | Weiter >> | Umfrage verlassen und löschen

Figure A.24. Questionnaire before FEELABUZZ use, welcome screen.

A.2 FEELABUZZ

There were only two FEELABUZZ questionnaires. One right before system use and and directly after the two-day study period. Both are very similar to the corresponding upstairs questionnaires. Some items were silently reworded when they were referring to the acoustic modality before so they apply to the sense of touch.

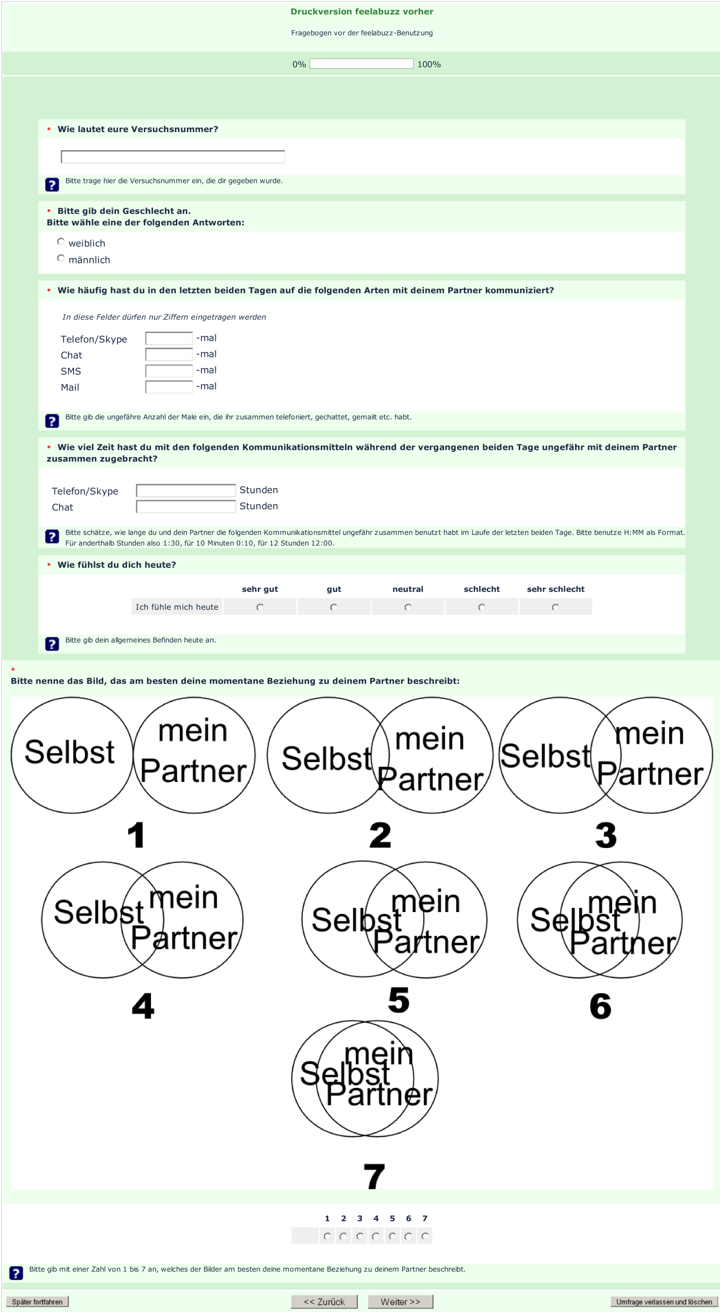

Druckversion feelabuzz vorher

Fragebogen vor der feelabuzz-Benutzung

0% 100%

* **Wie lautet eure Versuchsnummer?**

Bitte trage hier die Versuchsnummer ein, die dir gegeben wurde.

* **Bitte gib dein Geschlecht an.**
Bitte wähle eine der folgenden Antworten:

- weiblich
- männlich

* **Wie häufig hast du in den letzten beiden Tagen auf die folgenden Arten mit deinem Partner kommuniziert?**

In diese Felder dürfen nur Ziffern eingetragen werden

Telefon/Skype -mal
Chat -mal
SMS -mal
Mail -mal

Bitte gib die ungefähre Anzahl der Male ein, die ihr zusammen telefoniert, gechattet, gemailt etc. habt.

* **Wie viel Zeit hast du mit den folgenden Kommunikationsmitteln während der vergangenen beiden Tage ungefähr mit deinem Partner zusammen zugebracht?**

Telefon/Skype Stunden
Chat Stunden

Bitte schätze, wie lange du und dein Partner die folgenden Kommunikationsmittel ungefähr zusammen benutzt habt im Laufe der letzten beiden Tage. Bitte benutze H:MM als Format. Für anderthalb Stunden also 1:30, für 10 Minuten 0:10, für 12 Stunden 12:00.

* **Wie fühlst du dich heute?**

	sehr gut	gut	neutral	schlecht	sehr schlecht
Ich fühle mich heute	○	○	○	○	○

Bitte gib dein allgemeines Befinden heute an.

*
Bitte nenne das Bild, das am besten deine momentane Beziehung zu deinem Partner beschreibt:

Selbst mein Partner 1
Selbst mein Partner 2
Selbst mein Partner 3
Selbst mein Partner 4
Selbst mein Partner 5
Selbst mein Partner 6
Selbst mein Partner 7

1	2	3	4	5	6	7
○	○	○	○	○	○	○

Bitte gib mit einer Zahl von 1 bis 7 an, welches der Bilder am besten deine momentane Beziehung zu deinem Partner beschreibt.

Später fortfahren << Zurück Weiter >> Umfrage verlassen und löschen

Figure A.25. Questionnaire before FEELABUZZ use, page 1.

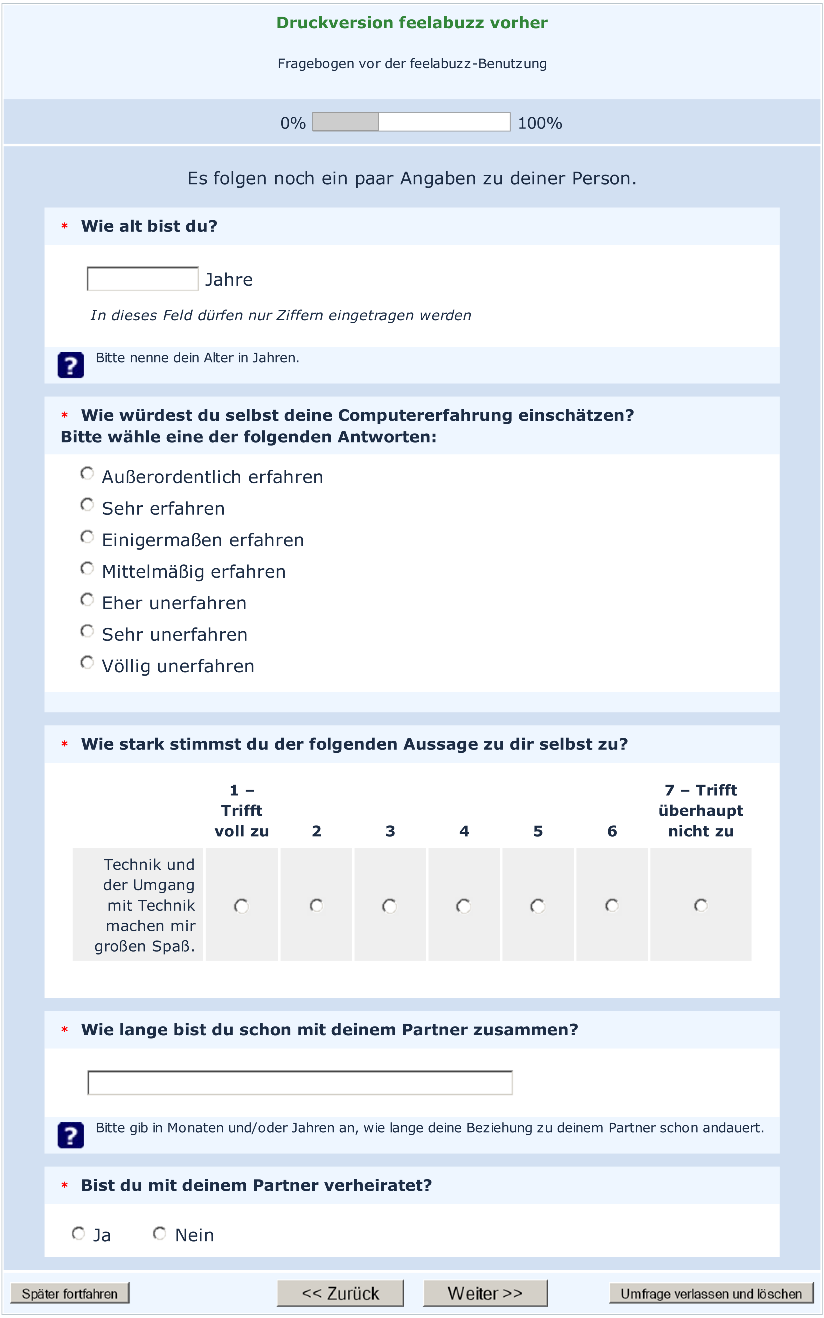

Druckversion feelabuzz vorher

Fragebogen vor der feelabuzz-Benutzung

0% 100%

Es folgen noch ein paar Angaben zu deiner Person.

* **Wie alt bist du?**

[ ] Jahre

In dieses Feld dürfen nur Ziffern eingetragen werden

? Bitte nenne dein Alter in Jahren.

* **Wie würdest du selbst deine Computererfahrung einschätzen?**
Bitte wähle eine der folgenden Antworten:

- Außerordentlich erfahren
- Sehr erfahren
- Einigermaßen erfahren
- Mittelmäßig erfahren
- Eher unerfahren
- Sehr unerfahren
- Völlig unerfahren

* **Wie stark stimmst du der folgenden Aussage zu dir selbst zu?**

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Technik und der Umgang mit Technik machen mir großen Spaß.	○	○	○	○	○	○	○

* **Wie lange bist du schon mit deinem Partner zusammen?**

? Bitte gib in Monaten und/oder Jahren an, wie lange deine Beziehung zu deinem Partner schon andauert.

* **Bist du mit deinem Partner verheiratet?**

○ Ja ○ Nein

Später fortfahren | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.26. Questionnaire before FEELABUZZ use, page 2.

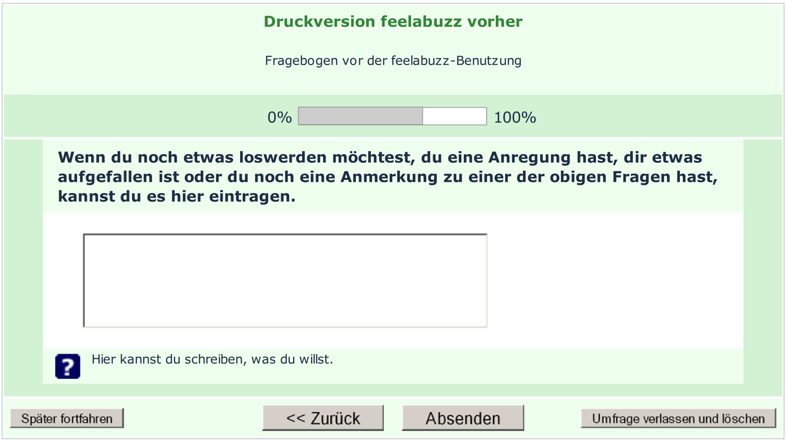

Figure A.27. Questionnaire before FEELABUZZ use, page 3.

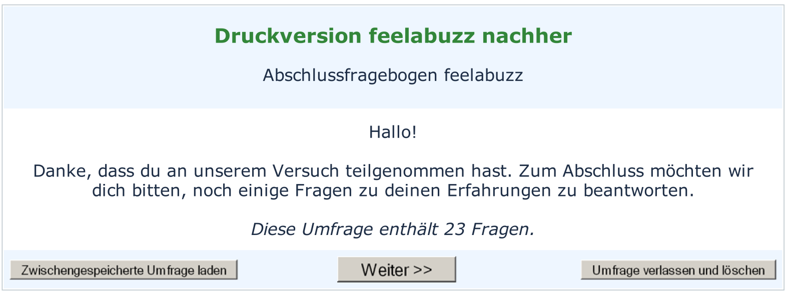

Figure A.28. Questionnaire after FEELABUZZ use, welcome screen.

Druckversion feelabuzz nachher

Abschlussfragebogen feelabuzz

0% [ ] 100%

* **Wie lautet eure Versuchsnummer?**

[ ]

Bitte trage hier die Versuchsnummer ein, die dir gegeben wurde.

* **Bitte gib dein Geschlecht an.**
Bitte wähle eine der folgenden Antworten:

- weiblich
- männlich

* **Wie häufig hast du in den letzten beiden Tagen auf die folgenden Arten mit deinem Partner kommuniziert?**

In diese Felder dürfen nur Ziffern eingetragen werden

Telefon/Skype	[ ] -mal
Chat	[ ] -mal
SMS	[ ] -mal
Mail	[ ] -mal

Bitte gib die ungefähre Anzahl der Male ein, die ihr zusammen telefoniert, gechattet, gemailt etc. habt.

* **Wie viel Zeit hast du mit den folgenden Kommunikationsmitteln während der vergangenen beiden Tage ungefähr mit deinem Partner zusammen zugebracht?**

Telefon/Skype	[ ] Stunden
Chat	[ ] Stunden

Bitte schätze, wie lange du und dein Partner die folgenden Kommunikationsmittel ungefähr zusammen benutzt habt im Laufe der letzten beiden Tage. Bitte benutze H:MM als Format. Für anderthalb Stunden also 1:30, für 10 Minuten 0:10, für 12 Stunden 12:00.

* **Wie fühlst du dich heute?**

	sehr gut	gut	neutral	schlecht	sehr schlecht
Ich fühle mich heute	○	○	○	○	○

Bitte gib dein allgemeines Befinden heute an.

*
Bitte nenne das Bild, das am besten deine momentane Beziehung zu deinem Partner beschreibt:

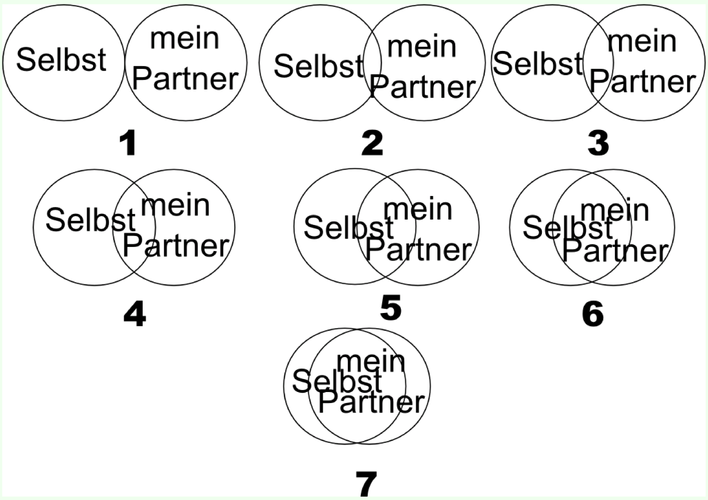

1	2	3	4	5	6	7
○	○	○	○	○	○	○

Bitte gib mit einer Zahl von 1 bis 7 an, welches der Bilder am besten deine momentane Beziehung zu deinem Partner beschreibt.

Später fortfahren | << Zurück | Weiter >> | Umfrage verlassen und löschen

Figure A.29. Questionnaire after FEELABUZZ use, page 1.

Druckversion feelabuzz nachher

Abschlussfragebogen feelabuzz

0% 100%

* **Welche verschiedenen Aktivitäten konntest du allein durch die Vibrationsmuster erkennen?**

* **Wenn du so etwas wie Gehen/Laufen erkennen konntest, wie gut konntest du verschiedene Geschwindigkeiten unterscheiden?**

Ich konnte gar keine Geschwindigkeitsunterschiede ausmachen.

Ich konnte Geschwindigkeitsunterschiede sehr fein bestimmen.

Später fortfahren

<< Zurück

Weiter >>

Umfrage verlassen und löschen

Figure A.30. Questionnaire after FEELABUZZ use, page 2.

Druckversion feelabuzz nachher

Abschlussfragebogen feelabuzz

0% 100%

Es folgen einige Fragen zu dem System allgemein, zu der Beziehung zu deinem Partner und dazu, inwiefern diese Beziehung durch das System beeinflusst wurde.

*

Uns interessiert, wie du das neu entwickelte System wahrnimmst. Bitte gib daher an, wie du das System als Kommunikationsmedium einschätzt und nutze für deine Einschätzung den Schieberegler. Es gibt keine richtige oder falsche Antwort, uns interessiert nur deine persönliche Meinung.

unpersönlich	1 – 4 – 7	persönlich
ungesellig	1 – 4 – 7	gesellig
gefühllos	1 – 4 – 7	einfühlsam
kalt	1 – 4 – 7	warm
nichtssagend	1 – 4 – 7	aussagekräftig
öffentlich	1 – 4 – 7	privat

* **Es folgen einige Aussagen zu dem System und zu der Beziehung zu deinem Partner, die einer 7-Punkt-Skala beurteilt werden sollen.**

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich denke oft an meinen Partner.							
Ich fühle mich selbst dann mit meinem Partner verbunden, wenn wir nicht zusammen sind.							
Ich mache mir meine Beziehung zu meinem Partner oft bewusst.							
Ich habe den Eindruck, dass mein Partner meine Interessen und Ideen teilt.							
Ich habe das Gefühl, dass ich viel mit meinem Partner gemein habe.							
Ich bin mit meinem Partner auf derselben Wellenlänge.							
Ich weiß oft, was mein Partner fühlt.							
Ich weiß oft, was mein Partner denkt.							
Ich habe das Gefühl, dass mein Partner oft weiß, was ich denke.							
Ich merke, dass mein Partner oft weiß, was ich denke.							
Meine Beziehung zu meinem Partner ist sehr eng, verglichen mit all meinen anderen Beziehungen (mit sowohl Männern als auch Frauen).							
Ich kann mit meinem Partner über alles reden.							
Abgesehen vom Kontakt über das System, fühlte ich mich oft mit meinem Partner verbunden.							
Ich fühlte mich häufig vollkommen allein.							
Ich denke, mein Partner fühlte sich häufig allein.							
Ich hatte häufig das Gefühl mich in derselben Umgebung wie mein Partner zu befinden.							
Mir war häufig bewusst, dass mein Partner und ich uns an unterschiedlichen Orten befanden.							
Ich habe meinen Partner kaum wahrgenommen.							
Ich denke, dass mein Partner oft das Gefühl hatte, dass wir uns gemeinsam in derselben Umgebung befanden.							
Ich hatte häufig das Gefühl, dass sich mein Partner meiner Anwesenheit bewusst war.							
Es war, als wären wir beide in derselben Umgebung gewesen.							
Ich wurde teilweise von den Stimmungen meines Partners beeinflusst.							
Ich hatte das Gefühl, mein Partner wurde teilweise von meinen Stimmungen beeinflusst.							
Was mein Partner tat beeinflusste, was ich tat.							
Was ich tat beeinflusste, was mein Partner tat.							

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Mein Partner hat sich stark an unserer Interaktion beteiligt.							
Mein Partner übermittelte generell eher Kälte anstelle von Warmherzigkeit.							
Mein Partner wirkte häufig distanziert.							
Ich fühlte mich meinem Partner nah.							
Ich fühlte mich mit meinem Partner verbunden.							
Mein Partner hat ein Gefühl der Nähe zwischen uns erzeugt.							
Das Erlebnis durch das System war sehr intensiv.							
Ich hatte das Gefühl, tatsächlich die Bewegungen meines Partners zu fühlen.							
Ich hatte sehr stark das Gefühl, bei meinem Partner zu sein.							
Ich hatte sehr stark das Gefühl, dass mein Partner bei mir war.							
Ich hatte den Eindruck, dass mein Partner tatsächlich über mir wohnt.							
Ich bekam einen guten Eindruck, wie mein Partner am anderen Ende reagiert.							
Ich bekam den echten Eindruck eines persönlichen Kontakts mit meinem Partner über die Verbindung durch das System.							
Man bekommt ein gutes „Gefühl" für die Person am anderen Ende.							
Ich würde Leute nicht gut kennenlernen, wenn ich sie nur über dieses System träfe.							
Mit meinem Partner über das System zu kommunizieren half mir herauszufinden, wie er sich an dem Tag fühlte.							
Mit meinem Partner über das System zu kommunizieren half mir ihn wissen zu lassen, wie ich mich fühlte.							
Das System half mir, mich mehr mit meinem Partner zu identifizieren.							
Das System half mir, mit meinem Partner in Kontakt zu bleiben.							
Das System ließ mich besser am Leben meines Partners teilhaben.							
Es war langweilig, das System mit meinem Partner zu benutzen.							
Es war aufregend, das System mit meinem Partner zu benutzen.							
Es machte Spaß, das System zusammen mit meinem Partner zu benutzen.							
Mit meinem Partner über das System zu kommunizieren half mir, mich ihm näher zu fühlen.							
Auch wenn wir nicht mehr über das System kommunizierten, dachte ich oft an etwas zurück, das mein Partner mit dem System übermittelt hatte.							

Figure A.31. Questionnaire after FEELABUZZ use, page 3, parts 1 and 2.

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Durch die Kommunikation mit meinem Partner über das System entstand eine Verbindung, die über die Dauer des Austauschs hinaus wirkte.							
Mit meinem Partner über das System zu kommunizieren half mir, mich mehr mit ihm verbunden zu fühlen.							
Mit meinem Partner über das System zu kommunizieren half mir, ihm soziale Unterstützung zu bieten.							
Das System half meinem Partner für mich da zu sein, wenn ich ihn brauchte.							
Mit meinem Partner über das System zu kommunizieren half mir mich besser zu fühlen, wenn ich einen schlechten Tag hatte.							
Mit meinem Partner über das System zu kommunizieren hat mir geholfen, mir über etwas weniger Sorgen zu machen.							
Ich machte mir Sorgen, dass mein Partner sich verpflichtet fühlt, mich über das System zu kontaktieren.							
Ich musste mit meinem Partner über das System kommunizieren, auch wenn ich gar nicht wollte.							
Ich fühlte mich schuldig, wenn ich auf etwas nicht reagiert habe, das ich über das System wahrgenommen habe.							
Wenn ich versucht habe mit meinem Partner in Kontakt zu treten, erwartete ich eine Reaktion.							
Ich war enttäuscht, wenn mein Partner zu lange zum Reagieren gebraucht hat, wenn ich versucht habe, ihn über das System zu kontaktieren.							
Ich machte mir Sorgen, dass ich die Erwartungen meines Partners nicht würde erfüllen können, was unseren Kontakt über das System angeht.							
Ich machte mir Sorgen, dass mein Partner etwas mit dem System hätte herausfinden können, das ich geheim halten wollte.							
Ich machte mir Sorgen um meine Privatsphäre, während wir das System benutzt haben.							
Ich machte mir Sorgen, dass Dritte etwas von dem würden belauschen können, das wir über das System austauschten.							
Ich machte mir Sorgen, dass ich die Privatsphäre meines Partners durch unsere Benutzung des Systems hätte verletzen können.							
Das System verriet mir mehr über meinen Partner als ihm vielleicht lieb ist.							
Das System lenkte mich von wichtigen Aufgaben ab.							
Mir ist meist nicht aufgefallen, dass das System lief.							
Es hat mich nicht gestört, dass das System lief.							
Es fiel mir leicht, das System in den Hintergrund zu schieben.							
Auch wenn ich nicht bewusst auf das System geachtet habe, bemerkte ich darüber Ereignisse, die wichtig für mich waren.							
Es war sehr schwierig, gleichzeitig etwas zu tun, das Aufmerksamkeit erforderte, während das System vibrierte.							
Ich hatte das Gefühl, dass die Vibrationen durch das System meine Konzentration gestört haben.							
Die Benutzung des Systems war mühelos.							

	1 – Trifft voll zu	2	3	4	5	6	7 – Trifft überhaupt nicht zu
Ich konnte einfach ignorieren, was mein Partner in einem bestimmten Moment tat.							
Ich habe meinem Partner große Aufmerksamkeit geschenkt.							
Das System hat häufig meine Aufmerksamkeit auf sich gezogen.							
Ich hatte oft den Eindruck, erst im Nachhinein zu realisieren, dass ich Vibrationen durch das System gespürt habe.							
Ich habe häufig aktiv versucht, mit meinem Partner in Kontakt zu treten.							
Ich hatte stark das Gefühl, dass ich eine echte Person gefühlt habe und nicht Vibrationen, die durch einen Computer erzeugt wurden.							
Ich mochte es, das System zu benutzen.							

Später fortfahren · << Zurück · Weiter >> · Umfrage verlassen und löschen

Figure A.32. Questionnaire after FEELABUZZ use, page 3, part 3.

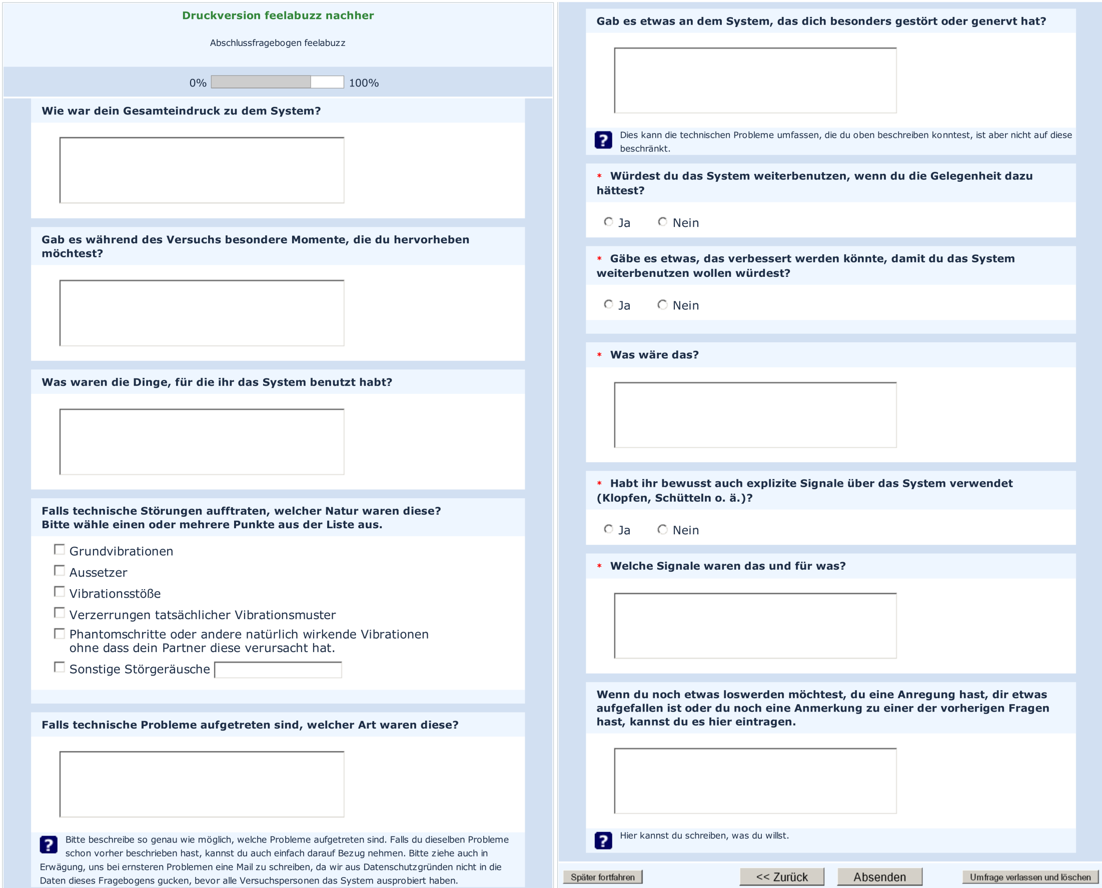

Druckversion feelabuzz nachher

Abschlussfragebogen feelabuzz

0% 100%

Wie war dein Gesamteindruck zu dem System?

Gab es während des Versuchs besondere Momente, die du hervorheben möchtest?

Was waren die Dinge, für die ihr das System benutzt habt?

Falls technische Störungen aufftraten, welcher Natur waren diese?
Bitte wähle einen oder mehrere Punkte aus der Liste aus.

- ☐ Grundvibrationen
- ☐ Aussetzer
- ☐ Vibrationsstöße
- ☐ Verzerrungen tatsächlicher Vibrationsmuster
- ☐ Phantomschritte oder andere natürlich wirkende Vibrationen ohne dass dein Partner diese verursacht hat.
- ☐ Sonstige Störgeräusche

Falls technische Probleme aufgetreten sind, welcher Art waren diese?

Bitte beschreibe so genau wie möglich, welche Probleme aufgetreten sind. Falls du dieselben Probleme schon vorher beschrieben hast, kannst du auch einfach darauf Bezug nehmen. Bitte ziehe auch in Erwägung, uns bei ernsteren Problemen eine Mail zu schreiben, da wir aus Datenschutzgründen nicht in die Daten dieses Fragebogens gucken, bevor alle Versuchspersonen das System ausprobiert haben.

Gab es etwas an dem System, das dich besonders gestört oder genervt hat?

Dies kann die technischen Probleme umfassen, die du oben beschreiben konntest, ist aber nicht auf diese beschränkt.

*** Würdest du das System weiterbenutzen, wenn du die Gelegenheit dazu hättest?**

○ Ja ○ Nein

*** Gäbe es etwas, das verbessert werden könnte, damit du das System weiterbenutzen wollen würdest?**

○ Ja ○ Nein

*** Was wäre das?**

*** Habt ihr bewusst auch explizite Signale über das System verwendet (Klopfen, Schütteln o. ä.)?**

○ Ja ○ Nein

*** Welche Signale waren das und für was?**

Wenn du noch etwas loswerden möchtest, du eine Anregung hast, dir etwas aufgefallen ist oder du noch eine Anmerkung zu einer der vorherigen Fragen hast, kannst du es hier eintragen.

Hier kannst du schreiben, was du willst.

Später fortfahren | << Zurück | Absenden | Umfrage verlassen und löschen

Figure A.33. Questionnaire after FEELABUZZ use, page 4.

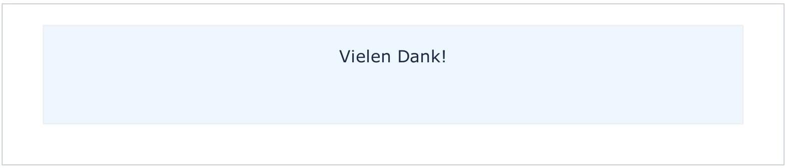

Figure A.34. Thank you message for FEELABUZZ questionnaires.

B Additional Study Material

B.1 Longitudinal Study

Scale Name; Main Source	Rev.	English Question	German Question
Relationship Salience; Connectedness Question-naire [550]	□	I often think of my partner.	Ich denke oft an meinen Partner.
	□	Even when we are not in each other's company, I often feel connected with my partner.	Ich fühle mich selbst dann mit meinem Partner verbunden, wenn wir nicht zusammen sind.
	□	I am often aware of the relationship with my partner.	Ich mache mir meine Beziehung zu meinem Partner oft bewusst.
	□	Aside from our contact via the system, I often feel connected with my partner.	Abgesehen vom Kontakt über das System, fühlte ich mich oft mit meinem Partner verbunden.

Scale Name; Main Source	Rev.	English Question	German Question
Shared Understandings; Connectedness Questionnaire [550]	☐	I feel that my partner shares my interests and ideas.	Ich habe den Eindruck, dass mein Partner meine Interessen und Ideen teilt.
	☐	I feel I have a lot in common with my partner.	Ich habe das Gefühl, dass ich viel mit meinem Partner gemein habe.
	☐	I feel on the same wavelength with my partner.	Ich bin mit meinem Partner auf derselben Wellenlänge.
Knowing Each Others' Experiences Connectedness Questionnaire [550]	☐	I often know what my partner feels.	Ich weiß oft, was mein Partner fühlt.
	☐	I often know what my partner thinks.	Ich weiß oft, was mein Partner denkt.
	☐	I feel that my partner often knows what I think.	Ich habe das Gefühl, dass mein Partner oft weiß, was ich denke.
	☐	I sense that my partner often knows what I feel.	Ich merke, dass mein Partner oft weiß, was ich denke.[1]

[1] A de facto doubling of items happened accidentally here (the German question should have been "*Ich merke, dass mein Partner oft weiß, was ich* fühle.") but we decided not to choose one over the other for the scale. We did not make any voluntary modifications to this scale that would change its α from what it was before.

Scale Name; Main Source	Rev.	English Question	German Question
Feelings of Closeness; Connectedness Questionnaire [550]	☐	My relationship with my partner is very close, compared with all other relationships.	Meine Beziehung zu meinem Partner ist sehr eng, verglichen mit all meinen anderen Beziehungen (mit sowohl Männern als auch Frauen).
	☐	I feel I can talk about anything with my partner.	Ich kann mit meinem Partner über alles reden.
Isolation/ Aloneness; Networked Minds [40] (Copresence)	☒	I often felt as if I was all alone.	Ich fühlte mich häufig vollkommen allein.
	☒	I think my partner often felt alone.	Ich denke, mein Partner fühlte sich häufig allein.
Mutual Awareness; Networked Minds [40] (Copresence)	☐	I often got the feeling of sharing a space with my partner.	Ich hatte häufig das Gefühl mich in derselben Umgebung wie mein Partner zu befinden.
	☒	I was often aware that my partner and I were at different places.	Mir war häufig bewusst, dass mein Partner und ich uns an unterschiedlichen Orten befanden.
	☒	I hardly noticed my partner.	Ich habe meinen Partner kaum wahrgenommen.

Scale Name; Main Source	Rev.	English Question	German Question
	☐	I think my partner often got the feeling of sharing a space with me.	Ich denke, dass mein Partner oft das Gefühl hatte, dass wir uns gemeinsam in derselben Umgebung befanden.
	☐	I often felt that my partner was aware of my presence.	Ich hatte häufig das Gefühl, dass sich mein Partner meiner Anwesenheit bewusst war.
	☐	It was as if we were both in the same environment.	Es war, als wären wir beide in derselben Umgebung gewesen.
Perceived Emotional Contagion/ Empathy; Networked Minds [40] (Psychological Involvement)	☐	I was influenced by my partner's moods.	Ich wurde teilweise von den Stimmungen meines Partners beeinflusst.
	☐	I had the feeling that my partner was influenced by my moods.	Ich hatte das Gefühl, mein Partner wurde teilweise von meinen Stimmungen beeinflusst.
Behavioural Interdependence; Networked Minds [40] (Behavioural Engagement)	☐	What my partner did affected what I did.	Was mein Partner tat beeinflusste, was ich tat.

Scale Name; Main Source	Rev.	English Question	German Question
	☐	What I did affected what my partner did.	Was ich tat beeinflusste, was mein Partner tat.
Perceived Other's Copresence; Nowak and Biocca [401]	☐	My partner was intensely involved in our interaction.	Mein Partner hat sich stark an unserer Interaktion beteiligt.
	☒	My partner communicated coldness rather than warmth.	Mein Partner übermittelte generell eher Kälte anstelle von Warmherzigkeit.
	☒	My partner created a sense of distance between us.	Mein Partner wirkte häufig distanziert.
	☐	I felt close to my partner.	Ich fühlte mich meinem Partner nah.
	☐	I felt connected to my partner.	Ich fühlte mich mit meinem Partner verbunden.
	☐	My partner created a sense of closeness between us.	Mein Partner hat ein Gefühl der Nähe zwischen uns erzeugt.
Telepresence; Nowak and Biocca [401]	☐	The experience with the system was very intense.	Das Erlebnis durch das System war sehr intensiv.
	☐	I felt as if I really was in the environment that I heard.	Ich hatte das Gefühl, mich tatsächlich in der Umgebung zu befinden, die ich gehört habe.

Scale Name; Main Source	Rev.	English Question	German Question
	☐	I had very much the feeling of being with my partner.	Ich hatte sehr stark das Gefühl, bei meinem Partner zu sein.
	☐	I had very much the feeling that my partner was with me.	Ich hatte sehr stark das Gefühl, dass mein Partner bei mir war.
	☐	I had the impression that my partner actually lived upstairs.	Ich hatte den Eindruck, dass mein Partner tatsächlich über mir wohnt.
Social Presence; Short et al. [494]	☐	I got a good idea of how my partner at the other end are reacted.	Ich bekam einen guten Eindruck, wie mein Partner am anderen Ende reagiert.
	☐	I got the real impression of a personal contact with my partner.	Ich bekam den echten Eindruck eines persönlichen Kontakts mit meinem Partner über die Verbindung durch das System.
	☐	I got a good "feel" for the other person.	Man bekommt ein gutes „Gefühl" für die Person am anderen Ende.
	☒	I couldn't get to know people very well if I only met them over this system.	Ich würde Leute nicht gut kennenlernen, wenn ich sie nur über dieses System träfe.

Scale Name; Main Source	Rev.	English Question	German Question
Emotional Expressiveness; ABCCT [584]	☐	Communicating with my partner using the system helped me tell how (s)he was feeling that day.	Mit meinem Partner über das System zu kommunizieren half mir herauszufinden, wie er sich an dem Tag fühlte.
	☐	Communicating with my partner using the system helped me let him/her know how I was feeling.	Mit meinem Partner über das System zu kommunizieren half mir ihn wissen zu lassen, wie ich mich fühtle.
Emotional Expressiveness (cont.); ABC-Q [257] (Recognition scale)	☐	Because of the system I could better identify with my partner.	Das System half mir, mich mehr mit meinem Partner zu identifizieren.
Staying in Touch; ABC-Q [257]	☐	The system helped me to stay in touch with my partner.	Das System half mir, mit meinem Partner in Kontakt zu bleiben.
	☐	The system made me feel involved in my partner's life.	Das System ließ mich besser am Leben meines Partners teilhaben.
Engagement &Playfulness; ABCCT [584]	☒	It was boring to use the system with my partner.	Es war langweilig, das System mit meinem Partner zu benutzen.
	☐	I was excited about using the system with my partner.	Es war aufregend, das System mit meinem Partner zu benutzen.

Scale Name; Main Source	Rev.	English Question	German Question
	☐	I had fun with my partner using the system.	Es machte Spaß, das System zusammen mit meinem Partner zu benutzen.
Presence in Absence; ABCCT [584]	☐	Communicating with my partner using the system helped me feel closer to him/her.	Mit meinem Partner über das System zu kommunizieren half mir, mich ihm näher zu fühlen.
	☐	After we were done communicating, I still kept thinking back to something my partner shared using the system.	Auch wenn wir nicht mehr über das System kommunizierten, dachte ich oft an etwas zurück, das mein Partner mit dem System übermittelt hatte.
	☐	Communicating using the system created a connection that lasts beyond the duration of the exchange.	Durch die Kommunikation mit meinem Partner über das System entstand eine Verbindung, die über die Dauer des Austauschs hinaus wirkte.
	☐	Communicating with my partner using the system helped me feel more connected to him/her.	Mit meinem Partner über das System zu kommunizieren half mir, mich mehr mit ihm verbunden zu fühlen.

Scale Name; Main Source	Rev.	English Question	German Question
Opportunity for Social Support; ABCCT [584]	☐	Communicating with my partner using the system helped me provide him/her with social support.	Mit meinem Partner über das System zu kommunizieren half mir, ihm soziale Unterstützung zu bieten.
	☐	Communicating with me using the system helped my partner be there for me when I need him/her.	Das System half meinem Partner für mich da zu sein, wenn ich ihn brauchte.
	☐	Communicating with my partner using the system when I was having a bad day helped me feel better.	Mit meinem Partner über das System zu kommunizieren half mir mich besser zu fühlen, wenn ich einen schlechten Tag hatte.
	☐	Communicating with my partner using the system helped me feel less worried about something.	Mit meinem Partner über das System zu kommunizieren hat mir geholfen, mir über etwas weniger Sorgen zu machen.
Feeling Obligated; ABCCT [584]	☒	I worried that my partner felt obligated to contact me using the system.	Ich machte mir Sorgen, dass mein Partner sich verpflichtet fühlt, mich über das System zu kontaktieren.

Scale Name; Main Source	Rev.	English Question	German Question
	⊠	I had to communicate with my partner using the system even when I didn't want to.	Ich musste mit meinem Partner über das System kommunizieren, auch wenn ich gar nicht wollte.
	⊠	I felt guilty if I didn't respond to my partner when I perceived something using the system.	Ich fühlte mich schuldig, wenn ich auf etwas nicht reagiert habe, das ich über das System wahrgenommen habe.
	⊠	When I tried to communicate with my partner using the system, I expected a response.	Wenn ich versucht habe mit meinem Partner in Kontakt zu treten, erwartete ich eine Reaktion.
Unmet Expectations; ABCCT [584]	⊠	I was disappointed when my partner wasn't there when I tried to contact him/her using the system.	Ich war enttäuscht, wenn mein Partner nicht da war, wenn ich es war.
	⊠	I was disappointed when it took my partner too long to respond over the system.	Ich war enttäuscht, wenn mein Partner zu lange zum Reagieren gebraucht hat, wenn ich versucht habe, ihn über das System zu kontaktieren.

Scale Name; Main Source	Rev.	English Question	German Question
	⊠	I worried that I was not meeting my partner's expectations for our contact using the system.	Ich machte mir Sorgen, dass ich die Erwartungen meines Partners nicht würde erfüllen können, was unseren Kontakt über das System angeht.
Threat to Privacy; ABCCT [584]	⊠	I worried that my partner might learn something using the system that I want to keep secret.	Ich machte mir Sorgen, dass mein Partner etwas mit dem System hätte herausfinden können, das ich geheim halten wollte.
	⊠	I worried about my privacy while my partner and I were using the system together.	Ich machte mir Sorgen um meine Privatsphäre, während wir das System benutzt haben.
	⊠	I worried that others might overhear something that my partner and I shared using the system.	Ich machte mir Sorgen, dass Dritte etwas von dem würden belauschen können, das wir über das System austauschten.
	⊠	I worried that I was violating my partner's privacy during our contact using the system.	Ich machte mir Sorgen, dass ich die Privatsphäre meines Partners durch unsere Benutzung des Systems hätte verletzen können.

Scale Name; Main Source	Rev.	English Question	German Question
	☒	The system revealed more about my partner than (s)he might like.	Das System verriet mir mehr über meinen Partner als ihm vielleicht lieb ist.
Cognitive Load	☒	The system distracted me from important tasks.	Das System lenkte mich von wichtigen Aufgaben ab.
	☐	I was mostly unaware of the system running.	Mir ist meist nicht aufgefallen, dass das System lief.
	☐	I didn't mind that the system was running.	Es hat mich nicht gestört, dass das System lief.
	☐	It was easy for me to push the system to the back of my mind.	Es fiel mir leicht, das System in den Hintergrund zu schieben.
	☐	I recognised important events with the system, even when I was not actively monitoring it.	Auch wenn ich nicht bewusst auf das System geachtet habe, bemerkte ich darüber Ereignisse, die wichtig für mich waren.
	☒	It was hard to concentrate while the system was producing noises.	Es war sehr schwierig, gleichzeitig etwas zu tun, das Aufmerksamkeit erforderte, während das System Geräusche produzierte.

Scale Name; Main Source	Rev.	English Question	German Question
	☒	I had the feeling that the noises produced by the system disturbed my concentration.	Ich hatte das Gefühl, dass die Geräusche durch das System meine Konzentration gestört haben.
	☐	Using the system was effortless.	Die Benutzung des Systems war mühelos.
	☐	It was easy for me to ignore what my partner was doing in any moment.	Ich konnte einfach ignorieren, was mein Partner in einem bestimmten Moment tat.
	☒	I payed high attention to my partner using the system.	Ich habe meinem Partner große Aufmerksamkeit geschenkt.
	☒	The system frequently drew all my attention to itself.	Das System hat häufig meine Aufmerksamkeit auf sich gezogen.
	☐	I often only realised in hindsight that I heard something through the system.	Ich hatte oft den Eindruck, erst im Nachhinein zu realisieren, dass ich Geräusche durch das System gehört habe.
Active Contact	☐	I frequently tried to actively contact my partner using the system.	Ich habe häufig aktiv versucht, mit meinem Partner in Kontakt zu treten.

Scale Name; Main Source	Rev.	English Question	German Question
Real Person; Casanueva [69]	□	I had the strong feeling that I heard an actual person over the system.	Ich hatte stark das Gefühl, dass ich eine echte Person gehört habe und nicht Geräusche, die durch einen Computer erzeugt wurden.
Pleasurability	□	I liked using the system.	Ich mochte es, das System zu benutzen.

Table B.1. Overview of which items make up which of the scales used. Also, the English translations are shown alongside the German questions used in the questionnaires. Where the German questions themselves stem from an English original, the original item is presented instead of a maybe more fitting back translation. Only the female version of the German questionnaire items is shown. In the male version, "*Partner*" is replaced by "*Partnerin*". The "Inv." column declares whether or not an item is inverted. Scale inversion is solely a function of the inversion of its constituents.

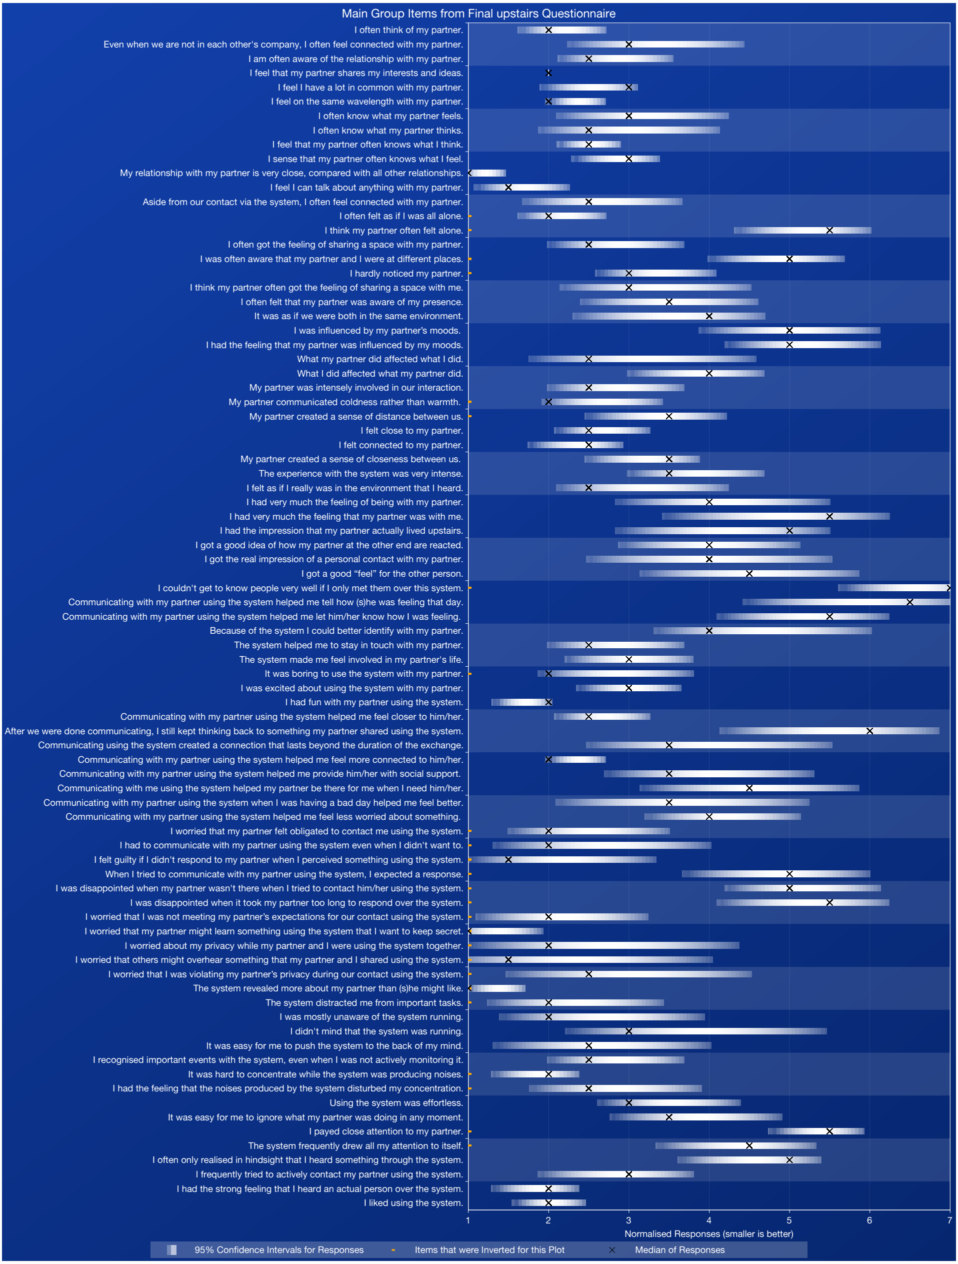

Figure B.1. Main group items from the final upstairs questionnaire $Q_{\beta\gamma}$. An expected 4.15 items have a population mean μ outside the confidence intervals. The items marked orange were inverted so that leftmost always means “better”.

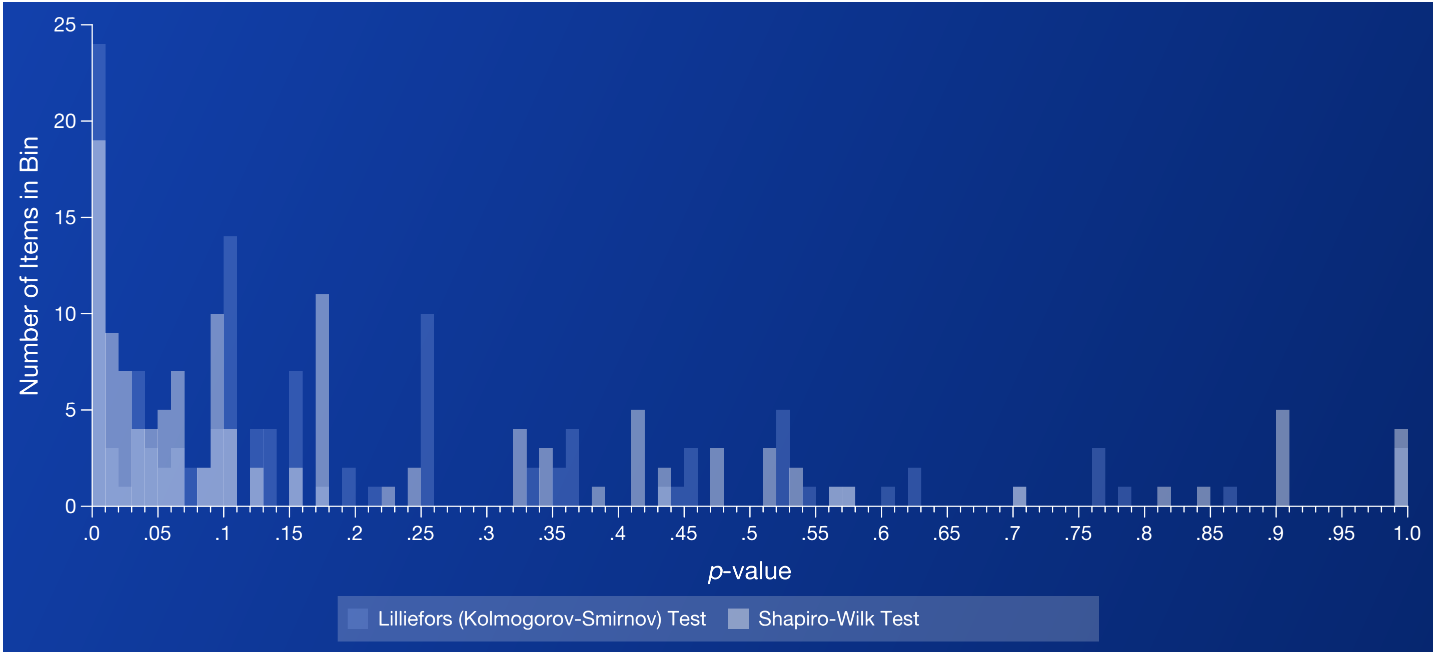

Figure B.2. Histograms for two normality tests on the items of the final upstairs questionnaire $Q_{\beta\gamma}$. The bins are of size .01 (100 bins in total). There are 127 items in total for which normality was tested with the Lilliefors [331] and Shapiro-Wilk [489, 490] tests. The null hypothesis these tests try to reject is that a given sample was drawn from a normally distributed random variable.

	CL1	CL2	CL3	CL4	CL5	CL6	CL7	CL8	CL9	CL10	CL11	CL12
The system distracted me from important tasks (CL1)	1	.2	−.4	−.02	−.3	.41	.48	−.22	.3	.38	.23	−.65
I was mostly unaware of the system running (CL2)	.2	1	−.22	.7	−.42	.1	.75	.09	.44	.51	.81	.28
I didn't mind that the system was running (CL3)	−.4	−.22	1	.27	−.01	−.38	−.5	−.77	.42	−.18	−.37	.4
It was easy for me to push the system to the back of my mind (CL4)	−.02	.7	.27	1	.06	.24	.63	−.26	.78	−.04	.29	.09
I recognised important events with the system, even when I was not actively monitoring it (CL5)	−.3	−.42	−.01	.06	1	.64	.1	.35	−.25	−.98	−.33	−.49
It was hard to concentrate while the system was producing noises (CL6)	.41	.1	−.38	.24	.64	1	.69	.33	−.03	−.54	.26	−.76
I had the feeling that the noises produced by the system disturbed my concentration (CL7)	.48	.75	−.5	.63	.1	.69	1	.28	.35	.06	.63	−.39
Using the system was effortless (CL8)	−.22	.09	−.77	−.26	.35	.33	.28	1	−.72	−.2	.36	−.07
It was easy for me to ignore what my partner was doing in any moment (CL9)	.3	.44	.42	.78	−.25	−.03	.35	−.72	1	.22	−.07	−.06
I payed close attention to my partner (CL10)	.38	.51	−.18	−.04	−.98	−.54	.06	−.2	.22	1	.42	.4
The system frequently drew all my attention to itself (CL11)	.23	.81	−.37	.29	−.33	.26	.63	.36	−.07	.42	1	.24
I often only realised in hindsight that I heard something through the system (CL12)	−.65	.28	.4	.09	−.49	−.76	−.39	−.07	−.06	.4	.24	1

Table B.2. Item correlation matrix for the Cognitive Load items of the $Q_{\beta\gamma}$ questionnaire.

Scale	Mean	Median
Relationship Salience	2.75	2.5
Shared Understandings	2.28	2
Knowing Each Others' Experiences	2.88	3
Feelings of Closeness	1.42	1
Overall Social Connectedness	2.47	2
Isolation/Aloneness	3.67	3
Mutual Awareness	3.56	3.5
Perceived Emotional Contagion/Empathy	5.08	5
Behavioural Interdependence	3.5	3
Perceived Other's Copresence	2.83	3
Telepresence	4.03	4
Social Presence (Direct Questions)	4.75	5
Social Presence (Semantic Differential Technique)	3.47	3
Overall Social Presence	4.04	4
Emotional Expressiveness	5.22	6
Staying in Touch	2.92	3
Engagement & Playfulness	2.5	2
Presence in Absence	3.63	3
Opportunity for Social Support	4.08	4
Feeling Obligated	3.04	2.5
Unmet Expectations	4.17	5
Threat to Privacy	2.13	1
Cognitive Load	3.38	3
Vexatiousness	2.74	2
Ability to Ignore	3.71	3
Overall Copresence	3.35	3
Networked Minds Copresence	3.58	3
Emotion Transmission	5.17	5.5

Table B.3. Scale results from $Q_{\beta\gamma}$ as visualised in Figure 6.6 on page 160. These scales are mainly adapted from the Networked Minds Questionnaire [39, 40, 41], Nowak and Biocca [401], the ABCCT [584] and ABC-Q [257].

Item	Item-Total Correlation	Cronbach's α without Item	Component 1	Component 2
I got a good idea of how my partner at the other end are reacted.	.39	.88	.52	.09
I got the real impression of a personal contact with my partner.	.76	.85	.89	.14
I got a good "feel" for the other person.	.88	.83	.76	.56
I couldn't get to know people very well if I only met them over this system.	.66	.86	.93	−.17
personal—impersonal	.28	.88	−.15	.98
sociable—unsociable	.61	.88	.47	.44
sensitive—insensitive	.85	.84	.77	.45
warm—cold	.65	.86	.35	.92
meaningful—meaningless	.77	.85	.83	.22

Table B.4. Combined items from both Social Presence scales with corrected item-total correlations, combined Cronbach's α without the item and factor loadings for the first two principal components after a varimax rotation.

Scale	Social Desirability Correlation	*p*-Value
Relationship Questionnaire	.09	.87
Relationship Closeness Inventory	.26	.612
Social Presence (Semantic Differential Technique)	.26	.626
Relationship Salience	.02	.968
Shared Understandings	.2	.704
Knowing Each Others' Experiences	.07	.901
Feelings of Closeness	−.92	.01
Isolation/Aloneness	.21	.688
Mutual Awareness	.17	.743
Perceived Emotional Contagion/Empathy	.33	.525
Behavioural Interdependence	.17	.741
Perceived Other's Copresence	.39	.439
Telepresence	−.03	.951
Social Presence (Direct Questions)	.25	.638
Emotional Expressiveness	.25	.635
Staying in Touch	−.44	.383
Engagement & Playfulness	−.24	.645
Presence in Absence	.03	.959
Opportunity for Social Support	.17	.743
Feeling Obligated	−.51	.305
Unmet Expectations	−.16	.762
Threat to Privacy	−.42	.41
Cognitive Load	.18	.732
Social Connectedness	−.13	.802
Ability to Ignore	.6	.21
Vexatiousness	.01	.988

Table B.5. Correlations with the Social Desirability scale [518].

Item	Anti-Image Autocorrelation
The system distracted me from important tasks.	.582
I was mostly unaware of the system running.	.141
I didn't mind that the system was running.	.203
It was easy for me to push the system to the back of my mind.	.532
I recognised important events with the system, even when I was not actively monitoring it.	.404
It was hard to concentrate while the system was producing vibrations.	.612
I had the feeling that the vibrations produced by the system disturbed my concentration.	.688
Using the system was effortless.	.396
It was easy for me to ignore what my partner was doing in any moment.	.264
I payed close attention to my partner.	.159
The system frequently drew all my attention to itself.	.351
I often only realised in hindsight that I had felt something through the system.	.223

Table B.6. Diagonal of the anti-image correlation matrix of the Cognitive Load scale on the final FEELABUZZ questionnaire.

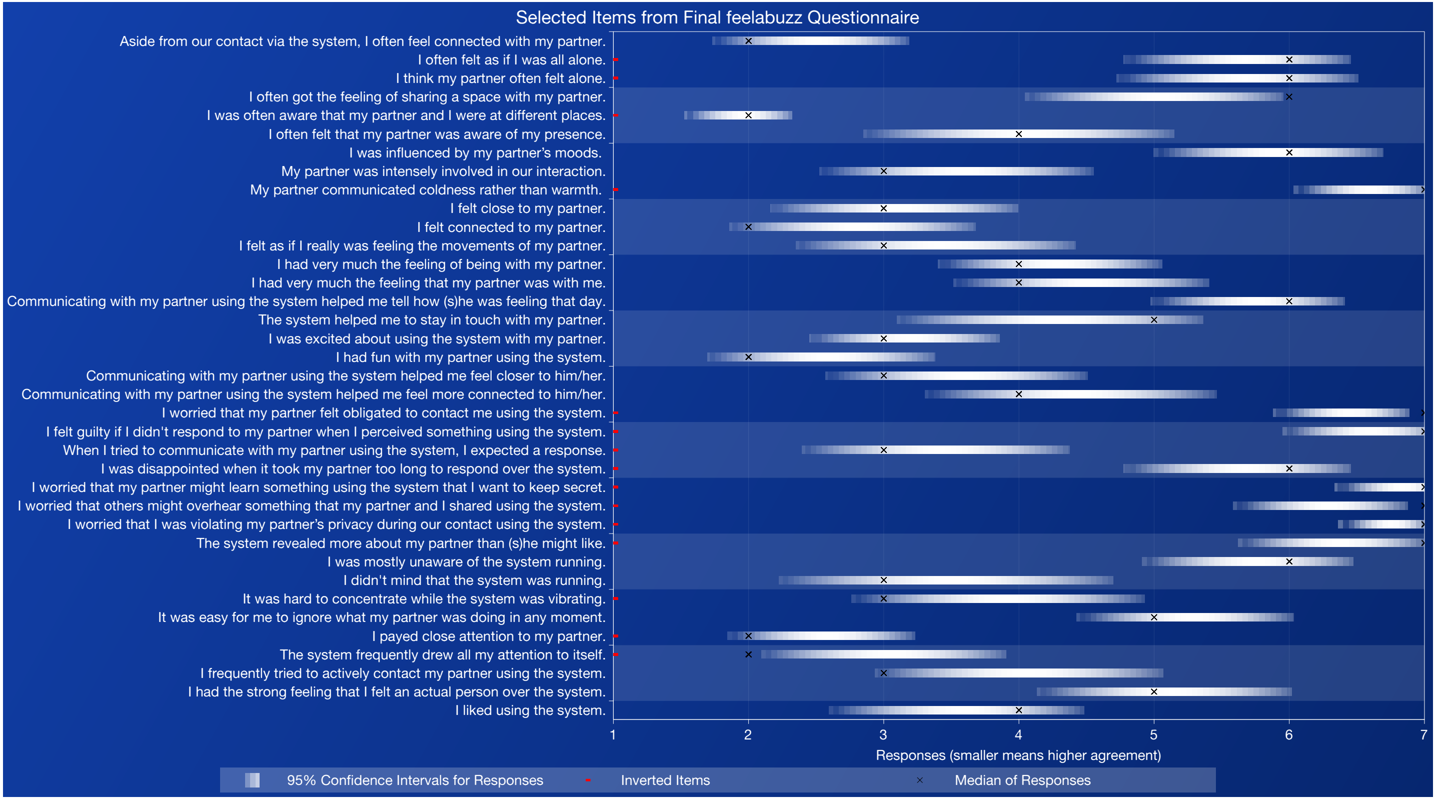

Figure B.3. Corresponding items to Figure 6.7 on page 161 for the FEELABUZZ final questionnaire.

Figure B.4. Corresponding items to Figure B.1 on page 433 for the FEELABUZZ final questionnaire.

Item	Item Number	Item-Total Correlation	Cronbach's α without Item
The system frequently drew all my attention to itself.	11	.58	.22
I often only realised in hindsight that I had felt something through the system.	12	.50	.25
It was easy for me to ignore what my partner was doing in any moment.	9	.49	.28
It was hard to concentrate while the system was producing vibrations.	6	.47	.24
It was easy for me to push the system to the back of my mind.	4	.46	.28
I was mostly unaware of the system running.	2	.42	.30
I had the feeling that the vibrations produced by the system disturbed my concentration.	7	.34	.30
The system distracted me from important tasks.	1	.30	.33
I didn't mind that the system was running.	3	−.24	.54
I payed close attention to my partner.	10	−.37	.50
I recognised important events with the system, even when I was not actively monitoring it.	5	−.41	.55
Using the system was effortless.	8	−.55	.53

Table B.7. Cognitive load items of the final FEELABUZZ questionnaire, sorted by their corrected item-total correlation. The Kendall τ distance [284] (number of swaps of adjacent pairs) to Table 6.9 on page 165 is 18 which normalises to .27.

Item	Item-Total Correlation	Cronbach's α without Item
warm—cold	.84	.68
sensitive—insensitive	.82	.69
meaningful—meaningless	.66	.75
personal—impersonal	.55	.8
sociable—unsociable	.35	.83

Table B.8. Social presence semantic differential technique items sorted by their corrected item-total correlation.

B.2 Pre-Study FEELABUZZ

On the following page, you will find a complete list of all the multiple choice items asked during the FEELABUZZ pre-study and histograms of the answers that were given. See Section 4.3 for more information on the study itself.

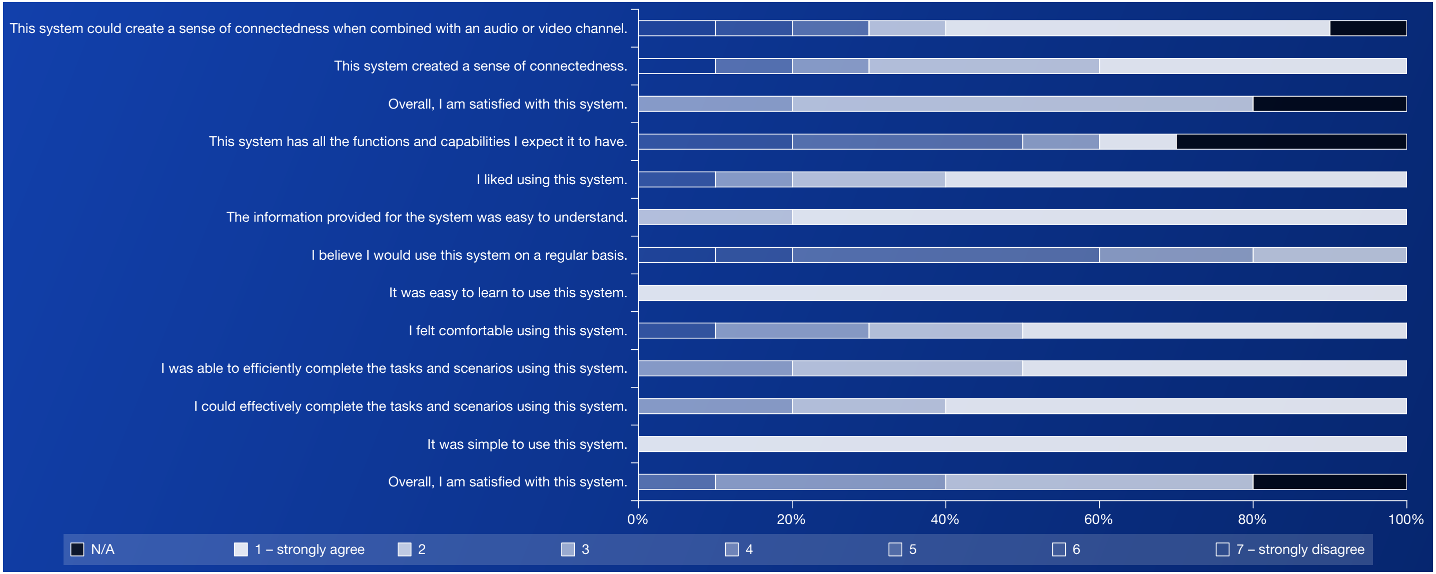

Figure B.5. All questions from the FEELABUZZ pre-study questionnaire and distribution of responses.

C Umber Additional Material

C.1 Two-Dimensional Example of Edge-Gaussians

We will illustrate the multivariate edge-Gaussians from Section 7.7.3.4 for $k = 2$, i. e. $\boldsymbol{x} \in \mathbb{R}^2$. We begin with the bivariate Gaussian where $\boldsymbol{\Sigma} = \boldsymbol{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $\boldsymbol{\mu} = \mathbf{0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$.

$$P(x_1, x_2) = \frac{1}{2\pi} \exp\left(-\frac{x_1^2 + x_2^2}{2}\right) \tag{C.1}$$

Dropping the normalisation, this becomes

$$P'(x_1, x_2) = \exp\left(-\frac{x_1^2 + x_2^2}{2}\right). \tag{C.2}$$

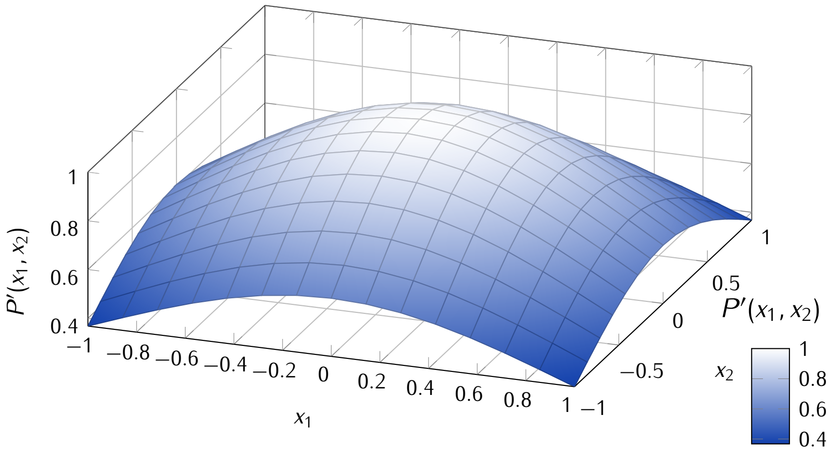

We apply two correction terms to straighten the four edges of the square.

$$\begin{aligned} f_1(x_1, x_2) =& P'(x_1, x_2) - P'(1, x_2) - P'(x_1, 1) \\ =& \exp\left(-\frac{x_1^2 + x_2^2}{2}\right) - \exp\left(-\frac{x_2^2}{2} - \frac{1}{2}\right) - \exp\left(-\frac{x_1^2}{2} - \frac{1}{2}\right) \end{aligned} \tag{C.3}$$

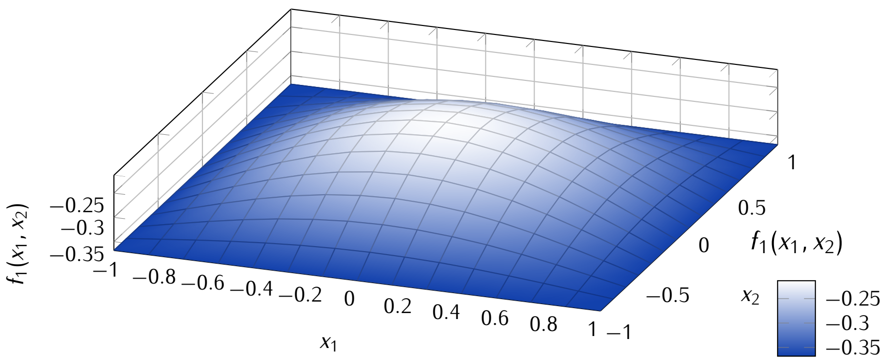

We then apply the scaling to get the desired edge length. Therefore for $f_2(x_1, x_2)$, $x_1, x_2 \in [-s_1, s_1] \times [-s_2, s_2]$:

$$\begin{aligned} f_2(x_1, x_2) &= f_1\left(\frac{x_1}{s_1}, \frac{x_2}{s_2}\right) \\ &= \exp\left(-\frac{x_1^2}{2s_1^2} - \frac{x_2^2}{2s_2^2}\right) - \exp\left(-\frac{x_2^2}{2s_2^2} - \frac{1}{2}\right) \\ &\quad - \exp\left(-\frac{x_1^2}{2s_1^2} - \frac{1}{2}\right) \end{aligned} \tag{C.4}$$

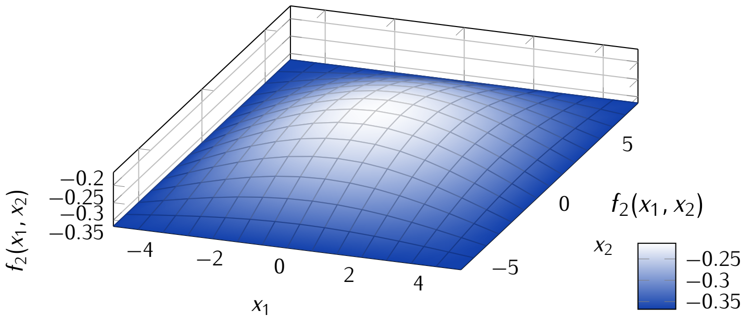

Then the translation is applied to centre the orthotope around $\boldsymbol{\mu}$, so for $f_3(x_1, x_2)$, $x_1, x_2 \in [\mu_1 - s_1, \mu_1 + s_1] \times [\mu_2 - s_2, \mu_2 + s_2]$:

$$\begin{aligned} f_3(x_1, x_2) &= f_2(\mu_1 - x_1, \mu_2 - x_2) \\ &= \exp\left(-\frac{(\mu_1 - x_1)^2}{2s_1^2} - \frac{(\mu_2 - x_2)^2}{2s_2^2}\right) \\ &\quad - \exp\left(-\frac{(\mu_1 - x_1)^2}{2s_1^2} - \frac{1}{2}\right) \\ &\quad - \exp\left(-\frac{(\mu_2 - x_2)^2}{2s_2^2} - \frac{1}{2}\right) \end{aligned} \tag{C.5}$$

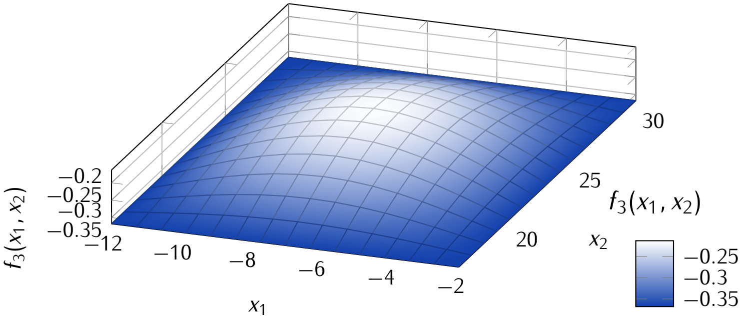

This is then flipped vertically, turning it upside-down.

$$\begin{aligned} f_4(x_1, x_2) =& f_3(\mu_1, \mu_2) - f_3(x_1, x_2) \\ =& 1 - \frac{2}{\sqrt{e}} + \exp\left(-\frac{(\mu_1 - x_1)^2}{2s_1^2} - \frac{1}{2}\right) \\ & + \exp\left(-\frac{(\mu_2 - x_2)^2}{2s_2^2} - \frac{1}{2}\right) \\ & - \exp\left(-\frac{(\mu_2 - x_2)^2}{2s_2^2} - \frac{(\mu_1 - x_1)^2}{2s_1^2}\right) \end{aligned} \tag{C.6}$$

which we can also write as

$$\begin{aligned} f_4(x_1, x_2) = & - \exp\left(-\frac{(2x_2 - x_{21} - x_{22})^2}{2|x_{21} - x_{22}|^2} - \frac{(2x_1 - x_{11} - x_{12})^2}{2|x_{11} - x_{12}|^2}\right) \\ & + \exp\left(-\frac{(2x_1 - x_{11} - x_{12})^2}{2|x_{11} - x_{12}|^2} - \frac{1}{2}\right) \\ & + \exp\left(-\frac{(2x_2 - x_{21} - x_{22})^2}{2|x_{21} - x_{22}|^2} - \frac{1}{2}\right) \\ & - \frac{2}{\sqrt{e}} + 1 \end{aligned} \tag{C.7}$$

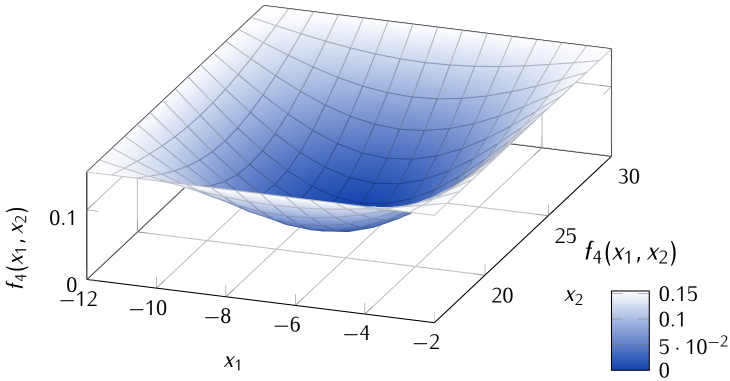

Finally, we normalise to unit integral. If you can do without this step in your implementation because you do not need this property, we would advise anyone implementing this to do so since this is the step where things get a bit unwieldy.

$$f_5(x_1, x_2) = f_4(x_1, x_2) \left(\int_{x_{11}}^{x_{12}} \int_{x_{21}}^{x_{22}} f_4(x_1, x_2)\, dx_2\, dx_1 \right)^{-1}$$

$$= \left(-2\sqrt{e} \left(\exp\left(-\frac{(2x_1 - x_{11} - x_{12})^2}{2|x_{11} - x_{12}|^2} - \frac{1}{2} \right) \right.\right.$$

$$- \exp\left(-\frac{(2x_2 - x_{21} - x_{22})^2}{2|x_{21} - x_{22}|^2} - \frac{(2x_1 - x_{11} - x_{12})^2}{2|x_{11} - x_{12}|^2} \right)$$

$$\left.\left. + \exp\left(-\frac{(2x_2 - x_{21} - x_{22})^2}{2|x_{21} - x_{22}|^2} - \frac{1}{2} \right) - \frac{2}{\sqrt{e}} + 1 \right) \right)$$

$$\left(4\sqrt{\pi}(x_{11} - x_{12})x_{21} - 4(x_{11} - x_{12})x_{22} \right.$$

$$
\begin{aligned}
&+\sqrt{e}\Bigg(\pi|x_{22}-x_{21}||x_{12}-x_{11}|\operatorname{erf}\left(-\frac{\sqrt{2}\,(x_{21}-x_{22})}{2|x_{22}-x_{21}|}\right)\\
&\operatorname{erf}\left(-\frac{\sqrt{2}\,(x_{11}-x_{12})}{2|x_{12}-x_{11}|}\right)-2(x_{11}-x_{12})x_{21}+2(x_{11}-x_{12})x_{22}\Bigg)\\
&+\sqrt{2}\Bigg(x_{21}|x_{12}-x_{11}|\operatorname{erf}\left(-\frac{\sqrt{2}\,(x_{11}-x_{12})}{2|x_{12}-x_{11}|}\right)\\
&-x_{22}|x_{12}-x_{11}|\operatorname{erf}\left(-\frac{\sqrt{2}\,(x_{11}-x_{12})}{2|x_{12}-x_{11}|}\right)\\
&+(x_{11}-x_{12})\,|x_{22}-x_{21}|\operatorname{erf}\left(-\frac{\sqrt{2}(x_{21}-x_{22})}{2|x_{22}-x_{21}|}\right)\Bigg)\Bigg)^{-1}
\end{aligned}
\tag{C.8}
$$

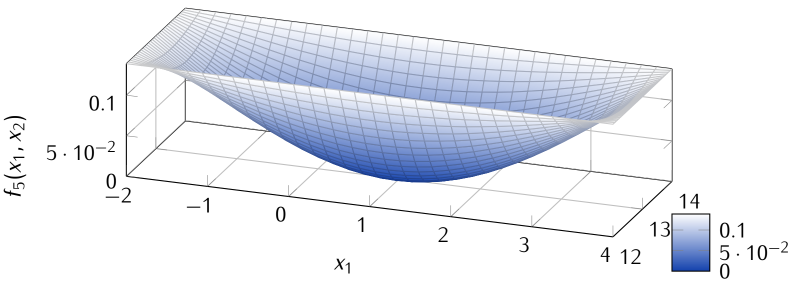

C.2 Lumber Features

Feature Name	Feature Description
`camera-Average Total`	Brightness in camera image
`camera-Faces Seen`	Number of Faces detected
`camera-Average Squared Total`	Movement in camera image (squared)

Feature Name	Feature Description
camera-Average Total_n-1	Brightness in camera image (1 time step earlier)
camera-Average Total_n-2	Brightness in camera image (2 time steps earlier)
camera-Average Total_n-3	Brightness in camera image (3 time steps earlier)
camera-Average Total_n-4	Brightness in camera image (4 time steps earlier)
camera-Average Total_n-5	Brightness in camera image (5 time steps earlier)
camera-Faces Seen_n-1	Number of faces detected (1 time step earlier)
camera-Faces Seen_n-2	Number of faces detected (2 time steps earlier)
camera-Faces Seen_n-3	Number of faces detected (3 time steps earlier)
camera-Faces Seen_n-4	Number of faces detected (4 time steps earlier)
camera-Faces Seen_n-5	Number of faces detected (5 time steps earlier)
camera-Average Squared Total_n-1	Movement in camera image (squared, 1 time step earlier)
camera-Average Squared Total_n-2	Movement in camera image (squared, 2 time steps earlier)
camera-Average Squared Total_n-3	Movement in camera image (squared, 3 time steps earlier)
camera-Average Squared Total_n-4	Movement in camera image (squared, 4 time steps earlier)
camera-Average Squared Total_n-5	Movement in camera image (squared, 5 time steps earlier)

Feature Name	Feature Description
camera-Average Total_mean_of_5	Brightness in camera image, averaged over the last 5 time steps
camera-Faces Seen_mean_of_5	Number of faces detected, averaged over the last 5 time steps
camera-Average Squared Total_mean_of_5	Movement in camera image (squared), averaged over the last 5 time steps
mouse-keyboard-Idletime	Time since the last keyboard or mouse activity
mouse-keyboard-Idletime_mean_of_5	Time since the last keyboard or mouse activity, averaged over the last 5 time steps
audio-Energy	Audio energy
audio-Spectral Centroid	Spectral centroid (average audio frequency)
audio-Spectral Spread	Spectral spread (audio frequency variance)
audio-Energy_n-1	Audio energy (1 time step earlier)
audio-Energy_n-2	Audio energy (2 time steps earlier)
audio-Energy_n-3	Audio energy (3 time steps earlier)
audio-Energy_n-4	Audio energy (4 time steps earlier)
audio-Energy_n-5	Audio energy (5 time steps earlier)

Feature Name	Feature Description
audio-Spectral Centroid_n-1	Spectral centroid (1 time step earlier)
audio-Spectral Centroid_n-2	Spectral centroid (2 time steps earlier)
audio-Spectral Centroid_n-3	Spectral centroid (3 time steps earlier)
audio-Spectral Centroid_n-4	Spectral centroid (4 time steps earlier)
audio-Spectral Centroid_n-5	Spectral centroid (5 time steps earlier)
audio-Spectral Spread_n-1	Spectral spread (1 time step earlier)
audio-Spectral Spread_n-2	Spectral spread (2 time steps earlier)
audio-Spectral Spread_n-3	Spectral spread (3 time steps earlier)
audio-Spectral Spread_n-4	Spectral spread (4 time steps earlier)
audio-Spectral Spread_n-5	Spectral spread (5 time steps earlier)
audio-Energy_mean_of_5	Audio energy, averaged over the last 5 time steps
audio-Spectral Centroid_mean_of_5	Spectral centroid, averaged over the last 5 time steps
audio-Spectral Spread_mean_of_5	Spectral spread, averaged over the last 5 time steps
time-DayX	Time of day on a 24h clock, horizontal position of the hand
time-DayY	Time of day on a 24h clock, vertical position of the hand

Feature Name	Feature Description
time-WeekX	Horizontal position in a circular week calendar
time-WeekY	Vertical position in a circular week calendar
title-Active Window Title-token([0-9]*)_(.*)	Active window title contains '\2' (0 = false, 1 = true)
title-Active Application Name-token([0-9]*)_(.*)	Active window name contains '\2' (0 = false, 1 = true)
title-Active Application Class-token([0-9]*)_(.*)	Active window class contains '\2' (0 = false, 1 = true)

Table C.1. Feature overview of all Lumber features. The number and names of the window title features (cf. Section 7.5.2.1) vary and are therefore given as the original regular expression (http://docs.python.org/2/library/re.html). Table rows were coloured in pairs of two for better readability only; this bears no particular connection between any two features whatsoever.