

Ambient Displays Supporting Environmentally-Conscious Behavior

by

Jan Hammerschmidt

A dissertation submitted to the Faculty of Technology at the Bielefeld University
in partial fulfillment of the requirements for the degree of

Doktor der Ingenieurwissenschaften
(Dr.-Ing.)

July 2018

Abstract

In the natural interaction with our environment, humans draw on a wealth of information without consciously focusing on it. In contrast, the vast majority of interaction with modern technological devices is accomplished through focused attention.

This thesis therefore researches the use of *ambient displays*, which convey information at the periphery of attention. As a use case that has been identified as particularly suitable for such interfaces but has not yet been researched very extensively, we have focused on supporting users to act in a more environmentally-conscious way. Furthermore, despite the fact that preliminary research in cognitive psychology and attention theory hints at significant advantages of employing the auditory modality for the peripheral perception of information, the use of sound for the design of ambient displays has been largely neglected so far, which is why we have deliberately explored this approach in our research, including the use of *blended sonification*, where existing environmental sounds are used as the basis for auditory representations of information.

In our practical work, we have developed four ambient displays as research vehicles, which cover three distinct settings in order to additionally evaluate the influence of an ambient display's context of use: The *InfoPlant* is a "living" interface, which we used in a longitudinal study to give feedback on the test subjects' electricity consumption in their apartments. Although we could indeed observe a reduction in consumption, this was only the case for appliances within sight of the feedback display and highly dependent on the users' initial motivation and interest. The *Sonic Shower* gives auditory feedback on the energy and water consumption while taking a shower, and in an online survey we found that our two blended sonifications were perceived as significantly less intrusive than for example a speech-based design. The *EcoSonic system* supports users in driving in a more fuel (or energy) efficient way, and in our study we could observe that our two types of auditory feedback led to a significant reduction of consumption as well as a reduced number of glances at the visual consumption display. Finally, the *Slowification system* provides feedback about the current vehicle speed in consideration of prescribed speed limits and common driving practices based on spatial panning of the car's audio system's sound signal, and in our study, employing a virtual reality driving simulator, we found that speed limits were adhered to significantly better and also observed less deviations from the traffic lane.

As can be seen from the individual results, the use of sound can indeed be highly advantageous for ambient displays, as it does not occupy the visual perceptual channel, which is predominantly used for primary activities. Furthermore, it allows for a higher variability in the context of use, which is also of particular importance for the design of ambient displays in general, as it influences the effectiveness of the peripheral information conveyance. Finally, incorporating, or even building on, existing sensory stimuli, as is done for blended sonifications, has shown to contribute to enhancing the unobtrusiveness of ambient displays.

Acknowledgments

First and foremost, I would like to thank my supervisor Thomas Hermann for his enthusiastic and always encouraging support during this whole work, for allowing me to work freely on this project while providing me with feedback and guidance when needed, and also for the personal support and empathy.

I also want to give thanks to the whole Ambient Intelligence Group for all their input and support, and for providing such a welcoming environment including Alexander Neumann for a lengthy and inspiring initial discussion on the overall topic, Andrea Lorena Aldana Blanco for having great tandem sessions, Carsten Schwede for everything related to three.js and the Oculus Rift, Christian Leichsenring and Marian Pohling for their help to integrate the InfoPlant into the CSRA, Eckard Riedenklau and Jiajun Yang for being such great office roomies, René Tünnermann for his work on Blended Sonification, and Sebastian Zehe for his always competent help with the InfoPlant.

A special thanks goes to my family and parents for supporting me wherever possible, especially my father for his very thorough reading of my English writing.

Finally, I want to thank Sophia for all her love and support. I am grateful for being able to discuss just about anything with you, for your patience and understanding, and for all the enjoyable distractions from writing.

Table of Contents

1. Introduction	1
1.1. Interacting with Technology	1
1.2. Calm Technology	2
1.3. Ambient Information Systems	3
1.4. Research Questions	5
I. Methods & Research Background	7
2. Attention Theory	9
2.1. Introduction	9
2.2. Selective Attention	10
2.2.1. Early Selection	11
2.2.2. Late Selection	11
2.2.3. Attenuation Theory	12
2.2.4. Context and Priming	12
2.2.5. Selective Attention and Calm Technology	13
2.3. Divided Attention	14
2.3.1. Attentional-Resource Theories	14
2.3.2. Practice and Automaticity	14
2.3.3. Crossmodal Attention	15
2.4. Insights Gained for Ambient Displays	15
2.4.1. The Periphery of Attention	16
2.4.2. Amount and Type of Conveyed Information	16
2.4.3. Context and Modality	17
2.4.4. Automaticities and Learning	17
3. Ambient Information Systems	19
3.1. Introduction	19
3.2. Common Themes	20
3.2.1. Advancing Display Technology	20
3.2.2. Providing Awareness	24
3.3. Design Dimensions	28
3.3.1. Event-like or Continuous Data	28
3.3.2. Intrusiveness and Notification Level	29
3.3.3. Information Capacity	31
3.3.4. Representation	33
3.3.5. Usage Type	34

3.3.6. Modality	35
4. Sustainable Human-Computer Interaction	37
4.1. Persuasive Technology	37
4.2. Ambient Awareness	39
4.3. Sustainable Interaction Design	41
4.4. Formative User Studies	42
4.5. Pervasive and Participatory Sensing	43
5. Auditory Display and (Blended) Sonification	45
5.1. Methods for Auditory Display Design	45
5.2. Blended Sonification	48
6. Research Agenda	53
6.1. Goals and Research Questions	53
6.1.1. The Use of Sound in Ambient Displays	54
6.1.2. Complexity and Variability of the Context	54
6.1.3. Unobtrusiveness and Acceptance	54
6.1.4. Feedback, Persuasion, and Affectiveness	55
6.2. Research Prototypes	56
II. Prototype Systems	57
7. The InfoPlant	59
7.1. Introduction	59
7.2. Related Work	60
7.3. Concept of the InfoPlant	63
7.4. Construction	66
7.5. Household Study	70
7.5.1. Measuring Electricity Consumption	70
7.5.2. Feedback Design	71
7.5.3. Comparative Feedback	75
7.5.4. Study Design	75
7.5.5. Goals and Hypotheses	79
7.5.6. Results	79
7.6. Discussion and Conclusions	87
8. InfoDrops: The Sonic Shower	91
8.1. Introduction	91
8.1.1. Challenges of Interface Design in a Bathroom Context	92
8.2. Related Work	93

8.3.	The InfoDrops System	95
8.3.1.	Implementation	96
8.3.2.	Feedback Design	98
8.4.	Sonification Design	100
8.4.1.	Pitch-Based Blended Sonification	100
8.4.2.	Blended Chords Sonification	101
8.4.3.	Transient-Triggered Blended Chords Sonification	102
8.5.	Online Survey	103
8.5.1.	Study Design	105
8.5.2.	Goals	105
8.5.3.	Results	106
8.6.	Discussion and Conclusions	112
9.	The EcoSonic System	115
9.1.	Introduction	115
9.2.	Related Work	116
9.2.1.	Efficacy of Eco-Driving	116
9.2.2.	Distraction by In-Vehicle Systems	118
9.2.3.	Comparison of Existing Eco-Feedback Systems	119
9.2.4.	Auditory In-Vehicle Interaction	120
9.3.	General Design Approaches	122
9.3.1.	Direct Sonification of Fuel Consumption	122
9.3.2.	Sonification of Secondary Parameters	123
9.3.3.	Gamification	123
9.3.4.	Sonifying Advanced Support Information	123
9.3.5.	Supporting Specific Qualities of Driving	124
9.4.	Sonification Designs	124
9.4.1.	Continuous Sonification of Fuel Consumption	124
9.4.2.	Metaphorical Sonification to Support Low Rpm Driving	125
9.5.	The EcoSonic Driving Simulator	125
9.5.1.	Graphical User Interface	126
9.5.2.	Internal Simulation	128
9.5.3.	Traffic Violations and Attention Task	130
9.5.4.	Track Editor and OSC Communication	131
9.5.5.	Hardware Setup	132
9.6.	Study Design	133
9.6.1.	Goals and Hypotheses	134
9.7.	Results	135
9.7.1.	Questionnaires	135
9.7.2.	Measured Data	137
9.8.	Discussion and Conclusions	140

10. The Slowification System	143
10.1. Introduction	143
10.2. Related Work	144
10.2.1. Spatial Panning to Guide Users	144
10.2.2. Dynamic Speed Assistance Systems	145
10.3. Interaction Design	145
10.3.1. Concept	146
10.3.2. Implementation	147
10.4. Study	148
10.4.1. A Virtual Reality Car Simulator	149
10.4.2. Study Design	149
10.4.3. Goals and Hypotheses	151
10.5. Results	152
10.5.1. Measured Data	152
10.5.2. Questionnaires	153
10.6. Discussion and Conclusions	154
III. Discussion and Conclusions	157
11.1. Complexity and Variability of the Context	158
11.2. Unobtrusiveness and Acceptance	159
11.3. The Use of Sound in Ambient Displays	160
11.4. Feedback, Persuasion, and Affectiveness	161
11.5. Ambient Information Systems Supporting Environmentally-Conscious Behavior	162
11.6. Outlook	163
A. Questionnaires, Images, and Sound Examples	164
Bibliography	165

1.

Introduction

When we take a step back and observe how we interact with our environment in everyday life, it soon becomes clear that we draw on a wealth of information without consciously focusing on it: Most of the time, we can tell how the weather is outside even without looking out of the window. In the kitchen, we can observe the water being heated for the pasta and very quickly notice when it is boiling, while simultaneously doing a multitude of other things. When on the bike, you can “feel” when a car is approaching, without actively listening to the traffic noise or looking over your shoulder. And you do not have to actively watch the door to notice when someone is coming into your room. Apparently, humans are quite good at being aware of things at the periphery of their attention – even while being engaged with some “primary” activity, which occupies most of their mental resources.

1.1. Interacting with Technology

Considering the modern use of technological devices, on the other hand, we can see that the vast majority of interaction is accomplished through focused attention: Checking your emails usually involves opening a dedicated application on your computer or smartphone and consciously scrolling through the inbox. In order to keep track of friends’ and family’s activities, most people rely on their favorite social network, which provides them with a constant stream of news that keeps them actively engaged – quite often more than they would prefer. And we all have seen people being so absorbed by what is happening on their smartphone’s screen that for the rest of the world they could as well be both blind and deaf.

In recent years, computers have not only become cheaper, but also much smaller in size, which has led to a growing number and pervasiveness of personal devices like smartphones, laptops and tablet computers, but also to computers being integrated



Figure 1.1.: “Life is what happens while you’re looking at your smartphone”: Traditional user interfaces can be both a chokepoint for and sometimes also a distraction from other (maybe more important) sources of information. (Image from: NPSAPPS, 2014)

into everyday objects and our environment, for example in the emerging field of smart homes. This development enables us to access and interact with digital information almost everywhere, anytime. In addition to that, the amount and availability of information is steadily increasing and, via the internet, we cannot only access general news and knowledge, but also personal data as well as communicate and exchange information with virtually the whole world. Furthermore, a growing amount of sensors integrated in smart watches, fitness trackers or energy-measuring smart meters provide us with a constant stream of information about ourselves and our environment.

While these developments can generally be seen as quite positive since they create new opportunities and extend the possibilities of supporting humans in their daily life, they also pose a challenge to interface design and to human-computer interaction (HCI) as a whole and can make traditional user interfaces both a chokepoint for and sometimes also a distraction from other (maybe more important) sources of information.

1.2. Calm Technology

The pervasiveness of computers was already foreseen more than 25 years ago by Weiser (1991), who predicted that every room would be equipped with a large number of interconnected devices of various sizes, which are aware of their location as well as the people using them. He also recognized that “even the most powerful notebook

computer, with access to a worldwide information network, still focuses attention on a single box” and envisioned “machines that fit the human environment instead of forcing humans to enter theirs”.

A few years later, Weiser and Brown (1996) concretized this vision by developing their concept of *calm technology*: In order to make the interaction with computers more natural and not to overburden the user with information, it should take place more in the user’s periphery rather than at the center of attention and easily move back and forth between the two – just as it is possible with most non-technological everyday objects.

1.3. Ambient Information Systems

Weiser’s concept of ubiquitous computing and calm technology has inspired the design and development of a wide range of innovative ambient (or peripheral) displays, which aim to present information in an unobtrusive way, “without distracting or burdening the user” (Mankoff et al., 2003). Systematic evaluation of such displays, however, is often quite expensive and for most of the published designs completely omitted, which makes an objective assessment of the efficacy of those displays rather difficult. This also means that there is only little general knowledge about which things “work” with ambient displays (and which do not) and what the instances are, where they actually perform better than traditional means of user interface design.

SUPPORTING ENVIRONMENTALLY-CONSCIOUS BEHAVIOR

As we will discuss in more detail in Chapter 3, ambient displays have already been used in several areas of application, e.g. at the workplace to inform users about communication events or upcoming appointments or to enable workers to keep aware of the availability and activities of their colleagues. Or, in a more personal setting, to provide an awareness of your significant other in a long-distance relationship or of an elderly person that might have to be taken care of.

Another use case that has been identified as particularly suitable for ambient displays is to support humans to act in a more environmentally-conscious way (Kappel, 2009). Given the fact that this topic has not yet been researched very extensively, we decided to concentrate our efforts on this particular area of application (cp. RQ1, Chapter 6). This decision is furthermore supported by the following considerations:

- In order to reduce energy consumption and CO₂ emissions, most efforts currently concentrate on the use of renewable energy sources and on efficiency increases of existing systems. The role of humans as energy-consumers, however, is often neglected, even though existing research hints at the large impact of our behavior in terms of energy consumption (e.g. Maréchal, 2009; Jackson, 2005).

- An important aspect of supporting environmentally-conscious behavior is to help users to obtain a more comprehensive and objective perspective on existing energy-consuming practices. In most situations, however, the user is already preoccupied with a certain primary activity (e.g. working on a computer or driving a car), which prevents him or her from actively focusing on additional information on energy consumption. And although this information can, in this context, be certainly assessed as important, it is obviously not critical for accomplishing this primary task and the interface should avoid a distraction of the user, making this a highly relevant use case for ambient displays.

Supporting environmentally-conscious behavior with technological artifacts can be seen as part of the emerging research field of sustainable HCI, which we will give a short overview of in Chapter 4.

THE USE OF SOUND IN AMBIENT INTERFACES

As can be seen from the introductory examples, *auditory* perception is of particular importance for our natural ability of being aware of our environment: Usually, we *hear* someone entering our room. We become aware of the boiling water through its sound. And by passively listening to the street's ambient noise, we can easily make out an approaching car as the soundscape changes. Moreover, preliminary research in the area of cognitive psychology hints at the auditory modality being highly suitable for use in an ambient display and for making a peripheral perception of information possible (cp. Chapter 2).

At the same time, the use of sound for the design of ambient displays has been largely neglected so far, and the majority of existing ambient interface designs focus on subtle *visual* cues being placed in the environment of the user to present information (cp. Chapter 3), which we think can only address a limited part of our ability to keep aware of our surroundings. Consequently, we specifically explored and researched the use of auditory cues in ambient displays to inform the user, based on the assumption that doing so can in many cases be superior to employing visual means to convey information (cp. RQ2, Chapter 6). Methodically, we could draw on a large body of research on sonification, i.e. the auditory representation of information using non-verbal sound, and auditory display design, which we will briefly review in Section 5.1. Furthermore, we have developed the concept of *blended sonification*, which we see as particularly suitable for use in an ambient display (Section 5.2).

INFLUENCE OF THE CONTEXT OF USE

Another aspect of ambient displays that has been largely neglected in preliminary research is the *context of use*, i.e. the users' overall perceptual situation, and if there is a clearly defined primary task or a variety of activities that users can be expected to be engaged with. Findings in the area of attention theory and our analysis of existing

systems suggest that the use context might be of particular importance for this type of interfaces, as it can influence the effectiveness of the peripheral information conveyance. Nevertheless, it has been completely disregarded in the design process and the evaluation of the majority of the ambient displays developed so far.

With our prototype systems and in the corresponding studies (as discussed in Part II), we consequently explored several settings with a range of different use contexts in order to evaluate the impact the context of use can have on an ambient display's effectiveness to convey information (RQ3).

1.4. Research Questions

In addition to the already described research goals, we were also interested in the *unobtrusiveness* of ambient displays, as this is arguably one of the defining features of such an interface. Furthermore, in the context of supporting a change in behavior, the role of *feedback* has been a focus of research, as well as the potential *persuasiveness* of ambient displays, especially when making use of affective cues.

All in all, the following research questions have guided the development and evaluation of the prototype systems, as we will describe in more detail in Chapter 6:

- RQ1: *How can ambient information systems support environmentally-conscious behavior? What are use cases that “work” for using such an interface? What are not?*
- RQ2: *Is using sound in a peripheral display a feasible alternative to (purely) visual ambient interfaces?*
- RQ3: *How does the variability of an ambient display's context of use affect the way it is perceived?*
- RQ4: *What are effective ways to make ambient interfaces less obtrusive and distracting?*
- RQ5: *Are more unobtrusive ambient displays better accepted by users?*
- RQ6: *Is providing feedback on consumption practices enough to induce a corresponding change in behavior?*
- RQ7: *Is it viable to use ambient displays as persuasive technology?*
- RQ8: *Is the use of affective cues beneficial for ambient displays aiming to persuade users towards a change in behavior?*

Part I.

Methods & Research Background

2.

Attention Theory

While most of the existing research on ambient displays is done in a more practical and exploratory fashion, we think that in order to understand how, why, or under which circumstances ambient displays work, it is important to have a good understanding of human information processing and especially the human attention abilities. Therefore in this chapter, we will provide an overview of critical research in cognitive psychology and try to give a working definition of the “periphery of attention”, which, in Weiser’s concept of calm technology, is a frequently used but only vaguely defined term.

2.1. Introduction

For most people, attention is something that they notice when it is missing: Almost all the time, we are bombarded by a multitude of different sensory information – and if we are unable to focus on the input that is, in this situation, relevant to us, it is almost impossible to process and, in consequence, to take action based on the overwhelming amount of information. In this way, attention can be understood as the opposite of distraction: It allows us to direct our limited mental resources towards the information that is, at any given moment, most important (or most salient) to us. In 1890, William James characterized attention quite fittingly as the “withdrawal from some things in order to deal effectively with others” (James, 1890, p. 404).

While for the understanding of ambient displays, our ability to attend to *external* sources of information is certainly the most important one, attention can also be devoted to performing a specific action or to internal thought processes. Furthermore, we can distinguish between several *functions of attention* (Sternberg, 2009, p.141):

Selective Attention: This is probably what is meant in most cases when the term attention is used in ordinary language: When we are reading a book and paying attention, we have made the choice to selectively focus on the visual input, the

words and their meaning, and to ignore such stimuli as, for example, the sound and the movement of children running around in the room or the thought, which groceries to buy for tomorrow's dinner.

Vigilance and Signal Detection: In other situations, we do not have something specific to focus on, but instead vigilantly direct our attention outwards so that we can take speedy action when we detect a specific signal. For example, walking through a dark alley in a neighborhood where there have been repeated muggings can make us watch out for suspicious movements or sounds.

Search: While vigilance can be seen as merely waiting for a signal stimulus, searching involves actively looking for something that we are not sure where to find, like for example searching for a particular term on this very page.

Divided Attention: In everyday life, there are many occasions, where we are engaged with more than one task at a time, as opposed to selective attention situations, and manage to divide and shift our attention between different sources of information. For example, when we are on the phone, we are in most cases able to simultaneously perform basic household tasks, like washing the dishes.

In our daily life, outside of specific work contexts, we only occasionally use our attention for searching and rather seldom engage in tasks that require vigilance. Since calm technology is grounded in everyday life contexts, the most relevant functions for our understanding of ambient displays are therefore selective and divided attention, which we will discuss in more detail.

2.2. Selective Attention

Probably the most well-known real life example to illustrate selective attention is the so-called "cocktail party effect" (Cherry, 1953): It describes the phenomenon that even in a crowded room full of people talking at the same time, we are able to distinguish and focus on a single conversation while filtering out others. It is also the basis for a type of experiments trying to replicate this setting under controlled laboratory conditions, called *shadowing*: A subject is instructed to listen to two different auditory messages at once, but to follow (or *shadow*) only one of them, i.e. the subject should try to repeat it as soon as possible after it is heard. The other message should simply be ignored and is commonly referred to as the rejected message.

With this setup, the experiment allows to observe how well the subjects are able to attend to the shadowed message. Even more interesting, however, is the question, how much of the *rejected* message is being processed while the subject is tracking the attended one: Cherry (1953) not only observed that the subjects found it remarkably easy to separate the two auditory channels and to focus on only one of them, but he also found out that when asked about the contents of the rejected message, they could say almost nothing about it and did not remember any words mentioned in it.




Figure 2.1.: Simplified illustration of an early selection model: Assuming a bottleneck in the perceptual process itself, Broadbent (1958) proposed a selective filter mechanism that only allows one stimulus to be processed at a time. Adapted from (Sternberg, 2009, p.154).

2.2.1. Early Selection

Based on this last observation, Broadbent (1958) suggested that at least part of our perceptual system is only capable of handling one stimulus at a time (Pashler, 1998, p.14). The central element of his proposed model of attention is therefore a selective filter, which only allows one channel of sensory information to pass and rejects all others, based on a preliminary analysis of simple physical attributes of the stimuli, e.g. pitch or volume (cp. Figure 2.1). More advanced information requiring higher perceptual processes is then only determined for the selected sensory stream, i.e. the selection takes place rather early in the process.

Broadbent's theory is supported by Cherry's findings that subjects were able to detect changes in the rejected message, if very basic features were changed, for example when the message was changed to a tone or the voice changed from a male to a female speaker (Sternberg, 2009, p.152). Semantic changes of the unattended message, however, were not detected, and the subjects even failed to notice when the language was changed or the message was played backwards.

2.2.2. Late Selection

Soon after Broadbent had proposed his model of attention, there was growing evidence that it must in fact be wrong: In experiments quite similar to the ones Cherry had conducted, Moray (1959) showed that even when subjects ignored most high-level (e.g. semantic) content of the unattended message, they still frequently noticed the occurrence of their names in it (Sternberg, 2009, p.155). Similarly, when presented with the request "you may now stop" in the unattended message, a significant number of subjects actually did so (Moray, 1959). These findings revealed that the meaning of at least some words must be processed *before* reaching the filter and were the basis for an alternative theory, first proposed by Deutsch and Deutsch (1963), which suggests

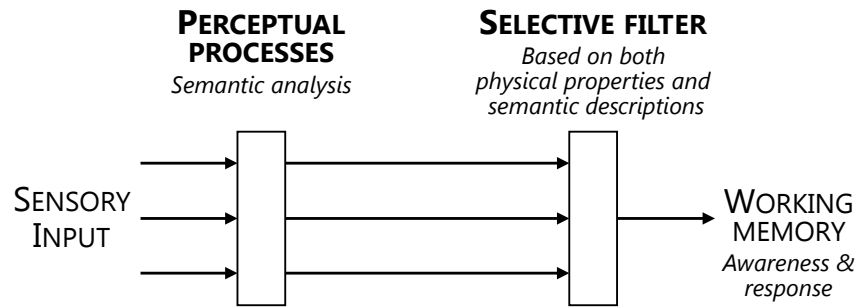


Figure 2.2.: Simplified illustration of a late selection model: Based on experimental findings, which showed that at least some of the unattended stimuli must be analyzed by higher level perceptual processes, Deutsch and Deutsch (1963) proposed that a selection of sensory input takes place *after* them. Adapted from (Sternberg, 2009, p.156).

that “recognition of familiar objects proceeds unselectively and without any capacity limitations” (Pashler, 1998, p.17). This effectively means that a significant part of the perceptual process is actually *preceding* the selective filter (cp. Figure 2.2). According to their theory, this part of the perceptual process happens involuntarily and below our level of awareness.

2.2.3. Attenuation Theory

Although both early and late selection theories have been highly influential and very well demonstrate a major point of controversy in attention research, we can see that both make rather extreme claims: Early selection theories, on the one hand, cannot explain why some words are recognized when present in the unattended message. A full processing of all sensory input, as it has been proposed in late selection theories, on the other hand, would be quite resource demanding and is considered rather unlikely as well, since most information does not even seem to be used. Therefore, most researchers agree that some form of compromise theory is necessary. One widely acknowledged theory has been proposed by Treisman (1964), who suggested that instead of blocking all sensory information other than the target one, the attention system’s filtering mechanism merely *attenuates* them and thus a weakened version of the stimuli is passed on to the next stage of perceptual processing, where a threshold determines if it should be identified (cp. Figure 2.3).

2.2.4. Context and Priming

Given this model, the central question is which factors determine the attenuation of unattended stimuli and their respective thresholds: Treisman had found out that subjects not only frequently recognized their name in the unattended message, but also

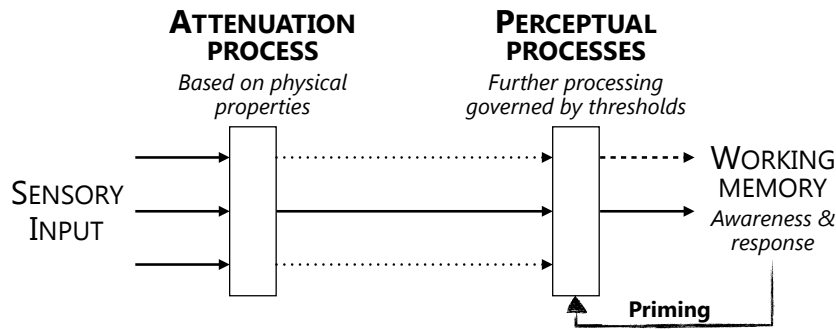


Figure 2.3.: Simplified illustration of Treisman's attenuation model: An attenuated sensory input can still be recognized and processed in the subsequent perceptual stage, depending on the threshold for this particular stimulus. The thresholds depend on the respective situational context and, through priming, can be modified by the perception of other stimuli.

words that were closely related to the information presented in the attended one (Treisman, 1960). Given these findings, she concluded that context plays a key role in determining the recognition threshold for incoming information and proposed that it does so by means of *priming*: Based on the (perceptual or semantic) information from the attended channel, related words or concepts are primed, i.e. the recognition threshold for them would be lowered. For example, when hearing the word “spaghetti”, semantically similar words such as “noodles” or “pasta”, or related words such as “food” or “eating” may be primed (Meyer & Schvaneveldt, 1971) and thus the likelihood for recognizing these words in the unattended message is increased.

Moving away from the experimental setting of shadowing, the priming stimulus does not necessarily have to be a word, but could also be an internal thought, performing an action, or a specific scent (Holland, Hendriks, & Aarts, 2005): Perceiving the smell of freshly baked cookies, for example, can lower the threshold for recognizing the corresponding word. Or conversely, hearing the word “cookies” can make us more receptive for the respective smell.

2.2.5. Selective Attention and Calm Technology

Interpreting selective attention situations with regard to ambient displays and calm technology, we would argue that the attended sensory information lies at the center of attention, while the periphery consists of all other, unattended stimuli. Following Treisman's model of selective attention we can say that not only basic, physical features of those unattended stimuli are analyzed, but to some degree higher-level information is extracted and processed as well. However, we must assume that this is done solely for the purpose of selecting the right input and that most unattended sensory information is ultimately discarded. Importantly, the context in which the perception takes place

as well as the type and content of the attended stimuli affects how much and which information from the periphery of attention is perceived.

2.3. Divided Attention

Different from research on selective attention, which is mostly concerned with the perception of sensory stimuli only, theories on divided attention take a broader approach and try to explain how we can simultaneously perform multiple attentional *tasks*. As we can see when considering a fairly complex task such as playing the piano, those can involve not only perceptions, but also physical actions or thought processes and may include multiple modalities.

2.3.1. Attentional-Resource Theories

Models explaining divided attention situations have moved away from the notion of a blocking (or attenuating) filter mechanism, but instead see attention as the allocation of a limited amount of mental resources, which can be assigned to perceptual or other cognitive processes, depending on what the respective task requires (Kahneman, 1973). Such a model, however, can obviously also be applied for selective attention situations and should be seen as complementing, as opposed to competing with, selective filter theories (Sternberg, 2009, p.158).

Considering the specific aspects that influence the allocation of resources, experimental findings suggest that the amount of attentional resources available to a specific person depends both on personal traits of the respective person and the overall arousal, e.g. if someone is tired or excited (Eysenck & Calvo, 1992). Another important aspect is, of course, the interest one has in a specific task or input compared to the interest in other (distracting) stimuli.

2.3.2. Practice and Automaticity

Probably the most important factor for the amount of resources needed for the execution of a task is the nature of the task itself, i.e. a complex or difficult activity will need more resources than a simple or easy one. However, the *subjective* difficulty is also highly dependent on how much practice a certain task has received: When learning how to drive, most people are overwhelmed by the combination of watching the road, steering the car, and changing gears. Then, after lots of practice, driving becomes rather effortless – and we manage to do so even while talking to a passenger or listening to the news.

CONTROLLED AND AUTOMATIC PROCESSING

Based on such observations, Schneider and Shiffrin (1977) have suggested the distinction between *controlled* processes, which require conscious control, and *automatic* ones, which seldom enter conscious awareness and are under the control of stimulation, i.e. they are executed based on rather static input-response mechanisms (Anderson, 1983). Depending on the task, those automaticities can be learned and thus a controlled process can become an automatic one. Although by now it has become clear that one cannot make such a clear differentiation between the two types and a continuum seems to be more likely (Styles, 1997, p.169), we still think that it is helpful to make this conceptual distinction. While most researchers agree that there is always a certain degree of interference for multitasking situations, we can observe that automatic processes generally require less conscious effort and can be combined with other tasks at only little cost, as opposed to controlled tasks, which can usually be executed only one at a time.

2.3.3. Crossmodal Attention

Although Kahneman's resource theory has given an influential perspective on attentional processes, some researchers have criticized it as too simplistic: People are much better able to divide their attention among tasks of different modalities, as opposed to tasks that require the parallel processing of only one modality (Sternberg, 2009, p.158). For example, when reading a book, it is much easier to observe your children with an acoustic baby monitor than to do the same with a video feed of their bedroom. Since this phenomenon cannot be explained by conventional resource theory, Navon and Gopher (1979) proposed a modified model, where each modality and the different perceptual processes have their own pool of attentional resources. Therefore in this model, allocation of resources for perceiving visual input does not interfere with those for auditory perception. Furthermore, it also explains why listening to the news while reading a book is also quite difficult: Despite employing different modalities, both tasks require the same resources for verbal processing.

2.4. Insights Gained for Ambient Displays

After this necessarily quite short and focused review of attention theory, we can now have a look at what this means for our understanding of ambient displays and calm technology. If not otherwise noted, we analyze selective and divided attention situations based on attenuation (Section 2.2.3) and multiple resource theory (Section 2.3.3), respectively. Weiser and Brown (1996) have suggested that a becalming way of interaction can be achieved by engaging not only the center, but also the periphery of attention. They subsequently describe the periphery as "what we are attuned to without attending to explicitly".

2.4.1. The Periphery of Attention

In cognitive psychology, the term *periphery* is commonly used to refer to the part of the vision that occurs outside of the center of gaze, i.e. everything that is not perceived by foveal vision. Although many ambient displays actually employ visual cues to convey information and therefore rely on our ability to see in the periphery, Weiser and Brown (1996) obviously use the term in a broader context, independent of any specific modality. Based on our review of fundamental research on attention, we argue that “the periphery” can actually mean two things:

- In a selective attention situation, the sensory input we are focusing on very clearly lies at the center of attention. The periphery would then consist of all other, unattended stimuli, or, put differently, the attenuated channels of perception when the focus of attention lies elsewhere. There is, however, no equivalent for what Weiser and Brown call “being attuned to”: While unattended information may very well enter our awareness, this is seen to be beyond our conscious control.
- For divided attention situations, there are several activities that receive various amounts of attentional resources. Based on the definition of peripheral vision, we would argue that the activity receiving *most* resources lies at the center of attention, while all other attended activities represent the periphery.

In her work on peripheral interaction, Bakker suggested that the periphery “consists of all potential activities that are not in the center, regardless of the number of resources being allocated to them” (Bakker, 2013, p.27). While we see this as a valid compromise definition, combining these two views, we also think that it is important to keep in mind the distinction between partially attended and unattended activities.

2.4.2. Amount and Type of Conveyed Information

Ambient displays can be seen to provide (non-critical) information for a user, who is simultaneously performing a higher-prioritized “primary” task. A key issue is therefore the question, how much and which type of information can be conveyed to and thus easily perceived by this user. In order to answer this question in the context of the discussed findings from cognitive psychology, we must again differentiate between selective and divided attention situations: In the latter, the user consciously tries to perceive information from the ambient display while performing the primary task. In this case, attention theory suggests that, if the processing of the provided information does not conflict with the primary task, the user is able to consciously perceive all information. While there is no final verdict on how exactly incoming stimuli are processed if their perception *do* conflict with a primary task, there seems to be a broad consent about the existence of capacity limits in the perceptual process, which can lead to only limited processing of stimuli as well as a certain interference between them.

In a selective attention situation, however, where a user is highly focused on the primary task, we must assume that most information perceived from the periphery never even enters conscious processing and, without automaticities to process the information on a subconscious level, is effectively discarded. The only possibility for an ambient display used in such a situation to be (consciously) perceived then is to become salient to the user, i.e. to function as an alarm. While this could indeed be used in specific circumstances, such a disruptive behavior should obviously not be the default mode of an ambient display.

2.4.3. Context and Modality

As we have seen, the context in which the ambient display has to function, and especially the primary activity are of special importance for how well and in which way an ambient display can and should function: If the primary task is an automatic one (cp. Section 2.3.2), the ambient display can be expected to receive a considerable amount of (attentional) resources. On the other hand, if the user is preoccupied with a highly demanding task, the display will quite likely be ignored most of the time.

Special consideration should be put into the choice of modality, as this can lead to particularly restrictive low-level conflicts. Since many tasks we do in our daily life are of visual nature, we assume that auditory ambient displays might lead to fewer conflicts than visual ones.

2.4.4. Automaticities and Learning

In Section 2.3.2, we have learned that under certain circumstances, humans are able to perform tasks without requiring a significant amount of attentional resources. Assuming that a user is performing a primary task of such nature would allow an ambient display to provide an enhanced amount of information without being perceived as overburdening. Conversely, if users were able to process the presented information in a similar, automatic way, this would allow them to simultaneously perform a rather demanding primary task.

The general prerequisite for such automaticities is assumed to be a rather static input-action scheme, i.e. there should be a rather easy way to encode how perceived stimuli translate into a specific action. Another, more involved perspective on this topic has been given by Neumann (1984, p.282): He proposed that a process is automatic if “its parameters are specified by a skill in conjunction with input information”. If this is not possible, i.e. if parameters are left unspecified, either due to lack of skill or due to unexpected input information, “one or several attentional mechanisms for parameter specification must come into play [...] and give rise to conscious awareness” (Neumann, 1984, p.282).

3.

Ambient Information Systems

3.1. Introduction

While cognitive psychology has given us a rather theoretical view on attention and its allocation, most work from the area of human-computer interaction (HCI) has focused on a more practical and exploratory way of researching Weiser's vision of calm technology. As a result, we can draw on a diverse range of prototype systems that have been designed to convey information in an encalming way. In order to give a structured overview of this area of research and to identify both major themes and essential design dimensions, we have reviewed six overview papers as well as 45 research prototypes and consumer-grade systems that can be classified as calm technology or even cite Weiser's original paper as motivation.

As diverse as the systems themselves are the names used to describe them: Among others, we have found *ambient displays* (Consolvo, Roessler, & Shelton, 2004), *notification systems* (McCrickard, Chewar, Somervell, & Ndiwalana, 2003), *ambient media* (Ishii et al., 1998), *peripheral displays* (MacIntyre et al., 2001), *informative art* (Redström, Skog, & Hallnäs, 2000), *ambient fixtures* (Dahley, Wisneski, & Ishii, 1998), and *information awareness interfaces* (Cadiz, Venolia, Jancke, & Gupta, 2002).

While those terms can be seen as highlighting specific important aspects of calm technology, it is difficult to make out clear differences between two groups of systems based on their names, as these seem to be chosen rather arbitrarily in most cases. Nevertheless, in our analysis we will also address emerging tendencies of differing characteristics that can be found for distinctly named systems as well as specific qualities that can be derived from the terms themselves. Considering the number of occurrences, the terms *ambient display* and *peripheral display* are most frequently used to describe the systems we are interested in and can also be considered as somewhat "established" in the research community. Pousman and Stasko (2006) have furthermore proposed to

use *ambient information system* as an umbrella term, which has also been adopted by Tomitsch, Kappel, Lehner, and Grechenig (2007). In order to account for the wide range of different types of systems and for reasons of continuity, we will likewise use this term to refer to *all* display technology that can be attributed to calm computing. Different to Pousman and Stasko (2006), however, who claim that ambient displays “have pointed aesthetic goals and present a very small number of information elements”, we see no fundamental difference in the use of *peripheral display* and *ambient display* in existing scientific work and think that these terms can and should be used synonymously.

3.2. Common Themes

In our analysis of existing ambient information systems, we first have focused on finding clusters of reoccurring themes that can be found in the goals and problem formulations of the respective systems and in how the work is motivated. On a very coarse level, we can distinguish between research projects that focus more on advancing display technology (which can then be the basis for designing ambient information systems) and those that aim to support people in a given situation by providing awareness for specific information.

Although most projects presented in this section also provide exemplary use-cases or even very specific descriptions of applications, they all try to advance display technology on a very basic level.

EXTENDING THE NOTION OF A PIXEL

One example is the *Hello.Wall*, which is a 1.8-meter-wide by 2-meter-high artifact with an array of 124 light-emitting cells (Prante et al., 2003). The authors envisioned it as part of a communicative area in a building, which can provide awareness to the members of an organization based on visual codes (Figure 3.1a). Since the perception of information is based on an understanding of the underlying codes, the display could communicate not only public, organization-oriented news, but also personal information, which only the recipient can understand. For visitors, the *Hello.Wall* might be considered simply as a decorative element. Finally, the presented information is conceived to be dependent on the distance between a user and the display: While public information would be displayed at all times, a personal notification would only be shown if the corresponding user is in the vicinity of the artifact.

3.2.1. Advancing Display Technology

In a different context, Gaver et al. (2006) have developed a tablecloth with a grid of hexagon-shaped patterns, which are printed on electroluminescent material and thus can be individually highlighted. With the help of additionally integrated weight sensors,




Figure 3.1.: **(a)** The *Hello.Wall* artifact provides information through dynamically changing light patterns (Streitz et al., 2005). **(b)** Design visualization of the *History Tablecloth*, whose patterns light up depending on the objects placed on it (Gaver et al., 2006).

the authors used the tablecloth to display the flow of objects in the home, i.e. when an object is placed on the table, a halo shows up and grows over a longer period of time and only slowly dims after it has been removed (Figure 3.1b).

Although seemingly quite different, these two designs share the idea of combining multiple objects that, taken together, exhibit pixel-like qualities and thus can be seen as a low-resolution screen. In consequence, the possible amount of conveyed information is already reduced at the display-level, which also makes it possible to move information off the traditional screens and better integrate it into our environment. Other projects have used bubbles of air (Heiner, Hudson, & Tanaka, 1999), table tennis balls (den Breejen & Deenstra, 2005), or the windows of an office building (Chaos Computer Club, 2001) as a collection of pixels.

WEARABLE (AND POCKETABLE) DISPLAYS

Other projects have tried to bring display technology *closer* to the user: Costanza, Inverso, Pavlov, Allen, and Maes (2006) have embedded two arrays of small LEDs at either side of an ordinary pair of glasses in order to deliver visual cues at the edge of the wearer's field of view, while avoiding to block peripheral vision (cp. Figure 3.2a). Due to the low fidelity of those cues, the authors proposed to use them primarily for notification purposes, e.g. to indicate incoming email. A few years later, Google developed a technologically greatly improved and advanced version of this idea, called *Google Glass* (Google, 2013a). The device employs a near-eye display to deliver information at the peripheral field of view (Figure 3.2b). Furthermore, due to the high resolution (640×360 pixels) of this display, it is also possible for the user to focus on the presented image in order to get more detailed information.




Figure 3.2.: Two peripheral head-mounted displays. **(a)** The *eye-q* eyeglass peripheral display notifies its wearer through LEDs embedded at each side of the glasses (Costanza, Inverso, Pavlov, Allen, & Maes, 2006). **(b)** *Google Glass* employs a near-eye display to convey information to the user (Google, 2013a). (Image from “*Google Glass*”, 2014)

While generally not considered as ambient information systems, both smart watches and the by now ubiquitous smart phones have certainly also brought display technology closer to the user and we would argue that they at least have the *potential* to be used as peripheral displays. In fact, the android mobile operating system has an “Ambient Display” feature, which provides the user with a glanceable overview of pending notifications.

AESTHETICALLY PLEASING AND ARTISTIC DISPLAYS

Quite a few projects emphasize the aesthetic value of a display or even see the primary goal of their work in creating an aesthetically pleasing object, which is *also* able to convey information (Fogarty, Forlizzi, & Hudson, 2001). Consequently, a term often used for this type of ambient information system is *informative art*. Redström et al. (2000), for example, have created a data visualization that is inspired by the artist Piet Mondrian (Figure 3.3a). Due to the abstract nature of the artwork and its clear geometric forms, it still resembles a traditional information visualization and one can easily imagine how the shape, position, and color of the individual fields could be changed to reflect a specific data source. Exemplary applications include the visualization of email traffic, the local weather forecast, and bus departure times (Skog, Ljungblad, & Holmquist, 2003).

Similarly, Shen, Eades, Hong, and Moere (2007) have used artworks inspired by Hans Heysen, an Australian watercolor painter of the early 20th century, to convey information by modifying specific features of the image. In the composition depicted in Figure 3.3b, for example, there are several statistics about a website being displayed: (1) The amount of fog in the mountains correlates with the number of people visiting the website, with heavier fog indicating fewer visitors. (2) The number of the trees on

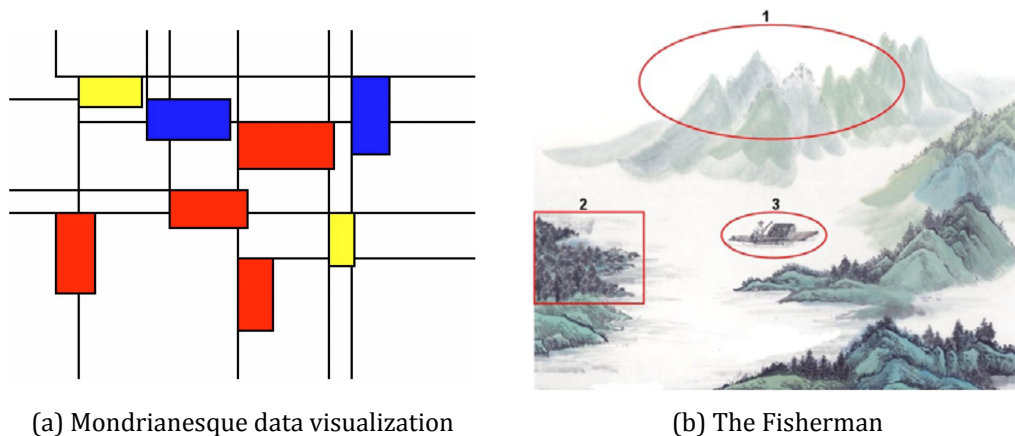


Figure 3.3.: Two pieces of *informative art*. **(a)** A Mondrian-inspired artwork is used as basis for an ambient display (Redström, Skog, & Hallnäs, 2000). **(b)** Several features of a traditional painting are modified to reflect changes in data (Shen, Eades, Hong, & Moere, 2007).

the left represents the number of pages viewed by each visitor. And (3) the vertical position of the boat shows the bandwidth used by the server.

Further displays that can be considered as informative art include the *Kandinsky* system, which aims to translate textual input into an “aesthetic information collage” by first finding representative images and computationally combining those while maintaining certain aesthetic properties (Fogarty et al., 2001), the *BlueGoo* display, which uses a similar approach to present news from an RSS feed in form of an animated photographic collage (Plaue & Stasko, 2007), and the *InfoCanvas* system, which focuses on providing a flexible, personalized way of representing data in a visual “scene”, consisting of several user-defined elements (Stasko, Miller, Pousman, Plaue, & Ullah, 2004).

What most of these systems have in common is that due to the strong focus on aesthetic value, the quality of information conveyance is considered only a secondary goal. Furthermore, similar to the Hello.Wall, the user has to know and learn how the information is represented and thus the initial accessibility of these systems is usually rather low. In consequence, it might be quite difficult to introduce them in a public setting, which has also been hinted at in a short field test of the visualization of bus departure times (Skog et al., 2003). However, when used as a personal display used over a longer period of time, this did not seem to pose a problem (Stasko, McColgin, Miller, Plaue, & Pousman, 2005).

INTEGRATION INTO URBAN AND ARCHITECTURAL SPACES

Finally, there is a cluster of work that aims at integrating display technology directly into our buildings: In the *AmbientROOM* project, several installations have been used




Figure 3.4.: Two examples of displays being integrated into an architectural context. **(a)** A window that is overlaid with a graphical display indicates bad weather in the next few hours (Rodenstein, 1999). **(b)** People in the street can send text messages in order to have them displayed in two speech bubbles projected onto a building (Notzold, 2006).

to augment a room, e.g. an “active wallpaper” consisting of several light patches projected on an inner wall, a “water lamp” that projects ripples produced by a solenoid being tapped into the water onto the ceiling of the room, and, most notably, a natural soundscape, which can be modulated based on input data (Ishii et al., 1998).

In a different project, Rodenstein (1999) tried to integrate a display directly into a room’s window. As a result, the outside view, which can be seen as already providing ambient information, was augmented with overlaying visualizations, e.g. a short term weather forecast (Figure 3.4a).

In a non-research project, Notzold (2006) created a public installation, where speech bubbles were projected onto the facade of a building, next to windows. People in the street could send text messages to a given phone number, which would then appear in the bubbles. While primarily conceived to be a “stage for creating spontaneous dialog”, this example shows how urban spaces could be augmented by display technology and thus potentially serve as a public ambient display.

3.2.2. Providing Awareness

The second type of work on ambient information systems is more concerned with solutions for application-specific problems that involve providing awareness for certain information. Although the range of different types of projects is too wide to be covered in its entirety, we have identified several clusters of reoccurring themes.

AT THE WORKPLACE

A quite frequently stated goal is the improvement of the productivity and quality of work. In one of the earlier works, McCrickard (1999) argued that most people at their workplace want or need to be aware of certain types of information, but that it is too time-consuming for them to regularly check for changes. Based on this assumption, he developed *Irwin*: an “information resource watching” system consisting of a small set of tools for desktop computers, which monitors email accounts, web pages, and weather data. The collected information is then visualized on a small portion of the screen and the user can be alerted in case of updates and modifications.

In a similar project, initiated by Microsoft Research, Cadiz et al. (2002) further explored the idea of having a central place for notifications and information the user might want to keep being aware of. Described as a peripheral awareness system, *Sideshow* occupies a small space at the side of the screen and can be configured to display a wide range of pieces of information, such as upcoming appointments, traffic status, or the availability of colleagues. Later on, a similar kind of sidebar was integrated into Windows and by now can be found in every major operating system.

Following a slightly different approach, MacIntyre et al. (2001) proposed to augment an office environment with large projected interactive surfaces and thereby to extend the screenspace offered by conventional desktop computers. The authors envisioned an interactive “wall display” produced by several beamers, which can be used to keep track of background and past activities of the user’s current working contexts, i.e. the documents related to a certain project or task as well as the on-going interactions with other people concerning this project.

Awareness of Workplace Activities. While the previous examples are more concerned with keeping aware of virtual sources of information, i.e. data that can be acquired from the internet or directly derived from the activities carried out at one’s personal workstation, there are also attempts to provide awareness for the presence, availability, and activities of co-workers. The *Nimio* system, for example, consists of a series of differently shaped, toy-like artifacts with both integrated sensors and LEDs, designed to give a group of co-workers a sense of awareness of each other (Brewer, Williams, & Dourish, 2005). The basic idea is that the movement of a Nimio artifact or sound in the vicinity of one will lead to all other Nimios belonging to the same group to glow in a specific color. Furthermore, users can respond to such an activity by shaking one of their own Nimios, which can lead to more advanced light patterns (Figure 3.5a).

With a stronger focus on utility, the *Notification Collage* is essentially a virtual bulletin board designed to foster interaction and interpersonal awareness between colleagues (Greenberg & Rounding, 2001). Users are able to post a wide range of media elements such as sticky notes, a screenshot of their desktop, images and web pages

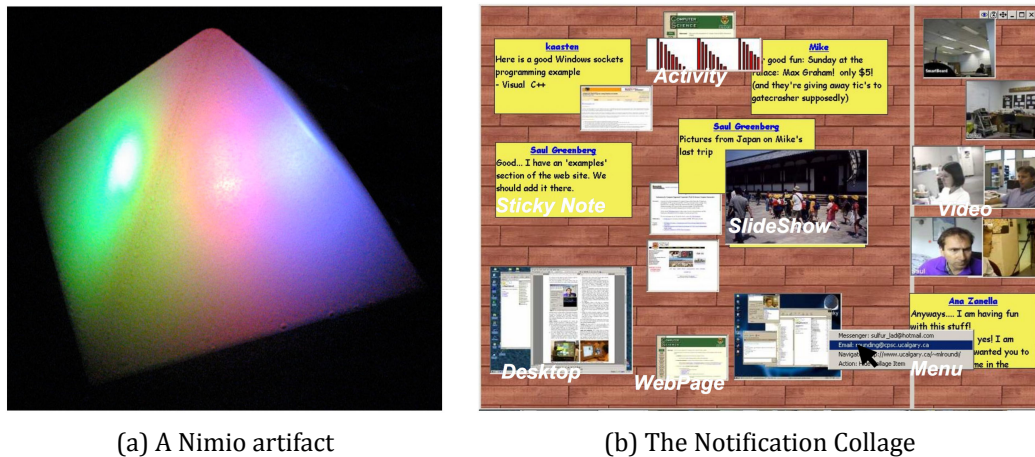


Figure 3.5.: Providing awareness of workplace activities. **(a)** A Nimio artifact reacts to action around other Nimios and can light up in several colors (Brewer, Williams, & Dourish, 2005). **(b)** The Notification collage is a bulletin board-like display, designed to keep users aware of their co-workers' activities (Greenberg & Rounding, 2001).

as well as a live video of themselves to signal their presence at the workplace (Figure 3.5b). The authors installed the Notification Collage both in a public setting and as a peripheral display for individual users, who could also move elements they felt important to the right side of the screen, where they are not replaced by subsequent posts by other users.

A completely different approach was taken in the *Audio Aura* project: In order to keep aware of the co-workers' activities, Mynatt, Back, Want, Baer, and Ellis (1998) proposed to use background auditory cues delivered by portable wireless headphones. For example, when a person wants to visit a colleague and only encounters an empty office, the *Audio Aura* system might give a qualitative auditory cue to indicate how long the respective person has been gone already. Furthermore, for geographically distributed teams, *Audio Aura* could offer a soundscape, described by the authors as a "group pulse", which reflects whether people are in the office that day or if a subset of the team is currently meeting face to face.

CO-PRESENCE SYSTEMS

A second cluster of work deals with providing awareness of another person that is located elsewhere. Contrary to the projects described in the previous section, this mainly involves personal relations, such as a family taking care for an elder or a couple in a long-distance relationship.

Care for the Elderly. Both the *Digital Family Portrait* (Mynatt, Rowan, Craighill, & Jacobs, 2001) and the *CareNet Display* (Consolvo et al., 2004) were designed to support

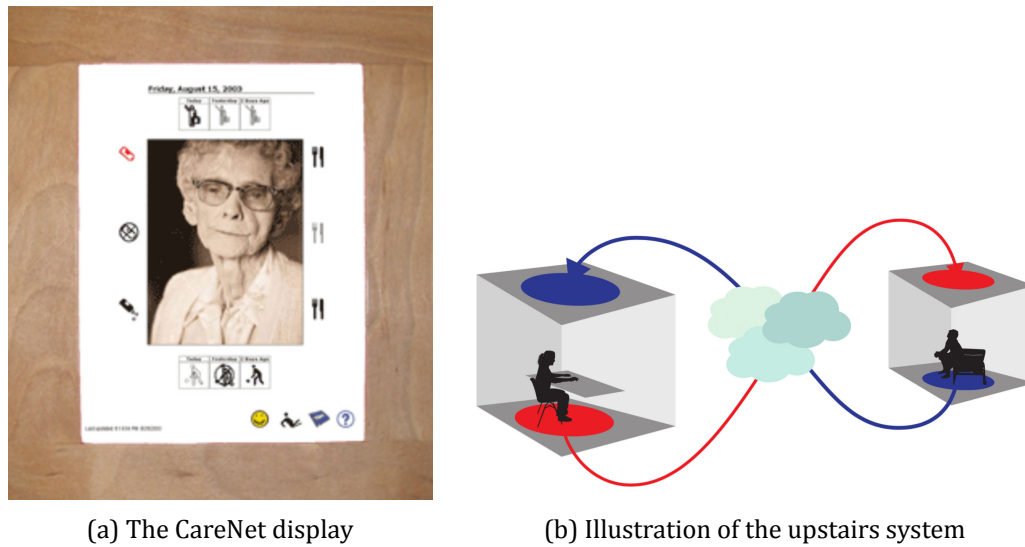


Figure 3.6.: Two systems designed to provide an awareness of another person. **(a)** The CareNet display shows information about elderly people that is important for the people taking care of them (Consolvo, Roessler, & Shelton, 2004). **(b)** The upstairs system acoustically simulates another person living one floor above (Tünnermann, Leichsenring, Bovermann, & Hermann, 2015).

the people who provide an elder with the care he or she needs to remain living at home. According to the authors, a major problem for these situations is the lack of awareness of the person being taken care of, which consequently led to the design of the respective displays. The information that was found to be vital for the caregivers includes if meals have been eaten or certain medications have been taken as well as more drastic events like the elderly person falling over. Furthermore, general activities and excursions outside the apartment were also deemed important. For both displays, the information is presented through iconic images, which are arranged around a photo of the respective person (Figure 3.6a). The CareNet display is furthermore integrated into a picture frame and can be seen as an augmentation of an artifact that could be found in a household anyway. The authors propose that the information could partly be collected by sensors installed in the elderly's home, in addition to the possibility of information being added by the caregivers or the elderly person him- or herself. The CareNet Display furthermore allows users to interactively obtain more detailed information such as a history over the past few days. Finally, for the Digital Family Portrait project, a similar display was proposed to be placed in the home of the elderly person to provide, for example, an awareness of the grandchildren's activities.

Long-Distance Relationships. Another situation where an awareness of another person is desirable emerges when two persons are in a long-distance relationship. Addressing this issue, the *LumiTouch* system consists, similar to the projects discussed

earlier, of a pair of interactive picture frames with both active and passive means of communication. It allows users to squeeze the frame, which is then displayed via color LEDs on the companion frame, depending on where, how hard, and how long the frame has been squeezed. Furthermore, the sole presence of a user, detected by a motion sensor, is also transmitted and displayed via a subtle glow of the other person's frame.

A quite different approach was taken by Tünnermann, Leichsenring, Bovermann, and Hermann (2015) in the *upstairs* project: Their work is based on the observation that being aware of another person in a shared space, e.g. when living together in the same apartment, relies heavily on auditory cues that occur when a person is interacting with the environment. The authors have developed a system that captures the sound created by footsteps and other interaction with the floor and plays them back in the remote location so that they appear to come from the ceiling, thereby simulating that the other person is living one floor above (Figure 3.6b).


3.3. Design Dimensions

While in the previous section, we have analyzed and structured selected existing systems in terms of differences of goals and problem formulation, i.e. giving an overview of *what* has been done in the field, we now want to focus on the *how* and have a look at the most important characteristics and design dimensions that differentiate the displays.


3.3.1. Event-like or Continuous Data

A first important distinction to be made is between systems that provide awareness of a continuous data stream and those that aim to notify the user about specific events. For example, while the upstairs system described above conveys a continuous audio stream, the eye-q eyeglass display is specifically designed for notifying a user of certain events (cp. Figure 3.2a). This kind of display can often also be found in work contexts, where incoming emails, tasks, or calendar events are of essential value. Besides systems like Sideshow (Section 3.2.2), this includes for example the *Scope* display, shown in Figure 3.7a: Described as a “glanceable notification summarizer”, it gives users an overview of pending events, with more urgent notifications being placed closer to the center of the display (Van Dantzich, Robbins, Horvitz, & Czerwinski, 2002).

Ambient information systems that are dealing specifically with event-like data are often referred to as *notification systems* (also cp. McCrickard et al., 2003). The main challenge for this type of display is to keep the user updated about (sometimes urgent) events while minimizing distraction. Note that most systems displaying continuous data are updated periodically and thus essentially have to deal with “update events” as well. If the goal is to keep the display as unobtrusive as possible, one must then think about



(a) The Scope notifications monitor



(b) The AuraOrb

Figure 3.7.: **(a)** The Scope Display provides users with an overview of pending tasks and emails, appointments, and other alerts (Van Dantzich, Robbins, Horvitz, & Czerwinski, 2002). The more urgent a notification is, the more centrally it is placed. Furthermore, visual cues inform the user about important properties of the notifications. **(b)** The AuraOrb notifies its users of incoming emails with a comparably subtle glow (not in the picture) and only if they are not looking at the picture, it displays more detailed information, i.e. the subject and sender of the email (Altosaar, Vertegaal, Sohn, & Cheng, 2006).

how to design the transition between two states. In other cases, it must be decided how important those updates are and thus how “visible” they should be.

3.3.2. Intrusiveness and Notification Level

Related to that, an important characteristic of an ambient information system is its intrusiveness, i.e. how much the display attracts attention from the user. For a peripheral display, this should obviously be avoided in most cases. However, there might be a type of information that can justify such a behavior.

While most overview papers focus more on the notification level, i.e. which method is used to present the input data and especially how it handles updates and notifications, we argue that the intrusiveness of the artifact itself should also be taken into account. The integration into the physical environment, for example, is an important aspect for many ambient displays. Among others, this is the case for the two pieces of informative art discussed in 3.2.1 (also cp. Figure 3.3), which are considered as calm artifacts mainly because of their resemblance to a real painting.

Regarding the notification level, we loosely follow Matthews, Rattenbury, Carter, Dey, and Mankoff (2003) and Pousman and Stasko (2006) and distinguish between four

categories of how a system might display (updates of) data. Note that these are mainly defined based on how the display is used, as opposed to how the information is presented.

User Poll: For some systems, the user has to actively poll information. This is for example the case for the OS X Notification Center, which is Apple's implementation of the sideshow concept (cp. Section 3.2.2).

Change Blind: In this case, the presented information should not grab any attention at all and the designer has made an effort for the system to update the data without the user noticing. This is usually the case for glanceable displays, like the one used for the CareNet system (cp. Figure 3.6a), where information is presented at all times, but users check the display at their own will. It is important to note that the notification level is information-dependent (and not an inherent characteristic of the display). For example, the CareNet display could also implement a mode to grab the user's attention to inform about the elderly person falling over.

Make Aware: Here, the display is used in a divided attention situation (cp. Section 2.3). The user wants to stay aware of the information presented by the display without being overly distracted by it. This is true for most displays in the work context, for example the Notification Collage (Figure 3.5b), which is designed to keep users aware of their co-workers' activities, but should not distract them from their task at hand. Naturally, the "change blind" and "make aware" modes to display information are the notification levels most commonly found in ambient information systems.

Interrupt: Like for the previous category, this mode would also be used in a divided attention situation, but with the explicit goal of attracting attention. For example, the LumiTouch system discussed in Section 3.2.2 emits an "ambient glow" as a subtle way to provide awareness of the remote user's presence, but will also light up in a more intrusive way when the companion frame is squeezed and thus might grab the user's attention. This mode is also commonly used in notification systems like Sideshow or Scope, for example to inform the user about an upcoming appointment.

REALIZATION

While the notification levels reflect how data is displayed from a conceptual point of view, the subsequent (and more difficult) question is how they should be realized. Research from cognitive psychology has shown that the human perceptual system is not only sensitive to the intensity of a stimulus, e.g. the perceived loudness of a sound or the brightness of a light, but mostly to the onset of stimuli and (abrupt) changes in the environment (Wickens, Hollands, Banbury, & Parasuraman, 2015, p.53). Therefore, a flashing light or the sudden appearance of a sound is rather likely to attract attention. In this case, we can also distinguish between a salient presentation of information itself

and an additional cue guiding the attention towards the display presenting information in a non-salient way. On the other hand, in order to prevent the distraction of a user, a static appearance and slow transitions between states are advisable.

CONTEXT

While it is quite clear how to display information in these two cases, ambient information systems are operating somewhere in between those extremes and therefore it is difficult to assess the right amount of attracting attention. For example, a notification system might want to alert the user about an incoming email, but still allow users to easily ignore the information so that they can (first) continue working on the task they are focusing on. Therefore, the notification's intrusiveness should be just enough for a user to notice, but not more. However, there are several factors that influence the perceived salience of a notification and its potential to be noticed by a user, which could be subsumed as the context of interaction. The context mainly consists of environmental stimuli, which could potentially mask or distract from a notification as well as the task the user is currently engaged in. For visual stimuli, the direction of gaze is obviously also an important factor that determines if a notification can be perceived at all.

Ambient information systems would therefore be most effective in terms of unobtrusive information conveyance if they were combined with technology to deduct the context of interaction. An example of such a system is the *AuraOrb* (Figure 3.7b), which integrates an eye contact sensor in order to determine if the user is looking at a computer screen (Altosaar, Vertegaal, Sohn, & Cheng, 2006). This way, an incoming email is indicated only via a subtle glow of the *AuraOrb* when the user is occupied with the computer, whereas it presents the subject and sender of the email on its ticker tape display, when the user is looking away from the computer, i.e. the system can distinguish between two states of user engagement and thus two contexts. However, for the vast majority of research projects, no context sensing capabilities are attempted to be integrated.

Alternatively, the context of interaction should be known beforehand and also be relatively stable so that the display can be optimized for one particular situation. For example, although the exact direction of gaze cannot be determined without additional sensors, an office worker sitting at a desktop computer can still be seen as a comparably static and predictable context. However, many ambient displays are not designed for such a clear usage scenario and have to deal with varying contexts, e.g. they might be integrated into various environments and utilized by users pursuing a range of different activities.

3.3.3. Information Capacity

An integral characteristic of an ambient information display is the amount of data that is being presented. For example, systems like *Sideshow* or the *Notification Col-*

lage (cp. Section 3.2.2) usually display a multitude of (rather information-rich) items, whereas the eye-q system (see Figure 3.2a) only indicates the occurrence of an event, without providing more detail on it. The advantage of a very information-rich interface is obviously the possibility to convey a sophisticated and detailed view on a large amount of data. However, a user would very likely find it difficult to keep aware of the data while being engaged with a primary task, due to the complexity of the presentation. A simple display, like the eye-q system, on the other hand, is more easily perceived in a divided attention situation, but of course cannot provide as much detail in information. It therefore seems as if there is a tradeoff between information-richness and the viability of a peripheral display to be used in conjunction with a primary task.

REACTION AND COMPREHENSION

In their overview paper on notification systems, McCrickard et al. (2003) proposed to differentiate between the design objective of a user being able to *react* to a given notification and the *comprehension* of presented information. For example, a user should be able to quickly react to the hints from a navigation system which exit to take, but does not need to understand *why* to do so. On the other hand, for a manager monitoring a complex workflow in order to find out how to improve it, comprehension of the provided data is obviously vital. While we think that these design objectives are helpful for a structuring of a broader range of displays, we argue that a system requiring true comprehension of the presented information can barely be considered peripheral: According to research in cognitive psychology, the essential characteristic of an ambient information system to require only a small amount of attentional resources is most likely to be achieved when its provided information can mostly be processed automatically (cp. Section 2.4.4). However, for the comprehension, i.e. consciously making sense, of some presented data, no automatic processing is possible and thus the attentional demand is rather high. Such a system can therefore only be used as a secondary display if the primary task is already highly automated, e.g. to listen to podcasts while driving a car.

VARYING DETAIL OF INFORMATION

A few systems we have reviewed offer the ability to switch between an overview of the data and a more detailed presentation. For example, if a user touches one of the icons of the CareNet display, the photo of the elder is replaced by a separate view providing more details on the respective item (cp. Figure 3.6a). Similarly, notification monitors like Sideshow are usually designed to provide a varying amount of detail, e.g. through a tooltip mechanism. While we think that such functionality can greatly increase the utility of a peripheral display by allowing users to obtain more detailed information in situations where they have the attentional resources to actively focus on and interact with the display, we would also argue that it should not be considered an essential feature of an ambient information system, since it merely combines the features of a

traditional display, which is designed to be perceived via focused attention, with those of a peripheral one.

However, since the user's ability to process peripheral information can change, depending on the context of interaction, even in a divided attention situation, we think that a varying level of detail would be valuable even for plain peripheral displays when used in a less predictable (or less stable) setting. Optimally, this would be coupled with context sensing capabilities as it has been done for the AuraOrb system (cp. Figure 3.7b). Alternatively, the level of detail can be made user-configurable, so that it can quickly be changed, e.g. depending on the primary task.

3.3.4. Representation

As we have argued earlier, for an ambient information system it is not only important how information is presented, but also how the system as a whole manifests itself in the physical world. Based on our analysis of existing systems and keeping in mind Weiser's original thoughts on calm computing, we distinguish between five types of representation:

Traditional Computational Devices: A rather straightforward and practical way of implementing a peripheral display, especially for use in work contexts, is to use the existing and widespread computing devices, such as desktop computers, laptops, or smartphones. The *ShutEye* display, for example, presents recommendations about when it is advisable to perform certain activities that can influence sleep quality on the wallpaper of the user's mobile phone (Bauer et al., 2012). Similarly, *Scope* (see Figure 3.7a) runs either on a desktop or laptop computer. The main advantages of using these devices are the comparably stable interaction context (cp. Section 3.3.2) and the rather easy deployability, which allows the systems to be evaluated quite easily.

Physical: Besides the idea of engaging the "periphery of attention" (also cp. Section 2.4.1), a central part of Weiser's vision of calm computing is the notion of disappearing technology and computers receding into and becoming part of our environment (Weiser, 1991). While one could argue that traditional computers *are* by now a part of our environment, Weiser actually meant the emergence of alternative manifestations of computational devices that are not as easily identified as such. Therefore, most realizations of calm technology are dedicated physical artifacts, such as the *Hello.Wall* (Figure 3.1a) or the *Nimio* system (Figure 3.5a).

Screen-based: Although designed to be part of our environment, many ambient displays still make use of pixel-based technology usually employed in traditional devices to present information. This can be a screen to display a dynamic piece of art or a projector used to present information on the wall of a building (see Figure 3.3 and 3.4b).




Figure 3.8.: Two examples of more uncommon usage types. **(a)** The system developed by Sturm, Herwijnen, Eyck, and Terken (2007) provides awareness of individual speaking time and gaze behavior for the participants of a meeting, using the table as a shared display for visualization purposes. **(b)** Depending on the probability for precipitation, the Forecast Umbrella glows in varying intensities to support the decision if it should be taken when leaving the house (Tharp & Tharp, 2005).

Integrated: In a few projects, the goal of integrating a computational system into our environment is achieved by augmenting an existing object and thus transforming it into an artifact with display-capabilities. This is for example the case for the eye-q system and the History Tablecloth (see Figure 3.2a and 3.1b). Like Moere and Offenhuber (2009), we think that for systems visualizing data, this form of ambient display has great potential to accomplish Weiser’s vision of calm computing. Furthermore, based on our review of existing systems, it seems to be a rather underexplored area of research.

Virtual: Finally, there are displays that are only virtually present and do not have a physical representation. Obviously, this is possible only for systems that do not rely on visual means to convey information, such as the upstairs project, where sound is used to augment a room in an apartment and thus no physical artifact is needed to represent the system.

While it might seem to be challenging to convey information in a non-visual way, one could argue that this virtual representation of a computational system might actually be closest to what Weiser called “disappearing technology”.

3.3.5. Usage Type

While ambient information systems have in common that they aim to create utility by providing information to the user, there are, from a conceptual point of view, several approaches to do so.

Providing Awareness: The most common usage type is to provide an awareness for a given set of information – in most cases while the user is additionally preoccupied with a (not further specified) primary task. A main design objective for these systems is to present information in such a way that it can be absorbed very quickly and easily (also cp. Section 3.3.3).

Support a Specific Primary Task: In contrast, a few displays are designed with a specific primary task in mind, which they aim to support or enhance. One example is a system developed by Sturm, Herwijnen, Eyck, and Terken (2007), which provides participants of a meeting with feedback on individual speaking time and gaze behavior in order to facilitate an evenly distributed group communication (Figure 3.8a). A distinct advantage of these systems is that the information conveyance can be optimized for a very specific situation and thus they do not suffer from varying interaction contexts, which easily lead to inappropriate presentation of information (cp. Section 3.3.2).

Offer Information Based on Context: Finally, there are a few displays that are directly embedded into a specific context and offer information for example at the time of a decision: The Forecast Umbrella (Figure 3.8b), for instance, features a handle that glows more intensely with the increased chance of rain in order to give users a clear indication if it is advisable to take the umbrella or not (Tharp & Tharp, 2005). Similarly, the Weather to Go system, designed by Tünnermann, Zehe, Hemminghaus, and Hermann (2014), provides awareness of the upcoming weather conditions by playing a sound that characterizes the local weather forecast when the user opens the door to leave the apartment.

Looking at existing work, it is obvious that up to this point, the last two approaches have been used only for a few projects and we would argue that they would be worth further exploration.

3.3.6. Modality

A final characteristic of ambient information systems is the (primary) modality that is used for information conveyance. A categorization of existing systems based on this design dimension reveals that the vast majority of displays employ visual means to do so, either via traditional screen-based technology, e.g. in the Notification Collage (Figure 3.5b) or the Scope Display (Figure 3.7a), or by employing more specialized display technology, as for example in the Aura Orb (Figure 3.7b) or the Nimio artifact (Figure 3.5a).

However, while we have found that other modalities have received far less attention by existing research, there are also a few examples of systems, where sound is used as an additional (or even the only) way to convey information, e.g. the Upstairs system

(Figure 3.6b) or the “knock’knock” prototype discussed in Section 5.2, which hint at the possible advantages of employing this modality.

4.

Sustainable Human-Computer Interaction

The field of sustainable HCI is concerned with how interaction technology can be designed so that it supports the goal of (environmental) sustainability. Although its origins lie in the HCI community, the field additionally draws on theories and concepts from social and environmental psychology, interaction design, and behavioral economics. Being a relatively young field, it is rather difficult to make out a clearly defined research direction (also see Silberman et al., 2014). However, following DiSalvo, Sengers, and Brynjarsdóttir (2010), we can distinguish between five subgenres, which both help to define the field as a whole and allow us to give a structured overview of the research activities in this area that have been done until now.

4.1. Persuasive Technology

A common approach in the field of sustainable HCI is to use technology to *persuade* people towards behaving in a more sustainable way. Many of those projects have originated from a psychological context; the work by Fogg (2002) on “using computers to change what we think and do” is frequently cited. According to Halko and Kientz (2010), there are four commonly used strategies for persuasive technology:

Instruction: Here, the user is persuaded simply by being told what to do, for example by a virtual agent. This obviously depends on users already having a rather specific goal they want to achieve.

Social Feedback: This strategy leverages the influence people have on each other to persuade them to behave in a specific way. For example, several users can be allowed to compete with each other in achieving a certain desirable behavior change. Similarly, technology can also be used to facilitate the notion of users cooperating as a team towards this goal.




Figure 4.1.: Illustration of the persuasive elements used in the UbiGreen Transportation Display (Froehlich et al., 2009). **(a)** Two possible progressions of the phone’s wallpaper – assuming that the user more and more often chooses to use “green” modes of transportation. **(b)** Secondary benefits (e.g. saving money, relaxation, exercising, or the ability to do other things while traveling) are emphasized by highlighted icons, addressing further intrinsic motivational factors of the users.

Motivation: Persuasive technology can be used to either address and appeal to intrinsic motivators, such as the good feeling of doing the “right” thing, or to make use of extrinsic ones, such as virtual trophies or rewards.

Reinforcement: Here, the user is persuaded by introducing positive (e.g. as a reward), or removing aversive stimuli. For example, an avatar of the user can be beautified by adding certain items to its appearance, or a scene depicting a stretch of dying nature can be transformed into a blooming field of green.

Although this list of strategies is certainly not exhaustive (nor mutually exclusive), it gives a good idea of the way persuasive technology can work. One example for this type of systems is the UbiGreen Transportation Display, which shows the user’s semi-automatically tracked transit behavior on a phone’s wallpaper (Froehlich et al., 2009). By choosing not to go by car and instead use alternative modes of transportation, such as walking, going by bike, or taking the train, the depicted scene is progressively beautified, for example by adding more leaves, blossoms, and finally some fruit to a tree, or by increasing the size of a glacier as well as the number of polar bears, seals, and fish living on and around it (Figure 4.1a). Furthermore, for the most recently chosen mode of transportation, the display highlights secondary benefits, such as saving money or getting exercise along the way, as an additional motivation (Figure 4.1b).

OTHER USES AND RELATION TO NUDGING

Persuasive technology is not only used in the field of sustainable HCI, but can be found in other domains as well – most prominently in the area of healthy living, where it is used to persuade users towards eating less or doing more sports (e.g. Consolvo, Everitt, Smith, & Landay, 2006). It is also closely related to so-called *nudging*, which likewise

aims to alter people's behavior (Thaler & Sunstein, 2008). Nudging has been conceived to be implemented on a rather large scale and might be used, for example, by governments to increase tax paying compliance. One of the most commonly used nudges are *default rules*, which have been shown to be effective due to people's reluctance to deviate from a default configuration. For instance, the energy-saving mode of a washing machine would be used far more often if it was implemented as the default program and was not another option that can be selected additionally by the user. While most persuasive technology – especially those developed in the academic sector – is based on the assumption that users effectively *choose* to be persuaded (e.g. by actively installing and using the UbiGreen smartphone application), nudges are usually implemented by a third party and often people are influenced without even knowing about it, which is why it has been emphasized that nudges “must be easy and cheap to avoid”, and “without forbidding any options” (Thaler & Sunstein, 2008, p.6). Finally, it can be observed that on the one hand, nudges are more and more often implemented through technology, and on the other hand, concepts and ideas from nudge theory are influencing both the deployment and research on persuasive technology (Adams, Costa, Jung, & Choudhury, 2015).

4.2. Ambient Awareness

As the name already implies, research in this area is closely related to those on ambient information systems (also cp. Chapter 3). Based on our understanding of the field, we see the work in this subgenre being based on three important observations:

- First, the ability of people to understand and estimate the impact of their behavior on the environment and other issues of sustainability is an important and necessary precondition for them to be able to change it.
- For many of those issues, however, it is very difficult for people to relate their actions to any consequences due to the temporal and geographical distance between cause and effect. For example, we would describe electricity consumption as being essentially *invisible* to users, both because of its centralized and intransparent generation, but also due to its intangibility: Most people's experience with electricity is limited to “plugging a chord into an outlet” (Pierce & Paulos, 2010), which then ensures the functioning of a device for a basically unlimited time.
- Furthermore, although there are many parts of our daily behavior that have an influence on sustainability issues and therefore a *continuous* awareness would often be the most adequate way to support their understanding, a consciously attending to those issues would obviously be somewhat overwhelming. Also, the impact on those issues is often a byproduct of a specific activity, from which the user should be avoided to be distracted. For example, while the energy consumption that takes



(a) The Power-Aware Cord



(b) The Flower Lamp

Figure 4.2.: **(a)** The *Power-Aware Cord* provides information about the electricity usage of the attached appliances via a glowing pattern that is displayed on the cable itself (Gustafsson & Gyllenswärd, 2005). The measured consumption controls both intensity and movement speed of the pattern. **(b)** Design visualization of the *Flower Lamp*, whose petals slowly open up to bloom when the household's electricity usage decreases (Backlund et al., 2007).

place when driving a car obviously has an impact on the environment, the primary goal of a driver should be to safely arrive at the target destination, with avoiding to use too much energy being an important, but only secondary goal.

Systems developed in this subgenre therefore aim to convey information and provide awareness to users about the impact of their behavior on issues of sustainability, while avoiding to overwhelm and distract them.

EXAMPLE SYSTEMS AND OVERLAP WITH PERSUASIVE TECHNOLOGY

One example for such a system is the *Power-Aware Cord*, which augments a conventional electrical power strip with additional electroluminescent wires being integrated into the cable (Gustafsson & Gyllenswärd, 2005). This setup allows to create dynamic glowing patterns that can be displayed on the cable itself (cp. Figure 4.2a). Depending on the electricity consumption of the connected appliances, both intensity and movement of those visual patterns are modified, thereby allowing users to easily estimate and keep aware of how much a device is currently consuming.

While informing users about issues of sustainability lies at the heart of ambient awareness systems, the ultimate goal behind these efforts is most of the time to induce a behavior change towards alleviating those issues. It is therefore not surprising that

there is a certain overlap between persuasive technology and ambient awareness research. This is also reflected by the fact that many systems actually combine concepts and ideas from both areas. For example, the *Flower Lamp* is a re-design of a common ceiling lamp, with the ability to transition from a closed position towards opening up like a flower (Backlund et al., 2007). The appearance of the lamp depends on the current trend of consumption: If the household has a decreasing trend of electricity use, the Flower Lamp slowly opens up to “bloom”, whereas an increased usage leads to the lamp folding its petals together (see Figure 4.2b).

In its overall design, the Flower Lamp is thereby on the one hand an effort towards designing a calm entity in the home. On the other hand, the beautification of a decreasing electricity consumption is also a deliberate use of the reinforcement strategy that is commonly found in persuasive technology.

4.3. Sustainable Interaction Design

In contrast to the two previous approaches, which *utilize* technology to support and interact with users, this subgenre of sustainable HCI is concerned with the design process of technology itself. Mankoff et al. (2007) define this as sustainability *in* design, as opposed to sustainability *through* design. Research on this topic, however, is only at a very early stage and mainly consists of design critique and a collection of rather general ideas.

One example is the proposition by Blevis (2007) to *link invention and disposal*, meaning that the design process for any technological artifact should always also take into consideration what will become of the objects and systems that are displaced or made obsolete by it. As Blevis points out, this might not even be a tangible artifact, but can also be a piece of software that effectively renders older hardware obsolete by increased performance demands. A second, frequently expressed goal is to increase the longevity of devices, which can, for example, be achieved by designing modular hardware platforms that allow devices to be both upgradable and easier to repair.

The *Phonebloks* concept, for example, describes the design of a modular smartphone (Hakkens, 2013): A mainboard with universal connectors would allow to add any needed functionality to the phone by attaching “bloks” of varying size (see Figure 4.3a). However, the technical and mechanical difficulties to actually implement such a design are quite substantial, and *Project Ara* (Google, 2013b), which explored the feasibility of a similar concept, was ultimately canceled. Although there are currently two consumer-grade phones that allow to attach a range of “mods”, such as a camera lens, a battery pack, or an image projector (Motorola, 2016), the devices themselves are still conventionally-built, monolithic pieces of hardware. As of June 2017, the *Fairphone*




Figure 4.3.: **(a)** Illustration of the *Phonebloks design concept* for a modular smartphone consisting of a wide range of easily replaceable parts (Hakkens, 2013). **(b)** Disassembled prototype of the Fairphone 2 (Fairphone B.V., 2015), consisting of seven major building blocks: the back cover, a removable battery, the main chassis, rear camera, receiver module, speaker, and the display (Image from Ars Technica, 2015).

2 is probably closest to the original idea of a modular device, allowing users to easily replace any parts in order to increase its longevity (Fairphone B.V., 2015).

4.4. Formative User Studies

As a fourth cluster of work, dedicated user studies aim to contribute to the field of sustainable HCI by generating knowledge that might be used to improve existing systems or serve as the basis for new approaches. Most work in this area belongs to one of the following two categories:

- Studies to understand how users think about sustainability, for example their attitude towards the environment and their motivation to act in an (un)sustainable way, and how this might be leveraged by interaction technology.
- Studies critically evaluating existing (also commercial) systems designed to support sustainable behavior in terms of usability and how the users understand and operate those systems, as well as their influence on the users' behavior and attitudes.

For example, Strengers (2011) studied three types of in-home displays designed to provide an enhanced awareness of resource consumption by conducting group interviews with the inhabitants of households where such a display had been installed. In the study, she found out that several assumptions that had implicitly been made during the design of the displays did not prove to be true, and users often had difficulties

to correctly interpret the conveyed information and to react accordingly: Many householders reported not to understand the units that were used in the charts and figures of the displays, such as kilowatts, greenhouse gas emissions, and liters. Also, since temporary spikes in consumption were indicated rather prominently on the display (e.g. with a signaling red light), users easily assumed that appliances such as an electric kettle used more electricity than appliances with a more moderate, but continuous consumption, such as a fridge – although this is of course not necessarily true. Finally, Strengers observed that while both men and children reportedly showed far more interest in the feedback than women, they were often less involved in relevant consumption practices, as women were in control of most household activities.

As an example for the first category, the *motivation* for acting in a sustainable way has been examined in three studies, each dealing with a different group of people: Evaluating 15 “typical” households, Chetty, Tran, and Grinter (2008) conclude that changes in behavior towards an increased resource efficiency are most commonly motivated by monetary reasons, while the desire for comfort often acts as a motivator for the opposite. Unsurprisingly, financial reasons were frequently mentioned also for inhabitants of households with a comparably low income (Dillahunt, Mankoff, Paulos, & Fussell, 2009). Interestingly, while the desire to be environmentally friendly was reported to be only of far lesser importance for the average household, the topic was mentioned more frequently both for the low-income group and, although slightly differently framed, in a study dealing specifically with people who showed significant efforts to optimize their environmental responsibility (Woodruff, Hasbrouck, & Augustin, 2008), indicating the need for further research in this direction. Furthermore, moral and religious reasons were reported for people of both of those two groups, although they were described as far more diverse for the “green” households (e.g. also including more spiritual or “New Age” tendencies), in comparison to a more “conservative” faith in (and responsibility to) God for inhabitants of the low-income households. Finally, the goal to be more self-reliant or simply to set oneself apart from others were additional motivators that were frequently observed in the study dealing with the ecology-minded households.

4.5. Pervasive and Participatory Sensing

A final cluster of work is concerned with developing systems to monitor and report on specific aspects of environmental conditions – often with the underlying goal of changing and improving these conditions. In the area of sustainable HCI, most efforts in this direction focus on so-called *participatory sensing*, which aims to include a larger number of lay persons in the gathering of data, which can then be consolidated and interpreted, usually by a small group of experts, e.g. to actively coordinate any action to improve the reported conditions.

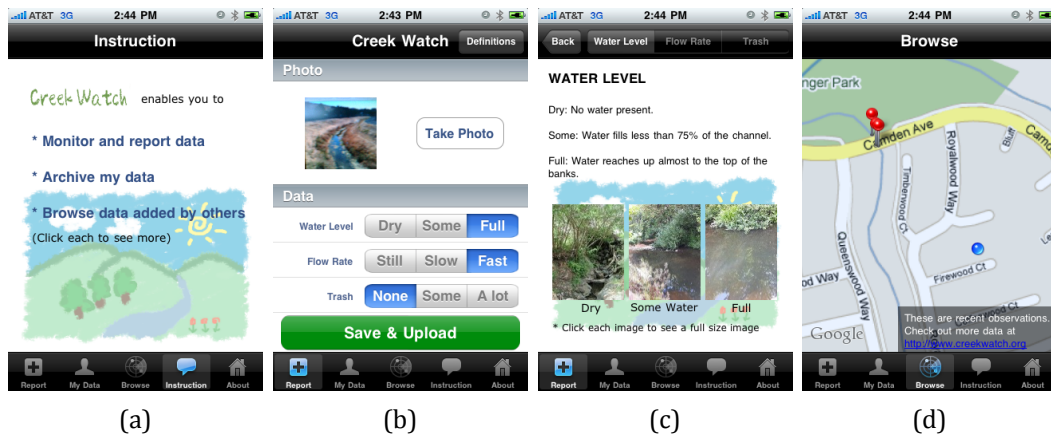


Figure 4.4.: Pictures from the *Creek Watch* iPhone App (Kim, Robson, Zimmerman, Pierce, & Haber, 2011). The images show (a) an introductory screen, (b) overview of an report, (c) a help screen for better categorization of the respective creek, and (d) recent observations in the user’s vicinity.

One example for such a project is the *Creek Watch* initiative, which aims to support water and trash management programs by enabling volunteers to report on the state of watersheds they encounter in their daily lives (Kim, Robson, Zimmerman, Pierce, & Haber, 2011). This is done with the help of a smartphone application, where users can indicate the water level, flow rate, and presence of trash for any creek they encounter and additionally take a complementary picture (see Figure 4.4). For this particular project, the researchers explored the use of HCI methods not only to design the capture interface, but also to ensure that the provided data is useful for the environmental protection programs receiving the reports.

5.

Auditory Display and (Blended) Sonification

5.1. Methods for Auditory Display Design

As we have discussed in Chapters 1 and 2, auditory information conveyance plays a major role in peripheral perception and in many cases its use should be advantageous to employing visual means to present information. At the same time, it has been largely neglected for the design of ambient displays, which is why we want to give a short overview of the most important methods in sonic interaction design and discuss the relevance and applicability for auditory ambient information systems as well as for the use in the area of sustainable HCI (for a more detailed description refer to Hermann, Hunt, & Neuhoff, 2011).

SPEECH

Probably the most widespread use of sound for communication purposes is the natural language: it is both highly flexible and has large expressive power. Consequently, there has been considerable research on speech recognition and synthesis to enable interactive systems to communicate with users in this way (e.g. O'Shaughnessy, 2003). The downside of its expressive power, however, is that language is a comparably complex form of communication and its understanding usually requires some cognitive effort and focused attention, which is why we think that it should only play a minor role for auditory displays designed for peripheral perception.

AUDIFICATION

Contrary to speech, audification is a very specialized and, in its essence, quite simple and straightforward form of sonification: As input, the method takes any time series with a wave-like shape, which is interpreted as a waveform signal that can be played back as a sound. For our purposes, the main disadvantage of this approach is that it

requires a very specific type of data and, despite the possibility of resampling the signal, a very large number of data points. In the context of sustainable HCI, audification might be used for example to review a household's energy consumption of the past month in a very short amount of time. The usefulness of such a representation as well as its suitability for daily use would still have to be critically evaluated, though.

AUDITORY ICONS

Similar to their visual counterpart, auditory icons are symbolic and metaphorical representations of an event or a discrete item in form of a short and easily identifiable environmental sound. For use in a peripheral display, we think that auditory icons should be suitable for a wide range of use cases – as long as they do not require a *continuous* representation of data – since they allow users to make use of their existing everyday listening skills, thereby making the auditory cues rather intuitively understandable and allowing for a peripheral perception without too much cognitive effort. Furthermore, auditory icons should be able to elicit an emotional response in users, which should be conducive to the design of persuasive technology in the area of sustainable HCI.

EARCONS

Quite similar to auditory icons, so-called earcons are commonly used to signal the occurrence of a certain event or to signify a discrete number of data values. They are short, symbolic, musical-like messages, which are used for example in modern operating systems to notify users about certain events, such as a USB device being disconnected. Different to auditory icons, they do not require an existing relationship between the produced sound and its meaning, which obviously makes the concept far more flexible. On the other hand, this also means that users must first *learn* the meaning of any introduced earcon to effectively make use of it, which we also see as the main disadvantage for use in an ambient display – especially when incorporating more than a few of them. However, contrary to auditory icons, earcons can be designed to be as short, simple, and distinct as possible, which could be used to make a peripheral perception of the provided auditory cues possible.

PARAMETER MAPPING SONIFICATION

Compared to the previously discussed methods, which are somewhat restricted in the type of data they can represent, parameter mapping sonification (PMSon) is a potentially quite complex technique that can be used to display any given multivariate data with a limited dimensionality. It is based on the fact that virtually every synthesis of sound is dependent on a number of parameters, such as amplitude, frequency, and waveform, and the central idea of the approach is to establish a mapping between one or several of the data set's variables and a selection of those parameters, which conse-

quently leads to a change in the perceived perceptual qualities of the sound, such as its loudness, pitch, or timbre.

Considering that a PMSon consists of both a specific mapping function and a particular type of sound synthesis, the design space of this type of sonification is comparably large – which also means that the suitability for use in a peripheral display highly depends on the actual implementation. Furthermore, users will first have to learn and understand the mapping function, which might be somewhat difficult for unexperienced listeners. Finally, the design process requires a certain knowledge about the perceptual qualities of sound, e.g. that the perception of pitch dominates (and potentially interferes with) the perception of other sound features. However, despite these difficulties we think that a well-designed PMSon can be suitable for peripheral perception, especially when replacing the synthesis of sounds with using existing environmental ones, as proposed in our framework of blended sonification, described in the following section.

MODEL-BASED SONIFICATION

The approach of model-based sonification is inspired by physical processes that govern the natural appearance of sounds, such as the acoustic response of an instrument or everyday object that we interact with. A specific instance of such a sonification therefore implements a model of a dynamic system, with a specification of a) the system's initial state, b) the evolution of the system in time, c) how the input data influence these model dynamics, d) the possibilities of exciting the system, i.e. how to feed new energy into it to produce a (sonic) reaction, and finally e) how the model's dynamic state is translated into producing an appropriate auditory response (Hermann & Ritter, 1999).

A practical disadvantage of this approach is that it is computationally far more expensive than the previously discussed methods. Furthermore, it has been developed primarily for use in exploratory data analysis, where a manual excitation of the system is used to purposefully examine a given data set, which is obviously not suitable for a peripheral display. However, these excitations can also be generated by certain events derived from the input data or, in accordance with the concept of blended sonification, be coupled to users' everyday interaction with objects. Furthermore taking into account that due to the similarity to physical sound generation users can potentially make use of their everyday listening skills, this sonification technique could very well be used in a peripheral display.

FURTHER APPROACHES AND COMBINATIONS OF METHODS

Besides these fundamental approaches, there exist several variations and combinations of the methods described above. For example, *parameterized auditory icons* exploit the fact that everyday sounds can be quite expressive and usually carry additional information about the sound-producing incident, such as the size of a bell that is being

rung, and allow an event to carry additional information based on the modification of a sound's acoustic property.

Similarly, *spearcons* utilize methods of speech synthesis to create a specific type of earcons that consist in highly sped up spoken phrases, which are not recognized as speech anymore, but can help to quickly identify items in a menu-based interface (also cp. Walker, Nance, & Lindsay, 2006).

5.2. Blended Sonification

As an auditory representation that should be particularly suited for use in an ambient display, we have developed the concept of blended sonification (Tünnermann, Hammerschmidt, & Hermann, 2013). Similar to the work on ambient displays, it is motivated by the observation that the user interaction with modern information systems is mostly still achieved through graphical, WIMP¹-style interfaces, which are increasingly becoming a chokepoint for the huge amount of information that is available to us and which frequently bind our attention to the narrow screens of our smartphones.

As a specific form of auditory representation of information, blended sonifications try to relieve this chokepoint not only by addressing a different modality, but also by putting special consideration on the *existing* environmental soundscape (cp. Figure 5.1). More specifically,

Blended Sonification describes the process of manipulating physical interaction sounds or environmental sounds in such a way that the resulting sound signal carries additional information of interest while the formed auditory gestalt is still perceived as coherent auditory event.

Furthermore, *“Blended Sonifications should be calm, well motivated and expectable by the user. They should stay in the periphery, but be ready at hand.”*

While this working definition already gives us a good idea of what a blended sonification might be, it also hints at the importance of a range of key characteristics and design principles, which will be discussed in the following.

USE OF THE EXISTING SOUNDSCAPE

In order to prevent an “acoustic pollution”, which can easily arise when creating auditory displays that have been designed without consideration of the context they are embedded in, the central idea of blended sonification is to take into account the existing environmental sounds for the design of an auditory representation of any kind of data. Moreover, those sounds should be used as the *basis* for such a sonification and be manipulated in a way that it can convey additional information to the user.

¹WIMP: Graphical user interfaces dominated by windows, icons, menus, and pointing

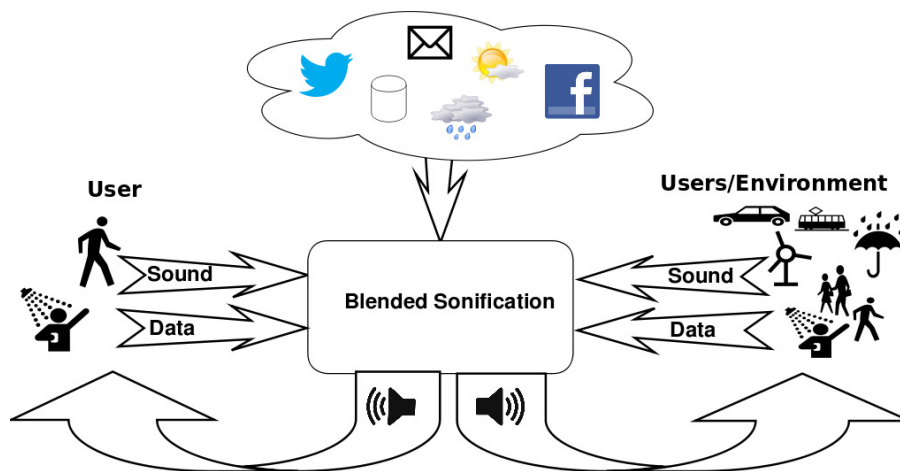


Figure 5.1.: General concept of a blended sonification (Tünnermann, Hammerschmidt, & Hermann, 2013). Taking into account both environmental and user-induced sounds, the resulting auditory output is coherent with the existing soundscape, thereby *blending* the presented information into the users' environment.

While the existing soundscape can be composed of a multitude of sources, physical interaction sounds caused by the user can be considered as particularly suited for a blended sonification, as those are most expectable for a user. For example, the augmented keyboard, developed by Bovermann, Tünnermann, and Hermann (2010), enriches the typing sound of a keyboard according to data of interest, e.g. the current weather situation outside (cp. Figure 5.2). More precisely, the interaction sounds are picked up by a contact microphone, as can be seen in Figure 5.2, and then filtered in such a way that they still share the basic characteristics of the original sounds, but additionally incorporate features of the external data. The resulting sonification, controlled by data-driven filter parameters, is played back in real time to the user by nearby loudspeakers, so that it blends together with the original typing sounds into a stream of coherent auditory events.

PERIPHERAL AND CALM INTERACTION

Exemplary for a blended sonification, the augmented keyboard enriches an existing interface without requiring users to change their behavior or to actively query for the provided information. Instead, the information is offered to them at all times when interacting with the keyboard and is readily available if a user chooses to pay attention to it. A blended sonification therefore makes use of our ability to tune out background noises of our surrounding soundscape (cp. Section 2.2). Moreover, an alteration in the augmented sounds, as would be caused by a change in the external data, can temporarily attract the users' attention towards the relevant characteristics of the sounds – before they become background noises again. Put differently, the augmented sounds



Figure 5.2.: The augmented keyboard (Bovermann, Tünnermann, & Hermann, 2010). The sounds of a user typing on a keyboard are picked up by a contact microphone, filtered, and fed back to the user. The filter parameters are determined by external data, e.g. the current temperature or humidity.

enable users to stay aware of any data changes at the threshold of their conscious attention.

As a second example of a blended sonification, the “knock’knock” system augments the interaction sound of someone knocking at an office door by a certain amount of reverberation, depending on the amount of time the sought-for person has been gone. Here, the visitor’s active query, which is primarily directed at the office’s owner, is enriched to convey additional information in case the office is empty (Tünnermann et al., 2013). Similar to the augmented keyboard, this example demonstrates that the blended sonification approach aims to prevent users from being confronted with any explicitly perceived technology, but rather to enrich the natural environment by additional information channels.

EXPECTABILITY AND COHERENCY OF SOUNDS

Another design principle demonstrated by this example is that the resulting sounds should be *expectable* by users, which is not only achieved by avoiding to change any important auditory features of the original sounds, but can also be supported by plausible metaphors employed in the auditory design: Knocking against the door of a completely empty room does indeed result in having a rather large amount of reverberation and thus the resulting sound should be somewhat familiar to the respective person.

It is important to note, however, that a blended sonification does not necessarily rely on filtering alone to enrich interaction sounds or the environmental soundscape, but can also use added sounds to do so, as long as the resulting auditory response is experienced as a coherent unit. More specifically, the added sounds should match both the

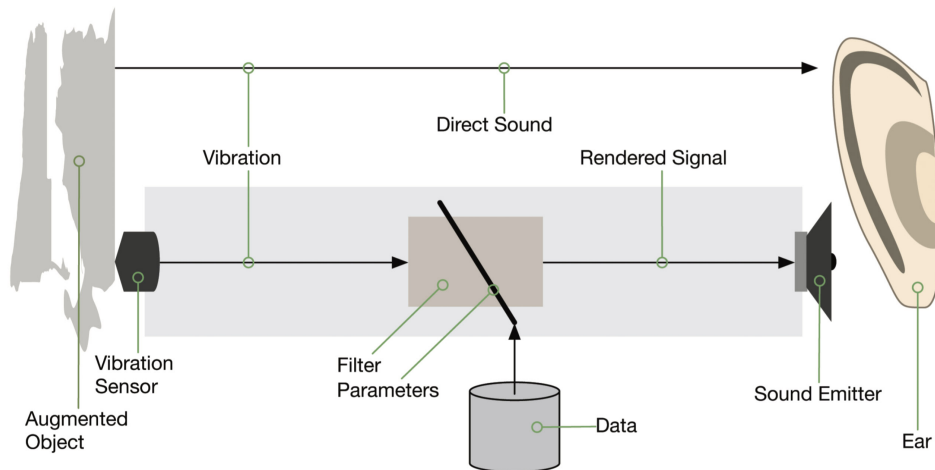


Figure 5.3.: Illustration of the concept of auditory augmentation (from: Bovermann, Tünnermann, & Hermann, 2010).

onset and the loudness of the original auditory events. With these considerations in mind, the concept of auditory augmentation, developed by Bovermann et al. (2010), can be seen as a special case of blended sonification, which restricts the auditory input to structure-borne interaction sounds caused by the user and solely relies on filtering to incorporate additional features according to the data of interest (cp. Figure 5.3).

In conclusion, different from the majority of sonification methods, which consider sound primarily as a consciously attended stream of information, the concept of blended sonification provides us with a framework for designing auditory displays that target the periphery of attention and blend into the users' environment by enriching their surrounding soundscape.

6.

Research Agenda

Having examined the most important related work in the previous chapters, we can now formulate a set of goals and research questions that we want to tackle in this thesis.

6.1. Goals and Research Questions

First, bringing back to mind our principal research direction, we want to investigate:

RQ1: *How can ambient information systems support environmentally-conscious behavior? What are use cases that “work” for using such an interface? What are not?*

In order to do so, we have designed several prototypical systems based on our insights gained from related work, in order to extend and improve upon the current state of the art. However, we also think that a critical limiting factor, both in the field of sustainable HCI and in ambient display research, is a lack of thorough evaluation and conduction of studies to expand the knowledge about those systems and, if possible, to generate generalizable insights that shape the overall direction of research and that can be used and built upon in subsequent projects.

We therefore see the development of additional systems not only as a way to systematically explore and extend the design space of ambient interfaces, but also as research vehicles that allow us to contribute to answering specific questions. Consequently, for each developed system, we have conducted extensive studies – both for evaluation purposes, but also to examine the users’ behavior and habits in the respective context in order to provide further knowledge to build upon when designing systems addressing the same domain.

6.1.1. The Use of Sound in Ambient Displays

In our survey of existing ambient information systems, we have seen that only very few of them have attempted to use auditory means to convey information to the user and that the vast majority is designed to address only the visual sense (cp. Section 3.3.6). On the other hand, we have learned that other modalities might often be better suited to be used in these systems: Discussing the information capacity of ambient displays in Section 3.3.3, we have concluded that a central and inherent problem of those displays is that conveying (lots of) information apparently interferes with performing an attention-demanding primary task.

Interestingly, according to findings in cognitive psychology, the reason for this interference might often be a conflict at the sensory level, i.e. a primary task requiring visual attention directly hinders absorbing information presented visually by a peripheral display (Section 2.3.3). Since most activities that would be considered a primary task are of visual nature, it therefore makes sense to avoid using the same modality to convey information. Finally, as discussed in Section 3.3.4, the purely “virtual” representation of a computational system that is achievable with an auditory peripheral display can be seen as being closest to Weiser’s vision of *disappearing* technology. Consequently, we want to know:

RQ2: *Is using sound in a peripheral display a feasible alternative to (purely) visual ambient interfaces?*

6.1.2. Complexity and Variability of the Context

Furthermore, in Section 3.3.2, we have discussed, from a more theoretical point of view, the relevance of the different aspects of the interaction context for a peripheral display and how knowing this context might be advantageous in that can allow to make better decisions about how salient a provided information should be. Focusing on the users themselves as what we think is the most important aspect of a display’s context, we therefore want to investigate:

RQ3: *How does the variability of an ambient display’s context of use affect the way it is perceived?*

In order to do so, we have deliberately developed our prototype systems to cover a range of settings with differing use contexts.

6.1.3. Unobtrusiveness and Acceptance

In Section 1.3, we have learned that unobtrusiveness is one of the core attributes of ambient information systems. We are therefore interested in:

RQ4: *What are effective ways to make ambient interfaces less obtrusive and distracting?*

On the other hand, considering the popularity of technical systems that require focused attention, such as smartphones or conventional computers, one might question the necessity of this quality, since it could be possible that, although unobtrusive interfaces are, in theory, advantageous for a user, they still might not be seen as something desirable. Consequently, we also want to know:

RQ5: *Are more unobtrusive ambient displays better accepted by users?*

6.1.4. Feedback, Persuasion, and Affectiveness

When discussing ambient awareness systems as part of sustainable HCI in Section 4.2, we have seen that most of them are based on the idea that the provided information will effectively *persuade* users to behave in an environmentally-friendly way. And in fact, many authors readily assume that eco-feedback more or less directly translates into a change in behavior – most of the time without adequately testing this hypothesis. For example, in their paper describing the Power-Aware Chord (cp. Figure 4.2a), Gustafsson and Gyllenswärd (2005) state that the users’ increased awareness of the different amounts of energy that is consumed by the various electric appliances will ultimately “lead them to question their energy behaviours”. However, while there are indeed several studies showing that feedback on energy consumption *can* lead to a reduction (e.g. Darby et al., 2006), there are also studies that cannot find any significant effect on consumption (e.g. Nilsson et al., 2014), suggesting that feedback does not *necessarily* lead to lower energy use. Therefore, we want to contribute to answering:

RQ6: *Is providing feedback on consumption practices enough to induce a corresponding change in behavior?*

Focusing on ambient displays as our primary subject, we also want to pose the more general question:

RQ7: *Is it viable to use ambient displays as persuasive technology?*

Although a connection is frequently made between ambient displays and persuasive technology (e.g. DiSalvo et al., 2010), we think that there is at least some reason to question this relationship, since the limited information capacity of those systems effectively prevents the implementation of most of the strategies commonly employed in persuasive technology (cp. Section 4.1).

Finally, there is preliminary research that hints at the effectiveness of using affective cues in order to persuade users (e.g. Obermair, Reitberger, Meschtscherjakov, Lankes, & Tscheligi, 2008). However, only limited work has been done in this particular area and thus we aim to contribute to answering:

RQ8: *Is the use of affective cues beneficial for ambient displays aiming to persuade users towards a change in behavior?*

6.2. Research Prototypes

In order to investigate these questions, we have designed and implemented four prototypical ambient information systems in three different environments:

The InfoPlant: This system has been designed to be used in a residential or office environment and is built around a living plant. In our study, it has been used to give feedback to users on their electricity consumption (Chapter 7).

The Sonic Shower: More focused on the bathroom environment, this peripheral display provides auditory feedback on the energy and water consumption when taking a shower (Chapter 8).

The EcoSonic System: In a different context, the EcoSonic system aims to support car users in driving in a more economical way by providing feedback either on the instantaneous consumption or, more specifically, when a high engine speed leads to an increased energy usage (Chapter 9).

The Slowification System: Similarly, this system provides feedback on driving too fast (or too slow) through spatial panning of the car's audio system's sound signal (Chapter 10).

We have selected to build this particular set of systems based on the following considerations:

- The respective use cases involve environments with differing complexity: While the InfoPlant aims to influence a wide range of activities that might be performed in a household, the Slowification system targets a very specific aspect of driving a car (cp. RQ3).
- In the field of sustainable HCI, the selected areas of application are mostly ones that have not yet been researched extensively.
- All of the above systems have been developed to explore the use of sound as means to convey information, with the last three even relying entirely on auditory cues to do so (RQ2).
- They also implement novel ways of unobtrusively conveying information to the user (RQ4) and explore the use of affective feedback (RQ8).

Part II.

Prototype Systems

7.

The InfoPlant

7.1. Introduction

As we have learned in Section 3.3.4, a central part of Weiser’s notion of calm computing is that of *disappearing* technology: devices that “weave themselves into the fabric of everyday life until they are indistinguishable from it” (Weiser, 1991). In preliminary research on ambient displays and calm computing, a common way to work towards this quality was to create new objects that were designed to blend into their surroundings.

When designing the InfoPlant, we have deliberately taken a different path: The proposed approach is to use an *existing* artifact that is well-established as an everyday object, or even appreciated as a beautiful addition to the environment, and technologically *augment* this artifact in order to enable it to function as an ambient display. When thinking about possible entities with these qualities, several objects might come to mind, which are in some way important to the respective person. However, we think that few are as universally appreciated as plants are. Being part of our natural environment, they can be found almost everywhere, including in our homes and our offices. Furthermore, they already have their own way of communicating to us some aspects of their overall state: Most of the time, one can easily tell if a plant needs more water, based on its general appearance, and due to their phototropic behavior, plants can also be seen to turn towards a specific direction (cp. Liscum, 2002). Finally, most people intuitively have a generally positive attitude towards plants, and as “primal” natural entities in our surroundings, they are able to elicit feelings of connectedness and care, as has also been hinted at in a recent study (Holstius, Kembel, & Hurst, 2004).

This chapter is, in parts, based on:

Hammerschmidt, J., Hermann, T., Walender, A., & Krömker, N. (2015). InfoPlant: Multimodal Augmentation of Plants for Enhanced Human-Computer Interaction. In *6th IEEE International Conference on Cognitive Infocommunications, 2015* (pp. 511–516). IEEE.

This last aspect is of particular importance when thinking about using ambient information displays as eco-feedback devices. Here, the presented information is supposed to lead to an adaptation of the user's consumption patterns, and ultimately to a decrease in overall expenditure. Due to the aforementioned reasons, it is reasonable to assume that an ambient display that is based on a plant can be expected to be more persuasive than an artificial one (also cp. RQ8).

7.2. Related Work

While there is a small body of work on using plants in a technological context and also as artifacts for interaction with humans, the aspect of unobtrusiveness and persuasion has rarely been studied empirically. The state of the art can roughly be categorized into four categories:

TECHNOLOGICAL INTEGRATION

Dealing with the broader topic of integrating plants into the technological world, Tanaka and Kuribayashi (2007) developed a toolkit for the design of artifacts for "human-plants-computer interaction". They propose two design strategies and several general patterns for plants being used in different ways (Kuribayashi, Sakamoto, & Tanaka, 2007). On the sensory-side, the toolkit focuses on measuring so-called "biopotentials", which are somewhat difficult to analyze, but can give an indication of the general condition of the plant. As an example, the authors have developed a demonstrator system, where measurements of the plant's biopotentials are mapped onto the color and intensity of a range of LEDs installed within the flowerpot (cp. Figure 7.1a). Changes in biopotential can occur due to the user watering the plant, changes in the environment, or the plant being touched by the user. In line with this example, the focus of their work lies less on the use of plants as ambient displays controlled by (external) data, but more on using the plant as a sensor. The authors also emphasize the notion of using the toolkit as a "versatile creative environment" for edutainment applications.

On a more abstract level, Christos Goumopoulos and Kameas (2004) transform ordinary plants into "ePlants", which feature an additional layer for interfacing with the surrounding (technological) environment. Here, the focus lies primarily on the integration of plants into a larger distributed network in order to establish "mixed societies of communicating plants and artifacts", with both sensory and interactive communication and distributed decision-making. As a conceptual framework, the authors have developed an ontology that defines the possible relations between ePlants, sensors and eGadgets.



Figure 7.1.: Pictures of related work dealing with plants being embedded into a technological context. **(a)** shows an augmented flowerpot with integrated LEDs to indicate changes in the plant's biopotentials (Kuribayashi, Sakamoto, Morihara, & Tanaka, 2007). **(b)** depicts the controlling of a plant's growing direction through phototropism, as used by Holstius, Kembel, and Hurst (2004).

ECO-FEEDBACK

To the best of our knowledge, Holstius et al. (2004) are the first to use a plant in an eco-feedback context. They designed a "living plant display" that is effectively able to lean to the side and can thereby function as an indicator for a one-dimensional variable. Using fast-growing corn seeds and two daylight bulbs, they took advantage of the plant's phototropic behavior, i.e. its tendency to grow towards a light stimulus, and by activating one of the two bulbs, it was possible to change the direction of growth and thus to produce a visible lean of the plant in a relatively short amount of time (cp. Figure 7.1b). Additionally, a robotic, biomimetic plant was developed, which emulates the looks of the natural one and is also able to tilt its leaves to the side.

The authors evaluated and compared both designs in a university's cafeteria, where they were put between pairs of recycling and trash containers, with the corresponding light being triggered when visitors threw away recyclables on the one side, or trash on the other side. As a result, they observed a slight increase in recycling behavior, which was largest for the natural plant display. According to a number of short interviews, this can be attributed to feelings of appreciation and caring for the plant.

Although the practical feasibility of this particular design can certainly be questioned, as it works only during the growing phase of a fast-growing plant and furthermore requires a considerable amount of lighting¹, the conducted study hints at an enhanced

¹According to the authors, a visible lean can be achieved with an exposure of 8 hrs/day of a 100W lightbulb in 3-4 days.

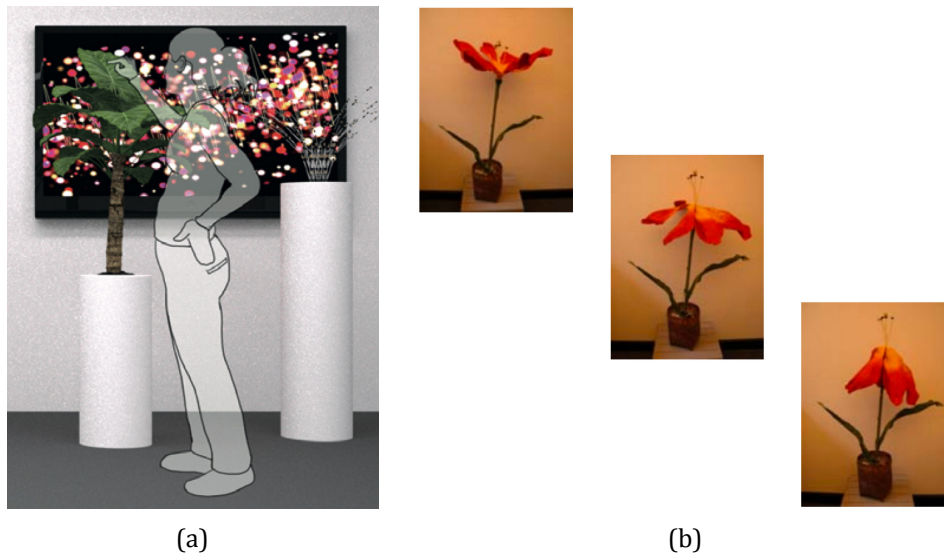


Figure 7.2.: Pictures of related work dealing with plants being embedded into a technological context. **(a)** illustrates a capacitive touch sensing method that can be used to detect a user's interaction with the plant (Poupyrev, Schoessler, Loh, & Sato, 2012). **(b)** shows an electromechanically augmented artificial lily that is able to let the blossom hang down or open it up to full bloom (Antifakos & Schiele, 2003).

efficacy of plants as persuasive entities, when compared to artificial artifacts. On the other hand, the authors point out that the measured results are not statistically significant and consequently these issues need further investigation.

PLANTS AS AN INTERFACE

Dealing more specifically with using a plant as interface technology, Poupyrev, Schoessler, Loh, and Sato (2012) employed the "Swept Frequency Capacitive Sensing" method developed by Sato, Poupyrev, and Harrison (2012) to detect any tactile interaction of a user with a plant's leaves. This is achieved by exciting the plant with an alternating current at several frequencies and simultaneously measuring the return signal to detect any changes in the frequency response and thereby, if (and where) the user has touched the plant (cp. Figure 7.2a).

ARTIFICIAL PLANTS

Finally, moving away from *living* plants, Antifakos and Schiele augmented an artificial lily (Figure 7.2b) with a mechanism to actuate the flower's petals, thereby being able to let the blossom hang down or open it up to full bloom (Antifakos & Schiele, 2003). Additionally equipped with a microphone, they used the flower as an ambient display in group meetings to show which participants were dominating the discussion – and which were left out of it. While no evaluation in terms of efficacy were conducted, the

authors observed that the flower was very quickly accepted as part of the environment.

In summary, we can see that, while various ideas and application scenarios have been proposed, research in this direction is still only at the beginning and needs further investigation. Also, no prototype has yet been developed to assess the long-term characteristics of plants used in a human-computer interaction context.

7.3. Concept of the InfoPlant

As we have seen in Section 3.3.2, an important aspect of an ambient display is its (potential) absorption of attention. Since such a display can be expected to rarely present information that is relevant for the *primary* activity of the user, a fundamental design goal is the overall unobtrusiveness of the system, and for the design of specific modalities, i.e. of ways to convey information to the user, one would usually aim for a low attention absorption (cp. RQ4). However, having the *potential* to alert the user in order to (temporarily) bring the attention to a certain issue obviously adds to the versatility of the system, for example when dealing with so-called horizon activities, which are monitored with the intent to become the user's primary activity in the near future (cp. Matthews, Rattenbury, & Carter, 2007).

Another design dimension that is of particular importance for a "natural" ambient display such as the InfoPlant is the possible speed of change of the different modalities: Although it is certainly attractive to use the plant itself to convey information, we must keep in mind that its changes in appearance must be expected to be relatively slow (also cp. Section 7.2). Finally, while this is obviously only possible to a certain degree, we want to keep the plant as natural as possible in order to retain its positive qualities, as discussed in Section 7.1. Consequently, we first have conceptualized a range of possible modalities to be used for the InfoPlant, each with a different potential attention absorption and speed of change. In Figure 7.3, we can see concept sketches of those modalities, which will be reviewed individually in the following.

OVERALL STATE OF THE PLANT

As we have discussed before, plants already have their own way of letting us know about their overall state in a rather unobtrusive way, e.g. by letting their leaves hang down, and we propose that this characteristic could also be used to represent the state of a (slow-moving) one-dimensional variable. Although there are also other ways of influencing a plant's overall appearance, for example by mechanically controlling its posture, as it has been done by Antifakos and Schiele (2003) for an artificial lily, the most authentic change in appearance can certainly be achieved by giving or withholding elementary elements for growth, which will then, indirectly, affect the plant's appearance

(cp. Leopold et al., 1964). While adjusting the intake of, for example, carbon and oxygen is largely unfeasible, as we would have to change the quality of the surrounding air, there are three elements that are commonly used to influence (mostly improve) a plant's growth. These are (a) light, (b) mineral nutrients like nitrogen, phosphorus, and potassium, and (c) water.

- (a) Although artificial light is certainly a possibility to boost a plant's growth (cp. Massa, Kim, Wheeler, & Mitchell, 2008) and, as we have previously seen, also to change its direction, the lighting itself must of course also be considered and can be seen as an unwanted (and rather salient) distractor, which is obviously something to avoid in an ambient display. Additionally, we would consider the amount of energy that is needed to achieve a visible change in appearance as unreasonably high for merely displaying a piece of information.
- (b) In order to control the availability of nutrients for the plant, such as nitrogen, phosphorus, and potassium, the soil would have to be precisely monitored, which is technically quite challenging (e.g. Hussain, Gondal, Yamani, & Baig, 2007). Moreover, not only would the visible changes in appearance be very slow, but due to the initial availability of nutrients in the soil, there would also be no possibility of withholding them for a rather long period of time.
- (c) On the other hand, controlling the water intake of a plant is comparatively easy, since soil moisture can be measured with a relatively simple sensor, and giving (or withholding) water can be achieved with the help of a water pump (also see Figure 7.3a). Although the capacity of both the soil and the plant itself to store water makes it difficult to change this parameter quickly, it should still be the fastest way to control the plant's overall appearance in a natural way.

ILLUMINATION

Different from the previous modality, augmenting the plant with artificial illumination would mean to definitely move away from a purely natural artifact towards an at least partially artificial one. However, as discussed earlier, this cannot completely be avoided when aiming for an ambient display with a certain expressive power. Also, it should be possible to install a number of LEDs in a way that they are barely visible, so that we can achieve a quite unobtrusive lighting of the plant. After discussing a number of possibilities of how to integrate a single or multiple LEDs into the InfoPlant, we decided that it would be best to install an RGB-LED-Stripe at the inside of the flowerpot, just above the soil, so that the plant can be illuminated from all sides from below, and with a range of different colors (cp. Figure 7.3b). With this, we would gain a quite versatile modality for communicating with the user – both in terms of subtlety and in speed of change.




Figure 7.3.: Illustrations of the modalities of the InfoPlant: **(a)** controlling the overall state of the plant by regulating the water intake, **(b)** illumination of the plant with a ring of RGB-LEDs, **(c)** tactile interaction with the plant, **(d)** rustling the leaves with the help of a fan, **(e)** changing the posture of the plant by tugging at a twig, and **(f)** integration of sound output.

TACTILE INTERACTION

While primarily designed as an ambient display, which should be able to unobtrusively communicate information *to* the user, a certain level of interaction would allow for a much greater flexibility when developing applications for the InfoPlant and can be helpful for switching between different display modes or to actively query specific information. Based on work by Poupyrev et al. (2012), we decided to integrate a capacitive touch-sensor into the InfoPlant, as this has the advantage of being relatively easy to install and allowing the user to actually get “in touch” with the plant, and at the same time keeps the plant itself as natural as possible (cp. Figure 7.3c).

LEAF RUSTLING

During the design process of the InfoPlant, we thought about which motions or movements of plants can also be observed for non-augmented species and therefore would be most natural for a “plant interface”. Something that most people have experienced before is the wind rustling through the leaves of a plant, generating both a soft sound

and a subtle motion. Usually, this process holds no other information than indicating the presence of wind. However, when generated for a plant in a living environment, this could be an additional, unobtrusive modality for signaling an event or displaying a specific status. The most straightforward way of recreating this rustling of leaves is by using small fans to produce the needed amount of wind, which is also what we did for the InfoPlant (Figure 7.3d).

POSTURE OF THE PLANT

While influencing the overall state of the plant as discussed at the beginning of this section would certainly be the most authentic change in appearance, it is also a very slow one. An alternative could be to actively modify the posture, so that the positioning of branches and twigs can be used to *mimic* a natural change in appearance. Furthermore, with a more precise control of the plant's posture, additional ways of conveying information would become possible, such as using a single twig as a gauge for a continuous variable.

Different to Antifakos and Schiele (2003), who used an actual shaft to open and close the petals of an artificial flower, we think that for a natural plant, attaching nylon threads to its twigs, which are then connected to small servos within the flowerpot, might be the best solution to implement this modality (cp. Figure 7.3e). This way, the augmentation of the plant would still be very subtle and barely noticeable – even when users are quite near it.

SOUND

Finally, the modality of sound would open up a completely new design space for the InfoPlant (also cp. RQ2). It should be easy to integrate, as we basically only need an additional loudspeaker, and it keeps the plant itself in its natural state (cp. Figure 7.3f). Furthermore, the modality is extremely versatile insofar as it allows the use of a very broad range of methods and sound designs, from short sonic events to extended, stationary and even continuous sounds to convey information. Also, sound can further accentuate and augment other events such as moving a twig or changes in illumination (cp. Tünnermann et al., 2013).

7.4. Construction

The construction of the InfoPlant can be divided into three phases:

Prototyping: In the first phase, we used the BRIX₂-platform, developed by Zehe (2017), to freely experiment with the different modalities discussed in the previous section. BRIX₂ is a framework for physical computing based on the Arduino project², with small programmable, LEGO[®]-compatible modules as its

²The Arduino Project: <http://www.arduino.cc>, last accessed: 2017-07-20

central components. Each module has several built-in sensors, a low-power radio transmitter for wireless communication between the individual modules, and additionally can be enhanced by a range of extension modules. During the prototyping phase, each modality was assigned its own BRIX₂-module in order to allow for a more flexible and independent implementation.

Integration into the first prototype: Subsequently, we integrated all functionality into a single Arduino and built a wooden casing to implement all modalities in one place. In addition to the low-level functionality implemented on the Arduino, we integrated a Raspberry Pi and added a high-level Python API for easier development of applications. An image of the first prototype can be seen in Figure 7.4.

Construction of the second prototype: Based on practical evaluation of the first prototype as well as a first study conducted to test its general feasibility and its appeal to potential users (cp. Hammerschmidt et al., 2015), we built a second prototype of the InfoPlant, which was then evaluated under real-life conditions in a number of test households (see Section 7.5). Pictures of the second prototype being installed in some of those households can be found in Figure 7.8.

In the following, for each of the conceived modalities presented in the previous section, we will discuss technical details of their implementation and our experiences in the different phases of construction.

OVERALL STATE OF THE PLANT

In order to control the overall state of the plant, we used a commercially available soil moisture sensor³ in combination with a low-cost industrial peristaltic pump, which has the advantage of keeping the water in the connecting tubes instead of letting it flow back, thereby making the water supply more precise. For the first prototype, the water tank was installed in a lower compartment of the casing, below the plant itself.

In our experience, a major difficulty with this modality is to determine when to give how much water in order to achieve a certain appearance of the plant: We discovered that the moisture level, where a plant is not completely drying out, but still lets its leaves hang down is quite narrow, and we obviously want to prevent the plant to wither. Furthermore, the control loop between water pump and moisture sensor has quite a large delay of several minutes, as the water accumulating in the flowerpot takes some time to be distributed in the soil. In order to have the most precise measuring of soil moisture, we collected the sensor's response to a range of moisture levels by putting a specific amount of water into several small containers filled with completely dried-out soil and logging the resulting voltage readings of the sensor after waiting for 15 minutes. These values were then used to generate a transfer function for the

³VH400 Soil Moisture Sensor Probe: <http://vegetronix.com/Products/VH400>, last accessed: 2017-07-20

moisture sensor, allowing us to better estimate the amount of water present in the soil. More difficult, however, is to identify a mapping from moisture level to the appearance of the plant, as it reacts to the given water with an even greater delay. While it was possible for us to define moisture thresholds for our setup, it seems to be difficult to find a set of universal parameters, which work for a broader range of plants.

With regard to the overall appearance of the InfoPlant, it must be noted that the integrated water tank required the casing to be rather voluminous, which can easily distract from the natural appearance of the plant (cp. Figure 7.4). Furthermore, in our first study we found out that, although participants indicated to easily accept it as a potential addition to their living environment and that this type of feedback could also bring them to change their consumption behavior, the fact that a living plant was displaying the information seemed to be of lesser importance to the participants (cp. Hammerschmidt et al., 2015). Our interpretation of this result is that this is not only due to the technical-looking overall impression of the system, but also because the interaction is being limited to *observing* the plant as there is no need to take care of it by giving it water .

Therefore, in order to allow the behavior towards the plant remain as natural as possible and thereby to elicit feelings of care, but also to reduce the size of the casing and due to the numerous problems associated with the modality, we decided to omit this functionality for the second prototype and completely rely on mechanically influencing the posture of the plant to mimic this natural change in appearance.

ILLUMINATION

For illuminating the plant, we used a strip of 24 RGB-LEDs, which we installed on the inside of the flowerpot. This way, users normally do not see the LEDs directly but only their reflection on the illuminated plant. In order to achieve a linear mapping between input values and the perceived brightness of the LEDs, we implemented a warping function to compensate for the nonlinearity of the LEDs' brightness. Furthermore, the implemented functionality allows to easily produce advanced lighting patterns, such as blending between two colors or generating a smoothly pulsing light.

TACTILE INTERACTION

Tactile interaction was implemented using an Arduino-compatible variant (Hoby, 2012) of the capacitive sensing method described by Sato et al. (2012). Due to the limited quality of the Arduino's frequency generator, the signal-to-noise ratio of the frequency response did not allow for further classification of the position or type of contact. However, this method was sufficient to detect a user touching the leaves of the plant, allowing for basic interaction. Since we have found that the activation level can vary greatly depending on how well the plant has been watered, touch detection for the second prototype was based on an adaptive threshold.

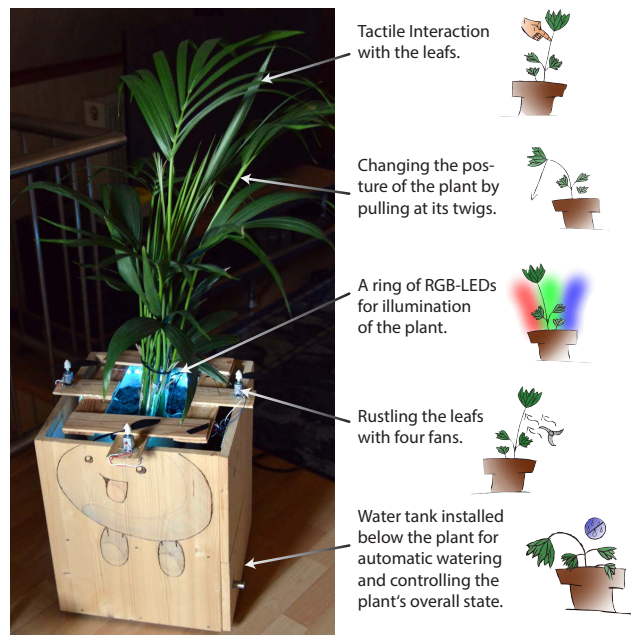


Figure 7.4.: First prototype of the InfoPlant, installed in a living environment. While the spacious construction of the casing allowed for almost all components of the original concept to be implemented in one place, it also makes the overall system look a little bulky.

LEAF RUSTLING

In the first prototype, we implemented this modality by installing four fans at the sides of the casing, which could be activated either individually or all together, and with different levels of intensity. However, the noise coming from the fans as a byproduct of generating a visible rustling of the leaves was louder than anticipated and was considered as slightly distracting by some of the participants in our first study. Furthermore, the fans at the side of the casing added to the overall impression of the InfoPlant being more of a technical system than a “natural” display, which is why we refrained from reimplementing this modality in the second prototype.

POSTURE OF THE PLANT

Changing the plant’s posture was achieved by attaching nylon threads to different twigs of the plant and connecting them to servos fastened to the inside of the casing. In the second prototype, we then moved the servos below the flowerpot in order to reduce the size of the overall system. Furthermore, to allow for an easy removal of the plant itself, we integrated magnetic couplings into the threads for temporarily removing the connection between the plant’s twigs and the servos.

SOUND

In comparison to implementing the other modalities, adding audio capabilities to the InfoPlant turned out to be a rather easy task: Since the Raspberry Pi that we chose to use as the central control unit is already equipped with an audio DAC, all we had to do was integrating a USB-powered loudspeaker into the casing. In order to keep the footprint of the system as small as possible, it was also installed below the flowerpot.

7.5. Household Study

Based on the second prototype described above, we conducted a study to evaluate the use of the InfoPlant for giving feedback on the overall electricity consumption as captured by a smart meter.

7.5.1. Measuring Electricity Consumption

Considering that only a few households are currently already equipped with smart meters, we decided to develop a portable solution for measuring electricity consumption that can easily be installed in any household. In order to do so, we employed a set of nine Plugwise *Circles*⁴, i.e. socket adapters that can be used to measure the consumption of connected appliances (cp. Figure 7.5). Although with such adapters, it is usually not possible to measure the *entire* consumption of a household, since for example ceiling lights are not connected via a socket, they have the advantage of allowing a more fine-grained analysis of the energy usage of individual (or small groups of) appliances.

AN ENERGY MEASURING FRAMEWORK

Since the proprietary tools provided by Plugwise do not allow to access any consumption data of the Circles directly, we have developed an energy measuring framework built on top of the Plugwise-2 library⁵, which uses the reverse-engineered Plugwise communication protocol to access the data provided by the socket adapters. Our primary goal for this framework was to enable a feedback display to obtain reliable consumption data and that it can be used in an everyday-context, i.e. is robust enough to deal with a wide range of scenarios and undesirable events, such as packet loss, power failure, or a user unexpectedly removing a socket adapter. To achieve this, the framework integrates two data sources:

- Every ten seconds, for each Circle, the short-term energy usage is polled and subsequently logged. Furthermore, this updates the current usage for each socket, which can then be used to update real-time feedback.

⁴Plugwise Circle: <https://www.plugwise.com/circle>

⁵Plugwise-2 project on github: <https://github.com/SevenW/Plugwise-2-py>



Figure 7.5.: A set of Plugwise Circles and the USB Adapter for receiving data from a PC or any other device with a USB port, such as a Raspberry Pi. Data is transmitted through a wireless ZigBee-Mesh network.⁶

- Additionally, we configured the Circles to internally log the accumulated energy consumption in ten-minutes intervals. This data can then be retrieved at any time and therefore does not rely on a constant connection to the socket adapter.

Combining these two data sources, the framework should be able to provide consumption data that reflects the actual electricity usage rather well even under suboptimal conditions. In addition to that, several summarizing properties of the data are calculated, such as the average daily consumption history as well as the variance of the energy usage, which can be used to decide when to inform the user about an unusually high consumption or to determine, in which range of deviation indicators for relative consumption should operate.

7.5.2. Feedback Design

The mapping of the measured consumption data on the InfoPlant's modalities is partly based on our experience with the eco-feedback demonstrator that we used in our first study (cp. Hammerschmidt et al., 2015). For example, we have found that a division of the plant into several sections to increase the amount of data that can be displayed seems to be somewhat difficult to distinguish, and thus we decided to synchronize the movement of all four installed servos, which also corresponds with our original idea of imitating the plant's natural change in appearance. Furthermore, we have reviewed preliminary research on goal setting in order to best support users in reducing their consumption.

⁶Wireless ZigBee network: <http://www.zigbee.org/what-is-zigbee/>

THE 4/5-FACTOR PRINCIPLE

Goal setting has long been identified as being highly beneficial to achieve behavior change (e.g. Latham & Locke, 1991). However, as the large majority of interventions in sustainable HCI are quite short, there has been little consideration about how to actually achieve a *lasting* reduction in consumption. Based on existing research on this topic, we argue that – just like when losing weight – a change in behavior should not be something to be forced in a short amount of time, but instead must occur in a slow pace, so that the danger of becoming frustrated or switching back to the old behavior is reduced (cp. Hill, 2009). This is also supported by recent work in empirical psychology, identifying a continuous and steady progress as one of the biggest motivating factors, even if it consists only of small “wins” (Amabile & Kramer, 2011). In order to support such an approach, Locke and Latham (2002) suggest to emphasize the relative changes in consumption, as those reveal the progress users make in relation to their goals. Furthermore, we have to consider the large differences in the overall energy usage of potential users, meaning that we have to effectively “meet” them at their individual current level of consumption.

Based on these considerations, we propose our 4/5-factor approach as a guiding feedback for a slow but continuous reduction in consumption. The basic idea for this approach is to always give feedback towards a short-term goal C_g that is only slightly lower than the (moving) average of the user’s consumption μ_C^{MA} , i.e.

$$C_g = v \cdot \mu_C^{MA}, \quad 0 \ll v < 1, \quad (7.1)$$

where v should usually be close to 1, e.g. $v = \frac{4}{5}$.

As a consistent reduction in consumption automatically leads to a lower moving average, which in turn influences the short-term goal, a gradual but always very moderate increase in goal difficulty is achieved. However, since a consequent adhering to the short-term goals conveyed by the feedback would eventually lead to a consumption level close to zero, the above formula must be slightly changed by implementing a given long-term goal:

$$C_g = v \cdot \mu_C^{MA} + C_{ltg} \cdot (1 - v), \quad (7.2)$$

where C_{ltg} could either be set by the user as a measure of difficulty, a fraction of the initial level of consumption, or even depend on the average electricity usage of the household’s neighborhood.

An illustration of this approach can be seen in Figure 7.6. Depending on the variability of the data, the electricity usage should first be observed for a few days in order to obtain a reliable initial estimate of the household’s consumption. Furthermore, the size of the averaging window can be adjusted to determine the desired adaptation rate,

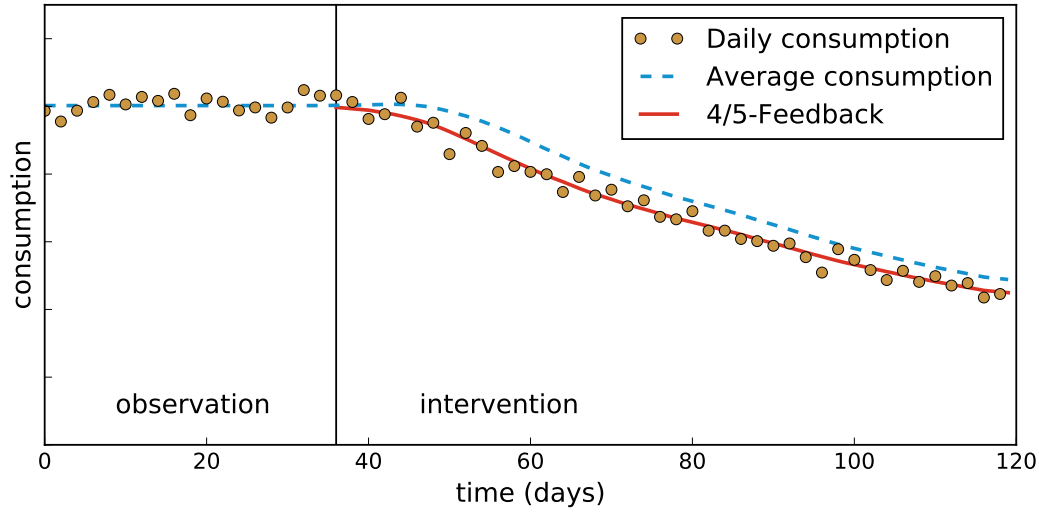


Figure 7.6.: Illustration of the 4/5-factor approach: The consumption goal is adaptively set to 4/5th of the moving average, causing a steady and mild feedback towards reaching an ultimate goal.

i.e. a smaller window translates into quicker adaption to the user's behavior and thus to a faster rate of change over iterations. In consequence, we can interpret the inverse window size as roughly corresponding to the user's ambition to change.

LONG-TERM CONSUMPTION

Based on the 4/5-factor approach, our feedback design maps several values derived from the measured consumption data on the different modalities of the InfoPlant: First, considering that every significant modification of the posture requires the servo motors to be activated, which in turn generates a low, but still potentially distracting noise, we chose to display the rather slowly-changing long-term consumption with this modality. More precisely, the twigs' current positions are determined by

$$\Theta_{\text{twigs}} = \frac{C_{24h} - \delta \left(\mu_{C_{24h}}^{\text{MA}} \right)}{\sigma_{C_{24h}}}, \text{ clipped to } -1 \leq \Theta_{\text{twigs}} \leq 1, \quad (7.3)$$

where $\Theta_{\text{twigs}} = 1.0$ indicates that the twigs are being completely relaxed and upright, and $\Theta_{\text{twigs}} = -1.0$ corresponds to their maximal deflection. C_{24h} is the household's cumulative consumption of the last 24 hours and $\mu_{C_{24h}}$ is the moving average of these values. Furthermore, C_g is the functionalized form of Equation (7.2), with a rather moderate $v = 0.95$ and a long-term goal C_{ltg} of half of the initial energy usage, and σ_C is the standard deviation of the consumption data, i.e. for calculation of Θ_t , the current consumption's deviation from the 4/5-factor short-term goal is normalized by

the overall variance in order to account for household-specific fluctuations of the daily consumption.

CURRENT CONSUMPTION

In order to additionally emphasize the *current* (relative) changes in consumption, the instantaneous electricity usage is displayed via the color of the plant's illumination. Analogous to the individual twigs being synchronized to control the overall posture of the plant, all LEDs are controlled in an identical fashion to avoid any confusion. Following the widely known traffic lights metaphor, the plant being illuminated in a red (or green) color indicates a rather high (or low) consumption, with a yellow light suggesting that energy usage is within expectable limits. Similar to the long-term consumption, this value is calculated based on previously measured data:

$$\Theta_{\text{LEDs}} = \frac{C_{\text{now}} - \delta \left(\mu_{C_t}^{\text{MA}} \right)}{\sigma_C}, \text{ clipped to } -1 \leq \Theta_{\text{LEDs}} \leq 1, \quad (7.4)$$

where C_{now} is the instantaneous consumption, σ_C is the observed standard deviation of these values, and $\mu_{C_t}^{\text{MA}}$ is the (moving) average of the consumption that was previously measured during the same time of the day. In order to prevent past consumption spikes from overly influencing the calculation, the values of μ_{C_t} are additionally smoothed by convolution with a 90-minute wide Hann window. With $\Theta_{\text{LEDs}} = -1$ mapped to red, 0 to yellow, and 1 to green, all other values lead to a smooth interpolation between these colors. Furthermore, considering that Θ_{LEDs} is approximately the first derivative of the long-term consumption, this indicator can also be interpreted in such a way that green lighting indicates the leaves moving up, while a red light implies that they are slowly being pulled down.

EXCESSIVE CONSUMPTION

A third component of informing users about their consumption is to indicate an exceptionally high electricity usage, which we defined as a consumption that is more than two standard deviations higher than the so-far observed average consumption during that time, i.e. when

$$C_{\text{now}} - \mu_{C_t} > 2 \cdot \sigma_C. \quad (7.5)$$

Being a quite unusual event⁷, it is indicated by a slowly pulsing red light in our design, which symbolizes its importance, but is still not overly intrusive.

AUDITORY FEEDBACK AND ACTIVE QUERYING

In addition to the visual modalities described above, the InfoPlant also uses auditory feedback to convey information about the user's electricity usage: If the current consumption is greater than previously observed during that time, i.e. if $\Theta_{\text{LEDs}} > 0$, a sound

⁷Assuming a normal distribution of the consumption data, the probability is less than 3%.

signifying high electricity usage is played about once per hour.⁸ Additionally, users can actively query this feedback by touching one of the leaves. In case of a comparably low current consumption, i.e. if $\Theta_{\text{LEDs}} \leq 0$, this interaction leads to a rather positively connoted sound being played back to the user.⁹

NIGHTTIME SCHEDULE

Finally, in order to keep potential distractions by the InfoPlant to a minimum during sleeping hours, we implemented a configurable and flexible nighttime schedule, which temporarily disables the LEDs and any updates of the display. Users can specify a time window, during which they usually go to sleep, where the InfoPlant continues to provide feedback, but can be disabled for the night by touching one of the plant's leaves. With no interaction from the user, this also automatically happens at the end of this period. Similarly, there is a time window in the morning, where the plant remains in its nighttime state by default, but can be activated in the same way.

7.5.3. Comparative Feedback

In order to evaluate if the InfoPlant as an ambient display would be better suited to provide feedback on energy consumption than a more conventional display that might be found in commercial solutions, we have additionally designed a screen-based visualization of consumption data, seen in Figure 7.7. It essentially consists of a bar chart displaying – in the default configuration – the hourly consumption of the past 24 hours. The rightmost column of the chart represents the current, ongoing hour, with the light blue color indicating that its final height has not yet been reached. Furthermore, the orange line displays the average consumption of the past days and thus allows users to determine if they have spent more or less electricity than usual.

INTERACTING WITH THE DISPLAY

In addition to the auto-updating, but otherwise static display of the household's consumption, users can interactively explore the data by scrolling to the left, i.e. reviewing the consumption of the previous days. The display furthermore allows to zoom out in order to get an overview over a longer period of time. For quick orientation, the days are separated by annotated vertical bars (as can be seen in Figure 7.7 on the left).

7.5.4. Study Design

As we only had one prototype of the InfoPlant and two sets of plug adapters available for the study, we were limited in the number of households that could participate in parallel. Each feedback was evaluated for one whole week to account for the variance

⁸The original sound can be found at <https://freesound.org/s/181131/>, where it is described as “Sound emitting from a large electrical box outside an industrial building”.

⁹The sound that was used here is from the R2-D2 robotic character from the widely known “Star Wars” movies.




Figure 7.7.: Bar chart of a household’s past consumption that was displayed on a laptop in the apartment of a study participant.

in consumption that could be expected over the course of one week. Furthermore, since the feedback was designed to be given in relation to the average household consumption, we first captured electricity usage for one week *without* any feedback. This is also the main reason why we chose to employ a within-subject design, where each household received feedback both by the InfoPlant and via the bar chart. In order to control for ordering effects between the two feedback conditions, a counterbalanced measures design was employed, i.e. for half of the households, feedback was provided first by the InfoPlant, whereas the other half received feedback first via the bar chart.

ACQUISITION OF THE PARTICIPANTS

In order to find suitable households, we designed a medium-sized poster with a call for participation, which we distributed at several public places, such as supermarkets, the library, a pharmacy, and some bulletin boards at the university. Additionally, we posted a similar call on the pages of a few regional groups on Facebook. To obtain a more comparable dataset, the call addressed one and two-person households only. The main requirement for participation was that the inhabitants were at home on a regular basis, i.e. they were not on holidays during the three weeks and also were at home on the weekends.

CONDUCTING THE EXPERIMENT

Over the duration of the experiment, we visited the participating households four times: At the beginning, the participants were given a short introduction explaining



Figure 7.8.: Pictures of the InfoPlant being installed in different study households.

the overall course of the experiment and signed a written consent agreeing with the gathering and evaluation of their consumption data. Subsequently, with the help of the inhabitants, we distributed the plug adapters to capture as much electricity consumption of the household as possible and annotated both the location of each plug and the appliances it measured. Finally, the laptop for continuously recording the data was installed at a central place of the apartment so that for each Circle, the necessary wireless connection could be established without any problems.¹⁰ Information about the study was left with the participants so that they could easily explain the experiment to other inhabitants or visitors. They were also instructed to just use the appliances whose consumption were measured to the same extent as they would usually do.

In the following two weeks, the setup was kept exactly the same, except for the added feedback provided either by the InfoPlant or the bar chart, displayed on a dedicated laptop. At the beginning of each week, the respective display was set up in the house-

¹⁰The Python script managing the data recording was configured to only start if all Circles could be reached during initialization in order to prevent a situation where the data from one or more Circles was not captured at all.

hold at a visible location, convenient to the inhabitants, while still allowing the system to reach all installed Circles. In order to support a more flexible placement, we brought a small shelf to put the laptop or the InfoPlant on, which also allowed the feedback to be perceived at eye level. Subsequently, the respective display was explained and interactively demonstrated to the inhabitants, i.e. we showed them how a specific consumption would look like and how they could interact with the display. Additionally, information about the feedback was left with the participants for them to refer to during the week. Also, both an email address and a cellphone number were provided in case of questions or unforeseen events. In order to prevent participants from feeling obligated to reduce their consumption when confronted with the feedback, just to adhere to the study's guidelines, we explicitly told them that it was not their *task* to do so, but that they should just react to it as they would do if a smart meter was installed in their apartment that provided them with feedback on their consumption. Finally, after the three weeks, all hardware was collected for use in the next household.

QUESTIONNAIRES

During the different phases of the experiment, there were four occasions, where each participant had to answer a specific set of questions:

- At the beginning, participants provided some general information about themselves, such as age and occupation, but also about perceptual deficiencies or if they are easily annoyed by visual or auditory stimuli.
- Then, after the first week, participants would indicate, if, despite being instructed otherwise, they had already tried to reduce their electricity usage or if they thought that the knowledge about their consumption being captured might have led to an unintentional reduction of consumption.
- In the second and third week, the questionnaire consisted of specific questions about the respective feedback, e.g. if it was understandable, helpful, or distracting.
- Finally, at the end of the study, the participants answered several questions directly comparing the InfoPlant with the graphical feedback as well as some open-ended questions and suggestions about feedback in general and other ways to support electricity consumption.

In addition to that, on each of these four occasions, we included questions about the participants' interest in their electricity consumption and their opinion on reducing it in order to see if the feedback might have influenced their general attitude towards the topic. The majority of questions were posed as statements that the participants would indicate their agreement to on a 7-point Likert-type scale ranging from "Not at all" (-3) to "Fully agree" (3).

7.5.5. Goals and Hypotheses

The primary goal of this study was to examine how and if the InfoPlant, being installed in a household as an ambient feedback display, is able to reduce the electricity consumption of the inhabitants and to compare it to a more conventionally designed, graphical feedback display. In addition to that, we also see the study as an opportunity for a more exploratory evaluation of the gathered information.

- First, we assume that the feedback is effective insofar as the study participants can estimate their *relative* consumption, i.e. if the electricity usage has increased or decreased compared to previous consumption (H1).
- Additionally – despite the ongoing controversy if feedback *is* able to lower energy usage – we expect to see a measurable reduction in consumption for both feedback conditions (H2), with a larger change for the InfoPlant (H3).
- Combining consumption data with the results of the questionnaires, we furthermore want to evaluate if there are any additional factors facilitating a reduction in consumption.
- Concerning how the feedback displays are perceived by the users, we expect the InfoPlant to be experienced as less obtrusive than the graphical feedback (H4).
- Finally, since we are capturing quite detailed consumption data, we want to examine how, where, and when electricity is consumed in the households.

7.5.6. Results

During its total runtime of 16 weeks, eight households consisting of either one or two persons were participating in our study.

PARTICIPANTS

The 13 inhabitants (5 male and 8 female) were between 20 and 61 years old. In the initial questionnaire (cp. Figure 7.9), all of them indicated that they are interested in a system that helps them to save energy (median $\mu_{\frac{1}{2}} = 2$). Interestingly, this was not entirely consistent with their interest in how much energy they use ($\mu_{\frac{1}{2}} = 1$), which we would interpret as the wish to consume less energy without one's direct involvement. In fact, when asked about other ways to reduce energy consumption, VP1's main suggestion was to simply use more efficient appliances. Furthermore, the vast majority indicated that they generally try to act in a sustainable way and not to waste energy ($\mu_{\frac{1}{2}} = 2$), with their motivation to cut energy consumption being environmental aspects ($\mu_{\frac{1}{2}} = 1$), and even more to save money ($\mu_{\frac{1}{2}} = 2$). However, when asked more specifically if they pay attention to how much energy they use or if they try to use not too much electricity, they did not agree that easily ($\mu_{\frac{1}{2}} = -1$ and $\mu_{\frac{1}{2}} = 0$, respectively), i.e. while there certainly exists the abstract goal to reduce one's energy consumption,




Figure 7.9.: Summary of the questionnaire that each participant answered at the beginning of the experiment. Answers could be given on a 7-point Likert-type scale indicating the level of agreement with the respective statement. In this chart, the width of each bar corresponds to how many people responded with the respective agreement, whereby the outer, more deeply colored bars represent a stronger reaction, while the inner, more lightly colored bars represent a weaker tendency. Only the responses that were not “neutral” are displayed.

there seems to be a lack of motivation to actually do so. Moreover, although most of the participants were quite sure that they do not consume more energy than the average person ($\mu_{\frac{1}{2}} = -1$), they did not even know how much electricity *they* use ($\mu_{\frac{1}{2}} = -2$).

CONSUMPTION DATA

As could be expected, the variance of the measured consumption data is quite high: With an average of 81.8W, the mean standard deviation of the hourly consumption for each household is more than 84W. A large part of this variance can be attributed to the different amounts of electricity used over the day (see Figure 7.10). On average, the lowest consumption (around 40W) can be observed in the night between 3am and 6am; the highest (more than 140W) in the evening between 8pm and 9pm. Although being lowest when seen in the course of the day, we were surprised to see such high consumption values during nighttime. Closer inspection of the data revealed four different types of electricity usage contributing to nighttime consumption:

- Appliances that need a constant energy supply, such as a fridge.
- Devices that are switched off but consume a certain amount of standby power, e.g. a television set or hi-fi equipment.
- Consumption that is deferred to the nighttime, for example a laptop being used in the evening that is charged in the night.
- Rare occurrences of “active” consumption, such as using the water boiler for a very late cup of tea.

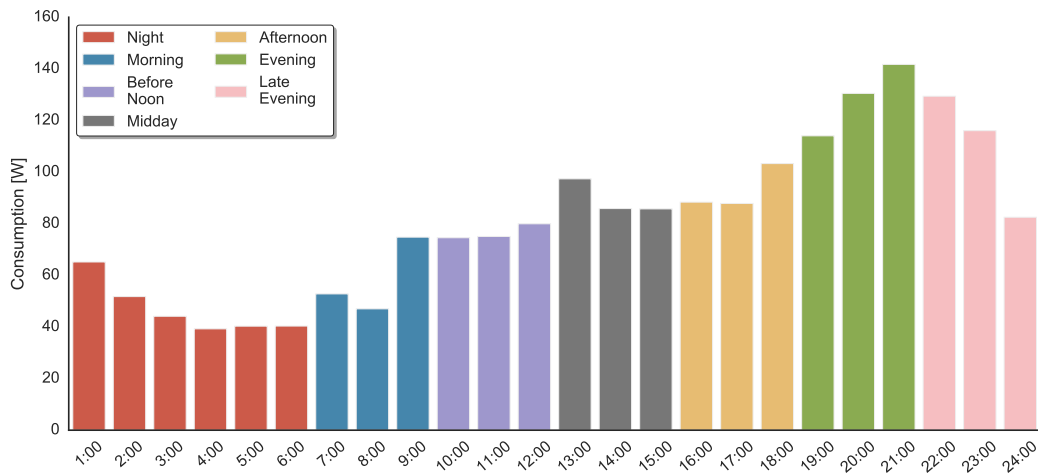


Figure 7.10.: Accumulated hourly consumption over the course of the day, averaged over all households. Although the participants each have their individual usage patterns (with differing amounts of variance), the increasing consumption towards the evening as well as the lowest consumption being measured during the night can be observed for the majority of households.

According to our data, the first two items constitute the main part of the nightly consumption, with the second one being an easy opportunity to save energy that still seems to be only seldom exploited, despite being publicly known for quite some time.

A second reason for the large variance in the consumption data are the differences between the individual households: Although the selection was restricted to single and two-person households only, there is a huge difference (almost 160W) between the lowest and highest average usage (also cp. Figure 7.11); the standard deviation is 58W. As expected, average consumption is higher for two (96W) than for one-person households (58W). However, the variance for the two-person households is surprisingly high ($\sigma = 70W$, as opposed to $\sigma = 16W$ for the single ones). In fact, the two households with the lowest consumption consist of two persons.

Finally, there is a certain variance when considering the accumulated *daily* consumption for each household, although with an average standard deviation of around 20W, it is relatively low in comparison. Moreover, while for some households, there are certain days of the week, where electricity usage is higher than on others, we could not find any systematic effects. For example, comparing the energy usage on weekends with the one during the week, there are only two households with a significant difference: one with a higher and the other one with a much lower consumption on the weekends.

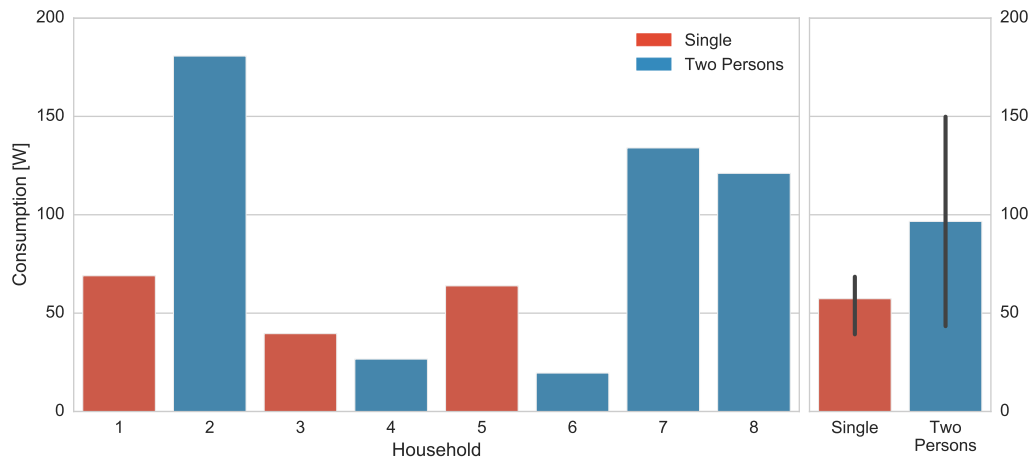


Figure 7.11.: Average electricity usage of each household and a comparison between single and two-person households.

COMPARISON OF OVERALL CONSUMPTION

To determine if the feedback had any effect on the overall consumption, we compared the electricity usage during the first week (i.e. without any feedback) with those when either the InfoPlant or the graphical feedback were installed in the respective household. Surprisingly, there are basically no differences between either of the conditions (see Figure 7.12a), and although with the InfoPlant as feedback display, the average consumption was marginally lower (80.6W) than for the baseline and graphical feedback (82.2W and 83.2W, respectively), those differences are not significant ($p > 0.5$ for both comparisons). Consequently, both H2 and H3 cannot be confirmed.

Considering Only Consumption Within Sight of the Feedback. As one part of the questionnaire as well as during each visit, the participants were invited to make comments and suggestions regarding their experience with the feedback. One issue that was mentioned several times was the *visibility* of the respective display:

- VP3, for example, commented on the InfoPlant that most of the time during his usual daily activities, it was simply outside of his field of view, which prevented him from keeping aware of the apartment's electricity consumption.
- In another household, VP4.1 suggested to give additional acoustic feedback that can be perceived in the whole apartment (and not only in the vicinity of the respective display).

Based on these comments, we evaluated how the consumption changed when feedback was installed specifically for the Circles that were located within sight of the feedback display in the respective households. Furthermore, in order to check for the change in the actual *usage* of the respective appliances, we analyzed the consumption relative

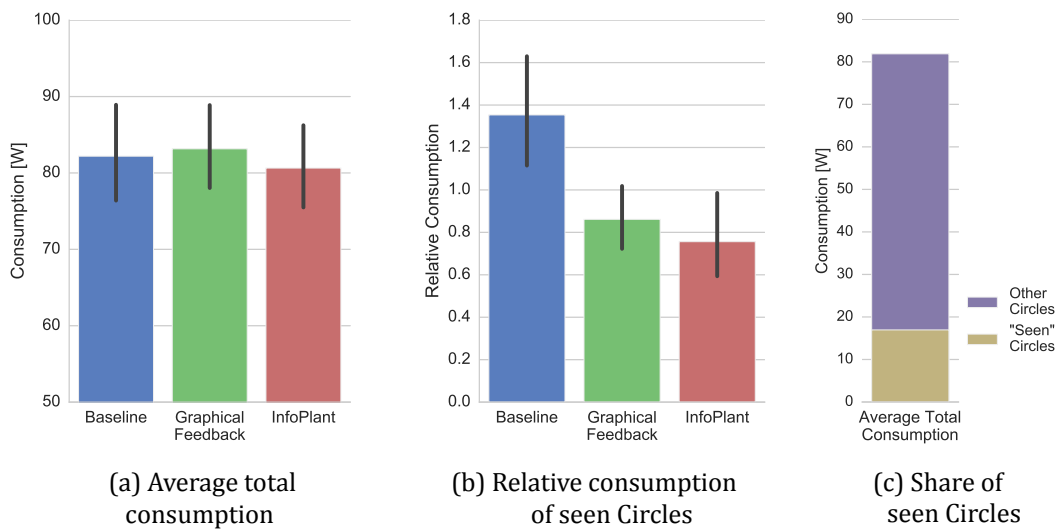


Figure 7.12.: **(a)** Accumulated consumption of all measured socket adapters (i.e. considering all Plugwise Circles), averaged over all test households. **(b)** Differences in the relative (i.e. normalized) consumption of the Circles that could be seen from the location of the feedback display. **(c)** Accumulated absolute consumption of those “seen” Circles in comparison to the average total consumption.

to its overall average. As we can see in Figure 7.12b, there is a considerable reduction in the use of the devices connected to those “seen” Circles, when a feedback display is used. For the InfoPlant, the decrease is around 60%; for the graphical feedback, it is still almost 50%. For both feedback conditions, the reduction is significant ($p < 0.01$). The difference (of about 10.5%) between InfoPlant and graphical feedback, however, is not ($p = 0.42$). Furthermore, the *absolute* reduction only amounts to an average of approximately 8-10W, since the consumption captured by the plugs within sight of the feedback is rather low (around 17W), and their share of the total consumption is only 20.7% (cp. Figure 7.12c), which also explains why there is no significant reduction in the overall consumption (Figure 7.12a).

OCCURRENCE OF CONSUMPTION SPIKES

As the displays were designed towards providing feedback in comparison to the past consumption and thus to emphasize unusually high electricity usage, we also analyzed if there were any changes in the frequency and duration of such consumption spikes when users were provided with corresponding feedback. Indeed, when considering only the devices in the vicinity of the feedback display, we can see a considerable reduction in the average daily duration (158 minutes for the InfoPlant and 165 minutes for the graphical feedback, compared to 227 minutes for the baseline) and there seem to be slightly less extreme durations, especially for the InfoPlant (cp. Figure 7.13a), both hint-

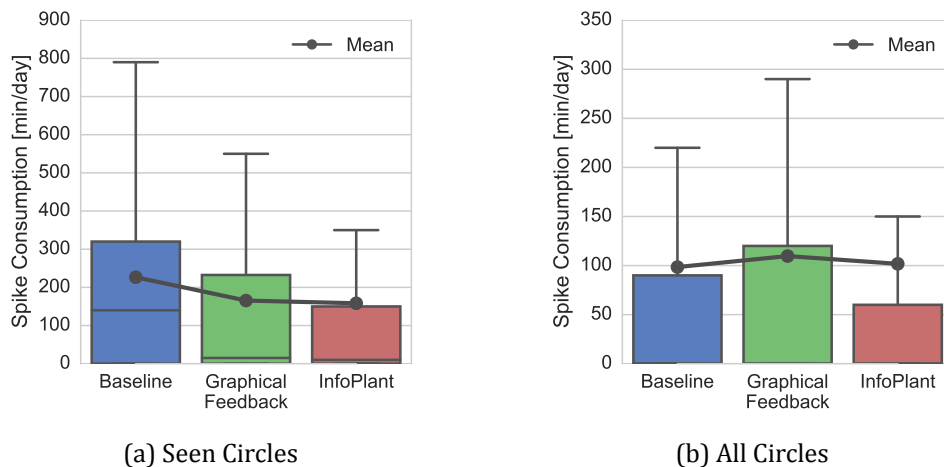


Figure 7.13.: Boxplot of the daily duration of excessive (spike) consumption for each observed power socket. A Circle’s consumption is considered excessive, if the captured devices use more than half of the average total consumption of the whole household. Not included are Circles without any spikes at all. **(a)** shows spike consumption for Circles in the vicinity of the respective feedback display, whereas **(b)** includes all Circles.

ing at an influencing effect of the feedback. However, the variance of the data is quite high (mean standard deviation $\bar{\sigma} = 270$ minutes) and there are no significant differences ($p = 0.11$ when comparing the InfoPlant with the baseline). Furthermore, when taking into account *all* captured appliances, there is basically no difference between the conditions (Figure 7.13b), similar to the overall consumption (cp. Figure 7.12a).

MOTIVATION AND INTEREST

Although a consistent reduction in electricity usage could only be found for a low amount of devices in the vicinity of the respective feedback display (cp. Figure 7.12), there are also households that reduced their *overall* consumption and thus a central question is if there are any preconditions or benefitting factors that enable users to profit from the feedback in a more effective way. When analyzing the change in overall consumption in relation to the responses of questions answered by the participants at the beginning of the experiment (cp. Figure 7.9), we have found a strong correlation between a reduction of electricity usage and answers indicating a motivation to actively deal with one’s energy consumption (Figure 7.14a). For example, participants who said to have a strong interest in their consumption also managed to reduce it when assisted by either of the feedback displays ($\tau = -0.45$; $p = 0.031$). Similarly, whether participants already knew about how much electricity they use per month – something that requires a certain effort and engagement to learn about – was a good predictor for a reduction in consumption ($\tau = -0.41$; $p = 0.048$). Interestingly, this correlation can also be found when considering the relative consumption of the devices in the vicinity

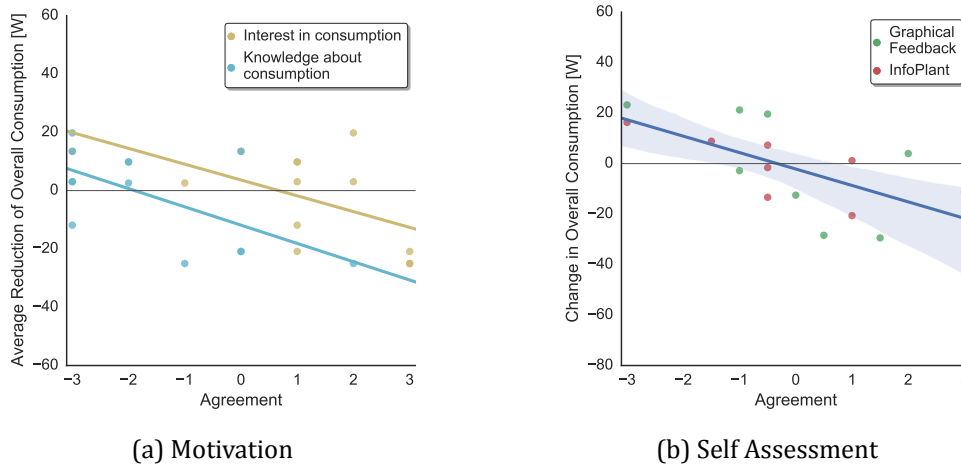


Figure 7.14.: Correlation between the average change in overall consumption and the agreement to statements of the questionnaire, as indicated by a linear regression. **(a)** Correlation with two initial questions (i.e. asked *before* the experiment) indicating a motivation to reduce energy usage: the interest in one’s own consumption and, similarly, already knowing about the household’s total monthly electricity usage. **(b)** Correlation with the assessment that one’s consumption has been decreased (or conversely: increased) during the respective week. The chart additionally shows the 95% confidence interval of the regression.

of the feedback display ($\tau = -0.6$; $p = 0.01$), i.e. the participants who reduced the usage of these devices were mostly those who initially indicated to be interested in their consumption.

EFFECTIVENESS OF THE FEEDBACK

Although to us the long-term objective of a feedback display is to achieve a reduction of energy usage, an intermediate goal is obviously to enable users to estimate their consumption. As we can see in Figure 7.14b, participants were quite able to correctly tell if their electricity usage had decreased or increased for a specific week ($\tau = -0.49$; $p = 0.008$), which is exactly the information we wanted to convey and which supports our hypothesis H1. However, combined with the previously discussed results, this finding also hints at the inadequacy of feedback alone – at least in this form – to induce a substantial reduction in consumption.

PERCEPTION AND COMPARISON OF THE FEEDBACK DISPLAYS

In the preceding analysis, we have seen that the measurable effects of the graphical feedback and the InfoPlant are very similar; what we have not covered yet is how they were perceived by the users. As we can see in Figure 7.15, the designs were generally quite well received. The only aspect the participants could not fully agree to was the




Figure 7.15.: Summary of the questionnaire about how graphical feedback and InfoPlant have been perceived by the participants. The answers were given at the end of the experiment (i.e. after they had experienced both feedback devices) on a 7-point Likert-type scale to assess one particular characteristic of the respective display. Only the responses that were not “neutral” are displayed.

motivational effect of the feedback displays ($\mu_{\frac{1}{2}} = 0$ and $\mu_{\frac{1}{2}} = 1$ for the InfoPlant and the graphical feedback), which also explains the apparent need for an intrinsic motivation to reduce one’s consumption (cp. Figure 7.14a).

While in general, both designs were assessed quite similarly, in terms of helpfulness, the graphical feedback ($\mu_{\frac{1}{2}} = 2$) was rated slightly better than the InfoPlant ($\mu_{\frac{1}{2}} = 1$). The only significant difference, however, can be found for the answers to how well the participants were able to estimate their consumption ($\mu_{\frac{1}{2}} = 1$ for the InfoPlant, as opposed to $\mu_{\frac{1}{2}} = 2$ for the graphical feedback; $p = 0.016$). This result is most likely due to the added information of the *absolute* electricity usage for the graphical feedback (see Figure 7.7), which is not displayed by the InfoPlant, as well as the ability to review past usage (cp. Section 7.5.3). The original goal of conveying the *change* in consumption, however, was achieved by both displays (see Figure 7.14b).

All in all, we can therefore conclude that the two designs performed quite similarly. On the other hand, this also means that there does not seem to be any significant advantage of using the InfoPlant over a conventional feedback display, either. For example, participants did not feel distracted from their usual activities by neither of the displays ($\mu_{\frac{1}{2}} = -2$), nor did they indicate to need to focus on either of the feedback displays in order to understand them (both $\mu_{\frac{1}{2}} = -2$), which also means that our hypothesis H4 cannot be confirmed.

AFFECTIVE RESPONSE

In part, these results can probably be attributed to the limited effect of the plant in terms of eliciting an emotional response from the participants: According to the ques-

tionnaire, they neither had the feeling that the plant providing the information was making the feedback more important to them ($\mu_{\frac{1}{2}} = -1$), nor did they indicate that it appealed to them on an emotional level ($\mu_{\frac{1}{2}} = -1$, which is the same result as for the graphical feedback).

On the other hand, VP4.1 spontaneously commented on the plant being replaced by the graphical feedback that it was almost sad that it had to go, because it had become almost like a pet to them, which you had to stroke every night in order to send it to bed. Interestingly, this household also had the second-largest relative reduction in overall electricity usage when receiving feedback from the InfoPlant. Apparently, however, not all participants felt the same, and the questionnaire results also do not suggest a relation between an appreciation of the plant and a reduction in electricity use.

7.6. Discussion and Conclusions

All in all, the construction of the InfoPlant and our experience with the two prototypes as well as the conducted study provide us with some valuable insights on “living” interfaces, ambient displays in general, and supporting users with feedback on their energy consumption behavior.

COMPLEXITY AND VARIABILITY OF THE CONTEXT

With regards to RQ3, the results of the study suggest that the large variety of activities taking place in a living environment clearly impairs the functioning of the InfoPlant as an ambient display: Depending on what the user is currently doing, the provided information might go unnoticed, or – although less often the case for the InfoPlant – it might distract or even annoy the user.

Based on our observations, this is not an issue of this particular design, but seems to be more of a fundamental problem of ambient displays as they have been conceived up until now. VP3 has rather pointedly described this issue by saying that during most of the activities carried out throughout the day, it was impossible for him to keep the InfoPlant in sight and thus to properly keep aware of the electricity consumption. With regards to ambient displays research as a whole, there are several approaches to dealing with this issue:

- First, we could limit the application of those interfaces to use cases, where no continuous feedback is necessary and the user requires only infrequent updates of information. However, this would reduce a peripheral display mostly to a (readily available) query interface and although this would alleviate the problems discussed above, it would also limit the appeal of using such a display.

- A second approach, already hinted at in preliminary research, is to attempt to deduce the user's attention level in order to dynamically adjust the feedback's intrusiveness. While we think that this approach will become more feasible in the future as the integration of sensors becomes ubiquitous, current attempts are still restricted in terms of the environments and contexts they can support (e.g. Palinko, Kun, Shyrovkov, & Heeman, 2010). Furthermore, considering the predominant design of ambient displays as stationary entities with a visual representation of information and given the users' limited field of view, this approach alone would be insufficient for dealing with the variety of activities that can take place in a living environment.
- In consequence, we propose that the design and development of a peripheral display should in most cases be explicitly targeted towards a rather specific usage scenario to be better able to anticipate the users' level of attention and field of view. While one could argue that this approach also limits the design space of those displays, it still retains the core characteristics of these interfaces, i.e. allowing users to keep aware of (secondary) information.
- Finally, we think that a more systematic use of the auditory modality will make it possible to overcome the problem of the users' limited field of view and changing direction of gaze and to provide information in more varying contexts (also cp. RQ2). This may either be in form of a purely auditory display or as a multimodal interface, where sounds might also be used for moderating attention, i.e. making aware of (visually presented) information.

UNOBTRUSIVENESS

Regarding RQ4, we have seen that the InfoPlant – despite its appropriateness for the living environment it was installed in – was not rated any better in terms of unobtrusiveness than the graphical feedback. This suggests that the perceived intrusiveness of ambient displays depends more on the way information is presented than on their overall appearance. Furthermore, the study has shown that the (admittedly rather sparse) use of sound has not been perceived as particularly intrusive by the participants, but has rather led them to suggest more extensive use of auditory feedback, which we see as a first indication that using sound in a peripheral display might indeed be a feasible alternative to a (purely) visual interface (RQ2).

SUFFICIENCY OF FEEDBACK

With regards to RQ6, the conducted study on the one hand suggests that feedback on energy usage can indeed help to reduce one's consumption, but on the other hand hints at the inadequacy of feedback alone to induce the necessary change in behavior. Further analysis of the questionnaire results has revealed that a preexisting intrinsic motivation seems to be a precondition for the feedback to work adequately. These findings are in line with a recent study conducted by Buchanan, Russo, and Anderson (2015), who conclude that the success of giving feedback “depends entirely on user engage-

ment". While this result may appear somewhat obvious, we do think that it is worth emphasizing, given the widespread assumption that providing feedback on energy usage more or less automatically leads to a lower consumption (also cp. Section 6.1.4). Consequently, we argue that research on eco-feedback should either

1. target users who can be expected to already have intrinsic motivation, i.e. to focus more on the aspect of *supporting* them, as opposed to trying to persuade them, or
2. concentrate on *engaging* users, for example through methods of gamification.

AFFECTIVE RESPONSE

Although we deliberately chose to use a plant as the basis for this ambient display to engage users on an emotional level, we could not find any evidence in the conducted study that the feedback provided by the InfoPlant has been perceived correspondingly by most of the participants. Consequently, it may not be the question if the use of affective cues is beneficial for a persuasive ambient display (RQ8), but more if the limited information capacity and interaction possibilities of such an interface allow to incorporate persuasive affective cues that can actually induce a behavior change. After all, while VP4.1 did seem to become emotionally attached to the InfoPlant to a certain degree, this was less due to its appearance, but more because of the interaction with it (cp. RQ7).

8.

InfoDrops: The Sonic Shower

8.1. Introduction

Although seemingly abundant for many of us, water is a precious resource, and two thirds of the global population (around four billion people) live under conditions of severe water scarcity for at least one month in the year (Mekonnen & Hoekstra, 2016). At the same time, taking a shower in the western world usually uses up more water than the typical person living in a developing country uses in a whole day (Watkins et al., 2009), which makes a reduction of water consumption both desirable and feasible. One might argue that in some regions, as for example in Germany, reducing local water usage is unnecessary due to frequent rainfall, or even detrimental to the environment as the oversized canal system has to be flushed with *additional* freshwater in order to prevent clogged sewer lines (Leist, 2001). However, we have to keep in mind that showering not only uses lots of water, but also a significant amount of energy to heat it: In the United States, domestic water heating accounts for between 15 and 25 percent of the energy consumed in homes (U.S. Department of Energy, 2001).

Similar to the InfoPlant discussed in the previous chapter, we have therefore developed the Sonic Shower to provide feedback on energy (and water) consumption to users taking a shower in order to support them in a reduction of their expenditure. However, despite the common goal of these two systems, there are two important differences:

- First, the environment the Sonic Shower is installed in as well as the activity that the system supports is more clearly defined than for the InfoPlant. As discussed in Section 7.6, we consider this an important precondition for an ambient display to be perceived by users in a predictable and consistent manner (also cp. RQ3).

This chapter is, in parts, based on:

Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2013). InfoDrops: Sonification for Enhanced Awareness of Resource Consumption in the Shower. In *Proceedings of the 19th International Conference on Auditory Display (ICAD-2013)*. Lodz, Poland.

- Second, as the name implies, the Sonic Shower is a purely auditory display: As previously argued, we think that the use of sounds can alleviate some of the problems found in visual ambient displays, and we also consider the bathroom environment as particularly suited for using the auditory modality, as we will discuss in the following.

8.1.1. Challenges of Interface Design in a Bathroom Context

Although certainly attractive from a functional point of view, the bathroom environment has been somewhat neglected in research on interface design and smart homes. Consequently, we have recently surveyed the distinctive challenges for user interfaces to be used in a bathroom context (Leichsenring, Yang, Hammerschmidt, & Hermann, 2016).

LEGIBILITY OF VISUAL DISPLAYS

The increased humidity, for example, can not only lead to problematic condensation on capacitive surfaces, but also to cause screens to become illegible, similar to a mirror becoming fogged up. The danger of splashes of water furthermore requires special protection of any electronics and tangible interfaces installed in a bathroom environment. Additionally, we cannot expect users to be in full possession of their vision, either – especially when taking a shower: Wearers of glasses will very likely leave them outside, and the water itself, possibly also in form of steam, can very easily interfere with a problem-free viewing of any visualization. Finally, putting the head under the water, as it is frequently done when taking a shower, automatically forces most users to close their eyes, which renders a visual display effectively useless for a significant amount of time.

These issues obviously lead to a rather limited design space of those interfaces, which is also reflected in the simplicity of existing visual eco-visualizations (also cp. Section 8.2).

CONFINED SPACE AND LIMITED ATTENTION

Another problematic aspect of most bathrooms is the rather confined space, which effectively restricts the number (and size) of any additionally introduced components. More specifically, for a visual display to be installed into a shower, it is rather difficult to find a place, where it is conveniently seen at all times. In a study evaluating such a display being installed near the shower head, a participant noted that “it’s very hard to see, you never look back behind you [while showering]” (Kuznetsov & Paulos, 2010).

Finally, the attention of potential users can be expected to be mostly focused on the main task of taking a shower, which makes it difficult for them to simultaneously attend to a visual display. This issue is perhaps best described by a participant of the same study commenting on a (numeric) visual display that it “seemed to jump to a high number every time she looked at it” (also cp. Figure 8.1a).

ACOUSTIC ENVIRONMENT

On the other hand, the distinctive characteristics of the acoustic environment, i.e. the prevalence of sound-reflecting surfaces and the potential masking effect of the noise of running water, also pose a challenge for auditory displays to be integrated into a bathroom context. However, we think that the benefits of employing the auditory modality clearly outweigh the potential difficulties of the acoustic environment.

8.2. Related Work

AMBIENT DISPLAYS IN THE SHOWER

To our best knowledge, Kappel and Grechenig (2009) were the first to develop an ambient display to support water conservation in the shower. Exploiting the metaphor of the drain being closed, thereby leading to a rising level of water, their display shows the consumption with an array of LEDs that are vertically assembled on a stick, with one additional LED being lightened up for every five liters of water used during a shower (cp. Figure 8.1c). Being conceived as an ambient display, it was kept intentionally simple to allow for an easy interpretation of the presented information. However, this also means that it does not account for the individual shower habits of users, e.g. by adapting the displayable range of consumption to match their usual water usage.

Slightly more advanced, Kuznetsov and Paulos (2010) used a previously measured average baseline consumption as a (static) reference point in their display design. After an initial prototype employing two LED bar graphs to indicate both individual and the day's cumulative consumption had been perceived as too confusing in a pilot study, the authors ultimately decided to simplify their design by utilizing only a single colored light indicating the current shower's water usage by leveraging a traffic light metaphor, i.e. showing a green color to indicate that the current water consumption is below average, yellow when it is above, and red if it reaches 150% (also see Figure 8.1a). In a subsequent study, the display was evaluated and compared to a numeric display showing both current and average water usage. The authors conclude that "the numeric display was ultimately less liked and less effective, despite participants' initial preference for the numerical modality".

CLOSELY-RELATED ENVIRONMENTS AND INTERFACES

In a similar context, Arroyo, Bonanni, and Selker (2005) explored the use of novel interfaces aimed at influencing the behavior and use of water in the domain of the sink. This includes a design that illuminates the water from a faucet based on its temperature, and one that displays the current and average household water consumption by means of two illuminated bar graphs. Notably, this installation also provides auditory feedback by playing chime sounds or prerecorded voice feedback when the tap has been closed.

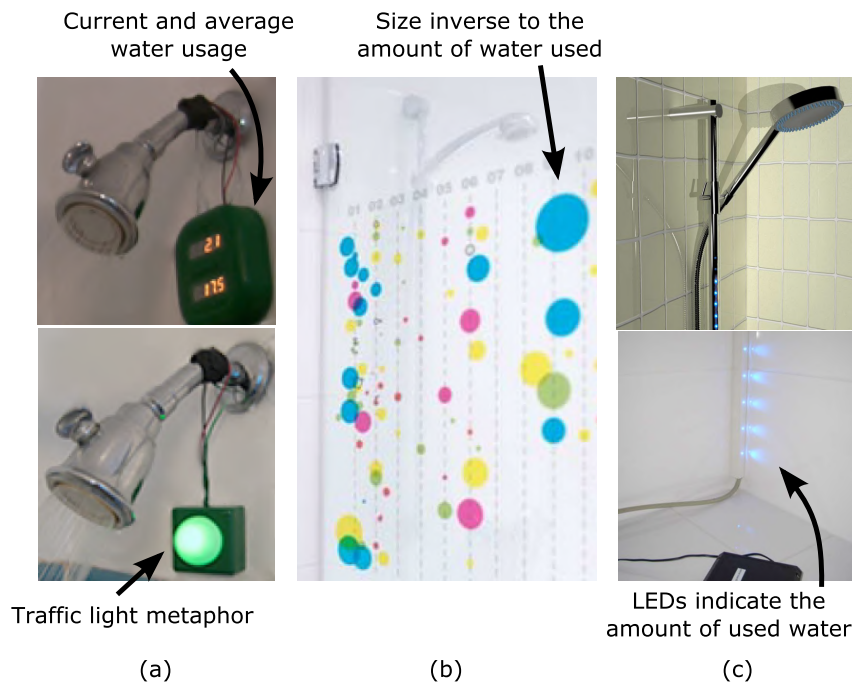


Figure 8.1.: Pictures of related work concerned with reducing water usage in the shower. **(a)** shows a numeric as well as an ambient display incorporating a traffic-light metaphor (Kuznetsov & Paulos, 2010). **(b)** illustrates the *Shower Calendar* developed by Laschke, Hassenzahl, Diefenbach, and Tippkämper (2011): Each dot represents the leftovers of using 60 liters of water. Finally, a design concept and prototype developed by Kappel and Grechenig (2009) can be seen in **(c)**: For every five liters consumed, an additional LED is switched on.

Pursuing a slightly different approach than the previously discussed projects, Laschke, Hassenzahl, Diefenbach, and Tippkämper (2011) introduced a “shower calendar”, displayed on a screen in the vicinity of the cabin, that shows the water usage per shower during a whole year (cp. Figure 8.1b). For each shower taken, a dot is displayed on the calendar, whose size corresponds to the leftovers of using 60 liters, i.e. the larger the dot, the less water has been consumed during the corresponding shower episode. For use in a family-context, the dots have different colors in order to distinguish between the consumptions of each of the family members. Different to the projects discussed above, this display gives feedback only after each shower und thus cannot give any guidance during it. Furthermore, although the overall concept of the shower calendar is rather simple and should be easily understandable, it is obviously not an ambient display, as it requires focused attention to comprehend the rather large amount of presented information on the screen.




Figure 8.2.: Illustration of a typical showering pattern; here depicted as function of time for a shower with two phases. The vertical axis depicts the instantaneous and integrated energy consumption (in arb. units).

In summary, we can see that, while preliminary research offers a range of design approaches that can serve as a basis for further exploration of this topic, existing displays almost exclusively rely on visual means to provide feedback to users, and the important aspect of energy consumption has been completely neglected so far.

8.3. The InfoDrops System

Although seemingly a mundane and simple routine, it is worth to first have a more detailed look at the process of showering itself. From a perspective of energy usage, a shower episode consists of several phases of continuous consumption, as hot water is running through the shower head. Normally, this is either (a) a single phase, during which the water is just kept running the whole time, (b) two showering phases, between which the water is switched off for soaping, or (c) three phases, if, for example, the hair is being washed individually.

CONTINUOUS FEEDBACK

Assuming that users set a goal to only consume (on average) less than a specific amount of energy and water, they should obviously receive feedback, once this amount is exceeded. However, it might well be the case that it is impossible for them to just stop the water at this point, as they might still need to rinse off. Therefore, a feedback system should not only (a) clearly indicate the recommended ending of one phase, but also (b) exhibit enough structure, and thereby orientation, *during* each phase, so that users can easily estimate their current consumption. Practically, it needs to be decided what display quantization will be both effective in stimulating the perception of change and fine-grained enough to give a sense of the absolute value.

Concerning the second requirement, a visual numerical display might seem to be the best representation, as it can show very detailed information. However, this would not only make it necessary to keep the display in sight for most of the time, which is both impractical and inconvenient, as we have seen earlier, and has led preliminary work to converge towards rather simple visual display designs (cp. Section 8.2). It would also require users to become an expert in interpreting raw data values concerning their consumption, which is a task that many of us are not familiar with (cp. Strengers, 2011). Instead, the system should provide this information in a way that is intuitively understandable to someone who has never experienced the feedback before.

USING BLENDED SONIFICATION

Furthermore, as we have argued in Section 8.1.1, we think that especially for the scenario of providing feedback in the shower, there are significant and considerable benefits in employing the auditory modality and taking a sonification approach, despite the challenges presented by the acoustic environment. In fact, we propose that the noise of the running water and the environmental sound that is inevitably produced by the water falling onto the bathtub can actually be used to our advantage and as a fundamental basis for the sound synthesis. We argue that making use of such a blended sonification (cp. Section 5.2) is, in this particular case, superior to an auditory display using conventional methods of sonification due to the following reasons:

- First, a sound synthesis that disregards the already existing environmental sounds would obviously lead to an interference between the sound that is produced by the auditory display and the noise of the water.
- Furthermore, we have to keep in mind that for most people taking a shower is something relaxing (cp. Berker, 2013). Consequently, we also want to minimize the interference of the auditory display on the act of showering itself. By using a blended sonification, we should be able achieve an unobtrusiveness that is adequate for an auditory display in such an environment (cp. RQ4).
- Finally, by incorporating the existing soundscape, we avoid to discard the information that it already conveys. In this case, the sounds can already give users a sense of the amount of water that is falling onto the bathtub and thus which is currently consumed for showering.

8.3.1. Implementation

Figure 8.3 shows the overall high-level hardware setup that we have used in our experiments: As the user is taking a shower, he or she will inevitably produce a sound of water falling onto the bathtub – a sound that most of us know and are accustomed to. With the bathtub as a resonating body, this sound is not only audible for the person in




Figure 8.3.: Hardware setup of our auditory display: Based on the current water flow and temperature data, the sound that is picked up by the contact microphone is processed and, through a speaker, directly fed back to the user.

the shower, but can also be captured by a contact microphone¹ and thereby be used as an (additional) input for a blended sonification.

MEASUREMENT OF CONSUMPTION AND SOUND PROCESSING

For the consumption sensing, we have installed a sensor that measures both the instantaneous water flow and the temperature of the water. More specifically, we are using a *Resol Grundfos VFS 1-12 l²* (also cp. Figure 8.4), which outputs these values as analog voltage signals that are read out and sent to a laptop with the help of an Arduino board. The (now digitized) measurements are captured and evaluated by a python application, which in turn controls a SuperCollider server via OSC. As the core element of our framework the SuperCollider-based sound synthesis performs a real-time processing of the water drop audio signal on the basis of the sensor readings (also see Figure 8.6). Finally, the audio output is projected to the user using off-the-shelf loudspeakers.

Interestingly, despite being an industrial-grade measurement element, we have found the flow sensor to have a surprisingly large latency, which sometimes leads to slightly odd results: While the sensor almost immediately detects a water flow, its readings need up to 15 seconds to converge to the correct value (see Figure 8.5a). When switching the water off, latency is around 7 seconds. Although the *total* measurement of

¹In our hardware setup, we use an AKG C-411.

²English manual of the *Resol Grundfos VFS 1-12 l*: http://www.resol.de/Produktdokumente/11201634_GrundfosDirectSensor_VFS.daten.pdf




Figure 8.4.: Picture of the installed Resol Grundfos flow sensor in the circuit between tap and shower head.

consumed water still seems to be rather precise, this behavior is obviously not optimal for a real-time feedback application.

DETERMINING THE ENERGY CONSUMPTION

Besides the water consumption, which can be calculated directly from the instantaneous flow measurements via numeric integration over time, in our setup it is also possible to approximate the energy that has been used to heat up the water. In order to do so, we first must determine the current energy consumption

$$P = C_p^w \cdot (T_{\text{now}} - T_{\text{ref}}) \cdot \dot{m}, \quad (8.1)$$

where $C_p^w \approx 4.2 \text{ kJ}/(\text{kg}\cdot\text{K})$ is the specific heat capacity of water, \dot{m} is the currently measured mass flow rate in kg/s, and $(T_{\text{now}} - T_{\text{ref}})$ denotes the difference between the measured temperature and those of unheated water, resulting in P , the power in Watts currently needed to heat up the water. Analogous to the water consumption, this value can also be used to obtain the accumulated energy usage.

8.3.2. Feedback Design

INPUT DATA

As we have identified the overall energy usage to be most important to reduce when taking a shower, at least in Germany (cp. Section 8.1), we have also chosen to use this value as the primary input for our sonification designs to provide feedback. Furthermore, the accumulated energy consumption is most likely also the most difficult information




Figure 8.5.: Real life data from the Resol Flow Sensor. From top to bottom, it shows a) the current water flow rate, b) the accumulated amount of used water, c) instantaneous temperature of the water, d) current energy consumption, and e) accumulated energy consumption. Note that this data represents a rather frugal consumption.

to keep track of for the user. However, as can be seen in Figure 8.5, the data on water usage is, in its character, very similar to that of the energy consumption, so that users should likewise be able to use our auditory display for keeping track of their water consumption. Furthermore, this correlation might be leveraged to simplify the technical demands of the system on the sensory side, e.g. by deriving one measurement from the other.

SHOWERING HABITS AND GOAL SETTING

The diversity of showering habits is another important aspect we had to consider for the decision on *how* the input data should be used in our sonification: While some people are done in five minutes, others can easily spend half an hour under the shower. We therefore argue that using a fixed scale for the total consumption, as has been done in previous projects (cp. Section 8.2), cannot do justice to the diverse shower habits that must be expected to be encountered in everyday life. Instead, we followed a similar approach as for the InfoPlant, employing the 4/5-factor principle in order to support a slow and continuous reduction in energy usage by revealing and emphasizing the progress users make towards an adapting short-term goal (see Section 7.5.2).

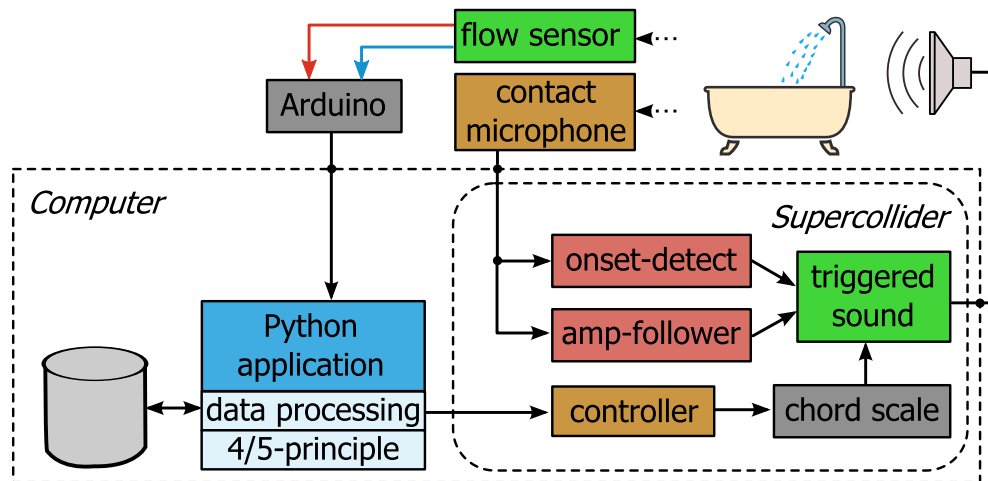


Figure 8.6.: Overall architecture of the Sonic Shower providing feedback via the transient-triggered sonification.

8.4. Sonification Design

In order to explore the design space of a blended sonification in a showering context, we have developed three different sound designs providing feedback on the energy usage while taking a shower. As discussed above, the input for each of these designs can be both the instantaneous and accumulated consumption, in relation to the current 4/5-factor goal.

8.4.1. Pitch-Based Blended Sonification

Our first sound design directly maps the accumulated energy consumption to the center frequency of a band-pass filter that is applied to the input signal of the environmental sounds as picked up by the contact microphone. More specifically, the frequency range from an A2 at 110 Hz to a f#6 at around 1480 Hz is used to indicate the consumption's progression from 0 kWh up until the current (short-term) goal C_g , using a dynamic frequency resonator filter. The main advantage of this design is that it represents a truly continuous feedback, and due to the rather easily discernible changes in pitch, it very well conveys the rising level of consumption. On the other hand, the continuous mapping also makes it rather difficult to convey a boundary (i.e. the current goal), as the only limits in pitch are given by the user's sense of hearing.

INTERACTION EXAMPLES

In order to provide an impression of a real-life usage of this and the following sonification designs as well as to illustrate the differences between the various approaches, we have recorded several sound examples, which can be found on the accompanying CD. All of them are based on the same input data, as shown in Figure 8.5, in conjunction






Semitones	Harmonic interpretation	Visual
0, 4, 7	major	
0, 5, 7	suspended (sus4)	
0, 3, 6	diminished	
0, 4, 6, 8	augmented (#11)	
0, 1, 2, 3	None	

Table 8.1.: Musical chords used for the affective dimension in the sonification.

with the corresponding auditory input from the contact microphone and the ambient sounds, as picked up by a dynamic microphone. The pitch-based blended sonification is illustrated in sound example S1. Note that since the sonification does not change between the two showering phases, we have cut out the (rather uninteresting) middle part of each recording in order to focus on what happens during the two actual showering phases. Also, the volume of the sonification was set to a relatively high level in order to make the sonification part more salient for the listener.

8.4.2. Blended Chords Sonification

With our second approach, we wanted to go beyond a neutral mapping of input data onto a specific sound parameter, like in the first design, but instead map the data on an *affective dimension* of sound in order to establish a (subconscious) association of a low consumption with a positive, and of a high consumption with a negative emotional response (cp. RQ8). Although work on affective sonification is clearly in its infancy and there is still a lack of theory and established methods, there have been relevant studies in the context of music performance (e.g. Bresin & Friberg, 2000), and we could find some evidence for an inherent emotional understanding of chords, i.e. the simultaneous occurrence of several tonal sounds within the twelve-tone system, for musicians and non-musicians alike (cp. Pallesen et al., 2005).

Consequently, we decided to use a set of differently colored musical chords to convey an affective association. The specific chords that we have used range from a major chord, which is universally perceived as consonant and harmonic, to a dissonant cluster chord, based on chromatic semitones with no meaningful harmonic interpretation (cp. Table 8.1). In our design, each chord corresponds to reaching a specific level of accumulated energy use, with the most disharmonic one indicating that the user's current consumption goal C_g has been reached.




Figure 8.7.: Spectrogram of the first minute of sound example S4.

Similar to the first sound design, a bank of dynamic frequency resonators is used to filter the input sound signal in such a way that only the environmental sound’s frequency components associated with the respective chord augment the existing soundscape (also cp. sound example S2). In addition to the chord’s fundamental tones, as shown in Table 8.1, we also used several transpositions in order to extend the filter’s overall frequency range and thereby to better convey the specific character of the respective chord.

8.4.3. Transient-Triggered Blended Chords Sonification

The filtering approach that we have used in the two previous sound designs is fully in line with the concept of auditory augmentation, which enforces the sonifications to stay very close to the existing environmental sound. Pure filtering, however, provides only limited design opportunities compared to the possibilities in computer sound synthesis. In order to further explore the design space of a blended sonification, we enhanced the previously discussed chord-based approach in such a way that the chords’ quality is conveyed by synthesized sounds that are tightly coupled to transients in the input audio signal. More specifically, we interpret the sound of the water falling onto the bathtub as impulses with differing energy and loudness, and use them for triggering sounds that are pitched according to randomly selected tones of the current chord. To stay in musical terms, we produce a sound-induced continuous arpeggio, where, similar to the previous approach, several transposed versions of the chord are used to construct a “chord scale” the arpeggio is based on.

In order to give users an even better impression of their energy consumption, we furthermore mapped the instantaneous water temperature onto the overall pitch of the

chord's notes. More specifically, since to us the higher pitched arpeggios managed to evoke thoughts of coldness, while the lower pitched ones made a rather "energetic" impression, reminiscent of a higher energy consumption, we also chose to represent a high temperature with a low-pitched sound. Furthermore, we discovered that the harmonic impression of the chords is fairly susceptible to *global* changes in pitch and that these can easily mask the changes in chord type. Consequently, only the existing range of the chord scale's tones is used to convey the impression of a pitch change, i.e. the input data effectively modifies the position of a frequency range, which in turn controls the selection of tones from the chord scale.

Concerning the triggered sounds themselves, we realized two different designs for this approach: The first one uses a rather short, percussive-sounding tone (cp. sound example S3), whereas the second design employs a number of recorded samples of water glass clinks in order to achieve a more lively and natural impression (cp. S4). Figure 8.7 shows a log-frequency spectrogram of this sound example, covering the first showering phase of approximately one minute, where some of the sound characteristics can also be observed visually: After the first seconds, with only ambient sound audible, the sonification slowly fades in, initially indicating a very low water temperature. A pitch drop, still within the first chord scale, then identifies a rising water temperature, after which it is regulated back and slowly decreases over time.

At 33 seconds into the sound example (label A in Figure 8.7), the chord scale changes, thereby indicating a rising level of consumption. This happens as well after 56 seconds (label B) and, in the second showering phase, at 01:41, after which the scale becomes the most disharmonic nearly at the end of the shower (at 02:06), which means that the user has missed his goal by only a little. At the beginning of the second phase, the small variations in temperature are, again, clearly audible.

8.5. Online Survey

In order to learn more about the use of an auditory information display for energy awareness purposes, we initially planned to conduct a study to critically examine the InfoDrops system under real-life conditions, similar to the evaluation of the InfoPlant, described in Section 7.5. However, since conducting such a real-life study turned out to be quite problematic due to privacy concerns for a further evaluation in the shower environment, we decided to instead focus on the auditory design and a more qualitative analysis of the InfoDrops system as well as further understanding the showering habits of potential users, and with the help of LimeSurvey³, we created an online questionnaire evaluating six different auditory designs.

³LimeSurvey: an open-source on-line survey application written in PHP (<https://www.limesurvey.org/>)

EVALUATED SOUND DESIGNS

Due to the similarities between the blended chords and the transient-triggered sonifications, we first selected the pitch-based and the blended chords designs to be evaluated in the survey, based on the results of a short preliminary study indicating a slight preference for the blended chords sonification. Furthermore, in order to examine a wider design space than covered by the already presented blended sonifications, we additionally developed a speech-based and three metaphorical auditory designs⁴ to give feedback on resource usage while taking a shower.

Speech-based design: To provide a contrasting example of how to give information about energy consumption during a shower, we created a purely speech-based feedback: For every 100 consumed watt-hours, a female voice – synthesized using Google’s text-to-speech engine – announces the accumulated energy usage up until that point. After reaching a certain threshold, i.e. towards the end of the shower, when having consumed more energy than during previous shower episodes, the announcements are made more frequently: In the provided example, they then occur every 50Wh.

Design emphasizing the monetary aspect of taking a shower: In order to address people who might be concerned about the costs of consuming energy to heat up the water, a second auditory design uses the sound of a cash register as an auditory icon representing the spending of money. Similar to the speech-based design, the sound is played back for every 100 consumed watt-hours, with more frequent occurrences towards the end of the shower.

“Emotional” Sonification: To evaluate whether creating a strong emotional response can be beneficial in persuading users to reduce their energy usage, this design uses the recorded sounds of a small child to indicate the level of consumption: While the accumulated energy usage is below one’s individual threshold, the sound of a happy, laughing child frequently provides positive feedback – analogous to the other designs every 100Wh. On the other hand, after exceeding this threshold, the sound of the child crying should give a strong incentive to come to an end.

Design emphasizing environmental issues: This auditory design aims to address people concerned about the environmental implications of taking a shower. Since the exact consequences of (excessive) energy consumption are not only rather difficult to determine, but also quite challenging to convey with auditory means, we decided to use the sound of a tree falling down as a relatively generic, but easily recognizable auditory icon. The sound is played back in the same pattern as the one in the second design.

⁴Corresponding sound files can be found on the accompanying CD.

8.5.1. Study Design

In order to be able to reach as many people as possible, we decided to conduct an online survey instead of a paper-based one. The survey was designed entirely within LimeSurvey, although some features had to be manually added as JavaScript extensions. The survey consisted of four parts:

- An introduction, where the participants were informed about outline and purpose of the survey and answered some general questions about themselves, such as age, musicality, and attitude towards saving energy.
- Questions about the participants' shower habits, for example how often and how long they are taking a shower, and the reasons for more frequent (or seldom) showering.
- The main part, where participants were listening to the six sound examples and subsequently could assess and give their opinion on them by answering several questions. For each sound design the same set of questions was used.
- At the end, the participants were asked to select their favorite sound design and, if possible, to say why. They were also given the possibility to make additional comments on the overall concept, one particular sound design, or the survey itself.

Most of the questions were posed as statements that the participants indicated their level of agreement to on a 7-point Likert-type scale. Additionally, there were some open-ended questions, which could be answered using free-form text. In order to convey the feeling of taking shower at least to some extent, a video of a shower head with water running was presented while playing the sound examples. Furthermore, to prevent ordering effects between the different sound designs, they were presented in randomized order. The survey was distributed mostly via mailing lists and social networks (e.g. Facebook).

LOUDNESS CALIBRATION

In order to allow the different auditory designs to be evaluated independently from their loudness, all presented sound files were normalized using a psychoacoustic analysis (Robinson, 2012). Furthermore, before listening to the various sound samples, all participants adjusted their volume in such a way that a prerecorded speech sounded as if a person was one meter away talking to them. Finally, we instructed the participants to make any changes in volume only through the presented player-interface (i.e. not via the computer settings) to enable us to track those changes.

8.5.2. Goals

Considering the nature of this study, the primary goal was less to test any concrete hypotheses, but more to learn about the topic of showering in general as well as to explore the design space for an effective auditory feedback system. More specifically,

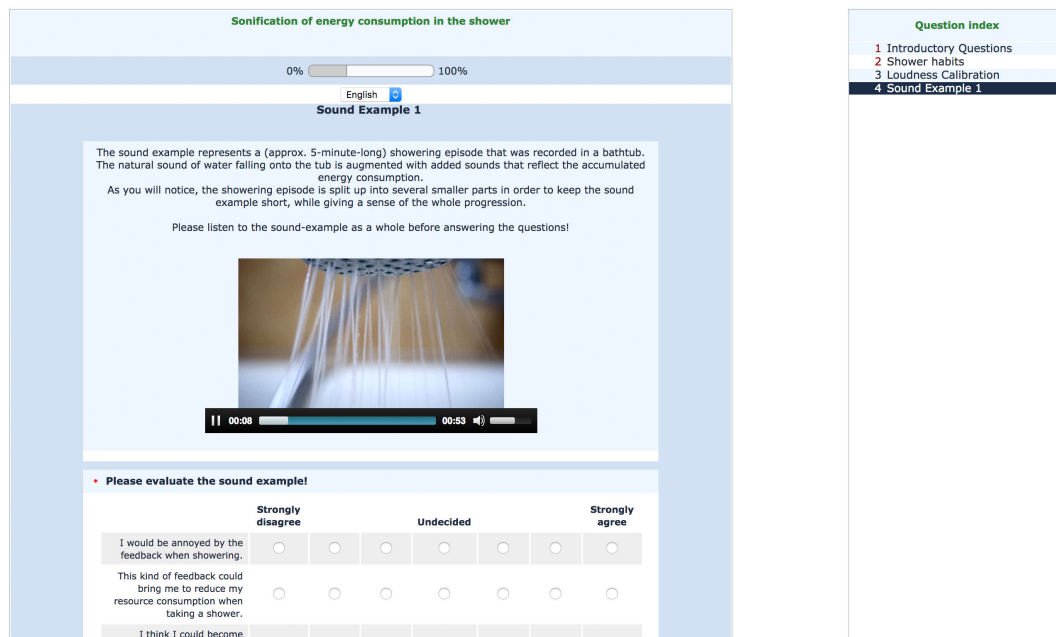


Figure 8.8.: Screenshot of the survey. The progress bar at the top allows participants to estimate how much time they still need to finish the survey. The navigation bar on the right can be used to have a look at (or edit) previously given answers.

we wanted to find out more about the people’s showering habits, e.g. how often and how long they are taking a shower, but also about the reasons why they behave this way, in order to find any information that can be used to improve the system, such as a frequently mentioned motivation that could be incorporated into a future feedback design. Furthermore, we wanted to know how the survey participants perceive and evaluate the six prototypical sound designs described above. Since taking a shower represents a rather private situation, we were most interested in the (un)obtrusiveness of the different designs and also expected the participants to be mostly concerned with this quality of the feedback. As it is important for an exploratory analysis not only to collect the responses to a range of predefined questions, we made sure to give participants on several occasions the possibility to freely express their opinion about the designs and to give unconstrained feedback, whose thorough analysis should give us further insights into how to improve or extend the InfoDrops system and the feedback designs.

8.5.3. Results

In total, exactly 100 people participated in the survey, with almost 60% being female. The participants were between 19 and 71 years old, with an average age of 27 years. Only two persons indicated to have a (slight) impairment of their sense of hearing




Figure 8.9.: Summarizing chart of how often the survey participants indicated to take a shower. The x-axis describes how often the participants shower per week, with “every other day” being interpreted as 3.5 times per week.

Due to the tendency to only select a rather coarse approximation of this information (i.e. the majority of participants indicated to shower either daily or every other day, with almost no answers in between), we additionally fitted a Gamma probability distribution function⁵(PDF), showing a more realistic probability distribution.

($\mu_{\frac{1}{2}} = -3$). Furthermore, the majority of participants said to be interested in a system that helps reducing their energy consumption ($\mu_{\frac{1}{2}} = 2$), which is in line with the results of the InfoPlant study.

SHOWERING HABITS

Of the 100 participants having completed the survey, 87 also answered questions about their showering habits, including why, how often, and how warm they usually take a shower:

Shower Frequency. On average, the participants indicated to shower 4.66 ± 2 times per week (cp. Figure 8.9), with women showering significantly more often (5.5 times) than men (4.12 times, $p < 0.01$ for a Mann-Whitney U test). Interestingly, only very few of the participants think that they shower “comparatively often” – and although there is a correlation between the shower frequency and the responses to this statement ($\tau = 0.28, p < 0.01$), even the people showering more frequently than the average do not generally agree to it ($\mu_{\frac{1}{2}} = 0$). For those who shower *less* often than the average, many participants stated health reasons (i.e. the fact that too frequent showering is bad for the skin) as a motivation to do so ($\mu_{\frac{1}{2}} = 1$). Far more seldom, the wish to save water or energy was given as a reason for showering less frequently (both $\mu_{\frac{1}{2}} = -1$).

⁵We chose to use the Gamma probability distribution function mainly because – different to the normal PDF, for example – it is explicitly supported only for positive numbers $x \in (0, \infty)$, which is obviously a sensible assumption for this question.

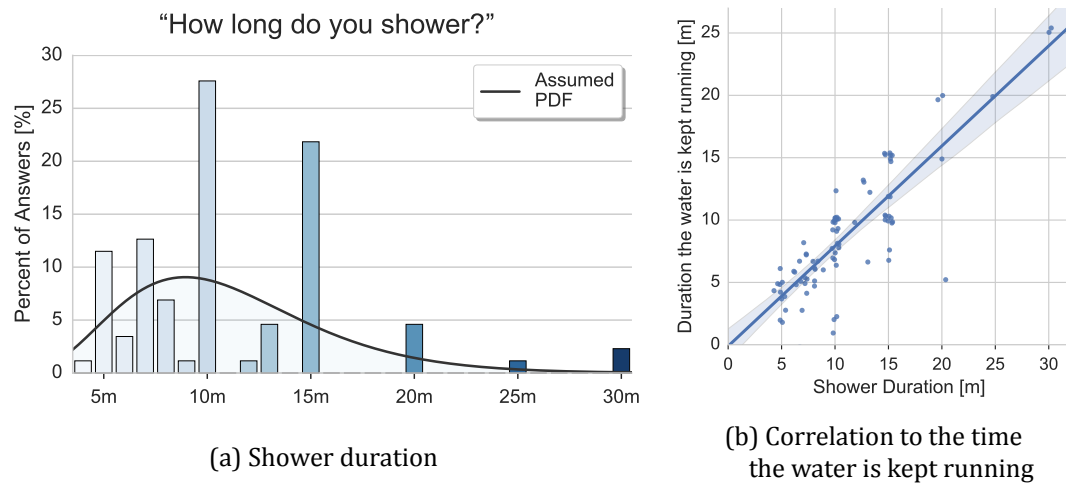


Figure 8.10.: **(a)** Summarizing chart of how long the survey participants indicated to take a shower. Similar to Figure 8.9, we additionally fitted a Gamma PDF to visualize a more realistic probability distribution. **(b)** Relationship between the total shower duration and the time the participants are letting the water run. In order to better visualize clusters of identical answers, a small amount of jitter has been added to the scatter plot.

Important Aspects of Taking a Shower. The majority of participants indicated that when taking a shower it is most important to them to simply be clean afterwards ($\mu_{\frac{1}{2}} = 3$). Furthermore, while certainly not seen as essential, some participants also stated the wish to relax or warm up during a shower (both $\mu_{\frac{1}{2}} = 1$). On the contrary, for most people, it does not seem to be particularly important to be done with the shower “as quickly as possible” ($\mu_{\frac{1}{2}} = 0$).

Shower Duration. Considering all survey participants, the average duration of a shower was stated to be 11.03 ± 5.2 minutes (cp. Figure 8.10a). Again, women indicated to shower longer (12 minutes) than men (9.5 minutes), although the difference is not significant ($p = 0.06$ for a two-tailed Mann-Whitney U test). Furthermore, the participants stated to keep the water running for about 8.75 ± 4.9 minutes, or $79.3 \pm 24\%$ of the total duration of the shower. Interestingly, this percentage is independent of the shower’s duration, i.e. according to those answers, the amount of water used can be seen as a linear function of the duration (also cp. Figure 8.10b).

Use of Warm Water. The majority of participants (almost 60%) indicated to use warm water when taking a shower, with another 32% even preferring hot water. This essentially means that, similar to the water consumption, the energy usage will most likely correlate strongly with the shower duration in most cases – although the amount of needed energy of course also depends on the exact temperature the water must be

heated up to. Of the 8% of the participants who indicated to shower with lukewarm or cold water, most stated health reasons (i.e. to improve blood circulation) as a motivation to do so ($\mu_{\frac{1}{2}} = 2$). Saving energy, however, does not seem to be an important aspect in this decision ($\mu_{\frac{1}{2}} = -1$), which is in line with the results found for the reasons to shower less frequently.

AUDIO SETUP AND LOUDNESS

Half of the participants followed our recommendation and listened to the sound files through various kinds of headphones, mostly closed (23%) and in-ear (20%) ones. The other half, despite relying on speakers, did not indicate to be impaired by background noises any more than the headphone users (around 20% for both groups). After volume calibration, there were only very few adjustments in loudness initiated by the participants: of all presentations of sound samples, changes occurred only for 5.5% of them. This result suggests that there are no significant differences in the perceived loudness of the different sound samples and can probably be attributed to the previously described loudness calibration of the sound files. On the other hand, the sparseness of the data does not allow any further interpretation. In fact, a Friedman test indicated that there are no significant differences between the auditory designs in terms of user-initiated loudness changes ($p = 0.32$).

COMPARISON OF AUDITORY DESIGNS

Understandability. Surprisingly, the sonic environment of a shower seems to be even more challenging than expected. On average, only 57% of the participants indicated to have understood what the sounds that were augmenting those of the shower represent. This is especially problematic for the metaphorical designs, since, different to the other, more abstract sonifications, one needs to identify the exact nature of the sounds in order to understand their meaning. Moreover, based on the participants' descriptions of what they think they heard, it became apparent that their impression was oftentimes not what they were intended to hear – even when being confident about it. For example, the sound of the cash register was frequently mistaken for a bicycle bell and only around 36% of the participants associated the notion of spending money through taking a shower with the presented sounds. The design most difficult to recognize was the one aiming to remind of the environmental consequences of consuming too much energy: The sound of the tree falling down was apparently not only too similar to other known sounds, but it was also masked by the sound of the shower, which exhibits a similar broad spectrum. Consequently, this sound was interpreted correctly only by 29% of the participants. The least problematic design in this respect is obviously the speech-based one, as there is little to no room for misinterpretation.

Unobtrusiveness. Although being easily understandable, the speech-based design was not considered as particularly unobtrusive by the survey participants ($\mu_{\frac{1}{2}} = -1$)




Figure 8.11.: Agreement to statements indicating the *unobtrusiveness* of the respective design.

and significantly less so than both blended sonifications ($\mu_{\frac{1}{2}} = 0$ and $p < 0.01$ for both comparisons using a Wilcoxon signed rank test; also cp. Figure 8.11). The only design perceived as more obtrusive by the participants than the speech-based one is the “emotional” sonification ($\mu_{\frac{1}{2}} = -2$). An analysis of the participants’ comments on this design indicates that this is most likely due to the fact that it was indeed successful in inducing an affective response ($\mu_{\frac{1}{2}} = 1.5$ agreement for corresponding statements, and significantly more than all other designs), which was, however, perceived as too disruptive and inappropriate in a showering context.

Considering the fact that the blended sonifications are not based on discrete occurrences of sounds like the other designs, but are continuous and thus able to provide a much tighter feedback loop, it is certainly remarkable that, together with the design using the sound of a cash register, they are rated best in terms of unobtrusiveness. However, we can also see that about half of the participants indicated to be skeptical about those designs being unobtrusive enough to be used in everyday life, which underlines the difficulty of providing additional interaction possibilities during such a private situation like taking a shower – especially giving feedback about potentially unpleasant information.

PARTICIPANT'S COMMENTS

In addition to indicating the level of agreement to a range of predefined statements, the participants were given the possibility to make individual comments and suggestions regarding each of the sound examples and about the concept as a whole, providing valuable insights into *why* the feedback was perceived in a certain way. In our analysis, we had a look at why participants chose to select a particular design as their favorite one, and we identified several clusters of topics that were mentioned repeatedly across all comments.

Favorite Sound Design. Asked about the design they liked best, most participants chose either (a) the one using the sound of a cash register (29%), (b) the speech-based

design (22%), or (c) the blended-pitch sonification (20%). A more in-depth analysis of the comments that were made in conjunction with this rating showed that the main reason for choosing (a) was that this design was perceived as least obtrusive by many people due to the shortness and simplicity of the sound. In contrast, (b) was chosen mostly for its clarity and perceived precision, i.e. for conveying exact numbers. Similar to (a), the blended-pitch sonification (c) was chosen for being perceived as more subtle and less salient than the other designs, but also for its continuous feedback and for being intuitively understandable and “embedded” into the shower’s soundscape.

The broad selection at this point shows that there is no optimal auditory design, but that each has its strengths and weaknesses, which are yet to be combined in a new design. Interestingly, only one participant suggested to use visual means to convey the energy consumption in the shower instead of our auditory approach.

Preserving the Calm of Taking a Shower. As expected, many of the participants’ comments are concerned with how much a particular feedback design interferes with the soothing effect of taking a shower: Although the preferences seem to be somewhat individual and there is no clear “winner” in this regard, many of the designs were judged solely based on this aspect and were either praised for their subtlety or criticized for their disruptive nature. Some participants also noted that the willingness of people to actually use such a feedback system critically depends on its unobtrusiveness, and that an “annoying” feedback, although probably even effective in terms of inducing a reduction in consumption, would probably be just switched off.

Comprehensibility of The Sounds. On the other hand, there were several comments on the intelligibility of the sounds: In addition to the issues discussed above, some participants criticized that the additional sounds were simply too quiet. These comments highlight the difficulty to strike a balance between comprehensibility and subtlety of the sound design. Furthermore, there were some comments about the speech-based design suggesting that, although in principle intelligible, the information about watt-hours is too abstract and not connected to any real-world experience and thus does not help in understanding one’s energy consumption.

Use of More Positive Sounds. Another important aspect of the feedback, which was frequently mentioned by the survey participants, is how it is perceived and interpreted by potential users. For example, several feedback designs were criticized for being too “negative” or patronizing. One comment points out that if a user decides to install a feedback system, there should be no need for “hidden” messages. Interestingly, even the speech-based feedback, which can be considered as relatively neutral, was perceived as criticizing and too negative by some of the participants. Consequently, several comments suggest to use more positive and encouraging sounds instead of those hinting at the negative aspects of taking a shower.

8.6. Discussion and Conclusions

Although the conducted online study only allows limited conclusions in terms of practical implications of the Sonic Shower, (i.e. a potentially reduced energy consumption and change in behavior induced by the system), the large number of responses and the frequently very detailed explanations provide us with some valuable insights on the design of feedback and how it is perceived by potential users.

AFFECTIVENESS

As one important aspect, we can review the different types of affective cues that have been incorporated into the feedback designs in order to estimate if their use is beneficial for an auditory ambient display providing feedback on the user's consumption (cp. RQ8). First, we can state that at least some of the designs *did* manage to evoke an emotional response for the participants – even on a rather conscious level. However, this seems to have a rather negative effect on the overall reception of these designs, both in terms of acceptance and unobtrusiveness of the system: Taking into account all feedback designs, there is a strong negative correlation between the perceived affectiveness and the unobtrusiveness of the feedback ($\tau = -0.2$; $p < 0.01$). However, to further analyze this issue we must distinguish between the different types of affective cues:

- As we have seen, the “emotional” feedback design was indeed successful in inducing an affective response, but consequently was perceived as least unobtrusive of the designs.
- Similarly, the participants who understood the meaning of the two symbolic designs hinting at the negative aspects of taking a shower also assessed them as rather affective (both $\mu_{\frac{1}{2}} = 1$) and not very unobtrusive, especially for the design emphasizing environmental issues ($\mu_{\frac{1}{2}} = -2$).
- Finally, the blended chords design aimed to subliminally convey emotional undertones without a specific underlying “message”. Unlike for the other designs, there is no significant correlation between the affectiveness and the acceptance of the blended chords sonification.

Interestingly, despite their impact on the feedback's unobtrusiveness, the emotional cues were also seen as somewhat conducive to the goal of changing one's behavior, as there is a slight correlation between the perceived affectiveness of a feedback design and the participants' estimate that it can bring them to reduce their consumption ($\tau = 0.1$; $p = 0.13$, for the designs employing emotional cues). Our conclusions from these results are as follows:

1. Very obvious attempts at creating a strong emotional response seem to be mostly incompatible with the general concept of a peripheral display, since such a response

is almost inevitably quite arousing and thus disruptive, therefore interfering with the unobtrusiveness as an essential element of such a display.

2. For ambient displays used as feedback devices, the emphasis on the negative aspects of one's behavior *can* lead users to change it, but might also create reactance as a similarly strong emotional response leading users to completely disregard the feedback. Furthermore, for users *consciously* deciding to install a feedback system, there should be no need for a "hidden" message, considering their given intrinsic motivation. However, when users *want* to be reminded of specific consequences of their behavior, this emphasis should still help changing it.
3. As an opposing approach, it might be reasonable to focus on evoking a *positive* emotional response. However, this would also mean to only emphasize positive changes in behavior, and it is questionable whether giving feedback only in those cases is as effective as comprehensively doing so. Nevertheless, this approach could be part of a larger feedback concept, e.g. by combining a neutral ambient feedback with a (potentially) positive and encouraging message at the end of a shower.
4. Finally, we see incorporating a subliminal affective cue as in line with the general concept of a peripheral display, which also seems to retain the unobtrusiveness of the system. However, the conducted survey cannot give any clear evidence on the practical superiority of this approach over a neutral feedback, as on the one hand the perceived affectiveness of the feedback seems to have an influence on the participants' estimate that it can bring them to reduce their consumption, but on the other hand, the blended chords sonification was not rated any better in this regard than the neutral blended pitch sonification.

UNOBTRUSIVENESS

The results of the survey furthermore give information on which qualities of a peripheral display make it less obtrusive and distracting (RQ4). In addition to providing insights into the impact of using affective cues on the unobtrusiveness of such a display, the results hint at a rather plain feedback design as being advantageous in terms of understandability and unobtrusiveness: As we have seen, the sonification using the sound of a cash register was rated best with regards to these aspects, and the users' comments clearly indicate that this was primarily due to shortness and simplicity of the sound. Second, using a blended sonification for a feedback display seems to be a good choice when aiming for unobtrusiveness – especially when providing *continuous* feedback – as both blended designs were rated second best in this regard. Finally, the participants' comments clearly indicate a preference for an unobtrusive design (RQ5), which is also reflected in the strong correlation between this quality and the acceptance of the different sonifications ($\tau = 0.4$; $p < 0.01$).

COMPLEXITY AND VARIABILITY OF THE CONTEXT

Regarding RQ3, the environment the Sonic Shower is installed in as well as the activity that it supports is rather clearly defined. The perceptual situation, however, is not completely fixed, as the user can easily move (or at least turn) around while taking a shower. Although these circumstances were reported to be a problem for visual displays (cp. Section 8.1.1), we see no indication that this is also the case for the InfoDrops system, as an auditory display only depends on the user's proximity to the sound source.

SUFFICIENCY OF FEEDBACK

Regarding RQ6, we must state that for all feedback designs, the participants were not too convinced that the feedback would bring them to reduce their consumption ($\mu_{\frac{1}{2}} = 0$), which we see as a further hint that feedback alone is not enough to change one's behavior (also cp. Section 7.6). This is also in line with the results of a study dealing with feedback on water consumption, where no consistent reduction could be found by simply installing a feedback display in the shower (Kuznetsov & Paulos, 2010). Similar to the results from the InfoPlant study, a preexisting intrinsic motivation, as indicated in the survey, seems to be the best predictor for a subsequent reduction ($\tau = 0.33$; $p < 0.01$). Nevertheless, for those who *are* motivated, being made aware of the consumption is assessed as highly important in supporting a reduction in consumption by the study participants, showing that feedback might not be sufficient, but still highly conducive to changing one's behavior.

9.

The EcoSonic System

9.1. Introduction

Eco-driving refers to a driving style that is both ecological and economical, thereby reducing the energy consumption when driving a car. Existing research highlights a range of opportunities and benefits that are associated with adopting such an energy-efficient driving style: First, as the energy consumption and pollution that is produced by both conventional and electric cars is one of today's major causes for greenhouse gas emissions (United States Environmental Protection Agency, 2015), eco-driving can have a major impact in terms of alleviating the negative effects of climate change (Barkenbus, 2010). Even more, contrary to current advances in building more efficient engines, a change in driving style does not require to buy a new car, but can be applied to any existing vehicle. Additionally, eco-driving is becoming even more important with hybrid and electronic cars, as the driving style has an even greater impact on the energy consumption when compared to conventional combustion engines (Romm & Frank, 2006). Finally, fuel-efficient driving generally also leads to a safer driving style (Young, Birrell, & Stanton, 2011) and can be done without a huge impact on trip time (Evans, 1979).

Despite these advantages, the behavior change towards adopting these techniques poses to be a challenging one: Currently, there exist a number of visual fuel efficiency displays providing feedback on instantaneous or long-term fuel economy to support

This chapter is, in parts, based on:

Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2014b). EcoSonic: Auditory Displays supporting Fuel-Efficient Driving. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational - NordiCHI '14* (pp. 979–982). New York, New York, USA: ACM Press

Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2014a). EcoSonic: Towards an Auditory Display Supporting a Fuel-Efficient Driving Style. In *Proceedings of the Conference on Sonification of Health and Environmental Data (SoniHED 2014)* (p. 56). York, England

Hammerschmidt, J. & Hermann, T. (2017). EcoSonic: Auditory Peripheral Monitoring of Fuel Consumption for Fuel-Efficient Driving. *Displays*, 47, 40–50.

drivers in achieving a lower level of fuel consumption (e.g. Manser, Rakauskas, Graving, & Jenness, 2010). Observing these displays, however, requires both visual attention and cognitive effort and can easily distract users from the actual driving task and thus be detrimental to a safe steering of the car (Brooks & Rakotonirainy, 2005). This is especially problematic insofar as it is precisely in situations when drivers should keep their eyes on the road (e.g. when quickly accelerating or approaching a street crossing or traffic lights) that the information from such a display becomes most relevant. Contrary to that, auditory displays have shown to be able to convey information in a less distracting way: Preliminary research on in-vehicle auditory interaction hints at a significantly reduced impact on attention as well as an improved effectiveness in terms of user performance (see Section 9.2.2), which is also supported by psychological research on multiple resources theory and dual task interference (cp. Section 2.3.3).

Based on these findings we have developed the EcoSonic system as an ambient display to support users in driving in a more economical way. To further contribute to the research question if auditory ambient interfaces are a feasible alternative to visual ones (RQ2), the display is – similar to the Sonic Shower – a purely auditory one. Different to the previously discussed two projects, however, we assumed that potential users will have a certain intrinsic motivation to change their behavior and consequently moved away from a feedback design that attempts to *persuade* users to change their behavior, towards the notion of an ambient display *supporting* them in doing so.

9.2. Related Work

While to our best knowledge, there is no research specifically studying the use of an *auditory* display for feedback on fuel consumption, the topic of eco-driving as well as the evaluation and comparison of visual fuel economy displays is well covered in literature (see Section 9.2.3). Additionally, more and more work is being done on the use of in-vehicle auditory displays in general (e.g. Jeon et al., 2015) as well as for specific use-cases, e.g. collision warnings, skill acquisition, or in-car entertainment systems (Section 9.2.4).

9.2.1. Efficacy of Eco-Driving

Although the findings for the precise amount of achievable reduction of fuel consumption vary to some degree, there already exist a number of studies evaluating the efficacy of eco-driving, which, taken together, clearly indicate a significant positive impact of an energy-efficient driving style. Gonder, Earleywine, and Sparks (2012), for example, evaluated the potential fuel savings that can be achieved by an “energy conscious” driving style: Using real-world data from trips collected with GPS devices and a vehicle model of a midsize car, they performed extensive simulations in order to assess the potential savings from specific behavior changes as well as the prevalence of inefficient

driving. The authors conclude that for aggressively driven trips, the adoption of efficient driving behaviors can result in fuel savings of approximately 20%, and even for more moderately driven trips, a 5%-10% reduction of fuel consumption is realistic.

Similarly, Johansson, Gustafsson, Henke, and Rosengren (2003) conducted a study with 16 driving school teachers (i.e. very experienced drivers), who were trained to additionally become eco-driving instructors. The researchers measured various parameters in the vehicle during two test-runs that were performed both before and after the participants had received instructions on how to drive in a fuel-efficient manner. Even though the authors point out that in the second run not all eco-driving recommendations were adhered, both fuel consumption and the emission of carbon dioxide were reduced by 10.9% on average. They furthermore note that the average speed was basically the same before and after instructions.

Finally, in a workshop organized by the International Transport Forum (2007), more than 20 eco-driving initiatives and pilot projects from several countries were reviewed. The workshop participants found out that the average fuel economy improvements were between 5-15%, with the best result for individual drivers being around 20-50%. Interestingly, available data for railways also indicate a significant 5% of possible savings of CO₂ emissions, and the potential fuel savings for inland-waterways were estimated to be 10-15% .

ELECTRONIC AND HYBRID CARS

While the influence of driving behavior on energy efficiency for electronic and hybrid vehicles has not been as intensively researched as for vehicles with combustion engines, there already exist a few studies indicating that there is a rather large impact for those types of cars, which might be even greater than those for conventional cars.

Knowles, Scott, and Baglee (2012) conducted a study with eleven test drivers, who were asked to drive a specific track with an electric vehicle. They conclude that even though the specific aspects of one's driving style that influence the energy consumption of electric vehicles seem to be somewhat different to the ones that are important when driving a conventional car, there still is a major impact on the energy consumption: In the study, the efficiency of the electric vehicle varied from 0.46 km to 1.89 km per percent of battery charge.

Finally, in an article by Romm and Frank (2006), dealing with hybrid cars in general and specifically with the fuel consumption that is achievable with these types of cars, the authors state that they are far more sensitive to how they are driven: While in principle, fuel consumption is somewhat lower than for conventional cars, aggressive driving can cause the fuel efficiency of hybrid cars to decrease by more than 30 percent.

9.2.2. Distraction by In-Vehicle Systems

As the number of in-vehicle systems that are available to drivers is steadily increasing, there is a justified apprehension that these systems might cause considerable distraction and reduced driving performance, which in consequence could lead to an increased incidence of traffic accidents. Several researchers have thus tried to quantify the effects of existing in-vehicle systems and find out the type of interaction that might best be used for those systems to alleviate the risk of distraction.

In a study with 23 participants, Lansdown, Brook-Carter, and Kersloot (2004) evaluated in which way distractions from in-vehicle information systems might affect drivers: The participants had to drive a test track in a driving simulator and were confronted with additional visual tasks, representing the interaction with in-vehicle systems. The authors found out that those tasks led to significantly reduced headways and increased brake pressure. Additionally, they observed compensatory speed reductions and an increased self-reported workload. Taken together, the authors see these results as evidence for a clear safety risk when using visual in-vehicle information systems. Similarly, Tsimhoni and Green (2001) found out that adding a secondary visual task led drivers to wander more in their lane and to depart from it more frequently. The increase in the standard deviation of the lane position was almost 80%.

AUDITORY SYSTEMS

While visual in-vehicle systems have been studied rather extensively, the impact of auditory interaction in the car is not a very well-researched area. Some existing work, however, hints at the advantages of employing the auditory channel: Wierwille (1993), for example, has analyzed the different modalities by which information can be conveyed to the user. He points out that visual attention is a limited resource that is used above all for the primary task of driving, and using it to provide additional information has a measurable negative impact on driving performance. On the contrary, the auditory sense is assessed to be used the least for the driving task and is therefore recommended to be utilized for prospective in-vehicle displays.

In line with these observations, Jensen, Skov, and Thiruravichandran (2010) conducted a study evaluating the impact of navigation systems on driving performance. They compared different audio-visual output configurations for such systems and found that a visual-only output not only led to a considerable amount of eye-glances, but also to a significant decrease in driving performance. While a user questionnaire indicated preference for an audio-visual configuration, the audio-only navigation system actually performed best in terms of these two aspects.

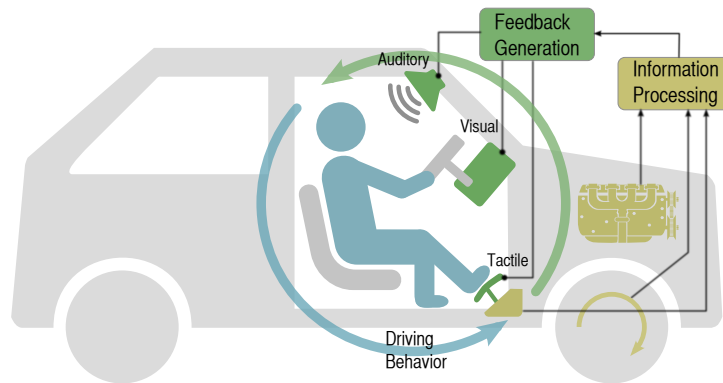


Figure 9.1.: Feedback loop for an eco-driving display: The driver controls the (acceleration of the) car, which in turn controls the feedback display. The feedback can be generated in visual, auditory, or tactile ways, and gives support in improving driving behavior.

9.2.3. Comparison of Existing Eco-Feedback Systems

As there already are a few (mostly visual) systems to provide feedback on eco-driving performance, a number of studies try to assess and compare these systems in terms of efficacy and user acceptance. For instance, Manser et al. (2010) examined ten commercially-available feedback systems according to a predefined set of criteria, including ease of comprehension, usefulness, and accessibility. In a second step, they derived common, more general design concepts from these systems, which then underwent a usability evaluation based on information from a previously conducted focus group (Jenness, Singer, Walrath, & Lubar, 2009). Finally, two designs were selected and evaluated in a driving simulator study in order to examine the utility of these designs. Interestingly, although receiving “behavioral” feedback on acceleration and deceleration led to a significantly smoother driving, it didn’t perform better in terms of fuel consumption than just telling the participants to drive more economically. Conversely, participants receiving direct feedback on how much fuel they consumed attained significantly better fuel economy.

Following a slightly different approach, Tulusan, Soi, Paefgen, Brogle, and Staake (2011) conducted a questionnaire and several semi-structured interviews in order to learn more about which eco-feedback types might be preferred most by car drivers. They found out that a comparison with an average consumption (e.g. of other drivers) in order to assess what could have been saved through a more ecological driving behavior was considered to be beneficial by the participants. Furthermore, the authors point out that most preferable were unobtrusive systems, which do not pose any additional workload (and frustration) to the drivers.

TACTILE FEEDBACK DEVICES

In addition to the prevalent choice of designing *visual* fuel economy displays, there has been some work towards using other modalities for eco-feedback (also cp. Figure 9.1). In a study by Staubach, Kassner, and Fricke (2012), three different interfaces were designed to advise drivers on acceleration and gear-shifting: a purely visual one, a haptic one employing an “active” acceleration pedal able to produce a certain counterforce, and a visual-haptic one. The designs were evaluated in terms of user acceptance and their impact on driving behavior, showing that although the study participants preferred the visual feedback, as it was less salient than the other two systems, it also had the worst results concerning objective performance measures, and the participants felt more distracted when using this interface. Concerning the low appreciation of the haptic interfaces, the authors hypothesize that the participants might have felt patronized by the system, i.e. it didn’t allow drivers to move the pedal as freely as would normally be the case.

AUDITORY FEEDBACK

Although to our knowledge there do not exist any studies specifically evaluating auditory feedback for supporting fuel efficient driving, there has been some work studying the influence of sounds coming from the engine of a car: In a driving-simulator study, Hellier, Naweed, Walker, Husband, and Edworthy (2011) evaluated how different levels of engine noise affect the driving style and perceived comfort. They found out that low levels of engine noise led to increases in driving speed and more traffic violations. Surprisingly, the low-noise feedback conditions were also associated with a decrease in driver comfort. The authors conclude that auditory feedback plays a major part in the ability of a driver to make judgements and choices about speed.

Similarly, Horswill and Plooy (2008) conducted a study with seven participants judging if a video-based driving scene where in-car noises were reduced in volume by 5 dB appeared to be faster or slower than the same driving scene with a realistic level of sound. The authors observed that the decreased level of sound led participants to consistently judge the speed as slower than the one with a normal level, i.e. a reduction in noise made vehicles appear slower than they actually were.

9.2.4. Auditory In-Vehicle Interaction

While auditory interaction has long played a minor role for use in automobiles, there has recently been considerable work towards using the auditory modality for in-vehicle interaction.

IN-CAR ENTERTAINMENT SYSTEMS AND MENU NAVIGATION

In order to simulate a dual-task scenario that would also occur when a driver had to interact with an in-car entertainment system, Jeon, Davison, Nees, Wilson, and Walker

(2009) conducted a study where 24 participants had to navigate through an alphabetized list of 150 song titles, while they played a simple visual ball-catching game on a computer. The menu was presented either with no sound or one of five combinations of text-to-speech (TTS) audio, spearcons (i.e. highly sped up spoken phrases; cp. Section 5.1), and spindex (an auditory cue based on the pronunciation of the first letter of a menu item (Jeon & Walker, 2009)). The authors found that all trials with additional auditory cues reduced the number of errors for the menu navigation task, with two conditions (TTS and spindex+TTS) also performing better in terms of speed of execution. Additionally, the performance of the primary task (i.e. playing the game) improved for all auditory feedback conditions, as compared to the non-audio one. Finally, participants perceived their workload to be lower with auditory cues, especially for the non-speech audio conditions.

In a follow-up study, Gable, Walker, Moses, and Chitloor (2013) evaluated eye tracking data, performance, workload, and user preferences in a study, where participants had to perform a list search on a cellphone, while doing a lane change task in a car simulator. With similar auditory enhancement as in the previous study, the authors showed that the spindex+TTS cue not only allowed a significantly longer visual fixation on the driving task than the visual-only condition, but was also associated with a decreased lane deviation. In conclusion, although designed only for a very special type of tasks, these studies show that auditory displays, used individually or to enhance existing interaction paradigms, have a high potential of improving the interaction of drivers with in-vehicle systems, especially in terms of both subjective and measurable impact on workload.

SKILL ACQUISITION

In a different context, Powell and Lumsden (2015) developed and evaluated an auditory display to support motorsport drivers in improving their racing skills. Based on a novel target matching design, the drivers were provided with tonal feedback on the lateral G-force on the one ear and with the target G-force (representing the limit of the car) on the other ear. Evaluation of the design showed promising first results, with the greatest efficacy of the system, when learning new tracks or familiarizing with new cars. Responses from the study participants also indicated a positive influence on the drivers' confidence.

AUDITORY WARNINGS

In a study conducted by Ho, Spence, and Gray (2013), different presentations of "looming" auditory warning signals (i.e. signals whose intensity increases over time) were compared with vibrotactile feedback. Contrary to the initial hypothesis, the authors discovered that the vibrotactile warning signals did not offer any benefits over the auditory ones and conclude that auditory warnings increasing in intensity (e.g. as a

function of the time to collision) represent a particularly promising means of alerting a driver in safety-critical situations.

Finally, Larsson and Västfjäll (2013) explored the design of in-vehicle auditory displays with regards to emotional reactions: They conducted a study with 30 participants in a simulator environment with several more or less imminent collision scenarios and evaluated the effect of auditory icons (cp. Section 5.1) in comparison to abstract earcon sounds. In the experiment, the auditory icons were perceived by the participants to be more activating and urgent and also resulted in faster response times (i.e. a quicker brake reaction) than the earcons.

In conclusion, this review of literature shows not only the high potential of measures to support eco-driving, but also the clear benefits of using the auditory modality for doing so.

9.3. General Design Approaches

While current scientific knowledge supports the notion of an auditory fuel economy display¹, the specific design of such a system is something that clearly needs further research. Since the possibilities for the actual realization are quite diverse, our first step has been to structure the design space by developing five general design approaches to guide the development process of an auditory fuel economy system. Note that these approaches are not mutually exclusive, but can also be combined to obtain a comprehensive display design.

9.3.1. Direct Sonification of Fuel Consumption

We see the first and most straightforward approach to designing an auditory fuel economy system in directly sonifying the available data on current fuel consumption. Conceptually, this can be seen as a direct translation of existing visual fuel efficiency displays to the auditory domain. Nevertheless, this approach raises several design questions that need to be considered and which ultimately determine if the display can convey the intended information and if it will be accepted by potential users: First, it must be decided if the display should provide a continuous sonification, allowing the driver to assess the fuel consumption at all times, or an event-based one, which is only emergent for a specific incident or situation, for example after consuming a specific amount of fuel, when the driver manages to achieve a comparably good fuel efficiency, or, conversely, when consumption is unusually high. For a continuous sonification, the main question is then how the data is mapped onto the characteristics of a sound, e.g. if

¹Although strictly speaking, a *fuel economy display* would only be used for cars with a conventional combustion engine, our use of the term also implies the inclusion of systems providing feedback on energy consumption in electric and hybrid cars.

pitch, brightness, and/or loudness should be modified based on the current fuel consumption. Finally, a concept like blended sonification (Section 5.2) should be applied in order to keep the sound as unobtrusive as possible.

9.3.2. Sonification of Secondary Parameters

As a second design approach, the sonification of “secondary” parameters can also be used to support drivers in attaining a fuel efficient driving style. More specifically, this includes all aspects of driving that are known to affect the fuel consumption of a car. An aggressive driving style, for example, is characterized by phases of quick accelerations followed by unnecessarily hard or too frequent braking, and an auditory fuel efficiency system could create an awareness for the negative effects of these behavior patterns. Another secondary parameter is the number of revolutions of the engine: Simply by keeping those low, users can implement a smoother and more fuel-efficient driving style.

9.3.3. Gamification

Creating an awareness for the car’s fuel consumption and providing feedback on how it is affected by driving behavior can be considered a necessary or at least very important part in achieving wide-spread adoption of eco-driving habits. However, the intrinsic motivation of some users might not be enough to pursue this goal for a longer period of time. Our third design approach therefore suggests to use gamification as a conceptual framework to keep an auditory display engaging. Gamification is an emerging area of research that deals with improving user experience and user engagement by using game elements in non-gaming contexts, and has shown to be able to motivate people over a longer period of time, most prominently for sports and health applications (cp. Deterding, Sicart, Nacke, O’Hara, & Dixon, 2011). However, to the best of our knowledge, the concepts of gamification have yet to be applied to an (exclusively) auditory display: Having originated from a video game context, most of the design principles of gamification do not seem to be directly applicable for an auditory display and would need to be adapted for the auditory domain. Nevertheless, we think that this is a promising research direction, which could also be explored for the design of an auditory fuel economy display.

9.3.4. Sonifying Advanced Support Information

A further design approach is to give information on *external* factors that influence driving performance, such as how far away an upcoming stop sign is, how much time a red traffic light will take to switch to green, or how fast the vehicle ahead is driving. Although already valuable in its own way, we think that such information is especially suitable to benefit and support users in driving in a more anticipatory and thus fuel efficient way. Also, it usually becomes most relevant in safety-critical situations, e.g. when

approaching a crossing or traffic lights, so that an auditory display would arguably be the best choice to convey this information, considering that it can avoid to divert the driver's attention away from the street (cp. Section 9.2.2).

9.3.5. Supporting Specific Qualities of Driving

Finally, supporting certain qualities that are associated with car driving could be a subliminal way to induce a more fuel economic driving style. For example, based on the assumption that for some people there is a certain desire for fast driving, we could strive to convey the *feeling* of speed and thereby reduce the need to actually *drive* fast. Although this approach certainly comprises a wider range of aspects of car design, we think that sound can play a major role in achieving such a feeling, e.g. by modifying the engine sound towards a more "sportive" one.

9.4. Sonification Designs

Based on the design approaches discussed in the previous section, we have developed a range of prototype auditory displays, two of which are discussed here in more detail and have also been evaluated in the study described in the following sections. Although the presented approaches can guide the development of quite elaborate support systems, we have decided to first concentrate on relatively basic designs based on the first two approaches, as those are better comparable to the existing visual counterparts.

9.4.1. Continuous Sonification of Fuel Consumption

Our first prototype auditory display is based on the approach of directly sonifying the fuel consumption: Similar to the pitch-based sonification of the InfoDrops system (cp. Section 8.4.1), the basic idea of this design is to map the instantaneous, relative consumption (i.e. liters per 100 km) to the frequency of a bandpass filter that is applied to a broadband noise signal, resulting in a high-resolution representation of the input data, which should at the same time be rather unobtrusive to the driver. Since the perceived loudness of the output signal obviously depends on its center frequency, we applied a basic psychoacoustic amplitude compensation to the generated sound so that a low-frequency sound is perceived approximately as loud as a high frequency one². Furthermore, the output level is additionally adjusted based on the loudness of the engine sound, meaning that a louder engine also leads to an increase of the sonification's output level.

²For amplitude compensation, we used Supercollider's AmpComp filter.

9.4.2. Metaphorical Sonification to Support Low Rpm Driving

Based on our second design approach (Section 9.3.2), we have furthermore developed an auditory display that gives feedback on the rpm³ range that is best for a fuel efficient operation of the engine: Although the driver generally has less ability for acceleration when driving with a low rpm as the engine is already operating with an increased torque in that case, the internal friction of the engine inevitably increases with higher rpm, which consequently leads to higher energy losses. This is also reflected in the fuel consumption map of a typical engine (cp. Figure 9.2a), which shows that fuel efficiency typically increases with lower rpm. To make this relation audible, we use an event-based sonification that is triggered when the driver leaves an “optimal” rpm range (see Figure 9.2b). In our current design, this range is defined based on the engine parameters, the selected gear, and the current slope of the street. Additionally, for use over a longer time, it could also be adjusted depending on the previous eco-driving performance of the user.

Based on preliminary work by Larsson and Västfjäll (2013), we have chosen to use a slurping-like sound as an indicator for increased fuel consumption, which should not only be easily recognizable, but also elicit an emotional response (cp. RQ8) and consequently be perceived as more activating and convincing than an abstract auditory indicator (also cp. Section 9.2.4). Furthermore, depending on how much the threshold is exceeded, the loudness of the auditory icon is gradually increased. Finally, similar to the continuous sonification described above, we compensate for the current loudness of the engine sound.

9.5. The EcoSonic Driving Simulator

In order to explore and assess the possibilities of auditory displays supporting fuel-efficient driving, we have developed the EcoSonic driving simulator for reproducible basic research. It is written in C++ and includes a detailed model of the engine and the car to simulate a physically correct driving behavior and to determine a realistic approximation of the instantaneous and accumulated fuel consumption. The system it runs on is an ordinary computer and can be operated with a steering wheel as well as two pedals for throttle control and braking. Gear selection is done with two buttons located on the steering wheel. To assess the effect of a specific fuel economy display, another component of the system logs the user’s actively-made input to the simulator, such as the pedal positions and selected gear, as well as several simulated parameters, including the car’s instantaneous fuel consumption. Additionally, the system captures gaze information with the help of an eye tracker that is attached to the simulator’s main

³the engine’s revolutions per minute

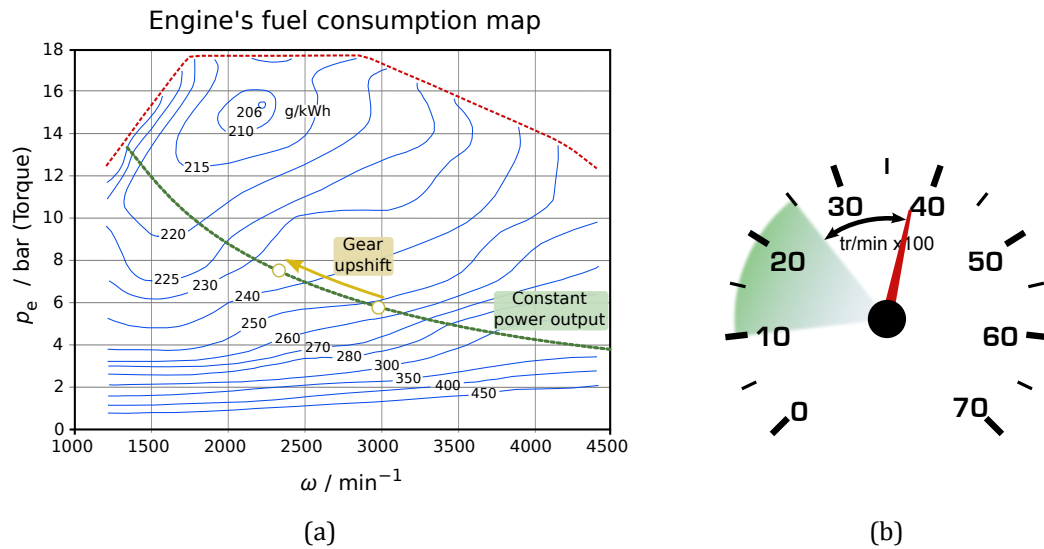


Figure 9.2.: **(a)** Exemplary representation of a fuel consumption map depending on number of revolutions and torque. The green line illustrates a typical level of constant power output, i.e. moving on this line does not change the effective thrust of the engine. We can easily see that altering rpm by shifting gears has a significant influence on the fuel consumption of the engine. (Original picture from Wikipedia, 2009) **(b)** Illustration of the basic idea of the metaphorical sonification display concept (Section 9.4.2).

display. Via an integrated OSC interface,⁴ the individual auditory fuel economy displays can be prototyped and designed with the help of an external sound synthesis tool like Pyo.⁵ For the implementation of the previously discussed sonification designs we have used Supercollider,⁶ which is also used to synthesize the car's engine noise. Finally, the simulator's track editor can be used to design individual test routes for studying specific aspects of driving.

9.5.1. Graphical User Interface

One of the first choices we had to make was how to visually convey the actual process of driving. There are basically three alternatives: (a) a first-person perspective, i.e. looking through the windshield of a car, (b) a view from above/behind the car, and (c) a side-view of the car. The first option is potentially the most realistic one, as this is how we normally experience driving a car. However, depending on the specific graphics as well as several other important factors (e.g. sound or input devices), the gap to realism can still be quite large. Furthermore, it can be difficult for users to overview the

⁴OSC: Open Sound Control (<http://opensoundcontrol.org>)

⁵Pyo: Python module for digital signal processing (<http://ajaxsoundstudio.com/software/pyo/>)

⁶Supercollider: A real-time audio synthesis language (<http://supercollider.github.io/>)

overall driving setting, such as the slope of a street or upcoming signs and traffic lights, especially if it is the first time to drive a specific test route. The second option (b) is best known from computer games, as it gives the viewer a better overview of the track (e.g. the curvature of the street), which is important for driving routes without knowing them beforehand. However, this is also the reason, why this view gives more the impression of an arcade game than of a driving simulator.

A SIDE-VIEW PERSPECTIVE

Finally, (c) is a rather abstract view on the track. The driving experience is insofar “reduced” as the users do not have to steer to keep on the street, which makes it a bad choice if the steering itself is of greater importance for the simulation. However, street signs and traffic lights can be seen from a distance and also the slope of the track is easily perceivable. As for our usage scenario, the important aspect of driving is the acceleration behavior (i.e. not the steering of the car), we decided to implement such a side-view perspective for graphically representing the driving experience. This choice also has the advantage that we are better able to reproduce the scenario of a person driving a track on a regular basis: In this case, the details of the route are relatively well known to this person, which is reflected by the high visibility range of the side-view perspective.

Taken as a whole, the graphical user interface of the simulator consists of a track view and a view of the dashboard. The track view (Figure 9.3a) basically uses only four elements to visually display the driving scenario: A simple line represents the street, which, at the same time, can be seen as the height profile of the track. The street moves beneath the car, which is always at the left side of the screen, resulting in a constant visibility range. In order to give the user a sense of the speed of the car, several trees move in a virtual plane very near the camera, realizing a parallax scrolling effect. Depending on the car’s speed, the images of the trees are filtered with a certain amount of motion blur, which adds to the impression of speed. Finally, there are several types of street signs (e.g. stop and speed signs) as well as traffic lights that can be placed along the track to model the characteristics of a specific route.

The dashboard consists of the typical elements that can be found in most cars currently in use, i.e. a speedometer and a display for the revolutions per minute of the engine (Figure 9.3b). Additionally, the instantaneous as well as the accumulated fuel consumption are displayed in a textual representation. As for our study most of the participants could be expected to be German, we used the in Europe more commonly used $L/100km$ as a unit for the instantaneous consumption. However, this is easily switchable to other units, like for example miles per gallon (MPG), which is more prevalent in English-speaking countries. Lastly, an indicator for the currently selected gear is displayed next to the dashboard. Note that this part of the user interface can either be




Figure 9.3.: **(a)** Screenshot of the driving simulator. Due to the motion blur of the tree (which is rather near to the virtual camera), even in a still picture the movement of the car is perceivable. Also, the height profile as well as approaching traffic signs are clearly visible. **(b)** The dashboard of the car, which is displayed beneath the track view: The main elements are a display for the revolutions per minute on the left and a speedometer on the right. The instantaneous and accumulated fuel consumption are displayed in-between. On the right, there is an indicator for the currently selected gear.

displayed below the track view on the same screen or on a separate display, which can then be positioned independently, e.g. to simulate two viewing planes.

9.5.2. Internal Simulation

The internal simulation of the car includes the computation of the car's movement depending on the user's input (e.g. throttle control and gear selection) and environmental factors (e.g. slope of the street). Of particular importance for the EcoSonic system is furthermore the calculation of fuel consumption, i.e. how much fuel is consumed depending on the torque and number of revolutions of the engine.

ENGINE MODEL

The basic function of a car engine is to produce a certain rotary force (i.e. torque τ), which, through the gearbox, drives the wheels of the car and thereby moves the vehicle. The power output of the engine then is $P = \tau \cdot \omega$, where ω is the angular velocity (which correlates to the number of revolutions). For an internal combustion engine, the torque output is considered to be usable only over a limited range of rotational speeds ω , typically between 1000 rpm and 6000 rpm. But even within this range, output torque varies not only based on user-adjusted throttle, but also depending on the rotational speed, and reaches its maximum only at a specific (range of) ω . In order to freely model the behavior of the engine, our simulator uses a torque map, which makes it possible to define several torque responses for different throttle positions (cp. Figure 9.4). The exact torque output is then calculated by linearly interpolating between the previously defined torque response curves.




Figure 9.4.: A typical torque map of an engine (Varbanov, 2013). Each of the differently colored curves represents a specific throttle position. Contrary to what one could expect, the output torque varies significantly even for the same throttle position, depending on the engine's speed ω .

While such a torque map allows for a quite detailed characterization of an engine, it does not account for energy losses, which naturally occur due to mechanical friction (e.g. of the crankshaft or the pistons) or pumping losses (i.e. the work required to move air into and out of the cylinders). Taken together, they are slightly more than linearly proportional to the engine speed and can be approximated by $\tau_L = L_f \cdot \omega^{1+L_e}$, where L_f is a coefficient determining the general extent of energy losses and L_e a (very small) number indicating a slightly exponential dependence on ω .

FUEL CONSUMPTION

A crucial aspect for the evaluation of an auditory fuel economy display in a simulated environment is of course the fuel consumption of the engine. The consumption primarily depends on the engine's power output P , and a first approximation of the instantaneous fuel consumption would be $C_{\text{fuel}} = c_E \cdot P$, where c_E is an engine-dependent coefficient. In reality, however, c_E also depends on the current speed and torque of the engine, so that

$$C_{\text{fuel}} = c_E(\tau, \omega) \cdot P, \quad (9.1)$$

where $c_E(\tau, \omega)$ is an engine-specific function that can be characterized by a fuel consumption map as shown in Figure 9.2a, which is also used internally by our driving simulator in order to model the engine's consumption characteristics.

Another aspect that can have a considerable influence on fuel consumption is the *over-run fuel cutoff*: Most modern cars automatically cut off the fuel supply when the driver does not press the gas pedal, the clutch is still engaged, and the car still has a certain speed, i.e. is able to keep the engine moving with its kinetic energy. For our simulator, this means that if $\omega > \omega_{\text{min}}$, the engine's torque output τ can actually be zero. Otherwise, a minimum fuel supply must be kept in order to prevent the engine from shutting

down. This behavior becomes most relevant when a driver for example approaches a stop sign and uses the engine to brake, as opposed to using the car's friction brakes.

SIMULATION OF THE CAR

Although the engine is certainly the most important part of a car to simulate, modeling the car itself is of course also crucial for simulating its movement. In this context, the car's mass m is a central attribute, as it plays a direct role in calculating the acceleration $a = F/m$, where

$$F = F_w - (F_r + F_u + F_d) \quad (9.2)$$

is the forward force of the car.

- Here, F_w is the force that is transmitted from the engine to the wheels and is dependent on the selected gear, the transmission, and the radius of the wheels.
- F_r is the rolling resistance force, which can be approximated by

$$F_r = C_{rr} \cdot \cos(\alpha) \cdot m \cdot g, \quad (9.3)$$

where C_{rr} is a rolling resistance coefficient, α specifies the current slope of the street, and g is the gravitational acceleration.

- The uphill resistance is only dependent on the street's slope, i.e. $F_u = \sin(\alpha) \cdot g$, and can become positive when going downhill (i.e. when α is negative)
- Finally, the air drag of the car is proportional to the squared speed v of the vehicle:

$$F_d = \frac{1}{2} \rho \cdot C_d \cdot A \cdot v^2, \quad (9.4)$$

where ρ is the air density, C_d the drag coefficient depending on the shape of the car, and A is the frontal area of the vehicle.

9.5.3. Traffic Violations and Attention Task

Besides the simulation of the car, the driving environment is another important aspect to consider. In addition to the track itself, i.e. the elevation of the street, the environment primarily consists of street signs such as speed limits, stop signs, and traffic lights. Each of these signs introduces a specific rule for the driver, which can, of course, be violated. Besides informing the user about such traffic violations, e.g. driving too fast or ignoring a stop sign, these incidents are also logged by the system, as they are an indication for how much attention the user is able to devote to the street.

Contrary to conducting real world studies, where distractions and needed attention are basically unpredictable, having a completely simulated environment enables us to introduce a fully controlled attention task, which is always the same for each study




Figure 9.5.: **(a)** Screenshot of the track editor of the EcoSonic driving simulator. The gray and white points represent the end and control points of the individual Bézier curves. The turn-left sign represents an attention task, where the iconic road is making a gradual left turn. The trigger distance of the traffic light is indicated by the vertical line on the left. **(b)** Two images of the iconic road, which is displayed above the track view as part of the attention task. In case of the lower image, the user would have to steer to the left.

participant. As the side-view scenario does not require the user to steer, we implemented a scripted steering task which requires the user to react to simulated curves: At the center of the track view, an iconic representation of a street is displayed, which can gradually turn into a curve, prompting the user to steer in the indicated direction (see Figure 9.5b).

9.5.4. Track Editor and OSC Communication

In order to make it possible to quickly design a specific driving scenario, the simulator includes a dedicated track editor, which allows the modification of the street's elevation profile as a cubic Bézier spline, where the two control points next to a curve's endpoint are kept collinear in order to maintain C1 continuity (cp. Figure 9.5a). Additionally, street signs and traffic lights as well as attention tasks can be placed along the road. Traffic lights by default start showing a red light and are triggered to switch to green when the user is approaching them at a specifiable distance. The exact moment of the switch is then determined either by an explicitly specified or a random delay time. Note that in our study, only traffic lights with a fixed delay time were used in order to keep the results as comparable as possible. For the attention tasks, the amount of added curvature as well as the duration of the curve (i.e. how long the modification of curvature will take) can be specified.

As the primary use of the driving simulator is to evaluate *auditory* fuel economy displays, it is mandatory to have an interface that allows for the prototyping of sonification design with the help of an external sound synthesis program. Since the majority of those programs are able to communicate via the specialized Open Sound Control




Figure 9.6.: Hardware setup of the EcoSonic driving simulator.

(OSC) protocol, the simulator offers an OSC-based interface that broadcasts the instantaneous fuel consumption as well as regular messages indicating the expenditure of a specific amount of fuel. Additionally, the interface is used for the synthesis of the engine sound and for that provides real-time updates on engine speed and the currently provided torque.

9.5.5. Hardware Setup

The required hardware for the EcoSonic driving simulator consists of only a few components and is therefore easily installed. For our setup, a Mac Pro was running the driving simulator, displaying the user interface on a 24-inch monitor. Attached to the monitor was an eye scanner to capture real-time gaze data of the user. For our purposes, we decided to use an *Tobii EyeX* tracker⁷, as the accuracy and response time are sufficient and the calibration procedure is quite fast. Given that it can only track gaze information for one screen, the dashboard was not shown on a separate monitor, but below the track view on the main display. Furthermore, we placed a pair of loudspeakers⁸ next to the monitor, as can be seen in Figure 9.6. In front of the display, a steering wheel was attached to the table and pedals for throttle control and braking were placed on the

⁷<http://www.tobii.com/en/eye-experience/>

⁸We used a pair of the Genelec 8020C.

floor.⁹ Users were sitting approximately 60 cm away from the monitor for the best eye tracking performance and so they could comfortably operate the steering wheel.

9.6. Study Design

With the help of the EcoSonic driving simulator, we conducted a study to evaluate the auditory fuel efficiency displays presented in Section 9.4. In order to reduce the number of necessary participants, we used a within-subject design, where for each condition the participants had to drive the same test track four times in order for them to familiarize with the respective display. Controlling for ordering effects between the different conditions, we employed a counterbalanced measures design, where all possible sequences of conditions (i.e. all 6 permutations) were evenly distributed over the study participants. In total, 30 people took part in the experiment.

At the beginning, each participant was given a short introduction, explaining the purpose of the experiment and a short questionnaire dealing with general questions about personal attitudes and preferences as well as previous experiences. Subsequently, the EcoSonic system, including the various elements of the driving simulator, were explained and the participants were told what they had to do during the experiment. Specifically, the participants were given four tasks:

- First, they should follow the universally known traffic rules. More specifically, they were told to adhere to the respective speed limits, to stop in front of a stop sign and not to drive through a red light. If they committed one of these traffic violations, they were notified about this incidence by a sound of a camera taking a picture in conjunction with an implied flash (the screen would go white and quickly fade back to normal), in order to give the impression of being caught speeding. Additionally, a short text would say, which traffic rule had been violated.
- Also, they were instructed not to obstruct the following traffic, i.e. the participants should not drive *too* slowly, depending on the current speed limit. If they did, this was indicated by the honking sound of a car, combined with an appropriate text message.
- Additionally, the participants should “keep on the street”. This referred to the attention task, discussed in Section 9.5.3, which was also explained to them in more detail. More specifically, they were instructed to keep the curve-indicator in a neutral position (also cp. Figure 9.5b).
- Finally, the participants should try to keep the fuel consumption as low as possible. It should be noted, however, that they were not told how to do so.

⁹For our setup, we used a consumer-grade steering wheel (Logitech Wingman Formula GP), which also has pedals included.

Before the actual experiment, the participants were given the possibility to get used to the driving simulator on a separate short test track, which included all important elements in a compact manner. Then, for each condition, the respective fuel efficiency display was shortly explained, after which they could independently conduct the four runs. Afterwards, individual questionnaires dealing with specific questions about the particular display were handed out. At the end of the experiment, the participants completed a final questionnaire with comparative questions and asking for feedback about fuel efficiency displays in general and the experiment as a whole. Altogether, the study lasted around 45-50 minutes for each participant, primarily depending on the time needed to complete the questionnaires.

9.6.1. Goals and Hypotheses

The primary goal of the experiment was to compare the two designs for an auditory fuel efficiency display to the well-established visual display found in most cars these days. More specifically, we wanted to find out, how they affect the driving behavior and fuel consumption, if there are any differences in terms of how much they distract the driver from the primary driving task, and how they are perceived by potential users.

Perception. As already discussed in Section 9.4.2, we hypothesized that the metaphorical sonification (MS) would be perceived as more affective (H1.1) and, in consequence, more convincing and activating (H1.2). Furthermore, as the auditory displays effectively add functionality to the driving experience, we expected them to be assessed as more helpful than the visual display alone (H1.3). On the other hand, based on previous experience with auditory displays, we also saw the danger of them being slightly annoying, i.e. less likely to be accepted as a part of an in-vehicle interface (H1.4).

Behavior. In terms of quantifiable influence on the driving behavior, we firstly expected the fuel consumption to decrease over the course of the experiment due to participants familiarizing with the test track and the driving simulator (H2.1). Furthermore, we assumed that the consumption would be lower when participants were supported by one of the auditory fuel efficiency displays (H2.2). Specifically for the continuous sonification (CS), we expected users to better be able to estimate the instantaneous consumption (H2.3) and to look less often at the corresponding visual display (H2.4). The metaphorical sonification was expected to support drivers in employing lower engine speeds in comparison to the other conditions (H2.5) and to cause less glances to the rpm display (H2.6).

Distraction. Based on modern attentional-resource theories, as discussed in Section 2.3.3, we hypothesized that the auditory displays might enable participants to perceive the presented information on a more subconscious level (H3.1). Additionally, we anticipated the displays to be perceived as less distracting and more unobtrusive

(H3.2), and in consequence to alleviate mental overload while driving (H3.3). Finally, these attributes were expected to provoke fewer traffic violations (H3.4) and to allow the drivers to focus more on the street (H3.5), thus enabling them to better “keep on the road” for the steering task (H3.6).

9.7. Results

9.7.1. Questionnaires

Based on the previously defined hypotheses, we designed several Likert-type scales consisting of up to 4 questions, which had to be answered either after each condition or at the end of the experiment as a comparative question. The majority of them were posed as statements the participants should indicate their agreement to on a scale ranging from “Not at all” (−3) to “Fully agree” (3). In the subsequent analysis, we only considered composite scales with a Cronbach’s alpha > 0.7 in order to only allow question groups that had been consistently answered. Furthermore, a Friedman test was used to determine if there are any significant differences between the three conditions. Then, for comparison between the individual conditions we used the Wilcoxon signed rank test to check if the differences between two scales are statistically significant. When testing for differences that were not hypothesis-driven, the Benjamini-Hochberg False Discovery Rate method was used to adjust the resulting p-values.

COMPREHENSION

The majority of study participants stated to have very well understood both the visual fuel efficiency display (VD) and the MS ($\mu_{\frac{1}{2}} = 3$ for both conditions). The continuous sonification was also well understood ($\mu_{\frac{1}{2}} = 2$), although the data shows a significant difference to the two other conditions ($p < 0.02$). When asked about being able to estimate the current level of fuel consumption, the two auditory conditions were rated only insignificantly better ($\mu_{\frac{1}{2}} = 1$) than the VD ($\mu_{\frac{1}{2}} = 0.5$), with Friedman $p = 0.24$, thus no statement can be made for H2.3.

AFFECTIVE RESPONSE AND HELPFULNESS

As expected (H1.1), the metaphorical sonification ($\mu_{\frac{1}{2}} = 0$) was perceived as significantly more affective than the visual display ($\mu_{\frac{1}{2}} = -3$; $p < 0.01$) and the CS ($\mu_{\frac{1}{2}} = -2$). Similarly, the participants assessed the MS as more convincing ($\mu_{\frac{1}{2}} = 2$) than the two other conditions (both $\mu_{\frac{1}{2}} = 0$, with $p < 0.01$), confirming H1.2. Although the correlation between these scales is surprisingly low ($\tau = 0.21$), it is still significant ($p < 0.01$). In terms of helpfulness, the MS ($\mu_{\frac{1}{2}} = 2$) was rated significantly better than both the visual display and the CS ($\mu_{\frac{1}{2}} = 0$) (Figure 9.7a), partly verifying H1.3.

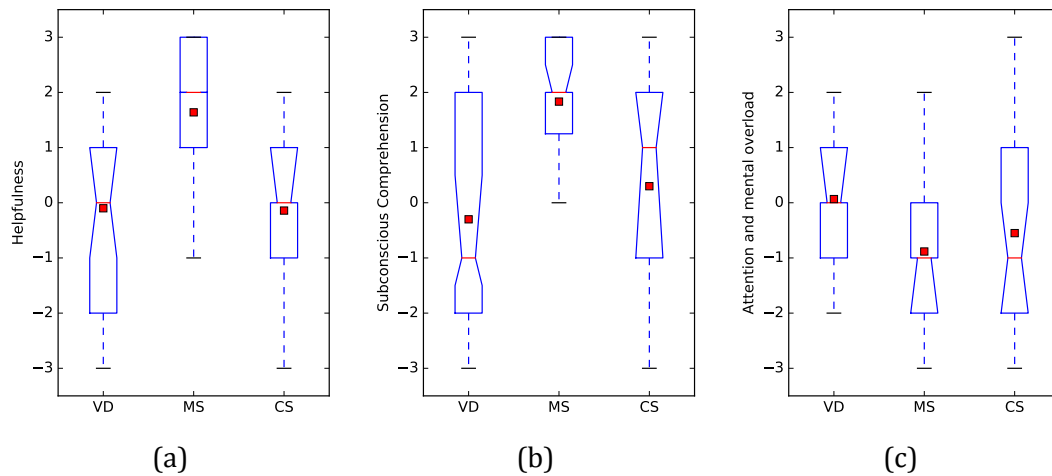


Figure 9.7.: Overview of some of the participants' assessments of the three fuel efficiency displays, namely **(a)** the perceived helpfulness, **(b)** if they had the impression to be able to subconsciously comprehend the provided information, and **(c)** if they felt distracted or overwhelmed when using the respective display. The notches represent the 95% confidence intervals of the median.

DISTRACTION AND ATTENTION

We also asked the participants to assess the unobtrusiveness of the different fuel efficiency displays, i.e. if they had been distracted by them while driving. Here, the continuous sonification ($\mu_{\frac{1}{2}} = 1$) was rated significantly better than the VD ($\mu_{\frac{1}{2}} = -0.5$, with $p = 0.02$). The metaphorical sonification was rated even higher ($\mu_{\frac{1}{2}} = 2$, thereby confirming H3.2), although the differences are significant only for the comparison to the visual display ($p < 0.01$). Similar results were obtained for the question if the participants felt that they had to concentrate on the fuel efficiency feedback or if they had been able to subconsciously absorb the presented information (cp. Figure 9.7b). Both CS and MS were rated significantly better ($p < 0.01$) than the VD, confirming H3.1. Again, the similarity between the scales can also be seen in the correlation between the corresponding variables ($\tau = 0.78$, $p < 0.01$).

The differences in terms of unobtrusiveness and subconscious comprehension may also explain, why the participants stated that they felt overwhelmed by the driving task and had difficulties in keeping their attention on the street significantly more for the visual condition ($\mu_{\frac{1}{2}} = 0$) than for the auditory ones (both $\mu_{\frac{1}{2}} = -1$, with $p \leq 0.03$), thereby supporting H3.3 (also cp. Figure 9.7c).

ACCEPTANCE

Contrary to what we expected, with regards to the acceptance of the respective designs, e.g. if they thought that they would be annoyed by them while driving, participants




Figure 9.8.: (a) Scatter plot showing the fuel consumption in dependence of the time the participants needed to complete the test track. Additionally, a least squares linear approximation of the data is overlaid. (b) A boxplot comparing how much fuel was consumed for the three conditions over the course of the four runs. The continuous lines connect the average values for each respective run.

rated the metaphorical sonification ($\mu_{\frac{1}{2}} = 2$) even slightly better than the VD ($\mu_{\frac{1}{2}} = 1$), and also than the CS ($\mu_{\frac{1}{2}} = -0.5$). Although the differences between the conditions were not significant (Friedman $p > 0.08$) and thus no statement can be made for H1.4, we see these results as a clear indication that auditory interfaces can be designed to appeal to a larger audience. In terms of influencing factors, we have found that the understandability ($\tau = 0.58$) as well as the perceived helpfulness and unobtrusiveness (both $\tau = 0.6$) of a fuel efficiency display were highly correlated with its acceptance.

9.7.2. Measured Data

FUEL CONSUMPTION

In Figure 9.8a we can see the total fuel consumption plotted against the time needed to complete the whole test track, which gives a rough impression on how these two variables depend on each other. Additionally, we performed a least squares fit of a linear function to this data. As could be expected, the participants that were faster consumed on average a little more fuel than the ones taking their time. However, the dependence is surprisingly low (Pearson correlation coefficient: -0.04) and statistically not significant ($p = 0.44$). These findings are in line with previous work examining the potential to reduce fuel consumption without a significant impact on trip time (Evans, 1979).

Comparing the measures of fuel consumption over the course of the four runs, we can see that for each condition the participants consistently improved their fuel efficiency (Figure 9.8b), which is consistent with our hypothesis H2.1. Due to the high variance of

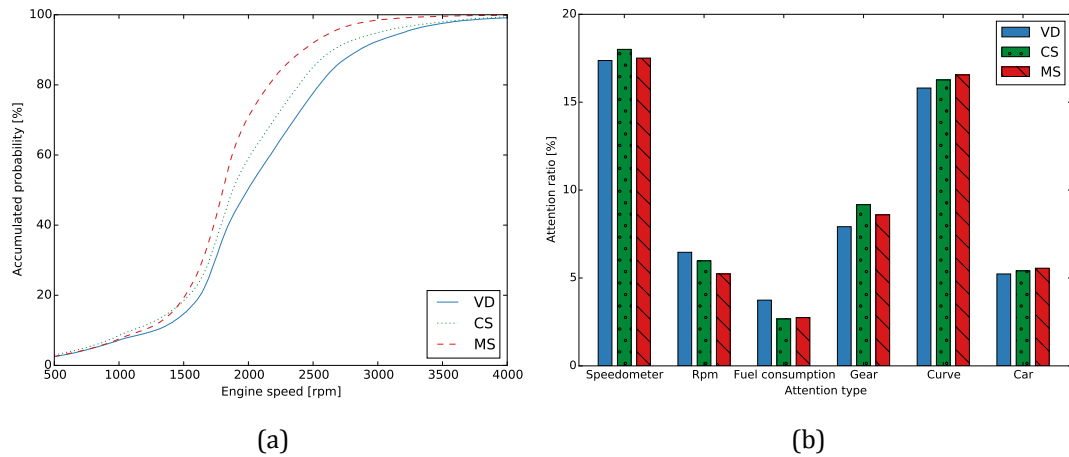


Figure 9.9.: **(a)** Cumulative distribution function of the average engine speed distribution for the different conditions. **(b)** Attention ratios for a selection of points of interest.

the data, however, the differences between the first and the last run are significant only for the MS ($p = 0.02$). Furthermore, the average fuel consumption differs between the three conditions (Friedman $p < 0.01$). Confirming H2.2, it is significantly higher for the visual condition ($1.37\text{dl} \pm 0.25$) when compared to the auditory ones ($p < 0.01$), but also lower for the MS ($1.22\text{dl} \pm 0.12$) than for the CS ($1.30\text{dl} \pm 0.22$) ($p < 0.05$), establishing a clear order of $\text{MS} < \text{CS} < \text{VD}$.

ENGINE SPEED DISTRIBUTION

Analyzing the participants' driving behavior, Figure 9.9a illustrates how much time during the test track the car was driven at a certain engine speed. We can clearly see that participants driving with the MS used the lowest engine speeds, which can also be statistically confirmed by comparing the respective average engine speeds for the whole track ($\text{MS} (1767\text{rpm} \pm 307) < \text{VD} (1994\text{rpm} \pm 323)$ and $\text{CS} (1889\text{rpm} \pm 307)$, both $p < 0.01$). This finding confirms the effectiveness of the metaphorical sonification, which was specifically designed to support driving with a lower number of revolutions, verifying our hypothesis H2.5. Furthermore, we can see that sonifying the fuel consumption (CS) enabled the drivers to reduce the average engine speed as well, when compared to the visual display, hinting at an increased effectiveness of the auditory representation of this data.

TRAFFIC VIOLATIONS AND STEERING TASK

On average, the participants committed 0.52 traffic violations for the whole test track, which can be considered quite few. Moreover, the differences between the individual runs were quite high ($SD = 0.72$), making it difficult to establish any statistical

statements. Thus, while it could be observed that drivers supported by the CS committed marginally fewer traffic violations (0.44 vs 0.58 for the VD), we could not find any *significant* differences between the three conditions and therefore cannot make any statements concerning hypothesis H3.4.

Similarly, data collected from the steering task contained a high variance that could not be attributed to a specific condition: As a measure for the level of distraction, we analyzed the instantaneous deviation from the neutral curve position $\delta(t)$ and calculated for each run the normalized integral

$$D_{\text{run}} = \left(\int_0^{t_{\text{run}}} |\delta(t)| dt \right) / t_{\text{run}}, \quad \text{with } -1 \leq \delta(t) \leq 1, \quad (9.5)$$

where t_{run} is the total time needed to complete the track, resulting in an indicator D_{run} for the deviation from the neutral curve position over a whole run. Surprisingly, the differences between the conditions are maximally 1% and have a high coefficient of variation ($c_v = 0.58$). Consequently, our hypothesis H3.6 cannot be confirmed. Considering that the mean $\mu_{D_{\text{run}}} = 0.24$ is surprisingly high, we can only hypothesize that the participants paid less attention to this comparably new and unfamiliar task, regardless of the respective condition.

EYE TRACKING DATA

In order to analyze gaze information that was provided by the eye tracker during the experiment, we first determined the part of the driving simulator the participants were looking at for each simulation step, based on a simple area matching algorithm. Then for each item, an attention ratio was calculated based on the accumulated time it was looked at, in relation to the total time of a run. Invalid or unclassifiable data (i.e. gaze points outside the application window) were discarded. In Figure 9.9b, the average attention ratios for each condition are shown.

One first observation we can make is that in general participants looked at the speedometer surprisingly often (i.e. 17.64% of the time), which might be attributed to difficulties in estimating their speed in the simulator environment. Furthermore, although the average attention ratio is comparably low, we can see that both for the MS (2.75% \pm 2.84) and the CS (2.67% \pm 2.44), participants looked less often at the fuel consumption display than for the VD (3.74% \pm 3.29). The differences are significant ($p \leq 0.01$ for both comparisons) and verify our hypothesis H2.4. We can also see that the total time spent looking at the display for the revolutions per minute was lowest for the MS (5.23% \pm 4.84) and significantly less than for the VD (6.46% \pm 5.15) ($p < 0.01$), which confirms hypothesis H2.6. For the CS, it was not as low (5.97% \pm 4.93), but still lower than for the VD ($p = 0.03$). While for the auditory conditions the users seemed to have slightly more time to keep an eye on the curve indicator as well as the car itself, these differences are not significant, preventing any statement on H3.5.

9.8. Discussion and Conclusions

With the EcoSonic system we have created two novel feedback designs to represent consumption-specific information in an unobtrusive way using real-time sonification. In terms of practical value of these displays, one of the main results of the conducted study is that both designs help drivers to significantly reduce their fuel consumption. Moreover, we have seen that this reduction is possible even without a significant impact on trip time.

USE OF THE AUDITORY MODALITY FOR AMBIENT DISPLAYS

The study has furthermore given evidence that the use of the auditory modality can indeed be beneficial for a peripheral display and therefore should be considered a viable alternative to visual ambient interfaces (RQ2). Not only did the two auditory feedback designs lead to a reduced fuel consumption, they were also perceived as more unobtrusive and less distracting than the visual display. This also supports our hypothesis that addressing a modality with a peripheral display that is not overly used to perform a main task is highly conducive to making ambient interfaces less obtrusive and distracting (RQ4). Furthermore, the results of the study indicate that the participants prefer an unobtrusive peripheral display in favor of a more involved one, which also supports our hypothesis that more unobtrusive displays are indeed better accepted by users (RQ5).

THE USE OF AFFECTIVE CUES

Different to the previously discussed projects, the metaphorical sonification used affective cues not to persuade users towards changing their behavior, but rather to *support* them in doing so by aiming to generate a quicker and more natural response in the user, which in turn should be beneficial for them. Comparing the participants' comments and the quantitative results of our study with those from the Sonic Shower, we can say that the reception of this approach in terms of unobtrusiveness, acceptance, and helpfulness is far better than the one aiming to persuade users to change their behavior (RQ8). Our interpretation of this result is that the limited information capacity of ambient displays is simply not enough to actually convince users to change their behavior, which would require a far more involved and information-heavy approach.

COMPLEXITY AND VARIABILITY OF THE CONTEXT

With regards to RQ3, the EcoSonic system is an example for a rather predictable environment, i.e. different to the InfoPlant, where a large variety of activities could take place in the living environment it was installed in, we are dealing with a very clearly defined primary task. Furthermore, in comparison to the Sonic Shower, the users' perceptual situation is even more predictable, as drivers can be expected to sit in their designated seat and most probably look towards the front, through the windshield of

the car. All in all, we see these conditions as particularly well suited for the use of ambient displays, which is also reflected in the results of our study. While for auditory interfaces, this level of invariability of the context would actually not be necessary (also cp. Section 8.6), such an environment should even allow visual ambient interfaces to be used due to the predictability of the user's gaze.

FUTURE WORK

Based on the two auditory feedback designs and the results of the conducted study, there are a few possible directions of future work that we see as worth exploring:

- As in our study the participants had the visual fuel efficiency display available even when being supported by the auditory feedback, it might be interesting to study the potential of an auditory display to effectively *replace* the, at least potentially distracting, visual indicator for fuel consumption, i.e. if a) drivers are able to achieve a similar fuel efficiency when relying entirely on the auditory feedback, b) users feel comfortable with the lack of a visual indicator, and c) if omitting the visual display can lead to a further reduction in distraction.
- Although the conducted laboratory study had the advantage of allowing us to control the test environment to a very high degree, this setting can obviously not reproduce the driving experience in an actual car, and we think that the implementation of the EcoSonic system in a number of vehicles and a corresponding longitudinal study in such a real-world setting would make it possible to further evaluate the feasibility of our designs for daily use as well as to examine how users adapt to it over a longer period of time.

10.

The Slowification System

10.1. Introduction

Especially when considering the mostly rather hectic urban traffic, car driving is not only a visually demanding task, but also one that is safety-critical for both the driver and other road users. Additionally, more and more in-vehicle systems are being integrated into the car, which almost exclusively rely on visual indicators for interacting with the driver.

For this reason, recent research efforts have targeted the *auditory* domain for in-vehicle interaction, and with the EcoSonic system, we have already demonstrated the effectiveness of two auditory ambient displays to support users in driving in a more energy-efficient way. The underlying data of these displays, however, is actually relatively simple, as they essentially convey information about a single one-dimensional variable. Given the generally low information capacity of peripheral displays, this is certainly a sensible choice, since presenting an increased amount of information might easily interfere with such displays' essential quality of unobtrusiveness. Nevertheless, in order to further extend the design space of ambient information systems, we wanted to explore, if and how it is possible to convey slightly more complex data to the user without compromising this quality.

Furthermore, it is important to acknowledge that the soundscape of a car is also a difficult environment to deal with, as we have to take into account a wide variety of background noises coming from the engine, the wind, and the tires. Additionally, many people listen to music or utilize a navigation system, which guides the driver by using speech notifications. In consequence, the majority of auditory cues used in the car are of rather salient nature, e.g. the sounds used in parking assistance systems or the distinct but admittedly fairly unpleasant noise to indicate that the driver should fasten the seatbelt. Similarly, indication that a driver is exceeding a prescribed speed limit,

provided for example by a navigation system, is commonly conveyed by quite salient auditory notifications.

Based on these observations, we have developed the Slowification auditory ambient display to support drivers in keeping a (continuously changing) “optimal” speed, as determined by prescribed speed limits, traffic lights, stop signs, and the overall traffic situation. The key conceptual idea of conveying this information is to assume that the sound of the car (i.e. the car’s audio system) is traveling with this optimal speed and to virtually position the driver into this space according to the current speed of the car, resulting in a sound which moves to the back as one drives too fast and catches up on slowing down.

For instance, when a driver has missed a speed sign and is driving too fast, the sound signal of the car’s audio system will gradually move from a centered position towards the back of the car. Conveying this information in such a way has three distinct advantages: a) The panning of a sound signal is rather easily perceived and should be difficult *not* to notice, matching the importance and urgency of the information. b) The meaning of the sound design should quite intuitively be understood, as you get the feeling of driving away from “your” sound. c) As the composition of sounds is not changed at all by this auditory display, it is very unobtrusive and thus should be easily accepted, which is of major importance when dealing with a sonic environment that so many people are exposed to as it is the case for automobiles.

10.2. Related Work

In addition to the relevant research we have already discussed in the previous chapter (especially Section 9.2.2 and Section 9.2.4), there has been some related work both on spatial panning and on speed assistance systems.

10.2.1. Spatial Panning to Guide Users

Although not used in lots of systems, there are a few instances where spatial panning has been incorporated into user interfaces to inform users about an event or point of interest in a certain direction. For example, Holland, Morse, and Gedenryd (2002) have developed a GPS navigation system with the goal of allowing users to be engaged in different activities while being guided by the system. To this end, they decided to use a non-speech audio interface to encode distance and direction of a location. In their prototype, the direction was represented by spatial panning of a tone across the stereo sound stage of a pair of headphones, based on the current moving direction of the user. Although seemingly coarse, this method yielded good enough results to discern the principal direction in an informal user trial.

In the context of automobiles, Fagerlönn, Lindberg, and Sirkka (2012) evaluated different ways of guiding drivers at the early stage of a dangerous driving situation like an imminent collision with another vehicle. In a study with 24 people, they compared using a) a mild warning sound, b) reducing the volume of the vehicle's radio, and c) panning the radio's signal. The authors conclude that panning the radio led to the lowest response times and was at the same time significantly better rated by users than the volume reduction.

10.2.2. Dynamic Speed Assistance Systems

Although currently the vast majority of speed limits are static, i.e. they consist of fixed signs that do not change in terms of position or limit, there are efforts to introduce more dynamic speed assistance systems, which take into account road geometry and vehicle characteristics (Jimenez, Aparicio, & Paez, 2008), or upcoming traffic signal information. For instance, Raubitschek et al. (2011) have developed an algorithm to calculate fuel consumption-optimized driving trajectories based on traffic lights information and other driving constraints such as stop signs. Although until now such projects mostly rely on simulated environments, there are efforts to make the necessary information available to in-vehicle assistance systems, e.g. via Car2X communication.

Those systems, however, will make the use of a traditional visual speed display far more challenging, as the drivers will have to deal with constantly changing and non-standardized speed limits, which, in turn, will require them to use additional or completely different user interfaces that can provide the necessary information in a way that does not interfere with the driving task.

10.3. Interaction Design

Even without a dynamic speed assistance system, keeping the speed is an important issue when driving, and too often the visual focus of attention is shifted to the speedometer and thus is distracted from the outside traffic situation where it should remain. However, speed limits are frequent: in cities, on country roads, close to railway crossings, and speeding is controlled and penalized. Furthermore, the speed of a car obviously has a direct influence on its fuel consumption, and as a "secondary parameter" in this regard (cp. Section 9.3.2), providing feedback on this information can therefore support a corresponding reduction.

In line with these observations, preliminary research on the effect of in-vehicle systems on the users' driving performance (cp. Section 9.2.2) suggests that the existing visual means for providing feedback on the car's current speed via a speedometer is not an optimal choice, as it leads to frequent visual distractions from the traffic situation. An

interaction design for providing this non-critical yet highly relevant information therefore needs to take the drivers' required focus for their primary task into account.

USE OF AN AUDITORY DISPLAY

As already discussed in the previous chapter, using an auditory ambient display should be advantageous in supporting this monitoring task, and indeed, some navigation systems already signal the exceeding of a speed limit by auditory alerts. However, not only can these be experienced as somewhat annoying, they also fail to represent details about the amount or significance of the deviation.

Analogous representations, in contrast, keep users informed at all times, provide a potentially less accurate, yet continuous cue about the underlying condition and leave the decision making in the hands of the user. The reason why such a continuous auditory display has not yet been considered for the speedometer is probably because a continuous sound would most likely be rather annoying in itself – even if we readily accept the permanent engine sounds and would even object if they were removed.

10.3.1. Concept

The preceding analysis hints at the potential of using one of the *existing* sounds that naturally emerge in a driving context – similar to the blended sonifications designed as part of the Sonic Shower (cp. Section 8.4). For the Slowification system, we chose to use the existing in-car audio system as source sound to be modified according to the available information. The fact that most users listen to music, audiobooks or radio while driving, and that a car's audio system is usually at least quadrophonic in order to allow a fine balance of sound between left/right and front/rear to meet the driver's preferences is the technical and conceptual basis for our sonification design.

A SOUND BUBBLE

Imagining that the sound of one's audio system is not fixed within the car, but instead travels at its own speed, the central idea is that, unlike the car itself, the sound travels exactly as fast as allowed (resp. as recommended), while still being elastically attached to the car's center of mass. One would further assume that the sound would be represented as a "sound bubble", which naturally encompasses the car and the driver. With this metaphorical setup, the following conditions can arise:

- If the driver exceeds the current speed limit, the sound bubble would fall back and be dragged by the car behind the user by means of the elastic attachment. This situation would naturally lead to the perception of the audio system's sound panning to the rear.

- On the other hand, if the driver goes slower than the allowed tempo and there is both traffic behind and no traffic in front, then the sound bubble would travel faster than the driver and lead to a spatial shift of the sound towards the front.
- Finally, if the car's speed is the same (or within tolerance) as recommended, the bubble would be perfectly centered, leading to no audible modification of the sound.

To further establish the metaphor of a sound bubble shifting either towards the front or the rear, and ultimately moving away from the car, the modification of the sound should not only include a pan to the corresponding direction, but also a change of other features, such as decreasing the sound level as the car's distance to the sound bubble's location increases, or to add reverberation, delay or other filtering plausible for distant sound sources. Such subtle cues might add to an enhanced sense of realism in this auditory display and thus improve its perception and also lead to an increased acceptance.

10.3.2. Implementation

As a first prototype, we implemented a rather straightforward version of the design concept described in the previous section. For this, we first defined a measure for driving faster (or slower) than a recommended speed:

$$d(\Delta_v, v_{\text{ref}}, \tau) = \max \left(\alpha \cdot |\Delta_v| + (1 - \alpha) \cdot \frac{|\Delta_v|}{v_{\text{ref}}} \cdot v_n - \tau, 0 \right), \quad (10.1)$$

where $\Delta_v = v - v_{\text{ref}}$ is the difference between the current and a reference speed, α is a weighting factor balancing relative and absolute speed difference, and v_n is a predefined neutral speed, where the (unweighted) relative and absolute speed differences would be the same. In our study (cp. Section 10.4), we used $\alpha = 0.8$ and $v_n = 70$ km/h. τ is a measure for the tolerated deviation from the reference speed and is used to define a "speed channel" around v_{ref} , with a separate lower and upper bound for going too fast (τ_u) or too slow (τ_l). In our current implementation, we have defined $\tau_l = 3$ km/h and $\tau_u = 5$ km/h.

Driving faster than v_{ref} would lead to a gradual spatial shift of the sound towards the back, while driving slower to a shift towards the front of the car. The amount of panning is determined by

$$P_{u/l} = \Phi(d(\Delta_v, v_{\text{ref}}, \tau_{u/l})) \quad \text{with} \quad \Phi(d) = \rho \cdot \sqrt{d}, \quad (10.2)$$

where ρ is a scaling factor and \sqrt{d} has been chosen to achieve a more noticeable spatial shift directly after crossing the threshold. In our quadrophonic speaker setup, we pan each stereo channel separately with Supercollider's *Pan2* UGen. Furthermore, if $P > 1$, the amplitude of the audio signal will be reduced by $\nu_{\text{dB}} \cdot (P - 1)$, indicating a further



Figure 10.1.: Picture of the car simulator.

movement of the sound bubble towards the respective direction (cp. Section 10.3.1). For our study, we have defined $\nu_{dB} = 25\text{dB}$ and $\rho = 0.2$.

Finally, when dealing with changing speed limits or traffic lights, the bounds of the speed channel further deviate: First, it is common practice for a driver to “coast” (i.e. to only slowly decelerate) when encountering traffic lights or a slower speed limit, which is also beneficial in terms of fuel economy. Therefore, we manage a separate upper and lower reference speed, and the lower one v_{ref}^l will drop by a deceleration constant $a_d = 0.1 \text{ km/h/m}$ well before passing red traffic lights or a sign indicating a slower speed limit, meaning that there will be no panning to the front if the driver chooses to coast. Additionally, the upper bound v_{ref}^u will drop a short distance before the sign by a braking constant $a_b = 0.8 \text{ km/h/m}$ to indicate the upcoming speed limit, if the driver has not reduced the speed by then.

10.4. Study

In order to assess the efficacy of our design in terms of a) drivers adhering to the prescribed speed limit, b) the subjective and measured distraction by the panning, and c) the acceptance of the general design, we have enhanced the EcoSonic simulator environment described in Section 9.5.2 by a more realistic driving experience.

10.4.1. A Virtual Reality Car Simulator

In order to do so, we have developed a car simulator conveying a virtual reality 3D environment with the help of an Oculus Rift¹ (also cp. Figure 10.1). It is written from scratch in three.js² and thus can be run in any browser. The portability of this approach has the distinct advantage of enabling us to use this high quality simulation environment also as part of an online study, therefore allowing us to reach far more participants than in a laboratory study. Alternatively, the possibility of experiencing the car simulator in a virtual reality environment allows for an even greater sense of realism.

The simulator includes the physics based engine model described in Section 9.5.2, including the torque map to model the engine's varying torque responses as well as the fuel consumption map for a realistic calculation of fuel consumption. It furthermore extends the simulation of the car to the 3D environment. Similar to the EcoSonic system, it has a dedicated interface to SuperCollider via OSC, which is also used to create the engine sound. For the study, we additionally implemented a way to stream internet radio into Supercollider via a virtual sound card in order to simulate listening to the radio while driving and as input for our Slowification system.

10.4.2. Study Design

With the help of our simulator environment, we conducted a study to evaluate the prototype implementation of the Slowification system discussed in Section 10.3.2 and to compare it to using a conventional speedometer. We employed a within-subject design, where participants consecutively experienced both conditions and for each of them drove the same test track three times in order for them to familiarize with the respective display. Controlling for ordering effects, we furthermore employed a counterbalanced measures design, where the two condition sequences were evenly distributed among the study participants.

For the study, we designed a circular track, with speed limits ranging from 30 km/h to 130 km/h. The lengths m_i of the individual segment belonging to a particular speed limit l_i were chosen in such a way that the time needed to drive through them was approximately the same, i.e.

$$t_i \approx t_j, \quad i, j \in \{1, \dots, n\}, \quad \text{with } t_i = \frac{m_i}{l_i} \quad (10.3)$$

Furthermore, the curve radius was adjusted depending on the respective speed limit so that segments with a high speed limit had a wider radius than segments with a lower one. The time to complete one lap was approximately 2 minutes.

¹Oculus Rift: A virtual reality headset (<https://www.oculus.com>)

²three.js: A JavaScript 3D Library (<http://threejs.org>)

ATTENTION TASK

When driving a car, a multitude of distractions, such as other cars, bicycles, and a lively surrounding, need to be attended to safely navigate to one's destination. In a simulator environment, however, those distractions could be easily ignored, and the participants' attention level would be difficult to assess. Consequently, we instead chose to implement a more abstract and plain attention task, whose performance can be measured more easily. In the spirit of the time, we designed a Pokémon-themed task that was both simple and engaging: While driving on the street, there will appear different kinds of Pokémon that you can, true to the original game, catch with a Pokéball (also cp. Figure 10.1). This is done simply by looking at the Pokémon and pressing a button located on the steering wheel.

HARDWARE SETUP

In Figure 10.2, we can see the actual hardware setup used in the experiment. Four loudspeakers³ were placed in a quadrophonic setup around the user, and as a virtual reality headset we used the consumer version of the Oculus Rift. As input device, the same combination of steering wheel and pedals was used as the one used in the EcoSonic study.

PROCEDURE

All in all, the procedure of the experiment was very similar to the one of the EcoSonic study, as described in Section 9.6. At the beginning, participants provided some general information about themselves, such as age and occupation, but also about perceptual deficiencies, if they were regular car drivers, or their general interest in a support system giving guidance on speed choices. They were also given a short written introduction explaining the basic concept behind the feedback provided by the Slowification system and telling them what they were expected to do during the experiment. Specifically, they were told to 1) keep on their lane, 2) not to drive through red traffic lights or ignore stop signs, and 3) to comply to the speed limits, i.e. basically to follow commonly-known traffic rules. As the last, secondary assignment, they were told to capture as many Pokémon as possible, including how to do so (cp. Section 10.4.2).

For the actual experiment, all participants were told to first familiarize with their real-world environment in order for them to be able to easily reach the pedals and the steering wheel, even when their view was obstructed by the Oculus Rift. Moreover, the participants could select any (internet streamable) radio channel to customize their soundscape to what they were accustomed to when driving a car. After familiarizing with the Oculus Rift and the car simulator, the participants had two driving sessions – one with and one without the Slowification system – where they would independently drive three laps of the track. After each session, they completed a questionnaire about

³We used four Genelec 8020A



Figure 10.2.: Hardware setup for the study. Two additional loudspeakers (not seen in the picture) were placed behind the participant. The computer monitor on the right was used only for controlling the application and could not be observed by the participants during the experiment. The head tracking sensor of the Oculus Rift can be seen between the two loudspeakers in the front.

the preceding driving session, followed by several comparative questions at the end of the experiment.

10.4.3. Goals and Hypotheses

The primary goal of the experiment was to evaluate the described design with regard to the following aspects:

- **Adhering to the prescribed speed limit:** As the participants are given the secondary task of catching Pokémon and the speed limit changes several times while they are driving the track, it can be expected that there is a certain amount of time where the respective speed limit will be exceeded. Our main hypothesis is that the Slowification system will help the participants to better adhere to the prescribed speed limits than without it (H1).
- **Distraction:** We furthermore assume that, in comparison to keeping an eye on the visual speed display, the participants will be less distracted by the panning of the radio's sound. We assume that this will not only be measurable by the amount of time the participants will deviate from their lane (H2), but also lead to the participants

feeling less distracted, as should be reflected by the answers in the questionnaire (H3).

- **Helpfulness:** Although the helpfulness of the Slowification system should as well be reflected by H1, we also expect the *perceived* helpfulness to be something that can be confirmed by the questionnaire (H4).
- **Acceptance:** A final important aspect of a user interface design that is meant to be installed in an automotive context is the user acceptance. Although most of the participants can be expected to be accustomed to the conventional speed dial and to the routinely glance to the dashboard, we hope that the Slowification system will at least be as comfortably to use for the participants as the speed dial (H5).

10.5. Results

In total, we invited 22 people to try out the Slowification system within our simulation environment. Three of them, however, had to abort the experiment as they were very soon feeling sick because of the VR environment (this is a common problem with VR Devices such as the Oculus Rift and has nothing to do with the Slowification system), leaving a total of $n=19$ fully evaluable data sets. The participants were 21-30 years old and balanced in terms of gender (9 male and 10 female participants). If not otherwise noted, we used a conventional t-test for comparing values from different conditions. For calculating the effect size, Cohen's d was used.

10.5.1. Measured Data

First of all, we can report that the attention task was successful in engaging the participants during the study for both the Slowification and the visual only condition, as there are no significant differences in the amount of Pokémon that have been caught per lap ($p > 0.5$).

Furthermore, In order to evaluate to what extent the prescribed speed limits were adhered to, we analyzed the percentage of time for each lap that a participant was driving more than 15 km/h too fast. As can be seen in Figure 10.3a, this was considerably less the case for the panning condition ($7.5\% \pm 9.5$) than for the baseline condition ($12.7\% \pm 15.7$), which confirms our hypothesis H1 ($p < 0.05$, Cohen's $d = 0.39$).

Finally, as a measure for being distracted, we compared the amount of time the drivers deviated from their own lane by more than 40 cm (Figure 10.3b). Although the differences are not as striking, there is a significant difference when considering our one-sided hypothesis ($p/2 < 0.05$, Cohen's $d = 0.34$) between driving with ($53.2\% \pm 11.0$) and without ($56.9\% \pm 10.6$) the Slowification system, confirming H2.

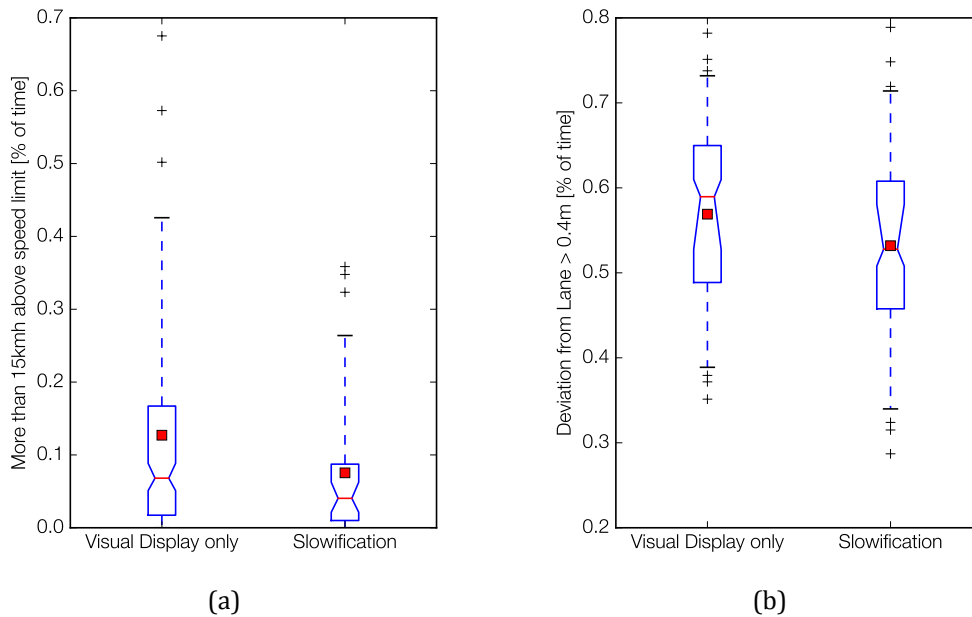


Figure 10.3.: Boxplot of **(a)** The percentage of time that a person was driving more than 15 km/h faster than the prescribed speed limit, and **(b)** the percentage of time that a person deviated too far from the street resp. the correct lane. The whiskers denote the 5% and 95% percentiles of the data, while the notches represent the 95% confidence intervals of the median. The mean values of the data are illustrated by the red boxes.

10.5.2. Questionnaires

This result is supported by the responses to the question how *distracting* the participants found the respective feedback. As can be seen in Figure 10.4, when being supported by the Slowification system, the users felt significantly less distracted ($\mu_{\frac{1}{2}} = -2$) than when not ($\mu_{\frac{1}{2}} = 1$), which clearly confirms H3 ($p < 0.01$).

Being asked about *helpfulness*, however, participants rated the two conditions almost the same (both $\mu_{\frac{1}{2}} = 1$), which cannot support our H4. Our interpretation of this result is that the participants, in the short amount of time they had to become accustomed to the system, could not *consciously* “grasp” it in a way that they could assess it as useful, although they were still able to at least partly process the provided information, as can be seen in the results on adhering to the speed limit (Figure 10.3a). This is also reflected by the answers to the question, how much the participants had to concentrate on the feedback, where there were no significant differences between the two conditions. However, attending to the Slowification system seemed to be less *stressful* ($\mu_{\frac{1}{2}} = -1$) than when attending to the speedometer ($\mu_{\frac{1}{2}} = 1$), which further supports our hypothesis H3 ($p/2 < 0.03$).




Figure 10.4.: Main results from the questionnaire of the study. Answers could be given on a 7-point Likert-type scale indicating the level of agreement with the statements that were given. In this chart, the width of each bar corresponds to how many people responded with the respective agreement, whereby the outer, more deeply colored bars represent a stronger reaction, while the inner, more lightly colored bars represent a weaker tendency. Only the responses that were not “neutral” are displayed.

Finally, as a measure for how well such a system would be accepted as an additional in-vehicle user interface, the participants stated that they could attend to the Slowification more comfortably ($\mu_{\frac{1}{2}} = 1$) than to the speedometer ($\mu_{\frac{1}{2}} = 0$), which confirms H5 ($p < 0.01$). Also, the acceptance of the displays is highly correlated with their unobtrusiveness ($\tau = 0.62$; $p < 0.01$), adding to RQ5.

10.6. Discussion and Conclusions

The conducted study gives a first indication for the efficacy of our Slowification concept, as described in Section 10.3.1. Notably, it clearly demonstrates that the use of sound in a peripheral display is a feasible, and in this case even advantageous, alternative to visual ambient interfaces (RQ2), as participants both managed to better adhere to the speed limits and simultaneously felt less distracted when supported by the auditory peripheral display.

Although the majority of participants (67%) indicated that they would prefer the Slowification over the speedometer, we think that it should not replace, but rather complement its precise display of the car’s current speed, so that the user always has a fallback option when the need arises. Also, we are aware that the chosen implementation as well as the rather subjective choice of parameters might not necessarily be the best possible one. For instance, the fact that the large majority of the participants was immediately feeling very comfortable with the system due to the general unobtrusiveness of the design leaves some room for making the indication of driving too fast

(or too slow) more distinct, thereby also taking into account the comments of a few participants who said to have difficulties to perceive the spatial shift of the sound.

A BLENDED SONIFICATION

As we have pointed out before, the soundscape of a car consists of several sounds and background noises that can lead to masking effects for any auditory display. Our approach with the EcoSonic system to deal with this issue was to slightly adjust the volume of the sonifications in dependence of the loudness of those noises. While this helps to retain the comprehensibility of the added sounds, it also increases the overall loudness of the soundscape, which can potentially lead to the sonification being perceived as somewhat obtrusive.

In contrast, with the Slowification system, we circumvented this problem by creating a blended sonification based on the usually most salient sound source in the car: its audio system. The Slowification system thereby exemplifies a more general approach of enhancing the unobtrusiveness of ambient displays by building on already existing sensory stimuli and altering these in order to convey additional information (RQ4). Although it must be noted that the current implementation of the Slowification requires the driver to listen to the radio (or other sources of audio) in order to work, it should also be possible to modify the engine noise in a similar way, when the driver is not.

DATA COMPLEXITY

Finally, the Slowification system demonstrates the unobtrusive conveyance of information in an ambient display based on a slightly more complex underlying data source than in the previously discussed projects, as the provided feedback not only depends on the car's current speed, but also on the occurrence of speed signs or traffic lights. In order to present this information in a non-obtrusive and peripheral way, two strategies were employed in our design approach: First, we focused on what information drivers might (and might not) need, leading to a feedback concept where the available data was heavily preprocessed and thereby reduced in complexity. Second, the chosen modality (i.e. the spatial panning of the car's audio system's sound signal) very well matches the resulting data (i.e. the deviation from an optimal speed), as it not only allows to represent a signed input variable, contrary to using a pitch-based mapping, for example, but also displays the data through a commonly known phenomenon of speed difference, i.e. the slowly moving away of one object from the other.

FUTURE WORK

Based on our initial design of the Slowification system, we see several viable directions of future work:

- During the study, one participant stated that “the panning is a really good idea” but felt that she needed more time to get accustomed to it and suggested “more time

for test drives”. A natural way to allow users to get accustomed to the Slowification would be to install the system in a small number of cars for people to experience the feedback over a longer period of time. While certainly more difficult to evaluate, as we would be dealing with a completely uncontrolled environment, this would give insight into how users would be using the system after really becoming accustomed to it and how well it is usable in real-life situations.

- Although the Slowification was designed to complement rather than replace the speedometer, it might be interesting to study how well it would work as the *only* available feedback and to evaluate if potential users would feel comfortable when not being able to rely on the speedometer’s precise display of the car’s current speed.
- Another way to further evaluate our speed indicator would be to compare it to a different type of (auditory) display, e.g. an alert-based system, which we would assume to be rated as far more annoying than the Slowification.
- Finally, it would be interesting to extend the use cases of the system by integrating an adaptive speed assistance system based on traffic light predictions, which we think would make the advantages of the Slowification even more distinct than with static speed signs only.

Part III.

Discussion and Conclusions

After having discussed in detail the individual results of each conducted study, we can now get back to our initial research questions, described in Chapter 6.

11.1. Complexity and Variability of the Context

First, we aimed to investigate:

How does the variability of an ambient display's context of use affect the way it is perceived? (RQ3)

As we have learned, one of the main differences between ambient displays and more conventional information systems is that it is usually not the users' main task to actively observe the display, but rather that they are engaged with another activity, which is complemented or directly supported by the display. The most important aspects of the context of use are therefore a) if there is a clearly defined primary task or a variety of activities users can be expected to be engaged with, and b) the users' overall perceptual situation, i.e. their field of view, which of their senses can be expected to be occupied, and if these aspects remain fixed during the use of the display.

With our prototype systems and in our studies, we consequently explored several settings with a range of different use contexts:

- The living environment the InfoPlant was installed in is an example for a context, where users can be expected to perform a large variety of activities in the vicinity of the ambient display, but also to frequently move out of sight and away from it, therefore leading to a quite unstable perceptual situation as well.
- Contrary to that, for the InfoDrops system, the users' primary activity is rather clearly defined: to take a shower. However, as they can easily move (or at least turn) around while doing so, the perceptual situation is obviously not completely fixed.
- Finally, the in-vehicle environment of the EcoSonic and the Slowification system represents a context, where users not only have a well-defined primary task (i.e. driving the car), but where the perceptual situation is rather predictable as well.

Based on our experience in the conducted studies we can now say that a clearly defined primary activity and a predictable perceptual situation are quite certainly conducive to the functioning of an ambient display: While on the one hand, the in-vehicle environment has shown to be particularly well-suited for a peripheral display, even for continuous monitoring purposes, the InfoPlant's household setting, on the other hand, has proven to be rather problematic, as users found it quite difficult to observe the display during their daily activities.

However, we have also seen that for usage scenarios between these two extremes, the question of suitability further depends on the employed modality, i.e. while a certain variability in the users' perceptual situation is perfectly acceptable for systems using auditory cues such as the Sonic Shower, this would already be problematic for a peripheral display solely relying on visual means to convey information, due to the users' limited field of view.

Prospectively, these problems might be somewhat alleviated by a more pervasive distribution of interconnected displays or by technological advances in deducing the users' level of attention. However, for displays where a purely auditory representation of information might not be the optimal choice, we instead advocate the use of *multimodal* displays, as we think that the strengths and weaknesses of the auditory and visual modality complement each other very well in a peripheral display.

11.2. Unobtrusiveness and Acceptance

Considering the importance calm computing places on the unobtrusiveness of user interfaces on the one hand and the popularity of technical systems that require focused attention on the other hand, we furthermore wanted to know:

Are more unobtrusive ambient displays better accepted by users? (RQ5)

Although we have to keep in mind that the use cases we have selected in our studies are of course compatible to a peripheral perception of the provided information, the results clearly show that unobtrusiveness is indeed a desirable quality for ambient displays and that they are most likely better accepted by users than those requiring focused attention. Consequently:

What are effective ways to make ambient interfaces less obtrusive and distracting? (RQ4)

First of all, in the InfoPlant study, we have seen a strong indication that the perceived intrusiveness of an ambient display depends more on the way information is presented than on its overall appearance. Therefore, a display represented by an artifact that is aesthetically pleasing or that blends well into its surroundings does not seem to particularly help in increasing its unobtrusiveness, and from a functional point of view, the emphasis on the appearance might even be detrimental to the display's ability to efficiently convey information.

On the other hand, it is certainly highly important *how* the information is presented:

- The results of our InfoPlant survey suggest that a rather simple and plain design should be advantageous not only for the understandability of the provided information, but also in terms of their unobtrusiveness. This also makes sense from a theoretical point of view, as more complex representations naturally require an increased cognitive effort to process.
- Furthermore, we have found that addressing a modality that is not overly used by the user to perform a primary task is highly conducive to making a peripheral display less obtrusive and distracting. This finding is also supported by theories and experiments from cognitive psychology (cp. Chapter 2) .
- Finally, we have explored building on already existing sensory stimuli and altering these in order to convey additional information, which, at least for the auditory modality in the form of blended sonifications, seems to contribute to enhancing the unobtrusiveness of peripheral displays.

11.3. The Use of Sound in Ambient Displays

With these results in mind, we can continue answering:

Is using sound in a peripheral display a feasible alternative to (purely) visual ambient interfaces? (RQ2)

First of all, we can state that the auditory displays evaluated in our studies have been rated very similar or even better than their visual counterparts in terms of understandability and other usability aspects, indicating that their usage does not pose a problem in this regard. Furthermore, taking into account that the majority of activities we would consider a “primary task” are of mainly visual nature, such as in our experiments taking a shower or driving the car, our results suggest that employing the auditory modality can even be beneficial, especially for *peripheral* interaction, as this allows the presented information to be processed via a comparatively free perceptual channel. Finally, we have seen that auditory displays are usable for a wider range of use contexts for ambient interfaces, since the peripheral perception of auditory stimuli does not require users to be in a perceptual situation as fixed as for visual ones.

In summary we can say that the use of sound in an ambient display is not only a feasible alternative to using visual means to convey information, but for a large number of use cases it even seems to be far more suitable to facilitate peripheral perception.

11.4. Feedback, Persuasion, and Affectiveness

Especially with regard to ambient displays that have been designed to induce a more environmentally-conscious behavior, we wanted to know:

Is providing feedback on consumption practices enough to induce a corresponding change in behavior? (RQ6)

According to our findings, feedback can certainly be considered a useful and important tool to *support* users in accomplishing behavior change, as it rather consistently improves their performance in this regard. On the other hand, we have also seen that users who lack intrinsic motivation for change are rather unlikely to be persuaded by feedback alone, and we argue that research in sustainable HCI should move away from the assumption that providing feedback will more or less automatically lead to a change in the recipients' behavior.

Furthermore, since preliminary research suggests that the use of affective cues might be conducive to inducing a behavior change, we were also interested in:

Is the use of affective cues beneficial for ambient displays aiming to persuade users towards a change in behavior? (RQ8)

Although based on our study results alone, this question cannot be answered conclusively, our findings suggest that:

- Rather obvious attempts at evoking an emotional response most likely increase the perceived intrusiveness of a peripheral display and, especially when used as part of a negative reinforcement design, might not only easily fail to induce a change in behavior, but also have a significant impact on the feedback's acceptance.
- Both the use of *subliminal* affective cues, as we have explored in several feedback designs for the InfoDrops system, and the emphasis on evoking *positive* emotional responses might be feasible approaches for ambient displays that should be further investigated.
- Similarly, while our findings suggest that affective auditory cues can generate a quicker and more activating response in users, we think that further research is needed to evaluate this specific aspect of auditory feedback.

Considering the more general question

Is it viable to use ambient displays as persuasive technology? (RQ7)

we see the limited information capacity of ambient displays to be a major restricting factor in their ability to actually *convince* users to change their behavior and to achieve a lasting change in attitude, since strategies to do so rely primarily on conveying information that should be consciously perceived and processed. However, when interpreting persuasion more as a way to *support* users in changing their behavior, ambient information systems can certainly be a useful tool to do so.

11.5. Ambient Information Systems Supporting Environmentally-Conscious Behavior

Finally, with these considerations in mind, we can have a look at peripheral displays from a more application-centered perspective by reviewing:

How can ambient information systems support environmentally-conscious behavior? What are use cases that “work” for using such an interface? What are not? (RQ1)

First, we have seen that ambient displays are particularly suited for providing relevant data to users who are, at the same time, engaged with a certain primary activity or task. However, due to their focus on peripheral perception, those displays must be considered inappropriate for conveying very detailed pieces of information. In the context of sustainable HCI, ambient displays might therefore

- provide feedback on everyday activities, such as helping users to gain a comprehensive view on their consumption practices, thereby enabling them to make better decisions in their conservation efforts.
- help users to improve the performance of a certain skill that is conducive to lowering the impact on the environment, such as eco-driving.
- support users in accomplishing a specific task or activity that can alleviate the environmental impact, for example providing cyclists with navigational cues for a faster or more convenient route to make this alternative to car driving more attractive.

On the other hand, peripheral displays would not be suited for providing extensive information on environmental issues, e.g. to facilitate *learning* about the impact of one’s action on the environment. Considering the fact that most issues of sustainability are actually rather complex, we therefore argue that in order to use ambient interfaces to their full potential, they should be used preferably in conjunction with other interaction technology.

11.6. Outlook

AUTONOMOUS DRIVING

As we have seen, in-vehicle environments are particularly suited for peripheral displays, and considering the work that is currently being done towards semi-autonomous cars, a challenging direction of research is to provide users with the necessary means to stay informed about the driving situation in order to intervene when needed. Since passengers of such vehicles can be expected to work in the car or to be engaged with other activities that can be accomplished in this environment, ambient interfaces should be perfectly suited to allow users to simultaneously stay aware of the driving process.

BEHAVIOR RETENTION AND HABIT FORMATION

Especially when considering ambient displays to provide feedback systems on energy usage, which frequently require a certain amount of energy to function themselves, one might question the effectiveness of such a system when being installed on a *permanent* basis: Assuming that users are able to adapt their behavior with the help of such a system, the central question here is to what degree this behavior change is dependent on the immediate availability of the feedback, or if the system might also be able to support a *sustained* change in behavior. Although examining this particular aspect of feedback systems would obviously require a longitudinal study over an extended period of time, we think that this would give further insight into the practical applicability of systems whose permanent usage might not be feasible.

LOAD SHIFTING FOR RENEWABLE ENERGY INTEGRATION

Finally, taking into account the current development towards renewable energy and the difficulties of providing a *consistent* electricity supply from the various energy sources, we argue that peripheral displays can not only support users in their conservation efforts, but should be equally suited to facilitate *shifting* consumption to a favorable time by providing feedback on the current availability or time dependent costs of electricity. Considering the difficulties of visual ambient displays with regard to the wide range of activities conducted in a household, we propose to do so with an easily installed distributed audio system to allow for ubiquitous monitoring, in conjunction with a central controlling unit for further and more detailed inspection of the available data.

A. Questionnaires, Images, and Sound Examples

All supplementary material of this thesis can be found on the accompanying CD.

This includes:

- The questionnaires and introductory material of the studies on the InfoPlant (Section 7.5), the EcoSonic system (Section 9.6), and the Slowification system (Section 10.4).
- Further images of the first and second prototype of the InfoPlant (cp. Section 7.4).
- The sound examples of the Sonic Shower we have discussed in Section 8.4, as well as those used in the online survey (Section 8.5).
- Two sound examples of both the metaphorical and the continuous sonification, as used in the EcoSonic system (Sections 9.4.1 and 9.4.2).
- The source code for the InfoPlant, EcoSonic, and Slowification projects.

Additionally, the source code is hosted on github to allow for subsequent development on the various projects:

- InfoPlant <https://github.com/JanHammerschmidt/InfoPlant>
- EcoSonic <https://github.com/JanHammerschmidt/EcoSonic>
- Slowification <https://github.com/JanHammerschmidt/car-simulator>

Bibliography

- Adams, A. T., Costa, J., Jung, M. F., & Choudhury, T. (2015). Mindless computing: Designing technologies to subtly influence behavior. In *Proceedings of the 2015 acm international joint conference on pervasive and ubiquitous computing* (pp. 719–730). ACM.
- Altosaar, M., Vertegaal, R., Sohn, C., & Cheng, D. (2006). Auraorb: Social notification appliance. In *Chi'06 extended abstracts on human factors in computing systems* (pp. 381–386). ACM.
- Amabile, T. & Kramer, S. (2011). *The progress principle: Using small wins to ignite joy, engagement, and creativity at work*. Harvard Business Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Harvard University Press.
- Antifakos, S. & Schiele, B. (2003). LaughingLily: Using a flower as a real-world information display. *International Conference on Ubiquitous Computing (Ubicomp)*. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.10.2743%5C&rep=rep1%5C&type=pdf>
- Arroyo, E., Bonanni, L., & Selker, T. (2005). Waterbot: Exploring feedback and persuasive techniques at the sink. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 631–639). CHI '05. Portland, Oregon, USA: ACM. doi:10.1145/1054972.1055059
- Ars Technica. (2015). *Fairphone 2 hands-on*. <https://arstechnica.com/gadgets/2015/10/fairphone-2-hands-on-modular-phones-are-finally-here/>. [last accessed: 26.06.2017].
- Backlund, S., Gyllenswärd, M., Gustafsson, A., Ilstedt Hjelm, S., Mazé, R., & Redström, J. (2007). Static! the aesthetics of energy in everyday things. In *Proceedings of design research society wonderground international conference 2006*.
- Bakker, S. (2013). *Design for peripheral interaction*. Eindhoven University of Technology.
- Barkenbus, J. N. (2010). Eco-driving: An overlooked climate change initiative. *Energy Policy*, 38(2), 762–769. doi:10.1016/j.enpol.2009.10.021
- Bauer, J. S., Consolvo, S., Greenstein, B., Schooler, J., Wu, E., Watson, N. F., & Kientz, J. (2012). Shuteye: Encouraging awareness of healthy sleep recommendations with a mobile, peripheral display. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1401–1410). ACM.

- Berker, T. (2013). "in the morning i just need a long, hot shower:" a sociological exploration of energy sensibilities in norwegian bathrooms. *Sustainability : Science, Practice and Policy*, 9(1), 57–63.
- Blevins, E. (2007). Sustainable interaction design: Invention & disposal, renewal & reuse. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 503–512). ACM.
- Bovermann, T., Tünnermann, R., & Hermann, T. (2010). Auditory augmentation. *International Journal of Ambient Computing and Intelligence (IJACI)*, 2(2), 27–41.
- Bresin, R. & Friberg, A. (2000). Emotional coloring of computer-controlled music performances. *Computer Music Journal*, 24(4), 44–63.
- Brewer, J., Williams, A., & Dourish, P. (2005). Nimio: An ambient awareness device. In *Proc. of ecscw 2005*.
- Broadbent, D. (1958). Perception and communication.
- Brooks, C. A. & Rakotonirainy, A. (2005). In-vehicle Technologies, Advanced Driver Assistance Systems and Driver Distraction: Research challenges. In *International conference on driver distraction*. Retrieved from http://eprints.qut.edu.au/2024/1/2024%5C_1.pdf
- Buchanan, K., Russo, R., & Anderson, B. (2015). The question of energy reduction: The problem(s) with feedback. *Energy Policy*, 77, 89–96.
- Cadiz, J. J., Venolia, G., Jancke, G., & Gupta, A. (2002). Designing and deploying an information awareness interface. In *Proceedings of the 2002 acm conference on computer supported cooperative work* (pp. 314–323). ACM.
- Chaos Computer Club. (2001). *Blinkenlights*. <http://blinkenlights.net/blinkenlights>.
- Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the acoustical society of America*, 25(5), 975–979.
- Chetty, M., Tran, D., & Grinter, R. E. (2008). Getting to green: Understanding resource consumption in the home. In *Proceedings of the 10th international conference on ubiquitous computing* (pp. 242–251). ACM.
- Christos Goumopoulos, N. D., Eleni Christopoulou & Kameas, A. (2004). *The PLANTS System: Enabling Mixed Societies of Communicating Plants and Artefacts* (P. Markopoulos, B. Eggen, E. Aarts, & J. L. Crowley, Eds.). Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/b102265
- Consolvo, S., Everitt, K., Smith, I., & Landay, J. A. (2006). Design requirements for technologies that encourage physical activity. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 457–466). ACM.
- Consolvo, S., Roessler, P., & Shelton, B. E. (2004). The carenet display: Lessons learned from an in home evaluation of an ambient display. In *International conference on ubiquitous computing* (pp. 1–17). Springer.
- Costanza, E., Inverso, S. A., Pavlov, E., Allen, R., & Maes, P. (2006). Eye-q: Eyeglass peripheral display for subtle intimate notifications. In *Proceedings of the 8th conference*

- on human-computer interaction with mobile devices and services (pp. 211–218). ACM.
- Dahley, A., Wisneski, C., & Ishii, H. (1998). Water lamp and pinwheels: Ambient projection of digital information into architectural space. In *Chi 98 conference summary on human factors in computing systems* (pp. 269–270). ACM.
- Darby, S. et al. (2006). The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*, 486(2006).
- den Breejen, J. & Deenstra, M. (2005). *Ping Pong Pixel*. <http://www.tupajumi.com/pingpongpixel/>.
- Deterding, S., Sicart, M., Nacke, L., O'Hara, K., & Dixon, D. (2011). Gamification. using game-design elements in non-gaming contexts. In *Chi ea '11*. doi:10.1145/1979742.1979575
- Deutsch, J. A. & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological review*, 70(1), 80.
- Dillahunt, T., Mankoff, J., Paulos, E., & Fussell, S. (2009). It's not all about green: Energy use in low-income communities. In *Proceedings of the 11th international conference on ubiquitous computing* (pp. 255–264). ACM.
- DiSalvo, C., Sengers, P., & Brynjarsdóttir, H. (2010). Mapping the landscape of sustainable hci. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1975–1984). CHI '10. Atlanta, Georgia, USA: ACM. doi:10.1145/1753326.1753625
- Evans, L. (1979). Driver behavior effects on fuel consumption in urban driving. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 21(4), 389–398.
- Eysenck, M. W. & Calvo, M. G. (1992). Anxiety and performance: The processing efficiency theory. *Cognition & Emotion*, 6(6), 409–434.
- Fagerlönn, J., Lindberg, S., & Sirkka, A. (2012). Graded auditory warnings during in-vehicle use. In *Proceedings of the 4th international conference on automotive user interfaces and interactive vehicular applications - automotiveui '12* (p. 85). New York, New York, USA: ACM Press. doi:10.1145/2390256.2390269
- Fairphone B.V. (2015). *Fairphone 2*. <https://www.fairphone.com/>. [last accessed: 26.06.2017].
- Fogarty, J., Forlizzi, J., & Hudson, S. E. (2001). Aesthetic information collages: Generating decorative displays that contain information. In *Proceedings of the 14th annual acm symposium on user interface software and technology* (pp. 141–150). ACM.
- Fogg, B. J. (2002). Persuasive technology: Using computers to change what we think and do. *Ubiquity*, 2002(December), 5.
- Froehlich, J., Dillahunt, T., Klasnja, P., Mankoff, J., Consolvo, S., Harrison, B., & Landay, J. A. (2009). Ubigreen: Investigating a mobile tool for tracking and supporting green transportation habits. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 1043–1052). ACM.

- Gable, T. M., Walker, B. N., Moses, H. R., & Chitloor, R. D. (2013). Advanced auditory cues on mobile phones help keep drivers' eyes on the road. In *Proceedings of the 5th international conference on automotive user interfaces and interactive vehicular applications - automotiveui '13* (pp. 66–73). New York, New York, USA: ACM Press. doi:10.1145/2516540.2516541
- Gaver, W., Bowers, J., Boucher, A., Law, A., Pennington, S., & Villar, N. (2006). The history tablecloth: Illuminating domestic activity. In *Proceedings of the 6th conference on designing interactive systems* (pp. 199–208). ACM.
- Gonder, J., Earleywine, M., & Sparks, W. (2012). *Analyzing vehicle fuel saving opportunities through intelligent driver feedback*. National Renewable Energy Laboratory (NREL), Golden, CO.
- Google. (2013a). *Google Glass*. https://en.wikipedia.org/wiki/Google_Glass.
- Google. (2013b). *Project Ara*. <https://atap.google.com/ara/>. [last accessed: 26.06.2017].
- Google Glass. (2014). <http://sites.psu.edu/internationalstudent/2014/03/28/google-glass/>.
- Greenberg, S. & Rounding, M. (2001). The notification collage: Posting information to public and personal displays. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 514–521). ACM.
- Gustafsson, A. & Gyllenswärd, M. (2005). The power-aware cord: Energy awareness through ambient information display. In *Chi '05 extended abstracts on human factors in computing systems* (pp. 1423–1426). CHI EA '05. Portland, OR, USA: ACM. doi:10.1145/1056808.1056932
- Hakkens, D. (2013). *Phonebloks: A concept for a modular smartphone design*. <https://phonebloks.com/>. [last accessed: 26.06.2017].
- Halko, S. & Kientz, J. A. (2010). Personality and persuasive technology: An exploratory study on health-promoting mobile applications. In *International conference on persuasive technology* (pp. 150–161). Springer.
- Hammerschmidt, J. & Hermann, T. (2017). EcoSonic: Auditory Peripheral Monitoring of Fuel Consumption for Fuel-Efficient Driving. *Displays*, 47, 40–50.
- Hammerschmidt, J., Hermann, T., Walender, A., & Krömker, N. (2015). InfoPlant: Multimodal Augmentation of Plants for Enhanced Human-Computer Interaction. In *6th IEEE International Conference on Cognitive Infocommunications, 2015* (pp. 511–516). IEEE.
- Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2014a). EcoSonic: Towards an Auditory Display Supporting a Fuel-Efficient Driving Style. In *Proceedings of the Conference on Sonification of Health and Environmental Data (SoniHED 2014)* (p. 56). York, England.
- Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2013). InfoDrops: Sonification for Enhanced Awareness of Resource Consumption in the Shower. In *Proceedings of the 19th International Conference on Auditory Display (ICAD-2013)*. Lodz, Poland.

- Hammerschmidt, J., Tünnermann, R., & Hermann, T. (2014b). EcoSonic: Auditory Displays supporting Fuel-Efficient Driving. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational - NordiCHI '14* (pp. 979–982). New York, New York, USA: ACM Press.
- Heiner, J. M., Hudson, S. E., & Tanaka, K. (1999). The information percolator: Ambient information display in a decorative object. In *Proceedings of the 12th annual acm symposium on user interface software and technology* (pp. 141–148). ACM.
- Hellier, E., Naweed, A., Walker, G., Husband, P., & Edworthy, J. (2011). The influence of auditory feedback on speed choice, violations and comfort in a driving simulation game. *Transportation Research Part F: Traffic Psychology and Behaviour*, 14(6), 591–599. doi:10.1016/j.trf.2011.07.004
- Hermann, T., Hunt, A., & Neuhoff, J. G. (2011). *The sonification handbook*. Logos Verlag Berlin, GE.
- Hermann, T. & Ritter, H. (1999). Listen to your data: Model-based sonification for data analysis. *Advances in intelligent computing and multimedia systems*.
- Hill, J. O. (2009). Can a small-changes approach help address the obesity epidemic? a report of the joint task force of the american society for nutrition, institute of food technologists, and international food information council. *The American journal of clinical nutrition*, 89(2), 477–484.
- Ho, C., Spence, C., & Gray, R. (2013). Looming Auditory and Vibrotactile Collision Warning for Safe Driving. In *7th international driving symposium on human factors in driver assessment, training, and vehicle design*. Retrieved from <http://trid.trb.org/view.aspx?id=1263140>
- Hoby, M. (2012). *Touche for Arduino: Advanced touch sensing*. <http://www.instructables.com/id/Touche-for-Arduino-Advanced-touch-sensing/>.
- Holland, R. W., Hendriks, M., & Aarts, H. (2005). Smells like clean spirit: Nonconscious effects of scent on cognition and behavior. *Psychological Science*, 16(9), 689–693.
- Holland, S., Morse, D. R., & Gedenryd, H. (2002). Audiogps: Spatial audio navigation with a minimal attention interface. *Personal and Ubiquitous computing*, 6(4), 253–259.
- Holstius, D., Kembel, J., & Hurst, A. (2004). Infotropism: living and robotic plants as interactive displays. *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques (2004)*, 215–221. Retrieved from <http://dl.acm.org/citation.cfm?id=1013145>
- Horswill, M. S. & Plooy, A. M. (2008). Auditory feedback influences perceived driving speeds. *Perception*, 37(7), 1037–43. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18773726>
- Hussain, T., Gondal, M., Yamani, Z., & Baig, M. (2007). Measurement of nutrients in green house soil with laser induced breakdown spectroscopy. *Environmental monitoring and assessment*, 124(1-3), 131–139.

- International Transport Forum. (2007). Workshop on eco-driving: Findings and messages for policy makers. <https://trid.trb.org/view.aspx?id=1123413>. [last accessed: 07.08.2017]. International Transport Forum.
- Ishii, H., Wisneski, C., Brave, S., Dahley, A., Gorbet, M., Ullmer, B., & Yarin, P. (1998). Ambientroom: Integrating ambient media with architectural space. In *Chi 98 conference summary on human factors in computing systems* (pp. 173–174). ACM.
- Jackson, T. (2005). Motivating sustainable consumption: A review of evidence on consumer behaviour and behavioural change.
- James, W. (1890). *The principles of psychology*. Henry Holt and Company.
- Jenness, J. W., Singer, J., Walrath, J., & Lubar, E. (2009). Fuel Economy Driver Interfaces: Design Range and Driver Opinions (Report on Task 1 and Task 2). Retrieved from <http://trid.trb.org/view.aspx?id=1135367>
- Jensen, B. S., Skov, M. B., & Thiruravichandran, N. (2010). Studying driver attention and behaviour for three configurations of GPS navigation in real traffic driving. In *Proceedings of the 28th international conference on human factors in computing systems - chi '10* (p. 1271). New York, New York, USA: ACM Press. doi:10.1145/1753326.1753517
- Jeon, M., Davison, B. K., Nees, M. A., Wilson, J., & Walker, B. N. (2009). Enhanced auditory menu cues improve dual task performance and are preferred with in-vehicle technologies. In *Proceedings of the 1st international conference on automotive user interfaces and interactive vehicular applications - automotiveui '09* (p. 91). New York, New York, USA: ACM Press. doi:10.1145/1620509.1620528
- Jeon, M., Hammerschmidt, J., Hermann, T., Bazilinskyy, P., Landry, S., & Wolf, K. A. E. (2015). Report on the in-vehicle auditory interactions workshop: Taxonomy, challenges, and approaches. *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications - Automotive'UI 15*.
- Jeon, M. & Walker, B. N. (2009). “spindex”: Accelerated initial speech sounds improve navigation performance in auditory menus. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 53, 17, pp. 1081–1085). SAGE Publications.
- Jimenez, F., Aparicio, F., & Paez, J. (2008). Evaluation of in-vehicle dynamic speed assistance in Spain: Algorithm and driver behaviour. *IET Intelligent Transport Systems*, 2(2), 132–142.
- Johansson, H., Gustafsson, P., Henke, M., & Rosengren, M. (2003). Impact of EcoDriving on emissions and fuel consumption. In *International scientific symposium on transport and air pollution, avignon, france* (1000). Retrieved from <http://trid.trb.org/view.aspx?id=719723>
- Kahneman, D. (1973). *Attention and effort*. Citeseer.
- Kappel, K. (2009). *Persuasive technology: Ambient information systems for residential energy awareness* (Doctoral dissertation, Vienna University of Technology).

- Kappel, K. & Grechenig, T. (2009). "show-me": Water consumption at a glance to promote water conservation in the shower. In *Proceedings of the 4th international conference on persuasive technology* (26:1–26:6). Persuasive '09. Claremont, California: ACM. doi:10.1145/1541948.1541984
- Kim, S., Robson, C., Zimmerman, T., Pierce, J., & Haber, E. M. (2011). Creek watch: Pairing usefulness and usability for successful citizen science. In *Proceedings of the 2011 annual conference on human factors in computing systems* (pp. 2125–2134). CHI '11. Vancouver, BC, Canada: ACM. doi:10.1145/1978942.1979251
- Knowles, M., Scott, H., & Baglee, D. (2012). The effect of driving style on electric vehicle performance, economy and perception. *International Journal of Electric and Hybrid Vehicles*. Retrieved from <http://www.inderscienceonline.com/doi/abs/10.1504/IJEHV.2012.050492>
- Kuribayashi, S., Sakamoto, Y., Morihara, M., & Tanaka, H. (2007). Plantio. In *Proceedings of the international conference on advances in computer entertainment technology - ace '07* (p. 139). New York, New York, USA: ACM Press. doi:10.1145/1255047.1255075
- Kuribayashi, S., Sakamoto, Y., & Tanaka, H. (2007). I/O plant. In *Chi '07 extended abstracts on human factors in computing systems - chi '07* (p. 2537). New York, New York, USA: ACM Press. doi:10.1145/1240866.1241037
- Kuznetsov, S. & Paulos, E. (2010). Upstream: Motivating water conservation with low-cost water flow sensing and persuasive displays. In *Proceedings of the 28th international conference on human factors in computing systems* (pp. 1851–1860). ACM.
- Lansdown, T. C., Brook-Carter, N., & Kersloot, T. (2004). Distraction from multiple in-vehicle secondary tasks: vehicle performance and mental workload implications. *Ergonomics*, 47(1), 91–104. doi:10.1080/00140130310001629775
- Larsson, P. & Västfjäll, D. (2013). Emotional and behavioural responses to auditory interfaces in commercial vehicles. *International Journal of Vehicle Noise and Vibration*. Retrieved from <http://www.inderscienceonline.com/doi/abs/10.1504/IJVNV.2013.053818>
- Laschke, M., Hassenzahl, M., Diefenbach, S., & Tippkämper, M. (2011). With a little help from a friend: A shower calendar to save water. In *Chi '11 extended abstracts on human factors in computing systems* (pp. 633–646). CHI EA '11. Vancouver, BC, Canada: ACM. doi:10.1145/1979742.1979659
- Latham, G. P. & Locke, E. A. (1991). Self-regulation through goal setting. *Organizational behavior and human decision processes*, 50(2), 212–247.
- Leichsenring, C., Yang, J., Hammerschmidt, J., & Hermann, T. (2016). Challenges for smart environments in bathroom contexts. In *Proceedings of the 1st workshop on embodied interaction with smart environments* (p. 3). ACM.

- Leist, H.-J. (2001). Requirements for a sustainable drinking water supply—part ii: Side effects of water-saving measures. *Wasser und Abwasser-GWF-Munchen*, 142(10), 712–718.
- Leopold, A. C. et al. (1964). Plant growth and development. *Plant growth and development*.
- Liscum, E. (2002). Phototropism: Mechanisms and outcomes. *The Arabidopsis Book*, 51, 1–21.
- Locke, E. A. & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American psychologist*, 57(9), 705.
- MacIntyre, B., Mynatt, E. D., Volda, S., Hansen, K. M., Tullio, J., & Corso, G. M. (2001). Support for multitasking and background awareness using interactive peripheral displays. In *Proceedings of the 14th annual acm symposium on user interface software and technology* (pp. 41–50). ACM.
- Mankoff, J. C., Blevins, E., Borning, A., Friedman, B., Fussell, S. R., Hasbrouck, J., ... Sengers, P. (2007). Environmental sustainability and interaction. In *Chi'07 extended abstracts on human factors in computing systems* (pp. 2121–2124). ACM.
- Mankoff, J., Dey, A. K., Hsieh, G., Kientz, J., Lederer, S., & Ames, M. (2003). Heuristic evaluation of ambient displays. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 169–176). ACM.
- Manser, M., Rakauskas, M., Graving, J., & Jenness, J. (2010). Fuel economy driver interfaces: Develop interface recommendation. report on task 3. *Fuel economy driver interfaces*.
- Maréchal, K. (2009). An evolutionary perspective on the economics of energy consumption: The crucial role of habits. *Journal of Economic Issues*, 43(1), 69–88. Retrieved from <http://mesharpe.metapress.com/openurl.asp?genre=article&id=doi:10.2753/JEI0021-3624430104>
- Massa, G. D., Kim, H.-H., Wheeler, R. M., & Mitchell, C. A. (2008). Plant productivity in response to led lighting. *HortScience*, 43(7), 1951–1956.
- Matthews, T., Rattenbury, T., & Carter, S. (2007). Defining, designing, and evaluating peripheral displays: An analysis using activity theory. *Human–Computer Interaction*, 22(1-2), 221–261.
- Matthews, T., Rattenbury, T., Carter, S., Dey, A., & Mankoff, J. (2003). A peripheral display toolkit. *University of California, Berkeley Technotes, UCB//CSD-03-1258*.
- McCrickard, D. S. (1999). Maintaining information awareness with irwin. In *Proceedings of the world conference on educational multimedia/hypermedia and educational telecommunications*.
- McCrickard, D. S., Chewar, C. M., Somervell, J. P., & Ndiwalana, A. (2003). A model for notification systems evaluation—assessing user goals for multitasking activity. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 10(4), 312–338.
- Mekonnen, M. M. & Hoekstra, A. Y. (2016). Four billion people facing severe water scarcity. *Science advances*, 2(2), e1500323.

- Meyer, D. E. & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of experimental psychology*, 90(2), 227.
- Moere, A. V. & Offenhuber, D. (2009). Beyond ambient display: A contextual taxonomy of alternative information display. *International Journal of Ambient Computing and Intelligence (IJACI)*, 1(2), 39–46.
- Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly journal of experimental psychology*, 11(1), 56–60.
- Motorola. (2016). *The New Moto Z Family with Moto Mods: Transform Your Smartphone in a Snap*. <http://blog.motorola.com/2016/06/09/the-new-moto-z-family-with-moto-mods-transform-your-smartphone-in-a-snap/>. [last accessed: 26.06.2017].
- Mynatt, E. D., Back, M., Want, R., Baer, M., & Ellis, J. B. (1998). Designing audio aura. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 566–573). ACM Press/Addison-Wesley Publishing Co.
- Mynatt, E. D., Rowan, J., Craighill, S., & Jacobs, A. (2001). Digital family portraits: Supporting peace of mind for extended family members. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 333–340). ACM.
- Navon, D. & Gopher, D. (1979). On the economy of the human-processing system. *Psychological review*, 86(3), 214.
- Neumann, O. (1984). Automatic processing: A review of recent findings and a plea for an old theory. In *Cognition and motor processes* (pp. 255–293). Springer.
- Nilsson, A., Bergstad, C. J., Thuvander, L., Andersson, D., Andersson, K., & Meiling, P. (2014). Effects of continuous feedback on households' electricity consumption: Potentials and barriers. *Applied Energy*, 122, 17–23.
- Notzold, P. (2006). *TXtual Healing*. <http://www.txtualhealing.com/blog/?p=49>.
- NPSAPPS. (2014). *How Many times Do People Look at Their Phones?* <http://npsapps.com/how-many-times-do-people-look-at-their-phones/>.
- Obermair, C., Reitberger, W., Meschtscherjakov, A., Lankes, M., & Tscheligi, M. (2008). Perframes: Persuasive picture frames for proper posture. In *International conference on persuasive technology* (pp. 128–139). Springer.
- O'Shaughnessy, D. (2003). Interacting with computers by voice: Automatic speech recognition and synthesis. *Proceedings of the IEEE*, 91(9), 1272–1305.
- Palinko, O., Kun, A. L., Shyrov, A., & Heeman, P. (2010). Estimating cognitive load using remote eye tracking in a driving simulator. In *Proceedings of the 2010 symposium on eye-tracking research & applications* (pp. 141–144). ACM.
- Pallesen, K. J., Brattico, E., Bailey, C., Korvenoja, A., Juha, K., Gjedde, A., & Carlson, S. (2005). Emotion processing of major, minor, and dissonant chords. *Annals of the New York Academy of Sciences*, 1060(1), 450–453. doi:10.1196/annals.1360.047
- Pashler, H. E. (1998). *The psychology of attention*. MIT press Cambridge, MA.

- Pierce, J. & Paulos, E. (2010). Materializing energy. In *Proceedings of the 8th acm conference on designing interactive systems* (pp. 113–122). ACM.
- Plaue, C. & Stasko, J. (2007). Animation in a peripheral display: Distraction, appeal, and information conveyance in varying display configurations. In *Proceedings of graphics interface 2007* (pp. 135–142). ACM.
- Poupyrev, I., Schoessler, P., Loh, J., & Sato, M. (2012). Botanicus Interacticus: interactive plants technology. *SIGGRAPH '12 ACM SIGGRAPH 2012 Emerging Technologies*. Retrieved from <http://dl.acm.org/citation.cfm?id=2343460>
- Pousman, Z. & Stasko, J. (2006). A taxonomy of ambient information systems. In *Proceedings of the working conference on advanced visual interfaces - avi '06* (p. 67). New York, New York, USA: ACM Press. doi:10.1145/1133265.1133277
- Powell, N. & Lumsden, J. (2015). Exploring novel auditory displays for supporting accelerated skills acquisition and enhanced performance in motorsport. In *Proceedings of the 21th international conference on auditory display (icad-2015)*. University of Music and Performing Arts Graz. Retrieved from http://eprints.aston.ac.uk/26263/1/ICAD%5C_Paper%5C_FINAL.pdf
- Prante, T., Röcker, C., Streitz, N., Stenzel, R., Magerkurth, C., Van Alphen, D., & Plewe, D. (2003). Hello. wall–beyond ambient displays. In *Adjunct proceedings of ubicomp* (Vol. 2003, pp. 277–278).
- Raubitschek, C., Sch, N., Kozlov, E., Bernard, B., Schutze, N., & Baker, B. (2011). Predictive driving strategies under urban conditions for reducing fuel consumption based on vehicle environment information. *Integrated and Sustainable Transportation System (FISTS)*, 13–19. doi:10.1109/FISTS.2011.5973609
- Redström, J., Skog, T., & Hallnäs, L. (2000). Informative art: Using amplified artworks as information displays. In *Proceedings of dare 2000 on designing augmented reality environments* (pp. 103–114). ACM.
- Robinson, D. (2012). *ReplayGain 2.0 specification*. http://wiki.hydrogenaud.io/index.php?title=ReplayGain_2.0_specification.
- Rodenstein, R. (1999). Employing the periphery: The window as interface. In *Chi'99 extended abstracts on human factors in computing systems* (pp. 204–205). ACM.
- Romm, J. J. & Frank, A. A. (2006). Hybrid vehicles gain traction. *Scientific American*, 294(4), 72–79.
- Sato, M., Poupyrev, I., & Harrison, C. (2012). Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, 483–492. Retrieved from <http://dl.acm.org/citation.cfm?id=2207743>
- Schneider, W. & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. detection, search, and attention. *Psychological review*, 84(1), 1.
- Shen, X., Eades, P., Hong, S.-H., & Moere, A. V. (2007). Intrusive and non-intrusive evaluation of ambient displays. In *Ambient information systems*.

- Silberman, M., Nathan, L., Knowles, B., Bendor, R., Clear, A., Håkansson, M., ... Mankoff, J. (2014). Next steps for sustainable hci. *interactions*, 21(5), 66–69.
- Skog, T., Ljungblad, S., & Holmquist, L. E. (2003). Between aesthetics and utility: Designing ambient information visualizations. In *Information visualization, 2003. infovis 2003. ieee symposium on* (pp. 233–240). IEEE.
- Stasko, J. T., McColgin, D., Miller, T., Plaue, C. M., & Pousman, Z. L. (2005). *Evaluating the infocanvas peripheral awareness system: A longitudinal, in situ study*. Georgia Institute of Technology.
- Stasko, J., Miller, T., Pousman, Z., Plaue, C., & Ullah, O. (2004). Personalized peripheral information awareness through information art. In *International conference on ubiquitous computing* (pp. 18–35). Springer.
- Staubach, M., Kassner, A., & Fricke, N. (2012). Driver reactions on ecological driver feedback via different HMI modalities. *elib.dlr.de*, (October), 1–13. Retrieved from http://elib.dlr.de/78202/2/VP%5C_FS%5C_ITS%5C_HMI%5C_Modalities%5C_120523.pdf
- Sternberg, R. J. (2009). *Cognitive psychology*. Wadsworth Cengage Learning.
- Streitz, N. A., Rucker, C., Prante, T., Van Alphen, D., Stenzel, R., & Magerkurth, C. (2005). Designing smart artifacts for smart environments. *Computer*, 38(3), 41–49.
- Strengers, Y. A. (2011). Designing eco-feedback systems for everyday life. In *Proceedings of the 2011 annual conference on human factors in computing systems* (pp. 2135–2144). CHI '11. Vancouver, BC, Canada: ACM. doi:10.1145/1978942.1979252
- Sturm, J., Herwijnen, O. H.-v., Eyck, A., & Terken, J. (2007). Influencing social dynamics in meetings through a peripheral display. In *Proceedings of the 9th international conference on multimodal interfaces* (pp. 263–270). ACM.
- Styles, E. A. (1997). *The psychology of attention*. Psychology Press.
- Tanaka, H. & Kuribayashi, S. (2007). Botanical interface design -creative kits, tools, and methods-.
- Thaler, R. H. & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
- Tharp, S. & Tharp, B. (2005). *The Forecast Umbrella*. <http://www.materious.com/projects/forecast/>.
- Tomitsch, M., Kappel, K., Lehner, A., & Grechenig, T. (2007). Towards a taxonomy for ambient information systems. In *Ambient information systems*.
- Treisman, A. M. (1960). Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology*, 12(4), 242–248.
- Treisman, A. M. (1964). Verbal cues, language, and meaning in selective attention. *The American journal of psychology*, 77(2), 206–219.
- Tsimhoni, O. & Green, P. (2001). Visual Demand of Driving and the Execution of Display-Intensive in-Vehicle Tasks. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 45(23), 1586–1590. doi:10.1177/154193120104502305

- Tulusan, J., Soi, L., Paefgen, J., Brogle, M., & Staake, T. (2011). Which eco-feedback types prefer drivers most ?, 1–8.
- Tünnermann, R., Hammerschmidt, J., & Hermann, T. (2013). Blended Sonification: Sonification for Casual Interaction. In *The 19th international conference on auditory display (icad-2013)*.
- Tünnermann, R., Leichsenring, C., Bovermann, T., & Hermann, T. (2015). Upstairs: A calm auditory communication and presence system. Georgia Institute of Technology.
- Tünnermann, R., Zehe, S., Hemminghaus, J., & Hermann, T. (2014). Weather to go-a blended sonification application. Georgia Institute of Technology.
- United States Environmental Protection Agency. (2015). *Transportation Sector Emissions in the United States*. <https://www.epa.gov/ghgemissions/sources-greenhouse-gas-emissions#transportation>. [last accessed: 06.08.2017].
- U.S. Department of Energy. (2001). *Water Heating*. http://apps1.eere.energy.gov/buildings/publications/pdfs/building_america/26465.pdf. [last accessed: 19.02.2013].
- Van Dantzich, M., Robbins, D., Horvitz, E., & Czerwinski, M. (2002). Scope: Providing awareness of multiple notifications at a glance. In *Proceedings of the working conference on advanced visual interfaces* (pp. 267–281). ACM.
- Varbanov, K. (2013). *Engine's Torque Map*. http://f1framework.blogspot.de/2013/03/f1-engine-maps_28.html. [last accessed: 01.09.2015], used with permission from the author.
- Walker, B. N., Nance, A., & Lindsay, J. (2006). Spearcons: Speech-based earcons improve navigation performance in auditory menus.
- Watkins, K. et al. (2009). Human development report 2006-beyond scarcity: Power, poverty and the global water crisis.
- Weiser, M. & Brown, J. (1996). The coming age of calm technology. *Xerox PARC*. Retrieved July, 8, 2007.
- Weiser, M. (1991). The Computer for the 21st Century. *Scientific American*, 265(3), 94–104. doi:10.1038/scientificamerican0991-94
- Wickens, C. D., Hollands, J. G., Banbury, S., & Parasuraman, R. (2015). *Engineering psychology & human performance*. Psychology Press.
- Wierwille, W. W. (1993). Demands on driver resources associated with introducing advanced technology into the vehicle. *Transportation Research Part C: Emerging Technologies*, 1(2), 133–142. doi:10.1016/0968-090X(93)90010-D
- Wikipedia. (2009). *Brake specific fuel consumption*. https://en.wikipedia.org/wiki/Brake_specific_fuel_consumption. released under the CC BY-SA 3.0 (<http://creativecommons.org/licenses/by-sa/3.0/>) [last accessed: 09.08.2017].
- Woodruff, A., Hasbrouck, J., & Augustin, S. (2008). A bright green perspective on sustainable choices. In *Proceedings of the twenty-sixth annual sigchi conference on*

- human factors in computing systems* (pp. 313–322). CHI '08. Florence, Italy: ACM. doi:10.1145/1357054.1357109
- Young, M. S., Birrell, S. A., & Stanton, N. A. (2011). Safe driving in a green world: a review of driver performance benchmarks and technologies to support 'smart' driving. *Applied ergonomics*, 42(4), 533–9. doi:10.1016/j.apergo.2010.08.012
- Zehe, S. (2017). *BRIX₂: The Xtensible physical computing platform*. <http://tiny.cc/brix2>.