ReflectOns:
Mental Prostheses for Self-Reflection

by

Sajid Sadi

M.S. Massachusetts Institute of Technology (2006)
B.S. Columbia University (2003)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2012 [FEBRUARY 2013]

Author ....
Program in Media Arts and Sciences
School of Architecture and Planning
September 10, 2012

Certified by ....
Prof. Patricia Maes
Alex Dreyfoos Professor of Media Arts and Sciences
Thesis Supervisor

Accepted by..
Prof. Patricia Maes
Associate Academic Head
Program in Media Arts and Sciences

ReflectOns: Mental Prostheses for Self-Reflection

by

Sajid Sadi

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
On September 10, 2012 in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Since the time of the first philosophers, logic and observed human behavior have stood somewhat in contradiction. More recently, scientist have started to delve into decision making to understand why the way we act differs from rational choice, and indeed from our own desires. We believe that it is possible to use just-in-time feedback drawn from machine-observable behavior to help align behavior with personal goals.

This dissertation presents *mental prosthetics*, a model for distributed, embodied, design-embedded, just-in-time interfaces that augment the human judgment process. Drawing information from the activity of the user around them, mental prostheses analyze behavioral patterns in a way orthogonal to human cognition. Unlike persuasive interfaces, mental prostheses attempt to align choices with personal goals by cueing the user with just-in-time information. Lastly, these devices provide calm yet understandable feedback to draw the user's attention at the correct time to the information available to them.

This dissertation provides several prototypes and design explorations as a means of sampling the various approaches to data collection, synthesis, and feedback. Focusing on self-reflection, these sample designs form a subclass of mental prostheses that we term *reflectOns*. We show through the studies carried out in the course of this dissertation that these systems are effective in changing behavior to be better aligned with user goals. Lastly, this dissertation provides a set of design guidelines that assist in the creation of new mental prostheses. While we discuss a variety of scenarios in this work, it is only the beginning of the exploration. The design guidelines provide insight into both the critical aspects of the design of such systems, as well as possible input and feedback methodologies. These guidelines, together with the reflectOns themselves, provide a basis for future work in this area.

Thesis Supervisor: Pattie Maes
Title: Associate Professor of Media Arts and Sciences

Thesis Readers

Thesis Reader ....

Prof. Drazen Prelec
Digital Equipment Corp. Leaders for Global Operations Professor of Management
Sloan School of Management
Massachusetts Institute of Technology

Thesis Reader ....

Prof. Alex Pentland
Toshiba Professor of Media Arts and Sciences
MIT Program in Media Arts and Sciences

Acknowledgements

I begin my acknowledgements by thanking my PhD advisor Pattie Maes for her constant encouragement, support and feedback. It is not an overstatement to say that without Pattie's input, the seed that became the core of this thesis would have fallen fallow. She saw the possibilities for the idea of reflectOns and mental prostheses, and provided keen analysis and thoughtful feedback through the years to guide the design and study of the prototypes and the overall underlying theory presented in this dissertation. Pattie has supported me not only in this research, but in the many other enterprises I have embarked. Her willingness to fold this divergent work into the group's story, as well as her support and encouragement, attests to her dedication to the pursuit of new ideas and approaches. Lastly, her focus on real-world impact has shaped not only the direction of this thesis, but my own evaluation of what is important in the work that I do.

I also drew great inspiration from my committee members: Sandy Pentland and Drazen Prelec. Sandy's unparalleled history of work in using machine observation to tease apart human behavior is a core component of this thesis, and his keen insights into human behavior have been the high mark by which I have considered the positions taken herein. Drazen's seminal work in repeated choice, as well as his input about the possibilities surrounding the study of perturbations of human behavior forms the cornerstone of the studies I conducted. Indeed, the idea for the longest study conducted for this thesis was planted in my mind in our very first conversation. I would also like to thank Joe Paradiso and Bill Mitchell for being part of my generals committee. Bill and Joe have been friends and inspirations as I have designed and engineered the prototypes that ultimately bore the weight of this work.

Special thanks go to my colleagues in Ambient Intelligence and Fluid interfaces. They have been the enlightened and close-knit family that has supported me through the years, and has provided me with the fertile intellectual grounding in which to grown an idea into thesis. A big thank you also to my officemates and co-conspirators: people who have worked with me on crazy projects on every scale, and have helped define the meaning of accomplishment in research and art.

Special thanks to:

- Orit and Amit for teaching me what Media Art and Science really meant.

- Amon, Karen and Stephanie for teaming up to keep Studcom afloat.
- Enrico, David B, Federico, Marcel, Retro, Gutz, Hector, Aaron and Orkan for the crazy adventures across the world (now we really can do anything).
- Pol and Amna for always having that special positive outlook.
- Karen for fighting the good fight (despite the exasperation).
- Orkan, Doug, Kayato, Amit, Aaron, and Karen for joining me in pitching crazy ideas and philosophical discussion.
- Amna and Stephanie for helping with the logistics of this thesis.
- My colleagues at Tikatok and Samsung, for their support, friendship and understanding while I worked on this dissertation.
- My UROPs for their hard work and assistance, with special mention to Eben for his amazing skill in making several projects happen single-handedly.
- Linda, Pat, Gigi, Aaron and Bill for making sure I actually could administratively graduate, and caring not only about student progress, but for their wellbeing and the community as well.
- The ML professors, staff and sponsors for making it possible to do this work, and for providing the foundation of open discourse and real-world feedback that makes ML what it is.
- David G, Assaf, Oren, Vincent, Enrico, David M, James, Marcelo, Pranav, Seth, Cory, Natan, Susanne, Peter, Nan-Wei, Neri, and all my other friends.

My previous mentors have set me on the path that brought me to this goal. Thanks to late Dr. Brritton Chance for giving me responsibilities when in truth I knew very little. From him I learned about dedication to science. Thanks also to Prof. Steve Feiner, Benko and Alex for mentoring me through my research at Columbia. They taught me to imagine the full possibilities of the field that brought me to this juncture.

My parents taught me to be industrious, to love knowledge, to be civically minded, and to explore without fear. My mother Rukhsana taught me that every obstacle can be overcome, and my father Abdus taught me to never lose faith in a better tomorrow. My sister Ruheena has been my friend and foil, and ceaseless supporter through thick and thin. My parents and sister have offered me only love, encouragement and trust, and have held the faith that I would go on to do great things. I hope I have done well by their faith in me in thus far.

Finally, I am blessed in having the friendship, love, and support of my wife Farhana and her wonderful family. She has encouraged me when I was at a low, pushed me when I ran slow, and stood by me when the end seemed unattainable. It is fair to say that without her endless support, encouragement, and sacrifice this thesis would never have been completed. She is the fire that ignites my ambition, and the light by which I find my way.

In ending, I am certain I have missed some of you, but know that the mark you have left on me is indelible. **Thank you.**

Contents

Introduction

Since the time of the first philosophers, logic and observed human behavior have stood somewhat in contradiction. For example, one would be hard-pressed to find someone who doesn't know that being overweight is a significant threat to wellbeing. Yet, 60% of the US population is either overweight or obese (RedOrbit.com n.d.). Many a person will tell you that they waste a significant amount of time online. One place often pointed out is Facebook, and we know that users spend an average of 12 minutes per user per day on Facebook, adding up to about 10.5 billion minutes spent on just that one site in January 2012 (Darwell n.d.). Lastly, despite a landslide increase in energy awareness, and despite the increasing cost of energy, US energy use has actually increased since 2009 (US Energy Information Administration n.d.). While we have clearly been capable of logical, causality, and exemplar aware reasoning for millennia, it is trivially obvious that our intentions and our actions are often out of alignment.

More recently, first psychologists, and then behavioral scientists started an attempt to delve into the decision and choice-making process at various scales in order to understand why the way we act differs from how we should act, and in fact, from how we want ourselves to act. The formalization of this chain of investigation, focused on creating an applied theory of choice and decision making, eventually birthed the field of behavioral economics (Laibson & Zeckhauser 1998). The knowledge gleaned in the process was eventually was distilled to forms suitable for popular consumption in books such as Nudge (R. Thaler & Sunstein 2008) and Predictably Irrational (Ariely 2009). While these authors investigated the cause, those doing research on possible counteraction approaches created the fields of captology and persuasive interfaces (Fogg 2002), attempting to "correct" the aberrance.

Despite the abundance of theories, pragmatic examples, heuristics, and even solutions for individual cases of less-than-rational behavior, there has been little overall guidance for converting the lessons learned from the study of human behavior into intervention devices that can help individuals make better decisions. Indeed, the statistical nature of our understanding of quantitative human behavior has promulgated mass solutions for mass behavior rather than focusing on individual goals. Some work has been done in the realm of persuasive interfaces that have attempted to provide individual solutions

(Obermair et al. 2008; Seligman & Darley 1977; Fitbit Inc. n.d.), but these solutions as a corpus do not yet point to a model for information, intervention, or the elicitation of rational self-reflection.

In seminal works of behavioral economics, such as by Kahneman and Tversky (1986), we already start to see the seeds of intervention. While applying the tenets of psychology in many ways, the behavioral economists attempted to understand the processes of choice by perturbing it. Thaler and Sunstein (2008) later term these perturbations of choice as the design of "choice architecture": the set of circumstances around a choice that can unconsciously prod a person into acting in a particular way. There is a long legacy of perturbing choice to promote a particular extrinsic goal (Brownell et al. 1980; Seligman & Darley 1977). Separately, marketing sciences have perfected this same art for different purposes, drawing people to and fro in support or opposition of various goods, services, and ideas.

However, in many ways the idea of extrinsic choice architectures is a bit claustrophobic and confining, regardless of the cooperation of the parties involved. While it achieves the intended goals, the people subjected to these interventions are not instilled with a sense of control over their actions and choices, and must operate on faith. While the rails of choice architecture remain *invisible*, people are happy to roll along their predestined paths. When these rails are made visible to them, however, they yearn for their freedom to choose — regardless of the knowledge that they will not choose in a rational or optimal way. As such, extrinsically driven choice architectures that are made apparent to a person are not likely to be effective in changing behavior, and because of lack of overt participation, a person is not likely to self-train in a way that allows them to continue to act rationally once the training wheels of the choice architecture are removed. Therefore, there is a need to understand how choice architectures can both be made accessible to the individual, and how they can be made palatable once they are made visible.

Despite these aforementioned caveats of existing work, the trend is clear: it can generally be claimed that choice architecture and persuasion, when applied to an individual, can be effective in forming new behavior that better aligns with the goals that individual holds (Fischer 2008). While our minds may not be well-adapted to continuous and encompassing rational though, it likewise cannot be said that we are incapable of rational thought altogether. Rather, we are in capable of rationality within the limits of our attention and knowledge (March 1978; Gigerenzer & Goldstein 1996).

This dissertation presents *mental prosthetics*, a model for distributed, embodied, design-embedded, just-in-time interfaces that augment the human judgment process. This class of interfaces encompasses a variety of embedded systems that are designed to become a part of the normal usage of devices and objects in the midst of everyday activity as diverse as eating, using the microware or the hair dryer, or even simply surfing the web. Drawing information from the activity of the user around them, mental prostheses analyze behavioral patterns in a way orthogonal to human cognition, focusing on integration, unit conversion, and memory — tasks that a computer is sublimely suited to, but that our own cognitive processes can be led astray on. Unlike many of today's persuasive interfaces, these mental prostheses attempt to focus the user on making rational decisions at specific infection points in choice architectures where a rational judgment is most effective, thus reducing the mental load and points of failure induced by attempting to logically reason every decision. Lastly, these devices provide calm (Weiser & Brown 1996) yet understandable feedback using embedded means to draw the user's attention at the correct time to the information available to them. The design of mental prostheses draws on existing literature in the behavioral sciences to understand where and why to intervene, and from ambient, pervasive, and ubiquitous interfaces to design data collection and feedback methodologies that meet the goals of the design.

As part of this work, this thesis provides several examples of design in various states of completion, from completely functional devices such as the utensil attachment that can help people eat better, to the conceptual design of a device that can give the user better awareness of their social circle and interactions. These prototypes and design discussions are meant to provide a sampling of the various approaches to data collection, synthesis, and feedback that can be taken to help users make better decisions. These sample designs focused eliciting self-reflection are referred to collectively as *reflectOns*.

Three of these reflectOn prototypes, focusing on web surfing, eating, and energy use behavior, have been physically constructed and studied for effectiveness, and we have shown that they are effective in changing behavior, and additionally capable of maintaining and improving upon their effects during the duration of the study. The first of these prototypes looks at the time spent surfing the web, as well as the "information intake" from each page, providing the user with both physical world cues through a modified mouse as well as on-screen feedback to allow them to better stick with a time budget and surf more effectively. The second prototype tracks a user's eating speed while they eat with a modified utensil, providing a guard against fast eating, which has been implicated in improper eating behavior. Lastly, the energy use prototype converts per-machine instantaneous energy use information into a visual indicator that makes it easier to reason about the energy use of particular devices, and indeed to find places to act proactively to change energy consumption behavior. Each of these devices has been shown to be effective in reaching their stated goals.

In addition to the physical prototypes, we also provide a set of design studies that discuss the application of the tenets introduced in this dissertation to a variety of other domains. Ranging from financial decision-making to health care, these design studies provide additional insight into how reflectOns can be designed and used in further scenarios. They also provide validation of the wide-ranging applicability of the ideas discussed herein.

Lastly, this dissertation provides a set of design guidelines that assist in the creation of new reflectOns and mental prostheses. While we discuss a variety of scenarios in this work, it is clear that there are many more possible scenarios where mental prostheses can be of help. The design guidelines provide insight into both the critical aspects of the design of such systems, as well as possible input and feedback methodologies, as well as the requirements of the on-board synthesis that should be done to present the data back in its most actionable form. These guidelines, together with the reflectOns themselves, provide a basis for future work in this area as well as an organizing rubric for past and future work.

It is the goal of this work to provide a means for designers to draw upon the research and science done by behavioral science researchers in incorporating aspects of mental prostheses into their designs. To that end, we provide a clear list of possibly caveats involved in the way we make everyday decisions, and thus open up the possibility that common decisions can be analyzed for such issues and reflectOns designed to counteract the flaws. Additionally we also provide clear guidelines on what types of use cases reflectOns are well- or ill-suited for. In total, we provide a toolkit for diagnosing behavior that deviates from personal goals, and provide inspiration for how designs can be created to counteract such behavior while supporting the goals of the individual user.

Thesis Contributions

Since the time of Plato and his contemporaries, the term *logic* derived from the root Greek term *logos* – a word with complex sense meaning approximately "reason," but also "a body of knowledge and means of dealing with this knowledge" (Stein 1988). In this dissertation, we hope to embrace the latter definition, providing a holistic approach to introducing a behavioral intervention into an interaction.

The primary contribution of this dissertation is the introduction and characterization of a new design space for interactive devices: *mental prostheses*. Mental prostheses are objects that fit in as part of interactions that already take place in day-to-day life, and interact with the flow of the existing interaction with the goal of providing just-in-time information and situational awareness that allows the logical judgment process to reinstate itself in the midst of mindless action — ultimately leading to a change in behavior or awareness of the conditions surrounding some target choice. Unlike existing persuasive interfaces, mental prostheses are by definition individually motivated and focused. Unlike choice architecture-based behavioral solutions, these systems operate with the explicit awareness of the user, and in fact become part of the user's overarching choice architecture in an attempt to change their own behavior.

In support of the design of mental prostheses, the dissertation provides examples focused on eliciting logical reasoning in application to various domains, from eating to energy use to web surfing: the *reflectOns*. These design studies are done variously as both prototypes and as design concepts, and provide some insight into how the design of a mental prosthesis can be slotted into an interaction that already exists. Furthermore, they offer a view into the various forms of information analysis and synthesis, as well as feedback methodologies that can be used to understand and communicate with the user.

Drawing from the above, the final conceptual contribution of this dissertation is a set of design guidelines for mental prostheses. These guidelines are application agnostic, and provide a set of delimitations as well as a design manifesto for the creation of new experience scenarios around behaviors that people wish to change in themselves. With these guidelines, it is our hope that designers of both products and persuasive interfaces will have a new tool for understanding the existing research in behavioral science already done in attempting to understand human behavior, while also gaining insight into the various parameters that affect the design of the experience. Lastly, the manifesto for the design of mental prostheses is meant to incorporate the end user and the choices and desires of the end users, with the intent not of simply providing rails on which their decisions may ride, but giving them some say in the steering of their unconscious choice processes.

Dissertation Roadmap

Chapter 1 of the dissertation provides the general motivation for the work on mental prostheses and reflectOns. The general contributions of the thesis, as well the outline of preceding work that informs the design and implementation of the prototypes and theories in this thesis are outlined here.

Chapter 2 provides a summary of work related to this thesis. The chapter draws from 3 separate disciplines: cognitive psychology and behavioral economics, ambient and ubiquitous interfaces, and persuasive interfaces. The first area, cognitive psychology and behavioral economics, provides the foundation on which this work stands. We do not presume to lay new theoretical groundwork in the area of cognitive psychology. Rather, this thesis builds on the existing rigorously tested theories and provides a framework for perturbing cognition based on what has been discovered by the pioneers of these fields. From ambient and ubiquitous interfaces we draw inspiration for the sensing of user activity and the formatting of output. Finally, we quickly inspect the state of the art in using technology to change behavior, and draw together the main exemplars in the field to provide a sense of the work that has been done, and how that fits with the understanding of cognition that we provide at the beginning of the chapter.

Chapter 3 crystallizes the information from the background section with the lessons learned in producing and studying the prototypes to generate a set of recommendations for creating reflectOns in the future. This chapter also provides a summary the different caveats of rationality as it applies to the creation of reflectOns, and ends with a discussion of opportunities for investigation into mental prostheses in general in various applications.

Chapter 4 provides the design motivation and details of several reflectOns. Some of these have been fully realized into working prototypes for the study that we discuss in Chapter 5, while others are conceptual or partially-working design studies that show how the tenets of reflectOn design may be applied to various application areas.

Chapter 5 goes into details of the studies performed for this thesis, and shows what aspects of the design choices were successful. This chapter provides details of the studies and findings, as well as a discussion of general trends and convergences we see as a result.

Finally, Chapter 6 presents the takeaway guidelines for developing new reflectOns, as well as an abridged guide to approaches that may be interesting in designing mental prostheses in general.

Background

In a very simplistic sense, this research begins from the claim that people have cognitive biases and behavioral irregularities that prevent them from achieving their goals in particular (and yet very common) choice scenarios. In addition, work in computer science and interfaces has shown that there are ways of automatically understanding user behavior, as well as subtle ways of providing feedback based on such behavior. Given this tableau, we propose that providing the right sort of feedback at the right place and time can not only allow users to gain a deeper understanding of the information, but transform that information into actionable choices that can affect goals. The challenge lies in gaining an understanding of how the information gleaned from decades of intensive study can be turned into practical systems that can actively affect goal alignment of behavior. In this section, we shall explore these underlying claims, as well as competing approaches and relevant related studies that support the designs presented in the applications section of the thesis.

The work presented in this thesis is a design approach that is informed by a large body of work. On one hand, it draws from psychology and behavioral economics to provide the detailed understanding of human behavior. In particular, the tenets of decision theory, cognitive artifacts, and theories concerning repetitive choices come into play. On the other hand, it draws from the disciplines of ambient interfaces and ubiquitous computing to create a set of tools that can used as a means of changing behavior. These disciplines inform both feedback design and general tenets of providing feedback in a data and feedback-dense environment. This section will begin with an exploration of both of these bodies of work. Proceeding from this overview, we shall present a number of studies that support the various design prototypes created for the thesis.

Understanding Human Behavior

Behavioral economics grew out of a reaction to a normative theory of economics. Normative theory effectively applied a rational (or logically sound) model to human behavior, and thus entertained several tenets that eventually came under criticism when applied to observations of common human behaviors. The four central aspects of normative theory are cancellation, transitivity, dominance, and invariance (Tversky & Kahneman 1986).We will begin with these tenets as well as a discussion of the effects that tend to counter such logical action in data-to-day life, drawing from descriptions in Tversky et al (1986).

Invariance. A tenet shared by all normative theories, invariance can be understood simply as the axiom that the way a choice is presented — the working or presentational circumstances of the choice — does not affect the choice if the various representations mean the same thing. That is to say, if after considering it logically the person making the decision could translate choice A into choice B, the person should choose either choice in equal proportion. It is tacitly assumed in most cases, including when one speaks of the choice as a random variable, since the statistical requirements of random variables requires invariance (Tversky & Kahneman 1986). Studies concerning the framing effect show that this is often not the case, and framing effects such as loss avoidance, or the propensity to choose a risky win over a sure loss, cause people to choose not only based on outcome, but also based on the framing of the choices (Tversky & Kahneman 1981). These effects are consistent and sizeable across all manner of choice scenarios.

Dominance. Dominance is a requirement that if a choice does better than the rest in one condition, and at least as well as every other choice in all other conditions, then it is the "dominant" choice and would be preferred. However, studies show that dominance applies to real choice scenarios only when the application is transparent. When obfuscated by mixed wins and losses, the dominant choice can be overlooked.

Transitivity. Transitivity can be understood as the requirement that all choices be convertible to a scale without concern for the other choices. Another way of approaching transitivity is to say that the utility of a choice is independent of other choices. Approaches such as regret theory (Loomes & Sugden 1982) posit conditions and studies which show that people often make choices based on not only the choice selected, but the utility of the "virtual loss" of the other choices as well, providing an alternative approach to loss of transitivity.

Cancellation. A basic tenet of utility-based theories, cancellation refers to the independence of the value of an outcome from any conditions within the world that ends up producing the same end result. Alternatively, the probabilities of choices should not affect the selection — only the value should. However, studies show that cancellation is sometimes ignored when the person perceives one choice as more of a certainty than another, which in turn triggers risk aversive behavior.

Behavioral Economics and the Cognitive Psychology of Choice

Behavioral economics, for the most part, modifies the tenets of normative theory to operate within certain conditions, and takes into account lessons from behavioral psychology to better understand the issues that affect the practical operation of people given various choice scenarios. Mullianathan and Thaler summarize the problem succinctly by noting that normative economics depends on three factors that cannot be rationally expected to exist: unbounded rationality, unbounded willpower, and unbounded self-interest (Mullainathan & R. H. Thaler 2000). Logically, one can reason that there are bounds to rationality simply because any single person has limited computational powers, time to acquire information, and time to reason about that information. More recently, work in behavioral sciences also shows that willpower is similarly limited, and may in fact be a limited resource (Vohs & Heatherton 2000). Research also indicates that there are times when blind self-interest is not the only guiding factor in the behavior of an individual. While it may appear that averaged across the population these individual variations would smooth out toward the optimal predictions of normative theory, the actual observations indicate that in fact these individual difference not only continue to exist, but are often amplified (Mullainathan & R. H. Thaler 2000). As a result, it is better in the large-sample case, and essential in the individual case, to reason about behavior with the limitations in mind.

At the inception of behavioral economics, Herber Simon wrote that while normative economics provided a baseline understanding and model for behavior at the macro scale, as we delve deeper into the individual actions and the minutia of the decisions, it was increasingly necessary to pull in aspects of observer micro-scale behavior and psychology of decision making in order for the model of the behavior to be valid (Simon 1959). Drawing from a large variety of practical examples, he argued that as our tools grew more fine-grained in both observation and analysis, the behavioral aspects of decision making and judgment also became more profound. Simon also suggested the use of digital computers as a means of analysis of the behavioral data, foreshadowing the future work that serves as a basis of this thesis. In speaking of behavior, we shall thus tread the line between psychological and experimental approaches in describing various cognitive artifacts.

A basic tenet of behavioral economics and behavioral psychology, bounded rationality divides human judgment into two forms: intuition and reasoning (Kahneman 2003). These forms of reasoning have also respectively been named reflexive and reflective thought according to their apparent means of activity. These systems are also referred to in a more neutral manner as system 1 (intuitive) and system 2 (reflective). For the purposes of this discussion, reasoning is the set of tools that lets us multiply two numbers, or follow the various machinations of a Rube Goldberg machine. On the other hand, intuitive thought can be used for example to select the largest of 12 apples at a glance. The table below outlines some of the aspects of intuition and reasoning. As can be seen, the reflective process is slow and effortful, while the intuitive system is fast and effortless. Systematic investigation indicates that most thoughts and actions are based on forms of intuitive thought (Kahneman 2003). There appears to be a level of monitoring from the reflective system that takes place to verify that the intuitive decisions are valid. This monitoring, however, is quite lax, and especially under time pressure and repeated application, the intuitive system can bypass the reflective monitoring and express itself regardless of actual correctness (Kahneman & Frederick 2002). Due to the difference in effort between the two systems, the attribution of a task to a system is proportionate to effort. Since the intuitive system is more effortless, it cannot be as easily "distracted" by additional tasks, and this can be tested with dual-task studies. The monitoring function has been shown to belong to the reflective system by such studies (Gilbert 1989).

Intuitive	**Reflective**
Automatic	Controlled (Requires conscious application)
Effortless	Effortful
Associative (Heuristic)	Deductive (Rule-based, stepwise)
Rapid and parallel	Slow and serial
Process is opaque	Self-aware (Process is visible)
Skilled action	Rule application

Table 1 The two cognitive systems, derived from (Kahneman & Frederick 2002).

The result of several decades of experiments and modeling indicates that most judgments are made intuitively and that the rules that guide intuitive judgment have more to do with rules of perception than rules of logic (Kahneman 2003). In his 2003 paper, Kahneman compared perception and intuitive thought using a diagram replicated in the figure below. In comparison, it becomes clear that the perceptual and intuitive systems have a great deal of similarity to each other. Additionally, it can be said that the perception-oriented intuitive system can come to impressions that do not have explicit vocalizable definitions, while the reflective system is always explicit about such definitions. In other words, intuitive judgments tend not to be able to explicitly answer "why did I do that?" However, the process used by the intuitive system is not limited to only current perceptions, but also previous perceptions, memory of those conditions, and heuristics about the choice scenario. This reliance on habits and memory is what makes intuitive judgments consistent over time and across similar choice scenarios.

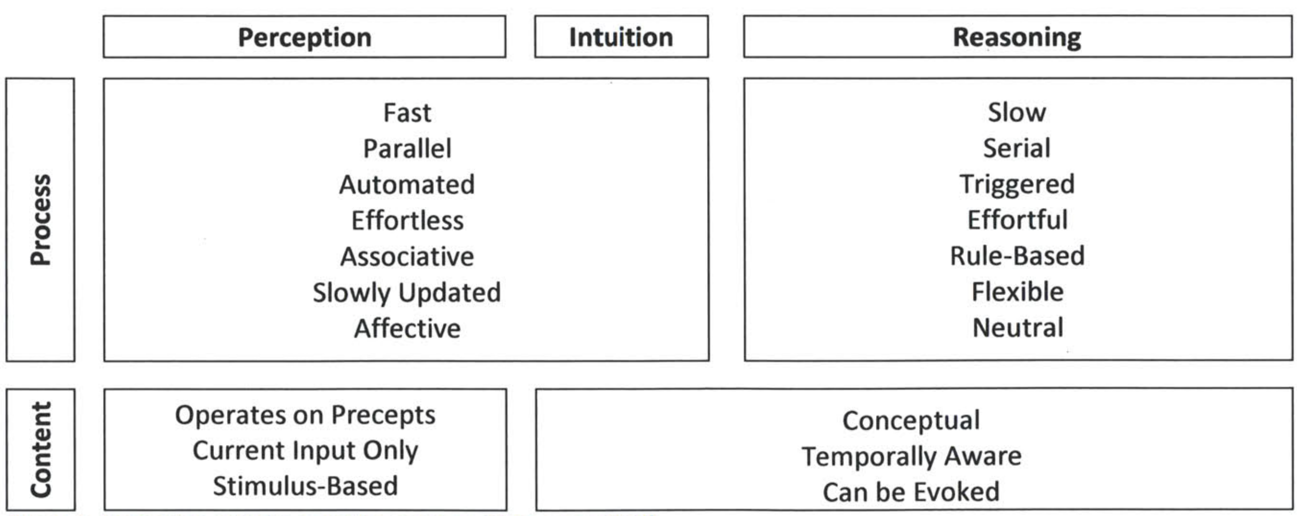

Figure 1. Perception and the cognitive systems (Kahneman 2003).

Given this dual system and the precept that for most of the time one reasons using the intuitive system to minimize effort and maximize speed, errors in any choice scenario can be thought of as two separate questions: (a) what error in intuitive reasoning caused the error in judgment, and (b) what error in the reflective system prevented it from noticing the error. Of course, given that reflective processes are possibly a limited resource (Vohs & Heatherton 2000), it may well be the case that the answer to latter is simply resource depletion or overuse, itself caused by the single-track nature of the reasoning system. Additionally, it has been suggested that heuristic use is a deliberate strategy used to reduce loading on the reflective systems (Gigerenzer & Goldstein 1996). Regardless of the answer to the latter part, the initial problem lies in determining the cause of the error of intuition.

One of the primary characteristics of intuitive thought is the spontaneity of the thought, which is called the *accessibility* of the content or concept (Higgins & Brendl 1995). Accessibility is guided by perception — a claim that can be easily tested. For example, a linear arrangement of 27 cubes gives increased accessibility to the number of items, but not their volume. The same 27 blocks arranged into a 3x3x3 cube gives greater accessibility to volume, but not the number. A random arrangement of the blocks on a flat surface makes neither number nor volume accessible. If a complex choice has an accessible attribute, we tend to use the accessible facet as a substitute for the overall complex question, and often without realizing that we have done so (Kahneman 2003). This process is known as *attribute substitution* — the act by which target attributes are replaced with heuristic attributes for the needs of a judgment. Some of the attributes are also automatically calculated, while others must be computed by the reflective system on demand. These automatic attributes include both physical properties like size and loudness, and mental states such as affective valence (good or bad), how surprising (or causal) an event is, the elicited mood, etc (Kahneman & Frederick 2002). These automatic attributes, called *natural assessments* by Tversky and Kahneman, are thus frequent candidates for becoming the heuristic attribute in a choice. A partial list can be found in Kahneman and Frederick, 2002.

In their discussion of attribute substitution, Kahneman and Frederick discuss a number of possible biases. The substitution and computation of heuristics introduce a variety of biases into the process. The most common of these is a weighting/duration bias, where the weight or duration (and thus probability) of a sample is misconstrued. Additionally, as the heuristic attribute is converted to the target attribute, there can be ratio errors introduced in the mapping between dimensions that may in fact have little direct correlation. This is further complicated when one of the scales has no upper or lower bound, and the other does. For example, there is no direct upper bound to authoritativeness, but there is to height. If one uses height as a heuristic attribute for authoritativeness, there is a mapping error that is amplified arbitrarily at the extremes, in addition to being moderated by the experiences of the person making the judgment. The role of affect has become increasingly apparent, and the affective valence of an attribute (is being tall good?) has an effect on the judgment independent of the direct mapping of attributes (Kahneman & Frederick 2002). In addition, factors such as salience and physicality of an attribute affect its selection as the heuristic attribute, but these factors can be compensated by priming or attention.

Lastly, it should be noted that while an attribute may be interpretable in multiple ways by the reflective systems, the intuitive system does not experience doubt about these sample attributes. As a result, whatever decision the intuitive system arrives at is generally absolute (Kahneman 2003).

While visiting the topic of heuristic attributes, it is also salient to touch upon the idea of *prototype heuristics*, more commonly understood as stereotypes (Rothbart et al. 1978). Psychological theories suggest that whenever a set of items exist for which a stereotype if formed, the attributes of the stereotype are pre-computed and automatically available. In fact, it is fair to say that the stereotype is in fact the collection of pre-computed attributes of the entire set given form. It is characterized by salient properties and averages, and allows us to quickly categorize new input. Since calculating average attribute matches in a question of the form "Is X an instance of category Y?" requires harder calculations, it follows from the above discussion that X may be replaced by its stereotype (Kahneman & Tversky 1972). As a result, data points that match the stereotype will implicitly increase the match. On the other hand, the matching for attribute similarity causes the neglect of the actual group attributes, and the subsequent lack of updates to the stereotype (Kahneman 2003).

In addition to the above, several other forms of information neglect commonly affect the intuitive decision process (Kahneman & Frederick 2002). *Scope neglect* causes a person not to scale values with the size of the effect. *Violations of dominance* cause additional state to otherwise equivalent choices that should not pragmatically affect values to still cause changes. Finally, *duration neglect* substitutes the actual integrated value over time with maxima and recent memory (Mehta et al. 2004). These artifacts all countermand the needs of normative theory at the individual scale. It should be noted that depending on the particular scenario, the reflective system can correct such artifacts. However, the reflective system needs to be cued to watch for such aberrations in order to be able to detect them with any regularity.

Having established a link between perception and intuition, we can now look to some of the characteristics of the perceptual system as a means of accessing characteristics of the intuitive system. One such characteristic is that the perceptual system is designed to enhance contrast. That is to say, it highlights difference and changes, and is thus reference-dependent since differences and changes are always relative to some perceived baseline (Kahneman 2003). Given this, it thus follows that the presentation or tableau of choice would affect the way the choice itself is perceived — a hypothesis that has been experimentally verified and codified as *prospect theory* (Kahneman & Tversky 1979). In cases where there is some uncertainty and risk, it is shown that people become loss-averse even if the long term gain suggests that the person take the risk. In fact, prospect theory generalizes to cases where there is no risk, but retains the characteristic that losses are counted more steeply than gains are valued. As such, a person with a free mug will demand more for the mug than they would have been willing to pay to get it. However, studies also show that the effect can be ameliorated with training (Kahneman 2003). It should be noted in closing that the apparently-opposite effect of aversion to temporary "out of pocket" loss also exists, where the virtual cost of the "lost opportunity" counts against the value of a choice, and is described as regret theory (Loomes & Sugden 1982).

Revisiting the concept of a perceptual reference, *framing effects* describe the violation of normative invariance. That is to say, framing effects describe the effect that extrinsic considerations such as the description of a choice have on the choice (Tversky & Kahneman 1981). For example, a choice with a negative affective valence certain choice will drive people to choose a risky choice, while the same problem stated with a positive valence certain choice will cause people to choose the less-risky choice. This would be the case even when there would be greater gain when choosing otherwise. An extension of this effect is the default effect, where the default choice enjoys a sizeable advantage simply by merit of being the default. As described by Kahneman:

> *The basic principle of framing is the passive acceptance of the formulation given. Because of this passivity, people fail to construct a canonical representation for all extensionally equivalent descriptions of a state of affairs.*

As previously mentioned, this violation of invariance raises significant questions about the applicability of normative theories to the realm of individual choices. In particular, normative rational agents are characterized by an all-inclusive framing that does not appear to exist is practical scenarios (R. Thaler & Sunstein 2008).

While the intuitive system provides judgment automatically in opposition to the explicit effort of the reflective system, the reflective system does to some extent automatically monitor the intuitive judgments. As mentioned previously, this monitoring is normally lax, but can be cued to monitor a situation more aggressively. The failure of the reflective monitoring elicit three possible vectors for the error to appear: (a) the person lacks sufficient information to make the judgment — for example if they are comparing the populations of two little-known towns, (b) the person perceives but does not apply the data provided by corrective cues — for example when a person misjudges distance not only because of perspective, but also atmospheric haze, and (c) the cues we believe to be the important are not the ones that are attended to by the person — or simply, the model for cueing is not matched to the accessibility of the cues (Einhorn & Hogarth 1981). Any judgment, as discussed previously, begins from perception, or more generally from acquisition of information from both memory and senses. Attention is central to the process of acquisition, and the way that cues are placed, and whether they are integrated (and thus likely to trigger stereotype-based reasoning) or separated (leading to per-attribute reasoning) has a large effect on how salient both reasoning systems consider the information (Lynch & Srull 1982). Thus, for example, in a Stroop test the word "red" written in blue will cause some mental effort, but the correction will be correctly triggered, since the salient information is immediately available and not hidden away. Therefore, a critical aspect of the triggering of the monitoring system is the bringing of cues within the attentional scope of the task (Einhorn & Hogarth 1981; Kahneman 2003).

While we previously mentioned that the update of stereotypes is generally lax, that is not to say that the effect of learning is negligible. In fact, the formation of stereotypes themselves is a form of learning that is facilitated by the intuitive system (Rothbart et al. 1978). In learning, the most concrete and verifiable results come from experimentation, while the least reliable results come from unaided judgment (Einhorn & Hogarth 1981). In the most common situations, factors such as the lack of control over variables and the lack of grounded results lead to the widespread use of unaided judgment. This situation is further exacerbated by the fact that most individuals are quite unaware of the pitfalls of learning from such unaided decisions. Brehmer further expands on this by noting that (a) we generally don't know when we should be learning, (b) given we do know, we do not know what we should be learning, and (c) when we have learned something, we are not necessarily aware whether we have learned, and what exactly we have learned (Brehmer 1980). Thus the importance of guidance and cues that trigger learning and provide mileposts to clarify the above concerns becomes paramount in attempting to use learning to overcome cognitive and judgment biases.

At the same time, experiments indicate that evaluation of judgments based on experience differ from evaluations of judgment based on information alone (Hertwig et al. 2004). In these experiments, subjects tended reliably to overweigh rare events when acting from descriptions, while underweighting rare events when acting from experience. Facets of uncertainty aversion as well as recency effects (Einhorn & Hogarth 1981) may form the underlying cause of this phenomenon. As a result, learning from descriptive sources alone may create a heuristic that is inapplicable to the practical scenario.

In addition to the importance of cueing for the triggering of the monitoring functionality, cueing also plays a large part in the accessibility of memory (Lynch & Srull 1982), which in turn powers the intuitive system. An otherwise salient heuristic thus may be easily obscured by a cue that causes other data to be "loaded" into the set of accessible memory, which is a small subset of the things actually remembered by a person. Since the accessible memory is relatively small, this provides an additional opportunity for distraction for the reflective system. It is important here to highlight the difference between recognition (based on stereotypical processing) and recall. While recall can be significantly affected by framing and organization, recognition is generally much less affected (Du Plessis 1994), as can be inferred by our previous discussion of prototype-based intuitive thinking. This provides an opportunity to use habituated recognition as a proxy for recall or reflection, since it is much harder to derail that process compared to the other two. Additionally, the framing of the task and the importance of recall as stated in the task is itself a factor in recall, while both this and associative cues do not affect recognition (Lynch & Srull 1982).This highlight the importance of the entire intervention experience as a while (both education and cueing) as a holistic approach to eliciting behavioral change.

While manipulations of the intuitive system can be used to elicit behavioral change, a more direct approach is to activate the reflective system directly. One way of codifying such activations is to approach it as reflective attention, in opposition to perceptual attention discussed previously. The primary matter of note in reflective attention is the limited reservoir of attentional capability available.

This has been discussed both in terms of "short term memory" and "limited attention", but for our purposes the terminology and approaches do not require separation. Kahneman's initial work in this area shows that the intensity and capacity are related, in that more capacity is allocated as task becomes harder (Kahneman 2003). More recent work shows that the reservoir is encompassing in nature, and is slowly refilled (Vohs & Heatherton 2000). As a result, it is not possible to apply maximum capacity to all tasks throughout the day. Even when capacity is available, reflective attention is selective (Lynch & Srull 1982). Lynch et al note that the most consistent result in this aspect is the importance of novelty, or a sensitization effect that causes reflective attention to diminish with repeated exposure. Thus in any intervention, maintaining some level of novelty is necessary in order to access reflective attention.

An interesting aspect of bounded rationality is the illusion of control: the belief that events — even statistically and causally unrelated ones — are under the control of the person making decisions. Situational factors such as personal involvement, familiarity, foreknowledge, and success at the task, as well as the mood of the person and their need for control over the situation, serve as modulating factors (Thompson 1999). As a result, most conditions where the person is attempting to overcome a pernicious habit fall under the rubric of illusory control, since with sufficient attention, the person can maximize all the modulating factors to produce the illusion, but due to depletion of reflective attentional resources, cannot maintain control. Due to various cognitive biases, including variants of loss aversion and heuristics-based prototype substitution, as long as chance aligns with goals with some critical probability (Langens 2007). Studies suggest that while illusory control is not a negative condition per se (Thompson 1999), it does regardless have consequences when the control is overestimated, preventing people from using cueing and other adaptive techniques to bolster reflective monitoring and increase actual levels of control.

In closing the section on behavioral economics and cognitive psychology, we would like to end with a short discussion of cognitive dissonance, a concept introduced by Festinger to describe to adjustment of beliefs to align with choices made (Festinger 1957). The preceding discussion of a behavioral view of judgment requires a task-oriented view of decisions over time, with various heuristics and belief sets being applied to decisions. As a result, it is natural that a person may have multiple cognitions that are internally self-consistent (with consideration to intuitive judgments), yet in opposition to each other. Studies show that we tend to like what we chose or have, and dislike those choices we rejected (Vroom 1966). This is consistent with propositions posited by regret theory, in that choices not made produce a utility drain on the choice actually made, and thus it is expected that we attempt to minimize the contradiction by realigning beliefs (Loomes & Sugden 1982). As Dan Ariely once noted while speaking of his book, "we will do whatever it takes to be able to look at ourselves in the mirror in the morning" (Ariely 2009). As a result of dissonance minimization, any intervention faces an uphill struggle in finding purchase within the belief system of the subject, and a large portion of the learning effects around an intervention thus may have to be directed not at behavioral effects, but simply on acknowledging the dissonance.

Theories of Repeated Choice

Thus far we have discussed choice as individual or independent actions. However, many choices in everyday life (and by definition, all habitual choices) are repetitive in nature. From the above discussion on monitoring functionality and limited reflective resources, it follows that repetitive choice has its own idiosyncrasies. In this section, we shall discuss some of the theories of repeated choice, as well as further effects that arise from such choice scenarios.

	Effect	Action	Example
Multiple gains	Since the gain function tapers off, additional gains feel smaller.	Multiple gains are segregated.	Multiple small gifts feel better than a single big one.
Multiple losses	The higher slope of the loss function near the origin makes many small losses considered together worse than a single loss of the same value.	Multiple losses are desegregated.	The aggregate loss of a credit card bill is better than the cash losses of the individual purchases.
Small loss, bigger gain	Given the steeper slope of the loss function, an equivalent gain and loss will feel like a loss, so the gain is used to "cancel" the loss.	Losses cancel larger gains.	Any trade: loss of money is cancelled against gain of purchase, but purchase counts as single unit of gain/loss.
Small gain, bigger loss	If the loss curve is flattened at the point of loss, then segregate and get a "silver lining". Otherwise, integrate for same reasons as "small loss, bigger gain".	"Silver linings" to loss situations are extracted out.	Points and rewards in credit cards. The rewards are small compared to the loss but compared separately.

Table 2. Accounting of various gain/loss scenarios, partialy derived from (R. Thaler 1985).

Repeated choices can be broken down into two categories: (a) choices that are each individual, with their own results, but of the same type, and (b) "distributed choices" where the aggregate of many smaller decisions lead to the final outcome (Herrnstein & Prelec 1991). Of course, type (b) choices also have individual feedback from the individual judgments, but in terms of goal adherence, these smaller judgments contribute to the overarching goal. For example, a purchase at a store has its innate utility, as well as the utility (or disutility) it contributes to the goal of saving money. In order to understand repeated choices, it is first necessary to first discuss how the human mind keeps track of the value of actions and events — the realm of mental accounting. As discussed before when touching on the concept of loss aversion, people tend to perceive losses as more injurious than they consider gains to be helpful (Kahneman & Tversky 1979). The theoretical curve of perception of losses and gains is shown in the figure below. As a result of the kink in the curve as well as the steeper curve of the loss function, combinations of gains and losses are treated differently, summarized in the table below. As a result of how the value of distributed choices is calculated, the sum of the distributed choice is rarely the sum of the value of the choices themselves.

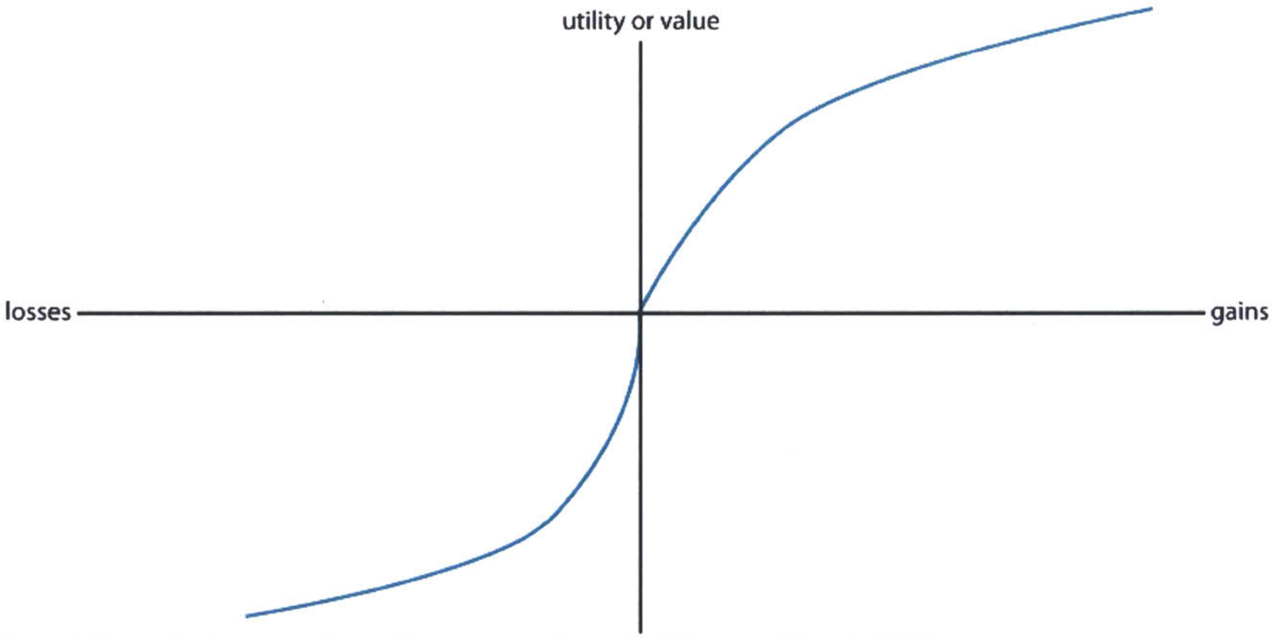

Figure 2. Theoretical curve of gains and losses vs. perceived cost (Kahneman & Tversky 1979).

Building on the idea of mental accounting, the theory of melioration gives insight into the effects of cognition on estimating value of repeated choices. This and related theories give insight into how the value propositions of distributed choices are occluded by cognitive processes. In order to understand the proposition of this theory, it is instructive to consider a distributed choice scenario. In this scenario, Alice must choose between two different (but equal-cost) meals each day: pizza and sandwiches. She prefers pizza by a factor of 2:1. If we were to look inside the intuitive system that determines that utility of a meal, we would find that the system holds on to the last n meals, and discounts repeated choices of either pizza or sandwich when it is chosen repeatedly. This is shown is the figure below. While this is a common — albeit simplified — scenario, it has some features that make maximization difficult, because the intuitive system, as discussed previously, is opaque to the rational system and indeed to itself (Kahneman 2003). Specifically, it is hard to determine what incremental change a choice will elicit, how it will affect long-term characteristics, and given these preconditions, what the curve of the value function may be (Herrnstein & Prelec 1991). Melioration theory states that while people know what they like or dislike, they don't know how much their selection of other options affect their likes and dislikes. Because they cannot see both distributions simultaneously, the result is generally a suboptimal distribution based on mental accounting principles. In experiments testing this scenario, it has been shown that the lack optimality can degenerate to the worst case, and the probability of choosing the choice that lowers overall value more often is high.

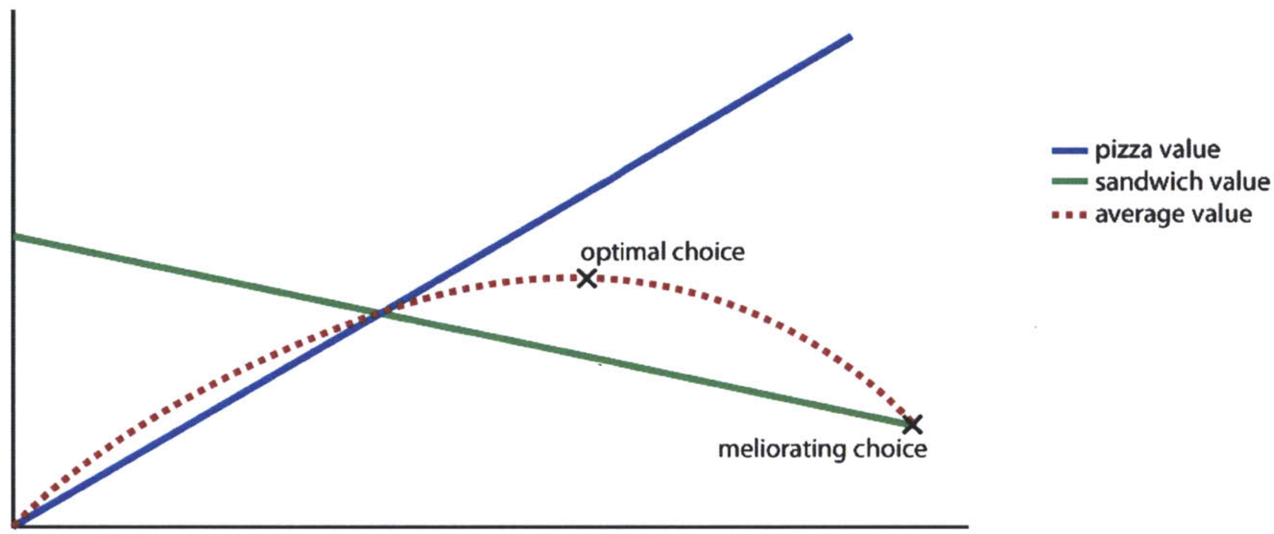

Figure 3. Utility of distributed choices (Herrnstein & Prelec 1991).

A meliorating consumer stabilizes their choices at a point where either one alternative absorbs the entire set of choices (corner case), or the value of all choices become equivalent (matching state). While this sounds rational a priori, there are a number of pathologies that affect meliorating choices. In cases where attentiveness to choices is a minimum (overeating, overworking, etc), these pathologies are at their maximum. In the corner cases they lead to addiction — that is, overindulgence in a choice — or in the more common cases lead to a general sense of unhappiness where the person is aware of the unhappiness, but not the source or corrective measure. A simple example of this is overworking — persons who overwork are unhappy that they work so much, but do not necessarily know how to reverse the situation. Simply choosing not to work temporarily causes the person to feel unhappier at getting nothing done, and thus makes the meliorating condition stable. The "correct" choice would be some combination of work and leisure depending of the person's value curve, but is difficult to reach without understanding of the overall value curve. It is a necessary condition of melioration that choices whose values increase with increase in action frequency will be under-consumed (Herrnstein & Prelec 1991), and thus leisure will be under-consumed if the person is already meliorating at an over-work state. In the corner case of addiction, the value of other choices may be reduced by the effects of the addictive choice, such that if overwork is an addiction, then the value of more work is nearly zero, but the value of leisure is lowered by the feelings that one is perhaps being lazy. In such conditions, an alternative such as a hobby — work that is leisurely — can help.

In closing this section, we would like to mention some additional experimental and theoretical data points. Experimental results suggest that the immediate reward is a strong trigger for melioration regardless of secondary delayed feedback (Tunney & Shanks 2002). This may be an effect of temporal discounting of distributed rewards (Rubinstein 2003), or simply a lack of self-control against the immediacy of the reward, but tend to exacerbate the selection of suboptimal distributions. It should also be noted that several other theories, including kinetic (Myerson & Hale 1988) and quasi-dynamic (Staddon 1988) models, have been proposed to address the observed meliorating behavior. While they differ in the exact formulation, all recognize the underlying behavioral anomaly of non-optimal distributed choice scenarios. The kinetic model provides an interesting insight into the reinforcement and learning, showing that the rate of reinforcement has a strong effect on switching rate, which in turn affects the probability of breaking from a meliorating condition. However, training effects are negatively accelerated, suggesting that it is harder to break a meliorating habit than to create one (Myerson & Hale 1988). The known certain cure to melioration (and perhaps an unreachable optimum) is to provide greater visibility of the value curve to the individual, thus allowing better rational supervision of the habitual meliorating choice.

Opportunities for Intervention

With the aforementioned background in behavioral theory, we can now see two avenues via which judgment errors make their way into our behavior. Figure 4 shows the process by which we convert perception to decision. In short, the world is perceived by out sense, and the intuitive reasoning system, in parallel, converts these perceptions to decisions using the means previously discussed. The reflective system monitors this process to the extent possible, though as discussed, the difference in processing speed between the parallel intuitive system and the serial reflective system assures that monitoring is not exhaustive (Simon 1959; Vohs & Heatherton 2000). The first opportunity for error comes from the conversion process itself. Due to the limitations of intuitive reasoning, the perception of the world may be converted incorrectly.

As we have discussed previously, the reflective monitoring of the intuitive system is not complete. The intuitive system makes many decisions in every moment, processing the vast flood of perceptual input into actionable items of action. The reflective system does not have the "bandwidth" to re-process the same input to verify correctness. In addition, recent research shows that willpower itself is a limited commodity. As we spend this reserve of willpower, out monitoring of errors becomes more lax, and ultimately more and more decisions go unchecked. For example, it is easy to decide at the beginning of the day to not eat fast food. While the decision is easy to impose in the morning, by the end of the day the continuous attention takes its tolls, and later in the day it is quite easy to choose to eat a candy bar or similar fast food.

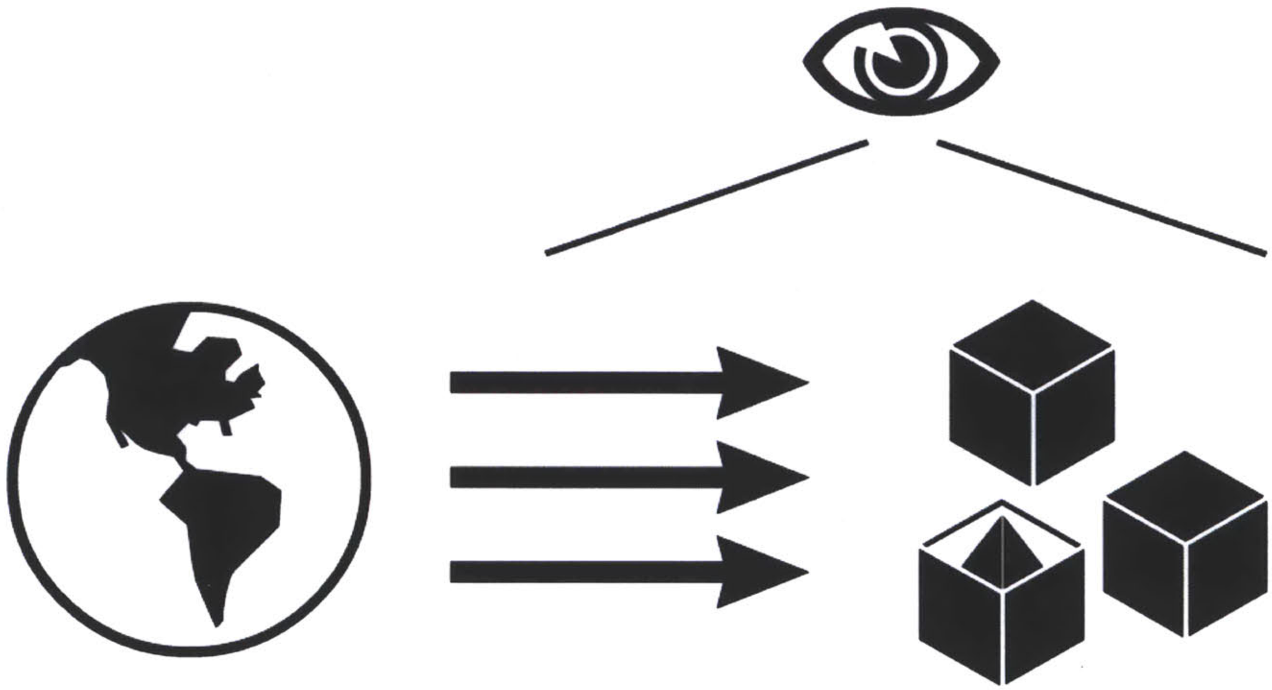

Figure 4. The perception-based intuitive judgment process, and reflective oversight.

In Table 3, we have summarized the major caveats affecting intuitive rationality (Kahneman & Frederick 2002; Mehta et al. 2004; Tversky & Kahneman 1981; Herrnstein & Prelec 1991; Simon 1959). By examining choice architectures through these lenses, it is possibly to identify one or more caveats affecting the reasoning. As an example, if we consider the case of balancing one's schedule, we can see that several of these caveats come into play. Time spent in meetings is misestimated due to duration neglect. Contrast enhancement may make the marginal gains from additional meetings seem better than they actually are. It may also be the case that incorrect valuation of gain and loss causes the person to feel that more time was "lost" than the value of the meeting, while in reality the value of the meeting may even have been sufficient to offset the loss, which may have itself been amplified by violation of dominance since the person feels the virtual lost value of the opportunity to better use the time.

Caveat	Reasoning Artifact
Scope neglect	Failure to scale decisions
Violation of dominance	Not basing choices of true overall utility
Framing effects	The statements surrounding a choice affect the result
Contrast enhancement	Things look either more different or similar than they are
Duration neglect	Time is not integrated
Incorrect valuation of loss/gain	Losses and gains are not computed equivalently
Balancing multiple choices	When multiple choices with interdependent values are mixed, it is difficult to identify balance

Table 3 Caveats of intuitive reasoning.

Relation to Conditioned Responses

An alternative approach to the information-centric approach we propose for mental prostheses is the conditioned response. The most commonly known work on behavioral conditioning was done by Pavlov, in the well-known study where dogs were conditioned to salivate (conditioned response) at the sound of a bell (conditioned stimulus), after they had been conditioned to associate the sound of the bell with salivation prior to feeing (unconditioned stimulus) (Pavlov 1927). However, recent studies have demonstrated that the connection between conditioned stimuli and unconditioned response is not as mechanical as supposed originally. Indeed, current thinking on this subject suggests that it is a form of environmental learning that takes place regarding the relationship between events (R. A. Rescorla 1988). This is quite similar to the habit formation that may be expected of mental prostheses, but we argue here that conditioned response approaches are inherently different from the perspective of the user.

We will begin with a short introduction to the various forms of conditioned responses. Conditioning exists in several forms. The fastest of these is forward conditioning: where the onset of conditioned stimulus (CS) precedes and signals the onset of unconditioned stimulus (US) (Chang et al. 2004). This can take several forms. There may be a delay between CS and US, though one of Pavlov's initial results was that this weakens the conditioning (Pavlov 1927). This delay may also be such that CS and US are not co-present, which is known as trace conditioning. Finally, CS and US may be of the exact same length (simultaneous) as well. In a reversal of this approach, inhibitory conditioning may be accomplished by providing the CS after the US has completed (Chang et al. 2004).

One of the most common models to explain classical conditioning is the Rescorla-Wagner model proposed in 1972 (R. Rescorla & Wagner 1972). This model mathematically maps the strength of the conditioning to the strength of the association to CS. As examples of CS to US are experienced, these records add to the strength of the association. The model thus predicts the extinction of the conditioned response as well, as the subject forgets some of these examples. One interesting aspect of the model is that its mathematical generalization suggests the inclusion of all stimuli present at the time of conditioning, not only the stimulus that is introduced for CS. As such, it accommodates the possibility that an extrinsic factor may block the creation of conditioned responses. However, recent work suggests that the model is too simple to account for all observed details. Alternative theories have been proposed, such as that by Gallistel & Gibbon (Gallistel & Gibbon 2000). This theory proposes that the learning due to conditioning is in fact a predictive rather than associative feature of cognition.

In addition to classical conditioning, where an unrelated stimulus is tied to a conditioned response, there also exists operant conditioning (also known as instrumental conditioning), which is used to strengthen or weaken an existing behavioral pattern, and differs from classical conditioning in that it controls voluntary behavior through the introduction of consequences (Skinner 1974). The core tools of operant conditioning are reinforcements (rewards) and punishments, and these may be applied either positively or negatively. That is to say, a negative stimulus may be applied after behavior to abate a negative behavior, or applied continuously until a positive behavior is elicited. The opposite applies for positive stimuli. Thorndike and later Skinner formulated a cognitively oriented model of operant condition that is based on the idea of consequences (Woodworth & Thorndike 1901; Skinner 1974). However, more recent theories show that operant conditioning may operate (at least in animals) simply by pairing of association of stimuli to end results, a process known as shaping (R. A. Gardner & B. T. Gardner 1988). As a result, it now appears that classical and operant conditioning are perhaps less apart than previously thought. Classical conditioning now takes a more cognitive approach (R. A. Rescorla 1988), while operant conditioning accommodates a more behavioral approach (R. A. Gardner & B. T. Gardner 1988). In the following discussion, we shall not separate these forms except as needed, since the end result continues to be effectively-involuntary change in behavior.

Conditioned response has been applied to human behavior in the field of psychology since the 1950s, and is known broadly and behavior therapy. Such classical conditioning techniques have been used to treat a number of conditions, including eating disorders, pain management, stress management, relationship therapy, depression, anxiety, and substance abuse (Morley et al. 1999). Conditioned responses are often helpful when the subject of treatment is uncooperative or unable to cooperate with the therapeutic goals (Cautilli et al. 2006). However, critics have noted that the underlying theory of cognitively motivated behavior is not strongly supported by experimental results (Beidel & Turner 1986). This critique is in line with the dual-rationality model of cognition espoused in this dissertation, though we do not directly deal with pathological behavior per se. Regardless of the critique of the approach, studies have shown that behavioral therapy is as effective as drugs in the control and abatement of negative behavior (Jacobson & Hollon 1996).

Within the context of this work, we must acknowledge that conditioned behavior is an effective approach to addressing the same set of problems. When considering only the feedback and response of the user, we can easily characterize a mental prosthesis as positive punishment form of operant conditioning, or negative feedback form of classical conditioning. However, there are several crucial differences when considering the system as a whole. First, as noted previously, current theories of conditioning suggest that it is not a particular stimulus, but the entire context of an experience that is used to form conditioning effects. For activities of the sort addressed by mental prostheses where the user is engaged on a regular basis, the creation of "clear contexts" for the inducement of conditioning is relatively difficult. Without such contexts, it has been shown that conditioned responses are formed exceedingly slowly or fail to form at all (Pavlov 1927; Woodworth & Thorndike 1901; R. Rescorla & Wagner 1972).

Second, conditioning requires a fairly strong adverse stimulus, which some consider a necessary requirement, in the sense that these theories posit that the creation of a fear response (or abatement thereof) is necessary. This is inherently opposed to the idea of calm feedback required of mental prostheses. However, it may be that in the pathological extreme cases where it is not possible to achieve user buy-in in addressing adverse behavior, behavioral conditioning is an excellent additional choice for achieving the same results.

However, a more subtle issue with the use of behavioral conditioning over mental prostheses lies in the class of behavior being targeted. While it can be said that we would all like to eat better, spent time more effectively, and have better control over our finances, few of us are willing to undergo the rigors of classical or operant conditioning to overcome these foibles. It is therefore necessary to "condition" the user in-situ in the contexts of their everyday lives. Therefore, it may be said that mental prostheses are in fact carrying out conditioning, but are in fact doing so automatically in contexts where such conditioning has not been considered previously, and using technology that brings more automation to the practice than has been considered as a framework for behavioral modification before.

Countering the position that mental prostheses are a new application of conditioning is the inclusion of the user in the process of behavioral change. In the upcoming sections we show that user buy-in has a significant effect as perceived by the users themselves. In addition, reflectOns as they have been designed and studies in the context of this dissertation have shown greater effectiveness and tools of information delivery than tools of decision transmission. In other words, they are perceived by the user to work better when the user makes the choice, rather than when the reflectOn makes the decision and tells the user what to do — perceivability is valued higher than actionability. Given the aforementioned description of conditioning, this approach does not align with the models of behavioral conditioning as understood by the psychology community. Therefore, while there may be aspects of conditioning in the feedback of the mental prostheses, the overall tableau is not compatible with pure conditioning.

That said, it is important to note that the underlying theories and feedback approaches provided in this dissertation, independent of the model by which we posit that these feedback should be applied, can be used for conditioning as well as mental prostheses. Indeed, as noted previously, it may be highly effective to apply the technological tenets to form a class of prosthetics that work based on evaluative feedback of behavior as conditioning does, rather than based on providing information to the user. Especially in pathological cases, or cases where the user finds themselves relapsing repeatedly, the ostensibly-stronger imprint left by conditioning approaches may be an excellent complement and extension of the work presented in this dissertation.

Tools and Technologies for In-Place Feedback

Interface techniques of relevance to this work can be divided into two broads groups: interface design approaches, and data gathering and presentation approaches. In this section, we shall provide a broad overview of each group. For our purposes, design approaches encompass the disciplines of ubiquitous computing, ambient interfaces, and augmented reality, while data gathering and presentation approaches encompasses sensor networks, agent-based systems, and persuasive interfaces.

Interface Design Approaches

Many of the interface design approaches we shall discuss in this section claim lineage stemming from the seminal work of Mark Weiser. At a time when personal computing was still nascent, Weiser and his colleagues as Xerox Parc envisioned computing as being invisible and omnipresent — woven into the fabric of everyday action and objects — a vision that is only coming to pass today after almost two decades (Weiser 1991). Weiser noted that the computer as it exists today is a focal point, demanding our continuous and unwavering attention. Instead, ubiquitous computing envisioned interfaces that reacted to everyday actions, and responded by manipulating the worlds around us subtle ways (Weiser & Brown 1996). Weiser and Brown summarized this changing relationship to computing in a table, an extension of which is reproduced below. The majority of changes to interfaces came from the realization the human attention is a limited resource, and cannot be used by many computers in the same way that a few computers now utilize it (Weiser 1999). As a result of this fundamental change, computers would have to become contextual to their use, and be designed around ever-more-specific usage scenarios while minimizing their "attention footprint" in order to attain transparency (Weiser & Brown 1996). The work presented in this document draws from this same concept of localized, context-relevant computing, and thus builds upon the insights provided by ubiquitous computing and its related fields.

The vision of ubiquitous computing has today become sufficiently established that the field has separated into several subfields. Pervasive computing focuses on the technical aspects, context-aware computing focuses on the use of context and environmental information to inform interaction and feedback, ambient interfaces and tangible interfaces focus on the difficulties of creating "invisible" interfaces, and ubiquitous computing itself has continued on to focus on design of vertically-integrated systems. We shall now shortly touch upon these various fields.

Period	Relationship between people and computers	Experience
Mainframe era	Many people : single machine	Technical
Personal computing era	Single person : single computer (or few)	Immersive
Internet era	Single people : connected computer	Social
Ubiquitous computing	Single person : many connected computers	Transparent

Table 4. The relation of computers to people.

Investigators of pervasive computing realized the technical challenges of the vision of ubiquitous computing, and realized that the computing model of the desktop and mainframe was not extensible to the encompassing interactions of ubiquitous computing. The first challenge therein is the distribution of computation to many small devices, necessitating communication and feedback that doesn't have to be "together" to act in concert, opening the doors to mobility (Saha & Mukherjee 2003). Pervasive computing also deals with many of the underlying technological and middleware issues that are a necessary requirement for creating ubiquitous, transparent computing systems: communication, fault tolerance, mobile sensing, energy-awareness, privacy, trustworthiness, and others (Satyanarayanan 2001). Pervasive interfaces are, as noted by Huang et al, "solutions in search of a problem" (Huang et al. 1999). While they are often discussed in the context of toy systems, they also inform the design of both hardware and software described in this thesis, and indeed many of the theoretical approaches that make mobile and distributed computing possible today.

Context aware computing, in contrast to the device and middleware focus of pervasive computing, builds upon that base by adding sensing, understanding, and awareness of the environment and use scenario of an interface — in short, its context. Context-aware computing introduced the idea of providing information to the user based on their actions at a given moment and environment (Dey 2001). Saliency is determined by the task at hand, while the task at hand is inferred from both the design placement of the system, and the sensors integral to the system itself. A large number of systems have been designed with these principles in mind, and provide a basis for understanding and evaluating the awareness a device has of its users as well as the means of modeling and responding to such information (Chen & Kotz 2000).

Effective	Use peripheral awareness correctly, and not occupy foreground attention
Efficiency	Gain and divest attention quickly
Safe	Not create ergonomic hazards by the placement of feedback — or alternatively, not disrupt existing rituals and practices
Utility	Inform the user of things that the user needs to know, and minimize accessory information content
Learnable	Mappings should be clear and easily learnable and memorable
Visible	Information is available to the user whenever she needs it without additional effort
Provide adequate feedback	When the ambient interface has manipulable controls, these controls should provide clear feedback about the setting changes effected
Constrained states	The state transitions of the interface should be clear and understandable
Consistent	The inputs that evoke an output should be logically understandable
Participatory design	Ambient interfaces have many open questions, and require testing and feedback from users to be correctly "tuned" to a context

Table 5. Some guidelines for designing ambient interfaces. Derived from (Gross 2003).

At a higher level, tangible and ambient interfaces focus on the merging of the worlds of bits and atoms, providing insights into the product design aspects of ubiquitous interfaces (Ishii & Ullmer 1997). Tangible interfaces are defined by the tight coupling of physical environment and digital information. Specifically, tangible interfaces are defined by embodiment and metaphorical linkage of interface and digital actions (Fishkin 2004). Since the interface is embodied, it fulfills the requirement of transparency. Additionally, cognitively the metaphorical relationship and physically-manipulable designs thus engendered allow for intuitive rather than reflective acceptance of the interface, reducing cognitive load and allowing faster parallel access. This feature fulfills the promise of "calm technology" originally proposed by Weiser by removing the interactions from the realm of the single-track and resource-constrained reflective processing system of the mind.

While tangible systems focus on collocated input and output, ambient interfaces have a greater focus on invisibility and legibility rather than interactivity. This extends to the special case where interaction is simply the observation of preexisting actions by the system, while feedback is still a matter of providing information. Ambient interfaces occupy or overlay the environment while providing salient, learnable cues about the state of the digital world in the physical world (Wisneski et al. 1998). In the work presented in this thesis, aspects of ambient feedback are used heavily, and the existing research provides not only inspiration for the design, but also guidelines for the effective generation and application of feedback (Gross 2003). This guidance is summarized in the table below.

In dealing with the multitude of calm yet watchful devices surrounding individuals in the ubiquitous computing model of interaction, input and output methods must be rethought. While visual feedback, as provided by the ambientRoom (Ishii & Ullmer 1997) or the Ambient Orb (Ambient Devices 2011a) can be used, other means have also been suggested. While the human visual system is capable of absorbing a vast amount of information in a physical sense, the processing of this information "compresses" the information based on salience and attention, rendering things within view either invisible or distracting (Navon 1977). A commonly used alternative is audio. The auditory mechanism, much like the visual system, is capable of distinguishing and concurrently processing a number of disparate stimuli, and provides a second means of output from an ambient system, while allowing the visual and physical attention to remain on a foreground task (Feldman et al. 2005). Spatialization and non-speech feedback have likewise been shown widely to be effective means of communicating information to the user (Lumsden & Brewster 2003).

A number of approaches have also been considered for input. In addition to sensing the environment directly to infer context or using explicit commands via buttons or other input devices, the most common form of input to ambient interfaces remains gesture. Gesture using hand symbols and head movement (Lumsden & Brewster 2003), as well as general activity sensing (Tapia et al. 2004) have all been considered in attempting to create mobile input methodologies, and are applicable to ambient interfaces. Lastly, there has also been investigation of approaches that change from ambient to interactive based on the user's pose relative to the system (Vogel & Balakrishnan 2004). Such approaches lend themselves well to systems that are initially or occasionally interactive, but generally ambient in nature.

The gesture sensing techniques used in this thesis are drawn from wide-ranging research in sensor networks. A number of platforms have been created by these researchers to investigate sensor properties and analytic frameworks. The systems have been variously used for motion analysis (A. Y. Benbasat & J. A. Paradiso 2005), pervasive health monitoring (Lo et al. 2005), wearable gesture recognition (Park & Chou 2006), and ad-hoc and general use-cases (O'Flynn et al. 2005). In addition to providing an understanding of the needs of mobile sensing system, additional literature also provides insights into pattern recognition techniques specific to mobile sensing (Hagras et al. 2004), as well as approaches that are well-suited to capability-constrained condition (A. Benbasat & J. Paradiso 2007). While a small subset of the literature is referenced herein, work done in the field of sensor networks directly informs the design of the electronic systems used in this thesis.

Persuasive Technology

The work presented in this thesis is a derivative of the idea of persuasive technology as introduced by B.J. Fogg, in that it attempts to formalize the concept of just-in-time feedback based on input from automated systems as well as the placement of feedback at the pace, time, and design point where a choice is made, while focusing on individual goals and informational feedback rather than conditioned responses. Fogg defines persuasive technology as technology that is designed to change a person's behavior or attitudes (Fogg 2002). Under this broad definition, everything from dialog boxes to agent-based systems can be considered to be persuasive (or at least to have persuasive effects). More directly, persuasion using technology can be broken down into persuasion by argumentation, persuasion by information, and persuasion by exposure to a state or environment (Atkinson 2006). The last of these has been the domain of marketing technologies (on-shelf coupon dispensers and checkout lane TV channels, for example) for some time now. The second is the aim for in this work, while the first is the form of persuasion seen in more rhetorically-oriented approaches, such as speech and manifestos. An additional form of persuasion can also be mediated by social mimicry (Fogg 2002), though it is not central to this work.

Persuasive technologies have been proposed and used for a number of applications. Chief among them are environmental awareness and health. Work by Intille et al. has suggested the use of this approach towards providing just-in-time feedback can improve health and compliance of healthy behavior (Intille 2004). However, the applications stemming from this initial foray have focused primarily on health care rather than systemic wellness. A number of researchers including Intille have also examined the use of persuasive approaches to moderating energy usage and building environmental awareness, whose work has been summarized by Wood et al (2003). These approaches have varied from simple cost and kilowatt usage — an approach that has since filtered to the commercial realm with devices such as the Energy Joule (Ambient Devices 2011b) — to more evaluative feedback that simply tells the user when the costs of kilowatt-hours used during an period are higher than some threshold. At the level of the entire household, these devices have shown some efficacy in moderating energy use.

One example of a persuasive interface is the perFrame prototype (Obermair et al. 2008), designed to improve the sitting habits of its users. It is a photo frame that shows a different photo or action by a pre-filmed actor based on whether the sitting habits of the user have been positive or negative. The form and the feedback sequences are shown is the figure below. This approach uses the social and empathic feedback provided by what is perceptually the judgment of one's behavior as a means of promoting behavioral change. Unlike the focus on reflectOns, the perFrames system quite unabashedly uses an evaluative approach, where the correctness of the decision is exposed. This however exposes a limitation of the ambient approach, in that the "bandwidth" of communication is quite small, thus limiting the authors on only a few axes of feedback. As a design choice, in this case the authors have chosen to only expose the decision, and only after choice have been made (a posteriori feedback). Lastly, the perFrames prototype serves as a sample of the social persuasion approach outlined by Fogg.

Figure 5. The perFrames prototype, usage scenario, and feedback frames (from positive +2, though neutral) to negative -2). Images from Obermair et al. (2008).

Another example is the TripleBeat system (de Oliveira & Oliver 2008), which uses the combination of a glanceable display and control of music playback in a mobile device to help the user reach exercise goals. The system tracks the workout pattern of the user in real time, selecting songs with higher or lower tempos based on whether the user should be going faster or slower during a segment of the workout plan. It also sets up a competition among users which encourages users to work towards higher goals. This system uses a very subtle form of feedback, acknowledging that during exercise periods users are unlikely to use more complex interfaces. The real time feedback also allows users to take in feedback directly. By combining the social aspect with the competition, it also sets up a means of comparison. However, the prototype does not attempt to generalize these strategies to further applications, and indeed heavy use of the fact that there already exists a common measurement mentality among people who exercise regularly. Nonetheless, a study of efficacy showed that users did like this sort of interface, which has since been commercialized in other form factors (such as the Nike+ (Nike 2011)). An example of the glanceable display for this prototype can be seen in the figure below.

In addition, there has been a significant trend in using persuasive interfaces in conservation. Midden et al. suggested the some of the first instances as early as 1983, putting in context the role of education, information, and feedback in monetary rewards as well as social/comparative forms (Midden Joanne et al. 1983). The work notes that in field trial, the primary driver was the availability of individual feedback, with even additional monetary incentives providing a smaller incremental benefit. A 2001 review of the field by Darby support this conclusion, suggesting that high-frequency, individual updates are by far the more effective across studies and approaches (Darby 2001).

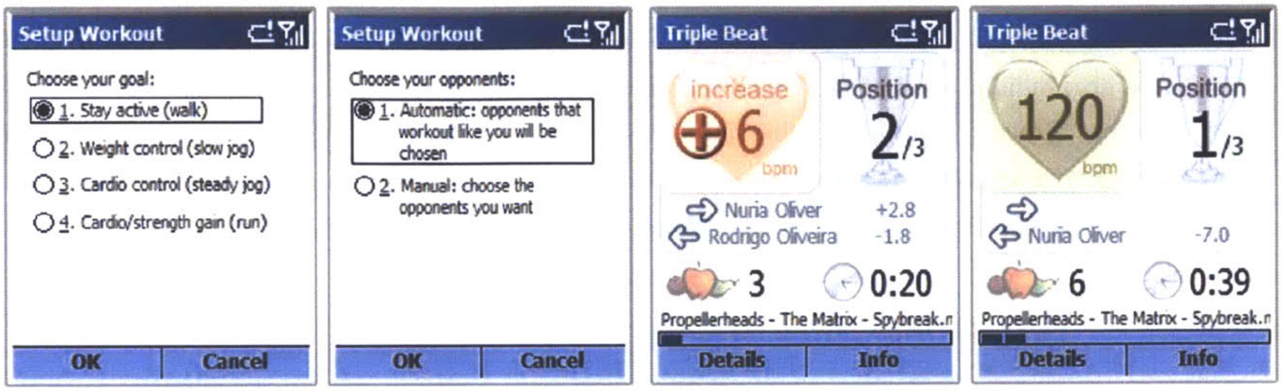

Figure 6. Setup and glanceable interface of TripleBeat system (de Oliveira & Oliver 2008).

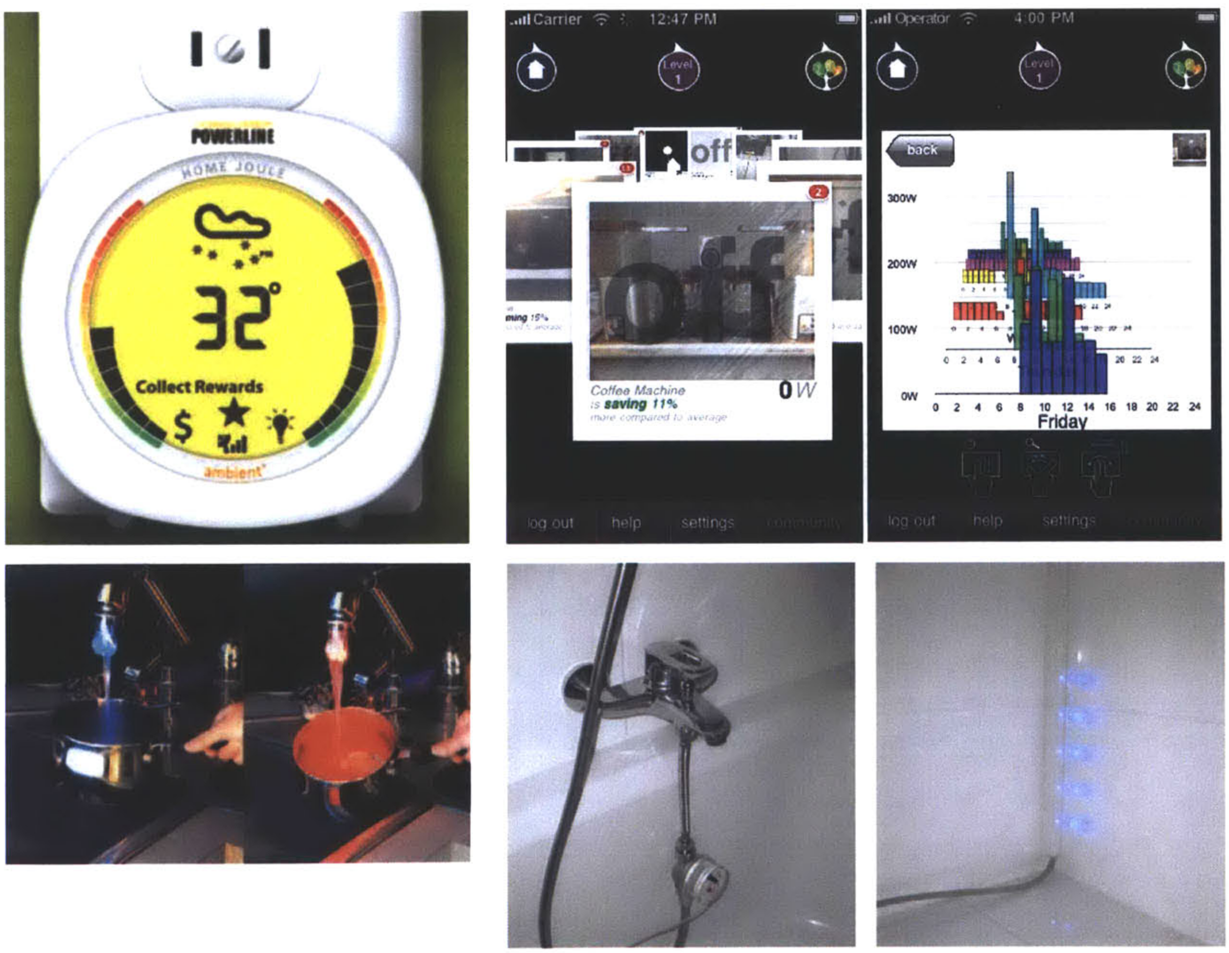

Figure 7. Various in-place energy/resource usage feedback methods. Clockwise from top left, EnergyJoule provides feedback about a single outlet (Ambient Devices 2011b); EnergyLife monitors and visualizes the entire home (Björkskog et al. 2010); HeatSink visualizes temperature at the faucet (Bonanni et al. 2005); and "show-me" shows the water usage behavior as a visual bar graph in the shower (Kappel & Grechenig 2009).

Many of these studies have suggested feedback approaches to provide such individual feedback. At the single-device end of the continuum, devices such as the Energy Joule (Ambient Devices 2011b) provide in-place feedback, while the EnergyLife (Björkskog et al. 2010) and WattBot (Petersen et al. 2009) approach of whole-house monitoring via mobile device and subtle ambient feedback attempt to approach the problem holistically. See figure below for details. Lastly, in addition to directly providing feedback about energy, similar approaches have been taken with other resource usage scenario. Selker et al. have done a number of projects related to energy and water use at the sink, while the "show-me" prototype has attacked the same problem at the shower. These projects are also shown in the figure, and as again take an individual, real-time approach to providing feedback as a means of persuasion. In all, these approaches underscore need for in-place, individual and actionable feedback in persuasive systems.

Summary

In concluding this chapter, we have visited three separate disciplines: cognitive psychology and behavioral economics, ambient and ubiquitous interfaces, and persuasive interfaces. The first area provides the foundation on which this thesis stands. This thesis builds on existing rigorously tested theories that provide a framework for perturbing cognition. We drew inspiration for sensing of user activity and the formatting of output from ambient and ubiquitous interfaces. Finally, we shortly touched on the state of the art in using technology to change behavior, and drew together the main examples in the field to provide a sense of the work that has been done.

Mental Prostheses

This dissertation focuses on the change and modulation of human behavior using automated feedback. In its general form, this type of system is termed a *mental prosthetic*: a distributed, design-embedded, just-in-time interfaces that augment the human judgment process. In other words, mental prosthetics modify or modulate behavior by augmenting the human judgment process by providing individual, in-place, just-in-time feedback that is embedded into the flow of the user's actions. This class of interfaces encompasses a variety of embedded systems that are designed to become a part of the normal usage of devices in the midst of everyday activity as diverse as eating, using the microware or the hair dryer, or even simply surfing the web.

The concept of mental prostheses is based on the hypothesis that it is possible to use just-in-time feedback drawn from machine-observable behavior to help align behavior with personal goals. Using just-in-time feedback, we hope to provide feedback that is correctly timed to affect decisions, situated to be within the user's frame or reasoning, and within the flow of the user's activities. Mental prostheses use machine-observable feedback to inherently limit the amount of direct input necessary. These devices are meant to reduce mental load on the reflective system and operate without direct manipulation, and provide feedback without polling the user constantly for input. As we discussed, reflective attention is a limited resource, and thus this approach requires us to think tangentially about problems in order to design reflectOns. However, it also allows us to avoid the many limitations of human perception and rationality that naturally make their way into input from the user.

Mental prostheses differentiate themselves from conditioning approaches by cueing the user not with "corrective feedback," but "corrective information" that helps the user self-correct their decisions to align with their goals. This approach to goal alignment attempts to balance feedback with personal intent and willpower, forming the basis for a greater sense of control over the experience of using mental prostheses. On the same note, mental prostheses differ from persuasive interfaces in that there is a much higher focus on personal goals. By focusing on intrinsic motivations with the user's permission, we hope to elicit a more engaged and personal response to these devices. As a result, the cultural and unconscious manipulations that are often used in persuasive interfaces are outside the purview of mental prostheses.

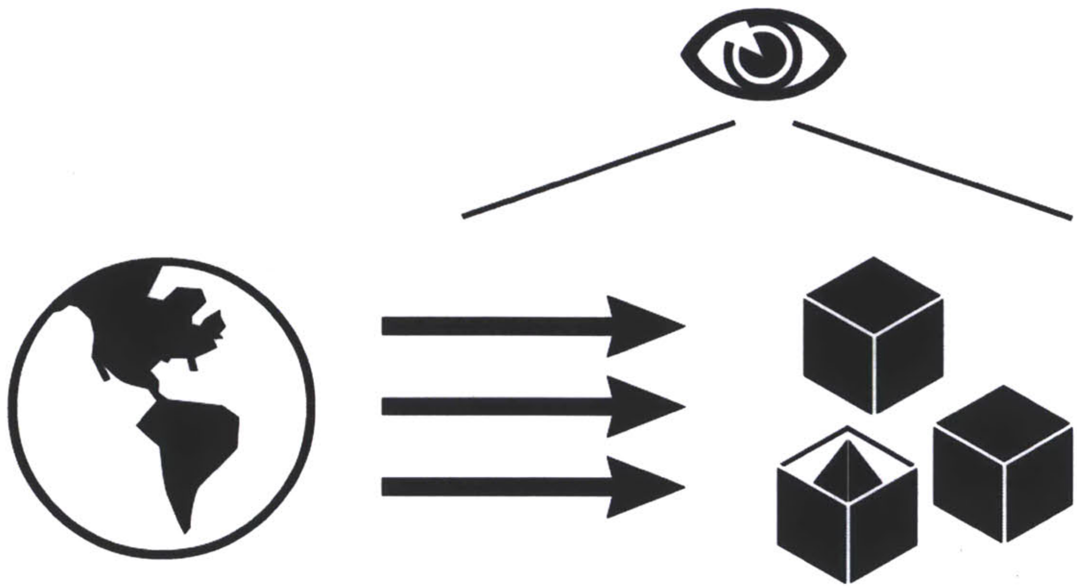

Figure 8. The perception-based intuitive judgment process, and reflective oversight.

Drawing on our previous discussion of behavioral theory, Figure 8 reiterates the process by which we convert perception to decision. We perceive the world by out sense, and the intuitive reasoning system converts these perceptions to a multitude of choices and decisions in parallel. The reflective system monitors this process to the extent possible, though the monitoring is not exhaustive (Simon 1959; Vohs & Heatherton 2000). Opportunity for error comes from the conversion process itself, as well as limitations of the oversight process. In this context, mental prostheses introduce cues and corrections into the environment in time to correct the biases or limitations that may be causing the intuitive system to make incorrect decisions (Figure 9). This may be done by introducing perceptual cues, which may be integrated into the flow of the decision, or perhaps separated but presented in a way that allows for inclusion in the choice. Mental prosthesis may also open up the opportunity to form new habits by repeated exposure to the corrected choice, or to learn the correct decision experimentally so that the user learns a personalized approach to avoiding the aberrance. A mental prosthesis may also tap into the fact that intuitive reasoning takes both past and present information into account, and remind the user of important cues (cue memory) that then serves to modulate the end behavior. Lastly, a mental prosthesis may operate on the reflective monitoring of the intuitive system. Systems that affect behavior via this approach are termed reflectOns.

ReflectOns: Mental Prostheses for Self-Reflection

ReflectOns are a class of mental prostheses that achieve the aforementioned results by using feedback to elicit self-reflection. Drawing information from the activity of the user around them, reflectOns analyze behavioral patterns in a way orthogonal to human cognition, focusing on integration, unit conversion, and memory — tasks that a computer is sublimely suited to, but where our own cognitive processes can be led astray on. Unlike many of today's persuasive interfaces, these mental prostheses attempt to focus the user on making rational decisions at specific infection points in choice architectures where a rational judgment is most effective, thus reducing the mental load and points of failure induced by attempting to logically reason every decision. Lastly, these devices provide unobtrusive and acceptable, yet easily understandable, feedback using embedded means to draw the user's attention at the correct time to the information available to them.

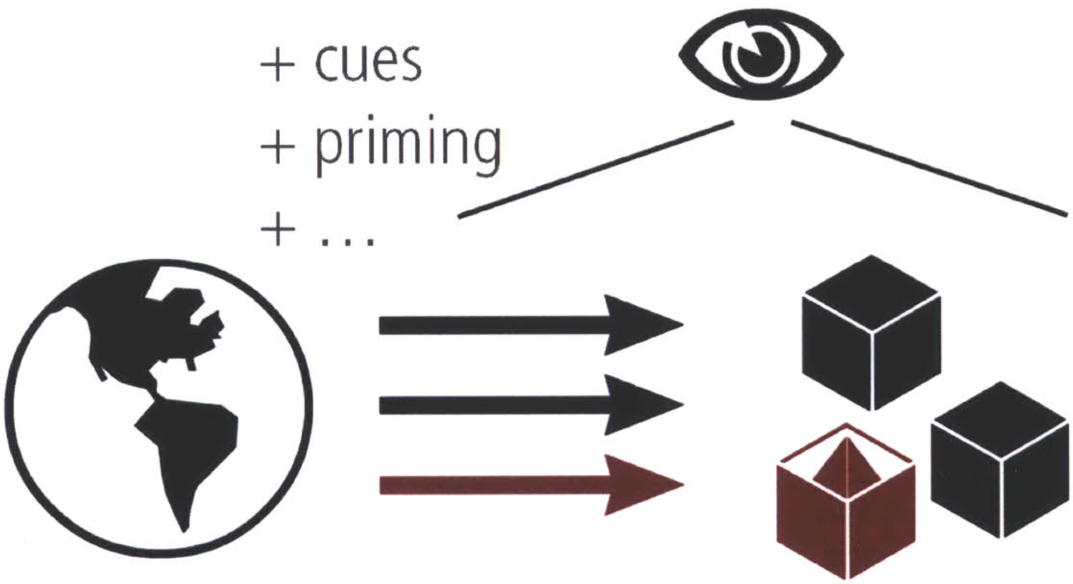

Figure 9. Mental prostheses: means of action.

As part of this work, this thesis provides several examples of design in various states of completion, from completely functional devices such as the utensil attachment that can help people eat better, to the conceptual design of a device that can give the user better awareness of their social circle and interactions. These prototypes and design discussions are meant to provide a sampling of the various approaches to data collection, synthesis, and feedback that can be taken to help users make better decisions. The studies performed in the course of this thesis show that the design approaches taken are capable of eliciting tangible behavioral change. They also underscore the importance of various design elements in creating a successful intervention.

Target Behavior for ReflectOns

As we have discussed previously, the reflective monitoring of the intuitive system is not exhaustive. The reflective system does not have the "bandwidth" to re-process the same input to verify correctness (Simon 1959). In addition, recent research shows that willpower itself is a limited commodity. As we spend this reserve of willpower, out monitoring of errors becomes more lax, and ultimately more and more decisions go unchecked (Vohs & Heatherton 2000). The goal of reflectOns is to provide just-in-time feedback that re-activates the monitoring system at times when monitoring is salient (Figure 10). This not only reserves the reservoir of attention for when it is needed, but also reduces overall load on the user in using the system.

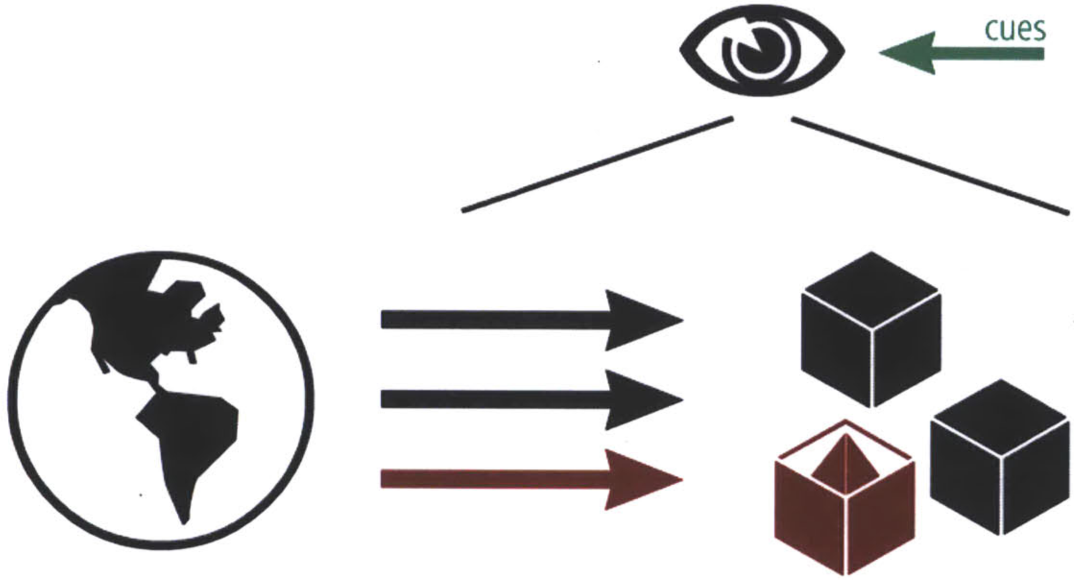

Figure 10. ReflectOns: means of action.

In Table 6, we have again summarized the major caveats affecting intuitive rationality, and suggested possible solutions a reflectOn may apply in order to overcome that limitation. Let us once again consider the earlier example of balancing one's schedule. As mentioned, we can see that several caveats come into play: duration neglect, contrast enhancement, incorrect valuation of gain and loss, and violation of dominance to name the previous examples. This example illustrates various approaches that a designer may take to ameliorate the situation, perhaps by showing total time spent in other tasks to reframe time used in meetings, or even providing means to capturing and displaying the value of the meeting to counter the apparent loss of time, which ultimately improves the utility the user feels. In summary, the work summarized herein provides a basis for exploring and identifying possible amelioration strategies and design possibilities for future reflectOns.

In addition, we can also draw upon some of the tenets of designing ambient interfaces (Table 7), and consider them in the context of reflectOns to provide some examples of how these principles may be applied. It should be remembered that ambient interfaces, while not focused on behavioral changes, have the same basic requirements as reflectOns in order to peripherally introduce information into a user's perceptual frame. Therefore, many of the requirements that have been articulated for ambient interfaces (Gross 2003) can also be translated to reflectOns, as we have done in Table 7. By tying the design principles of reflectOns to the superset of principles that drive design of ambient interfaces, we can now rely on the bevy of research in ambient interfaces to further inform the design of the output provided by reflectOns.

Caveat	Reasoning Artifact	Possible Amelioration
Scope neglect	Failure to scale decisions	Sense and show scaled effects
Violation of dominance	Not basing choices of true overall utility	Enhance loss or gain based on goals
Framing effects	The statements surrounding a choice affect the result	Provide compensatory framing
Contrast enhancement	Things look either more different or similar than they are	Make decision variable more accessible to highlight errors
Duration neglect	Time is not integrated	Integrate for the user
Incorrect valuation of loss/gain	Losses and gains are not computed equivalently	Project into alternative framing
Balancing multiple choices	When multiple choices with interdependent values are mixed, it is difficult to identify balance	Display integrated effects of each choice

Table 6 Caveats of intuitive reasoning and sample ameliorations offered by a reflectOn.

Design Principle	Description	Use in ReflectOns
Effective	Do not occupy foreground attention	Provide feedback cues that merge with the perceptual frame of an activity
Efficiency	Gain and divest attention quickly	Cues must be easily legible and peripherally noticeable
Safe	Must not disrupt existing rituals and practices	Cues must not provoke additional cognitive aberrations
Utility	Inform the user of things that the user needs to know	The cue should be framed to only apply to the reasoning task
Learnable	Mappings should be clear and easily learnable and memorable	Cues should not need excessive reasoning to process
Visible	Information is available to the user whenever she needs it without additional effort	Cures are placed in perceptual frame of activity
Constrained states	The state transitions of the interface should be clear and understandable	The feedback provided should have a clear scaling to avoid triggering additional cognitive issues
Consistent	The inputs that evoke an output should be logically understandable	The cue should make sense to the user, and not be overly processed to the point that user loses trust in it
Participatory design	Ambient interfaces require testing and feedback from users to be correctly "tuned" to a context	Cues should be tuned to the design context, and we must have the user's approval to intercede

Table 7 Ambient interfaces as used in reflectOns. Derived from (Gross 2003).

What Makes a ReflectOn?

The application of the lessons learned in the background section, combined with the guidelines provided below, provide us with a basis for producing new reflectOns that correct intuitive reasoning aberrations. In addition to the mappings provided in Table 6 and Table 7, reflectOns share a few core design principles that define them as a class. ReflectOns subscribe to 5 general principles: they are a situated, personal, acceptable, just-in-time, and reasoning-based. In this section, we discuss these principles, and their respective importance.

Situated ReflectOns are by definition situated: they exist for a particular interaction, in a particular place and time within an interaction the user engages in, and able to sense and provide feedback into the environment such that their feedback becomes part of the perception of the flow of activity for the user. Situatedness is important for two reasons. On one hand, it allows the device to be located as close to the point where the user is engaging in the activity of interest as possible, thus giving the most unbiased information to the system for computing feedback. Additionally, situated sensing overcomes the problem of self-reporting bias in the user, as well as unintentional loss of data. This also leads to a non-disruptive collection of data – an important factor in making the system acceptable to the user.

On the other hand, situatedness allows the reflectOn to provide feedback in the context of the activity. It should be noted that when we are engaging in irrational behavior, the behavior is often driven by intuitive reasoning that is heavily tied to perception. By moving the feedback into the context of the activity, we can attempt to change the perceptual valence of an action, and thus either promote or dissuade the particular activity.

Personal ReflectOns focus exclusively on the personal realm, whether to select goals, preferences, or behavior. The mass change in behavior is outside the purview of reflectOns. Instead the focus is on understanding the behavior of a single user in the context of an activity, and providing feedback based on the user's own actions, and in alignment of the user's own goals. This personalization is an important aspect of behavioral change based on the cueing of rational oversight of activity, with the goal of reducing cognitive overhead. Transforming feedback from general to specific scenarios only adds to the cognitive load, while overgeneralization also reduces the trust the user can place in the system.

Acceptable The design of reflectOns as situated personal devices requires that the devices fit the social, cultural, and personal norms of the user. The designer thus must be sensitive to the presentation of the feedback, especially when social norms may be connected to the monitored behavior. In such cases, the user is in effect entrusting the system with privileged information, and if the system broadcasts the information collected along with the feedback, it undermines the trust the user has in the system and makes it less likely that the user will continue to use the system. As shown by the studies, acceptability modulates helpfulness and overall performance of the feedback in the test scenarios.

The feedback must also be unobtrusive enough to the user so that it does not feel like a "guilt trip" or a distraction to the everyday flow of the activity. If the feedback is too judgmental, the user may associate negative affective valence with a positive action, thus countermanding the goal of the system. Unobtrusive feedback also prevents the reflectOn from causing fatigue to the user, and optimally allows for better cueing.

Just-in-Time ReflectOns should be designed to provide the contextually correct data in the contextually relevant location. The goal of these systems is to "wake up the reasoning-based oversight" of intuitive decision-making, and that is best done when feedback is provided judiciously to avoid overwhelming the reserve of high-awareness time available within a use context. By providing just-in-time feedback, we can conserve this reserve as much as possible, which in turn maximizes the chance that feedback will be effective.

Reasoning-based ReflectOns, as suggested by the name, operate by cueing the user to reflect on their actions at the right place an time. In the following sections we will discuss a number of possible approaches that do not necessarily use this mechanism. However, this is the particular mechanism that has been verified by the work done in this thesis, and provides one of the most straightforward (to the user) means of producing behavioral change. By focusing on rational behavior, this approach feels more understandable to the user, and has less of a flavor of trickery, which can undermine the effectiveness of an intervention.

These defining characteristics also provide a general outline for the design of new reflectOns in that they generalize upon the aspects that helped the prototype reflectOns successfully affect the perception and behavior of users. While each reflectOn must be tuned to the context it will be used in, we feel that these five features provide the basis for a successful design in the majority of cases.

What Are ReflectOns Good For?

While reflectOns can have a wide variety of applications, it is also necessary to keep in mind that it is not panacea. The design guidelines proposed above also place limits on the applications where reflectOns may excel. In this section, we touch on aspects of choice scenarios that may affect a successful design.

Repeated Choice ReflectOns are ill-suited to single-choice scenarios, but can be effective when a series of choices – either as a cluster, or connected by some measure of similarity – are targeted. That is to say, it is ill-suited to prevent a person from consuming dessert on a particular day, but may be effective in modulating their eating behavior as a whole by reminding them not to have dessert in general. While reflectOns do provide personalized feedback in a just-in-time manner, the goal of the feedback is to change habitual action rather than to act as traffic lights for behavior. However, effects such as melioration suggest the need for caution. While maximization or minimizations are more easily approachable, the contrast-enhancing nature of perception makes it difficult to train for balance. In the case that a multivariate choice must be balanced, it may be necessary to take into account the difficulty of gaining and retaining a balanced state.

Change over Time ReflectOns produce a learning effect by producing a corrective "force" against the vector of habituated behavior. However, habitual actions and heuristics take time to form and change, and are based on the summation of many pieces of feedback, as demonstrated in the studies.

Committed User ReflectOns are meant to be personal. That means the feedback is meant for an individual, but also that the individual is accepting of the feedback. If the user is inimical to the feedback, then quite logically effects are lessened relative to approaches that "trick" the user into doing the right thing.

Extrinsic Factors One may imagine a ReflectOn design and a force that tries to change the trajectory of a moving object (Figure 11). However, that does not mean that additional (and stronger) forces are not also present. One example of such a force may be social pressure. For example, in Mauritius it is expected that women be (by medical standards) overweight to be considered attractive. A reflectOn that encourages its user to eat less would be of limited effect there, since there is a strong social pressure to eat more.

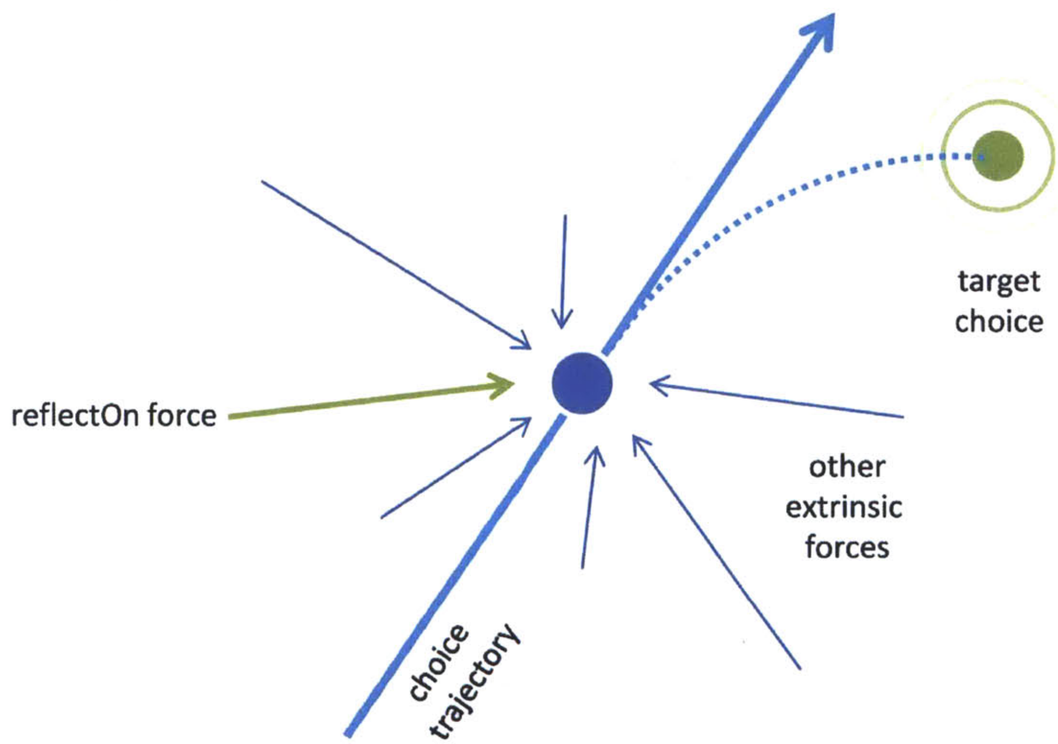

Figure 11. Force representation of the activity of reflectOns.

The example is also an example of lack of situatedness. The design of a reflectOn must be situated not only in the user's flow of activity, but also in the social and cultural norms of the realm of deployment. It is possible that a reflectOn that acts on eating habits operates very differently in Mauritius than it does in the United States.

The same also applies when there is a strong immediate feedback due to an action, or the action is heavily pre-committed. In either case, the "velocity" of the action along the choice trajectory is too high to overcome for that particular choice. However, even in these cases it may be possible to design a reflectOn that modifies the starting trajectory over time by changing the perceived end value of the action. In such cases, the designer and the user need to understand and set proper expectations for what is possible.

Feedback Placement In many cases where a reflectOn may be used, the reason that the intuitive choice is incorrect in the first place is that the results of the action are temporally or physically removed. Thus it is easier to overeat than not (tastes good now, weight gain later), and easier to spend on credit (get goods now, pay later). If it is not technically, socially, or attentionally possible to place feedback within the perceptual frame of the behavior, the reflectOn is unlikely to be effective in changing the behavior. Indeed, as the studies indicated, the more immediate the placement, the more helpful the intervention. It should also be noted that unobtrusiveness is also a necessary factor in this case to prevent the reflectOn from overloading the user. Again, if a very intense feedback is necessary to reach the user, the feedback is likely ill-placed or ill-suited, and the reflectOn itself may need to be reconsidered.

Clarity The intuitive reasoning system operates on easily accessible feedback. If the feedback is highly encoded and requires significant mental resources to process, then it is the equivalent of requesting resources from the reasoning system to begin with. Since many reflectOns will likely provide some form of continuous feedback, this is another way where the intensity of the feedback provided by the reflectOn becomes artificially amplified, and thus inappropriate for the user.

Garbage In, Garbage Out There are conditions where what the reflectOn really needs to sense is the internal state of the user. With the current state of technology, such inference is generally fraught. ReflectOns are meant to be extensions of one's decision process, and thus require a high degree of trust in the system. Therefore, overly-statistical approaches with insufficient (for the situation) accuracy are ill-suited to reflectOn designs. In such situations, it is better to approach the problem tangentially with information that is available and discernible. As with immediacy and clarity of feedback, the immediacy and clarity of sensing are important drivers in providing feedback in a timeframe that allows the user to tie the action to the projected consequence. Like the credit card bill, feedback that is not tied to individual actions is unlikely to be helpful in changing behavior.

This is also true of user goals and preferences. It is generally better to create single-purpose systems that have a fixed goal to which the user can align (or not). However, if the system does have multiple modes, the designer must ensure that the mode is strongly aligned to user goals to create the correct level of trust that allows long-term use. If the user's goals and preferences need to be gathered, these should also be considered for feasibility. If the user doesn't know how to express their goals and preferences, or if the process is too cumbersome before the feedback becomes valid, as again the design should be considered for revision.

Lastly, it is the duty of the designer to communicate the limits of the system to the user. While certainly it is easier to "promise the moon," such expectations inevitably lead to destruction of the trust in the system that is essential to the operation of mental prostheses of every kind. It should be acknowledged that human behavior is complex and that mental prostheses – much like their physical prostheses counterparts – are assistive rather than replacement technologies with inherent limitations. However, within their operational bounds, such technologies still hold the promise of improved quality of life.

Further Approaches for Mental Prostheses

While thus far we have focused primarily on reflectOns, which in turn focus on change of behavior by cued reasoning, research done in the course of this thesis also opens up the possibility of manipulating intuitive reasoning directly. While such approaches are less generalizable, we would like to note the possible approaches here. While one may characterize reasoning as situation and data-oriented, intuition is best characterized as perception-triggered (Kahneman 2003). Therefore, this alternative approach can be loosely termed "intuition tuning."

When speaking of intuition or intuitive decision making, it is important to keep in mind that intuitive reasoning has a specific meaning in the context of this discussion. The intuitive reasoning system is a portion of our reasoning framework that creates impressions of affective and decision valence based on high-speed, prototype-based processing of input stimuli. In other words, it converts perceptual stimuli to conceptual (ie, prototypical) representations based on past experience. Thus intuition does not spring from nothing, and nor does intuition become triggered in isolation. As discussed in the background section, many of the behavioral aberrations we discuss can be seen as aberrations of the intuitive reasoning system. Thus, if one wishes to perturb this system, it is crucial to understand why the results are out of tune with reality, and additionally, why the monitoring system provided by rational reasoning does not trigger in such cases.

If we wish to tune the intuitive reasoning system, we must either change the prototype or prototype selections of the intuitive system, or trigger the monitoring provided by the reasoning system as done by reflectOns. Since reasoning appears to be a limited resource (Vohs & Heatherton 2000), it is necessary to trigger this system judiciously. The alternative approach is to target intuition directly by changing perception. Below we shortly outline some possible avenues for this approach which may be approached in future work.

Remapping The intuitive reasoning system remaps perceptual information into conceptual frames. In other words, it extracts the meaning of perception. This remapping is done with the help of heuristics to reduce the processing required. Naturally information is lost in the process, and generally it is the differential information that is lost first. Given a good understanding of the heuristics (and inputs) for a given situation, it is possible to proactively show information that modifies the heuristic outputs. This may be done by highlighting or attenuating information, or possibly by showing additional information.

However, this approach requires the sensing of action to move very close to the initiation of an action, or be permanently present, and may be well suited to approaches such as augmented reality.

Prototype Heuristics If remapping is the process of converting perceptions to meaning, then prototype heuristics (or stereotypes) are the output of the remapping process. They are pre-packaged sets of affect, information, concepts and reactions that are used as shortcuts in the decision process. It is possible to conceptualize the creation of such stereotypes. In fact, marketing campaigns have successfully done so repeatedly. For example, Toyota created the Scion brand and marketed it to break free of its stolid stereotype, and did so quite successfully (Microsoft MSN n.d.). It may be possible to manipulate information to create such stereotypes on a personal basis. However, this approach is rather complex to apply predictably, since it is impossible what additional affects will be tied to the stereotype in unpredictable ways. Nonetheless, a strong affect tied to an action can be an effective tool in manipulating behavior.

Prescaling From the discussion in of existing work, it is clear that we do not carry out scaling and integration while engaged in intuitive thought. While reflectOns approach this problem by externalizing the processing, it is also possible to manage the issue perceptually by manipulating the framing of the value. If it is possible to model the perceptual nonlinearity of a value, it may be feasible to change the perception to account for the nonlinearity. This is of course less scalable to scenarios where the "perception" of the amount happens internally, but it is an important tool to keep in mind when designing visualizations for mental prostheses.

Tweaking Perception Perception as used in intuitive reasoning hinges on accessibility. If the information is not easily gleaned, it is less likely to be used in decisions. Therefore, if we can show that certain information is helpful or inimical to a decision flow, manipulating that information to change its accessibility is an effective way of "selecting" the heuristics by which a situation is judged.

Perception (whether physical or mental) inherently enhances the true contrast between values. This is a double-edged sword. It means that the designer does not have to specify whether some value is at an extreme. The value would be automatically rounded for intuitive processing. However, this also means that intermediate values are difficult to perceive, and may need to be remapped. Along the same vein, immediate results are generally higher-valued, while long-term results are discounted. The function mapping results to valence can be quite complicated, but once mapped provides a rather direct way of manipulating perceptual behavior.

Lastly, perception is framed. Unbounded values are also framed, though with arbitrarily high errors at the limits of the frame. The framing of values not only determines the valence of the value itself, but also the affective connotations. This affective valence can have a large impact on how the perceived value is weighted. As noted before, the framing thus may be manipulated to modify or even invert the overall valence of a decision.

The approaches listed above only touch on broad areas where perception and intuitive thought may be directly affected. This list is not meant to be exhaustive in details, and much future work remains in identifying, designing to, and testing the possible ways that the perceptual, intuitive, and reasoning systems can be used to craft behavioral change.

Applications

Before delving into the various applications we have designed in order to understand how the design goals we have set forth in the earlier chapters may be applied, we would like to begin by recapitulating some of the key design parameters that we will be referring to throughout this section. First and foremost, reflectOns are person-facing in their goal orientation, in that they observe the individual, and provide feedback in alignment with the goals of the individual. As shown in Figure 12, we do not approach the entirely of human behavior in this thesis, but rather the smaller but significant group on behaviors that can be considered "impulse-driven" or "unaware actions." With respect to these behaviors, we want to provide users with an understanding of the long- and short-term impacts of such cognitively-invisible micro-actions.

We proceed from the basis set forth in the background section that individuals have difficulty reasoning about the effects caused by the integration of micro-actions, and that it is generally difficult in such conditions for individuals to balance the long terms costs against short-term rewards. It is our goal to design reflectOns to allow users to understand how their actions affect real-world outcomes in valuations the user can understand. We hope via this intervention to allow users to make choices that better reflect the real-world utility of their choices, and to realize personal goals by being better armed in micro-choice scenarios with realistic valuations of the choices.

In the following sections, we will discuss a number of prototypes and design solutions that approach various choice classes with a palette of tools that allow us to design reflectOns that can be effective in a given choice environment and use case. The first of these is the use of in-place automated sensing of micro-choices. The difficulty lies in the fact that most of the sensors we have available today are ill-suited to the direct observation of human action. As a result, design of reflectOns involves using tangential sensing approaches combined with holistic analysis of use cases that may open up new opportunities for sensor deployment.

Once a micro-choice is sensed, it is necessary to design in-place analysis that can convert what is sensed into a form that can be used to provide feedback. While offline analysis can be useful in some circumstances, and can provide external observers with interesting insights into behavior, action and feedback must be closely coupled in the case of an individual user for the user to mentally connect actions with feedback and make informed decisions. However, we set the additional goal with reflectOns of not only providing a binary output, but also sufficient information to allow users to understand and evaluate a choice. As a result, the system designs must convert units such as watts, which are often difficult to visualize or compare in a real-world sense, into more understandable units where possible. Finally, the systems must provide this feedback to the user in a form that is easily legible and noticeable, but not so attention-grabbing as to be a detraction from everyday interactions.

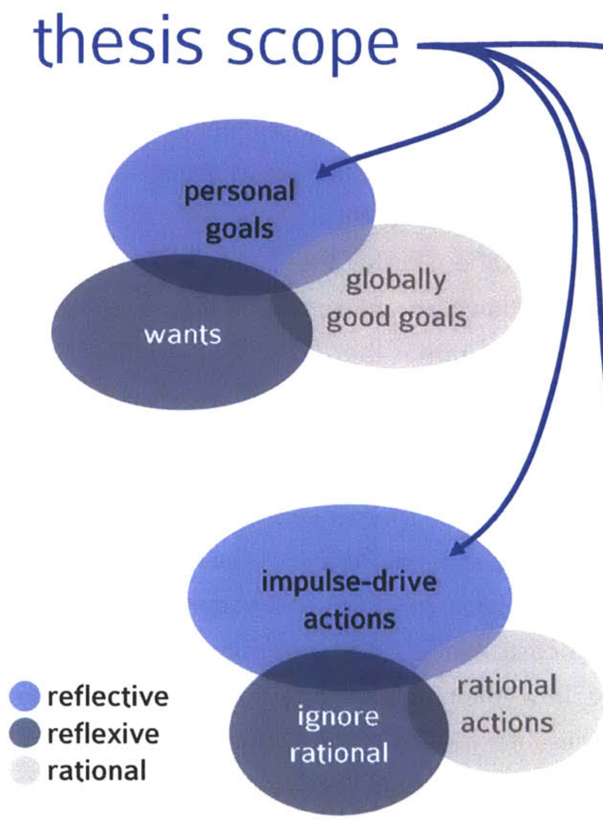

Figure 12. Scope of thesis relative to the behavioral artifacts

With these factors in mind, we would like to introduce the applications that will be used to illustrate various aspects of reflectOn design. The first three application scenarios we discuss below are systems for which we have developed physical prototypes. Following this, we dedicate a section to design solutions: outlines of system design which have not been implemented, but are in a state where they may be implemented from the outlines provided given resources and time.

Application 1: Web Surfing

Web surfing behavior is an excellent subject for the study of micro-decisions, in that it is central to many of our lifestyles, and yet has both positive and negative consequences on productivity. Much has been written about the benefits and sociological impact of the internet (Rheingold 2000; Turkle 1995). However, the drivers for web surfing are complex and multifactorial, and cannot necessarily be deconvolved into specific bins of usage without having an understanding of the intent driving the choice. For example, email is central to the workflows of most professionals, and thus may be fall under the umbrella of productivity. However, when the user avoids work by checking email repeatedly, the same action become inimical to productivity, and in fact there is a general consensus that limited or regimented access to email is optimal for productivity (Jackson et al. 2003). Some work has even been done investigating the addictive potential of web surfing, with Young setting forth a set of criteria to diagnose it as a medical issue (Young 1999). Due to the multifactorial nature of web surfing, even such extreme behavior can prove difficult to detect from the perspective of automated observation.

Approach

In order to gain insight into a multifactorial behavior such as web surfing, it is first necessary to consider three facets:

1. How the is web being used.
2. Possible amelioration strategies.
3. Ease of practice of various amelioration strategies.

To the first point, it is nearly impossible to understand a priori the intent of the user when she visits a site. While it is possible to know the site and the time spent, and possibly even the user's evaluation of the utility of or affinity to a site, it is not possible for a particular use event to know whether the intents were aligned with the user's goals at the time of the use without continuous polling, which out clearly be rather intrusive to the user's experience. Additionally, it is often the case that surfing a particular site is a "guilty pleasure" in that the user is completely aware that the time may not be appropriate for surfing the site, but may nonetheless draw pleasure (or even productivity) from the event. Therefore it is important to keep in mind the limitations of external observations in attempting to design a reflectOn to ameliorate web surfing behavior.

Once we know that a particular event may not be in alignment with the user's goals, it is necessary to consider possible coping or amelioration strategies that allow the user to regain control of their behavior. In this case, it is interesting to look at counter-aligned events as obsessive or addictive behavior in order to expose the possible approaches, as discussed by Young (Young 1999). These approaches can be broken down into 3 groups: goal setting, awareness, and supportive intervention. We will not examine supportive interventions such as family therapy in this document, since such interventions are better suited to more pathological cases of unwanted behavior than we are targeting. These groups and examples are summarized below.

Approach	Examples	
Goal Setting	Abstinence	Stop using the web entirely
		Stop using particular sites or applications
	Setting Time Goals	Limit amount of time on particular sites or applications
		Limit undesired activity to a particular time of day
Awareness	Reminder Cards	List and positive effects of improved behavior
		List and post consequences of negative behavior
	Personal Inventory	Keep track of the extent of improvement, and list positive effects of reductions
	Time Opposition	Ban target activities during times when the user normally engages in such activities
		Use scheduling approaches to disrupt negative activities (eg, schedule a meal during target times)
Supportive Interventions	Supportive Practices	Family therapy: engage family in preventing unwanted behavior
		Support groups: provide a social network of support

Table 8. Interventional approaches suggested in treating behavioral anomalies.

Goal setting can be done by abstinence – either from all surfing activity, or specific applications or sites. While effective in the immediate time frame, it is crude and may quickly lead to rebound since the behavioral control is fully extrinsic. Additionally, such approaches may prevent the user from accessing sites and applications that are necessary for normal daily activities, and thus may be impractical. A more nuanced approach is to set time goals or stop times to limit negative behavior. Here however the diversity of the internet works against the approach, since the time can be spent in any number of ways with varying alignments to user goals, and thus can be justified in whatever way the user chooses.

A clinical approach that tries to resolve the limitations of goal-setting is awareness-generating activities. The most simple of these simply bring to the user's attention the benefits of abstinence, or perhaps the costs of overindulgence. By posting these costs and benefits in clear view, the goal is to create intrinsic awareness or motivation. However, as discussed in the background section, unwanted or uncontrolled behavior is not merely the result of lack of motivation except in deeply pathological cases. Indeed, even when the user is highly motivated, there needs to exist tools that allow the user to measure and reflect their activities.

With these strategies in hand, we can now examine the extent to which such techniques can be automated or assisted by technological means, thus informing the creation of reflectOns for such purposes. Foremost, it is important to note that it is quite difficult without external aid to measure the amount of time spent surfing. Additionally, melioration effects and the mixed-use nature of the internet complicate unaided attempts to calculate time spent for various purposes. However, it is just as important to note that the external observer does not always (or even often) know the user's intent in using a site. As noted previously, something as simple as email may have multiple divergent use-cases based on the context. Given the difficulty in integrating use over time mentally, it may appear attractive to consider simply start and end times. Many users today however leave their browsers open all the time, and thus do not note the start time of a surfing session. Additionally, it is quite likely that over any appreciable period of time the user engages in actions that they themselves may consider positive or negative with respect to use of time. If one were so inclined, it would be simple to generate some reasoning for the use of time, thus making it an ineffective form of self-feedback. Lastly, it is also necessary to acknowledge that since the internet does have mixed uses, it is quite likely that even actions that the user may objectively consider a waste of time may provide some benefit to the user that, once again, is not knowable without context.

Since it is difficult to design a reflectOn that can successfully put web surfing behavior into different bins without user intervention, we have chosen to focus on the underlying difficulty the user faces when surfing the web: keeping track of how the flow of micro-decisions affects the use of time, and how quickly the decisions are being made. We recognize that full abstinence is generally too extreme an option for many users who only casually (rather than pathologically) misuse their time while surfing the web. At the same time, doing fine-grained self-control is generally too complex. As such, we decided to provide the user with a simple visualization of their time budget. By providing a user-configurable blacklist of sites that are in need of additional vigilance on the part of the user, we attempt to accommodate the ambiguity that is inherent in web surfing intent. At the same time, we expose only an aggregate measure over many site visits and clicks, thus reducing the decision making demands placed on the user. Additionally, we also provide a whitelist that allows the user to exclude "known good" sites, thus again minimizing the ambiguity in the measures to the extent possible without extensive polling of user intent.

In addition to time budgeting, we also recognize that by the time the budget is exceeded, the user may well have used more time on a particular site that he or she might have chosen to. We thus provide a system of focus tracking that examines the melioration aspect of web surfing by measuring the number of page transitions within a time window. This measure allows the user to back away from behavior that may be relatively low value by providing an ongoing measure of the dwell time (and by proxy, the interesting information content) of pages visited.

Lastly, while on-screen indicators are sufficient when the user is directly engaged in browsing, it is reasonable to assume that the pre-commitment to browsing often happens while the browser is not visible. As such, we have designed a modified mouse that physically provides the feedback in a manner that is attached to the computer, yet without requiring the user to bring the browser to the foreground. By forestalling some of the pre-commitment to web surfing, we hope to provide an avenue for the user to be mindful of their activity before they engage fully with it. The mouse also provides a haptic feedback when the user reaches their limit, thus providing a less intrusive warning when the limit is reached without obscuring the content of the page, thus further decoupling the time used from a particular action that may be justified away.

Design

The web surfing reflectOn has a two-part design. One part operates within the browser and provides observation, configuration, and in-browser feedback capabilities, while the second component provides feedback via a modified mouse. While we recognize that a multitude of browsers exist in the wild, supporting each of them sufficiently is outside the scope of this work. We thus chose the Mozilla Firefox browser as the target platform due to its combination of popularity and developer-friendly capabilities. The observational component was created as an "add-on" for the browser: a self-contained piece of JavaScript code that operates within the "chrome" or UI of the browser. This implementation allows the add-on access to user actions while maintaining sufficient isolation to prevent other add-ons from interfering with proper operations (sandboxing).

The add-on observes the loading of pages, tab creation, focus changes, as well as any mouse of keyboard activity taking place within the browser window. Thus it is possible to know when the user has begun and ended access to a page, as well as when the user has stopped actively using the browser. A self-resetting counter is used to create a timeout that effectively treats inactivity beyond two minutes as a loss-of-focus event and stops the browser from doing further timing. Additionally, the tab focus is tracked so that the user is not penalized for opening many tabs at the same time. This also allows us to track when the user is accessing various pages by switching between tabs rather than by clicking through pages one at a time. The combination of observations allows us to calculate both dwell time and total time per page.

The observation add-on maintains several different histories of user actions. To simplify selection of whitelist and blacklists, the add-on tracks sites the user visits, as well as the total amount of time spent on those sites since the installation of the add-on. Additionally, time spent on every site visited in the prior 24 hours is tracked in order to calculate time budgets. Entries in this queue are removed on a rolling 24-hour basis based on the start time of the entry. This ensures that time in the budget is reclaimed in a way that accommodates any temporal pattern of user behavior. The add-on additionally also tracks all page transitions within a 10 minute window, again removing items based on the time of the event on a rolling basis. This gives a local estimate of the "information value" of the pages the user is consuming, and identifies when the user is slipping into meliorative behavior. The observation behavior of the add-on may be disabled either by selecting the appropriate entry in the add-on's menu, or by simply entering Privacy Mode in Firefox itself. These options are shown in the figure below.

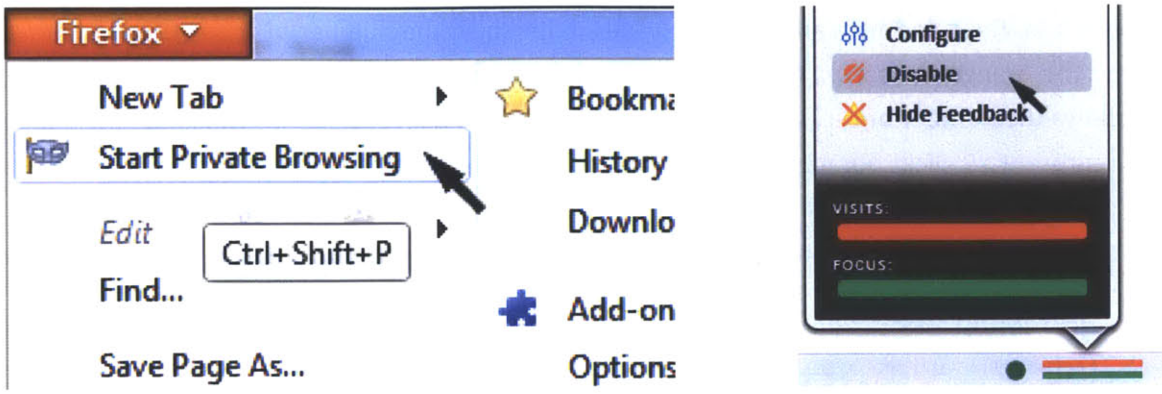

Figure 13. Methods of disabling observation functions of the web surfing reflectOn

The add-on provides a simple point-and-click interface for adding and removing sites from the blacklist and whitelist, which can be seen in the figure below (Figure 14). The user may choose to either budget time, or view focus information. Domains may be added to the lists either from recent history, or by typing in the domain. In order to accommodate the widest gamut of technical adeptness among users, we chose to limit choices to domains rather than paths or wildcard matches. This also has the side effect of simplifying the interface and presentation in general, since it is impossible to enter overlapping entries accidentally into the black and white lists.

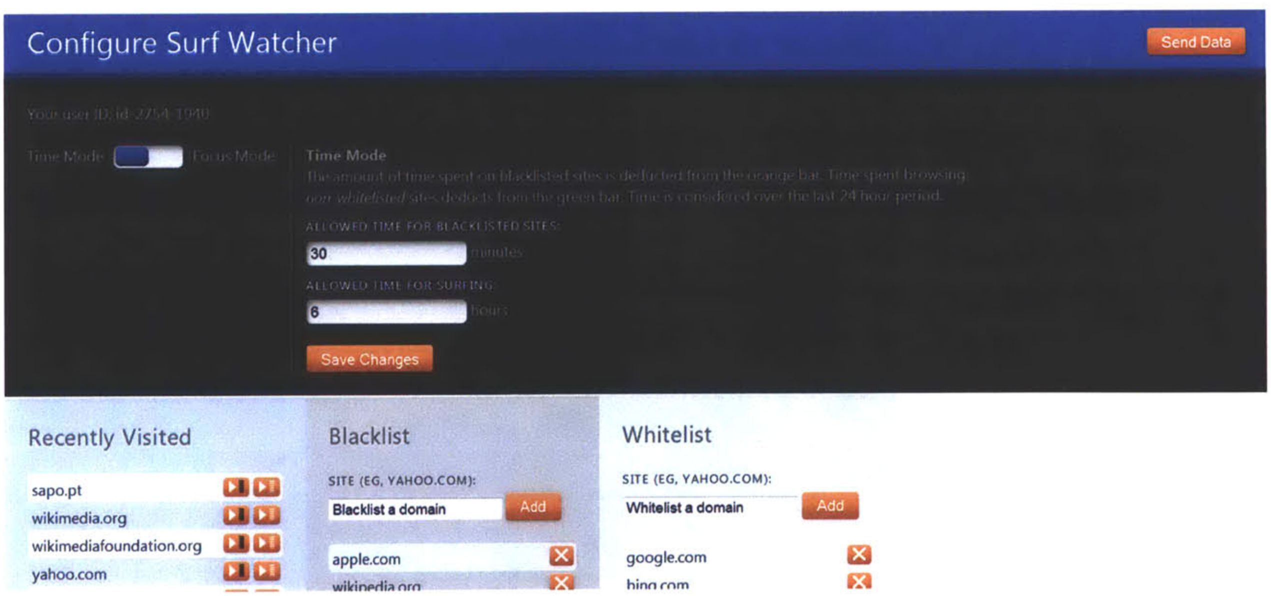

Figure 14. Configuration screen of web surfing reflectOn

In time mode, the add-on operates as a time budgeting tool. The user specifies the number of minutes to allow in the blacklisted sites, and total number of hours per day for surfing. The amount of time used is shown as two bars (Figure 15) both in the always-visible add-on bar, and in the add-on menu as a larger representation. The orange bars indicate time spent on blacklisted sites, while the green bar decreases as time is spent on any non-whitelisted site.

In focus mode, the same bars are used. However, both show the number of sites visited in the last 10 minutes. The orange bar shows the number of clicks among blacklisted sites, while the green bar once again shows the number of clicks in non-whitelisted sites. Via pilot testing, we chose to limit the orange bar to represent 20 clicks per 10 minutes, while the green bar represents 30 clicks per 10 minutes. The lower limit on the orange bar allows users to elect sites to the blacklist that they wish to be more vigilant about. The user may switch between the modes at will, though the add-on tracks both metrics in the background at all times.

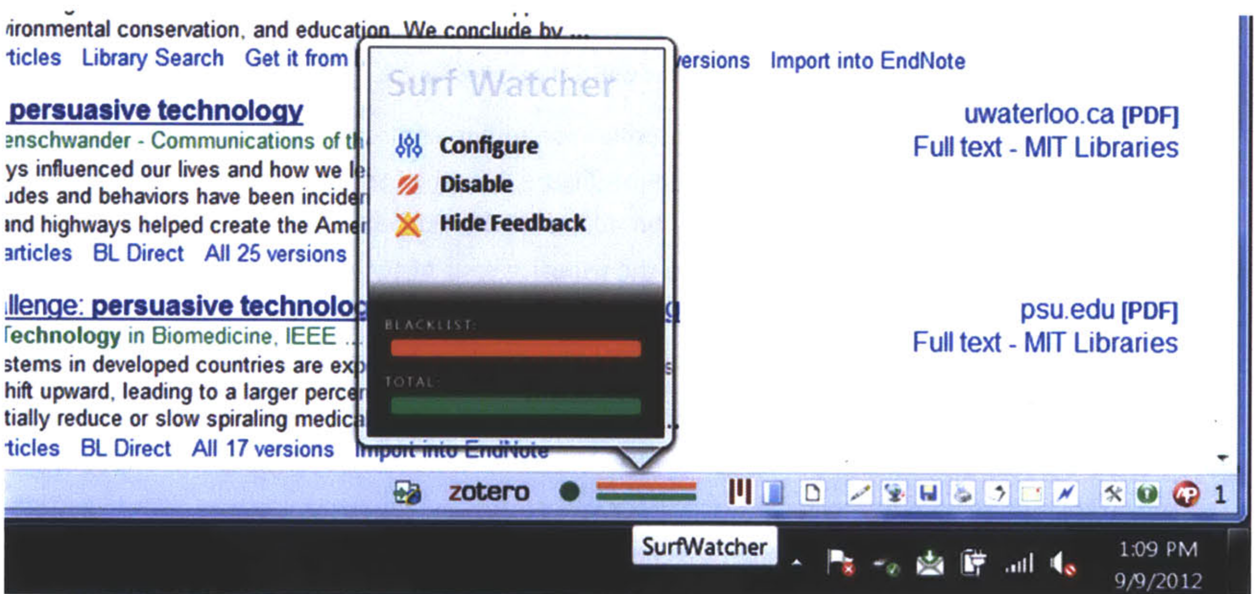

Figure 15. Web surfing reflectOn on-screen feedback.

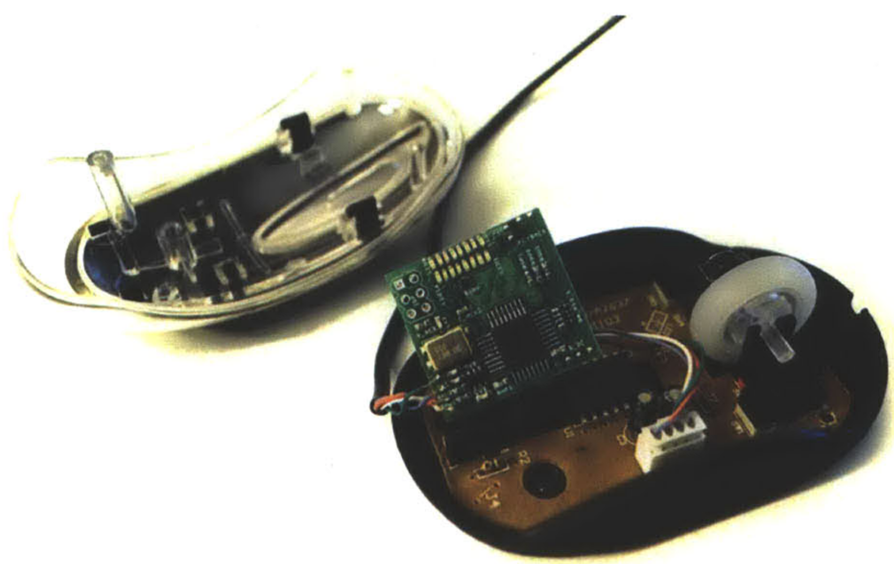
Figure 16. Internals of first mouse prototype showing the circuit board with LED bar display.

The add-on has the capability to communicate with the modified mouse (Figure 17), which embodies the bars shown on the add-on as 2 LED displays of matching colors. The mouse also has a haptic actuator that allows it to vibrate gently when either bar is reduced to empty. In order to avoid irritating the user, the alerts are rate-limited to once per minute. The mouse retains the last output from the browser even when the browser is out of focus or closed, giving the user a snapshot of the browser's display in a glanceable form.

Figure 17. Modified mouse for use with web surfing reflectOn

Feedback Modality	Feedback	Feedback Tradeoff
Bars on browser	Shows remaining time or level of focus	Glanceable and close to content, but only visible when the browser itself is in focus.
Popup warning in browser	Appears when bars are depleted. Dismissible with a click.	Immediate. Intrusive.
Bars on mouse	Mirror bars on browser	Glanceable. Can affect browsing goals before the user starts surfing. Removed from direct perceptual frame.
Haptic feedback on mouse	Vibrate when bar reaches zero	Immediate, but transitory. Removed from direct perceptual frame.

Table 9. Web surfing addon: feedback modalities.

We have designed the system with four separate feedback modalities to allow us to test various feedback strategies. The strategies trade off simplicity of feedback with directness of information (see

Table 9). Two of the forms of feedback live exclusively on the screen. The bars provide a glanceable feedback without requiring foreground attention. Once the bars are depleted, and dismissible popup appears. Due to the persistence and positioning of the popup, this feedback is highly immediate and intrusive. The bars on the mouse echo the bars on screen, but are insulated from loss of foreground focus on the screen since the display is physically instantiated. Additionally, it provides a view that can be accessed without interacting with the browser, which allows the user to set initial goals before she starts browsing. Finally, haptic feedback on the mouse is unobtrusive but immediate, allowing the user to self-regulate without disrupting workflow in the same way a popup may. We hope in the user studies to use these various output methodologies to gain a better understanding of the level of abstraction needed in reflectOns.

Application 2: Food Intake

Food intake is an excellent example of micro-choice, in that we must engage in the act of eating throughout our lives, multiple times per day. Even though we decide literally millions of times within a lifetime whether to take one more bite of food, the individual choice is highly multi-factorial and difficult to consider individually. The costs of the choices over a lifetime, however, are manifold and readily apparent. More than 60% of Americans are overweight (Ogden et al. 2006), and the problems with excess weight are apparent worldwide. The side effects of excess body mass range for hypertension to increased chances of cancer. While food intake is not the only contributor, it is a significant factor alongside exercise towards this trend. This information is widely known, but it is extremely difficult for an individual to consider these negative effects caused by long term aggregate behavior in making a choice that must be made so often. It is also unrealistic for the human cognition to cope with reasoning about matters in such detail so often.

Approach

For this scenario we describe a reflectOn that provides timely reminders to allow the user to gain an opportunity to reconsider their micro-decision at a point where consideration of wider goals can lead to a realistic change in behavior. By focusing on a discrete decision points within the act of eating a meal rather than on every bite taken, we can reduce the number of choices necessary significantly from the number of morsel-associated micro-decisions.

While it was possibly to at least approximate traditional caloric counts, we realized early in the design process that such systems were not tenable in the context of human cognition. On one hand, calorie-counting has to deal with the pragmatic difficulty of determining the caloric value of food. While we may be able to know what the caloric contribution of a part of even a whole meal is, it is difficult to determine how that value is distributed within a meal, and two persons consuming the same mass of food from the same meal may end up with very different caloric intake. There is also the pragmatic difficulty of sensing caloric intake without significant user intervention, which disrupts the flow of the user in a flow-sensitive task such as eating. On the other hand, even if the user were able to know the number of calories taken (or we were able to let the user know the amount), calories are difficult to reconcile with the effects of taking in too many calories.

Considering the above, we felt that it was a better approach to provide the user with feedback that allows for the creation of new habits and intuitions that lead to overall improvements in food intake behavior. We focus in particular on rate of food intake (time between bites) and total number of bites. These feedback approaches take advantage of the ways in which we determine when to stop eating.

Other than the overt and active decision to stop at a certain point, the body incorporates two physical mechanisms to determine when a person stops eating. The first is the pure distention of the stomach from the volume of food. This is the "feeling of fullness" after a meal. However, this is generally a poor indicator for a stopping point: today's highly processed foods can allow us to exceed caloric requirements without being full. For example, the Dairy Queen Caramel Cheesequake® Royal™ Shake can supply up to 1000 calories in a single serving (Dairy Queen n.d.). Additionally, the signal is delayed in most people prone to overeating, thus making the signal moot. The second signal is received directly in the brain by a blood sugar level sensor built into the brain itself. When we have digested incoming food to the point where the blood sugar rises, the immediate "hungry" sensation recedes. However, this signal is delayed by the processing rate of the digestive system, with the signal of satedness arriving approximately 15 minutes after food intake begins (Brand-Miller et al. 2006). This allows for plenty of time for the person to take in too many calories.

In this design, we can approximate or make use of both signals. To use the first signal, we can provide the user notification when a certain amount of food has been eaten. Over time, this has the possibility to give the user an understanding of how "full" their stomach should be when they receive the signal, thus providing a second stomach distention set point that allows the user to stop eating at the appropriate time. This can also give the user an estimate of how much of meal should be eaten (since one cannot see the food already taken in at that point). This additional intuition could allow the user to estimate a proper portion size later when the device is not at hand. To use the second signal, we can provide feedback when the user is eating too quickly. By urging the user to slow down, we can attempt to affect the cadence of food intake, slowing it to the point where the number of calories taken in when the brain receives a satedness signal is still small. As a result the person would physiologically become less inclined to take in more calories.

Design

In order to approach the problem of food intake, we chose to design a device which can attach to an eating utensil for the duration of a meal (Figure 20). We chose this form factor for a number of reasons.

While it is possible to computationally determine when and how a user is eating in a lab scenario, the problem is considerably more complex in the wild, where it is impossible to instrument all possible eating areas that a person may use even during a single day, which bars camera-based approaches in most cases.

Figure 18. First version of utensil attachment.

Figure 19. Final version of augmented utensil used for studies.

On-body devices can be used to determine when the user is eating, such augmentation can be cumbersome to use, and a device located on the body is generally ill-positioned to provide non-disruptive feedback. As a result, we chose to look at the overall ecosystem of items used during food consumption. Discounting eating by hand, one of the most common items is the eating utensil (often a fork or spoon). By providing an attachment for utensils, we gain a usable sensing vantage point. More importantly, since the utensil is held in hand, we gain the ability to provide feedback in reasonable proximity to the locus of food intake via a number of means. Additionally, since most users use a variety of utensils throughout a day, be it for convenience, availability or simply because the other has not been washed yet, the attachment form factor allows us to continue monitoring the user without having to modify every utensil that the user may encounter.

Rather than attempting to sense eating directly, we use an acceleration sensor to determine the user's movement when s/he takes a bite using the instrumented utensil. From the count of bites, we gain aggregate knowledge about when and how much (in approximate volume) the user has eaten. By providing feedback on eating speed and byte count, we can give the user discrete or continuous decision points at which to evaluate eating habits against longer term goals.

Implementation

The attachment is implemented a set of hardware and software modules. **Error! Reference source not ound.** This section provides an overview of the various high-level components. The hardware allows for an optional non-volatile storage subsystem, a wireless communication subsystem, and a display subsystem. Modular drivers within the firmware of the MCU allow for selection of output modality. The system is powered by an on-board battery. The battery can be charged using an external charger.

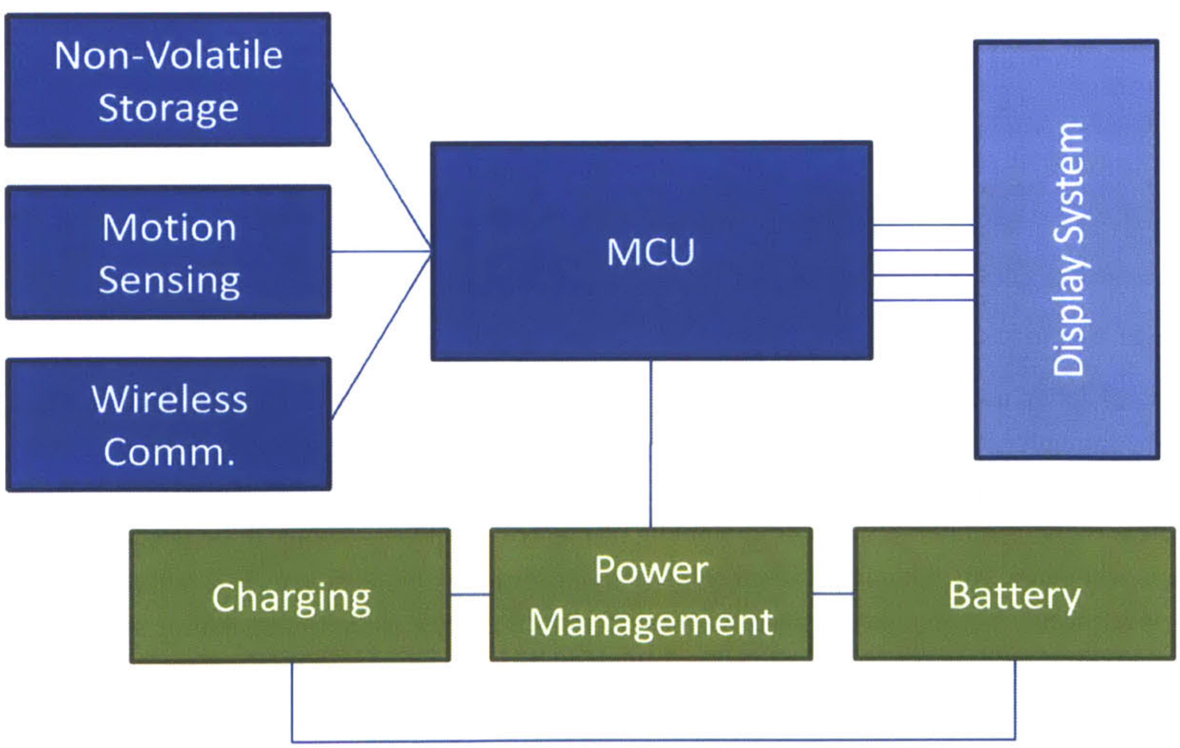

Figure 20. High level diagram of utensil attachment

The data and control flow can be seen in Figure 21. The sensing module provides four outputs for consumption by the recognizer: a timebase, raw x/y/z acceleration, mean windowed acceleration (approximating velocity when multiplied by time window), and windowed variance. These inputs are consumed by a recognition system which in turn produces a time stamped message whenever a bite is detected, as well as a quiescence message whenever the recognizer infers that the device has been out of use for some time. The messages are passed on to the display drivers, which take the recognition messages and determine the best means to display the results. Additionally, each driver can be switched to either show feedback based on number of bites or intake rate. The quiescence message causes the driver to disable display to conserve power. Additionally, the microcontroller activates power saving once quiescence is detected for some time. In this mode, communication, display, and storage modules are suspended in low power states, and the sensor is checked at a reduced rate to determine if there is sufficient motion to resume normal operations.

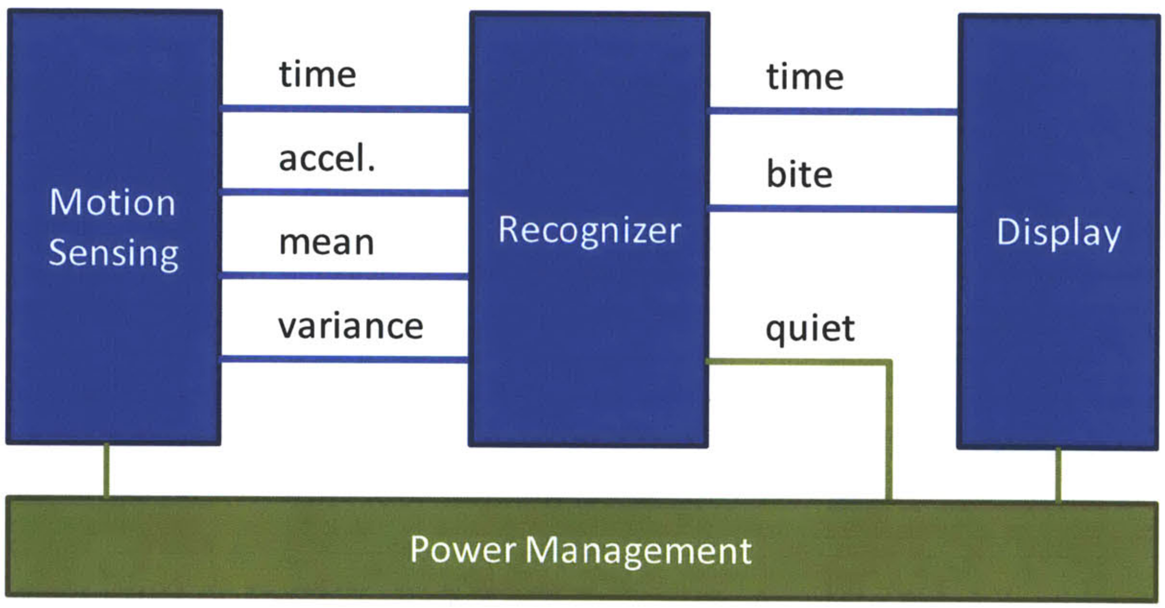

Figure 21. Data and Control Flow

Initial validation uses the radio in streaming mode, allowing the data from the sensor to be received by a nearby computer for recording and analysis. Figure 22 shows the physical layout of once such prototype. As can be seen, these initial prototypes included a number of production-oriented featuees such as wireless charging and breakaway programming connections that were meant to be used in the context of a longer study. The data from these prototypes was retrieved in real time over wireless link, manually labeled, and the results are classified to produce a decision tree classifier capable of segmenting out bites. This is then further tuned to improve separation of the events. The classifier is then converted to a C structure and code and compiled into the microcontroller's firmware as part of the recognition module to produce the production firmware used in the user studies.

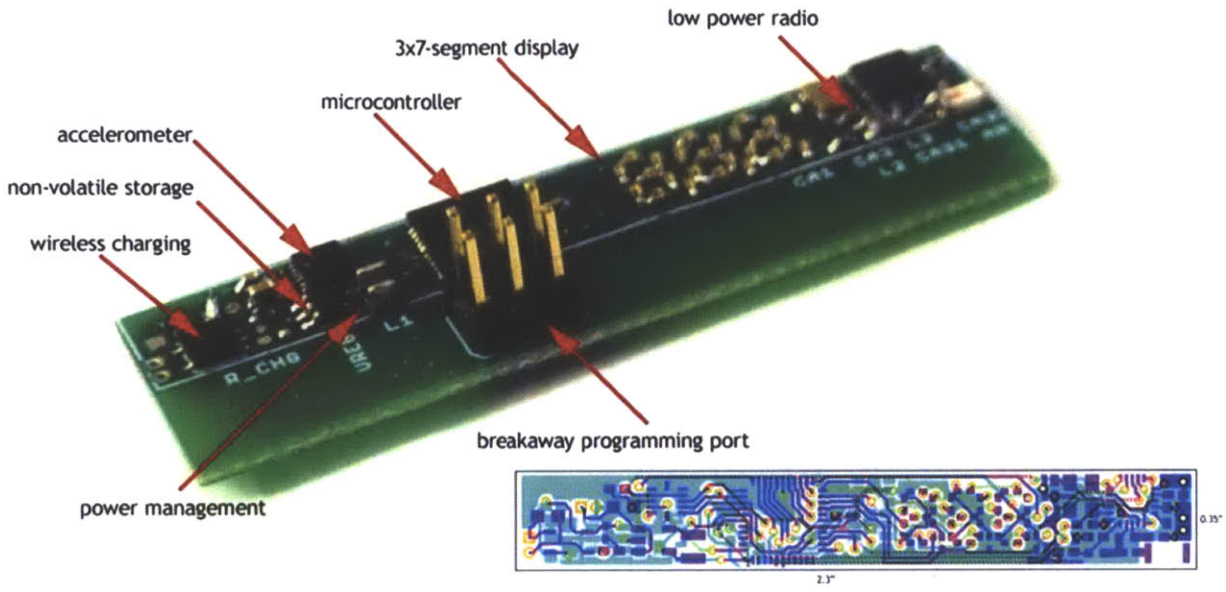

Figure 22. Layout of components in one variant of the utensil augmentation board.

Application 3: Spending Behavior

Finance provides several interesting venues where micro-choices abound. The average person makes many daily choices that have a financial component. Choosing to have breakfast at home our outside, choosing to take the car or ride the bike to work, choosing whether to get a coffee after lunch, or choosing between oatmeal or bran cereal while grocery shopping: these choices all affect finance in some direct or indirect way. On the other hand, our feedback regarding finance is relatively shallow. Some awareness of bank balance, credit card balance, and overall size of a bill make up the bulk of our "data" for reasoning. While this combination is itself sufficiently treacherous for reasoning about finances, the situation is made worse by our inherent aversion to risk (Hertwig et al. 2004). This causes us to see small purchases as benign, and large ones as harmful. For example, it can be expected that one would check if a $100 purchase is real, but not if the gas station added $5 to the bill at every visit to the station. Additionally, we do not assign causes to aggregate expenditure. Thus, most people would have difficulty estimating monthly cost of having lunch out, but not monthly cost of cable TV, though the prior is likely more expensive for those who do not bring lunch from home. These factors combined with the fact that we currently receive input about the financial consequence (credit card bill, wallet running empty) much later that the actual stimulus and actions that precipitate the subsequent result (Figure 23) make it difficult for a person to cope with reasoning about micro-choices.

A common cognitive bias is to perceive past misbehaviors as aberrations, while viewing future misbehaviors as unlikely (R. Thaler 1985). This bias is perpetuated by the decoupling of the cause and effect of an action. Therefore, it is further imperative that action and consequences are tied together, even if the action itself is difficult to reverse. This provides a natural juncture in the choosing process for the user to reevaluate the utility of a choice. On the other hand, it is common for users to over-focus on micro-choices with little actual consequence (Kahneman & Tversky 1979). These choices suffer from low utility even when they are completely non-aberrant in the context of the behavior of peers. ReflectOns can be designed to provide positive feedback in such cases and help lower the cost of micro-decisions with low "value" relative to long-term goal-oriented activity.

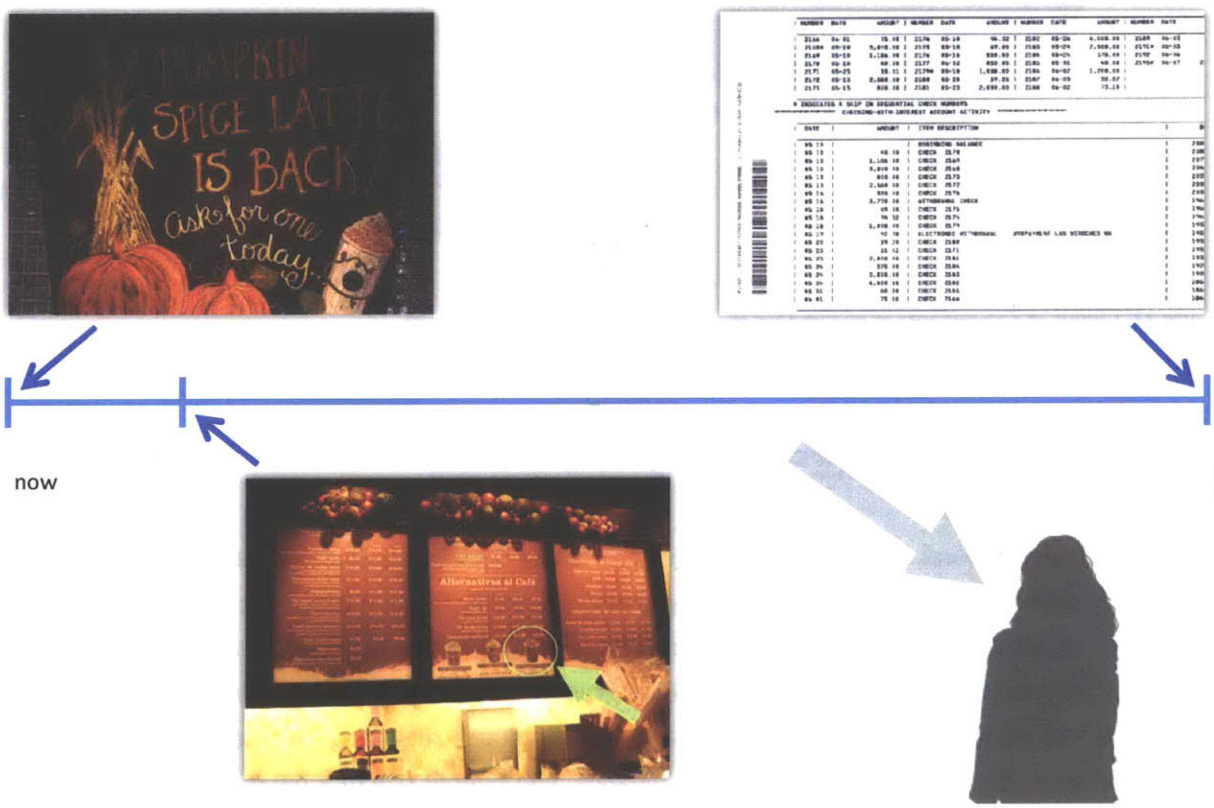

Figure 23. Feedback decoupling in finance

Approach

In designing reflectOns that operates on financial behavior, it is tempting to start with the goal of "preventing" the user from making a choice that is not in alignment with their goals. This seems tenable at first glance, but in fact requires the equivalent of perfect information about user intent to work correctly. It is very difficult, for example, to separate the fact that a person goes to the lobby of a nearby hotel to get coffee most of the time, but to attend a meeting occasionally. Using existing technologies, we can only guess (but not *know*) the intent of the person. If a system based on such premises acts prematurely, it would severely degrade the trust the user has in the system. This trust is essential in the functioning of a system that the user is essentially using to augment their cognition, and thus the loss of trust can be considered a fatal flaw.

In an attempt to avoid that trap, the system may choose to err on the side of caution, and not provide feedback until the user is at the counter of the coffee shop. However, now the system runs aground in the trap of pre-commitment (Ariely & Wertenbroch 2002). While it may have been possible to convince the user that it would be better to avoid going out for a cup of coffee to save some money, if the feedback arrives after the user has already chosen the coffee and can almost taste it, the negative impact (or conversely, the amount of self-control required) becomes quite high. As a result, the user is liable to feel very negatively about the system. Once the system is perceived as "no fun" or "the bag guy," the user is once again unlikely to pay attention to it.

In light of these observations, it is unrealistic to believe that we are well equipped technically to prevent a particular episode of negative behavior. However, this does not mean that it is impossible to modulate the aggregate utility of a purchase in order to nudge the user in a direction that aligns with their goals. By looking at trends of expenses, it is possible to estimate the existence of expenditure habits, as well as the relative costliness of the habits. By providing feedback on each action that contributes to the habit after the choice has been made, it may be possible to lower the utility of negative habits so that the user becomes less prone to the actions. Conversely, it may be possible to provide reassurance for positive habits, which can further bolster the user's perception of the system and its feedback. It is also important to note that financial responsibility and health are aggregate long-terms results that are minimally affected by individual deviations. As such, if it is possible to have an effect on behavioral patterns, it may lead to a better long-term result than the prevention of a particular negative episode.

Design

As discussed above, we hope to provide feedback on trends of purchases which can add up to have larger effects. Additionally, we hope to affect the perceived utility of micro-choices. Much spending today happens via electronic means, which help to further conceal the nature of purchases. We thus chose to build the system within the form factor of a credit card (Figure 25**Error! Reference source not found.**). Whenever the card is used, a standard magnetic stripe sends credit card data for the purchase. However, the system also uses a special reader to gain information about the exact details of the purchase. In its final incarnation, the prototype carries out all processing on board. Based on the results, the system then can display a light color based on the trend of spending. Feedback in provided right after the user pays for an item or items, thus getting the "last word" in the total utility of a choice. We call this design concept SpendTrend.

The first version of the SpendTrend prototype (Figure 26) relied of processing systems completely embedded in the reader that communicated with a central system in order to create feedback, while the device itself was fairly simple, and carried only a system. However, this required remote systems to retrieve user information very quickly (during a card swipe) and come back with a result. This proved to be a complicated proposition, especially given the slowness of current credit transaction processing systems, which may take up to a day to notify issuing banks about transactions. Additionally, since the reflectOn accesses much more data than a simple credit authorization contains, there is an additional concern regarding privacy. As such, the version 2 device (**Error! Reference source not found.**) was esigned with all processing completely on board. With this device, the reader communicates an itemized receipt to the card directly. The card contains the necessary on-board processing and trend tracking systems to decide what display to provide.

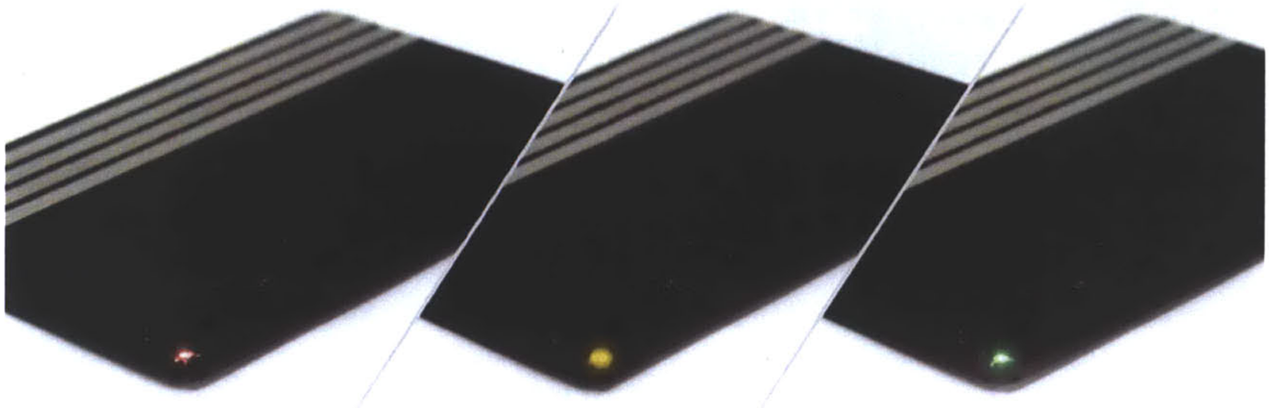

Figure 24. SpendTrend version 1 prototype

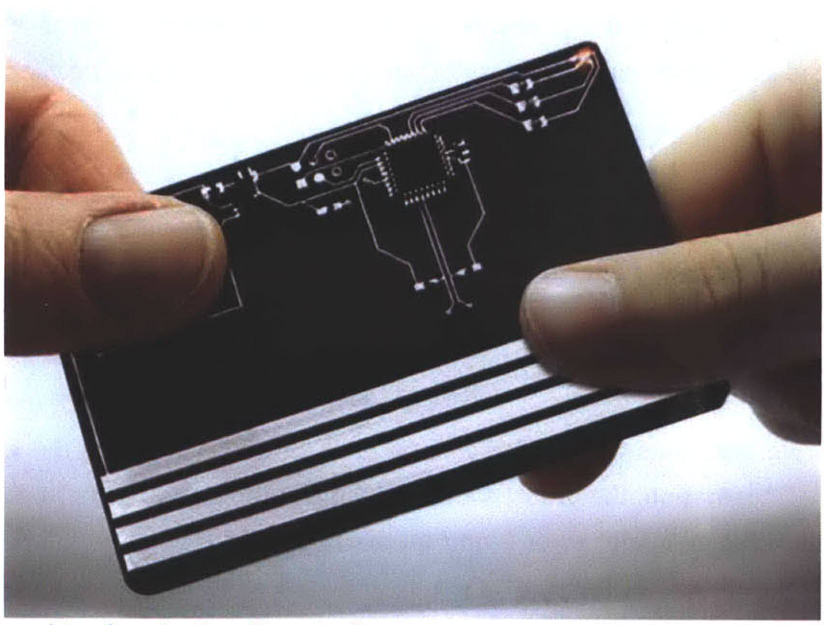

Figure 25. Prototype SpendTrend version 2 credit card reflectOn

Every time the SpendTrend card is used to pay for a purchase, it receives a categorized list of items on the receipt from the reader. These categories (such as eating out, car-related, grocery, etc) are used to update a set of bins of expenditure. At the same time, older entries are discarded, and a time-weighted average rate of growth for each bin is calculated. The outcome of the bin that shows the most change due to the current receipt is chosen and shown as the color expressed by the card. As a result, the user can reason about the purchase to some degree and adjust their perception of the utility of the purchase accordingly.

While the output modality of the card is limited, it is important to remember that aggregate purchases (such as at a supermarket) are generally less common than individual purchases of goods. Additionally, the information is stored by the card and can later be accessed for a more detailed review or use in concrete budgeting approaches where the additional data allows the user much greater insight into their own behavior. This can also be combined with a simple savings goal (e.g. $100/week) where the user received feedback based on how purchases affect that goal due to deviations from expected behavior. Though the feedback becomes more complex, SpendTrend can thus be repurposed to promote savings as well.

Implementation

The first version of SpendTrend (Figure 26) outsourced the bulk of the processing task to the reader system, which collected the receipt data and submitted it to a user data store, which was then queried by a separate module which calculated per-user feedback. This feedback was used by the controller of the reader to directly program the display output of the card. As previously noted, this required the user to leave the card in the reader for some time while the system connected with the server which stored the user data. In practice, this would be a significant impediment to adoption. In addition, the control of the display was less than optimal. Since the display lacked power management electronics, it could not turn itself off. In addition, this caused the display color to drift as the system ran out of energy and shifted color. This left open the possibility that the user would receive the "wrong" message from the device. Despite these shortcomings, the system benefited from a central data store. This central store could easily calculate social metrics and dynamic limits based on use of many users, and provide more complex temporal integrations than can be reasonably calculated on a low-power MCU.

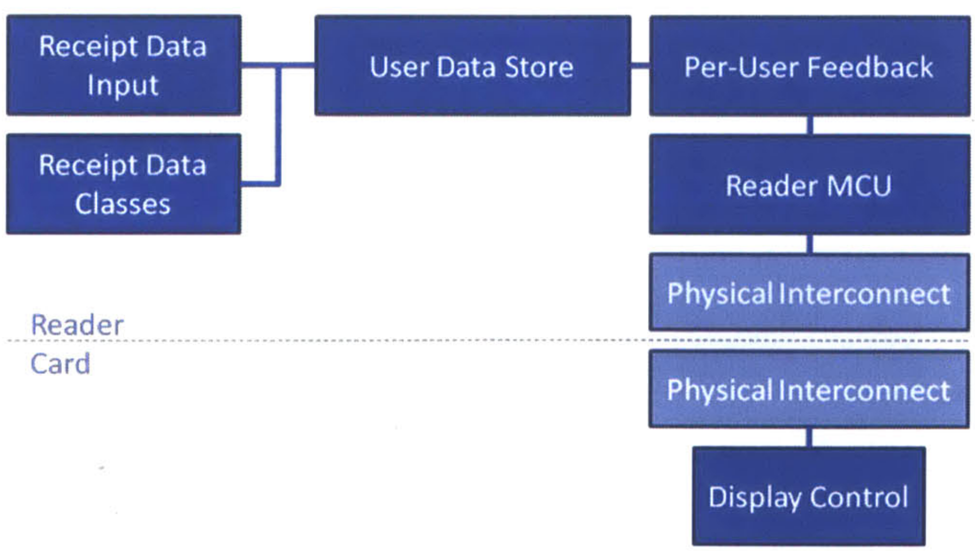

Figure 26. SpendTrend architecture, version 1

To address the shortcomings of the first generation prototype, the second version of SpendTrend was created (**Error! Reference source not found.**). This version uses a fairly simple reader and packages most of the complexity f the device onto the card itself. The card is powered by a set of super-capacitors, which are charged with pulses from a power supply on the reader. Sufficient power is stored during a swipe of the card for the MCU to categorize the receipt data, calculate feedback, and operate the display for several seconds. However, due to the limits of on-board processing and power systems, it's difficult for the card to show social metrics and do complex temporal trend calculations. This is offset by greatly increased privacy and ease of use, since the card can be swiped quickly and does not require network access at the reader to provide feedback. Additionally, the card can be programmed to maintain a record of usage that can then be read by a special reader at home to fill in budgets, support concrete budgeting, or provide greater insights into spending trends and causes of various feedback generated by the card.

The SpendTrend cards of both versions are designed to be usable with a standard swipe-type card reader. Electronics are confined to the top two centimeters of the card, while the bottom portion is made of printed circuit board substrates of the same thickness as a standard credit card with traces on only the front side. The bottom back portion of the card is reserved for a traditional magnetic stripe that can be read with a standard card reader. Even if the reader is unable to provide the additional information needed by SpendTrend for tracking and feedback purposes, this allows the card to still be used.

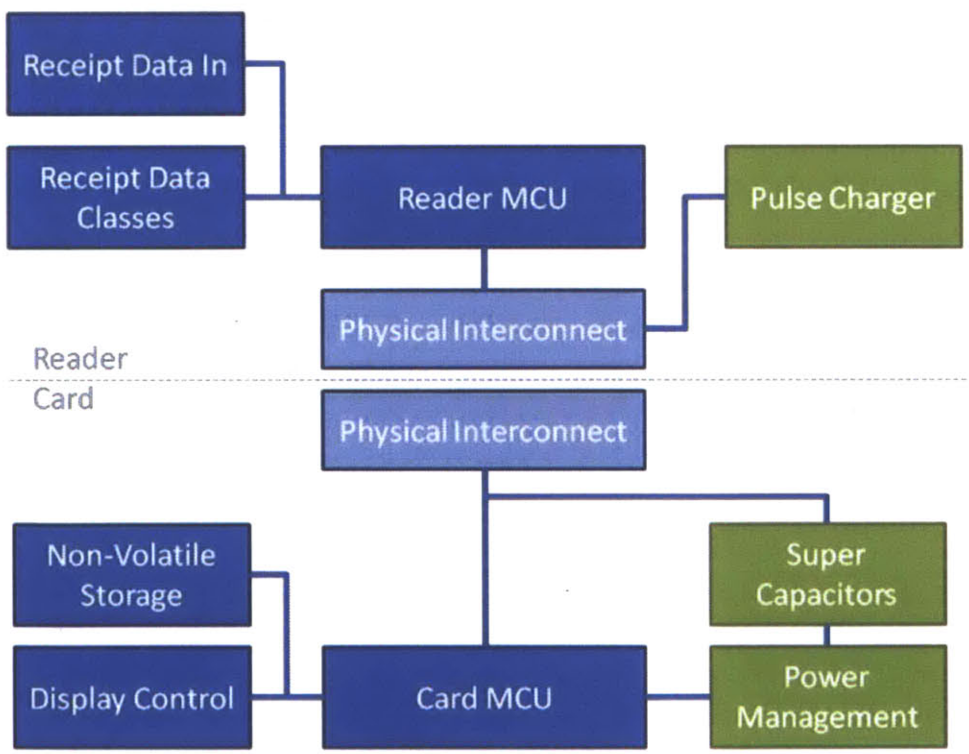

Figure 27. SpendTrend architecture, version 2

When the SpendTrend card is used, the processor first takes the receipt data and stores it for logging and future use. It then discards any line items that constitute a "free" category (uses such as grocery items, fuel in case of commuters, etc) as specified by the user. The remaining line items are then compared with the internal bins of expenses maintained by the system. These bins are configured to currently conform to the credit card usage bins (shopping, travel, services, food, etc). However, it is simple to support much more granular bins. For all affected bins, the time weighted average expenditure over a week is calculated, and if any of these rates is higher than the expected threshold for that bin, calculated as a percent of the time-weighted remaining budgeted amount for that bin or a superset of the given bin, the largest absolute difference is passed on to the display module for output.

Application to other payment platforms

While the above discussion focuses on the SpendTrend prototype, the feedback methodology and design principles can be applied to other payment modalities. Applied to near-field communication (NFC) based payment systems, it is possible apply the basic software unaltered. The point-of-sale transfers the purchase information over the NFC channel. The device used to handle the NFC payment request then carries out the same processing as the version 2 SpendTrend card and provides feedback on the NFC-enabled device. The primary advantage is that the user need not use a specialized card to achieve the same effects as SpendTrend, which reduces cost of deployment as well as the complexity of processing that can be done.

Application 4: Power Usage

Use of energy provides quite a challenge to human cognition. While we may understand at an intellectual level that energy is being used to power a lamp, a refrigerator, or a car, it is quite difficult to assess the relative amounts of energy needed for such tasks. There is no common unit of measure that suffices, since each device expresses output differently. We may be able to compare to incandescent bulbs by measuring output, but it is hard to imagine the "output" of a refrigerator or a hair dryer, because we lack convenient frames of reference (Tversky & Kahneman 1986). As shown by studies in cognitive science, without such frames of reference, it is difficult for us to grasp a metric even at a rational level. Further, each device we own uses energy in a different ways. An incandescent bulb tends to draw a continuous flow, while a refrigerator draws large gulps of energy. Since users lack good metrics of measurement, it is difficult for them to act upon their intentions.

Additionally, the energy requirements of the devices that are easily observable (such as light bulbs and microwaves) are eclipsed by the energy use of appliances that provide services such as TVs, water heaters, refrigerators, etc. As a result, the energy footprint of a person is highly correlated to a variety of other actions, and the correlation is difficult to deduce. Nor are all of these factors directly influenced by use. A refrigerator will consume energy regardless of use, though one can modulate it by opening it less often, or choosing to buy a more efficient model. However, the same reasoning does not directly apply to an entertainment setup, where a receiver and plasma TV may easily consume kilowatts of energy, but more sporadically.

The problem is further complicated by the fact that we receive feedback on energy use in such diverse ways. All of the energy use in the average house (sometimes sans heating, if it is electric) is aggregated into a single kilowatt-hour metric and shown to the end user. Moreover, while energy prices vary all the time, this is not made apparent, and thus users have no good way of knowing the cost of energy in the way they know the cost of gasoline, for example. The feedback is also highly delayed, as with finance. This leads to decorrelation between cause and effect, and leads the user to feel unempowered to change their energy usage behaviors. These circumstances provide fertile ground for the design of reflectOns that can provide users with better understanding of both their behavior, and the consequences of their behavior.

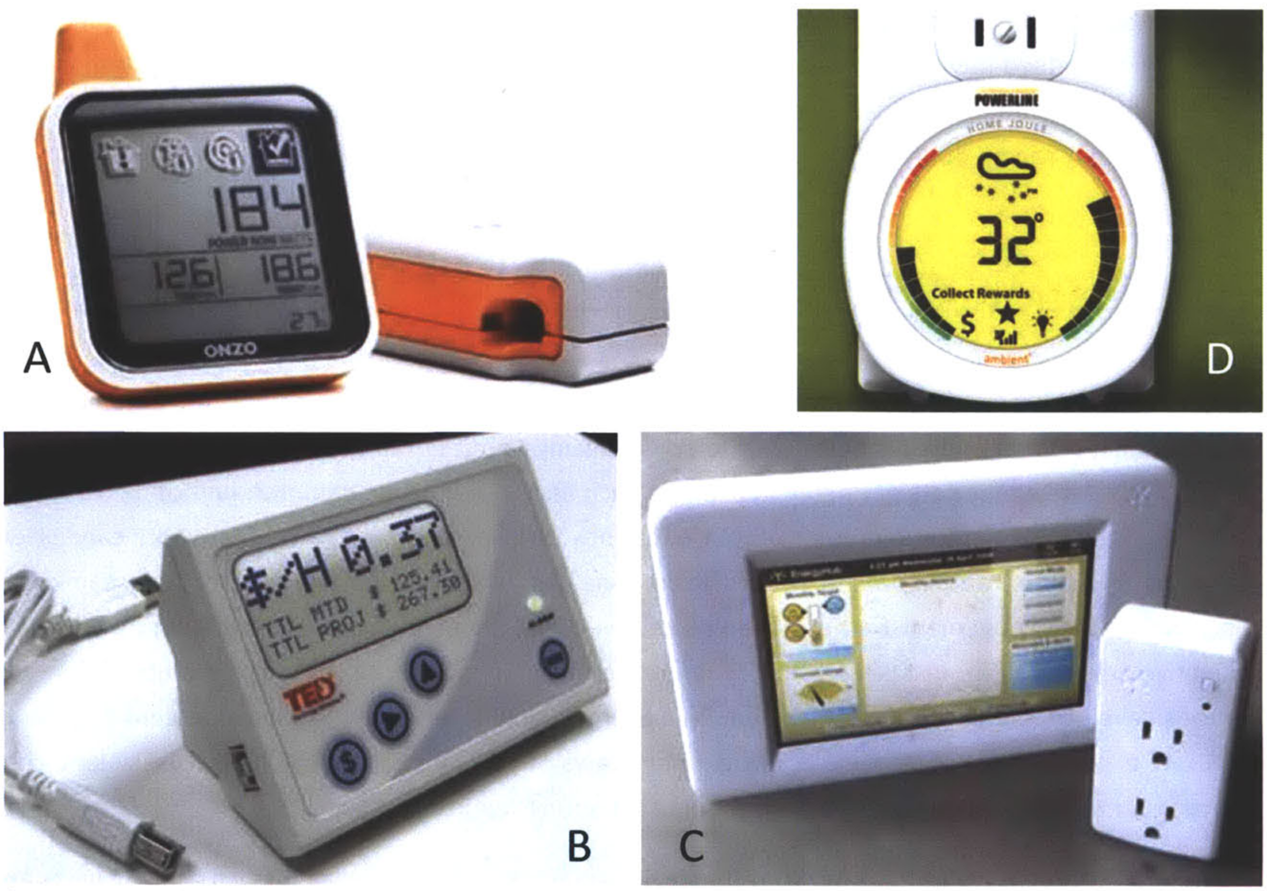

Figure 28. Energy metering solutions on the market. Counterclockwise: (a) Onzo smart energy kit (Onzo 2008), (b) The Energy Detective or TED (Energy, Inc 2008), (c) EnergyHub (EnergyHub Inc 2011), and (d) Energy Joule (Ambient Devices 2011b).

Approach

There have been an influx of devices to the market that allow the basic measurement of energy use (Figure 28), and very much like the power meter used outside homes or on the test bench, there are well-documented approaches to measuring energy use. However, energy needs are such a basic part of the daily life of a person living in the developed world that holistic behavior change is quite difficult. A person's energy profile is made of literally hundreds of both active (close the fridge quickly) and passive (where or when the fridge was bought) micro-choices. While most of the energy meters have done a good job of bringing feedback about energy closer to the time and place where the energy use occurs, we feel that products currently or expected to be in the market do not yet attack the problem of converting these values to reasonable units, or providing good options for comparing the energy use at different locations.

An additional concern with any Design is the units to be used. Energy is by definition invisible to the human senses, being only the potential for work rather than work itself. Units such as watts and horsepower apply this potential to a known task, providing a uniform measure, but one that has little bearing on everyday life. Finally, for the reason of price fluctuations mentioned previously, cost is often not a good metric, because it confounds two different integrated amounts. Long term behavioral patterns are best supported by repetition(Geier et al. 2006), and the confusion of the two variables makes the overall result less sensible for creating behavioral change by reducing the consistency of feedback for a particular action taken by the user.

In order to address these issues, we chose to use a measure that is easy to compare mentally, but not necessarily meaningful, but allows the user to compare two devices at a glance. While color was initially a choice for an ambient display (Ishii & Ullmer 1997), it limited the legibility of the system in uncontrolled a variable lighting situations. We eventually chose to use speed as the indicator, much as in the traditional power meter. It has the advantage of a large dynamic range, but provides the ease of legibility we sought. Additionally, we acknowledge that there is need for objective measures at times, and chose to expose a secondary, less visible measure that can be examined when the user chooses. In order to tie cause and effect, we chose to design the initial prototype in a form factor that fits on the plug itself. While this provides some limitation, we have ameliorated this shortcoming by providing remote readout capabilities to the devices that allow a remote system to query, process, and display all metrics collected by a remote unit.

Design

The Watt Watcher prototype is an inline module with a plug and a built-in outlet (**Error! Reference ource not found.**). The device is designed to occupy only one receptacle of a dual outlet, though the final prototype version is a little too large to allow two Watt Watchers to be placed on a dual outlet. The device draws minimal amounts of power when not in use, and is self-powered. The outlet in located centrally, and is surrounded internally by a set of side-firing LED lamps in a circular pattern along the edges of the device. These lamps act as the first display of the Watt Watcher, and can be set to rotate at any desired speed around the device. We use the rate of spinning of the light around the device as the ambient indicator of energy use. Since the movement is perceived as a throbbing from the sides, it can be used even when the front of the device is not directly visible. The second display on the device is a tiny 7-segment LED display on the top of the device. This can be used to show more specific information such as kilowatts per day. We also have internal capability to convert this to a comparison unit (such as number of 60W light bulbs left on for a day). This can allow the user to gain more detailed information about the energy use of the device connected to the Watt Watcher.

Due to the high voltages present internally, we chose to make the Watt Watcher completely sealed. Access to the various modes of operation of the device is instead handled by a radio interface. This interface allows the user to query the device remotely, change its operational mode, and even use one Watt Watcher to view the results of a different device that is perhaps located in an occluded position. Additionally, this opens the possibility of whole-house monitoring systems using the Watt Watcher as a measurement device.

The Watt Watcher can be used in three modes. The first mode offers a view of the instantaneous energy use. This can be used to quickly gain an understanding of how much energy a device that is plugged in uses without having to allow the Watt Watcher to "train" against the device to some time. The numeric display in this mode shows the equivalent number of 60W bulbs that would consume the same amount of instantaneous power. This mode of operation may be good for an outlet where many devices are plugged in at different times.

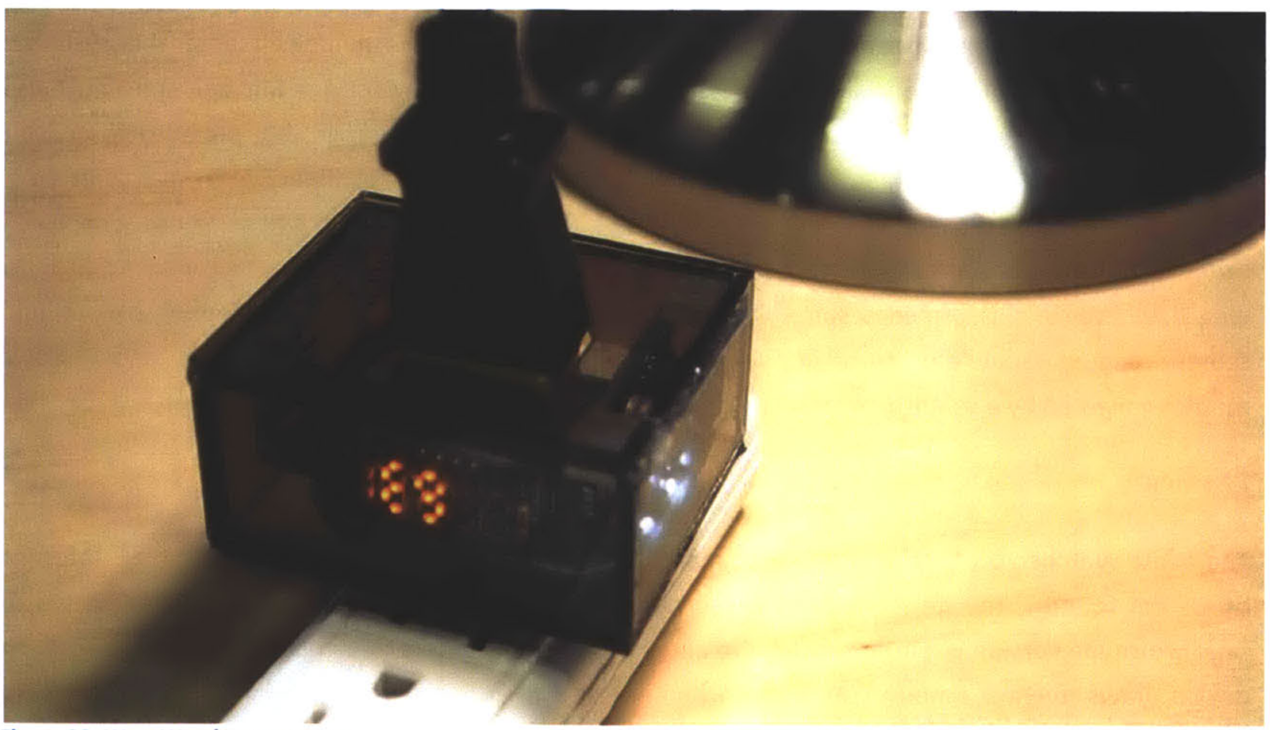

Figure 29. Watt Watcher prototype

The second mode uses an aggregate use model. The Watt Watcher must observe the device's use pattern for several hours for to understand average power use. After the average energy use has stabilized, the rotational display shows average energy use, and responds when a local average (for the last 60 minutes) changes relative to the average use over a longer period. This mode is helpful for devices such as heaters, refrigerators, and televisions that use energy sporadically, but may draw more energy when the use pattern changes. The goal of this mode is to provide feedback of actions that lead to higher than normal energy use. The numeric display in this case shows the number of 60W light bulbs that would have to be left on for a day to match the energy use of the device.

The third mode of the device is remote display. In this mode, the device is paired via the radio to another device. The device acts as if the load connected to the remote device is connected to it instead. This allows the user to monitor an occluded device without having to peek behind an obstruction regularly. This may often be the case for water heaters, refrigerators, TVs, etc. In addition, it is also possible for a Watt Watcher to be used to display aggregate energy use across all Watt Watchers in a home. While we do not feel that this is a valuable display, it is sometimes useful to gain an overall measure of the efficiency of use within a home from day to day.

Implementation

The Watt Watcher is implemented as a self-contained module with integrated power measurement and supply. As such, it can be used a "plug-and-play" device without any configuration. The Watt Watcher is designed as a set of modules that connect together to form the device itself. The circuit is fabricated as a single board that is then segmented and mounted together to form the compact package. The base module incorporates power supply and the rotational display, as well as measurement electronics. The MCU module accommodates the on-board processor and radio, while the display module has its own microcontroller dedicated to controlling the numeric display.

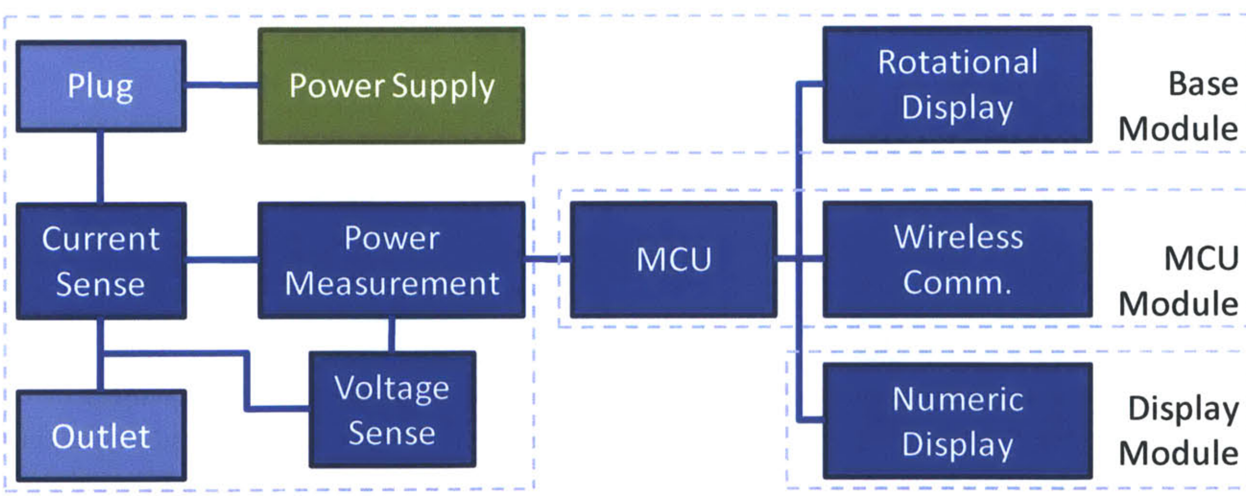

Figure 30. Watt Watcher system diagram

Due to the form factor constraints we chose, as well as the need to accommodate a plug as well as a receptacle in the device while not occluding nearby outlets, we designed a fully-online power supply for the Watt Watcher that allows the entire circuit to operate from mains voltages. The power supply is a non-isolated half-wave rectifier operating on mains voltage. Drawing inspiration from a design by Dave Johnson, the supply keeps the supply ground and device ground at the same potential as input ground while providing 5V at 40mA for the current sensor chip. The grounding requirement is imposed by the current sensor chip, and results in the DC voltage being referenced to the "hot" side of the input AC power. As a result, almost all portions of the circuit is at or very near mains voltage, since the 5VDC seen by the device internals is in fact referenced to a sine wave with the same amplitude as the mains voltage. This led us to design for a fully-enclosed design with no buttons or external controls that can become a point of failure and expose the user the mains voltage. While optical and capacitive input remained options for the design, we felt that the radio-based control discussed previously provided better flexibility while simplifying to device significantly from an on-device UI perspective.

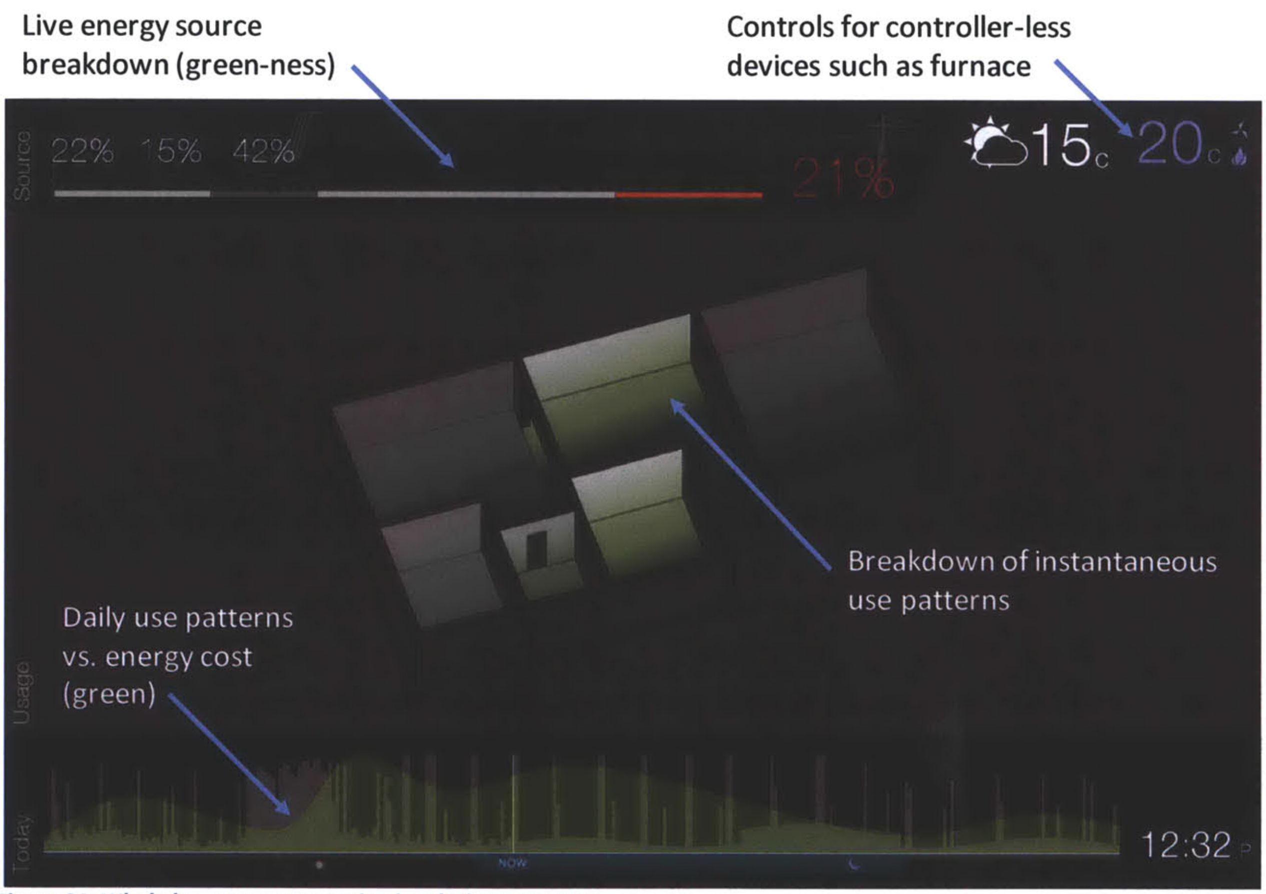

Figure 31. Whole house energy monitoring design prototype

The measurement system is calibrated after construction, and depending on the exactness of the calibration device can achieve approximately 1% error rate. The prototypes are calibrated with a much more crude system and have an error rate better than 10%. While this could be made better with more accurate and expensive parts, we felt that for the task of giving the user a sense of relative energy use, better error rates were not the highest priority. However, we did budget for sufficient accuracy so that "vampire power usage" – the power used when a device is turned off but waiting for input, much like a TV on standby – is still measurable and accounted for in the total use pattern.

Application in Generating Systemic Understanding

The Watt Watcher platform opens up the possibility for granular analysis of energy usage foci. While the on-wall module provides insight into the behavior of a particular device, it doesn't have the necessary interface to give the user an understanding of energy use patterns throughout the home. This is important because energy use is highly time and activity-correlated. By leveraging the connectivity component, we can also provide much more detailed analysis.

In Figure 31 we present a design prototype for such a whole-house monitoring solution. By aggregating energy demands from a number of Watt Watchers and combining it simple "whole-house" monitoring data as well as basic information from the utility company, it is possible to give the user an understanding of when and how they use energy. Additionally, the granular energy use of a room can be seen in the context of other uses, and reasoned about. For example, an increased energy use in the evening in the kitchen area can be understood as "I use more energy while I'm cooking" rather than an abstract value. If the user deems the use to be a significant contributor to their energy profile (easily seen from the whole-day graph), the user is now equipped to attack specific activities for reduction in energy use.

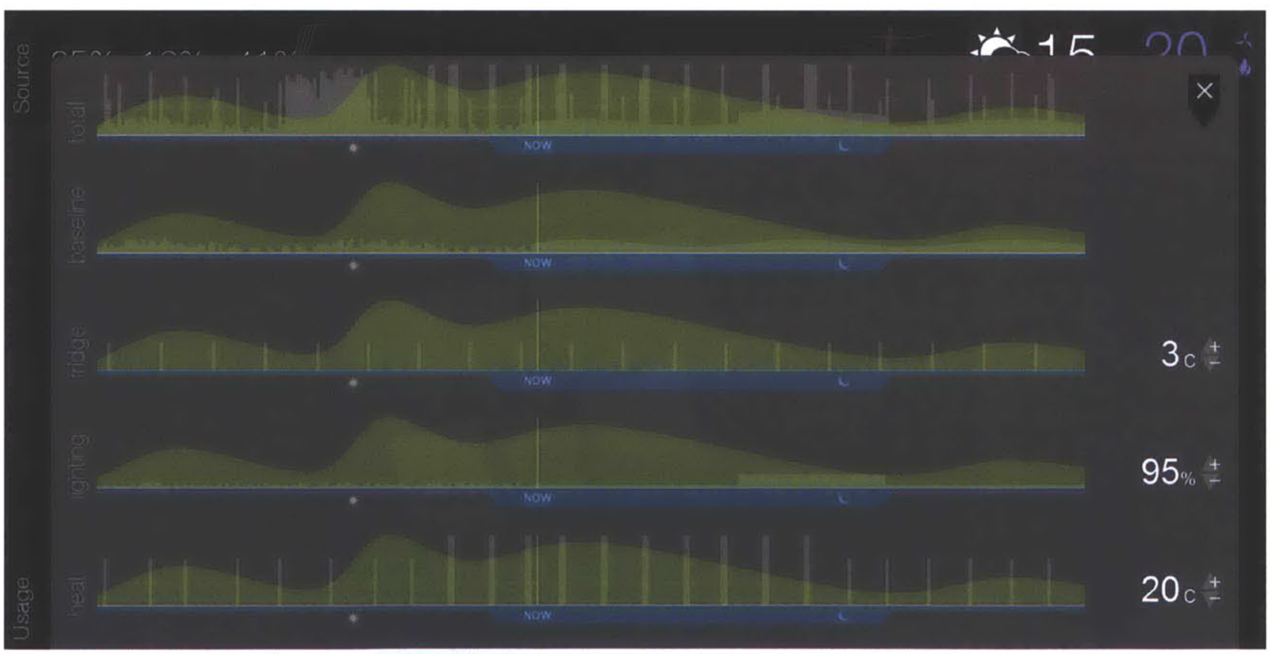

Figure 32. Granular view of use on a per-room basis

Additional Design Studies

While in the above sections we have described systems which we have implemented to a large extent to probe various aspects of reflectOn design, the overall design strategies used above apply to a variety of other problem spaces. In this section, we outline the application of the design principles underlying all reflectOn designs to a number of other problem spaces. While we have not completed the physical prototypes for these designs to the same extent as the solutions, we hope to show the generality of the approach in the following sections.

Health Monitoring: Blood Glucose

While acute symptoms caused by diseases such as sore throats can be easily treated using the current model of doctor-patient interactions where the doctor interacts with the patient at specific points for a limited time, chronic diseases such as diabetes require long-running interventions where the patient's role in the monitoring and treatment is significant. Under such conditions, the patient must understand and reason about long term cause-effect conditions that are affected by micro-choices much as finance or use.

A prime example of this is diabetes – a chronic progressive disease with few initial externally visible effects that can causes significant systemic damage in the long term if left unmanaged (Demand Media 2011b). According to the Center for Disease Control (CDC), approximately 24 million individuals in the United States alone are affected (Neale 2011). Patients must regularly test their blood glucose levels, but such testing provides a very narrow window into a multi-factorial issue, since instantaneous blood glucose readings depend on time of last meal, stress levels, time between meal and test, mean contents, and many other factors. We feel that a reflective approach to providing feedback by showing the patient trends in progress (or regression) of their own actions can possibly provide a better outcome than the paper workbook approaches taken today.

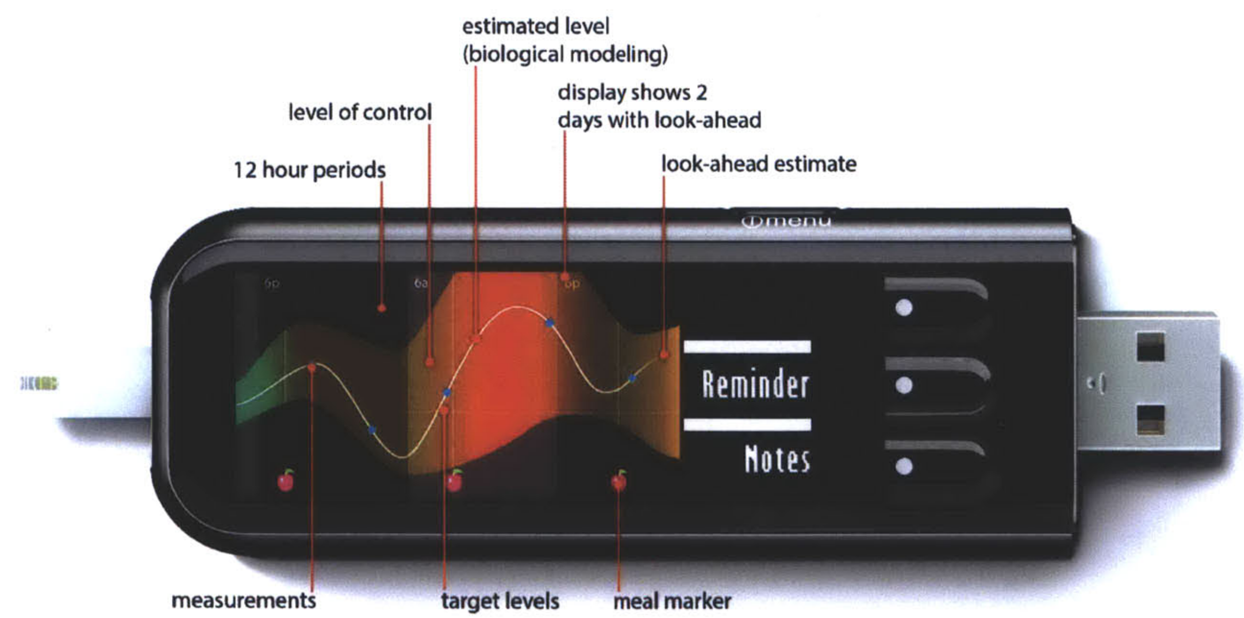

Figure 33. Sample blood glucose visualization. Meter shape based on Bayer Contour USB Blood Glucose Meter (Bayer HealthCare LLC n.d.).

While modern blood glucose meters give the illusion of accuracy with exact milligrams/deciliter readings, they provide no more understandability than the color-chart-based readings provided by older chemical tests. This claim can be easily verified by looking at online discussions, where the reading is effectively used without units, and questions about what a "good" or "bad" range is (and whom to trust about the answer) abound. The values are provided ostensibly for logging purposes, but one is hard-pressed to find a meter on the market that does not store many months of data. The accurate testing remains useful to the care provider properly equipped to understand what the values mean, but for daily reasoning about the values, a normative green-to-red scale is often more than sufficient. We feel that simplification of the results can provide significant gains in understandability in this arena.

A sample design, based on the body of an existing high-end blood glucose meter, shows a possible visualization that provides the user with actionable information regarding their glucose readings (Figure 33). In the sample design, numbers are discarded in favor of biological model-based estimation of continuous sugar levels, combined with a display of the level of "control" the user had over blood sugar at a given time. This provides the user with understanding of how the glucose level is affected by meals and time of day. Additionally, the readings from the latest test are used to estimate future control levels, which gives the user a chance to take ameliorative measures (such as limiting snacks, or conversely carrying extra food if hypoglycemia is possible) to correct the trend. By incorporating real-world events such as meals, it additionally encourages the user to keep the readings up to date.

Additionally, current meters provide no trend information to the user, leaving them to write the same data that exists in the system's memory down on paper, and then attempt to analyze the results themselves. This is not only tedious, but does nothing to compensate for factors such as time of test, days where testing was missed, etc. There are well-known and simple statistical techniques for compensating for such irregular datasets, but they are not implemented by meters, leaving a significant task that is trivial to a computer to a human being instead. Additionally, the most important information in steady-state blood glucose levels is the variance (level of control over blood sugar levels) as well as deviations from the norm (information about cause-effect relationships to actions taken by the user). This data is almost impossible for a person to calculate easily without using a spreadsheet, and yet trivial to handle programmatically. Additionally, a mean trend can be used to give the user an understanding of long-term control – a feature that can be useful in determining when to be concerned about testing results.

Lastly, blood glucose meters today do not take advantage of the fact that they have computational capability to act as the proxy of the care provider under crucial circumstances. For example, if the meter took the opportunity to ask when the user last ate, it would allow the system to compensate against known information about digestion and blood glucose changes after a meal. Much like the utensil attachment, such a reflectOn can enrich the avenues of attack against a problem condition, such as eating meals too far apart, testing too soon or late after a meal, etc. by collecting ancillary metadata about the central concern.

Social Wellness

Given the high level of mobility individuals have in the western world, it is easy for a person to become "disconnected" from a physical social network. While the initial presumption was that contemporaneously developed electronic communication methods would compensate for many of these issues, more recent commentary seems to indicate that electronic communication is at best an incomplete substitute (Kraut et al. 1998). Individuals often have hundreds of "friends" on social networking sites such as Facebook.com, but do not see any of them of a regular basis. Since the human mind treats digital "contact" somewhat similarly to face-to-face contact, but much less information can be exchanged digitally along some axes, it leads to a protracted sense of loneliness without any obvious (to the user) causes. In addition to consequence at a personal level, such isolation also has cognitive and psychological side-effects that affect society in aggregate.

Current technologies such as personal area networks and near-field communication give us the ability to know when a person is nearby (Laibowitz & J. A. Paradiso 2004). Additionally, it is possible to use statistical signatures extracted from voice data to similarly gain an understanding of proximity and communication. Such techniques have been used to extract the social connectivity graph (Eagle & Pentland 2006), understand information flow, and detect the type of conversation (Madan 2007). However, such techniques can also be used to give users an understanding of their own social interactivity level. One of the advantages of the mechanistic nature of reflectOns is the apparent objectivity of results presented by it, and can be used to motivate at-risk individuals to change their behavior socially. In doing so, there is no significant need to use complex metrics, since the reflectOn would in fact reflect a fact that the user already knows, but does not necessarily confront.

Figure 34. Blossoms prototype, by Sajid Sadi.

A number of projects have variously made efforts to make such connectivity more visible in an ambient way. One such project is the Blossoms prototype (Figure 34), which allows individuals to maintain a physical collection representative of their friends, and to transmit a more playful signal reflecting though or contact between remotely connected units. Paired flowers can be shared between friends, and can retain their connectivity via the internet while separated, allowing for this form of connectivity. However, this solution is still an overt display of digital communications, and lacks the fact-to-face components of contact that are also essential to social wellbeing. However, such a system, combined with the approaches mentioned previously, can create a physical display of real contact, thus bringing to light neglected in-person contacts. As with the other designs discussed here, the goal is not to force the user to make contact, but rather to elucidate an invisibly-made decision so that it can be reflected upon. The intent is that once the state of these relationships is visible to the user, it will become far easier for them to be reminded to act accordingly or perhaps to change the expectations they might have otherwise held of that relationship.

Wellness Badge

While pedometers have existed for some time now, advances in signal processing and the availability of low-cost accelerometers based on microelectromechanical systems (MEMS) principles has made them commodity. One of the interesting aspects of pedometer use is the tension between public and private information. Those who wear a pedometer often discuss this information socially, but the device itself is usually a pager-like module that invites a private use. Given the clear correlation between weight gain (or loss) and social support, visibility of the metric can easily promote the creation of an environment that is both conducive and supportive of greater levels of physical activity. This has been well-supported by the practical lessons gleaned Nike+iPod (Nike 2011) and Livestrong (Demand Media 2011a), where the support from online social networks promotes greater activity.

Figure 35. Sample social wellness badge based on the Fitbit device (Fitbit Inc. n.d.). Device uses a persistent ePaper-like display to constantly show an easily externally-legible reading of mobility levels. A monthly aveage is shown in reduced intensity.

A very simple reflectOn may merge the pedometer readout and daily goal into a visible badge in a way that becomes a physical and public display of accomplishment (Figure 35), allowing those nearby to more constantly provide encouragement, support, and to some extent competition. While this provides a daily incentive to the wearer to improve themselves in the eyes of their cohort, especially given the high social value of being fit, it also reminds the user more subtly that their activity level is in fact public information, and those close to them often know how active they are, regardless of their announcement of the information. On the other hand, the observer is reminded to provide social support to the wearer, thus reinforcing the circle of social motivation. A publicly visible on-person display is much like a marathon band in that it both broadcasts information and invites inquiry, thus overcoming the social obstacle to discussion of the topic. Lastly, it also "outsources" the need to remember that a person is attempting to be more active, and makes itself conspicuous in its absence: the observer doesn't need to remember to ask, and when the person is in danger of stopping their activities, the lack of the badge becomes itself a conspicuous reminder for friendly social intervention and support. While this design outline is necessarily simple, it highlights the opportunities available to the field to make significant contributions with relatively minimal effort.

Summary

In this chapter, we provide as examples a number of designs for reflectOns, applied to fields as diverse as eating behavior and medical wellness. We hope to show through these examples a trajectory or design approach that can be generalized to further fields, and can be used to present information to users in ways that allow them to better reach their self-set goals. In the following chapter, we study the effectiveness of some of these prototypes to better understand how the design principles apply in realistic usage scenarios. In latter chapters, we shall draw together the lines that connect these prototypes and design sketches into a set of guidelines that outline what is possible with the design of reflectOns, and mental prostheses in general.

Study and Discussion

This chapter discusses the studies we conducted to better understand the possibilities and limitations of the reflectOns, and to distill some design guidelines that can help inform further exploration of the concepts presented in this thesis. Both quantitative and qualitative data was collected to inform the discussion.

Evaluating Performance vs. Efficacy

When we consider evaluation of a system in a traditional context, we often default to discussing the *efficacy* of the system in quantitative terms. Efficacy of the system can only be realistically examined relative to some existing system with sufficient similarity. This becomes critical when studying a novel intervention or interaction, simply because the new system may be *different* in significant ways. A number of papers have acknowledged this difficulty (Liebermann 2003; Greenberg & Buxton 2008). In such cases, one option is to quantify the *performance* of the system along relevant axes. This allows us to numerically explore the strengths of the system, rather than forcing a comparison which may be ill-fitting. This is the approach we have taken below.

Furthermore, the impact of a system, especially one that operates on mental processes rather than physical manipulations, may not be captured well in a numeric way. Therefore it is important to consider the qualitative aspects of the experience to bolster the numeric measurements (Creswell 2009). In light of this, we explore the efficacy of reflectOns both qualitatively and quantitatively to gain a more comprehensive understanding of the impact of these systems.

A Note about the Figures

In the following sections, we use a uniform convention for figures. Given the number of items, it is efficient to note the conventions before diving into the details of each study. The quantitative data comes either from Likert-type questions, or from event logs that provide direct numeric measures. While there is considerable discussion regarding the particular use of such data in the applied fields (Vigderhous 1977; Jamieson 2004), it generally agreed that presentation of ordinal data using averages or other statistical measures may be highly misleading.

As such, we show straightforward numeric or percentage measures that expose the mode, but do not use averages or other aggregate measures, since it is not possible to assert that the distance between ordinal entries are consistent from one to another, or from one user to another. Namely, the difference between a 1 and 2 rating for helpfulness in not guaranteed to be the same as the difference between 5 and 6, so the numbers in this case are independent labels. The graphs of these data are meant to show the numerical power and overall clustering of the data. For these sorts of data, there are cases where individual results (eg, overall effect from a number of sources) are either cumulative or discrete. In cases where the data is cumulative, we use stacked area graphs to show both the contribution of each measure, as well as the overall trend. Figure 36 shows such a graph, where the total time spent online is the accumulation from 2 separate components (entertainment and productivity). Since the "time spent with a goal" and "total time" are not cumulative, they are shown as lines. Line graphs are also used when the data series are fully independent, as is the case for Figure 37 and Figure 47. The latter shows numeric data from events, but the presentation is inherently a histogram of events per day, and is thus presented to show trends over time.

We acknowledge that the data is traditionally presented as bar graphs, but find that the legibility suffers when multiple data series are present using bars. We thus present the data as line or area graphs for the sake of clarity. Where appropriate, we will present both the distribution of the data and ordinal or cumulative statistics side by side. Additionally, some of the data collected, though incidentally numeric in nature, shows only a qualitative distribution of user opinion or located on a scale which is too nonlinear to even tentatively offer traditional statistics upon. Such data has been accordingly marked.

Web Surfing Behavior Study

The study of the reflectOn designed to intervene in web surfing behavior is the largest of the studies presented in this thesis. The purpose of the study is to determine the effects of feedback on behavior that individuals are generally unaware of, and resultant behavioral changes from feedback. It is also necessarily to ascertain baseline behavior expressed by the user to ascertain changes in behavior. A second goal is to gain some insights into the efficacy, acceptability, and legibility of various feedback approaches.

As discussed previously, the research platform is a browser extension that can track domains (URLs are not specifically tracked due to privacy concerns) visited by the user. The extension tracks time between page transitions as well as domains visited. It can be configured with a black list and white list of sites that dictate behavior for various domains, as well as to provide feedback over a short time interval (10 minutes) or an entire 24 hour period. The extension can optionally show the results in the status bar of the browser as bar graphs and produce a notification when a threshold for activity is exceeded. It can also optionally be used with an augmented mouse that can show the user the data using dual embedded bar graphs and provide haptic feedback.

The study plan is designed to gain an understanding of baseline awareness, as well as an understanding of which form of feedback would be most effective for them, and then test feedback modes. Since the web pages we visit have great privacy implications, information about particular sites were not exported. Rather, data regarding the triggering of alerts was saved to our servers. The alerts are triggered when the measurement bars are fully emptied, which is indicative of times when the user is not reaching their goals.

Stage 1: Users were given a survey about their current behavior, and provided with the browser add-on. The add-on was used to monitor user behavior with no feedback.
N = 28, T = 1 week

Stage 2: Feedback in the browser itself is enabled, and users are allowed to enable the feedback and configure the add-on.
N = 28, T = 1 week

Stage 3: Users are provided with the augmented mice, and allowed to use the system for one more week. This stage examines the effect of pushing the feedback into the environment, and additionally examines the difference between feedback modes.
N = 24, T = 1 week.

User Demographics

The participants in the study were equally male (48%) and female (52%). 95.5% of the participants used the web every weekday. The distribution of reported hours spent for various online activity groups is shown in **Error! Reference source not found.**. It is interesting to note that given the user estimates of he amount of time spent doing various tasks (which may not be comprehensive for the total measurement), user underestimated the total by 32% (discounting <1 and >6 hours). Since one would expect that the total number of hours spent surfing include categories beyond those listed here, this is an estimate of the lower bound on misestimation of time spent.

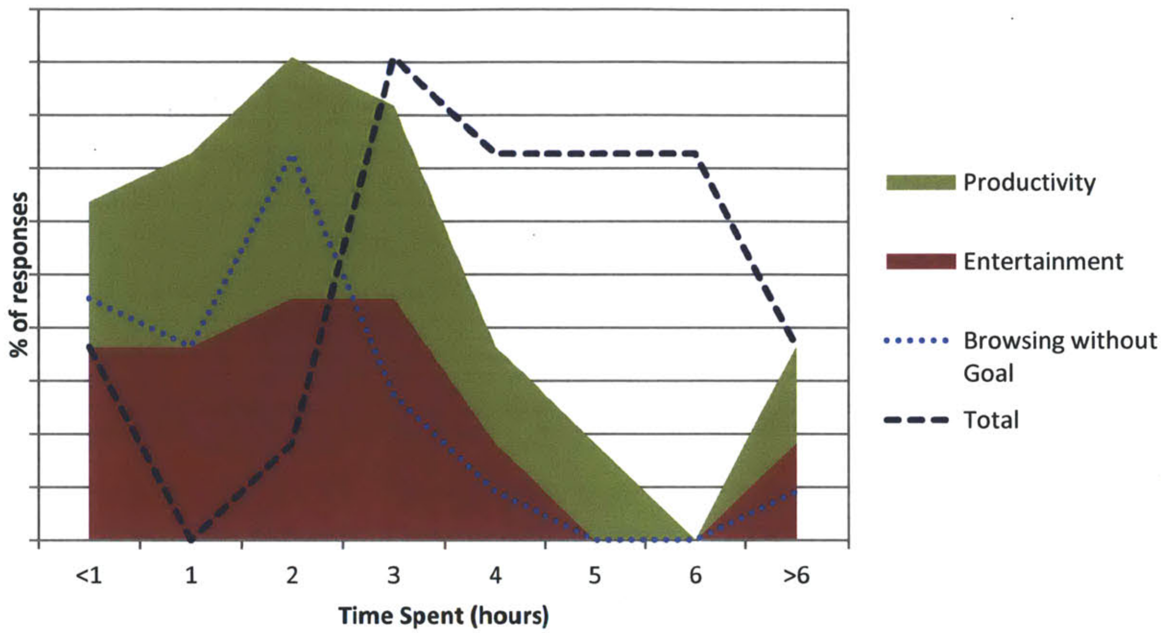

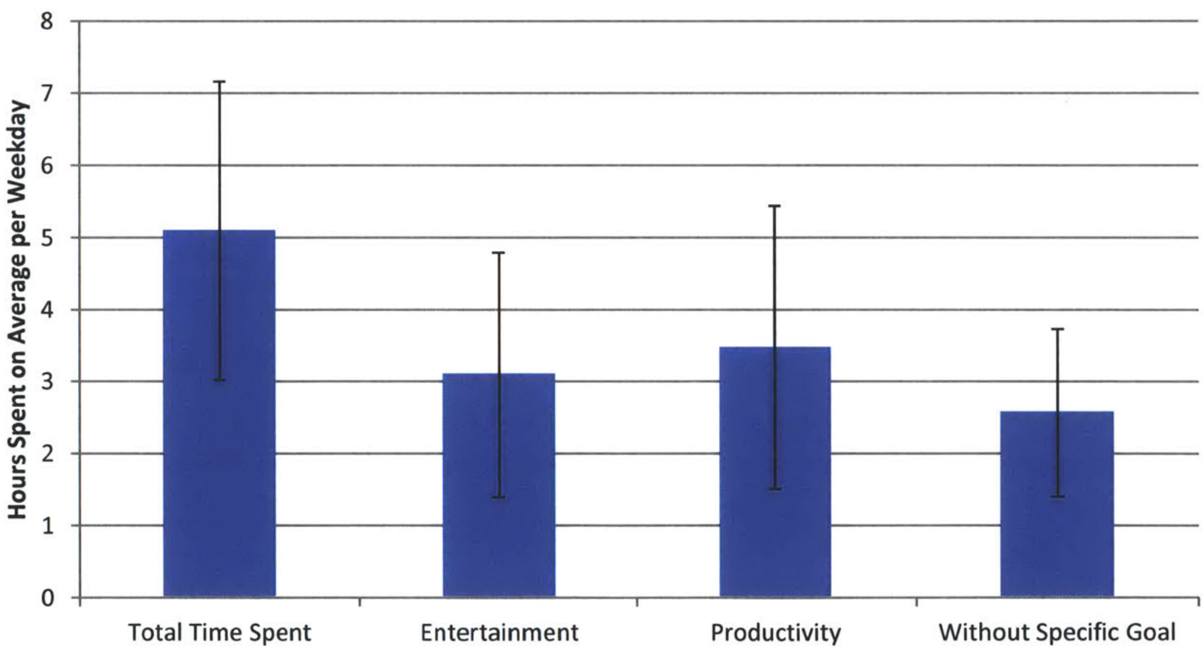

Figure 36. Reported time spent (hours). Note that given the information present, the Total Time Spent graph should be at centered at twice the value garnered from the study.

In Figure 37 we show the distribution of expectations of time spent. Lower values on the horizontal axis indicate that the user spent less time than she expected or wanted to on a particular category, while higher values indicate that they spent more time than expected. Perhaps expectedly, users expressed an overall propensity to rate their time expenditure as excessive in all areas except productivity, though it was interesting to note that users felt that tasks where they had a fixed goal also tended to take longer than expected.

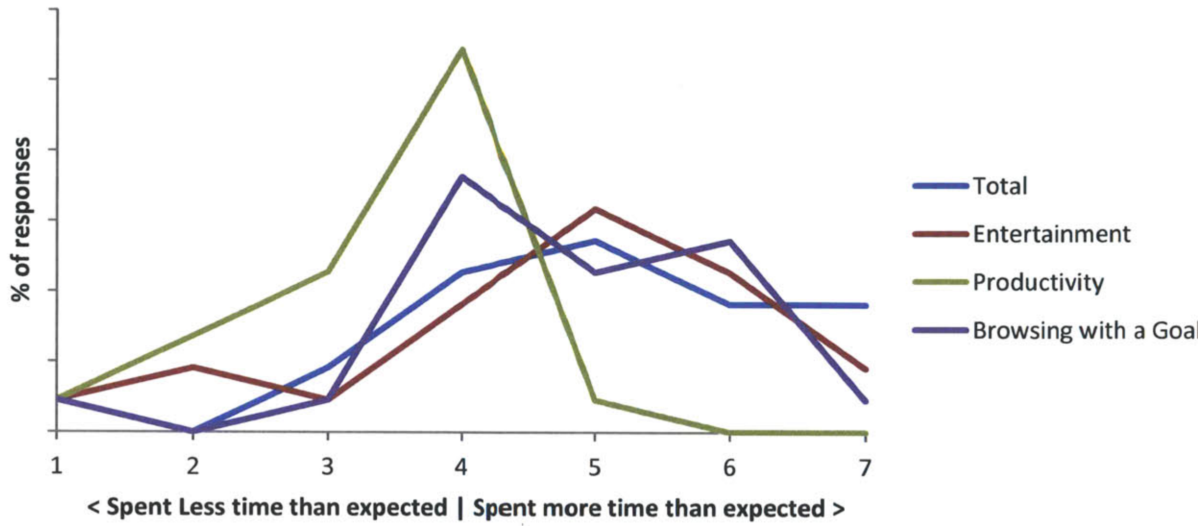

Figure 37. Distribution of expectations of time expenditure. Higher values indicate that the user spent more time on a particular type of activity than they expected to spend. Total: median 5, mode 5; Entertainment: median 5, mode 5; Productivity: median 4, mode 4; Surfing without goal: median 5, mode 4.

We also asked the participants about their time budgeting behavior. As shown in Figure 38, a majority of users did not budgeting time before starting surfing activity. However, the distribution of responses indicate that when the users did budget, the users were relatively equally likely to stick to (or not stick to) the budget. Additionally, other than 28% of users using mental notes as a time tracking tool, users did not in general use any tools to enforce their budgets. 60% of users claimed to have less awareness of the passage of time while surfing than what they would consider to be normal. Finally, 85% of users felt that their balance of time spent trended towards unfocused activities.

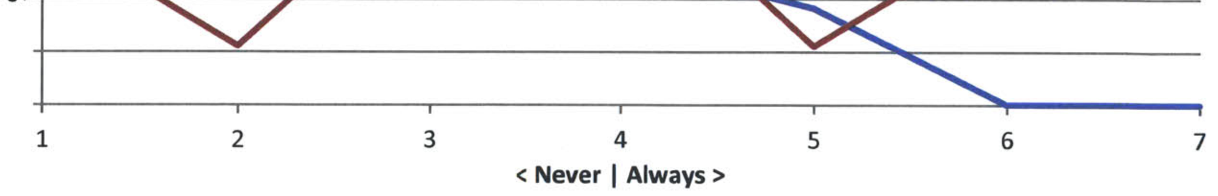

Figure 38. Distribution of budgeting behavior. Note that the percent of users who stuck with their set budget is expressed as percentage of users who budgeted in the first place. Median & mode = 2 for budgeting time, Median & mode = 4 for sticking with budget (excluding those who never budget).

After the first phase of the experiment, we provided users with totals of time spent on various sites. 87% of users found this data useful in configuring their blacklists and whitelists. It is interesting to note that contrary to expectations, users did not find the totals to be surprising (Figure 39), with most finding the values to be either in line with expectations or below expectations. This may be explained in two ways. On one hand, it is possible that the users experienced anxiety about their time expenditure, and were thus pleasantly surprised at the result. On the other, it is possible that the users justified the time expenditure a posteriori to match the values shown in the interface. There exists a propensity in these situations to re-imagine historic usage in favor of the desired outcome (Ariely 2008), and that may play a significant role in self-reports such as this.

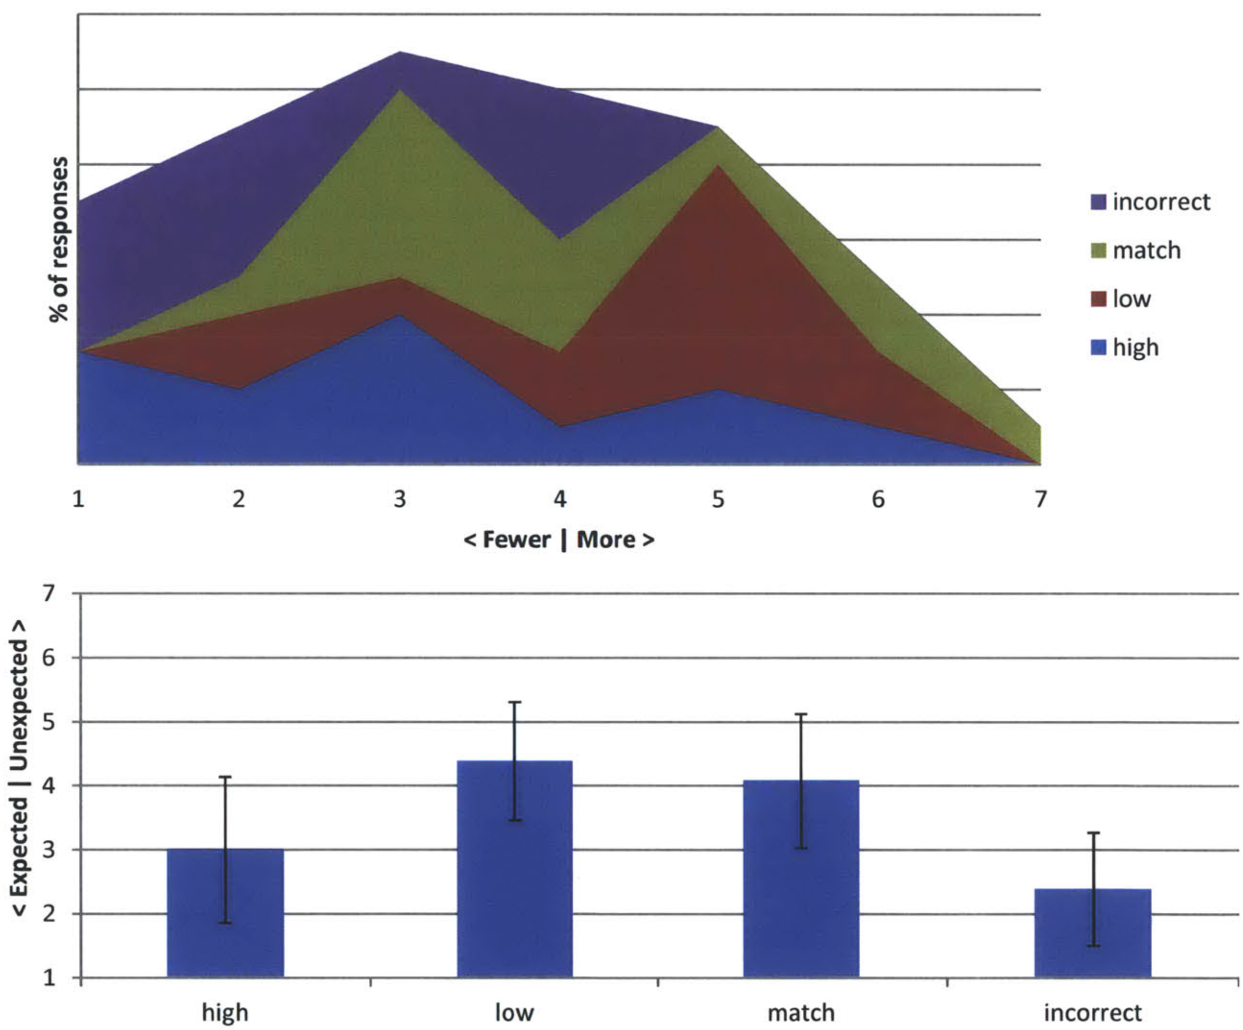

Figure 39. Extent to which actual usage matched user expectations of amount of time spent.

Results

The design of the study took into account several areas of inquiry, noted below.

- Time to settle into a steady state of usage.
- Aware of the feedback provided.
- Understandability and acceptability of feedback.

- User preference towards types of feedback (continuous vs action-oriented).
- Change in awareness over time.
- Self-awareness of behaviors.

With regards to settling time to steady state, over a two-week period, the number of changes to the settings dropped off significantly. As shown in Figure 40, the number of times the users reconfigured the system dropped significantly with length of use. There is a rise at the beginning of week two, since many users checked their configurations when they received the mouse. Since user data was anonym zed before transmission, the graph below infers day of week from the date of initiation, thus introducing an unavoidable source of error. However, the overall trend is nonetheless clear in broad strokes, and confirms that the configuration system provided a sufficient level of clarity to allow users to stabilize on their configuration quickly over time. This suggests that a properly designed configuration system can reach a steady state without having to preset the parameters of the system.

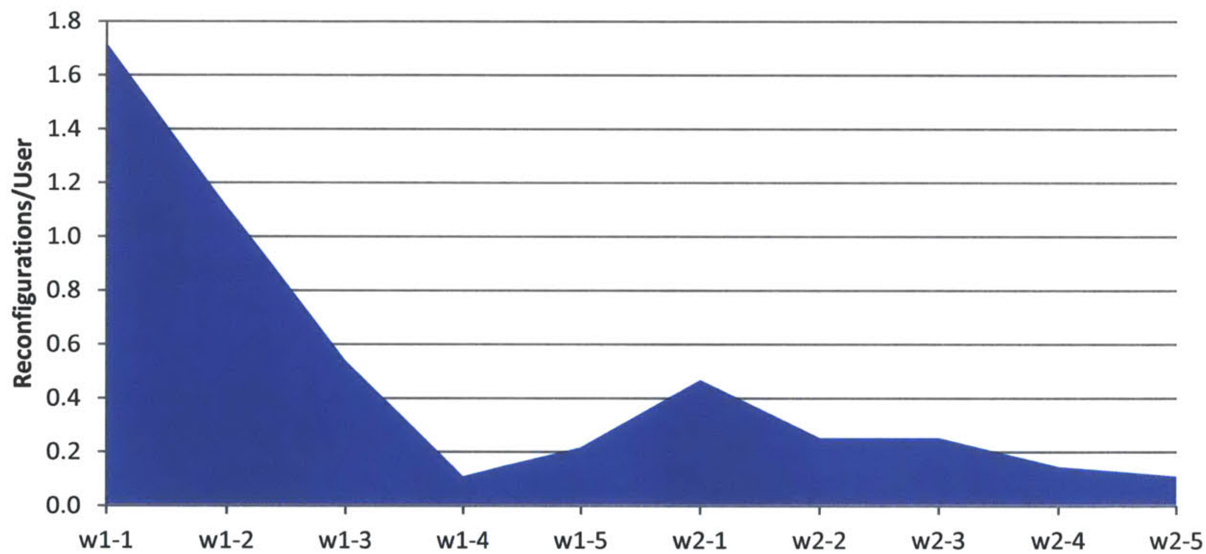

Figure 40. Distribution of configuration events over time. W1 denotes week one (visual feedback only), while w2 denotes week two (with mouse).

With respect to awareness of feedback, two approaches of approximating how visible the feedback is have been explored. We have used glance rate to get an approximation of how *available* the information provided is, and then added an obtrusiveness measure to gauge the extent to which the feedback is mentally processed by a user, on the theory that less noticed feedback is less likely provoke change in action. Figure 42 shows the results of these inquiries, and suggests a correlation with information availability. The mouse, being outside the field of view, is less available in this design, and thus less noticed. However, the isolated presentation of the mouse makes the information on it more obtrusive, thus increasing the likelihood that the use will not only see but also consider the information provided. The popup produced by the add-on is rated as highly obtrusive since it is designed to intentionally disrupt the user's flow. This may be a less desirable approach than the haptic feedback, which treads a middle course between in obtrusiveness while still having significant noticeability. This supports our design hypothesis that the feedback provided is noticed by the user, even when changes may be slow, or relative visible area of effect for feedback may be quite small (as in the case of the bars displayed by the add-on).

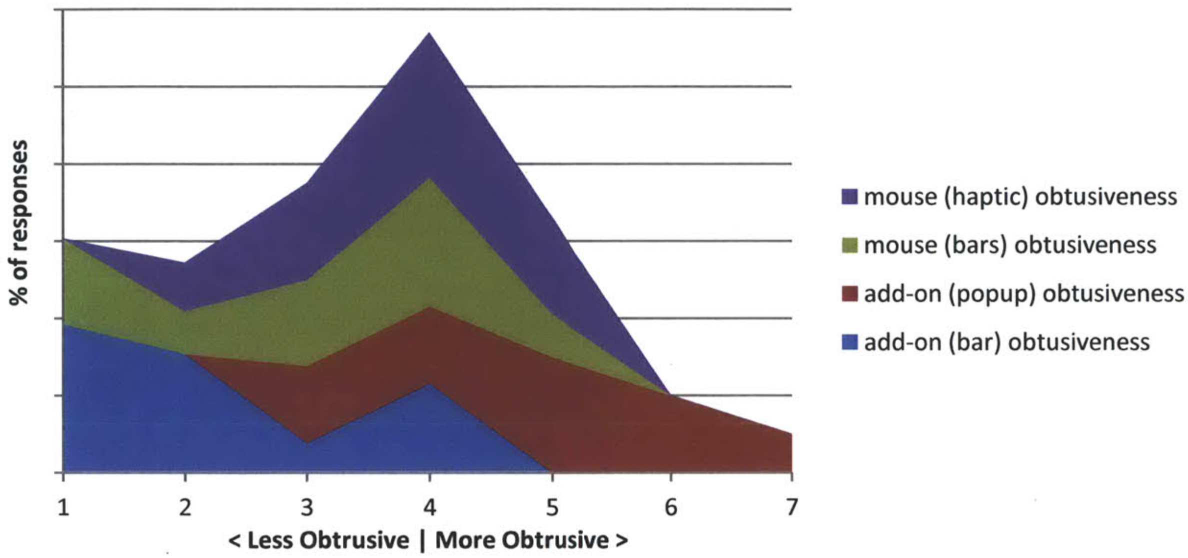

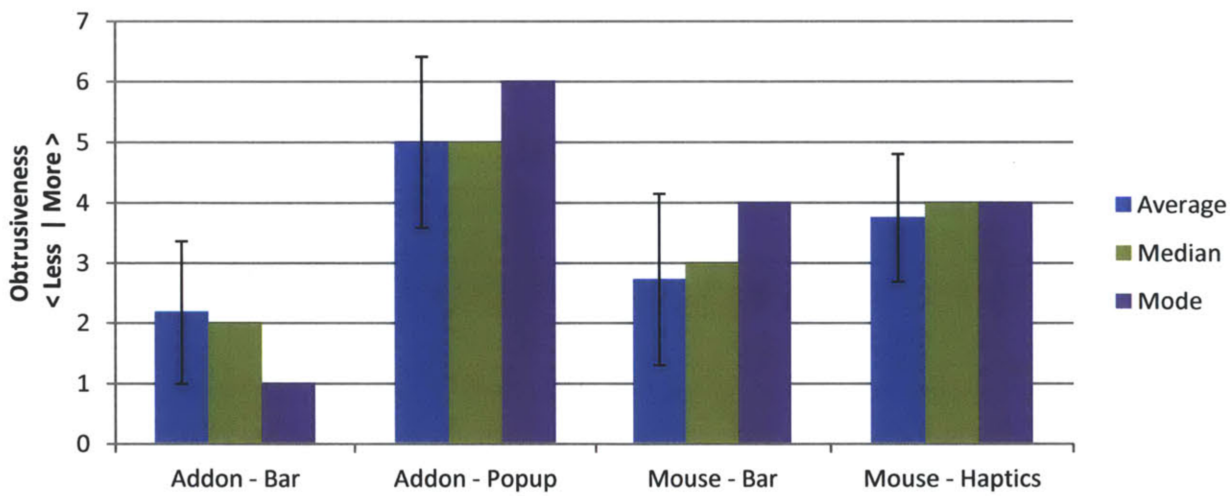

Figure 41. Obtrusiveness of feedback. Note that data is ordinal, and average and standard deviation should treated with caution.

In order to gauge the understandability and acceptability of feedback, we asked about the four possible avenues of feedback. As shown in Figure 43 and Figure 44, understandability mapped to helpfulness fairly well in all feedback types. It is interesting to note that when considered alongside glanceability, there is an inversion with regards to continuity of feedback. The mouse, which is less in view, is found to be more helpful when it provides actionable, infrequent feedback. The on-screen display from the add-on, on the other hand, is considered more helpful when it is providing continuous feedback via the time bars.

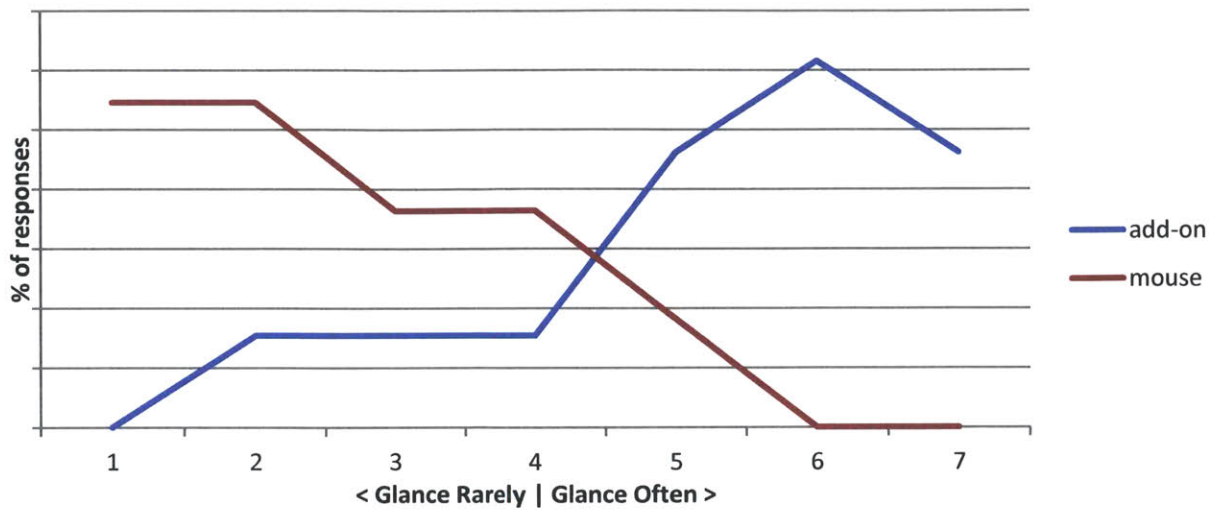

Figure 42. Glance rate and obtrusiveness measures for add-on and mouse. Add-on glance rate: average 5.3±1.5, median & mode = 6. Mouse glance rate: average 2.5±1.4, median = 2, mode = 1. Note that data is ordinal, and averages and standard deviations should be considered with caution.

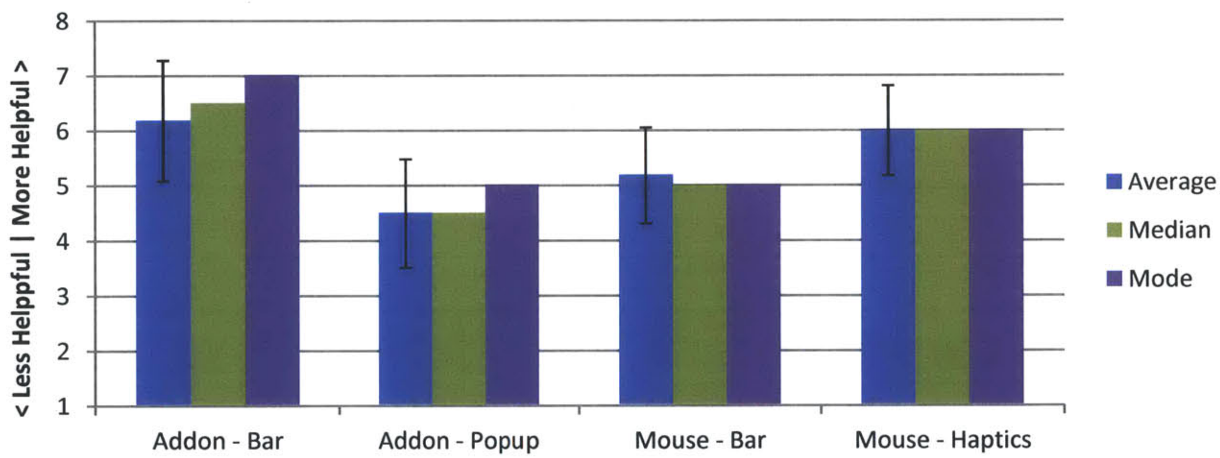

Figure 43 Helpfulness of feedback types towards change in awareness.

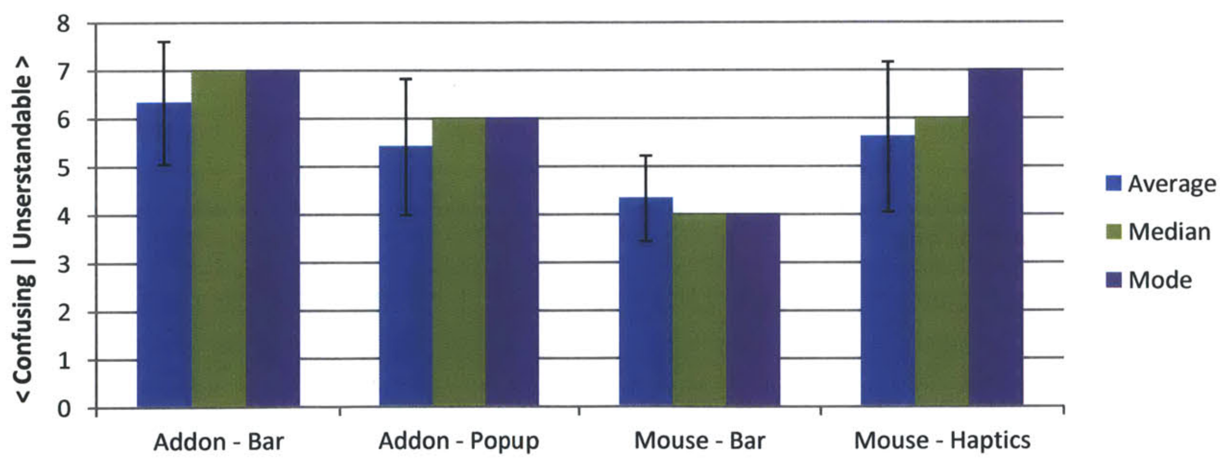

Figure 44 Understandability of feedback. Note that data is ordinal and average and standard deviation should be considered with care.

We examined the performance of the feedback both directly and indirectly, in the hopes of flagging any artifacts of self-report. On one hand, we looked directly at the temporal placement of the feedback within a browsing session. Figure 45 shows the relative positioning of the feedback provided in various forms. A lower value indicates that the feedback arrived later than the actionable moment, while higher values indicate that the user was apprehensive of the feedback. A design goal of the reflectOns is provide feedback that allows users to self-correct at the right moment. The importance of situated feedback becomes clear in short terms interactions in examining the results of the feedback: the bars provided by the add-on performed well in providing feedback that is relevant to the immediate browsing session. It is interesting to note that the users felt that the feedback with the higher informational content (namely the bars) were more effective in allowing them to self-correct than the more actionable feedback (eg, the popup message) when the results were situated on the screen within the direct focus of the users. However, the results were the opposite when the feedback was situated more peripherally on the mouse, wherein the users felt the haptic feedback provided better outcomes. Also of note, the popup and haptic feedback in this case was provided when the user had already exceeded their limits. It may be interesting to study various set points for the triggering of these actionable forms of feedback to deconvolve the performance of the feedback from the fact that the user may wish to self-correct with some margin remaining.

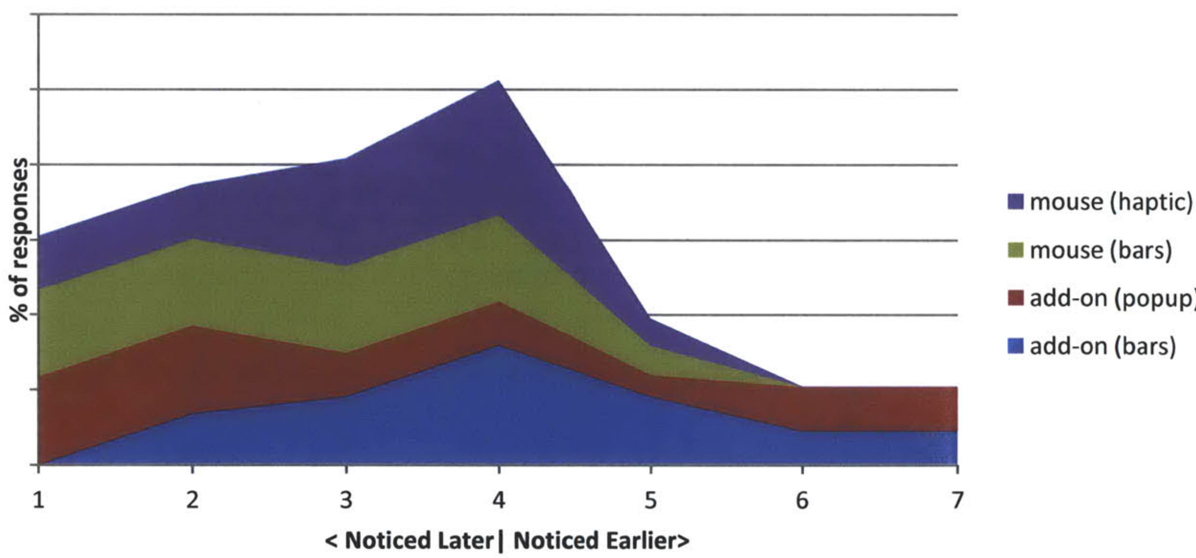

Figure 45 Temporal positioning of feedback within browsing session. The scale for this result is considered too subjective and nonlinear for traditional statistics.

In addition to the temporal performance of the feedback, we also queried users regarding the effect on self-awareness that they attributed to the various phases of the study. The results are summarized in Figure 46. Since self-awareness is a multi-factorial and cumulative effect, we have shown the results as stacked areas. In this graph, 1 represents no effect, while 7 represents an intense self-awareness. While we certainly would like to see effects from the study, it is also important to note that excessive awareness is often a conditioned response to negative stimuli, and can lead to long-term rejection of the goals of the feedback. While each stage contributed moderately, there is a clear trend towards a balanced level of effect from each stage. It is interesting to note that the initial commitment to the study had a more significant effect on self-awareness than did the presence of the mouse. As before, the highly situated on-screen bars, as well as the haptic response via the mouse, had higher perceived performance than did the popup message and the bars on the device, echoing the results discussed previously.

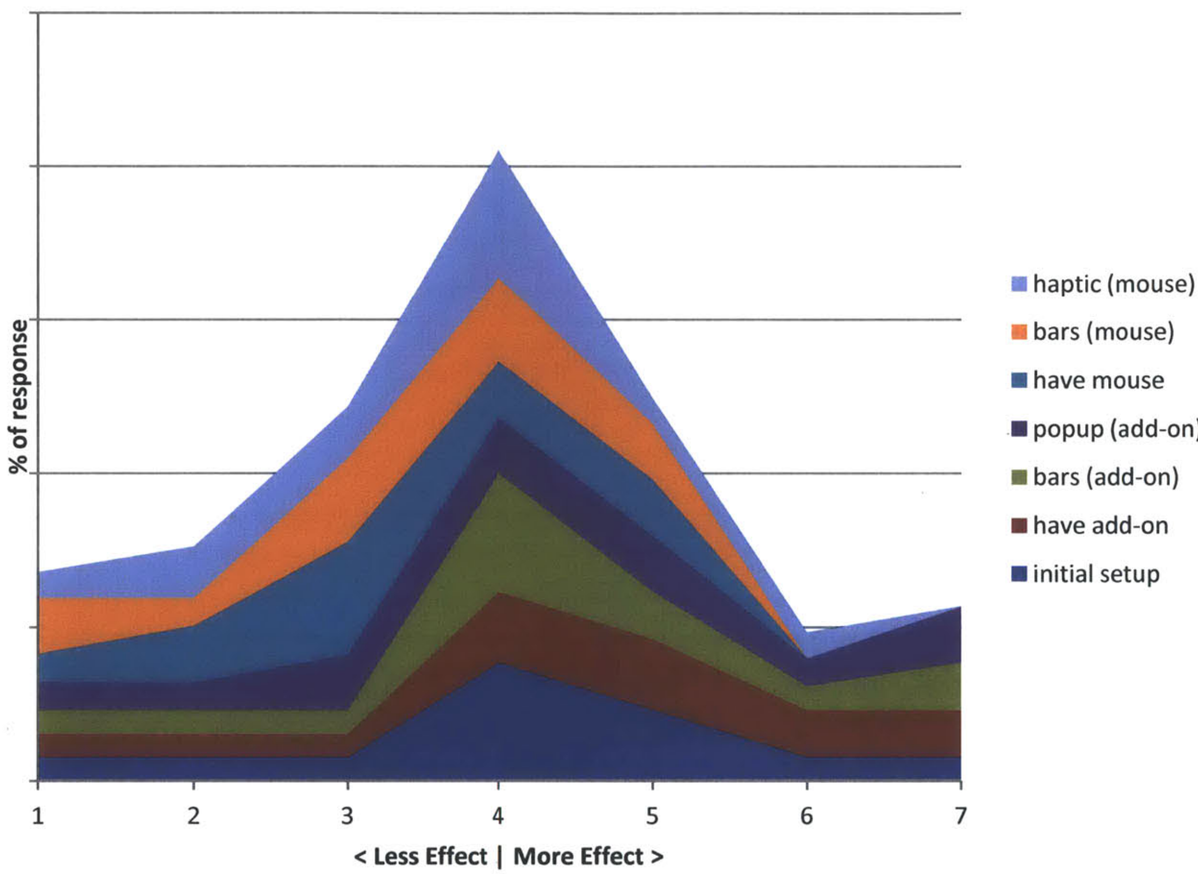

Figure 46 Distribution of effect on awareness due to the stages of the study. This data is considered too qualitative to be subjected to traditional statistics.

Lastly, we wished to identify any change in behavior over time from this study. Understandably, behavioral change or change in awareness involves many internal states that are not visible to external parties, and perhaps also invisible to the study participants themselves. As such, we chose to automatically record the frequency with which the displayed bars dropped below a threshold. The results are shown in Figure 47. Since there are far more events in the first few hours of use, we have shown the horizontal axis in $\log_2$ format. There is a clear and rapid downward slope, demonstrating that as the users spent time with the system, there was a clear learning effect that allowed them to have a better estimate of their time expenditure and thus avoid triggering the events. The dip in the graph at the approximately 24 hour mark is likely due to overlap in the sleep schedules of all subjects, though due to the variety of lifestyles represented this dip is minimal. This confirms the basic tenet of this thesis that subtle and glanceable feedback can have a significant effect on behavior and provide a means of changing long-term behavior.

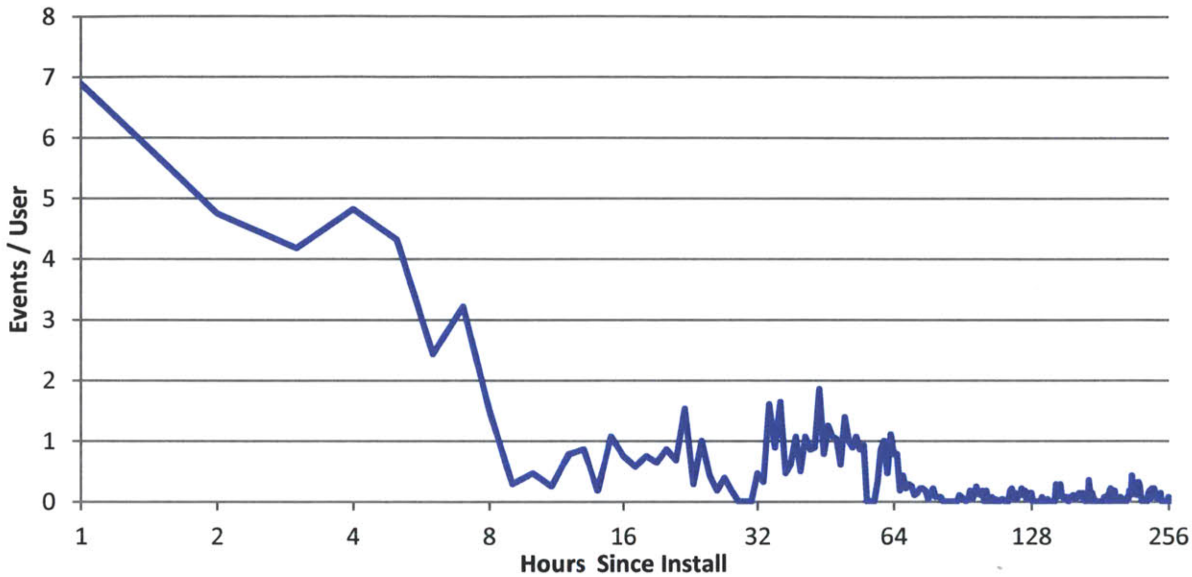

Figure 47 Number of limit events per hour (note logarithmic x-axis).

Eating Behavior Study

The study on eating behavior was conducted to address a shortcoming of the web based study: the web based study already has a focus (the screen) that is always in the attentional center when the user is engaged in the behavior of interest. In eating this focus is on the food rather than the utensil, which allows us to examine the effect of introducing additional hardware with a glanceable display to the tableau.

Since we are most interested in how the introduction of the hardware affects the user and differs from the reactions we observed in the web surfing study, we chose to offer the same feedback (2 colored LED bar displays along with haptic feedback) in the utensil as well. In order to accommodate the extreme variety of eating habits and the social correlates of eating behavior, the study involves a cohort using a special spoon only for breakfast. Breakfast is an ideal meal for this study since in many families it is an informal and individual event, and the participant is usually in a rush, thus promoting the sorts of rapid eating behavior we hope to forestall.

For the study, each user was given a modified utensil for use at home for breakfast for a week. Nine participants completed the study (N = 9). The participants were equally male and female (55% female, 45% male) with a range of professions. Due to the limited time and scope of study, a requirement for inclusion was that the person generally ate breakfast every day. Each person used the device for 5 days, and filled out an entry and exit questionnaire.

Demographics

The demographics of the cohort are summarized in the graphs below. The majority of the cohort spent 15 minute or less consuming their meal, felt at least a moderate time stress, and consumed a mix of prepared and homemade items during their meal. This anecdotally common scenario distracts the user and causes them to be more prone to overeating and unmindful eating.

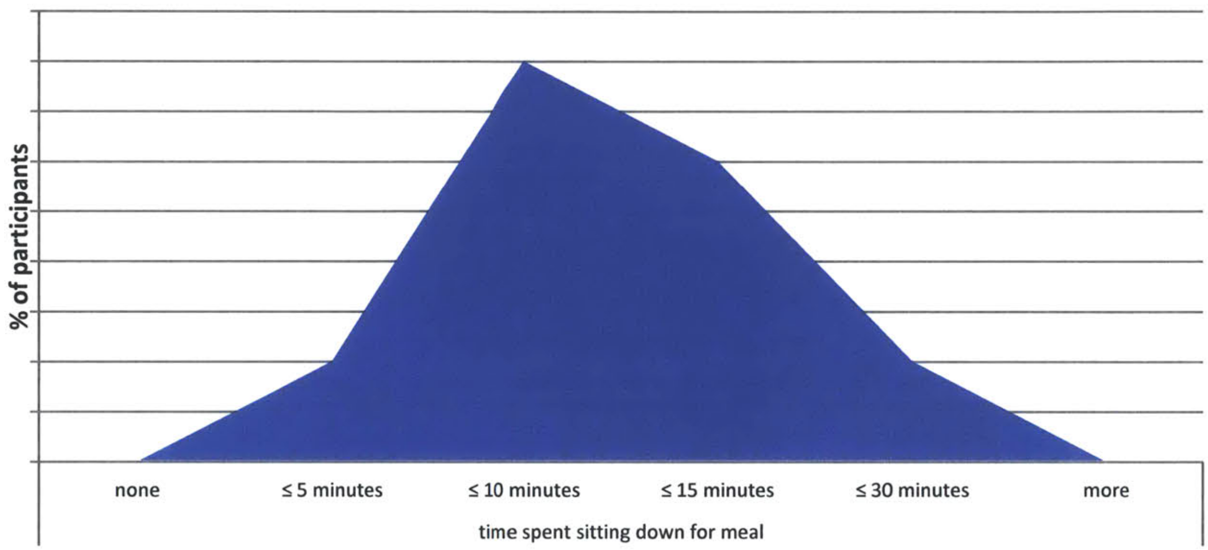

Figure 48. Time spent consuming meal. Average 10.2 minutes presuming users average to center of bin. Note that data is ordinal and statistics should be treated with caution.

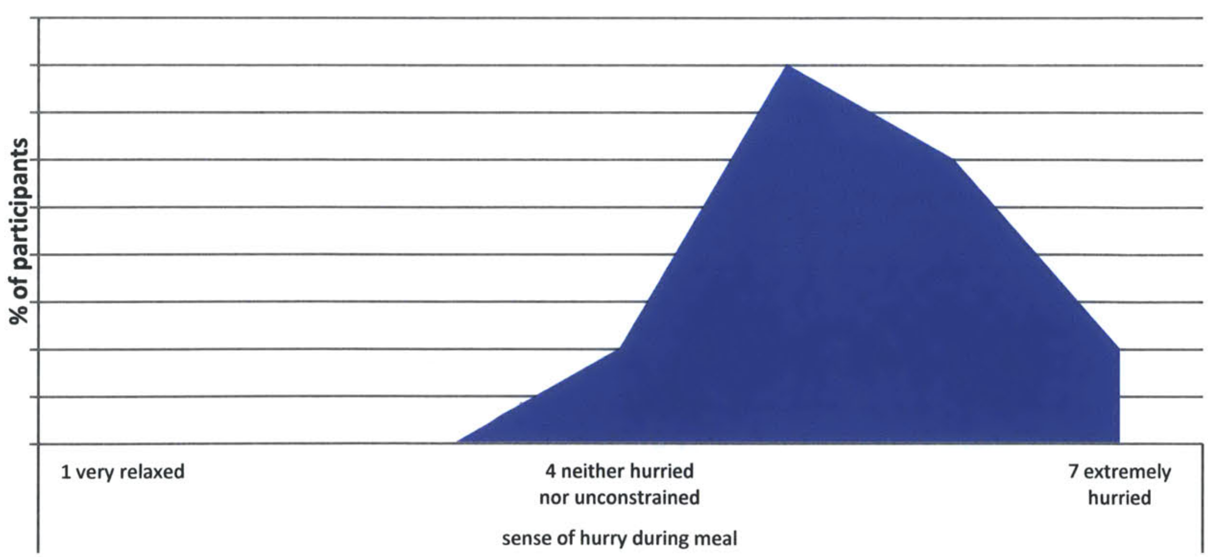

Figure 49. Sense of expediency felt by participants during breakfast. Average 5.4 ± 0.9. Note that data is ordinal. Treat statistics with caution.

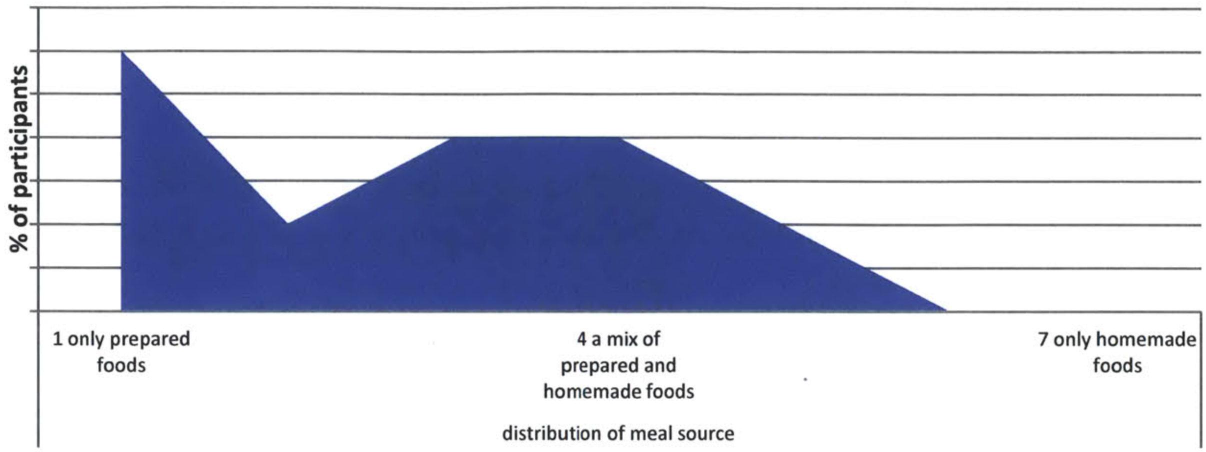

Figure 50. Mix of food sources consumed by participants. Average 2.6 ± 1.3. Note that data is ordinal. Treat statistics with caution.

Results

As noted previously, the spoon has three channels of feedback. It shows the number of bites taken since the beginning of the meal, as well as the proportion of bites that were taken faster than 5 seconds after the previous bite. Both of the bars are set to 30 bites. Whenever the user eats quickly, the spoon vibrates with a short pulse. This feedback (bite rate feedback below) is limited to once per minute. Once the count bar (end of meal below) depletes to zero, the mouse produces two longer buzzes.

We were interested in both the understandability and the obtrusiveness of the system. As shown in Figure 51, while all 3 channels scored acceptably for understandability, the bite rate alert, being the most immediately actionable, was considered somewhat more understandable. It is also the feedback that is the most closely tied to immediate actions, and therefore the easiest to analyze. However, the effectiveness of the bar graph was tempered by its relative obtrusiveness compared to the haptic channels (Figure 52). It is interesting to see that in a social and normatively charged activity like eating behavior, the visibility of feedback is in fact an issue. This is in comparison to the mouse, where greater visibility was a desired factor. The glance rate of the bar graph is also relatively high compared to the mouse, indicating both the efficacy of situated feedback in improving the effectiveness of the feedback.

Another method of looking at the performance of the feedback is to look at where the feedback was temporally positioned relative to the decision point. These results are shown in Figure 54. The end of meal feedback was generally considered to arrive too late to be helpful to the decision process. On one hand, this may be because the feedback does not "build up" in the same way as the bars, and thus the participant feels pressured to end a meal at a point where they were not yet mentally ready to do so. The bar graph itself has a wide spread of results, echoing the wide range of glance rate. With a low glance rate, the bars are similar to the end of meal feedback, which with high glance rate it provides a sense of anticipation. Finally, the bite rate again performed well due to the temporal immediacy of the feedback.

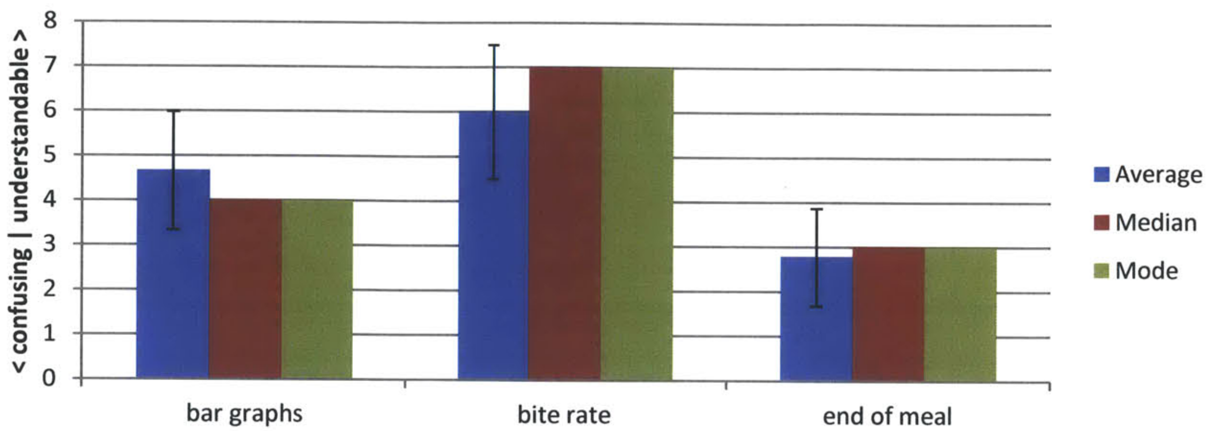

Figure 51 Understandability of feedback by feedback type for eating study. Highly understandable feedback requires less thinking on the part of the user to consider.

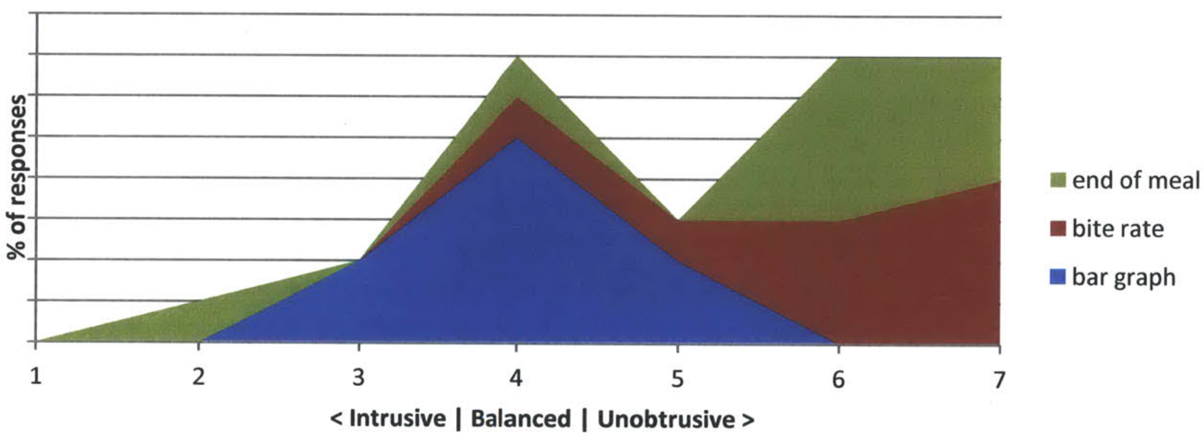

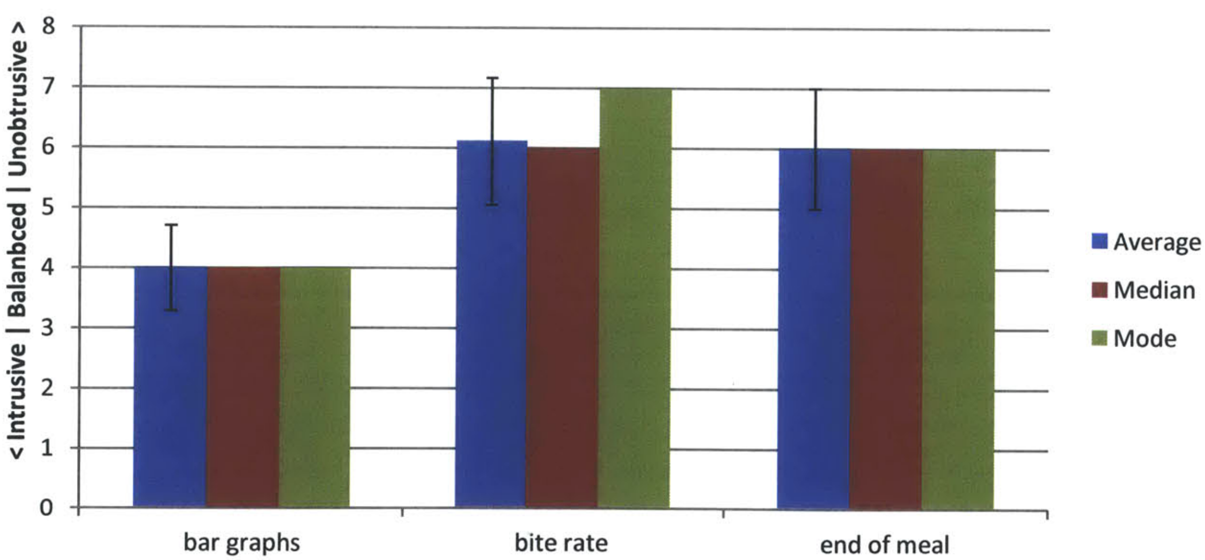

Figure 52 Obtrusiveness of feedback by feedback type for eating study. Less obtrusive feedback feels should be less disruptive to the normal flow of actions.

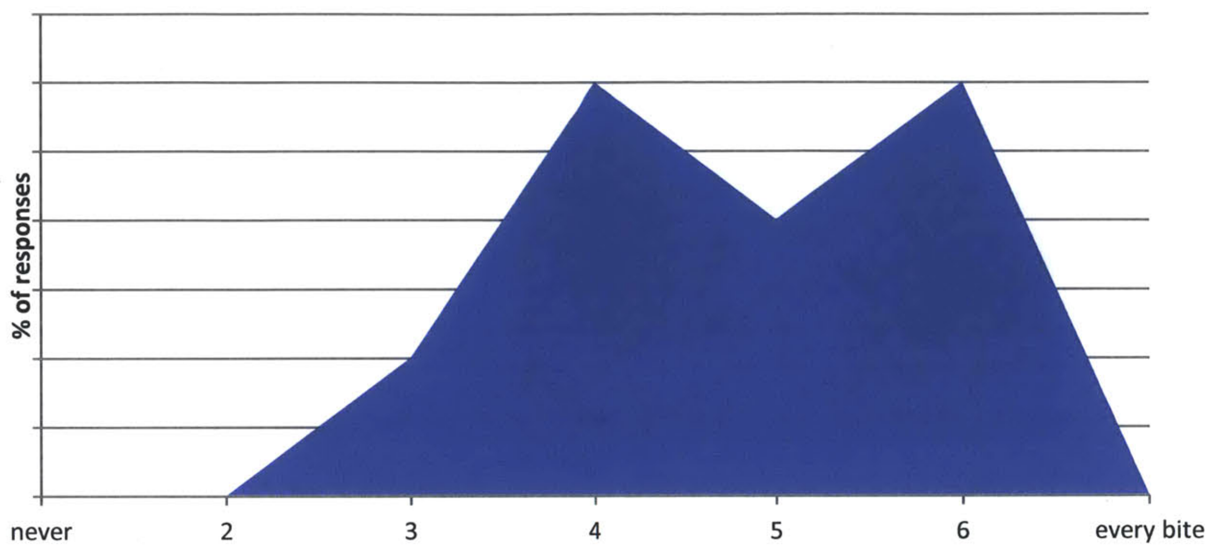

Figure 53 Spoon bar graph glance rate. Average 5.3 ± 1.1. Note that data is ordinal, and statistical measures should be considered with caution.

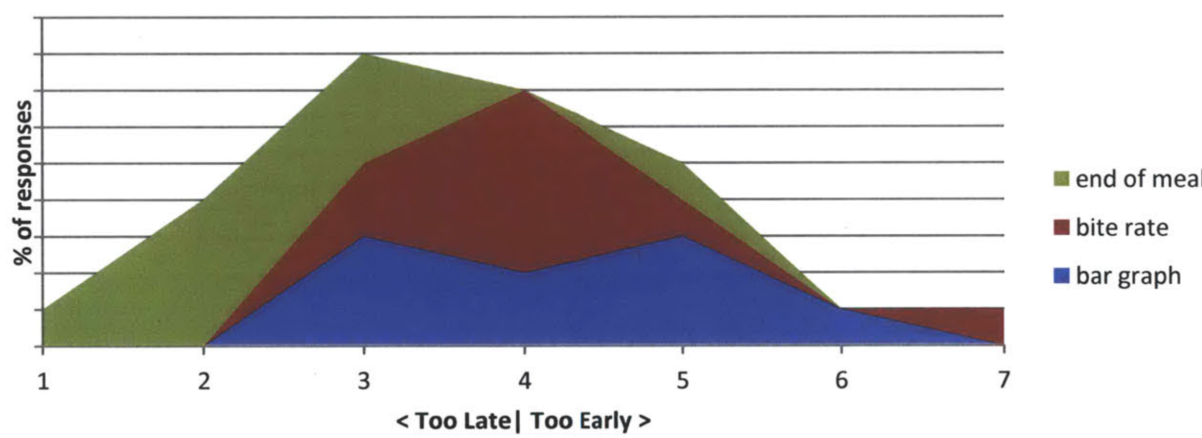

Figure 54 Temporal positioning of feedback relative to the decision point. Later feedback arrives too late to inform a good decision. The scale for this result is considered too subjective and nonlinear for traditional statistics.

Finally, we examine the effectiveness ratings of the various feedback channels and actions in creating awareness within the user. As before, the effect of simply using the spoon has a strong impact on awareness. The bar graph and bite rate alerts performed approximately equally as well, providing effective feedback on eating behavior control. However, the end of meal feedback continued to perform less well. It may be that allowing the user to choose a set point for this bar may be effective for those who felt the setting was simply incorrect. Additionally, a midpoint or 80% warning might have been better for those who simply felt that they were broadsided by the reminder to consider the end of a meal. It also highlights the difficulty of providing highly actionable feedback that borders on the normative. While the subjects were that the end of meal feedback was simply there to remind them to stop eating once they were full, it is not inconceivable that the feedback was consciously or unconsciously perceived as a command to stop eating.

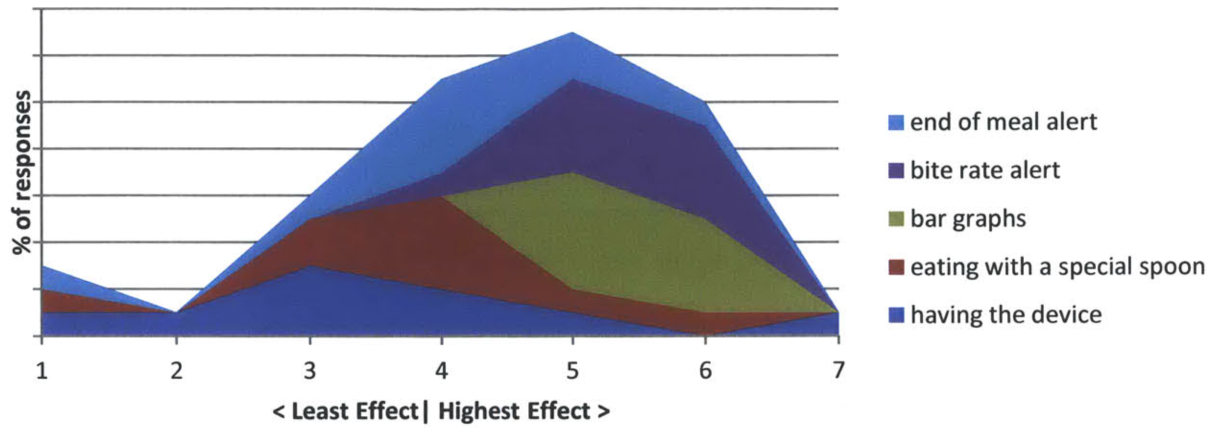

Figure 55 Distribution of effect on awareness by type of feedback. This data is considered too qualitative to be subjected to traditional statistics.

While this study does not have the numeric power or duration of the web surfing study, we did see a distinct downward trend in the aggregate number of alerts per day. This data is shown in Figure 56. The number of events in this study is far lower than the surfing study, since the user only engages in the activity of interest for a limited time each day. Nonetheless, there is a significant downward slope that suggests that the feedback was effective in changing behavior.

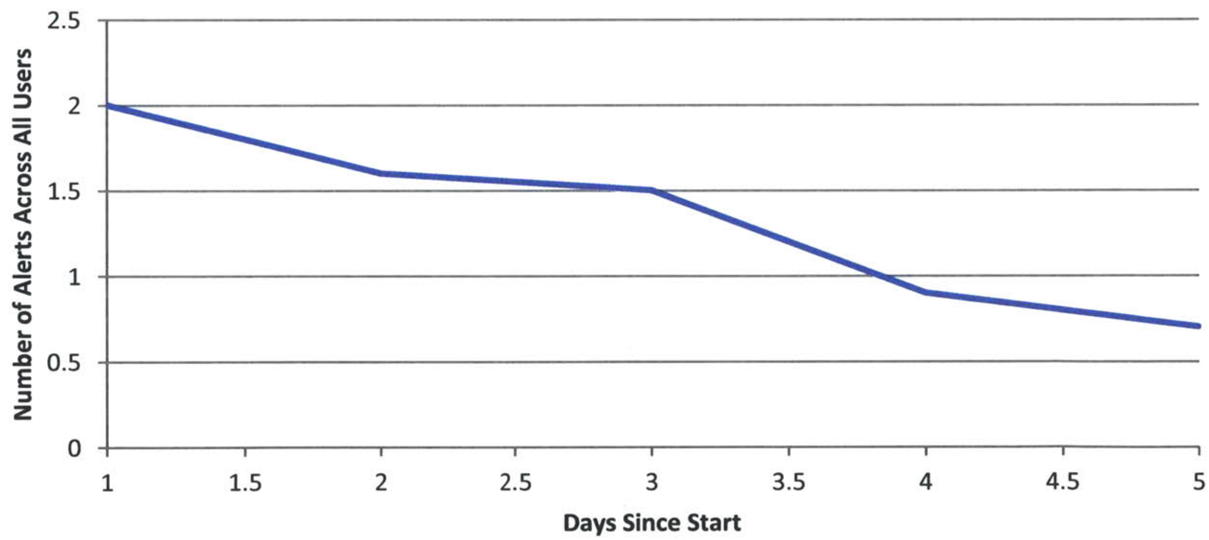

Figure 56 Aggregate number of bite alerts per day since start of study.

All in all, the study on eating behavior provides generalization to the approaches taken in the web surfing study. Both studies showed clear change in behavior, and provided counterpoints to each other by using the same feedback in different contexts. Particularly, the situated physical display used in the eating study underlines the need for contextual presentation in the design of these artifacts. As well, the importance of social considerations comes to light, as shown by the relative rejection of the visual display in preference to haptic output in this normatively charged scenario.

Energy Usage Study

The energy usage study was done in a qualitative fashion, attempting to delve in the usage of the WattWatcher prototypes in real world scenarios. Since only a limited number of these prototypes were created, we allowed users to use two of the devices for a day, and then participate in an interview to gain some insight into how the users deployed, used, and understood these devices.

As a short review, each WattWatcher module can be plugged into the wall, and in turn has one outlet available for which the energy use is metered. The WattWatcher shows instantaneous energy use using a ring of lights that emulate a rotating motion. Additionally, there is a 3-digit 7-segment display on the side of the device that allows users to see total energy use since the device was first plugged in. There is no direct means of user interaction with the device. While the devices can be controlled using a local peer-to-peer network, we did not test those aspects in the study.

Outcomes

All of the users were initially intrigued by the rotating motion-based displays. The preliminary commentary focused on the ease with which such a display can be read at a glance. However, one of the users noted that the displays because somewhat intrusive since they drew attention to themselves whenever the devices were on. The ability of the eye to perceive motion even at the periphery made it relatively difficult to ignore the devices. However, others noted that the output of the system was effective in drawing attention to the appliances drawing (or occasionally, slowly siphoning) energy. A design tradeoff exists here between users who felt that this was an opportunity for discovering energy use data, and those that used the devices as if they were fixtures in everyday life. One suggestion was to only show the rotational display when a button is pressed, though this counteracts the benefits of continuous comparison between energy uses.

Each user chose a different pairing of devices to compare. Interestingly, most of the users chose to select a pair of devices that they felt were worthy of closer inspection, while only one chose to move the devices around to sample energy use from various devices. Many of the users noted interest in having feedback over entire rooms rather than on a per-appliance basis. This is in fact a capability provided by the whole-house monitoring software which was not part of the study. The devices were generally deployed with a clear view, though in some cases the device had to be placed out of direct line of sight. Nonetheless, users who were forced to use the device out of sight were satisfied by the fact that they could "peek behind" the appliances to gain a quick look at energy use without having to move the appliance. The choice of a light-emitting display (as opposed to a traditional reflective LCD display) was a design win in this case, since it enhanced the glanceability of the display. All users were in consensus that the display was sufficiently unobtrusive in a normal kitchen environment, though some did express slight concern that the displays would glow while they slept. None of the users had a living situation that placed them in proximity to the kitchen while asleep (eg, in a studio apartment), so the concern was not further examined.

With regards to the actual use of the device, users felt that the rotating display was the most compelling output since it didn't require comparing numbers over time or plugging in both units at the same time to understand which device was using more energy. Additionally, it was felt that the users received a better "feel" for the energy use cadence of a device from the rotational display. For example, one user chose to compare his computer and TV set by plugging an extension cord into each. The user noted that this allowed for a much better understanding of where the power went. There was a surge of usage related to the laser printer powering up, or waking up the computer (which caused the display to turn on). This user was not aware of residual power use during standby, and was surprised to note the entertainment console's meter ticking away even when the devices were all on standby. This user felt that the energy use in that case was a more compelling case because it clearly showed energy being wasted. Others also reported similar reactions to residual use of energy, since in these cases energy was being used without providing a direct benefit to the user.

One interesting set of data came from the user who used the device to sample many different appliances and devices. This user felt that the results he received were closer to in every outlet, receptacle, or switch had the same sort of output. Though the user could not see the cadence of energy use over time, the user felt that the benefit of fully being able to compare all the devices helped in the process. In doing this sort of "energy survey" of the space, the user felt the total kilowatt-hour meter on the side of the device was helpful in giving accurate measures. It would have been desirable to have enough prototypes to fully instrument a space to better understand this use case.

All of the users felt that the visual indication of energy use provided a strong impetus to consider energy use when engaging in an activity. A majority also stated that such a device can serve as a reminder to, for example, turn off the TV when the device is only a view when leaving the area. The possibility of decoupled "door meters" that show information about a room as the user passes out of a space is an interesting area for future investigation. Additionally, this suggests that users need not have the device continuously displaying their information. Combined with the distracting nature of the animation when it is in the periphery continuously, this suggests that there may be an avenue of investigation where the displays are activated for a limited time on a fixed interval, showing a few cycles of animation before switching to standby. Lastly, combined with an activity sensor, it may make sense to show feedback in a space based on the occupancy or activity in a space such that it is timed for the beginning or end of occupancy.

Users felt that the watt-hour (Wh) display was less useful not only because of the lack of glance ability but also because there is no way to synchronize the devices so that they all measure on the same interval. One solution to this is to show the usage per 24 hour segment and automatically reset the on-device display to zero daily so that their totals remain in synchrony. Additionally, many noted that their usage pattern was more casual than intentional, and that the Wh display was better suited to an intentional usage pattern. However, users who applied the devices to appliances like refrigerators that draw energy in pulses felt that the instantaneous readings were somewhat misreading. It might be interesting to introduce an averaging metric that sets the rotational display to show a windowed estimate of energy usage over a time period that would be sensitive to immediate changes in pattern of use (eg, leaving the fridge door open) while still showing overall expenditure so that the user can use the device equivalently between continuous but pulsed uses such as refrigerators and air conditioners, and on-demand uses such as TVs and microwaves.

Overall, users reported a greater awareness of their usage, though the devices were considered better for gaining an understanding of immediate rather than long terms usage patterns. Users noted that they considered immediate actions more carefully more often than without the devices, but that the devices did not inform the long term usage patterns sufficiently. This was an expected outcome, since the devices are meant to be used with the whole-home monitoring solution that acts as the long-term counterpart to the immediacy of the in-place feedback provided by the WattWatcher devices. Users also rated the devices are well-balanced between usability and unobtrusiveness, and felt that a less-noticeable device also provided less effective and actionable feedback.

Takeaways

From the three studies outlined here, we have shown that the design principles for just-in-time information embodied in the reflectOns discussed above are effective in provoking behavioral changes in cases where the person is not highly aware or mindful of the actions they engage in. The reflectOns successfully increased awareness of target activities as reported by the users, and additionally produced a learning effect which shows clear trends of reduced dependence on these mental prostheses over time as behavior becomes better aligned with goals. Additionally, the studies have helped validate and inform several design guidelines can be extracted for use in the design of future reflectOns. We summarize the essential points of these guidelines below.

First and foremost, the studies underscore the importance of situated feedback. More than the form or the actionability of the feedback, the positioning of the feedback drove the sense of legibility and usefulness. This was underlined by the fact that the same feedback, placed differently in the first two studies, showed different levels of acceptance based on placement. Likewise, the user interviews conducted regarding energy usage also supported the importance of placing the feedback in view for maximal effect.

While we had initially theorized that the actionability of feedback would be the driver for the helpfulness of that feedback, this was not supported by the study. Rather, users preferred when they had plenty of foreknowledge of an approaching limit in order to have time to modify their feedback to suit the goals. Thus as an example, the value of feedback showing how time is passing during a browsing session is considered more useful by the user than the fact that the target amount of time has been spent. This forms an interesting counterpoint to examples used in the discussion of existing theories of repeated choice. While we tend to see actions being relatively independent in such descriptions, it appears that in micro-choice scenarios the flow of choices, rather than a particular end choice, must be changed for the user to perceive the highest performance from the feedback.

When the feedback drops out of the immediately observable realm, the mode of usage also appears to change. For example, the haptic feedback when the user has used their allocated time for surfing, or the display on the energy use prototype that is installed out of the line of sight, are used like alarms to redirect attention or bring forth moments of awareness. Despite the relatively lower level of success encountered by such feedback, we feel that non-just-in-time feedback can serve as an important secondary component of reflectOn designs by reinforcing the feedback provided just-in-time.

It should also be noted that the simple presence of tools to enhance awareness of particular actions themselves serves to enhance awareness. While it is easy to dismiss this simply as a placebo or novelty effect, it is important to note that behavior is a multifactorial outcome, and conditions that affect behavior should not be dismissed solely based on their passive nature. Indeed, by making the device itself an object of desire, it is conceivable that both compliance and overall behavioral change would be significantly enhanced. Indeed, it was noted by several subjects that the "finished" look of the mouse and the WattWatcher prototypes made the feedback seem more important.

Lastly, it is important to note the role that social acceptance plays in generating and environment where the feedback of the device is well-accepted. For example, both the mouse and the spoon use visual bar graphs on the device to convey information. However, surfing for a long time is not considered taboo to the same extent as overeating is, and this had a clear effect on the acceptability and obtrusiveness of the device. Since one may expect that such devices will be used in the long term, it is necessary to take into consideration the social acceptability of the device. On the other hand, this can also serve as an excellent driver for compliance. Since "green technologies" and energy saving are in the vogue as of the writing of this thesis, users of the energy consumption prototypes felt good about talking about the prototypes and how they were used. That in turn caused the user to spend more time considering the issue of energy use, which can have an enhancing effect on the overall behavior.

Conclusion & Future Work

The investigation of mental prostheses was done with the understanding that this work would open a door, and that the outcomes of the work done for this thesis would confirm the validity of the basic tenets and provide guidelines and possible fertile areas of exploration. Indeed, the applications of mental prostheses are strongly tied to their context, domain and goals, and are thus by definition varied and manifold. In this section, we summarize the design principles of reflectOns – mental prostheses designed to trigger reasoning and self-reflection. We also provide a concise overview of possible approaches to intervention in human behavior.

Contributions

This dissertation offers a number of concrete contributions to the state of the art:

Articulate need and working definitions for mental prostheses & reflectOns: We show that existing literature suggests that there exist a number of caveats to human rationality, and that these caveats cause us to experience difficulty in achieving the goals we set forth for ourselves in a variety of applications. We define *mental prostheses* as technological interventions that use just-in-time feedback based upon machine-observable behavior to help align behavior with personal goals, and provide guidance on the creation of these interventions. Additionally we define *reflectOns* as mental prostheses that operate by cueing the rational cognitive process to detect errors in intuitive cognition.

Created the first instances of reflectOns and demonstrated efficacy of approach: We offer four working prototypes of reflectOns, as well as a number of design studies outlining applications to additional domains. We have conducted user studies with three of these prototype systems, and have shown that the interventions thus created are effective in increasing mindfulness or creating new behavior within their respective domains. We additionally show that these interventions are acceptable and helpful to the users of the systems.

Distillation of prior work and mapping of areas of exploration: We provide an overview of the critical aspects of behavioral theory and the state of the art in our understanding of how intent and behavior diverge, as well as technical approaches to providing feedback. Pulling together these domains, we provide guidance on analyzing and designing for situations where a reflectOn may be effective, as well reflections on areas where other approaches may be necessary.

Guidelines for the application of this approach to additional areas: While this dissertation focuses on the particular subset of mental prostheses that operate on the reflective monitoring of the intuitive system, we additionally provide some possible areas where direct manipulation of the intuitive system may be effective, and some reflection on how such an attempt may be made. Additionally, we provide sufficient background in the thesis to allow future investigators to use it as a guide in attempting such interventions.

Future Work

As noted previously, this dissertation opens a door into an array of possible interventions that allow a person to change their own behavior to better align with their intentions. While this work verifies the basic tenets and provides general guidance, much work remains to be done in fully understanding the capabilities and limitations of this approach.

First and foremost, it is desirable to prototype additional designs and conduct studies to more deeply understand the limits of the approaches posited in this dissertation. While the results thus far suggest that the effects of these systems are not based purely on novelty, a lengthier study would help confirm this suggestion. We also focused significantly on preventing undesired behavior for the studies in this work. The theory underlying mental prostheses clearly extend to positive as well as negative feedback, and a combination of both is certainly possible as well. It would be instructive to carry out studies with only positive reinforcement, as well as congratulatory positive feedback at the end of a sequence of negative feedback in order to tease out the effective differences between the various feedback approaches. Finally, we touched on the idea of distributed feedback in carrying out the study on energy use. However, quantitative studies with more prototype units per household would be instructive in learning more about distributed passive feedback.

Additionally, many of the wider philosophical questions surrounding the design and use of mental prostheses beg a deeper discussion than is the scope of this dissertation. Being an initial exploration, we have not fully answered the question of where reflectOns — and more generally mental prostheses — are most effective. We have discussed some of the general requirements for reflectOns the subsection titled "What Are ReflectOns Good For?" of chapter 3. However, there remains an open question around what else could fall under the purview of mental prostheses. Certainly, deliberated and one-off choices are difficult to design for given the constraints we have placed on the definition of mental prostheses, since the effect is learned rather than conditioned. Additionally, it is unclear how mental prostheses can be extended to more creatively inclined activities. Certainly it is within the realm of possibility to create a system which looks for something akin to uniqueness in ideas, or assists in creative activities by providing a monologue of tangential pseudo-misinterpretations of ideas. However, such interventions most likely exceed the boundaries of mental prostheses as defined here and enter into the realm of mental augmentation. Nonetheless, it would be greatly interesting to investigate these possibilities.

The ethical repercussions and responsibilities surrounding the design of mental prostheses is also an area that is of significant long-term concern. On one hand, we attempt to sidestep the issue that by focusing on personal goals. Unlike persuasive interfaces and marketing and culturally oriented approaches, where the ends tend to justify the means, mental prostheses are meant to be conversant with the user and operate with the permission and cooperation of the user. However, when these technologies are used by external parties and then forced upon the user, it remains unclear whether we can be absolved of responsibility in the situation. For example, Progressive Insurance now is piloting a device that connects to the engine management system to monitor driving habits, and lowers insurance rates for "safer" drivers (Progressive Insurance n.d.). Some of the units beep when the user is outside of the defined deceleration levels, and there is an additional penalty for driving late at night. Clearly the enticement of lower rates is a strong driver, but the same data can be used to create a reflectOn that helps users drive more safely. As mentioned above, such dual-use reflectOns have not yet been studied, and their effect, personal and sociological, positive and negative, remains unknown. In all, there remains much need for an extended discussion around the issue of maintaining control over technology that affects our behavior, and whether it happens with, without, or with grudging ascent of our personal intent. It may very well be that there is no inherent safeguard other than vigilance on the part of users that insulates them from the negative uses of this technology, much as with advertising, pricing pressures, placement, or policy.

Extending the discussion of moral issues, it is also possible that the use of mental prostheses itself leads to a form of dependence where the user forms a pathological relationship with the intervention in question. This may take two forms. For negative reinforcement, the user may ultimately become dependent on the input of the system to activate rational oversight, thus removing by disuse a capability that she may otherwise have enjoyed. Since mental prostheses are single-purpose in general, it is possible that the reduction is oversight may "spill over" into other related behavior that are not guarded by the prosthetic, thus leading to an overall worsening of behavior. While this is a theoretical possibility, it is possible that the generalized nature of rational oversight can prevent any wasting of the capability, but this hypothesis should be confirmed. Alternatively, it may be necessary to have an on-off schedule of feedback that explicitly guards against this possibility. On the other hand, positive feedback or gamification of feedback can lead to a positive feedback addition that makes the normal case without the feedback seem lackluster, or cause a meliorative condition where the user overindulges in a choice due to the effects of the positive feedback. While operating on the long term habit-formation time scale provides significant protection from this possibility, as again there is need for independent verification of such effects.

Lastly, it is necessary to note that mental prostheses, despite being designed to allow users to pursue their individual goals, are normative to some extent in that they have some measures of utility built into them. In the naïve analysis, if a reflectOn prevents the user from eating a candy bar, the reflectOn is making a utility judgment about the value of the candy bar, and the value of the health benefits, and thus establishing some equivalence between the two choices. One response to this concern is that reflectOns are generally designed for long-term goals, and that the metric being optimized is inherently long-term as a result. A temporary deviation is allowable as a result, and indeed may be quite necessary to prevent the strong negative reaction to the denial of a choice to which the user has a strong precommitment. However, this is a design principle that must be enforced by the designer of the system. As is often the case, in conditions such as the example above the most obvious path (prevent consumption of the candy bar) is not necessarily the correct design choice. However, this requirement needs to be better vetted and articulated as part of the design principles of mental prostheses.

There also remain significant questions surrounding both habituation and habit formation. Habituation is an effect by which a discernible cue becomes indiscernible over time. A longitudinal study focused on habit formation would indicate whether this is a significant concern. It is important to note that such effects did not become apparent after two weeks and rather constant feedback in the web surfing study. As such, it is quite possible that the design constraint of limiting of feedback to prevent tiring of the reflective system may be sufficient to prevent habituation to a large degree. However, this hypothesis needs to be rigorously verified.

While this dissertation does not directly deal with the systemic or social manipulations that are more common to persuasive interfaces, it would be very interesting to see if mental prostheses can be dual-use in some cases. That is to say, the device could be used to provide extrinsically motivated feedback to the user in order to motivate a personal goal – for example, by showing that neighbors are more energy efficient and thus create social pressure to perform at least as well. The same device could then use the observations it garners to provide feedback that allows the user to align with these socially motivated goals, thus allowing the user to gain the satisfaction of achieving their goals in an easier way. Such approaches may exceed persuasive interfaces in effectiveness by giving actionable goals in addition to social pressure to perform.

It is clear from the discussion of the prototypes that the underlying mechanisms of data analysis and feedback can be generalized or at least modularized between reflectOns, since the cognitive artifacts being addressed are actually the same across all reflectOns. This brings to mind the possibility of a modularized technical "toolkit" that allows designers to more easily experiment with creation of mental prostheses across a variety of areas without having to gain the expertise to create the dedicated hardware platforms. Such standardization of prototyping approaches would allow for a more rapid growth in the corpus of knowledge surrounding the idea of mental prostheses and would allow the vetting of more ideas more succinctly.

Conclusion

This dissertation presents *mental prosthetics*, a model for distributed, embodied, design-embedded, just-in-time interfaces that augment the human judgment process. Drawing information from the activity of the user around them, mental prostheses analyze behavioral patterns in a way orthogonal to human cognition. Mental prostheses attempt to align choices with personal goals by cueing the user with just-in-time information, and provide calm yet understandable feedback to draw the user's attention at the correct time to the information available to them.

This dissertation provides several prototypes and design explorations as a means of sampling the various approaches to data collection, synthesis, and feedback. Focusing on self-reflection, these sample designs form a subclass of mental prostheses that we term *reflectOns*. We show through the studies carried out in the course of this dissertation that these systems are effective in changing behavior to be better aligned with user goals. Lastly, this dissertation provides a set of design guidelines that assist in the creation of new mental prostheses. While we discuss a variety of scenarios in this work, it is only the beginning of the exploration. The design guidelines provide insight into both the critical aspects of the design of such systems, as well as possible input and feedback methodologies. These guidelines, together with the reflectOns themselves, provide a basis for future work in this area.

As living beings, it is in our nature to adjust to the inputs present around us. We feel that the world needs a set of tools that counters the simple flood of information that we have seen thus far, and helps us make better sense of the information around us. These tools will combine the increasing capabilities that machines have gained to observer us with the long history of our endeavor to understand our own minds to extend our rationality in new and exciting ways. As designer-technologists, we must creatively blend the insights of many disciplines to craft experiences that go beyond simply presenting information to becoming extensions of our minds, thus allowing us to change and grow in not in accord with the pressures of the exterior world, but with the will of our own needs and desires.

Bibliography

Ambient Devices, 2011a. Ambient Devices : Orb. Available at: http://www.ambientdevices.com/cat/orb/orborder.html [Accessed August 7, 2011].

Ambient Devices, 2011b. energyjoule. Available at: http://www.ambientdevices.com/products/energyjoule.html [Accessed August 8, 2011].

Ariely, D., 2009. *Predictably irrational: The hidden forces that shape our decisions*, HarperCollins.

Ariely, D. & Wertenbroch, K., 2002. Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13(3), p.219.

Atkinson, B., 2006. Captology: A critical review. *Persuasive Technology*, pp.171–182.

Bayer HealthCare LLC, Bayer's Contour USB Meter - home. Available at: http://www.bayercontourusb.us/ [Accessed September 17, 2012].

Beidel, D.C. & Turner, S.M., 1986. A critique of the theoretical bases of cognitive behavioral theories and therapy. *Clinical psychology review*, 6(2), pp.177–197.

Benbasat, A. & Paradiso, J., 2007. A framework for the automated generation of power-efficient classifiers for embedded sensor nodes. In *Proceedings of the 5th international conference on Embedded networked sensor systems*. ACM, pp. 219–232.

Benbasat, A.Y. & Paradiso, J.A., 2005. A compact modular wireless sensor platform. In *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*. IEEE, pp. 410–415.

Björkskog, C.A. et al., 2010. EnergyLife: pervasive energy awareness for households. In *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*. ACM, pp. 361–362.

Bonanni, L. et al., 2005. Exploring feedback and persuasive techniques at the sink. *interactions*, 12(4), pp.25–28.

Brand-Miller, D.J., Foster-Powell, K. & Mendosa, D., 2006. *The New Glucose Revolution What Makes My Blood Glucose Go Up . . . and Down?* 2nd ed., New York, NY: Da Capo Press.

Brehmer, B., 1980. In one word: Not from experience. *Acta Psychologica*, 45(1-3), pp.223–241.

Brownell, K.D., Stunkard, A.J. & Albaum, J.M., 1980. Evaluation and modification of exercise patterns in the natural environment. *The American Journal of Psychiatry; The American Journal of Psychiatry*.

Cautilli, J. et al., 2006. Resistance Is Not Futile: An Experimental Analogue of the Effects of Consultee" Resistance" on the Consultant. *International Journal of Behavioral Consultation and Therapy*, 2(3), p.13.

Chang, R.C., Stout, S. & Miller, R.R., 2004. Comparing excitatory backward and forward conditioning. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*.

Chen, G. & Kotz, D., 2000. *A Survey of Context-Aware Mobile Computing Research*, Technical Report TR2000-381, Dept. of Computer Science, Dartmouth College, November 2000.

Creswell, J.W., 2009. *Research design: Qualitative, quantitative, and mixed methods approaches*, Sage Publications, Inc.

Dairy Queen, Carmel Cheesequake Royal Shake. Available at: http://www.dairyqueen.com/us-en/Menu/Drinks/Carmel-Cheesequake-Royal-Shake/ [Accessed August 3, 2012].

Darby, S., 2001. Making it obvious: designing feedback into energy consumption. In *Proceedings, 2nd International Conference on Energy Efficiency in Household Appliances and Lighting. Italian Association of Energy Economists/EC-SAVE programme*.

Darwell, B., Users spent more than 10.5B total minutes per day on Facebook in January, not including mobile use. *InsideFacebook.com*. Available at: http://www.insidefacebook.com/2012/03/27/users-spent-more-than-10-5b-total-minutes-per-day-on-facebook-in-january-not-including-mobile-use/ [Accessed September 14, 2012].

Demand Media, 2011a. LIVESTRONG.COM - Lose Weight & Get Fit with Diet, Nutrition & Fitness Tools. Available at: http://www.livestrong.com/ [Accessed August 22, 2011].

Demand Media, 2011b. Long-term Effects Of Diabetes Type One. Available at: http://www.livestrong.com/article/77488-longterm-effects-diabetes-type-one/ [Accessed August 22, 2011].

Dey, A.K., 2001. Understanding and using context. *Personal and ubiquitous computing*, 5(1), pp.4–7.

Eagle, N. & Pentland, A., 2006. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4), pp.255–268.

Einhorn, H.J. & Hogarth, R.M., 1981. Behavioral decision theory: Processes of judgment and choice. *Journal of Accounting Research*, pp.1–31.

Energy, Inc, 2008. TED 1000. Available at: http://www.theenergydetective.com/ted-1000 [Accessed August 22, 2011].

EnergyHub Inc, 2011. EnergyHub Home Base. Available at: http://www.energyhub.com/our-products/home-base/ [Accessed August 22, 2011].

Feldman, A. et al., 2005. *ReachMedia: On-the-move Interaction with Everyday Objects*. Massachusetts Institute of Technology, School of Architecture and Planning, Program in Media Arts and Sciences.

Festinger, L., 1957. *A theory of cognitive dissonance*, Stanford Univ Pr.

Fischer, C., 2008. Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency*, 1(1), pp.79–104.

Fishkin, K.P., 2004. A taxonomy for and analysis of tangible interfaces. *Personal and Ubiquitous Computing*, 8(5), pp.347–358.

Fitbit Inc., Fitbit. Available at: http://www.fitbit.com/ [Accessed September 14, 2012].

Fogg, B.J., 2002. Persuasive technology: using computers to change what we think and do. *Ubiquity*, 2002(December), p.5.

Gallistel, C.R. & Gibbon, J., 2000. Time, rate, and conditioning. *Psychological review*, 107(2), p.289.

Gardner, R.A. & Gardner, B.T., 1988. Feedforward versus feedbackward: An ethological alternative to the law of effect. *Behavioral and Brain Sciences*, 11(03), pp.429–447.

Geier, A.B., Rozin, P. & Doros, G., 2006. Unit bias - A New Heuristic That Helps Explain the Effect of Portion Size on Food Intake. *Psychological Science*, 17(6), p.521.

Gigerenzer, G. & Goldstein, D.G., 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, 103(4), p.650.

Gilbert, D.T., 1989. Thinking lightly about others: Automatic components of the social inference process. *Unintended thought*, pp.189–211.

Greenberg, S. & Buxton, B., 2008. Usability evaluation considered harmful (some of the time). In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*. ACM, pp. 111–120.

Gross, T., 2003. Ambient Interfaces: Design Challenges and Recommendations. *Proceedings of the 10th International Conference on Human-Computer Interaction-HCII*, pp.68–72.

Hagras, H. et al., 2004. Creating an ambient-intelligence environment using embedded agents. *Intelligent Systems, IEEE*, 19(6), pp.12–20.

Herrnstein, R.J. & Prelec, D., 1991. Melioration: A theory of distributed choice. *The Journal of Economic Perspectives*, 5(3), pp.137–156.

Hertwig, R. et al., 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), p.534.

Higgins, E.T. & Brendl, C.M., 1995. Accessibility and applicability: Some" activation rules" influencing judgment. *Journal of Experimental Social Psychology*, 31(3), pp.218–243.

Huang, A.C., Ling, B.C. & Ponnekanti, S., 1999. Pervasive computing: what is it good for? *Proceedings of the 1st ACM international workshop on Data engineering for wireless and mobile access*, pp.84–91.

Intille, S.S., 2004. A new research challenge: persuasive technology to motivate healthy aging. *Information Technology in Biomedicine, IEEE Transactions on*, 8(3), pp.235–237.

Ishii, H. & Ullmer, B., 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp.234–241.

Jackson, T.W., Dawson, R. & Wilson, D., 2003. Understanding email interaction increases organizational productivity. *Communications of the ACM*, 46(8), pp.80–84.

Jacobson, N.S. & Hollon, S.D., 1996. Cognitive–behavior therapy versus pharmacotherapy: Now that the jury's returned its verdict, it's time to present the rest of the evidence. *Journal of Consulting and Clinical Psychology*, 64(1), p.74.

Jamieson, S., 2004. Likert scales: how to (ab) use them. *Medical Education*, 38(12), pp.1217–1218.

Kahneman, D., 2003. Maps of bounded rationality: Psychology for behavioral economics. *American economic review*, pp.1449–1475.

Kahneman, D. & Frederick, S., 2002. Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, pp.49–81.

Kahneman, D. & Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pp.263–291.

Kahneman, D. & Tversky, A., 1972. Subjective probability: A judgment of representativeness. *Cognitive psychology*, 3(3), pp.430–454.

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*. Claremont, California: ACM, pp. 1–6.

Kraut, R. et al., 1998. Internet paradox: A social technology that reduces social involvement and psychological well-being? *American psychologist*, 53(9), p.1017.

Laibowitz, M. & Paradiso, J.A., 2004. The UbER-Badge, a versatile platform at the juncture between wearable and social computing. *Advances in Pervasive Computing, Oesterreichische Computer Gesellschaft*, 363, p.368.

Laibson, D. & Zeckhauser, R., 1998. Amos Tversky and the ascent of behavioral economics. *Journal of Risk and Uncertainty*, 16(1), pp.7–47.

Langens, T.A., 2007. Regulatory focus and illusions of control. *Personality and Social Psychology Bulletin*, 33(2), p.226.

Liebermann, H., 2003. The tyranny of evaluation. In ACM CHI Fringe. Fort Lauderdale, Florida. Available at: http://web.media.mit.edu/~lieber/Misc/Tyranny-Evaluation.html.

Lo, B. et al., 2005. Body sensor network–a wireless sensor platform for pervasive healthcare monitoring. In *The 3rd International Conference on Pervasive Computing*.
Loomes, G. & Sugden, R., 1982. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368), pp.805–824.
Lumsden, J. & Brewster, S., 2003. A paradigm shift: alternative interaction techniques for use with mobile & wearable devices. *Proceedings of 13th Annual IBM Centers for Advanced Studies Conference CASCON*, pp.97–110.
Lynch, J.G., Jr. & Srull, T.K., 1982. Memory and attentional factors in consumer choice: Concepts and research methods. *Journal of Consumer Research*, pp.18–37.
Madan, A., 2007. *Jerk-O-Meter: Speech-Feature Analysis Provides Feedback on Your Phone Interactions*,
March, J.G., 1978. Bounded rationality, ambiguity, and the engineering of choice. *The Bell Journal of Economics*, pp.587–608.
Mehta, N., Rajiv, S. & Srinivasan, K., 2004. Role of forgetting in memory-based choice decisions: A structural model. *Quantitative Marketing and Economics*, 2(2), pp.107–140.
Microsoft MSN, The Scion Story. Available at: http://autos.msn.com/browse/Scion.aspx [Accessed August 20, 2012].
Midden Joanne, F. et al., 1983. Using feedback, reinforcement and information to reduce energy consumption in households: A field-experiment. *Journal of Economic Psychology*, 3(1), pp.65–86.
Morley, S., Eccleston, C. & Williams, A., 1999. Systematic review and meta-analysis of randomized controlled trials of cognitive behaviour therapy and behaviour therapy for chronic pain in adults, excluding headache. *Pain*, 80(1), pp.1–13.
Mullainathan, S. & Thaler, R.H., 2000. *Behavioral economics*, National Bureau of Economic Research Cambridge, Mass., USA.
Myerson, J. & Hale, S., 1988. Choice in transition: A comparison of melioration and the kinetic model. *Journal of the experimental analysis of behavior*, 49(2), p.291.
Navon, D., 1977. Forest before trees: The precedence of global features in visual perception. *Cognitive psychology*, 9(3), pp.353–383.
Neale, T., 2011. Number of Patients with Diabetes Climbs to 24 Million. Available at: http://www.medpagetoday.com/Endocrinology/Diabetes/9927 [Accessed August 22, 2011].
Nike, A., Inc., 2011. Apple - Run or workout with Nike + iPod. Available at: http://www.apple.com/ipod/nike/ [Accessed August 13, 2011].
O'Flynn, B. et al., 2005. The development of a novel minaturized modular platform for wireless sensor networks. In *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*. IEEE, pp. 370–375.
Obermair, C. et al., 2008. perFrames: Persuasive picture frames for proper posture. *Persuasive Technology*, pp.128–139.
Ogden, C.L. et al., 2006. Prevalence of overweight and obesity in the United States, 1999-2004. *JAMA: the journal of the American Medical Association*, 295(13), p.1549.
de Oliveira, R. & Oliver, N., 2008. TripleBeat: enhancing exercise performance with persuasion. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*. ACM, pp. 255–264.
Onzo, 2008. Onzo - Solutions. Available at: http://onzo.com/solutions/ [Accessed August 22, 2011].
Park, C. & Chou, P.H., 2006. Eco: Ultra-wearable and expandable wireless sensor platform.
Pavlov, I.P., 1927. Conditioned reflexes: An investigation of the activity of the cerebral cortex. *Trans. GV Anrep*.
Petersen, D., Steele, J. & Wilkerson, J., 2009. WattBot: a residential electricity monitoring and feedback system. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*. Boston, MA, USA: ACM, pp. 2847–2852.

Du Plessis, E., 1994. Recognition versus recall. *Journal of Advertising Research*, (3), pp.75–91.
Progressive Insurance, How It Works: Snapshot Discount, Usage-Based Insurance. Available at: http://www.progressive.com/auto/snapshot-how-it-works.aspx [Accessed September 16, 2012].
RedOrbit.com, Percentage of Overweight, Obese Americans Falling. Available at: http://www.redorbit.com/news/health/1112397597/percentage-of-overweight-obese-americans-falling/ [Accessed September 14, 2012].
Rescorla, R. & Wagner, A., 1972. Variations in the Effectiveness of Reinforcement and Nonreinforcement. *New York: Classical Conditioning II: Current Research and Theory, Appleton-Century-Crofts*.
Rescorla, R.A., 1988. Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43(3), p.151.
Rheingold, H., 2000. *The virtual community: Homesteading on the electronic frontier*, The MIT Press.
Rothbart, M. et al., 1978. From individual to group impressions: Availability heuristics in stereotype formation* 1. *Journal of Experimental Social Psychology*, 14(3), pp.237–255.
Rubinstein, A., 2003. "Economics and Psychology"? The Case of Hyperbolic Discounting. *International Economic Review*, 44(4), pp.1207–1216.
Saha, D. & Mukherjee, A., 2003. Pervasive computing: a paradigm for the 21st century. *Computer*, 36(3), pp.25–31.
Satyanarayanan, M., 2001. Pervasive computing: vision and challenges. *Personal Communications, IEEE [see also IEEE Wireless Communications]*, 8(4), pp.10–17.
Seligman, C. & Darley, J.M., 1977. Feedback as a means of decreasing residential energy consumption. *Journal of Applied Psychology*, 62(4), p.363.
Simon, H.A., 1959. Theories of decision-making in economics and behavioral science. *The American Economic Review*, 49(3), pp.253–283.
Skinner, B.F., 1974. An Experimental Analysis of Behaviour. *London: Methuen and Co.*
Staddon, J.E.R., 1988. Quasi-dynamic choice models: Melioration and ratio invariance. *Journal of the experimental analysis of behavior*, 49(2), p.303.
Stein, H., 1988. Logos, logic, and Logistiké: Some philosophical remarks on the Nineteenth Century transformation of mathematics. *History and philosophy of modern mathematics*, 11, pp.238–259.
Tapia, E.M., Intille, S.S. & Larson, K., 2004. Activity recognition in the home using simple and ubiquitous sensors. *Pervasive Computing*, pp.158–175.
Thaler, R., 1985. Mental accounting and consumer choice. *Marketing science*, pp.199–214.
Thaler, R. & Sunstein, C., 2008. *Nudge: The Gentle Power of Choice Architecture*, New Haven, CT: Yale University Press.
Thompson, S.C., 1999. Illusions of Control. *Current Directions in Psychological Science*, 8(6), p.187.
Tunney, R.J. & Shanks, D.R., 2002. A re-examination of melioration and rational choice. *Journal of Behavioral Decision Making*, 15(4), pp.291–311.
Turkle, S., 1995. *Life on the Screen: Identity in the Age of the Internet*, Simon and Schuster.
Tversky, A. & Kahneman, D., 1986. Rational choice and the framing of decisions. *Journal of business*, pp.251–278.
Tversky, A. & Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science*, 211(4481), p.453.
US Energy Information Administration, Total Energy - U.S. Energy Information Administration (EIA). Available at: http://www.eia.gov/totalenergy/ [Accessed September 14, 2012].
Vigderhous, G., 1977. The Level of Measurement and" Permissible" Statistical Analysis in Social Research. *Pacific Sociological Review*, pp.61–72.

Vogel, D. & Balakrishnan, R., 2004. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*. ACM, pp. 137–146.
Vohs, K.D. & Heatherton, T.F., 2000. Self-regulatory failure: A resource-depletion approach. *Psychological Science*, 11(3), p.249.
Vroom, V.H., 1966. Organizational choice: A study of pre-and postdecision processes. *Organizational Behavior and Human Performance*, 1(2), pp.212–225.
Weiser, M., 1999. Some computer science issues in ubiquitous computing. *ACM SIGMOBILE Mobile Computing and Communications Review*, 3(3).
Weiser, M., 1991. The Computer for the Twenty-First Century. *Scientific American*, 265(3), pp.94–104.
Weiser, M. & Brown, J.S., 1996. *The coming age of calm technology*, Xerox PARC, October.
Wisneski, C. et al., 1998. Ambient displays: Turning architectural space into an interface between people and digital information. *Cooperative buildings: Integrating information, organization, and architecture*, pp.22–32.
Wood, G. & Newborough, M., 2003. Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design. *Energy and Buildings*, 35(8), pp.821–841.
Woodworth, R.S. & Thorndike, E.L., 1901. The influence of improvement in one mental function upon the efficiency of other functions.(I). *Psychological review*, 8(3), p.247.
Young, K.S., 1999. Internet addiction: symptoms, evaluation and treatment. *Innovations in clinical practice: A source book*, 17, pp.19–31.

Appendix A: Questionnaires for Web Surfing Study

Initial Questionnaire

Web Surfing - Starting Survey

Thinking about your most usual weekday (ie, work day, regular schedule, etc), please answer the following questions. Please consider what you believe the daily average is for each question.

1* **How often during a week do you surf the web?**

0-5 days ▾

2* **How may days do you surf for entertainment (eg netflix, web based games, etc) per week?**

0-5 days ▾

3* **How much time do you spend surfing per day (on average)?**

0-15 hours ▾

4* **How does the time you spend surfing compare to how much time you would like to spend?**

1 Much less	2	3	4 Expected amount	5	6	7 Much more	N/A
○	○	○	○	○	○	○	○

5* **How much time per day do you spend surfing for entertainment (on average)? Examples may be Netflix, YouTube, Last.fm, social networking sites, dating sites, etc.**

0-15 hours ▾

6* How does the time you spend surfing for entertainment compare to how much time you would like to spend?

	1 Much less	2	3	4 Expected amount	5	6	7 Much more	N/A
	○	○	○	○	○	○	○	○

7* How much time per day do you spend surfing for productivity (on average)? Examples may be email, research, finance or banking, etc.

0-15 hours ▾

8* How does the time you spend surfing for productivity purposes compare to how much time you would like to spend?

	1 Much less	2	3	4 Expected amount	5	6	7 Much more	N/A
	○	○	○	○	○	○	○	○

9* How much of the time do you feel you spend surfing without a particular goal? This may include spending time on entertainment or productivity site, but with no particular goals.

0-15 hours ▾

10* How does the time you spend surfing without specific goals compare to how much time you would like to spend?

	1 Much less	2	3	4 Expected amount	5	6	7 Much more	N/A
	○	○	○	○	○	○	○	○

11* How aware are you of the passage of time as you surf?

	1 Not aware at all	2	3	4 Normal awareness of time	5	6	7 Highly aware of time
	○	○	○	○	○	○	○

12* How often do you set a time goal or budget for the amount of time you spend surfing the web? For example, you may tell yourself you will only surf for 10 minutes before starting.

	1 Never	2	3	4 Sometimes	5	6	7 Every time
	○	○	○	○	○	○	○

13* If you set a goal, how often (when you have set a goal) do you stick successfully to that goal?

	1 Never	2	3	4 Sometimes	5	6	7 Every time	I never set a goal
	○	○	○	○	○	○	○	○

14* How much value (eg, satisfaction about productivity or entertainment) do you get from your time online? In other words, to what extent do you feel you spent the time well for the results you received?

1 Much less	2	3	4 What I expect	5	6	7 Much more
○	○	○	○	○	○	○

15* How well do you feel you balance online time between focused (with goals) and unfocused (just surfing without any particular goals) time?

1 A lot of unfocused time	2	3	4 Balanced equally	5	6	7 A lot of focused time
○	○	○	○	○	○	○

16* Do you user any tools to limit your time online?

- ☐ Timer
- ☐ Alarm
- ☐ Mental note
- ☐ Brower-based plugin
- ☐ Other software
- ☐ Reminder from another person
- ☐ I don't use any tools

Other (Please Specify):

17 Please tell us any other information you feel is relevant to your web surfing habits:

Initial Configuration Questionnaire

Feedback Start Survey

Please upgrade and configure your add-on before taking this short survey. We have moved these questions to an earlier point in time so that you may answer them at a time when the memory of your choices and reactions are still fresh.

1* **How helpful was the time information in configuring the add-on?**

	1 - Not Helpful	2	3	4 - Neutral	5	6	7 - Very Helpful
	○	○	○	○	○	○	○

2*

	1 - Very Few	2	3	4	5	6	7 - Almost All
How many of the time estimates are HIGHER than you expected?	○	○	○	○	○	○	○
How many of the time estimates are LOWER than you expected?	○	○	○	○	○	○	○
How many of the time estimates MATCHED you expectations?	○	○	○	○	○	○	○
How many of the time estimates SEEMED INCORRECT?	○	○	○	○	○	○	○

3 **Comments on this change:**

Final Questionnaire

Web Surfing Study Final Survey

Thank you for completing the study.

Now that you have had a chance to use the various prototypes, please answer the questions below with regards to your experience.

Thank you!

1* **What portions of the study did you complete?**

- Initial Survey
- + Used add-on for calibration period
- + Configured add-on and enabled feedback
- + Used add-on with visual feedback
- + Used add-on with mouse (without problems)
- Other (Please Specify): ____________

2* **How often did you reconfigure the add-on to your preferences?**

- Only at the beginning
- 1-2 times
- 3-4 times
- 5-7 times
- More than 7 times
- I did not know that I could reconfigure the add-on

3* **How often did you change the feedback mode?**

- Only at the beginning
- 1-2 times
- 3-4 times
- 5-7 times
- More than 7 times
- I did not know that I could reconfigure the add-on

4* **How do you feel you split your time between the time and focus modes?**

1 Used focus mode all the time	2	3	4 Used both modes equally	5	6	7 Used time mode all the time
○	○	○	○	○	○	○

5* How helpful did you find...

	1 Not helpful at all	2	3	4 Neither helpful not un-helpful	5	6	7 Very helpful	Not Applicable or did not use
... changing modes?	◯	◯	◯	◯	◯	◯	◯	◯
... changing blacklist or whitelist?	◯	◯	◯	◯	◯	◯	◯	◯
... changing settings for time mode?	◯	◯	◯	◯	◯	◯	◯	◯
... ability to disable add-on?	◯	◯	◯	◯	◯	◯	◯	◯
... the time feedback mode itself?	◯	◯	◯	◯	◯	◯	◯	◯
... the focus feedback mode itself?	◯	◯	◯	◯	◯	◯	◯	◯
... the bar graph on the add-on?	◯	◯	◯	◯	◯	◯	◯	◯
... the popup reminders from the add-on?	◯	◯	◯	◯	◯	◯	◯	◯
... the bar graph on the mouse?	◯	◯	◯	◯	◯	◯	◯	◯
... the vibration warning on the mouse?	◯	◯	◯	◯	◯	◯	◯	◯
... the mouse itself?	◯	◯	◯	◯	◯	◯	◯	◯

6* How understandable did you find...

	1 Not helpful at all	2	3	4 Neither helpful not un-helpful	5	6	7 Very helpful	Not Applicable or did not use
... changing modes?	◯	◯	◯	◯	◯	◯	◯	◯
... changing blacklist or whitelist?	◯	◯	◯	◯	◯	◯	◯	◯
... changing settings for time mode?	◯	◯	◯	◯	◯	◯	◯	◯
... the time feedback mode itself?	◯	◯	◯	◯	◯	◯	◯	◯
... the focus feedback mode itself?	◯	◯	◯	◯	◯	◯	◯	◯
... the bar graph on the add-on?	◯	◯	◯	◯	◯	◯	◯	◯
... the popup reminders from the add-on?	◯	◯	◯	◯	◯	◯	◯	◯
... the bar graph on the mouse?	◯	◯	◯	◯	◯	◯	◯	◯
... the vibration warning on the mouse?	◯	◯	◯	◯	◯	◯	◯	◯

7* How obtrusive (attention-getting to annoying) do you feel the following were:

	1 Completely unobtrusive	2	3	4 Neither unobtrusive nor obtrusive	5	6	7 Highly obtrusive	N/A or did not use
The bar display of the add-on	○	○	○	○	○	○	○	○
The popups from the add-on	○	○	○	○	○	○	○	○
The bars on the mouse	○	○	○	○	○	○	○	○
The vibration function of the mouse	○	○	○	○	○	○	○	○

8* How often did you glance at:

	1 I rarely looked	2	3	4	5	6	7 I looked very often	N/A or did not use
The bars on the screen	○	○	○	○	○	○	○	○
The bars on the mouse	○	○	○	○	○	○	○	○

9* When you did look at the feedback, when do you feel you noticed the feedback from the system with respect to how you were using the system?

For example, did you feel that when you looked at the bars on the screen, it was already too late to change how you were acting, or that you had enough time to change your actions?

	1 I noticed the feedback too late	2	3	4 I noticed the feedback at the right time	5	6	7 I noticed the feedback all the time	N/A or did not use
The bars in time mode	○	○	○	○	○	○	○	○
The bars in focus mode	○	○	○	○	○	○	○	○
The popups in time mode	○	○	○	○	○	○	○	○
The popups in focus mode	○	○	○	○	○	○	○	○
The bars on the mouse	○	○	○	○	○	○	○	○

10* How do you feel the following affected your awareness of how you spend your time?

	1 It had no effect	2	3	4	5	6	7 It had a large effect	N/A or did not use
The process of configuring the add-on	○	○	○	○	○	○	○	○
Having the add-on installed	○	○	○	○	○	○	○	○
The bar graphs on screen	○	○	○	○	○	○	○	○
The alerts on screen	○	○	○	○	○	○	○	○
Having the mouse installed	○	○	○	○	○	○	○	○
The bar graphs on the mouse	○	○	○	○	○	○	○	○
The vibration of the mouse	○	○	○	○	○	○	○	○

Additional open-ended questions regarding each phase of the study were included at the end of this questionnaire.

Appendix B: Questionnaires for Eating Study

Initial Questionnaire

1 How many days (of weekdays) a week do you have breakfast on an average day?

2 How much time do you spend exclusively for breakfast (eg, sitting at a table):

	None	< 5 minutes	< 10 minutes	< 15 minutes	< 30 minutes	More than 30 minutes
	◯	◯	◯	◯	◯	◯

3 To what extent do you feel hurried on the average weekday during breakfast?

	1 Very relaxed	2	3	4 neither hurried nor unconstrained	5	6	7 extremely hurried
	◯	◯	◯	◯	◯	◯	◯

4 Where kind of food you have for breakfast?

	1 only prepared foods	2	3	4 a mix of prepared an homemade foods	5	6	7 only homemade foods
	◯	◯	◯	◯	◯	◯	◯

Final Questionnaire

1 How many days did you use the utensil?

▾ 0-5 days

2 On how many days did the utensil function correctly?

▾ 0-5 days

3 How helpful did you find:

	1 Not helpful at all	2	3	4 Neither helpful nor un-helpful	5	6	7 Very helpful	Not Applicable or did not use
... the utensil in general	○	○	○	○	○	○	○	○
... bar graphs?	○	○	○	○	○	○	○	○
... bite rate alert?	○	○	○	○	○	○	○	○
... end of meal alert?	○	○	○	○	○	○	○	○

4 How understandable did you find:

	1 Very confusing	2	3	4 Neither confusing nor understandable	5	6	7 Very understandable	Not Applicable or did not use
... the utensil in general	○	○	○	○	○	○	○	○
... bar graphs?	○	○	○	○	○	○	○	○
... bite rate alert?	○	○	○	○	○	○	○	○
... end of meal alert?	○	○	○	○	○	○	○	○

5 How often did you glance at:

	1 Rarely	2	3	4 Every once in a while	5	6	7 I looked at it all the time	Not Applicable or did not use
... bar graphs?	○	○	○	○	○	○	○	○
... bite rate alert?	○	○	○	○	○	○	○	○
... end of meal alert?	○	○	○	○	○	○	○	○

6 When you did look at the feedback, when do you feel you noticed the feedback from the system with respect to how you were using the system?

	1 I noticed the feedback too late	2	3	4 I noticed the feedback at the right time	5	6	7 I noticed the feedback all the time	Not Applicable or did not use
... bar graphs?								
... bite rate alert?								
... end of meal alert?								

7 How obtrusive (attention-getting to annoying) do you feel the following were:

	1 Completely unobtrusive	2	3	4 Neither unobtrusive nor obtrusive	5	6	7 Highly obtrusive	Not Applicable or did not use
... bar graphs?								
... bite rate alert?								
... end of meal alert?								

8 How do you feel the following affected your awareness of your eating behavior?

	1 No effect	2	3	4	5	6	7 Very large effect	Not Applicable or did not use
... having the device?								
... eating with a special spoon?								
... bar graphs?								
... bite rate alert?								
... end of meal alert?								

Additional open-ended questions regarding each phase of the study were included at the end of this questionnaire.