

**SUPPORTING USER UNDERSTANDING AND ENGAGEMENT
IN DESIGNING INTELLIGENT SYSTEMS FOR THE HOME**

by

Rayoung Yang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in The University of Michigan
2016

Doctoral Committee:

Associate Professor Mark W. Newman, Chair
Professor Mark S. Ackerman
Professor Jodi L. Forlizzi, Carnegie Mellon University
Assistant Professor Emily M. Provost

Copyright © 2016 Rayoung Yang

DEDICATION

To my parents, Hyun-im Kil (길현임) & Tae-soo Yang (양태수)

ACKNOWLEDGEMENTS

“The Lord is my shepherd, I lack nothing.

He makes me lie down in green pastures, he leads me beside quiet waters, he refreshes my soul.

He guides me along the right paths for his name’s sake.” Psalm 23:1-3

I feel deeply grateful to many people from whom I have received love, support, and expertise, and without whom I could not have finished this dissertation. Specifically, my thanks to:

Mark W. Newman, my advisor and dissertation chair, for his persistent encouragement and invaluable advice throughout my intellectual journey. This research would not have happened at all without my advisor. He trusted and supported me all the way from when I was a Master’s student to this finish line. He gave me confidence in the direction I wanted to take regarding my dissertation. I would like to offer my deepest thanks to my advisor for his guidance.

My committee members – Mark Ackerman, Jodi Forlizzi and Emily Provost – for their intellectual support of this dissertation and for their thoughtful feedback. I am privileged and honored to have had these great scholars on my committee. In particular, Mark Ackerman has been always generous in taking time to provide feedback on my research. Every time I talked with him, I was challenged, inspired and encouraged. I am very lucky to have collaborated with Jodi Forlizzi. She provided me an excellent guidance and feedback. I am grateful to have had Emily Provost as my cognate member. Thank you for offering your expertise and insights.

Inteco members – Stanley Chang, Tao Dong, Gaurav Paruthi, Jeff Huang, Manchul Han, I-chun Hsiao and Shriti Raj – for providing me great reassurance and insightful advice on my work. Tao and Stanley have been a great source of support from the very beginning of my Ph.D. journey. Having them as my colleagues, lab mates, and friends has been a blessing as we worked our way through the program.

Karina Kervin, Ayse Buyuktur, and Chrysta Meadowbrook, for all the precious memories that we shared, and our joy and struggles; Liz Kaziunas, for constant reinforcement and needed encouragement throughout my graduate study; Jina Huh and Jessica Hullman for inspiring me and fostering my confidence and strength. You are my role models; Xuan Zhao, Jasmine Jones, and Youyang Hou for being dear friends who kept me company whenever I needed someone to rely on; Melissa Chalmers, Rebecca Frank, and Andrea Barbarin for giving me comfort as I worked through the final stages of my dissertation.

Thanks to all of the Ph.D. students at UMSI. My wonderful colleagues and friends, with whom I shared the emotional roller coaster of graduate study, who inspired and encouraged me to keep my scholarly journey going. I give my deepest thanks to all who provided comments on my work and more importantly made happy memories during my time with you all.

Thank you, as well, to the faculty and staff members at UMSI. Elizabeth Yakel and Tom Finholt for giving me the self-assurance to follow my curiosity at the beginning of my dissertation work; Tawanna Dillahunt for being my mentor and friend; Soo Young Rieh provided me professional advice during my program. I want to thank members of the UMSI staff, including Allison Sweet, Glenda Bullock, Jay Jackson, Dennis Hogan, and Mike Doa. Thank you for all the support you have provided to me.

I want to convey special thanks to my dear friends at UMSI: Sangseok You, Grace Jeon, Jiyeon Yang, Jamin Koo, Souneil Park, Sungjin Nam, and Jungwon Yang. My life in Ann Arbor was never lonely thanks to your friendship.

My friends for their prayers and spiritual support: Eunice Shin, Ayoung Yoon, Woojeong Shim, Na Hyung Kim, Arum Han, HyunJin Cho, Kyoo Hye Lim, and Sangeun Park. You're my beloved sisters!

Jim Knox, the first Director of the University of Michigan Adaptive Technology Computing Site, for being a great mentor and an excellent example of advocating for the user. His encouragement and kindness during my first research experience made a huge difference in my life. Thank you, Jim.

The Rademaker family – Laura, Patrick, Willem, and Owen Rademaker – for being my family in the United States and the Netherlands. Thank you, Laura and Patrick, for bringing adorable human beings into the world and making me a proud Emo.

My beloved family – my parents and my brother – has been always with me regardless of the physical distance between us. Thank you for your love and endless support. I love you!

And finally, my Heavenly Father for His grace, mercy and love.

RESEARCH ACKNOWLEDGEMENT

Throughout the thesis, I use “we” while describing the conducted research because much of the presented work is the result of collaborations of varying degrees. I would like to highlight the contributions of others.

1. I conducted three studies, reported on in Chapter 3, Chapter 4, and Chapter 5, in collaboration with my advisor, Mark W. Newman.
2. **Eco-Interaction study (Chapter 4):** While I conducted all interviews and diary studies, I collaborated with Dr. Jodi Forlizzi, Carnegie Mellon University, on the analysis.
3. **ThermoCoach study (Chapter 5):** I collaborated with Devika Pisharoty and Dr. Kamin Whitehouse, University of Virginia. They designed, developed and deployed the ThermoCoach prototype. I conducted the evaluation of the prototype during the field deployment. Dr. Soodeh Montazeri conducted 25 interviews for this study.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
Research Acknowledgement	v
LIST OF TABLES	ix
LIST OF FIGURES	x
ABSTRACT	xii
CHAPTER 1. INTRODUCTION	1
Thesis Problem	5
Thesis Statements	9
Thesis Approaches	9
Thesis Contributions	10
Thesis Overview	12
CHAPTER 2. LITERATURE REVIEW	14
Part 1: Intelligent Technology for the Home	14
Part 2: Intelligent Technology for Sustainability	21
CHAPTER 3. LEARNING FROM A LEARNING THERMOSTAT: LESSONS FOR INTELLIGENT SYSTEMS FOR THE HOME	30
Introduction	30
Related Work	32
The Nest Thermostat	34
Study Methods and Participants	35
Findings	40
Improved design leads to greater engagement	41
The learning system fails to understand user intent	42
The system's behavior is hard to understand	43
Users found ways to work with the 'intelligence'	45

The Nest did not clearly lead to energy savings	48
Discussion	50
Bridging the intention gap: Exception flagging	52
Bridging the understanding gap: Incidental intelligibility	53
Widening the interaction: Constrained engagement	54
Conclusion	56

CHAPTER 4. MAKING SUSTAINABILITY SUSTAINABLE: CHALLENGES IN THE DESIGN

OF ECO-INTERACTION TECHNOLOGIES **57**

Introduction.....	57
Background and Related Work.....	59
Methods.....	63
Study 1: Conventional thermostat study	64
Study 2: The Nest thermostat study	64
Data collection and analysis	65
Common Problems with Thermostat Control.....	66
People do not use a setback temperature.	66
People do not use schedules.	67
People fail to reassess existing control patterns.	67
Product Ecology Analysis	68
Initial Experience with the Nest.....	69
Increased engagement and awareness	69
Changing practices and interaction with the Nest	70
Settling into a Routine	73
Less interaction with the Nest over time	73
Reduced effort in improving energy performance	74
Changes in behavior were not sustained	77
Discussion	78
Designing eco-interactions	78
Use mixed initiative to balance competing concerns.	79
Bridge the gap between awareness and control	80
Reframe interactions around reflection and reassessment	81
Conclusion.....	83

CHAPTER 5. ASSESSING A RECOMMENDATION SYSTEM FOR ENERGY-EFFICIENT THERMOSTAT SCHEDULING	84
Introduction.....	84
Background and Related Work.....	86
Eco-Coaching Design Approach.....	89
Methods.....	90
Participants	91
Field Deployment Study.....	91
Findings.....	95
Supporting user agency	96
Encouraging experiments with a setback strategy	99
Shortcomings of ThermoCoach	103
Discussion.....	106
Creating recommendations: Improving personalization	107
Assessing recommendations: Multi-phase assessment	108
Conclusion.....	112
CHAPTER 6. CONCLUSION	113
Limitations and caveats.....	119
Concluding Remarks.....	123
REFERENCES	127

LIST OF TABLES

TABLE 1. SUMMARY OF PARTICIPANTS * P13 SUBMITTED ADDITIONAL DIARY ENTRIES AFTER HER DIARY STUDY COMPLETED. ** P16 AND P17 WHO PARTICIPATED IN AN INTERVIEW STUDY IN FEBRUARY 2012 PARTICIPATED IN A FOLLOW-UP INTERVIEW IN AUGUST 2012. PT: PROGRAMMABLE THERMOSTAT, H: HEATING, C: COOLING.....	37
TABLE 2. KEY CHARACTERISTICS OF THE 36 HOUSEHOLDS.....	91
TABLE 3. THE PARTICIPANT HOMES WERE DIVIDED INTO THREE GROUPS. ALL GROUPS RECEIVED ECO-FEEDBACK EMAILS (FIGURE 9). GROUP P COULD ONLY MANUALLY PROGRAM THEIR THERMOSTAT. GROUP N USED NEST’S AUTOMATION FEATURES. GROUP TC RECEIVED ECO-COACHING RECOMMENDATIONS (FIGURE 10).	92

LIST OF FIGURES

FIGURE 1. EVERYDAY TECHNOLOGY FOR THE HOME BECOMES MORE CONNECTED, AUTOMATED, AND INTELLIGENT.....	2
FIGURE 2. USERS CAN CONTROL THE NEST VIA THE WALL-MOUNTED DISPLAY (A), A MOBILE APP (B), OR A WEB APP (C) THE MOBILE AND WEB APPS PROVIDE ACCESS TO ENERGY HISTORY (D).	4
FIGURE 3. USERS CAN CONTROL THE NEST VIA THE WALL-MOUNTED DISPLAY (A), A MOBILE APP (B), OR A WEB APP (C) THE MOBILE AND WEB APPS PROVIDE ACCESS TO ENERGY HISTORY (D).	35
FIGURE 4. P9'S NEST SCHEDULE SHOWED FREQUENT TEMPERATURE CHANGES ON CERTAIN DAYS. TIME IS PLOTTED ON THE X-AXIS AND WEEKDAYS ARE PLOTTED ON THE Y-AXIS. THE DOTS SHOW THE TEMPERATURE SETTING AT THE PARTICULAR DAY AND TIME.	46
FIGURE 5. P2 POSTED SCREENSHOTS OF HIS SCHEDULE BEFORE AND AFTER HE MODIFIED IT.....	47
FIGURE 6. USERS CAN CONTROL THE NEST VIA THE WALL-MOUNTED DISPLAY (A), A MOBILE APP (B), OR A WEB APP (C). © NEST LABS.....	61
FIGURE 7. P4'S ENERGY HISTORY SHOWS THE A/C WAS RUNNING DURING THE WEEKEND WHILE HE WAS NOT AT HOME (BOTTOM TWO ROWS). WHEN HE LEFT HOME, HE SET THE NEST TO 'AWAY' MODE, WHICH WAS SET TO 80°F, A TEMPERATURE LOWER THAN HIS NORMAL DAYTIME SETBACK OF 85°F (SHOWN IN THE UPPER THREE ROWS). P4 WAS SURPRISED TO DISCOVER THAT HIS SYSTEM WAS WASTING ENERGY WHILE HE WAS AWAY.....	75
FIGURE 8. P14'S ENERGY HISTORY IS SHOWING THAT THE COOLING WAS ON DURING THE WEEKEND WHEN NOBODY WAS AT HOME.	76
FIGURE 9. THE NEST MONTHLY REPORT EMAIL PROVIDES ECO-FEEDBACK. IT DISPLAYS NUMBER OF HOURS FOR THE COOLING OR HEATING USED IN COMPARISON OF THIS PAST MONTH TO THE PREVIOUS ONE, WHAT FACTORS IMPACTED MORE OR FEWER HOURS OF COOLING OR HEATING WERE USED IN THIS PAST MONTH COMPARED TO THE PREVIOUS ONE, A TIP TO HELP YOU SAVE MORE ENERGY, AND THE NUMBER OF A HOME EARNED A GREEN LEAF BY ADJUSTING THE TEMPERATURE TO SAVE ENERGY.....	93

FIGURE 10. THERMOCOACH EMAILS 4 OPTIONS TO EACH USER: THEIR CURRENT SCHEDULE, A HIGH COMFORT SCHEDULE, AN ENERGY SAVER SCHEDULE, AND A SUPER ENERGY SAVER SCHEDULE. USERS CAN ACTIVATE A SCHEDULE BY CLICKING A BUTTON TO SET IT AUTOMATICALLY TO BE PROGRAMMED INTO THEIR THERMOSTAT.94

FIGURE 11. JIM'S PREVIOUS SCHEDULE KEPT THE TEMPERATURE MOSTLY AT 76°F ALL TIMES. 101

FIGURE 12. JIM ADOPTED THE 'HIGH COMFORT' OPTION, WHICH INCLUDED TWO NEW SETBACKS; ONE WAS 8°F DEGREES HIGHER AND THE OTHER WAS 4°F HIGHER THAN HIS NORMAL TEMPERATURE. 102

ABSTRACT

With advances in computing, networking and sensing technology, our everyday objects have become more automated, connected, and intelligent. This dissertation aims to inform the design and implementation of future intelligent systems and devices. To do so, this dissertation presents three studies that investigated user interaction with and experience of intelligent systems. In particular, we look at intelligent technologies that employ sensing technology and machine learning algorithm to perceive and respond to user behavior, and that support energy savings in the home.

We first investigated how people understand and use an intelligent thermostat in their everyday homes to identify problems and challenges that users encounter. Subsequently, we examined the opportunities and challenges for intelligent systems that aimed to save energy, by comparing how people's interaction changed between conventional and smart thermostats as well as how interaction with smart thermostats changed over time. These two qualitative studies led us to the third study. In the final study, we evaluated a smart thermostat that offered a new approach to the management of thermostat schedule in a field deployment, exploring effective ways to define roles for intelligent systems and their users in achieving their mutual goals of energy savings.

Based on findings from these studies, this dissertation argues that supporting user understanding and user control of intelligent systems for the home is critical allowing users to intervene effectively when the system does not work as desired. In addition, sustaining user engagement with the system over time is essential for the system to obtain necessary user input and feedback that help improve the system performance and achieve user goals.

Informed by findings and insights from the studies, we identify design challenges and strategies in designing end-user interaction with intelligent technologies for the home: making system behaviors intuitive and intelligible; maintaining long-term, easy user engagement over time; and balancing interplay between user control and system autonomy to better achieve their mutual goals.

CHAPTER 1. INTRODUCTION

With advances in computing, everyday systems and devices in the home are becoming more connected, automated, and intelligent. This trend follows the trajectory of the “smart home” that has been forecasted and researched in the HCI and Ubicomp communities for the past two decades. This vision describes a home which seeks to adapt to its inhabitants and respond to their informational and comfort needs (Weiser & Brown, 1997), and there is increasing evidence that the vision is poised to become a reality.

Many home appliance manufacturers are introducing new generations of digitally enhanced home appliances, which promise the benefit of reducing manual work, operating efficiently with little or no user intervention, and providing new types of information to the user that was not available previously. Examples of these devices include, but are not limited to, applications such as dishwashers that select energy efficient cycles depending on the load, robot vacuum cleaners that autonomously clean the house, as well as video recorders that are able to recommend movies.

Managing home energy consumption represents a particularly rich domain for smart, domestic technologies, especially when 21% of the total energy consumed in the United States is used by home appliances (EIA, 2011). In particular, there are two good reasons to study home heating, ventilation, and cooling (HVAC) systems—one is that they are important from the energy perspective, and the other is that they are the first to gain traction as a mass-market smart home technology. In the United States, for example, residential HVAC systems account

for roughly 50% of all household energy consumption, which equates to about 9% of the nation's total energy budget (EIA, 2011). Designing technologies for energy savings has been the focus of numerous research projects as well as commercial offerings. The advancement of thermostat control over the past 60 years illustrates this trend well.



Figure 1. Everyday technology for the home becomes more connected, automated, and intelligent.

A simple manual thermostat is easy to set to maintain a temperature, and it will remain at that temperature unless someone changes the setting. However, it is more inconvenient for people to manually adjust the temperature throughout the day as they get up, leave and return home and go to bed. Often people forget to change the temperature setting and waste energy by running a heating or cooling system when they are not at home.

A more advanced programmable thermostat automatically changes the temperature according to a schedule its user defines. This reduces the inconvenience of walking up to the thermostat or forgetting to change the temperature before going work. However, programming the thermostat is difficult for lay users in the home (Peffer, Pritoni, Meier, Aragon, & Perry, 2011), as the temperature schedule often does not match people's changing schedules. Sometimes, it becomes more inconvenient and annoying to change or override the temperature schedule. Thus half of programmable thermostats run inefficiently or are no longer used (ibid.).

To address the problem of residential HVAC systems not being operated efficiently by their users (Peffer et al., 2011), leading to unnecessarily wasted energy, a number of researchers have investigated ways to improve the operation of HVAC systems. Research into eco-feedback (e.g., (Froehlich, Findlater, & Landay, 2010)) has focused on ways to provide information to people about their resource usage in order to motivate them to change their usage patterns. However, there is little evidence that obtaining information reliably causes people to take action or change behavior (Strengers, 2011).

Another approach that has been investigated is predictive heating control, which uses sensing and machine learning to try to learn the occupancy patterns of a house's residents in order to automatically adjust the temperature. Work in predictive control seeks to reduce or even eliminate the role of user choice in controlling HVAC systems by automating temperature adjustments based on occupancy predictions (e.g., (Gupta, Intille, & Larson, 2009; Koehler, Ziebart, Mankoff, & Dey, 2013; Scott et al., 2011)). Systems in this category have been built and tested in limited deployments. It remains to be seen what issues would arise in a more general deployment with people who vary more widely in terms of geographic mobility, schedule predictability, tolerance for error, and desire for control.

Many manufacturers have been interested in adding new energy saving functions to energy consuming devices. Recently, a new generation of thermostats has become available. These aim to solve the problems of programmable thermostats by automatically creating and updating temperature schedules (to eliminate the trouble of programming to adapt to users' changing schedules) and adjusts temperature based on occupancy in the home in order to save energy.

In late 2011, the Nest thermostat was introduced to the market and received a great deal of media attention. It was considerably more advanced than other available thermostats, with novel features such as schedule learning, remote access, occupancy sensing, and eco-feedback.

The Nest features an attractive wall-mounted device, as well as smart phone and web-based control capabilities (Figure 2). In addition to providing access to the temperature schedule and real time control, its web and phone apps provide a feature called “Energy History,” which offers a detailed history of when and how long the heating and cooling system ran.

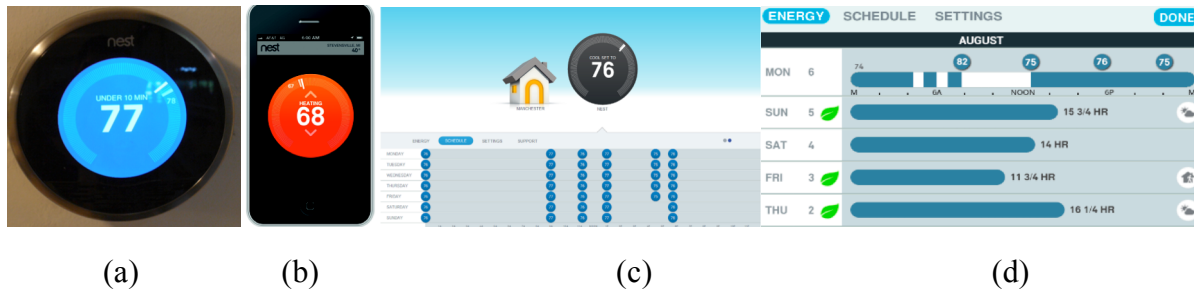


Figure 2. Users can control the Nest via the wall-mounted display (a), a mobile app (b), or a web app (c) The mobile and web apps provide access to Energy History (d).

The Nest includes a pair of intelligent features that utilize machine learning and motion sensing: Auto-Schedule and Auto-Away. The Auto-Schedule feature automatically generates a schedule based on temperature changes users make. The Nest takes about a week to generate its initial schedule and thereafter continually adapts the schedule according to users’ temperature adjustments. Users can manually revise the schedule via the wall-mounted device or through the web or mobile applications. Users can also turn off this feature and use the Nest as a regular programmable thermostat. The Nest has an embedded motion sensor on the wall-mounted unit that detects the movement of occupants within a certain range. If the Nest does not sense movement for some time, it goes into “Auto-Away” mode, which automatically adjusts the temperature to a user-defined level to avoid heating or cooling an empty home.

The Nest represents an intriguing phenomenon for study, as it is the first mass-market thermostat in the U.S. to feature motion sensing and machine learning. With the Nest’s ‘smart’ capabilities of recognizing human movement and learning human thermostat control behavior,

it promises to generate a personalized heating and cooling schedule that will promote comfort, energy savings, convenience, and more enjoyable interaction. We are starting to interact more with intelligent systems in our daily home environment, and this trend will almost certainly increase.

THESIS PROBLEM

These intelligent systems are often depicted as easy and simple to use and promise to bring convenience and comfort to the home by autonomously working on behalf of their users. However, questions remain regarding how well these intelligent systems will work in the everyday home environment and how users will interact with and experience these novel systems.

These new capabilities bring great potential, but also great concern — are “smart” devices going to make our lives easier, more productive, or more enjoyable? Or are they going to bring a new set of frustrations, expectations, and responsibilities that will outweigh their possible benefits?

As everyday intelligent systems continue to evolve and play a larger role in the management of daily tasks such as keeping oneself healthy and making home environments comfortable, more research is needed to examine the user experience of such technologies and emerging problems and issues users encounter in using and interacting with them over time. As such, the Nest Learning Thermostat provides an excellent opportunity to study the user experience of living with a ‘smart’ domestic appliance in the wild, particularly one that seeks to learn and adapt to consumers’ behavior, and help people to save energy.

In this thesis work, we are interested in investigating various challenges and tensions that arise when users interact with intelligent systems in everyday home environments. Using the Nest Learning Thermostat as a lens, we aim to better understand the real-life, long-term user

experience of living with a smart thermostat in order to inform the design of intelligent systems for the home more broadly. In the following paragraphs, we highlight key challenges facing designers of everyday intelligent systems. These are informed by previous research on ‘smart homes,’ ‘adaptive systems,’ and ‘interactive intelligent systems’ that tried to address these and related issues.

1) User Understanding and System Intelligibility

First, it is difficult for users to understand how intelligent systems work, especially when these systems gather multiple sources of implicit data and utilize complex algorithms to act on that data (Eagle & Pentland, 2006; C. D. Kidd et al., 1999). As these intelligent systems and devices aim to learn more by gathering both explicit and implicit data about us (location, activities, behavior patterns, preferences and interests) and act on our behalf, they become so much more complex and unpredictable that it becomes difficult to understand what they are doing and are going to do (Edwards & Grinter, 2001). This lack of understanding often leads users to lose control and trust, and ultimately, disuse the system (Lim, Dey, & Avrahami, 2009). Bellotti and Edwards (2001) emphasize the importance of “intelligibility” for adaptive systems, which allows users to know how the system learns about its users’ changing contexts and thus to understand why the system behaves in certain ways. However, most research on supporting intelligibility has investigated interactive machine learning systems and employed one-off, lab-based studies.

Many interactive machine learning systems learn user preferences based on inputs and feedback users provide, but do not use sensors to observe and learn user behaviors, which are rather implicit and unconscious. Therefore, there is a need for more exploration to support intelligibility of intelligent systems that seek to learn and adapt to everyday user behaviors in the home.

II) System Limitations and User Control

Another significant challenge for intelligent systems is to make accurate inferences about users' status and nuanced contexts based on limited data the system gathers from sensors and other data sources (Suchman, 2006). Peoples' everyday lives in the home are full of unpredictable and nuanced situations and events (changing routines, situations, preferences and expectations). There is a mix of diverse individuals with different characteristics (different levels of technical aptitude, varying preferences and motivations) living in heterogeneous homes. However, machines can only interpret limited types and ranges of sensory inputs to a fixed set of states and responses (Suchman, 2006). This might not be too problematic if the intelligent system works in a closed environment such as a lab or a purposefully built home like those used in early smart home research studies. However, dynamically changing, unpredictable and nuanced everyday situations make standard homes a more difficult domain for intelligent systems to function in. Indeed, even for other humans, it is not simple and straightforward to infer the reasoning behind other human's actions and behaviors (Bellotti & Edwards, 2001).

One strategy to mitigate these limitations of intelligent systems is to simply ask users about their intentions, preferences, plans, and goals. Human input is necessary to provide proper and necessary data to the system in order to improve its performance and prevent malfunctions (Kapoor, Lee, Tan, & Horvitz, 2010; Rogers, 2006). Supporting such communication requires that users understand system states and decisions, and know what types of feedback would be helpful, thus tying back to the work on intelligibility. It also requires effectively managing users' attention so as not to exhaust their patience with the system. Mozer (2005) noted the paradox that more information from the user would improve the performance of the smart home, but a system which requests less information directly from users would be considered more successful.

III) Understanding Lived Experience in Everyday Home Environment

Finally, when technologies enter the house, not only do they infiltrate the home environment, they also change daily domestic life (Dourish & Bell, 2011). For example, when washing machines became prevalent in the home, they changed the standard of cleanliness and paradoxically increased the domestic workload (Cowan, 1993). In order to design technology that fits into peoples' daily lives at home, it is crucial for designers of devices to understand characteristics of daily activities, tasks, and unique physical and social aspects of a home environment (Crabtree, Rodden, Hemmings, & Benford, 2003). There has been little investigation of the quality of user experiences and interaction with intelligent systems in *everyday life situations*.

Technical demonstrations of intelligent environments have illustrated the feasibility and desirability of adaptive systems for the home (e.g., (Cook et al., 2003; Intille, 2002; Kidd et al., 1999; Mozer, 2005), but few projects have provided insight into the lived experience of occupants. Previous research on user experience of intelligent systems have been mostly conducted in laboratories (e.g., (Intille, 2002; Kidd et al., 1999)), or with prototypes in experimental settings (e.g., (Gupta et al., 2009; Scott et al., 2011)), and lasted for relatively short periods. The particular nature of the everyday sphere suggests that even established approaches to supporting user interaction with intelligent systems need to be reexamined, *in situ*, in everyday environments. In real daily life situations, human behavior is more unpredictable; preferences change over time; everyday routines are unstable and contingencies are too rare to form a pattern (Suchman, 2006).

The goal for this dissertation is three-fold: First, we seek to better understand users' lived experience with intelligent systems, particularly what challenges they face when interacting with and using the systems, and how the systems' intelligent features influence users' interaction with and use of the systems over time in everyday home environments. Second, by

observing the ways in which users interact with everyday intelligent technologies and identifying challenges and problems that arise over time, we aim to provide design guidelines and considerations for desirable goals and design properties for end-user interaction with everyday intelligent technologies. Finally, we focus on the design of intelligent systems to help people manage their home energy consumption more effectively.

THESIS STATEMENTS

The claim this thesis makes is summarized in the following statements:

This dissertation argues that supporting user understanding and control of intelligent systems is necessary to promote user engagement and improve system performance for long-term use.

THESIS APPROACHES

This thesis consists of several studies informing and evaluating the design of intelligent systems to support energy savings in the home. We summarize the methods and approaches that we employed in this thesis work below.

1) Identifying challenges in user understanding and interactions through lived experience

In order to better understand the challenges of deploying intelligent systems in the home, we studied the experience of living with a commercially available intelligent thermostat, the Nest Learning Thermostat. The Nest utilizes sensing, machine learning, and networking technology, as well as eco-feedback features. It learns users' behavior patterns, and then adapts and automates its operation to control home heating and cooling systems. We conducted interviews with 23 participants, ten of whom also participated in a three-week diary study. We collected empirical data eliciting challenges and problems users encountered when using intelligent systems for the home.

II) Describing changes in practices and user interaction over time

Whereas our first study drew on early-stage usage experiences to inform the design of usable intelligent systems for the home, our goal in the second study was to observe changes in user interactions between conventional and intelligent thermostats, as well as changes in their interactions with intelligent thermostat over time. We compared user interactions and practices around thermostat control with conventional thermostats, intelligent thermostats other than the Nest, and the Nest, and also demonstrated how users' interactions with intelligent thermostats changed over time through a longitudinal study. This study also aimed to identify challenges and opportunities in the design of eco-interaction technologies, by which we mean the study of interaction between humans and energy-consuming systems, with an eye towards minimizing energy use while preserving an acceptable level of user-perceived benefits.

III) Evaluating a design approach to balance system autonomy and user control

With findings from two previous qualitative studies that resulted in a set of design approaches to address challenges in user interaction with everyday intelligent technologies, the following step was to evaluate design strategies that we proposed. In our third study, we deployed a prototype for a smart thermostat that employed a mixed-initiative approach and evaluated how users respond to and interact with the system. In particular, we investigated the impacts of recommendations for thermostat scheduling and eco-feedforward features.

THESIS CONTRIBUTIONS

The following contributions arise from our studies of user interaction with intelligent systems in their everyday environments, and the design and field deployment of a prototype system:

- 1) Empirical evidence that describes users' lived experience of everyday intelligent technologies over time and demonstrates problems and challenges that users encountered in their daily environment.

- I. Lack of support for intelligibility and user control in everyday intelligent technologies hinders users from understanding how the system interprets and adapts to users' behavior and situations, and thus deters them from intervening to guide or correct the system's behavior. (Chapter 3)
 - II. Users' engagement with the system helps to address system shortcomings and improve performance. However, maintaining users' engagement over time is difficult when users have little motivation to go through the effort of understanding and assessing the system's behavior. (Chapter 3, Chapter 4)
 - III. Users' reliance on intelligent systems and diminished interactions results in missed opportunities for energy savings. Sustaining user interaction and engagement with intelligent system is critical to achieve the goal of energy savings. (Chapter 4)
- 2) We develop design recommendations for end-user interaction with everyday intelligent technologies.
- I) We propose three avenues for future development of everyday intelligent technologies to support user understanding and control of the system: *Exception flagging, Incidental intelligibility, and Constrained engagement*. (Chapter 3)
 - II) We propose that employing mixed initiatives is a promising direction for balancing system autonomy and user control. We create a set of design implications for eco-interaction, the design of features and human-system interactions with the goal of saving energy, which includes: *Providing actionable recommendations, Providing eco-feedforward and Stimulating reflection and reassessment rather than control and convenience*. (Chapter 4)
 - III) To build user trust with recommendation-based eco-coaching systems, systems should support users to assess the quality and performance of recommendations over time:

Assessing the actual performance of recommendations after use with consideration of real-world factors and conditions, Providing hindsight evidence with post-hoc simulation of alternative recommendations, and Performing assessment for the schedule in use. (Chapter 5)

3) We provide contributions to sustainable HCI.

IV) We found that the combination of eco-feedback and machine learning-based personalization led to increased engagement with energy-saving features of the system in the short term, but that such engagement was not sustained over the long term. (Chapter 3 and Chapter 4)

V) We found that the eco-coaching approach 1) made it easier for users to implement an effective thermostat schedule, 2) supported user agency in negotiating trade-offs between energy savings and comfort, 3) facilitated learning different scheduling strategies as well as weighing pros and cons of different options, and 4) challenged users' beliefs about how well they were doing. These outcomes, in turn, were successful in getting users to employ and experiment with more efficient setback strategies. (Chapter 5)

THESIS OVERVIEW

This thesis is divided into six chapters, as follows.

Chapter 2 reviews and synthesizes background work and previous research that informs this thesis. We describe related work about users' interaction and experience with intelligent systems as well as the design and implementation of technologies to support sustainability. We explain the research gap and how this dissertation work addresses the gap.

Chapter 3 presents a formative study that we conducted. In order to better understand the challenges of deploying intelligent systems in the home, we studied the experience of living with an advanced thermostat, the Nest. The Nest utilizes sensing, machine learning, and networking technology, as well as eco-feedback features. We identify challenges and opportunities for designing intelligent systems for the home.

Chapter 4 focuses on findings regarding long-term user experiences with intelligent systems for energy savings. It compares user interaction with conventional manual and programmable thermostats with user interaction with an intelligent thermostat. Then, it presents a one-year follow-up study of intelligent thermostat users. The findings of this study generate design guidelines to inform the design of ThermoCoach, which we evaluate in the following chapter.

Chapter 5 describes the results of the evaluation study of the ThermoCoach system. We first introduce the design approach of eco-coaching and detail the design features of ThermoCoach. Then, we describe how these influenced user thermostat scheduling and energy savings outcomes. We also describe design implications for future eco-coaching system design and implementation.

Chapter 6 reflects on and draws conclusions about the implications of future design of intelligent systems for the home. We highlight the importance of supporting user control as a design principle, and discuss design implications of supporting user agency in energy conservation in the home. This chapter also identifies areas for future work.

CHAPTER 2. LITERATURE REVIEW

In this literature review, we will examine background work informing this thesis, reviewing previous research on user interaction with intelligent technologies and identifying issues that users face as they use these novel intelligent systems. The larger goal of this literature review is to provide insights into understanding how **users** interact with **intelligent systems**, which learn and respond to users' actions and behavior, and operate autonomously in their **everyday home environment**. First, we will provide an overview of the development of “smart home” technology; we will then review key challenges such technology presents to users. Finally, we will focus on applications of intelligent technology to reducing energy consumption, particularly in the area of home heating and cooling.

PART 1: INTELLIGENT TECHNOLOGY FOR THE HOME

Weiser (1991) described the vision of “ubiquitous computing.” He used it to describe the future in which computers become more invisible as they are integrated into our daily life and make daily tasks easier and our everyday environment more informative (Weiser, 1999). To do so, ubiquitous computing focuses on designing “calm technologies” (Weiser & Brown, 1997) that make the home environment smart by learning about its inhabitants and responding accordingly without disrupting their ongoing activities or daily routines.

Visions of the smart home have played a central role in ubiquitous computing research since its inception (Weiser, 1999). Broadly speaking, a smart home can be described as one that adapts to its inhabitants (Brush et al., 2011) and responds to their varying informational and comfort

needs (Rogers, 2006). Different versions of smart home projects have been driven by both academia and industry, and their focuses vary. Different terms, such as “pervasive,” “invisible,” “calm,” “context awareness,” “ambient intelligence,” and “information appliances” illustrate diverse characteristics promulgated by different approaches (Abowd & Sterbenz, 2000; Rogers, 2006).

In the human-computer interaction (HCI) and ubiquitous computing (UbiComp) fields, various research studies have been motivated by Weiser’s vision. Smart home research has been pursued through a variety of large-scale initiatives, including Mozer’s adaptive house (2005), the Georgia Tech Aware Home (Kidd et al., 1999), and MIT’s House_n (Intille, 2002). These projects were valuable for demonstrating the potential of smart home technology, but their comprehensive approach to redesigning the home technical environment has meant that many of their envisioned applications and interactions remain impractical to implement. Edwards and Grinter note that a more likely scenario is the “accidentally smart home,” (2001) in which smart, connectable devices enter the home piecemeal over a long period of time, without conscious planning on the part of the inhabitants. In sync with this vision, a number of projects have looked at more specific application opportunities within the broad umbrella of the smart home, including adaptive heating and cooling systems (Gupta, Intille, & Larson, 2009; Peffer, Pritoni, Meier, Aragon, & Perry, 2011; Scott et al., 2011), health monitoring (Consolvo, Roessler, & Shelton, 2004; Kaushik, Intille, & Larson, 2008; Rowan & Mynatt, 2005), and reminder systems to help people manage their plans and schedules (Davidoff, Zimmerman, & Dey, 2010; McGee-Lennon, Wolters, & Brewster, 2011).

With advances in computing, long envisioned by HCI and UbiComp researchers, the promise of a home that can learn its occupants’ needs, desires, and behaviors — and adapt itself appropriately — is being realized. Many home appliance manufacturers are introducing new generations of digitally enhanced home appliances, which promise to reduce manual work,

operate efficiently on behalf of users with little or no user intervention, and perform a variety of different roles in the home, including managing entertainment, health, security, and home automation.

These new capabilities bring great potential, but also great concern — are “smart” devices going to make our lives easier, more productive, or more enjoyable? Or will they bring a new set of frustrations, expectations, and responsibilities that will outweigh their possible benefits?

Such concerns underlie a number of challenges that face designers of intelligent systems in the home. Firstly, it is difficult for users to understand how adaptive systems work, especially when these systems gather multiple sources of implicit data and utilize complex algorithms to act on that data (Eagle & Pentland, 2006; Kidd et al., 1999). Secondly, it is difficult for the system to make accurate inferences about users’ status and nuanced context based on limited data the system gathers from sensors and other data sources (Bellotti & Edwards, 2001; Brush et al., 2011; Eagle & Pentland, 2006). Therefore, users’ input is often necessary to improve the performance of the system, but this requirement can be at odds with typical notions of the smart home based on visions of calm (Weiser & Brown, 1997) or invisible computing (Rogers, 2006). Finally, more research is required to guide designers in creating systems that will be successfully adopted and integrated into home life (Davidoff, Lee, Zimmerman, & Dey, 2006; Sadri, 2011).

In the following sections, we will cover how previous research has tried to address these and related issues and seek to identify still unresolved questions that require further exploration.

Improving Intelligibility

When a system is working properly, as expected by its users, knowing ‘how-to-use-it’ may suffice. However, when a system behaves in an unexpected and erroneous way, a user’s understanding of ‘how-it-works’ becomes more crucial if the user is to be able to identify

errors and fix problems (Fein, Olson, & Olson, 1993). When a user does not understand the system's behavior—when he or she lacks an accurate mental model of the system—the result is often inefficient use, confusion, dissatisfaction, and abandonment of some features of the system (Lim & Dey, 2010). Supporting users' understanding is especially important because there is evidence that novice users form mental models of system operation during the early stage of their interaction with the system, and that those mental models are unlikely to change based on subsequent interactions even if disconfirming evidence is encountered (Tullio, Dey, Chalecki, & Fogarty, 2007).

The first challenge of supporting users to build a proper mental model of a system has been studied extensively under the topic of “intelligibility,” which covers user interface techniques that seek to help users understand the behavior of complex, often intelligent, systems. Bellotti and Edwards (2001) emphasize the importance of “intelligibility” for adaptive systems, which allows users to know how the system learns about its users' changing contexts and thus to understand why the system behaves in certain ways.

One approach to promoting intelligibility is enhancing the system's interface with additional data to help promote users' understanding of how the system works (McNee, Lam, Guetzlaff, Konstan, & Riedl, 2003). It has been shown that even simple explanations can help to increase users' trust in a system and improve their overall satisfaction with using it (Herlocker, Konstan, & Riedl, 2000). For example, providing explanations can help users understand the strengths and weaknesses of the system, and lead them to use the system in a more accurate, efficient way (Herlocker, Konstan, & Riedl, 2000). This, in turn, can provide an overall boost to users' satisfaction with intelligent systems. Explanations can be provided through special displays and/or interactions that users can engage in order to gain a better understanding of system states or behaviors.

One notable approach to provide explanations was developed by Ko and Myers (Ko & Myers, 2004). They proposed “Why?” and “Why Not?” dialogs as a technique for end-user debugging. Subsequent research has investigated the application of this technique to machine learning-based systems and found them to be effective there as well (Kulesza et al., 2009; Lim & Dey, 2009; Stumpf et al., 2006, 2009). For example, Lim and Dey (2010; 2009) explored presenting multiple explanations for context-aware applications and found that explanations improve user understanding, trust, and control in different context-aware applications. Individual differences among users may suggest the need for multiple, perhaps personalized approaches to generating and presenting explanations (Kulesza et al., 2009; Lim & Dey, 2009; Stumpf et al., 2006; Tintarev & Masthoff, 2007).

While many research studies have focused on providing interactive explanations for how the system works and why it behaves in certain ways, there are some drawbacks—lack of longitudinal studies and assumptions about motivation. Nearly all of the research on supporting intelligibility has employed one-off, lab-based studies. Evaluation in everyday situations, in particular longitudinal study is much needed in order to thoroughly understand the impact of the intelligibility of a system’s behavior, or lack thereof (Lim & Dey, 2009; Stumpf et al., 2009) with respect to supporting diverse users. Another drawback is that all of these approaches to intelligibility assume that the user has a conscious interest in understanding the system and is willing to invest time in doing so. However, in everyday settings, lay-users may not be motivated to learn about the technical specifications of a system, but are more likely to be interested in what it can do (Paepcke & Takayama, 2010). Therefore more work on lightweight approaches to intelligibility may be required.

Eliciting user input

The second significant challenge for adaptive systems is to elicit input from users in order to train and correct the systems’ inferences about the world and users’ status (Eagle & Pentland,

2006). As machines can only interpret limited types and ranges of sensory inputs to a fixed set of states and responses, it is difficult for them to make accurate inferences about users' status and nuanced contexts (Suchman, 2006). Moreover, today's sensing technology is often imprecise in detecting the data it is supposed to detect. This produces inaccurate or skewed data, leading to misinterpretations of a user's activities (Youngblood & Cook, 2007). Finally, even when more reliable and accurate sensors become available, there are various human behavioral characteristics for which it is nearly impossible for systems to make inferences (Bellotti & Edwards, 2001). Indeed, even for other humans, it is not simple and straightforward to infer the reasoning behind other human's actions and behaviors (Bellotti & Edwards, 2001).

With complete accuracy out of the question, the question becomes, "How can we develop an *adequate* model of users and the real world for specific applications?" One strategy is to simply ask users about their intentions, preferences, plans, and goals. Keeping the human in the loop can improve system performance across the lifecycle (Kapoor et al., 2010; Rogers, 2006). Supporting such communication requires that users understand system states and decisions, and know what types of feedback would be helpful, which ties back to the work on intelligibility. As mentioned earlier, supporting users' understanding of intelligent systems helps users to provide more and improved feedback, and results in better performance of the system for users' needs (Kulesza, Stumpf, Burnett, & Kwan, 2012). It also requires effectively managing users' attention so as not to exhaust their patience with the system. Mozer noted the paradox that more information from the user would improve the performance of the smart home, while a system that requests less information directly from users would be considered more successful (Mozer, 2005).

Promoting engagement between users and machine learning systems has been explored under the aegis of interactive machine learning (Kulesza et al., 2009; Kulesza, 2012; Kulesza, Burnett, Wong, & Stumpf, 2015; Stumpf et al., 2009). Based on studies of music

recommenders and email classification tools, Stumpf *et al.* (Stumpf et al., 2006;) and Kulesza *et al.* (Kulesza et al., 2009) argued that machine learning systems should support users' ability to fix their logic and system behavior when problems occur. They both stressed the need for two-way communication between end-users and machine learning systems, and demonstrated the benefits of allowing users to provide feedback on system performance (Kulesza et al., 2009; Stumpf et al., 2009).

Many interactive machine learning systems learn user preferences based on inputs and feedback users provide (i.e., you can indicate your preferences to a music recommender system such as Pandora¹ by explicitly clicking a 'like' or 'dislike' button when you hear a song that you like or dislike). However, intelligent systems that passively observe user behavior using sensors often do not provide an interface like a music recommender. As Bellotti et al. (2002) addressed there are new design challenges for designing novel systems that use "sensing" user-interfaces. While traditional graphical user interface design have provided interaction techniques to support communications between the interactive systems and the users, more research is needed to gain insights into user interaction and engagement with intelligent technologies that seek to learn and adapt to everyday user behaviors, which are implicit and often unconscious behavior.

Fitting into everyday environments

Designing technology for the home requires understanding the characteristics of people's daily activities and tasks in the home and the particular constraints and opportunities related to the

¹ <http://www.pandora.com/>

physical and social environment of the home. Crabtree and Rodden (2004) have stressed the importance of understanding how domestic routines operate and impact technology use, while Friedewald, Da Costa, Punie, Alahuhta, and Heinonen (2005) have emphasized that technology must be designed to adapt to users' changing situations while still remaining under users' control Davidoff *et al.* (2006) noted the dynamicity of household routines, and presented design principles to better support flexibility for variable daily routines and the “conflicting goals” of multiple users in the home.

However, there has been little investigation of the quality of user experience and interaction with intelligent systems in *everyday life situations*. Technical demonstrations of intelligent environments have illustrated the feasibility and desirability of adaptive systems for the home (e.g., (Cook et al., 2003; Intille, 2002; Kidd et al., 1999; Mozer, 2005), but few projects have provided insight into the lived experience of occupants. Previous research on user experience of intelligent systems has mostly been conducted in laboratories (e.g., (Intille, 2002; Kidd et al., 1999)), or with prototypes in experimental settings (e.g., (Gupta et al., 2009; Scott et al., 2011)), and has also lasted for relatively short periods. The particular nature of the everyday life sphere suggests that even established approaches to supporting user interaction with intelligent systems needs to be reexamined, *in situ*, in everyday environments. In real daily life situations, human behavior is more unpredictable than it is in a lab; people's preferences change over time; everyday routines are unstable and contingencies are too rare to form a pattern (Suchman, 2006).

Next, we proceed to the second section of the literature review.

PART 2: INTELLIGENT TECHNOLOGY FOR SUSTAINABILITY

A particular area of domestic technology use that has received attention within the Ubicomp and HCI communities is that of managing energy consumption. Given that 21% of the total

energy consumed in the United States is used by homes, such attention is clearly warranted. Increased interest in efficient energy consumption has prompted research investigating the use of home appliances as well as new designs to support energy savings in the home. For a more detailed description, Coskun et al. provide a comprehensive list of smart home projects that promote sustainable behaviors (Coskun, Zimmerman, & Erbug, 2015).

One area where smart home devices promise to deliver great benefits is in the control of home heating, ventilation, and cooling (HVAC) systems. HVAC control is an important domain from the perspective of environmental sustainability. For example, residential HVAC systems account for roughly 50% of all household energy consumption in the United States, which equates to about 9% of the nation's total energy budget.² Moreover, it is known that residential HVAC systems are not operated efficiently by their users (Peffer et al., 2011), leading to unnecessarily wasted energy.

Home Heating and Cooling Control

Previous literature has shown that proper control of the thermostat in summer and winter can save energy without sacrificing thermal comfort (Al-Sanea & Zedan, 2008). One strategy for saving energy in heating and cooling is to use a “setback” temperature. A setback temperature is used to reduce the temperature at certain times, such as when no one is at home during the day or when people are sleeping during the night.

However, commonly available thermostats do not provide adequate support for people to operate their HVAC system in the most energy efficient ways (Peffer et al., 2011). Thermostats that require manual control are frequently left on all the time, even when the house is empty.

² <http://www.eia.gov/consumption/residential/data/2009/>

People with a manual thermostat often forget to adjust the temperature or keep it running at the same temperature all the time to maintain a comfortable temperature upon arriving at home.

Programmable thermostats allow users to input schedules of temperature changes to reduce the likelihood that an empty house will be heated or cooled to an unnecessary level. Programmable thermostats automatically operate their HVAC systems and can help save energy by scheduling the thermostat settings according to daily pattern.

While the ability to support setbacks is one of the supposed benefits of programmable thermostats, programmable thermostats are no longer considered energy-saving appliances due to the lack of proper use (Meier, Aragon, Peffer, Perry, & Pritoni, 2011). While 42% of households in the U.S. have programmable thermostats, only 56% of these are actually used on a regular basis (Peffer et al., 2011).

Meier et al. (2010) investigated what problems and complaints people have regarding programmable thermostats and concluded that poor usability of programmable thermostats is the main barrier to efficient use of programmable thermostats. One common reason for this is that their programming interfaces are very difficult to use (Dey, Hamid, Beckmann, Li, & Hsu, 2004; Meier, Aragon, et al., 2010). Common programmable thermostats only allow people to create rigid schedules for their heating and cooling that often cannot accommodate irregular or unexpected changes in their household routines (Meier, Aragon, et al., 2010). As a result, a large number of houses are heated or cooled while no one is home, resulting in wasted energy.

Responding to these problems with commonly available manual and programmable thermostat control, a number of researchers have investigated ways to increase energy conservation in the home. Here we discuss two approaches. First, research into eco-feedback has focused on ways to provide information to people about their resource usage in order to motivate them to

change their usage patterns. Second, automation approach such as predictive control tries to reduce the workload for users by automating tasks that users otherwise need to do manually.

Eco-feedback

Eco-feedback systems have been proposed as a way to promote greater awareness of energy use, which could in turn, motivate people to save more energy (work in this area is extensive; Froehlich, Findlater, and Landay (2010) provides a survey). For example, eco-feedback displays real-time data to inform users about their energy consumption of various resources, such as electricity, gas or water, and thus seeks to motivate users to change their energy use behaviors. As people better understand how much they are consuming in real-time, or even detect problems that arise (e.g., leaking water pipe), they can act to reduce their energy waste. It has been shown that providing real time, dynamic feedback information regarding energy usage can lead to energy savings of 5-15%, but only if people are already motivated to reduce energy consumption (Darby, 2006).

While eco-feedback has been shown to increase awareness of energy consumption when people are motivated, several studies have found that obtaining information did not actually trigger people to take action or change behavior (Pierce, Fan, Lomas, Marcu, & Paulos, 2010; Strengers, 2011). Pierce et al. (2010) and Strengers (2011) investigated everyday practices of people consuming energy and pointed out that obtaining information does not actually cause people to take action or change behaviors (Pierce, Schiano, & Paulos, 2010; Strengers, 2011). Pierce et al. (2010) claimed that people do not always make rational decisions; rather they habitually consume energy and pursue convenience.

Considerable motivation and engagement on the part of consumers is required for eco-feedback to lead to behavior change. Strengers (2014) warned against a common assumption that the eco-feedback approach holds, depicting the user as “a resource man” who makes

rational choices and acts accordingly when provided information. Even when people are aware and motivated, it can be still difficult for them to effectively control their systems. As mentioned earlier, poor usability was a significant barrier for the efficient use of programmable thermostats (Meier, Aragon, et al., 2010).

Given the challenges of persuading people to change their behavior, Pierce et al. (2010) suggested designing interfaces to “nudge” people to save energy by default, thereby reducing the need for consumers to make conscious decisions or enact behavior changes. Hazas, Brush, and Scott (2012) take this argument further, proposing that technology-centered approaches, rather than user-centered approaches, offer the greatest promise for saving energy.

Intelligent systems for thermostat control

Managing home energy consumption represents a particularly rich domain for smart domestic technologies and has been the focus of numerous projects. Many smart home research looked at a more specific application space of adaptive heating and cooling systems within the broad umbrella of the smart home (Gupta et al., 2009; Peffer et al., 2011; Scott et al., 2011).

A promising approach for helping people maintain an acceptable level of comfort while attaining greater energy efficiency is to automate the operation of the system to some degree. For example, work in predictive control seeks to reduce or even eliminate the role of user choice in controlling HVAC systems by automating temperature adjustments based on occupancy predictions. There has been an emerging wave of “intelligent” thermostats that seek to learn occupants’ preferences and adjust the temperature based on sensed conditions such as householders’ geographic location (Gupta et al., 2009) or home occupancy (Scott et al., 2011).

By tracking occupancy patterns using GPS (Gupta et al., 2009; Koehler, Ziebart, Mankoff, & Dey, 2013) or RFID and motion sensing (Scott et al., 2011), it is possible to build reasonably accurate models that can predict occupancy and make sure a house is heated or cooled to a

desired temperature when people are home and to a less energy-intensive level otherwise. Gupta et al. (2009) used GPS data to predict the arrival time of a home's residents to control the thermostat to reach the pre-defined temperature on the person's arrival. Scott et al. (2011) gathered occupancy data through RFID and motion sensors, and used the data to predict occupancy patterns and operate the thermostat accordingly.

These systems have shown promise. However, systems in this category have been built and tested in limited deployments. While these automation-based approaches promise to relieve the programming burden for users, there are reasons to believe that the benefits of a fully automated approach will not be realized in a meaningful and straightforward manner.

It remains to be seen what issues would arise in a more general deployment with people who vary more widely in terms of geographic mobility, schedule predictability, tolerance for error, and desire for control. 100% accuracy will be unattainable, and it is not clear how much error consumers will tolerate. Removing control from users makes it more difficult for users to understand how the system works, detect errors, and fix problems. This is especially problematic when there are unexpected and nuanced changes occurring in someone's daily life. These changes are not understood or managed well by an automated system.

In recent years, interest has grown in applying smart technology in the area of reducing energy consumption. Many manufactures have also been interested in adding new energy-saving functions to energy consuming devices. This interest has resulted in the introduction of novel consumer devices like the Nest Learning Thermostat.³

In late 2012, the Nest thermostat came out with advanced features, such as schedule learning, remote access, motion sensing and eco-feedback. It was introduced to the market and received

³ <https://nest.com/>

a great deal of media attention. The Nest gained the spotlight due to promises of convenience and energy savings. It is an interesting example of advancement in digital technology for the home since, prior to the Nest, there had been no major changes in the basic thermostatic controls in the previous sixty years (Peffer et al., 2011).

While the Nest thermostat has been received positively by early adopters, and generally perceived to be a huge improvement in terms of usability, there remains a need to understand how users will interact with “smart” features. For example, the Nest thermostat tracks user behavior and uses the information to project future heating and cooling patterns. If a user makes the same kind of changes repeatedly – i.e. raises the temperature two days in a row – then the system starts adjusting the thermostat the same way going forward (i.e., by adding an entry to the daily schedule to make the same change at the same time on future days). If the user’s changes were prompted by an aberration or a one-time event, (for example, when there were guests staying over in the home) and are not intended as long-term changes, the automated adjustment would cause more inconvenience as it may not be as easy for users to manually undo the learned pattern to meet their thermal comfort level once the need for adjustment disappears – e.g. the guests leaving. More sophisticated learning is possible, but as systems become more complex, it becomes more difficult for people to understand how they work and to predict what they will do (Edwards & Grinter, 2001).

The Nest provides an excellent opportunity to study the user experience of living with an intelligent system in the wild, particularly one that seeks to learn and adapt to consumers’ behavior in order to help people to save energy. Our interest in this literature review is to understand general issues related to the integration of intelligent systems into the home.

In the following chapters, we describe our research studies that aim to better understand the lived experience of an intelligent device for managing home energy use. Before we turn our

attention towards the Nest, a novel mass-market thermostat that utilizes machine learning, sensing, and networking technology to control home heating and cooling systems, we revisit the challenges that we addressed in first sections of literature review.

These challenges for designing intelligent systems for the home informed and guided our research questions and approaches in our studies that consist this dissertation work.

With advancements in computing, users are starting to interact more with intelligent and autonomous systems in their daily home environments, and this trend will almost certainly increase. However, using and interacting with these intelligent systems is a challenge for novice users. The foregoing literature review highlighted key challenges.

Users' lack of understanding: It is difficult for the user to understand how adaptive systems work when these systems gather multiple types of implicit data and utilize algorithms to operate themselves (Edwards & Grinter, 2001). In turn, users do not know how to control the systems to make them work as desired. Previous research shows that it is crucial for users to develop sound mental models of their systems to reduce the mismatch between their expectations and a system's actual capabilities. Sound mental models—or functional, how-to-use, as well as structural, how-it-works, knowledge—help users develop trust in their systems, which in turn fosters continued use of those systems.

Need for soliciting user input: It is difficult for a system to make accurate inferences about users' status and nuanced contexts based on limited data the system gathers from sensors and other data sources (Edwards & Grinter, 2001). Therefore, previous research indicates that users' input is necessary to provide proper and necessary data to the system in order to improve its performance and prevent malfunctions.

Messiness of home environment: Peoples' everyday lives in the home are full of unpredictable and nuanced situations and events (changing routines, situations, preferences and expectations). There is a mix of diverse individuals with different characteristics (different levels of technical aptitude, varying preferences and interest) living in heterogeneous homes. Thus, it is necessary to better understand the unique characteristics of a home environment in order to design intelligent systems that fit into peoples' daily lives.

Designing for sustainability: An area where intelligent systems promise to deliver great benefits is in the control of HVAC systems. Research into eco-feedback has focused on ways to provide information to people about their resource usage in order to motivate them to change their usage patterns. However, there has been little evidence that obtaining information reliably causes people to take action or change behaviors (Strengers, 2011). Predictive heating control, which uses sensing and machine learning to try to learn the occupancy patterns of a house's residents in order to automatically adjust the temperature. While these systems have shown promise in limited field trials, there remains a need to understand how such intelligent systems works in everyday home environments.

Taken together, the challenges we discussed in this literature reviews demand more research to inform the design of intelligent systems in the home.

CHAPTER 3.
LEARNING FROM A LEARNING THERMOSTAT:
LESSONS FOR INTELLIGENT SYSTEMS FOR THE HOME

INTRODUCTION

With advances in computing, everyday systems and devices in the home are becoming more connected, automated, and intelligent. This trend follows the trajectory of the “smart home” that has been forecasted and researched in the HCI and Ubicomp communities for the past two decades. This vision describes a home which seeks to adapt to its inhabitants and respond to their informational and comfort needs (Weiser & Brown, 1997), and there is increasing evidence that the vision is poised to become a reality. Many home appliance manufacturers are introducing new generations of digitally enhanced home appliances, which promise to reduce manual work, operate efficiently on behalf of users with little or no user intervention, and provide new types of information which were not available previously.

Managing home energy consumption represents a particularly rich domain for smart, domestic technologies, and has been the focus of numerous research projects (e.g., (Froehlich, Findlater, & Landay, 2010; Gupta, Intille, & Larson, 2009; Scott et al., 2011; Strengers, 2011)) as well as commercial offerings. In late 2011, the Nest thermostat was introduced to the market and received a great deal of media attention. The Nest represents an intriguing phenomenon for study, as it is the first mass-market thermostat in the U.S. to feature machine learning. The Nest’s learning promises to generate a personalized heating and cooling schedule that will promote comfort, energy savings, convenience, and more enjoyable interaction. Studying the

adoption and use of the Nest, then, provides an excellent opportunity to study the user experience of living with a ‘smart’ domestic appliance in the wild, particularly one that seeks to learn and adapt to consumers’ behavior.

Previous research on the user experience of smart, adaptive home technology has mostly been conducted in laboratories (e.g.,(Intille, 2002; Kidd et al., 1999)), or with prototypes in experimental settings (e.g., (Gupta et al., 2009; Scott et al., 2011)). As mainstream domestic technologies become smarter and more complex, more research is required to better understand the real use and adoption of such systems in the context of everyday life, where different individuals and families reside and behave. In order to better understand real-life, long-term experience with the use of such ‘smart’ digital technology in the home, we studied households that had installed a Nest. Using the Nest as a lens, we draw on our in-depth examination of users’ experience living with a smart thermostat to inform the design of intelligent systems for the home more broadly.

Our study findings provide valuable insights into how people perceive, use, and interact with intelligent systems, and what challenges lie in making intelligent systems work in real homes. In particular, we saw that people were surprised and frustrated by the Nest’s inability to distinguish between routine behavior (that the Nest ought to remember) and temporary adjustments (that it ought to forget). More generally, users also struggled to understand what the Nest was attempting to learn about them and how it was using its acquired knowledge to control their home’s temperature. In addition to leading to user frustration, these difficulties led to confusion about whether the Nest was actually helping users save energy—a goal that had originally motivated many of them to acquire the device in the first place.

Based on our analysis of these observations, we derive three promising avenues for future research on intelligent home systems: *exception flagging*, *incidental intelligibility*, and *constrained engagement*.

RELATED WORK

Even as the full realization of the smart home vision remains elusive, a number of studies have sought to understand the opportunities and challenges of the smart home through examining interaction with existing home technologies and prototyping future environments.

Programmable digital technologies such as VCRs, thermostats, and set-top boxes have been present in typical homes for many years, and their adoption and use have been studied fruitfully (e.g., (O'Brien, Rodden, Rouncefield, & Hughes, 1999; Rode, Toyne, & Blackwell, 2004)). More extensive forms of home automation have been pursued in small communities of users, however, these communities have been dominated by highly-engaged hobbyists and/or households wealthy enough to afford high-end professional installation and maintenance.

While studies of home automation adopters have yielded insights into the technology's barriers and benefits (e.g., (Brush et al., 2011; Mennicken & Huang, 2012; Takayama et al., 2012)), they have not provided insights into the mainstream user experience of *adaptive* home technologies that seek to learn about occupants' behaviors and preferences and change their operation accordingly.

Technical demonstrations of intelligent home environments have illustrated the feasibility and desirability of adaptive systems for the home (e.g., (Cook et al., 2003; Intille, 2002; Kidd et al., 1999; Mozer, 2005)), but few such projects have provided insight into the lived experience of occupants. A notable exception is Mozer's Adaptive House (2005), in which the researcher deployed adaptive systems in his own home across several months. An important conclusion from this study was that adaptive home systems needed to be designed to "educate" their occupants about their operation, so that they can act appropriately in the face of partial or

complete failures. This conclusion echoes Edwards and Grinter's observation that a fundamental challenge for smart homes is to offer advanced functionalities, yet still be manageable for users (Edwards & Grinter, 2001). When considering adaptive home systems that utilize sensing and machine learning, issues of intelligibility and control become central to the concept of "manageability" (Bellotti & Edwards, 2001; Edwards & Grinter, 2001). It has been noted that the gap between users' mental models and the actual system model can cause inefficient use, confusion, dissatisfaction, and abandonment of some features of the system (Tullio, Dey, Chalecki, & Fogarty, 2007). While extensive research has been done into how to design interfaces that render system behavior more *intelligible* (Bellotti & Edwards, 2001; Lim, Dey, & Avrahami, 2009; Stumpf et al., 2009), such research has yet to be pursued in the context of everyday domestic life.

A particular area of domestic technology use that has received attention within the Ubicomp and HCI communities is that of managing energy consumption. Given that 22% of the total energy consumed in the U.S. is used by home (EIA, 2011), such attention is clearly warranted. For the design of systems to promote sustainable lifestyles, numerous research projects have investigated eco-feedback systems as a way to promote greater awareness of energy use (e.g., (Froehlich et al., 2010)), which will, in turn, motivate people to save more energy. Strengers et al. (Strengers, 2011), however, pointed out that obtaining information did not always cause people to take action or change their behaviors. Previous studies (Peffer, Pritoni, Meier, Aragon, & Perry, 2011) investigated how people use their thermostat and concluded that poor usability of programmable thermostats is a critical barrier for their efficient use. Automation-based approaches have been proposed as a way to relieve the burden from users, implementing machine learning and sensing technology to automate system operation to some degree (Gupta et al., 2009; Scott et al., 2011). While these systems have shown promise in limited field trials,

there remains a need to understand how such ‘smart’ features will interact with users’ desire for control and predictability.

To better understand the lived experience of an intelligent device for managing home energy use, we turned our attention towards the Nest, a novel mass-market thermostat that utilizes machine learning, sensing, and networking technology to control home heating and cooling systems.

THE NEST THERMOSTAT

The Nest was released in October 2011 and was offered for sale for an initial price of US\$249; at this time, a standard programmable thermostat could be purchased in the U.S. for around \$30-\$40. At the time of its release, the Nest was considerably more advanced than other thermostats on the market, with novel features such as schedule learning, remote access, occupancy sensing, and eco-feedback. Here we describe the main features of the original (v1.0) Nest based on the description available on the Nest website ⁴.

The Nest features an attractive wall-mounted device, as well as smart phone and web-based control capabilities (Figure 3). In addition to providing access to the schedule and real time control, the web and phone apps provide the Energy History, which is the detailed history of when and how long the heating and cooling system ran. Additionally, the Nest includes a pair of intelligent features that utilize machine learning, and motion sensing: Auto-Schedule and Auto-Away.

⁴ <https://nest.com/>

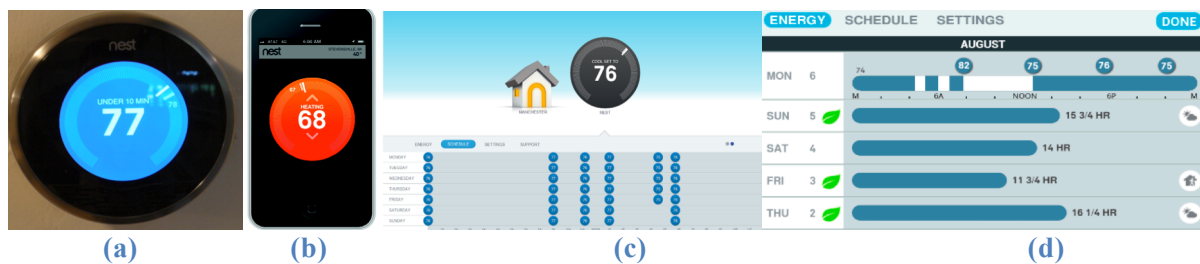


Figure 3. Users can control the Nest via the wall-mounted display (a), a mobile app (b), or a web app (c) The mobile and web apps provide access to Energy History (d).

Auto-Schedule: The Auto-Schedule feature automatically generates a schedule based on temperature changes users make. While the manufacturers of the Nest do not provide details of the algorithm, it can be said that the Nest takes about a week to generate its initial schedule and thereafter continually adapts the schedule according to users’ temperature adjustments. Users can manually revise the schedule via the wall-mounted device or through the web or mobile applications. Users can also turn off this feature and use the Nest as a regular programmable thermostat.

Auto-Away: The Nest has an embedded motion sensor on the wall-mounted unit that detects the movement of occupants within a certain range. If the Nest does not sense movement for about two hours, it goes into “Auto-Away” mode, which automatically adjusts the temperature to a user-defined level to avoid heating or cooling an empty home. Separately from the “Auto-Away” function, users can manually set the Nest to “Away” mode.

STUDY METHODS AND PARTICIPANTS

We interviewed 23 participants from nineteen households between February and September, 2012. All 19 households participated in interviews, and ten of them also participated in a diary study. All interviews were conducted by phone except one, which was conducted via video chat. Interviews lasted 45 minutes on average. During each interview, we asked participants

how they used their previous conventional thermostat compared to the Nest, as well as their overall experience and understanding of the Nest. While overall experiences and opinions were reported in the interviews, we learned more details about the individual situations, decision-making processes, and changes in users' perception and their understanding of the system over time from the diary study. For the diary study, we asked participants to report daily routines, changes made to the thermostat, and reactions to the Nest. We recruited participants using various methods, including email, Facebook, and Twitter messages, as well as contacting individuals who publicly posted about their experiences with the Nest. The resulting households were located in eight different states across the U.S. Demographic details are shown in Table 1.

Table 1. Summary of Participants

*** P13 submitted additional diary entries after her diary study completed. ** P16 and P17 who participated in an interview study in February 2012 participated in a follow-up interview in August 2012. PT: Programmable Thermostat, H: Heating, C: Cooling**

House hold	Number of Interviews (Diary entries)	State	Participant(s)	Adults (Child ren)	Occupation	Months of Nest usage by study end	Number of Nest and other thermostats
H1	2 (25)	MI	P1	3	Aerial Photographer	1 (C)	1 Nest
H 2	3 (21)	MI	P2	2 (1)	Interaction Designer	1 (H)	1 Nest
H 3	3 (4)	AZ	P3	3 (3)	Software Developer	1 (C)	2 Nests + 1 PT
H 4	3 (21)	AZ	P4	1	Software Developer	1 (C)	1 Nest
H 5	3 (12)	TX	P5	2 (2)	Software Developer	1.5 (C)	1 Nest + 1 PT
H 6	3 (7)	TX	P6	3	Municipal Program Professional	1.7 (C)	1 Nest + 1 PT
H 7	4 (20)	AZ	P7, P20	2	Software Developer, Accountant	1 (C)	1 Nest + 1 PT
H 8	1	MI	P8	2	Software Developer	1 (H)	1 Nest
H 9	1	MA	P9	2	Software Developer	1.5 (H)	1 Nest
H 10	1	CO	P10	2 (2)	Professor	2 (H)	1 Nest
H 11	1	CA	P11	2 (2)	Sales Manager	2.5 (H)	1 Nest
H 12	2 (19)	MI	P12	2	Web Designer	2.5 (C)	1 Nest
H 13	3 (37) *	MI	P13, P21	2 (1)	Interaction Designer, Cost Analyst	4 (H and C)	1 Nest
H 14	4 (21)	TX	P14, P22	2	Optometrist, Office Manager	6 (C)	2 Nests
H 15	2	CA	P15, P23	2 (2)	Software Developer, Stay at home mom	8 (C)	2 Nests
H 16	2 **	CA	P16	2	Software Designer	9 (H and C)	1 Nests + 2 PTs
H 17	2 **	MN	P17	2	Software Designer/Developer	9 (H and C)	1 Nest
H 18	1	TX	P18	2	Sales Manager	9 (H)	1 Nest
H 19	1	DC	P19	2	Marketing Consultant	Abandoned	1 Nest

In each household we studied, we identified the individual who was primarily in charge of thermostat control. This “primary” participant was generally the person who had taken the initiative to acquire and install the Nest. In 15 households, we interviewed only the primary participant. In another four households, we additionally interviewed a “secondary” participant, i.e., a Nest user who was not primarily responsible for integrating the Nest into the home.

Out of 19 primary participants, 18 were male and only one was female. Three of the secondary participants were female, and one was male. We endeavored to recruit a more balanced sample, but had difficulty finding women who had initiated the purchase of the Nest for their home, or who self-identified as the primary user in their household. In addition to being disproportionately male, our participants tended to be technically skilled and highly interested in new technology. The relatively high cost of the Nest meant that our participants were fairly affluent. While it would be valuable to study the voluntary adoption and use of an intelligent system like the Nest among a more diverse population, we were unable to recruit an appropriate sample given the timing and constraints of our study.

As noted, ten households participated in a diary study in addition to interviews. In all cases, the primary participant completed the diary entries. Eight of these ten households had obtained their Nest fewer than three weeks before they started the diary study. The remaining two households had been using their Nest for two and six months, respectively. Participants were asked to report diary entries for three weeks, and were interviewed at the beginning, during, and the end of the study period.

Participants submitted diary entries using Catch ⁵, a free web-based application that allows users to share pictures, text, and voice notes. We asked participants to describe their comings and goings, changes made to the thermostat, and reactions to the Nest. We provided example diary entries but did not provide prompt questions. Once a week, we asked participants to upload screenshots of the Nest schedule and the Energy History from their web or smartphone app. Occasionally we left comments on diary entries to encourage participation and to clarify what they reported in their entries.

Analysis

All interviews were audio-recorded and transcribed. The Nest schedule and energy history screenshots were reviewed and compared with the diary entries to find explanations for changes that were observed. The interviews and diary data were coded and analyzed using an iterative process of generating, refining, and probing the themes that emerged. Codes were initially drawn from research questions and then supplemented with those that emerged from the interviews and diary entries.

Our interest in this study was to understand general issues related to the integration of intelligent systems into the home. However, the Nest's users do not experience the 'intelligent' aspects of the Nest separately from its other features, so we sought to understand our data at multiple levels. At the highest level, we tried to understand users' overall experience with the Nest, including their judgments about its benefits compared to previous thermostats, changes to their household routines and thermal control patterns, and perceived improvements to their home's energy efficiency. This level serves as a backdrop to our analysis of the phenomena

⁵ Catch service was discontinued in 2013 and no longer available. <http://techcrunch.com/2013/07/31/evernote-competitor-catch-com-shuts-down-its-note-taking-apps-company-heading-in-different-direction/>

related to users' interactions with the Nest's intelligent features (principally the learning and sensing features)—including problems and successes encountered with these features, users' mental models of their operation, and users' subjective perception of the usefulness and desirability of these features.

From this it should be clear that it is not the goal of our study to proclaim the Nest a “success” or a “failure.” Stated differently, this paper is not intended to serve as an evaluation of the Nest, *per se*. Indeed, it is worth noting that from a commercial standpoint, there is ample evidence that the Nest is a reasonably successful product. From a viewpoint that is concerned with sustainability, though, we might assess success based on whether a product maximizes energy savings, or whether through automation or encouraging more energy efficient behaviors. Our particular concern in this paper is to gain insights into how to successfully deploy intelligent systems in the home. From this vantage point, we might look to a product like the Nest to assess how well users are able to take advantage of the system's advanced features, including its support for automatic scheduling and occupancy sensing. From these latter perspectives the Nest's success is decidedly less clear, as we shall see.

FINDINGS

Preliminary findings from seven of the households in our study were previously presented at the HomeSys workshop (Yang & Newman, 2012). Here we present a more detailed analysis based on the full set of 19 households, with special attention paid to participants' interaction with the Nest's intelligent features.

Based on our interviews and diary study, most of our participants were satisfied overall with the Nest, due in large part to the huge improvement over previous thermostats they had owned. So, first, as a way to set the context, we will describe the positive aspects of participants' experience of the Nest, namely increased engagement and greater awareness of energy usage

patterns. We discuss the particular features that changed our participants' interaction with the Nest as compared to conventional thermostats. Next, we will focus on the issues related to the Nest's intelligent functions, such as automatic scheduling and occupancy sensing, followed by a discussion of practices that emerged for dealing with these functions' shortcomings. Finally, we discuss the consequences of these shortcomings by considering whether the Nest led to energy savings.

Improved design leads to greater engagement

Participants found the Nest to be far more enjoyable to use than the thermostats that had been replaced. This perceived improvement derived largely from the elegant industrial and interactive design of the device and its remote control applications. Many participants liked the Nest lighting up as they passed by it, appreciated the intuitive graphical interface, and enjoyed being able to simply open their laptop or tap on their phone to control their thermostat.

For example, P22 was reluctant to change the temperature setting of her previous thermostat because *“it was really confusing to use.”* Instead of raising the temperature when she was uncomfortable, she would wear a sweatshirt at home, even during the summer. However, with the Nest, she found it easy to adjust the temperature: *“I love that it's so easy to track ...from your phone what the temperature is in our house. ...That way we look online and we're like, oh, we're not going to be here for the next five hours, and the air conditioning is on. We can change it.”*

Most participants also found the Energy History useful. It allowed some participants to remain engaged and make informed decisions, like P14: *“It kind of keeps me engaged on it. I think the engaging process of the machine is probably part of the reason why the energy savings come in because you pay more attention to it and you make sure it's running properly.”*

The learning system fails to understand user intent

While the interactive features, graphical interface, remote control, and energy usage information were all received positively and contributed to increased user engagement, participants had a different experience with the ‘intelligent’ aspects of the Nest, such as schedule learning and occupancy sensing.

When we first interviewed P16 in February 2012, he said that his Nest worked well and seemed to understand his desired comfort level. However, when we interviewed him again in August 2012, he was considering uninstalling the Nest. He found the learning was not successful and he was not satisfied with the changes the Nest had made to the schedule. Controlling the Nest was difficult for him, as the system continued to learn his temperature changes without recognizing the situations or intent behind his inputs: *“I’m not really happy with it anymore. The problem is, it is too controlling and not enough adaptive to our immediate needs. ...I had a pregnant daughter [visiting], and she doesn’t like hot weather, so we turned it down for her. Once you turn it down, then it learns that, and it says, “Okay, you’re going to want to do this every day.” It just becomes a very complex thing to adapt. ...It makes assumptions, and I don’t like the assumptions, and I can’t train it to make different assumptions. I feel like I’ve lost control over it. ...It only is able to see ...the clock schedule, and we don’t live by the clock.”*

Participants who were actively managing the temperature according to changing situations tended to have more problems, as the Nest could not detect *why* the user was setting different temperatures. It therefore could make erroneous assumptions about their intent, ultimately making unwanted changes to the temperature schedule.

While some participants felt that the Nest was overly eager, others felt it was not sufficiently sensitive to their input. P13 described his Nest as ‘*arrogant*,’ feeling that it would do whatever

it thought was right, regardless of his attempts at control. He wanted the Nest to follow his directions: *“There might be settings that we can decide to make it less arrogant? ...If I set in the evening to 75, then I want it at 75 and definitely for this night, ...I decided I want it 75. Don't turn it back to something else.”*

The system's behavior is hard to understand

The fact that the Nest often failed to recognize the reason behind temperature changes the user made was compounded by the fact that participants had trouble understanding how the Nest interpreted their input when creating a schedule and how the Nest sensed their movement or occupancy.

For example, P7 thought, *“Everything else [about the Nest] was straightforward but learning.”* He was uncertain about how much data were necessary to input for the Nest to create a schedule. He wondered whether changing the temperature every hour would confuse the Nest and how long it would take for the Nest to learn a new pattern. He lived with two other people and was curious about the impact of multiple adjustments.

As participants did not understand how the intelligent features work, such as Auto-Schedule and Auto-Away, they had difficulty to make the Nest work as they desired. P2 expressed his confusion about Nest in a diary entry: *“It's unclear to me whether [the learning] is done or if it is continuing to learn patterns. ...I'm also not sure of the time resolution of the Away calculation. ...Does it resume the regular schedule as soon as someone's presence is detected, or can it predict this event in advance if the pattern of home/away is regular enough? The very minimal Nest instructions do not discuss these decision-making parameters, but basically ask for trust, (perhaps before trust is earned).”*

In an interview, P2 said, *“Without knowing very much more about the parameters, I don’t really expect it to do that effective of a job in matching the schedule I prefer. Doing the schedule manually seems to be the easier course.”*

Another intelligent feature most participants expected to help them save energy was Auto-Away. Participants expected Auto-Away would save energy when they are not at home. However, many participants felt that Auto-Away was not working accurately. P4 wrote in his diary that Auto-Away turned on while he was at home: *“2:10 PM: While working, it was getting increasingly warm. Didn’t know what was going on. I checked on temp and noticed that it was at 80ish degrees. Set temp back down to 73 at the thermostat. Turned off Auto Away functionality.”*

After this entry, P4 walked past the Nest once every hour for the next six hours even though he had turned off Auto-Away. He wanted to make sure the Nest knew he was there and he was uncertain if turning it off would solve the problem. A week later, he regretted disabling Auto-Away after he found the A/C was working all day when he was not home. Regardless, he kept Auto-Away turned off because he suspected that it would work inaccurately again if he turned it back on.

Another participant, P16, who had the Nest stuck in “Away” mode, expressed his frustration: *“I would like to see it work. It just wasn’t working for us. ...The Nest is doing its own [thing] and doesn’t tell you what it is doing. It just doesn’t. So you really don’t know. ...It’s very hard to do anything but what it wants to do quietly.”*

In P14’s case, he speculated that Auto-Away stopped turning on because he was telling the Nest that he was actually at home when it turned on Auto-Away: *“[Auto-Away] was not turning on as much as I wanted it to. That was a problem that I was trying to address over the last couple months. ...The Auto-Away ...had turned on in the first couple weeks when we didn’t*

want it to. ...It's really easy you just go up and you press it and tell it that you're still there. I think we may have done that too much. ...[T]hat's probably why the Auto-Away stops turning on."

Months later, he concluded that the location of the Nest was not ideal for detecting people's movement.

Participants were surprisingly reluctant to give up intelligent features and displayed a willingness to work around some of their shortcomings. However, efforts to 'fix' the situation or 'take back' control in most cases were either discouraged or undermined by the participants' lack of understanding of how the learning actually took place. P17, whom we interviewed after nine months of Nest usage, found that the Nest stopped learning his temperature settings after he deleted all the unnecessary temporary changes the Nest remembered. He did not understand why and thought it was his fault: *"I thought when I started using the Nest that it was going to do a better job of tracking my changes, ...and just automatically updating the schedule. It was for a while and then it stopped. I haven't figured out why yet. Everything you see on that schedule now I entered manually, which I didn't ...have to do that. I don't know what happened. ...It's just stopped doing something that it should be doing and that's probably my fault ...because it was working up until I deleted the settings."*

Users found ways to work with the 'intelligence'

Despite the limitations of the Nest's learning, participants came up with strategies that could take advantage of certain intelligent features and make the Nest work better for them.

Overall Experience with Learning

More than half of the participants (P1, P3, P7, P11, P12, P14, P15, P18 and P23) reported the Nest remembered their temperature settings 'well enough.' Many of them kept a regular schedule or maintained consistent temperature settings. When these participants found the

learning was not successful or they did not like the adaptive changes the Nest had made to the schedule, they were willing to modify the schedule manually. They were content with the Nest since the improved graphical user interface and remote applications made it relatively easy for them to control it. Other participants (P2, P5, P8, P9, P13, P16, P17 and P21) found the learning did not work well and some were even annoyed by the adaptive changes the Nest had made to the schedule. In both cases, the learned schedule needed to be revised by participants, but as long as the Nest did not make drastic changes to the schedule they set manually, they still kept the learning function active.

Correcting the schedule

Several participants felt that the Nest merely memorized their adjustments. They were disappointed when the Nest appeared to simply remember their input rather than do something more ‘intelligent’ like generate a good average schedule. P9 found that the schedule the Nest generated (Figure 4) was “*probably more crazy and detailed than it really need[ed] to be.*” P8 also revised the schedule so that the Nest would not be making small changes: “*I just went through and sort of cleared it up so that it won't be making all those little changes all the time.*”



Figure 4. P9’s Nest schedule showed frequent temperature changes on certain days. Time is plotted on the X-axis and weekdays are plotted on the Y-axis. The dots show the temperature setting at the particular day and time.

Three days after he installed his Nest, P2 found that an initial schedule had been learned. Three days after that, he determined that the learned schedule was unsatisfactory, so he modified it. He posted before and after screenshots in his diary, which are shown in Figure 5.

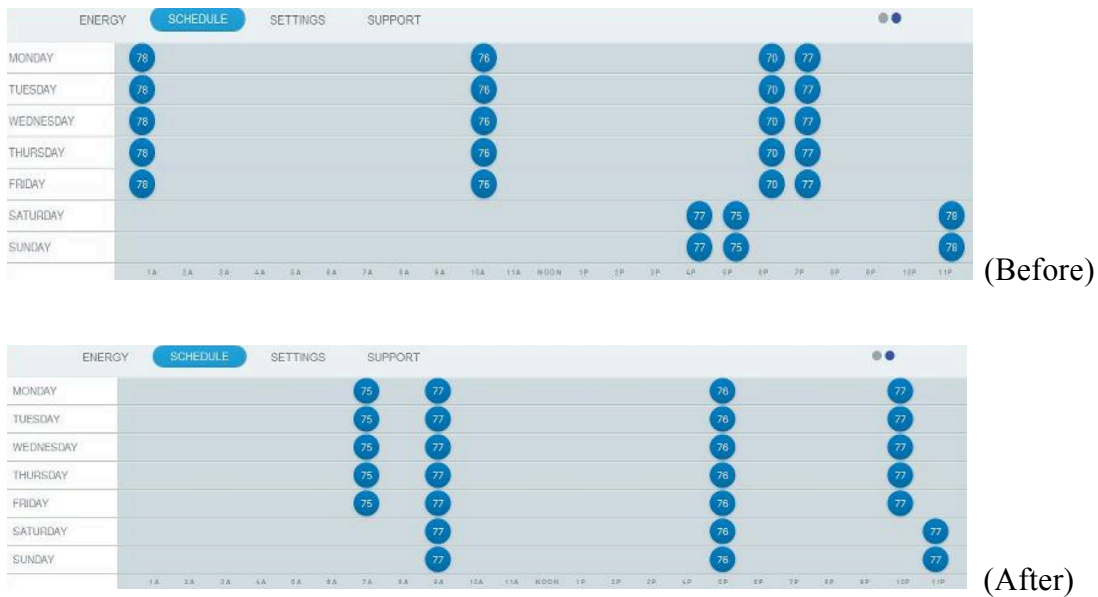


Figure 5. P2 posted screenshots of his schedule before and after he modified it.

Teaching and guiding the learning

Once several participants realized the Nest’s machine learning limitations, they changed the way they interacted with it. For example, P17 intentionally gave limited input for the Nest to memorize. He described how he managed the Nest schedule once he concluded that the Nest simply memorized his input: *“For the first week we had it, I was adjusting it all the time, because it was fun to do. But then after about a week, I looked at the schedule that it had memorized and it was crazy, it was all over the map. So, I erased the whole schedule and we started again. And at that point, basically, not more than three times a day.”*

Monitoring

With the Nest creating the initial schedule and updating it as the patterns changed, many participants said that they monitored the schedule the Nest was generating. Several participants actively checked to see if it was reasonable. They reviewed the Energy History to look for any abnormalities in how the heating and cooling system had been running. When participants

noticed an improper or inefficient temperature setting, they made adjustments and deleted the undesirable temperature setting.

The Nest did not clearly lead to energy savings

Most participants expected the Nest to be helpful for energy savings. However, except for some participants who said that they were very conscious about energy savings, many participants were uncertain about whether the Nest saved energy. P11 said, *“I will not say if it saved me any electricity at this point.”* P9 was not sure if he saved money with the Nest, explaining his doubt: *“In reality, it might be that I played with Nest so much, it cost me an extra 300 bucks.”* As we described, Auto-Schedule and Auto-Away each displayed shortcomings and therefore may not have directly contributed to participants’ energy savings.

Users pursue convenience and comfort

The expected benefit of remote access is to enable users to control their thermostats when they are away from home. Interestingly, most participants used the remote control at home frequently, sometimes more often than the wall-mounted device. Participants said that having the remote control is convenient since it allowed them to check their thermostat more frequently and make changes without even getting up. For example, P9 used the remote control in his bed: *“If I wake up and I'm freezing, I'll just grab the iPad next to the bed and crank up the heat. Then I haven't even gotten out of bed yet.”*

Learning may not generate an energy efficient schedule

Participants initially expected that the Nest would be smart enough to figure out the ideal schedule for the heating and cooling system to achieve comfort and save energy. However, several participants (P2, P8, P9, P13, and P16) found the Nest simply memorized their input, but it did not generate an energy efficient schedule. P16’s Nest generated a higher heating temperature setting than he would have set, *“It seems like it stays warmer longer than what we*

would've done it if we left it purely manually.” P10 intentionally set up a schedule manually since he did not want the Nest learning an undesirable schedule based on his family members’ input. He believed that his family members set the temperature unnecessarily high or low, and often forget to adjust the temperature before going out.

The Nest’s learning might have created a less-than-ideal schedule, since it learned participants’ patterns of temperature adjustment and many participants were likely to make adjustments for comfort rather than efficiency. Several participants (P2, P3, P5, P13 and P16) explicitly stated that they preferred comfort to energy savings, and thus did not change their behavior to save energy after getting a Nest. As mentioned earlier, many participants found it easy to change the temperature via remote control. With a conventional thermostat, they might well have stayed with a less comfortable schedule they had initially programmed due to the difficulty of changing it. If users make capricious changes and do not monitor how they affect the schedule, the Nest schedule may stay inefficient.

Auto-Away failure led to wasted energy

Another intelligent feature most participants expected to help them save energy was Auto-Away. Participants expected Auto-Away would save energy when they were not at home. Several participants reported that they did not obtain much benefit from it since Auto-Away often either turned on when they were at home or did not turn on when they were not at home. From our diary study, we observed that four households out of ten had occasions when they wasted energy since Auto-Away did not turn on while they were away. For example, two months after P13 installed the Nest, she discovered that Auto-Away had not been working for several days. She wrote in her diary: *“Auto away feature is broken!!! It no longer senses when we are not home. That was my favorite thing about the nest, so this is annoying. ...It happened during the hottest week too. My A/C was on a LOT without needing it! Aargh...”*

She felt that she could not rely on Auto-Away to function properly and created a schedule to prevent the Nest from cooling the house during the day.

Users' motivation is the key to savings

Despite the intelligent features of the Nest that promised energy savings, such savings seemed to largely result from participants' motivation and engagement with monitoring their Energy History and making necessary changes to save energy. Many participants who were actively monitoring their thermostat usage were confident that they saved more energy by making a conscious decision to change the schedule to a more energy efficient setting. For example, P12 mentioned that one day he checked the Energy History and noticed that the air conditioner was running ten or more hours a day. He raised the temperature setting by one degree and saw the air conditioner ran only six or seven hours a day after the change. He was okay with being less comfortable because it was his "conscious decision." However, we also observed that participants' excitement and engagement faded over time. Once most participants settled down with the Nest schedule, they paid less attention to the schedule or the Energy History.

To sum up, we found that participants were most satisfied with the Nest's user interface and remote control; the intelligent features of the Nest, such as Auto-Schedule and Auto-Away were less successful. We also observed new practices of user control emerged to address the Nest's limitations. It is notable that participants' workarounds reflected their willingness to employ intelligent features despite their shortcomings; even so, users had trouble determining whether they were saving energy.

DISCUSSION

At a high level, the findings just reported will not be surprising to many readers who are conversant with the issues surrounding interactive intelligent systems. The fact that systems struggle to understand human context and intent, and that users cannot orient their actions with

system appropriately without an adequate understanding of how the system operates have been often discussed in the literature. Indeed Suchman (2006) classically identified a pair of key challenges for the design of *interactive machines in general* as being 1) the machine's limited access to a user's actions and circumstances of the user and 2) the user's difficulty in recognizing the machine's constraints. Clearly these challenges are magnified when discussing *intelligent* interactive systems, as the system seeks to learn *patterns* of user behaviors, preferences, and decision making, and users seek to understand and control *complex* and *malleable* system behavior.

It would be tempting to conclude that our findings, then, are simply a reflection of poor design decisions by the Nest. In the versions of the Nest that we studied, the subsystem that learned user preferences was only capable of detecting one aspect of user behavior (control changes) and the system provided no convenient mechanisms for indicating which inputs ought to be remembered by the system and which ought to be forgotten. Additional relevant dimensions of user behavior such as occupancy, the presence of particular household members and guests, and household activity levels as well as contextual dimensions such as humidity, external temperature, and sun exposure—all of which could be relatively easily sensed and incorporated into a predictive model—were simply not included, and there was no mechanism for compensating for their absence.

Additionally, the Nest made no attempt to explain or account for its behavior, leaving users little information with which to build an effective mental model. We argue, however, that the issues uncovered in our study reflect deeper challenges in designing intelligent systems for the home that cannot be addressed by collecting more data, building better models, or applying existing approaches to making system behavior intelligible.

Bridging the intention gap: Exception flagging

Suchman's challenges articulate a fundamental gap between what computing systems can sense and the user's intentions. That is, no matter how many sensors we include or how elaborate our models become, there will be gaps in the system's knowledge. Our data supports the view that some amount of human behavior is unpredictable, some preferences change, some routines are unstable, and some contingencies are too rare to form a pattern. Yet, intelligent systems can provide benefits by automating the aspects of life that are predictable, enduring, stable, and regular. A key design challenge, then is to elicit input from users to help the system differentiate the data that represents regular, stable preferences or behavior from input that does not. Existing approaches to correcting system inference focus on giving feedback on the system's output (e.g., (Kulesza et al., 2009; Stumpf et al., 2009)) or on eliciting more and higher quality input from the user (e.g., (Dey, Rosenthal, & Veloso, n.d.)). However, neither of these approaches seem well suited to the type of system represented by the Nest. Such systems are characterized by mostly invisible output (system-initiated control changes will only be noticed after the fact in most cases, and in many cases may not be noticed at all), and user input is not solicited, but rather passively observed.

The promise of the Nest that it will learn users' preferences based on their behavior and build a suitable schedule is clearly appealing to end-users. It is unclear whether users would be able or willing to endure a special "training mode" of any duration, or whether they would be willing to inspect system outputs and provide feedback with any regularity. The nature of domestic life and the relative unimportance of thermostat control would suggest that neither approach would be appealing. An alternative approach would be to develop interactive techniques that require intentional user input only in the case of exceptions. Techniques for *exception flagging* would allow implicit user input to be collected and used for learning in the normal case, but allow users to identify, or *flag*, exceptional inputs (i.e., inputs that should not be learned), triggering

the system to ignore such inputs when building models and making predictions. While such a mechanism would be simple to implement technically, it would present challenges in terms of interaction design, as it is not clear that users would always be able to articulate at the time of execution when an action was exceptional. It might be easier to identify exceptions in retrospect, but it is not clear how or when it would be best to ask users to review previous inputs and label them appropriately. We believe the further research will be required to develop and test effective techniques for eliciting exception labels from users across different domains in the smart home.

Bridging the understanding gap: Incidental intelligibility

A different but related challenge is helping users to understand how the system is interpreting and acting upon the data it receives from users. This challenge (loosely captured by Suchman's second challenge noted above) has been studied extensively under the topic of "intelligibility," which covers user interface techniques that seek to help users understand the behavior of complex, often intelligent, systems. A major focus of intelligibility research has been on providing interactive explanations for how the system works and why it behaves in certain ways (e.g., (Kulesza et al., 2009; Lim et al., 2009; Stumpf et al., 2009)). Such approaches to intelligibility, however, assume that the user has a conscious interest in understanding the system, and is willing to invest time in doing so.

Our observations of Nest users suggest that the desire to understand the system arises infrequently (only when something goes wrong), and that there is little motivation for exploring or developing one's understanding of the system's learning capabilities as an independent activity. While users may not see the value in understanding the system's behavior, it would clearly be beneficial to the system's operation—and ultimately to the user—if they did. It would also allow users to head off problems of misunderstanding before they become dire, thus reducing frustration at a later date. Thus finding ways to increase users'

understanding of how the system learns and makes decisions is a valuable goal, even if the users might not place a high value on it.

Moreover, as we saw in our study, users were able and willing to adapt their behavior based on even a partial understanding of how the Nest operated. Such co-adaptation has been observed among users of configurable systems (Mackay, 2000) and collaborative systems (Orlikowski, 1992), and perhaps ought to be expected among users of intelligent systems as well.

Supporting co-adaptation requires helping users gain a practical understanding of the system's operation. To foster system understanding without requiring explicit interaction dedicated to the task, we suggest that intelligibility ought to be delivered opportunistically, in small pieces commensurate with the relatively small, occasional, incidental interactions that characterize users' interactions with the Nest. Such *incidental intelligibility*—interaction elements that increase users' understanding of the system's intelligent behavior embedded in the tasks they consciously seek to accomplish—could build understanding that would help users orient their behavior over the long term while not asking users to attend to learning how the system “thinks” as a discrete task.

Widening the interaction: Constrained engagement

Both exception labeling and incidental intelligibility demand users attention, even if that demand is minimized as much as possible. Conventional thermostats, both manual and programmable, are designed largely with the goal of reducing demands on user attention to nearly zero, in accordance with both longstanding cultural trends in home automation and, coincidentally, with Weiser's visions of disappearing and calm computing (Weiser & Brown, 1996).

As Rogers points out, however, a strong stance on making computing invisible runs counter to visions of “smart” technologies that learn about and understand their users (Rogers, 2006).

While Rogers goes on to suggest that UbiComp move away from its emphasis on smart systems and towards the design of engaging experiences, we suggest that home control systems like the Nest present a venue where intelligence and engagement ought to co-exist. Specifically, we note that the effective application of intelligence to problems like temperature control will require user engagement in the form of (at least) periodic, thoughtful input from the user along with consideration of and monitoring of system outputs. People know about the situations (e.g., Mary is pregnant and likes to be warm) and plans (e.g., we are having five guests over for dinner in an hour) that impact the behavior observed by the system and so it is important to not just provide mechanisms for input but to engage users to interact the system.

Such engagement, however, must be dramatically constrained, given that the interaction between user and system is necessarily sparse and peripheral yet continuous and long-lived. Assuming that we are evolving towards a world in which users engage with dozens if not hundreds of intelligent services like the Nest, a challenge faces UbiComp researchers to come up with ways of designing technologies that engage but do not overwhelm—a goal that we refer to as *constrained engagement*.

Here, actually, we feel that the Nest got it mostly right. Many participants enjoyed having more control over their thermostat. Indeed, we observed that new practices of user control emerged to address the Nest's limitations. It is notable that participants' workarounds reflected their willingness to employ intelligent features despite their shortcomings. Moreover, energy savings we observed with the Nest are did not come from automation such as auto-learning or auto-away, but resulted from participant's engagement to save energy. The Energy History feature increased awareness about energy consumption, supported informed decisions, and motivated green behavior, mainly by making it easy and enjoyable to monitor system performance. Also, ease of use enabled users to put their thoughts into action. By providing a baseline of user engagement through attractive and thoughtful design, systems like the Nest

can more easily gain needed access to the user for confirming inputs, explaining outputs, and supporting the process of productive co-evolution.

CONCLUSION

In this paper, we present an account of the user experience of adopting an intelligent thermostat drawn from interviews and diary study of 23 participants regarding managing the temperature in the home and energy saving as a result. Our study results reveal challenges and opportunities of intelligent systems, particularly those that utilize machine learning and motion sensing. Based on our findings, we provide a set of design implications for intelligent systems for the home.

CHAPTER 4.

MAKING SUSTAINABILITY SUSTAINABLE: CHALLENGES IN THE DESIGN OF ECO-INTERACTION TECHNOLOGIES

INTRODUCTION

The smart home is here. Long envisioned by HCI and Ubicomp researchers, the promise of a home that can learn its occupants' needs, desires, and behaviors — and adapt itself appropriately — is being realized. Networked digital devices and services are being manufactured and marketed in ever-increasing numbers, performing a variety of different roles in the home including entertainment, health, security, and home automation. These new capabilities bring great potential, but also great concern — are “smart” devices going to make our lives easier, more productive, or more enjoyable? Or are they going to bring a new set of frustrations, expectations, and responsibilities that will outweigh their possible benefits?

An area where smart home devices promise to deliver great benefits is in the control of home heating, ventilation, and cooling (HVAC) systems. HVAC control is an important domain from the perspective of environmental sustainability. In the United States, for example, residential HVAC systems account for roughly 50% of all household energy consumption, which equates to about 9% of the nation's total energy budget ⁶. Moreover, residential HVAC systems are not

⁶ <http://www.eia.gov/consumption/residential/data/2009/>

operated efficiently by their users (Peffer et al., 2011), leading to unnecessarily wasted energy. A number of researchers have investigated ways to improve the operation of HVAC systems. Research into eco-feedback (e.g., (Froehlich, Findlater, & Landay, 2010; Pierce, Schiano, & Paulos, 2010; Strengers, 2011)) has focused on ways to provide information to people about their resource usage in order to motivate them to change their usage patterns. Another approach that has been investigated is predictive heating control (e.g., (Gupta, Intille, & Larson, 2009; Koehler, Ziebart, Mankoff, & Dey, 2013; Scott et al., 2011)), which uses sensing and machine learning to try to learn the occupancy patterns of a house's residents in order to automatically adjust the temperature.

Both eco-feedback and predictive control can be seen as approaches to a broader problem we call *eco-interaction*, by which we mean the study of interaction between humans and energy-consuming systems with an eye towards minimizing energy use while preserving an acceptable level of user-perceived benefits. Eco-interaction includes eco-feedback and predictive control, but also includes the design of control interfaces, infrastructures, and basic functionality required to facilitate user interaction.

In this paper, we seek to inform the design of future eco-interaction systems by investigating users' experiences with the Nest Learning Thermostat (hereafter "The Nest"), a commercially available smart home device. The Nest combines elements of eco-feedback and predictive control with networked remote control to allow users to create a custom heating and cooling schedule that matches their preferences and helps them save energy.

The work in this paper builds on an earlier study (Yang & Newman, 2013) that looked at users' initial experiences with the Nest, including problems encountered with the Nest's learning and sensing capabilities, and users' strategies for dealing with the Nest's limitations. Here we look

at a different aspect of users' experiences with the Nest, namely how the features of the Nest changed users' interactions with HVAC systems in the home over both the short and long term.

We do this by first examining how people interact with “conventional” thermostats — i.e., the ubiquitous manual and programmable thermostats that can be found in the vast majority of North American homes (Peffer et al., 2011). We then report users' initial experiences upon acquiring a Nest by re-analyzing the data originally collected for (Yang & Newman, 2013) from the perspective of the Product Ecology Framework (Forlizzi, 2008). The Product Ecology Framework allows us to more easily see changes in consumers' perception of and interaction with a novel product like the Nest, and to tease out different threads that impact the user experience. Finally, we report a new follow-up study that was conducted to learn about how their interaction with the Nest had changed after owning it for 12-21 months.

Our study found that the Nest impacted users' pattern of HVAC control, but only for a while. During the first few months after installing a Nest, many of the users we studied were more engaged, interacted with their thermostat more, and sought ways to save energy more actively than did users of conventional thermostats. After a period of time, however, the engagement with the Nest, along with the frequency of interaction, diminished and users' interactions settled into patterns that resulted in missed opportunities for energy savings. Based on these findings, we identify a set of implications for the design of eco-interaction systems.

BACKGROUND AND RELATED WORK

Our work builds upon and informs existing approaches to promoting energy savings, principally eco-feedback and predictive control. As our particular focus in this work is on the changes in use of eco-interaction systems over time, we also draw upon prior studies of long-term interaction.

Eco-Interaction

The goal of eco-feedback is to promote greater awareness of energy use, which could in turn, motivate people to save more energy (work in this area is extensive; (Froehlich et al., 2010) provides a survey). However, there has been little evidence that obtaining information reliably causes people to take action or change behaviors (Strengers, 2011). Rather, considerable motivation and engagement on the part of consumers is required for eco-feedback to lead to behavior changes. Moreover, even when people are aware and motivated, it can be difficult to effectively control their systems. Previous studies showed that poor usability was a significant barrier for the efficient use of programmable thermostats (Peffer et al., 2011).

Given the challenges of persuading people to change their behavior, Pierce *et al.* (Pierce et al., 2010) suggested designing interfaces to “nudge” people to save energy by default, thereby reducing the need for consumers to make conscious decisions or enact behavior changes. Hazas *et al.* (Hazas, Brush, & Scott, 2012) take this argument further, proposing that technology-centered approaches, rather than user-centered approaches, offer the greatest promise for saving energy. Following this thread, work in predictive control seeks to reduce or even eliminate the role of user choice in controlling HVAC systems by automating temperature adjustments based on occupancy predictions. By tracking occupancy patterns using GPS (Gupta et al., 2009; Koehler et al., 2013) or RFID and motion sensing (Scott et al., 2011), it is possible to build reasonably accurate models that can predict occupancy and make sure the house is heated or cooled to a desired temperature when people are home and a less energy-intensive level otherwise. Systems in this category have been built and tested in limited deployments. It remains to be seen what issues would arise in a more general deployment with people who vary more widely in terms of geographic mobility, schedule predictability, tolerance for error, and desire for control.

The Nest Learning Thermostat⁷ features an attractive wall-mounted device, as well as smart phone and web-based control capabilities (see Figure 6).



Figure 6. Users can control the Nest via the wall-mounted display (a), a mobile app (b), or a web app (c). © Nest Labs

In addition to providing access to a schedule and the ability to control the temperature in real time, the web and phone apps provide a detailed Energy History, an eco-feedback display showing when and how long the heating and cooling system ran and providing feedback about whether the day's performance was energy efficient. A green leaf icon appears when users set the Nest to a temperature that is considered energy efficient by the Nest's algorithms⁸.

Additionally, the Nest includes Auto-Schedule and Auto-Away, intelligent features that utilize machine learning and motion sensing, respectively, to implement a form of predictive control. In contrast to the systems mentioned earlier that seek to predict occupancy, the Nest's Auto-Schedule feature generates a schedule based on temperature changes that were previously made. While the manufacturer of the Nest does not provide details of the algorithm, it claims

⁷ The description of the Nest presented here is based on the version of the Nest our users had at the time of the study, as described at <http://www.nest.com> (Accessed: 2012-09-24.).

⁸ <https://nest.com/downloads/press/documents/efficiency-simulation-white-paper.pdf>

the Nest generally takes about a week to compute an initial schedule and thereafter continually adapts the schedule based on users' temperature adjustments. Users can also use the web-based control interface to manually revise the schedule created by Auto-Schedule, and Auto-Schedule can also be turned off. The Nest has a motion sensor embedded in the wall-mounted unit that detects the movement of occupants within a certain range. If the Nest does not sense movement for some time, it goes into Auto-Away mode, which automatically adjusts the temperature to a user-defined setback level to avoid excessively heating or cooling an empty home.

As a precursor to the work presented in this paper, Yang and Newman (2013) studied households that had installed a Nest, focusing on users' initial impressions of the Nest and their experiences with the Nest's "smart" features, namely Auto-Schedule and Auto-Away. Whereas (Yang & Newman, 2013) drew on early-stage usage experiences to inform the design of usable intelligent systems for the home, our goal here is to identify challenges and opportunities in the design of eco-interaction systems for long-term use.

Understanding Product Use and Long-Term Interaction

Numerous approaches exist for understanding how and why technological products are acquired, adopted, and used. Models such as the Technology Acceptance Model (Venkatesh & Davis, 2000) and Orlikowski's duality of technology (Orlikowski, 1992) help explain how particular features of a product or system, along with contextual factors such as user expectations, social dynamics, and temporal trajectories impact how the product is integrated into the life of an individual or collectivity.

In this work, we draw on the Product Ecology framework to look specifically at how the Nest changed users' relationship to their HVAC system over both the short and long term. The Product Ecology is a theoretical framework that describes social product use (Forlizzi, 2008).

It is informed by social ecology theory, which is broadly concerned with the dynamic relationship between an individual and the physical and social environment. From the Product Ecology viewpoint, the product is the central unit of analysis. The functional, aesthetic, symbolic, emotional and social dimensions of a product, combined with other units of analysis, or *factors* in the ecology help to describe how people make functional, social, and symbolic relationships with products. These include the product; the surrounding products and other systems of products; the people who use it, and their attitudes, disposition, roles, and relationships; the physical structure, norms and routines of the place the product is used; and the social and cultural contexts of the people who use the product (Forlizzi, 2008). The Product Ecology has been used in the long-term study of adoption of semi-autonomous products in the home to understand how they change interactions in the household (e.g., (Forlizzi, 2007)).

Other work has looked at long-term interaction with interactive products. Several papers, for example, have looked at the change in user satisfaction and perception over time (e.g., (Karapanos, 2013; Kujala, Roto, Väänänen-Vainio-Mattila, Karapanos, & Sinnelä, 2011)), but these did not focus on how users' interaction patterns changed or what impact those changes had on outcomes enabled by the product, such as comfort or energy savings. Our work contributes to both the literature on Product Ecology and long-term interaction by investigating the changes in users' relationship to and usage of a novel device over both the short and long term.

METHODS

We conducted two qualitative studies, one with manual and programmable thermostat users and another with the Nest thermostat users. Both studies consisted of a diary study augmented by semi-structured interviews.

Study 1: Conventional thermostat study

We conducted a three-week diary study with 16 participants between September and December 2011 in order to understand how people use their thermostat to manage their thermal comfort in their homes. We recruited participants using personal networks, mailing lists, and the snowball sampling method. Eight participants had manual thermostats, and the other eight had programmable thermostats. Our participants lived in six different states in United States. Participants made daily diary entries for three weeks and participated in two interviews. We employed a diary study to capture participants' day-to-day heating and cooling control and to avoid limitations of interview data such as participants' inaccurate memory of their actual behaviors and perceived comfort for each day. Each participant reported the arrivals, departures, and sleep times of their household members, how they felt about their comfort, and what changes they made to their thermostat and why. The initial interview focused on participants' household routines, general thermostat control practices, and thermal comfort preferences. The exit interviews elicited additional details related to diary entries.

Study 2: The Nest thermostat study

To understand users' experiences with the Nest, we drew from a re-analysis of data collected for a previous study (Yang & Newman, 2013) and a new set of 15 follow-up interviews conducted with members of the households that participated in the original study. For the initial study reported in (Yang & Newman, 2013), 23 participants from 19 households who had purchased a Nest were interviewed between February and September 2012. Ten of these also participated in a three-week diary study and two additional interviews, which took place during and after the diary entry period. Diary entries described participants' comings and goings, changes made to the thermostat, and reactions to the Nest that they had, positive or negative. Participants also submitted periodic screenshots of the Nest's Energy History and schedule.

Follow-up interviews were conducted between August and September, 2013. Fifteen participants from nine of the original households agreed to participate for the follow-up (two household members participated in the follow-up who did not participate in the initial phase). Participants sent us updated screenshots of their Nest schedule and the Energy History prior to the interview and we asked them about their long-term experience and use of the Nest.

Data collection and analysis

In total, we conducted 90 interviews and collected a total of 508 diary entries. All interviews were audio-recorded and transcribed. The conventional thermostat study data was coded and analyzed using an iterative process of generating, refining, and probing the themes that emerged. The data from the initial-phase and follow-up Nest usage studies were analyzed using the Product Ecology Framework (Forlizzi, 2008). We chose to employ the Product Ecology Framework in order to investigate how dimensions of the product influenced thermostat control activities that took place around the use of thermostat. Our interest in this study was two-fold: first, to better understand how product features influence users' interaction with a thermostat and second, how users' interaction with a semi-autonomous thermostat changes over time within the home. We were also interested in aspects of behavior that supported energy savings.

In the following sections, we draw on the conventional thermostat study data to describe people's current practices with regard to thermostat control as well as existing problems. Then, we discuss how the Nest changed the practices, interactions, and relationships associated with the thermostat, how it addressed existing problems, and what new issues it presented. Finally, we present the breakdowns that occurred with the Nest over time. In the Discussion, we reflect on these findings to extract a set of implications for the design of future eco-interaction systems.

In the findings below, we refer to participants by thermostat type and subject number, for example, PT1 is the first participant interviewed who had a programmable thermostat. MT is used for a manual thermostat. For the Nest study, we use same participant codes, P1-P23, that were used in (Yang & Newman, 2013), adding P24-P25 for additional follow-up study participants. We indicate whether source was an interview (I) or diary entry (D), and note the number of months the participant had been using the Nest at the time.

COMMON PROBLEMS WITH THERMOSTAT CONTROL

In our study of conventional thermostat usage, we observed common problems in participants' thermostat control patterns that echoed those described in previous work (Peffer et al., 2011). With manual thermostats, people often forget or find it inconvenient to manually adjust temperatures to increase energy savings. While programmable thermostats ought to make it easy to reduce energy consumption, people find it difficult to program their thermostats due to usability flaws (Peffer et al., 2011). Here, we look beyond the well-documented usability problems with existing thermostats to shed light on more fundamental reasons that efficient management of thermostats is challenging. Specifically, we show how practices surrounding thermostat control are tightly related to people's comfort and convenience as well as frequent changes in routines and schedules in daily life.

People do not use a setback temperature.

Making effective use of a setback temperature—i.e., an energy-efficient temperature setting to be used when the house is unoccupied — is one of the most important steps people can make to reduce the energy used for heating or cooling their homes (Peffer et al., 2011). Many of our participants did not consistently use a setback temperature, and cited various reasons. Many participants wanted to avoid the long wait time until the house heated or cooled upon returning home, while others simply forgot or were unaware of the potential energy savings.

People do not use schedules.

Interestingly, some participants did not even try to figure out how to program their devices. PT1 did not use a schedule even though it meant he had to frequently wait up to two hours for the house to cool down to his desired temperature. He said, *“There’s a button called PRG, which I figure is probably for Program, ... I was too lazy. I never really bothered to figure out how to use it.”* Rigid scheduling options were another reason that participants avoided scheduling. Many programmable thermostats offer limited options for scheduling, allowing only a “weekday” and a “weekend” schedule, each with limited preset times when temperature changes can occur (e.g. morning, day, evening, and sleep time). This rigidity made it difficult to effectively set a schedule for more complex and nuanced daily routines. More importantly, inflexible scheduling options combined with a difficult scheduling process hinders participants from accommodating frequently changing schedules and temperature preferences.

People fail to reassess existing control patterns.

Some participants kept non-optimal temperatures that were “more comfortable” than they actually needed. Several participants programmed their thermostats for the season and stayed with the schedule throughout the season. PT2 referred to a programmable thermostat as *“a little more maintenance-free”* than a manual thermostat, as she only needed to adjust the schedule twice a year: *“I kind of do an assessment, if you will, before winter starts and before summer starts to make sure my temperatures are kind of where I want them to be.”* However, we believe that this “set-and-forget” approach will not be optimal for energy saving because both weather and people’s schedules change frequently during the season.

In addition to failing to reassess their schedules, participants used temperature settings that were not optimal for *either* saving energy *or* achieving thermal comfort. This was revealed accidentally, for example, when MT1 at one point forgot to change the temperature back to 70°F from 64°F as she usually did upon arriving home during the winter. She only realized her

oversight when, two days in a row, her husband came home and said; “*It’s kind of cold in here isn’t it?*” She wrote in her diary: “*I guess that people do adapt to the temperature and can tolerate a wide range (more than we initially recognize).*”

In the next section, we describe how our participants with the Nest used and interacted with it differently from those participants with manual or programmable thermostats.

PRODUCT ECOLOGY ANALYSIS

As described earlier, we used the Product Ecology framework to analyze how the Nest impacted people’s interactions with their heating and cooling systems as mediated by their thermostats and whether and how this, in turn, affected behavior, roles and relationships in the household. We coded for three factors in the product ecology of the thermostat, *people*, *activities*, and *products*, with special attention to the dimensions of the Nest as a product (*functionality, aesthetics, symbolism, emotion, and social attribution*).

As a *product*, the Nest was well-received by most users, especially in terms of its *symbolic, aesthetic, and functional* aspects. Symbolically, the Nest was seen as a “cool,” “stylish” gadget that reflected its owners’ good taste and technical savvy. Most participants thought that the Nest was designed for anyone because it was easy to use. However, a few mentioned that it was designed for young, technically-savvy users because they felt they were not taking full advantage of all of the Nest’s features. It also held the promise of saving energy, which was a significant factor in many participants’ decision to purchase the Nest in the first place, and it reinforced self-images of the purchaser as an energy-conscious consumer.

All participants mentioned energy savings as one of their motivations for getting a Nest, but ultimately, they were most motivated by the perception of Nest being a cool, beautifully-designed product. *Aesthetically*, the Nest was seen as a huge improvement over the dull plastic thermostat it most frequently replaced, and it was praised for its appealing appearance. P9(I-

1.5m) expressed that “*just having something on the wall that’s not an ugly piece of plastic from [brand name] is also totally worth it.*” The smart phone and Web interfaces were similarly seen as elegant and attractive, leading to an overall positive *emotional* experience for most of our participants. While some participants used anthropomorphic language when discussing the Nest (e.g., talking about what the Nest “knows” or “thinks”), *social attribution* did not seem to be a dominant dimension in users’ experience.

The Nest’s novel *functionality* was, however, very prominent in participants’ minds. The Nest’s functional aspects, along with a heightened sense of engagement due in large part to the positive emotional response and changes in the relationships of household members vis-à-vis thermostat control, impacted the *activities* that users performed with and around the Nest.

INITIAL EXPERIENCE WITH THE NEST

Through our analysis, we noted several ecological changes that occurred among adopters of the Nest as compared to those using a traditional thermostat. Here we focus on two. We first provide insight into how key dimensions of the Nest (particularly *functionality, aesthetics, and symbolism*) led to greater engagement, which, in turn, led to increased awareness and interaction with the system. Second, we provide details about how changed interactions with the Nest impacted activities related to thermostat control and energy savings.

Increased engagement and awareness

Conventional manual or programmable thermostats were not seen as exciting to use. Conversely, several dimensions of the Nest combined to promote greater engagement, and more importantly, stimulated our participants’ interest in thermostat control and energy savings. The Nest’s novel features (*functionality*) such as the Energy History and its interactivity led participants to be more aware of their heating and cooling system. P15(I-8m)’s interaction with a thermostat control changed in the following way after getting a Nest: “*I’m much more in tune*

with what my heating and cooling systems are doing. I'm much more aware of their presence and their function. I know it sounds kind of silly because it's a heating and cooling system. But, before I just avoided. ... I only dealt with it when I had to, but now I just like to see what it's doing when I walk by."

The Energy History feature and appealing graphical interface motivated and facilitated to assess and improve their settings. P6(I-1.7m) had previously bought a programmable thermostat for about US\$20 and never programmed it during the four years he had it. He stated that he "*didn't want to bother with it*" because it was "*old technology.*" With the Nest, though, he explored the Nest's different functions and was motivated to optimize its schedule: "*I was sort of messing around with [the schedule] and then got excited about it. Then, suddenly I was adjusting all the temperature... It's really fun... It's almost a game like, 'OK, let me see if I can make it a little bit warmer on this day and try to save a little energy there.'*"

Here, we find that the Nest's *symbolic, social, and functional* aspects successfully engaged participants in performing tasks that were problematic for conventional thermostat users, namely using schedules and reassessing them for energy savings.

Changing practices and interaction with the Nest

Along with increased engagement and awareness, the Nest's functional aspects also impacted thermostat control activities in key ways, as we now describe.

Scheduling becomes an interactive and iterative task.

The Nest participants' scheduling activities became more ***interactive and iterative***, as compared to the tedium of conventional thermostats. Once the Nest generated a schedule based on its learning of a participant's input, many participants reviewed and revised their schedules, often repeatedly. Thus scheduling was not a one-time task but instead was ongoing, as the Nest automatically kept updating the participant's schedule over time and the participants kept

reviewing the changes. P18(I-9m) described how he interacted with the Nest's schedule and why the scheduling could not be solely left up to the Nest: "*Reading [the schedule] now, it says, 'On Wednesday 4:30, sent from my Nest thermostat' is when I put it at 75°F, but then it shows that Monday, Tuesday, Thursday, and Friday were set to 75°F because of the learning feature... I look at it occasionally to see why it's set like that... If it has added something in there that I didn't think... was something good, then I would change it back to something else.*"

For other participants as well, keeping an eye on the schedule became necessary once they realized that the Nest remembered temporary changes and added them to the schedule. During the diary study, P2(D-2m) "*noticed a few aberrations in the schedule.*" Once he found that the Nest quickly responded to temporary adjustments and made them part of a regular schedule, he monitored how the Nest was changing his schedule. P8(I-1m) said, "*I look at it [the schedule] every now and then to see if it has added something crazy in there.*"

Temperature control becomes more fluid and adaptive.

While participants found that conventional programmable thermostats offered limited options for scheduling, and thus, it was hard to accommodate temporary changes, many participants appreciated the fact that the Nest was **flexible and adaptive**. The Nest's functional aspects, such as Auto-Away, remote control, and flexible options for scheduling enabled participants to accommodate frequently occurring changes in their schedules. P5(I-1.5m) felt that he gained more control with the Nest and pointed out that the Nest's flexibility and adaptivity in scheduling actually allowed him to save more energy: "*The previous one wasn't very flexible so you were kind of at its mercy. You didn't have a lot of control over your energy usage ... [N]ow, I can a lot more proactively manage [my energy usage]... [The Nest] is very flexible ... [T]he best thing so far is being able to set the temperature from when out of the house, being able to set it Away. I like the Auto-Away and just the ability to manage it on a day-to-day hour-to-hour basis is helpful.*"

Establishing a setback temperature becomes easier.

P5(D-1.5m) further described in his diary that having remote access led him to employ a setback temperature. He wrote, “*Really what changes our behavior is setting it to Away or turning it off while we’re gone.*” The Nest’s ability to adjust the temperature based on occupancy and to control the temperature remotely enabled the participants to employ more flexible, temporary, and immediate temperature setback strategies. With Auto-Away and remote control, it was okay to forget to change the temperature before leaving. P13(I-9m) described having the remote control as freedom and empowerment: “*The freedom that the Nest gives you from having to... remember to turn it down. You’re empowered wherever you are to make those changes.*” Adjusting the temperature was no longer on her husband’s vacation to-do-list. In addition, they did not need to worry about wasting energy because they forgot to turn down the temperature. The couple P15(I-8m) and P23(I-8m) did not need to call their housekeeper to check whether she had adjusted the temperature after cleaning. Instead, they could check remotely and avoid wasting energy.

Monitoring the Nest emerges as a new task.

Supervising and monitoring the Nest emerged as a new task, as participants were intrigued by how it learned from their temperature adjustments and operated autonomously.

In addition to monitoring changes to the schedule as described above, many participants monitored their energy histories to see how the Nest was performing. While traditional thermostats do not provide means for users to see how their heating and cooling systems have been working, the Nest’s Energy History allows users to track how the Nest has been operating on a daily or weekly basis. A few participants drew on the Energy History to reassess whether their behaviors and existing schedules were aligned with their comfort and energy-saving goals. For example, P12(I-2.5m) noticed that the A/C was running ten or more hours a day, based on the Nest’s Energy History. After he raised the temperature setting by just one degree, he found

that the A/C ran only six hours or seven hours a day. By monitoring the Energy History, he found ways to save and experienced the impact of making minor changes to his schedule.

SETTLING INTO A ROUTINE

The changes in engagement and awareness described in the previous section did not, however, persist over the long term. Our follow-up study identified major changes in the use of the Nest that were apparent after a year. First, far less interaction with the product was cited. Over time, the device became mundane, and people rarely interacted with features like the Energy History and the schedule. Second, a decreased effort to improve energy performance was noted. While participants were initially interested in monitoring their energy use, they came to rely on the Nest's automatic functions more over time. In essence, participants' thermostat control practice changes — such as monitoring the Energy History and fine-tuning the schedule for energy efficiency — did not last over time. The effort and engagement required to maintain such behaviors was not sustainable.

Less interaction with the Nest over time

Our follow-up interviews revealed that by 12–21 months after installing the Nest, many participants did not remember when they last checked their schedules or energy histories. Many had stopped reviewing their energy histories, checking or adjusting their Nest schedules over six months or more prior. Only two participants explicitly said they kept paying attention to the Nest in order to save energy. Most participants in the follow-up study interacted with their Nest thermostats only to directly adjust the temperature. P1(I-13m) said, *“We don't interact with the Nest, really. ... We just use it as a regular thermostat.”*

Over time, the Nest schedule became static, and the system made few adjustments.

Accordingly, some participants became less concerned about what the Nest was doing. P4(I, 13m) said, *“The Nest takes care of all the changes in temperature... It doesn't require me to*

babysit it.” P6(I-15m) did not check the schedule or the green leaf anymore to see if the temperature settings were energy efficient: “I haven't checked the schedule or the green leaf because mainly, I think, the Nest has learned pretty well what we like so I don't really think about the thermostat too much anymore. Maybe, it's almost working too well because I don't think about looking for the green leaf. On the occasion that I do go up and adjust the thermostat, like let's say it just gets too hot or too cold, then I do look for the green leaf on those occasions. But in general, I don't interact with the thermostat that much. I don't even think about checking it anymore. It kind of faded into the background for me.”

As long as the Nest did not set the temperature abnormally, it did not call for our participants' attention, and he or she did not need to actively get involved in controlling or interacting with it.

Reduced effort in improving energy performance

During the follow-up interview, participants were asked to take a look at their current schedules and energy histories. To five participants' surprise, the Nest thermostats were not working as they had expected. For example, P4(I-13m) failed to reassess his default settings, resulting in wasted energy. He was away during the weekend, however, the Away temperature setting was set to 80°F, and the Nest was cooling the house during the weekend, as shown in Figure 7. P4's Energy History shows the A/C was running during the weekend while he was not at home (bottom two rows). When he left home, he set the Nest to 'Away' mode, which was set to 80°F, a temperature lower than his normal daytime setback of 85°F (shown in the upper three rows). P4 was surprised to discover that his system was wasting energy while he was away.

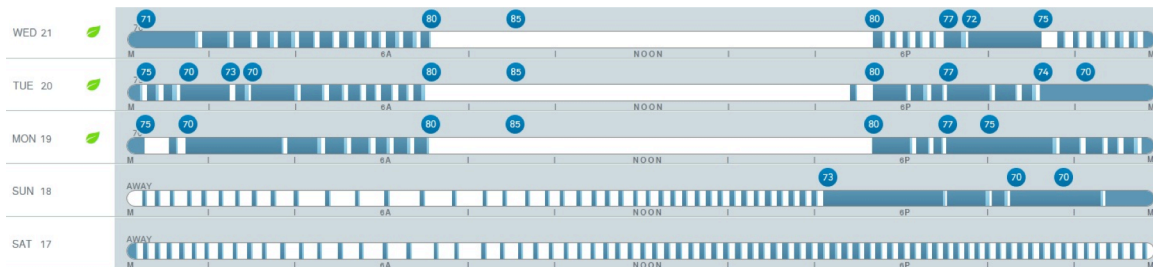


Figure 7. P4's Energy History shows the A/C was running during the weekend while he was not at home (bottom two rows). When he left home, he set the Nest to 'Away' mode, which was set to 80°F, a temperature lower than his normal daytime setback of 85°F (shown in the upper three rows). P4 was surprised to discover that his system was wasting energy while he was away.

When asked why he set the Away temperature to 80°F while he had an 85°F setback temperature, he answered, *“Oh, I have no idea. I think [the] Away [temperature] was already set at 80°F [about a year ago when he installed the Nest]. I just didn't change the setting. I just turned on the 'Away.'”* He continued, *“I guess left to my own devices, it would have stayed at 80°F. It makes sense to turn it to 85°F. I really didn't even pay attention to what it was at.”*

Reduced engagement with monitoring and over-reliance on the Nest functions caused surprises regarding its automation routines. P14(I-18m) trusted the Nest and rarely looked at his schedule or Energy History. He believed that the Nest was saving money due to its functions, such as Auto-Schedule and Auto-Away. However, during the follow-up interview, he found that the Nest was actually running during the weekend while he and his wife were away. He had checked the Nest on his phone and had seen that the Nest was in Away mode after leaving home on Friday. He thought the Nest would maintain the Away mode the entire weekend. To his surprise, the Nest somehow turned on and cooled the house, as shown in Figure 8.




Figure 8. P14's Energy History is showing that the cooling was on during the weekend when nobody was at home.

One couple, P15 and P23(I-20m), thought the Nest was adjusting the temperature autonomously as it understood their needs. P23(I-20m) said, *“It senses when we’re home and it knows what temperature we’d like it to be at various times of the day and so it adjusts it on its own.”* However, her husband P15(I-20m) found that one of their two Nests did not have a schedule even when its learning function was active. He was *“actually surprised”* since *“[he and his wife had not] noticed that it didn’t have a schedule.”*

Participants also felt that their schedules could be improved but usually lacked the motivation to do so. P24(I-14m) explained he could make two adjustments to save energy. First, he would create an additional setback temperature of 81°F at 1 p.m., which was earlier than the 81°F setting he already had at 3 p.m. *“because I know it gets above the temperature in the day before then so there’s no reason to keep it at 81°F until that late in the afternoon.”* He would also raise the 81°F setting at 3 p.m. to 83°F to save more.

Note that in all of these examples, had we not conducted follow-up interviews, participants would not have seen problems with the Nest. They all believed that their Nest thermostats were working as they had expected. These anecdotes call into question the desirability of having the Nest fade into the background even though this might be desirable from a customer satisfaction standpoint. The perceived promises of the Nest’s energy-saving features might have led participants to rely on the Nest and allowed them to be less active in controlling their heating and cooling systems.

Changes in behavior were not sustained

A number of factors could explain why people's interest in interacting with the Nest faded over time and why changes in energy-saving behavior were not sustained. The first was simply that the novelty effect wore off, a phenomenon commonly seen with technology products (Hekkert, Snelders, & Wieringen, 2003). For example, P6(I-1.7m) was explicitly motivated to save energy, but it took less than two months for his excitement to wane: *"When we first got it, it was really exciting. A new gadget, we're trying to figure out what it can do. Now we're sort of used to it ... [T]he novelty ... has kind of worn off."*

Second, many participants began to rely on the Nest. As it turns out, they might have overestimated the Nest's capabilities. P1(I-13m) and P25(I-13m) had not checked the schedule and the Energy History for about a year. Nevertheless, P1(I-13m) explained why he believed the Nest was doing its job: *"I just have faith in it. I assume that it's doing its job, but I don't really know. I haven't checked up on it."* When we asked what caused him to have such faith in it, he answered: *"Well, because it's a computerized thing, and it's fancy and it lights up when you put your hand near it."* The Nest's features, such as recognizing them when they passed by, might have played a role in the trust that the participants felt with the system.

Third, participants often became forgetful in their interactions with the Nest. As long as the Nest did not drastically change a participant's schedule and maintained the user-guided/revised schedule, most would not bother to reassess it in order to make it more efficient. This may have led to wasted energy and missed opportunities to save energy and money. For example, both P14(I-18m) and P24(I-14m), who mentioned finding ways to improve their schedules, forgot about the idea because they were doing something else at that time. P24(I-14m) gave more fundamental reasons, *"I have to do so little adjusting to the Nest that I did not remember to doing with it ... We have not had unusual activity on other fronts, so I was not motivated to check on it and make the necessarily change."*

Finally, the motivation to save energy might not be strong enough to overcome the inertia of existing behavior. After learning that the Nest was working while she was not home, P22(I-18m) mentioned that she might start *“taking a more active role, at least checking [the Nest] before [she] leave[s] for the day,”* and she even expressed a desire to learn how to check the Nest schedule and use other Nest features. However, she quickly admitted that she would not actually do this.

DISCUSSION

In our study of Nest usage, we observed that participants’ thermostat control practices changed immediately after the installation of a Nest. Many of our participants actively tried to save more energy when they first got the product. However, after time passed, their engagement with saving energy decreased. In many cases, participants showed a tendency to trust the Nest and neglect active monitoring or decision-making for energy savings as long as they did not notice any problematic issues.

In our follow-up interviews, several participants (four out of nine households) were surprised that the Nest was not operating as they had originally believed. It is possible that participants might never have discovered incidents of energy loss and stayed with the current schedule when the ability to set a more optimal one existed. This highlights a central tension with the Nest—its success from a user experience standpoint (it performed well enough that people felt they did not need to pay it much attention) impeded its success from a sustainability standpoint (users’ trust and resulting inattention led to missed opportunities for energy savings).

Designing eco-interactions

As we saw in our findings, some degree of active involvement by participants happened early on but needed to happen iteratively over time to sustain or increase energy savings. A challenge for designers, then, is to preserve the benefits of system autonomy and automation

while facilitating interaction to promote and sustain users' engagement for achieving desirable energy efficiency over time. In the remainder of this section, we reflect on the tensions elucidated by our study and propose a set of design implications for eco-interaction systems, emphasizing the design of mixed-initiative systems that invite participation and reflection with the goal of saving energy at home.

Use mixed initiative to balance competing concerns.

In discussing the potential energy savings that can be obtained using the Nest, Nest Labs notes that the Nest's goal is not "solely energy savings," and that the Nest "places a high priority on the user's comfort"⁹. This prioritization of user comfort and control is reflected in the fact that Auto-Schedule attempts to learn the pattern of users' manual temperature changes rather than occupancy or some other implicit signal indicating intent behind users' inputs. As noted, in Yang and Newman's initial study of Nest usage (Yang & Newman, 2013), participants were not always sure that learning from their temperature adjustments would result in an optimal schedule for energy savings as, oftentimes, temperature adjustments were made to improve comfort. Even though many of our participants had a high-level goal of saving energy, the more immediate goal of achieving comfort would often win out. It follows that a "learning" thermostat that receives all of its initiative from users could end up optimizing for comfort rather than savings, resulting in undesirable outcomes over the long run.

An alternative design might be to create a mixed-initiative system (Horvitz, 1999) wherein the system primarily pursues the goal of energy savings and the user is free to pursue their goal of immediate comfort within certain system-defined bounds. The general notion of a mixed-initiative thermostat was proposed in (Keyson, de Hoogh, Freudenthal, & Vermeeren, 2000;

⁹ <https://nest.com/downloads/press/documents/efficiency-simulation-white-paper.pdf>

Koehler et al., 2013), and here we extend this notion to articulate a clear goal for the system: to maximize energy savings while respecting users' expressed comfort preferences and desire for control. To balance these needs, it will be necessary for the system to *push* information, requests, and suggestions to the user rather than allow the user to initiate all interactions. As we saw, users' initial engagement with the system, which included active monitoring of system performance and fine-tuning the Nest schedule, waned after a few months. Thus, over the long term, a thermostat with an agenda may need to be assertive in getting the user's attention. The question remains: how can a smart device assert and pursue its goals without annoying or alienating the user?

We propose that designing spontaneous, enjoyable interactions to prompt users to engage with the system sporadically over time might be a valuable direction to explore. Here we emphasize that sustaining user engagement while not requiring constant attention is an important goal. As an example of a possible opportunity, our participants enjoyed seeing the Nest light up as they passed by, briefly attracting their attention. Perhaps such moments could be leveraged to alert users to situations that require attention, or to remind them to re-engage and reassess existing settings. However, merely alerting users to problems or reminding them to reassess may not be enough, as the challenge of converting information into action remains. We address this challenge next.

Bridge the gap between awareness and control

Horvitz suggested that mixed initiative systems should ultimately leave the user in control (Horvitz, 1999). In particular, *allowing direct invocation and termination* of system services and *employing socially appropriate behaviors* (e.g., informing users of actions that will affect them), systems can maintain users' trust while providing significant value. In the case of the Nest, we saw situations where users recognized an opportunity for savings but were unable to follow through and take the required action.

Designers should consider ways to generate concrete plans for increasing energy savings that leave users in control but are easy for users to implement. As an example, consider a recommendation for an improvement to the user's schedule that appeared on the home screen of a thermostat control app or in an email. This recommendation could include an option that allows the user to implement the recommended change instantly. To help users decide whether such recommendations ought to be followed, systems could further provide *eco-feedforward* messages or visualizations to convey the projected impacts of the recommended changes. Providing actionable recommendations along with information about the projected benefits of those recommendations would enable systems to suggest courses of action that align with system goals while allowing users to stay in control. As a system does more prompting to assist with goal setting, it may also prompt a person's curiosity and motivation. Designers should thus provide opportunities for deeper interaction and reflection alongside the simple courses of action presented for easy invocation.

Reframe interactions around reflection and reassessment

In addition to drawing users' attention to potential energy saving opportunities, it would also be valuable to maintain lightweight engagement between users and the system on an ongoing basis. Smart devices like the Nest are not as 'smart' as users might expect. Limitations of current intelligent systems require users to monitor and remain involved in order to maintain and improve performance (Yang & Newman, 2013). When the Nest was actively creating a schedule early in the study, participants were more curious and engaged with the Nest. Participants willingly paid attention and felt their interaction with the Nest was necessary. It seems likely that by the end of the first few of months of interaction, participants had taught the Nest a reasonable set of temperature changes that reflected their household routines and preferences, yet saved energy where possible by using a scheduled setback temperature or Auto-Away.

The normal ebb and flow of the household, however, combined with changes in people's needs, caused changes in ideal heating and cooling schedules. Our Nest participants often failed to negotiate making these changes, leaving the thermostat schedule as it was and reducing the potential to save energy. To overcome this, eco-interaction systems need to *stimulate reflection and reassessment*. Doing so requires rethinking the interaction design to emphasize reflection and reassessment rather than control and convenience. As an example, designers might consider designing ways to periodically perturb the user's routine interactions with the system. Mechanisms could be designed for the system to periodically initiate the evaluation and reassessment of the schedule, perhaps by expiring schedules after a period of time or asking users to choose between an existing schedule and a more efficient one proposed by the system.

People thought the Nest, with its clean aesthetic appearance and friendly UI, worked well enough. However, this was problematic because the Nest did not necessarily seek out optimal control patterns or adjust its control patterns to changing circumstances in the home. The resulting control patterns were often not as efficient as those that could be achieved by human intervention, yet users did not know when and how to enact changes to improve performance.

For the successful adoption of eco-interaction systems like the Nest, and to achieve the goals of energy savings for early adopters and the general population, we need to design more cooperative, collaborative and coordinated interactions between semi-autonomous systems like the Nest and their users, and figure out how to sustain those interactions over time. We suggest that tighter feedback loops between eco-interaction systems and their users can help them to develop and maintain more sustainable practices while users achieve their desired benefits such as comfort and energy efficiency. Employing mixed-initiative designs, providing actionable recommendations, and stimulating reassessment may be starting points for designing more effective eco-interactions in the future.

CONCLUSION

The availability of smart home devices offers great promise in multiple arenas, most significantly in reducing consumers' energy usage through more efficient HVAC control. To inform the design of eco-interaction technologies—i.e., technologies that help people save energy while meeting comfort goals—we investigated how the Nest Learning Thermostat situates in the home and affects human behavior. We conclude that such systems can be better designed to better project their benefits, and to help users realize their goals in saving energy. We hope our research on eco-interaction will inspire the community to understand and improve upon products that work on behalf of people in everyday life.

CHAPTER 5.

ASSESSING A RECOMMENDATION SYSTEM FOR ENERGY-EFFICIENT THERMOSTAT SCHEDULING

INTRODUCTION

Many believe that technology can play a key role in helping people consume less energy, but there are competing approaches to achieving this goal. Eco-feedback techniques inform people about their own energy usage in order to empower and motivate them to make better decisions and consume less energy (Froehlich, Findlater, & Landay, 2010). For example, some systems tell people how much they are consuming (Jiang, Dawson-Haggerty, Dutta, & Culler, 2009) and in some cases present a breakdown of how they are consuming it (Ranjan, Griffiths, & Whitehouse, 2014). However, information alone is not always enough. Some energy saving actions are too complex or time consuming for people to do them regularly.

Therefore, automation techniques take action to save energy on behalf of the user. This approach has been explored extensively for things like lighting control (Mozer, 1998), thermostat control (Lu et al., 2010), and even vehicular route navigation (Ganti, Pham, Ahmadi, Nangia, & Abdelzaher, 2010). While automated systems have shown promise in limited field trials, recent studies found evidence that showed that autonomous systems did not work as successfully as expected in the real world (Shih, Han, Poole, Rosson, & Carroll, 2015; Yang & Newman, 2013).

In this paper, we examine a different approach that we call “eco-coaching”: giving personalized suggestions for specific actions that would reduce wasted energy. Eco-coaching

systems assist users with energy savings, but leave users in control. We propose that eco-coaching should go one step farther than eco-feedback: it should not only provide feedback on energy consumption in the past, but also leverage user behavior in order to identify waste and recommend actions to prevent energy waste. However, eco-coaching should stop one step short of automation: it should identify actions that can reduce waste but should not take them on behalf of the user. Using mixed-initiative (Horvitz, 1999), eco-coaching extends approaches to balance system autonomy and user control in thermostat scheduling (Koehler, Ziebart, Mankoff, & Dey, 2013; Lu et al., 2010; Pisharoty, Yang, Newman, & Whitehouse, 2015; Yang, Newman, & Forlizzi, 2014).

In this paper, we build upon the previously published work by focusing on evaluating the features of eco-coaching that led to the success of ThermoCoach in saving energy.

ThermoCoach (Pisharoty, Yang, Newman, & Whitehouse, 2015) was designed to provide eco-coaching for thermostat control. It first monitors users' behavior and energy use patterns over time and identifies areas for improvement to reduce energy waste. It then generates and emails personalized and actionable schedule recommendations to users and makes it easier for them to take action. To assist users in balancing their energy savings and comfort goals, it presents eco-feedforward messages to provide information about projected savings and comfort expectations for each recommendation. Finally, ThermoCoach allows users to customize the recommended schedules it provides. This is especially useful in cases where the system could not identify particular needs or situations, such as preferences for sleep temperature or the presence of pets.

To evaluate the impacts on energy savings, a 12-week field deployment was conducted, comparing the energy saving outcomes of ThermoCoach with two other approaches: manual programming (i.e., Users program their thermostat schedule manually) and automatic

scheduling (i.e., Users use a smart thermostat to automatically program their thermostat schedule). As previously reported, results indicated that eco-coaching saved 4.7% more energy than manual programming and 12.4% more energy than automation (Pisharoty et al., 2015).

The present paper is based on a new analysis of interview data collected as part of the field deployment study. In particular, we focus on reactions to and experience of eco-coaching from the user's perspective, and examine how ThermoCoach's eco-coaching features influenced users to save energy.

We found that the eco-coaching approach 1) made it easier for users to implement an effective thermostat schedule, 2) supported user agency in negotiating trade-offs between energy savings and comfort, 3) facilitated learning different scheduling strategies as well as weighing pros and cons of different options, and 4) challenged users' beliefs about how well they were doing. These outcomes, in turn, were successful in getting users to employ and experiment with more efficient setback strategies.

While our initial results are promising, we also find room for improvement, especially in supporting users to assess recommendations. In particular, evaluating the fit and performance of the recommendations is important for building user trust, thereby increasing acceptance and maximizing benefits of eco-coaching recommendations.

BACKGROUND AND RELATED WORK

We summarize key approaches in designing systems to promote energy conservation. First, eco-feedback aims to help users to be aware of their energy use and make informed decisions to save energy. Second, automation tries to reduce the workload for users by automating tasks that users need to do manually. Third, mixed initiative approaches seek to balance system capability and human control to address shortcomings of previous approaches.

Eco-feedback displays data to inform users of their consumption of various resources, such as electricity, gas or water, and thus seeks to motivate users to change their behavior. While eco-feedback has been shown to increase awareness of resource consumption, several studies that investigated everyday practices of consumption (Pierce, Fan, Lomas, Marcu, & Paulos, 2010; Y. A. A. Strengers, 2011; Yang et al., 2014) found that obtaining information did not actually trigger people to take action or change behavior. Strengers warned against a common assumption that the eco-feedback approach holds, depicting the user as “a resource man” who makes rational choices and acts accordingly when provided information (Y. Strengers, 2014).

Automation approaches seek to relieve user burdens by automating users’ tasks. For example, research on occupancy-based thermostat control (e.g., (Gupta, Intille, & Larson, 2009; Scott et al., 2011)) seeks to detect and/or predict when the home is unoccupied so that the thermostat can be set to an optimally efficient level, and return to the occupants’ preferred comfort level before they are likely to return. The Nest thermostat takes a different approach by learning users’ adjustments and automatically generating a schedule instead of modeling the occupancy pattern (“Nest | The Learning Thermostat | Home,” n.d.). With the advancement and availability of wearable and smart home sensing devices, Huang et al. (Huang, Yang, & Newman, 2015) investigated the possibility of developing new models of thermal comfort to generate a real-time predictive comfort model (Feldmeier, Paradiso, & others, 2010) instead of or in addition to the occupancy pattern or temperature adjustment pattern.

While these systems have shown promise in limited field trials, there remains a need to understand how such “smart” features will interact with users’ desire for control and predictability. Recently, several studies investigated users’ lived experience of smart systems and observed evidence that autonomous systems did not work as successfully as expected in the real world (Shih et al., 2015; Yang & Newman, 2013). These studies showed that systems often fell short of anticipating or responding to dynamically changing everyday life situations.

Some users became frustrated and even abandoned the smart devices when they did not understand how those devices worked or why they did not function as expected (Brush et al., 2011, 2011; Lazar, Koehler, Tanenbaum, & Nguyen, 2015; Mennicken & Huang, 2012; Yang & Newman, 2013).

In response to smart system shortcomings such as lack of understanding of nuanced and dynamically changing situations in the real world and users' loss of control, several projects have proposed mixed-initiative approaches (Horvitz, 1999) to balance system autonomy and user control in thermostat scheduling (Koehler et al., 2013; Lu et al., 2010; Pisharoty et al., 2015; Yang et al., 2014). More recently, several research projects have designed and evaluated agent-based systems that provide suggestions to conserve energy by utilizing dynamic pricing and renewable energy (Bourgeois, Van Der Linden, Kortuem, Price, & Rimmer, 2014; Costanza et al., 2014; Simm et al., 2015). The notion that systems could recommend energy saving actions has been proposed before, including in the home environment (Intille, 2002). These systems process a large amount of data such as weather forecasts, peak loads on the power grid, and renewable energy generation to figure out the best time to perform energy-consuming activities. Then, these systems prompt users to use appliances such as laundry machines or dishwashers at times when the electricity pricing is cheaper or when they can use renewable energy.

Costanza, et al. (Costanza et al., 2014) conducted a simulation study in which a washing machine agent allowed users to book a time to do their laundry to save cost. They found that some participants who had structured routines for doing laundry found it easier to fit the system into their existing practice. However, other participants found it challenging to plan for their laundry as they usually ran their washing machine as needed. Bourgeois, et al. (Bourgeois et al., 2014) employed a series of interventions such as energy feedback, proactive suggestions and direct user control to support users with photovoltaic solar energy generation to plan their

laundry on daily basis. Bourgeois et al. found that participants perceived proactive suggestions to be more useful than feedback messages, although they did not necessarily follow the suggestions. Simm, et al. (Simm et al., 2015) designed a system to forecast renewable energy generation for a local community. They found that participants were able to make use of the information and that some participants actively shifted times when they did laundry or dishwashing to maximize their use of green energy.

In this paper, we study ThermoCoach (previously published in (Pisharoty et al., 2015)), which provides multiple personalized and actionable suggestions for thermostat scheduling. It differs from the aforementioned systems (Bourgeois et al., 2014; Costanza et al., 2014; Simm et al., 2015) in that it provides personalized recommendations based on individual homes' characteristics instead of making general suggestions based on power grid load or green energy generation. ThermoCoach represents an embodiment of the Smart Thermostat algorithm (Lu et al., 2010) which uses occupancy data to automatically calculate optimal thermostat schedules but allows users to choose between a set of schedule options.

Next, we introduce the eco-coaching approach that guided the design of ThermoCoach (Pisharoty et al., 2015).

ECO-COACHING DESIGN APPROACH

Here, we define *eco-coaching* as giving personalized suggestions for specific actions that would reduce wasted energy. Below are specific eco-coaching features of ThermoCoach:

Personalized recommendations: The system monitors occupancy patterns of a home over time using Bluetooth-based occupancy sensing and motion sensing within the home. It builds a model indicating the probability that the home is occupied at each time of the day. Then, it generates schedule options that fit the occupancy pattern of the house and reduce energy waste. The recommended schedules maintain the typical temperature settings preferred by each

household, but generates setback strategies by adding new setback temperatures for times when the house tends to be unoccupied. Three scheduling choices differ in the amount of time the system remains at the setback temperature level and therefore represent different tradeoffs between predicted comfort and energy savings. Further details about the recommendation algorithm can be found in (Lu et al., 2010) and (Pisharoty et al., 2015).

Eco-feedforward: To assist users to make informed decisions about whether such recommendations ought to be followed, the system provides *eco-feedforward* messages to convey the projected impacts of the recommended changes (Yang et al., 2014) in terms of energy savings (i.e., 5%, 7% and 10%) and expected comfort level (i.e., Barely change, May decrease slightly and Will decrease). This supports users in exercising their agency to negotiate priorities between energy savings and comfort preferences.

User control: Instead of letting the system automatically change the schedule, the system leaves users in control by asking them to review schedule options and make decisions about which recommendation to follow (if any) and how schedule options should be implemented. Users can adjust schedule options to complement what is not considered by the system and better accommodate their preferences.

Easy invocation: Oftentimes users do not have the time or willingness to program their thermostats. The system sends a ‘push’ email containing different recommendations. Users can instantly initiate a new schedule by clicking a button in the email.

METHODS

A previous paper reported on an analysis of the sensor data, thermostat usage, and ThermoCoach interaction to determine whether ThermoCoach had an impact on energy savings (Pisharoty et al., 2015). That paper, however, did not investigate *why* ThermoCoach impacted users’ choices. For that, we turned to the interview data to study how different types

of mechanism affect participants' energy saving and thermostat scheduling practices.

Participants

Over 27,000 flyers and door hangers were distributed through local newspapers and by manually placement on doorknobs to recruit participants. After screening and dropouts, 36 households participated in the interview associated with this study. All homes were located within 30 miles of each other and were subject to similar weather conditions throughout the study. Table 2 below describes characteristics of those households.

Table 2. Key characteristics of the 36 households

Housing type	32 homes own a single family house, 4 own/rent an apartment, condo or town house.
Education	25 homes with Master's or Ph.D., 11 homes with B.A/B.S or Associate's degree.
Number of occupants	21 homes have adults with children under 18 years old. 13 homes have only adults. 2 homes are single person homes.
Pets	25 homes have pets (e.g., dogs, cats, fish, rabbits).
Average summer energy bill	24 homes: ranged between \$100~\$200; 4 homes: above \$200; 4 homes: under \$100; 4 homes did not respond.

FIELD DEPLOYMENT STUDY


A 12-week field study instrumented 40 homes with over 190 data collection endpoints, over 250 motion sensors, over 135 Bluetooth low energy transmitter tags, and 40 Nest thermostats. Participants were divided into three groups to compare the energy saving impacts among three thermostat schedule approaches: manual programming, automation, and eco-coaching.

Table 3. The participant homes were divided into three groups. All groups received eco-feedback emails (Figure 9). Group P could only manually program their thermostat. Group N used Nest’s automation features. Group TC received eco-coaching recommendations (Figure 10).

	Group P	Group N	Group TC
Eco-feedback	V	V	V
Auto-Schedule, Auto-Away	-	V	-
Recommendations	-	-	V

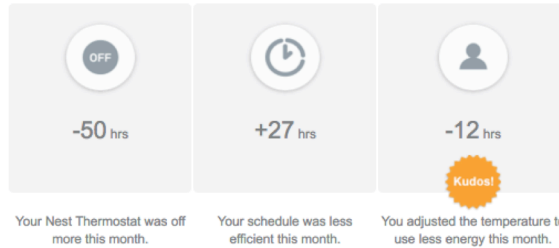
As shown in Table 3, all homes received Nest Energy Report emails (Figure 9), which included eco-feedback elements such as the number of hours the air conditioner ran and the number of days that the Nest detected the home was unoccupied and activated an energy-saving mode called “Auto-Away.” Each group used a different thermostat schedule approach. Group P manually programmed their Nest thermostat. Group N used their Nest thermostat with automation features that automatically programmed the schedule (Auto-Schedule) and adjusted the temperature when motion was not detected for a certain amount of time (Auto-Away). Group TC used their Nest thermostat and received ThermoCoach email that provided schedule recommendations (Figure 10).

Here's how you did:



Why did your energy use change?

We look at a lot of reasons your energy use can change — from weather to Auto-Away — and these are the ones that made the biggest difference this month. [Learn more >](#)



A tip for you:

Spooked by your energy bill?

Line the driveway with solar lights, and help the little monsters find their way.

...

A look at your Leafs:

You get a Leaf when you choose an energy-efficient temperature. And now, see your Leafs add up all year long.



Let your friends know how many Leafs you earned.

[Share](#) [Tweet](#)

Figure 9. The Nest monthly report email provides eco-feedback. It displays number of hours for the cooling or heating used in comparison of this past month to the previous one, what factors impacted more or fewer hours of cooling or heating were used in this past month compared to the previous one, a tip to help you save more energy, and the number of a home earned a green leaf by adjusting the temperature to save energy.

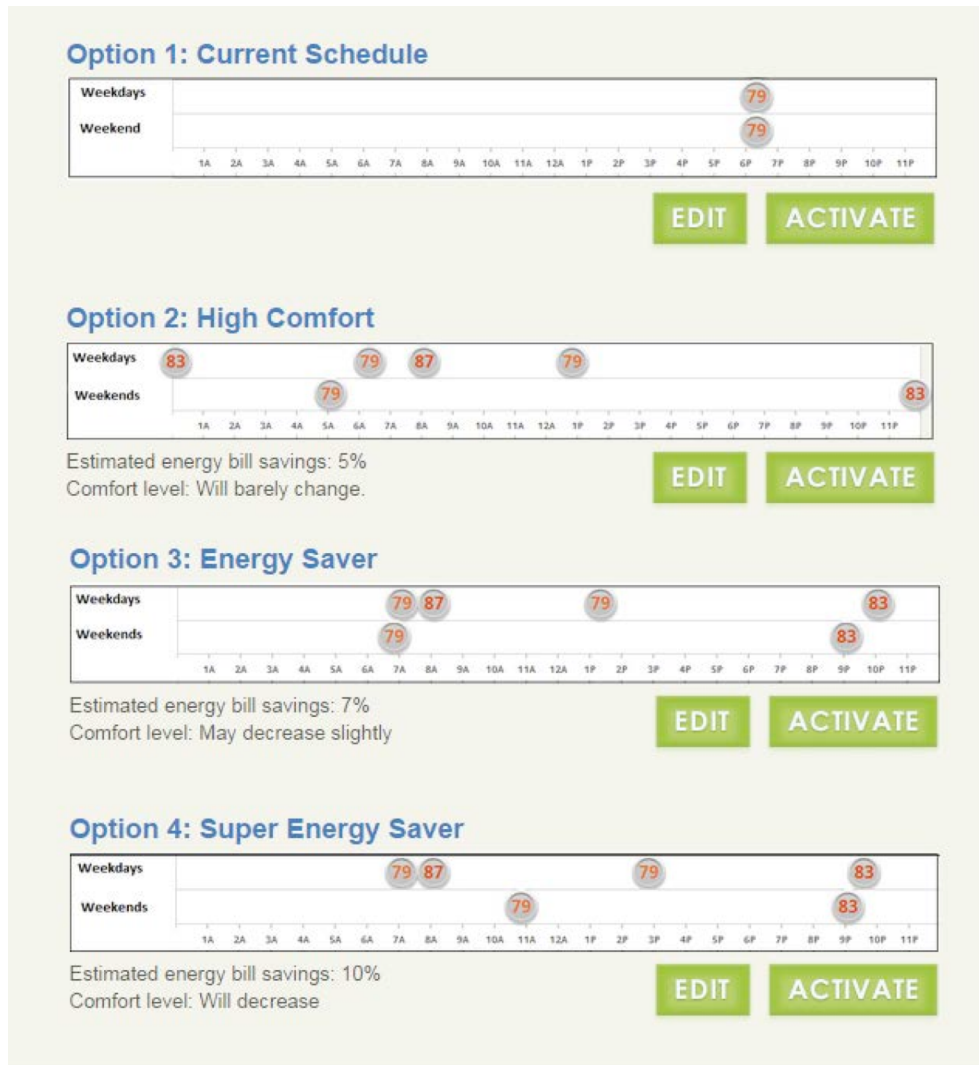


Figure 10. ThermoCoach emails 4 options to each user: their current schedule, a high comfort schedule, an energy saver schedule, and a super energy saver schedule. Users can *Activate* a schedule by clicking a button to set it automatically to be programmed into their thermostat.

As a part of the field study, a series of interviews were conducted. Each participant home participated in three interviews during the study period (at the beginning, in the middle, and at the end of the study). During the interviews, researchers asked participants about how their scheduling practices changed during the study and how they used and reacted to the scheduling mechanism they were assigned. We analyzed 108 interviews with 36 homes that completed

their participation in the interview study. All interviews were conducted by phone and audio-recorded. The interview data were transcribed, coded and analyzed using an iterative process of generating, refining, and probing the themes that emerged.

FINDINGS

Before we report findings from the interview study, we briefly summarize the findings from the quantitative comparison of the three approaches, which were previously published in (Pisharoty et al., 2015).

- Group TC saved more energy than the other two groups. It received eco-coaching recommendations and was able to save 4.7% more energy than Group P, which manually programmed its schedule, and 12.4% more energy than Group N, which used the Nest automation.
- Eight out of 12 homes in Group TC adopted a new schedule. Six of these eight homes had a schedule without setbacks prior to the intervention and activated a schedule with new setbacks afterwards. Two homes already had a schedule with setbacks, but activated a more energy efficient schedule.

This indicated that Group TC homes adopted new and more efficient setback temperatures than homes in the other groups and thus saved more energy. In the following sections, we draw on interview data to describe Group TC participants' experiences and reflections on ThermoCoach's eco-coaching features. We refer to participants by pseudonym. Unless indicated otherwise, participants were from Group TC.

Actionable recommendations allowed users to take action.

Participants liked that the recommendations provided them actionable and concrete plans for future savings. Most participants reported that the recommendations allowed them to recognize opportunities to save energy and indicated what to do, making it easier to take action. Many,

like Amy, considered prospective, proactive plans useful: *“The Nest monthly report reflects on the past month, which is nice. ...The [ThermoCoach] recommendations are nice because they're thinking in the future, and they're looking at your patterns and saying, “Well, based on what you've got it set at now, these options would be totally doable. Three percent savings that you could consider, depending on how much you want to change it.”*”

Personalized recommendations increased credibility and reduced uncertainty.

Generating an effective schedule required reconstructing the daily ins and outs and the nuances of users' comfort needs. When participants recognized that ThermoCoach recommendations were tailored to their particular household's situations, they considered the recommendations credible and relevant to follow.

When participants saw that the recommendations reflected their occupancy patterns and followed their temperature settings, they liked the recommendations since they were tailored to their specific home. Tom described ThermoCoach as somebody who works for him:

“I loved it 'cause somebody was thinking for me. Absolutely. ...We could have done it, but ...it was just really nice for somebody else to evaluate how we used the house, and then to make suggestions.” Participants appreciated recommendations because they not only reduced their workload, but also offered recommendations that were relevant and credible. This made participants more willing to adopt an option. They felt ThermoCoach was not asking them to follow a random schedule as it reflected their existing patterns and temperature preferences.

Supporting user agency

A core design principle of eco-coaching is leaving users in charge of making decisions about how to balance their energy savings and comfort needs while providing assistance to guide their decision making. Group TC participants described ways in which ThermoCoach supported user agency in their thermostat scheduling.

Having a choice provided sense of control in negotiating energy savings and comfort goals.

Three schedule recommendations allowed participants to review and compare different options and make decisions about whether to opt-in and which option to choose. Several participants reported that having a choice between different schedule options gave them a sense of control. Jessica explained that she felt she was in charge: *“You guys weren't just telling me, ‘Here's the best way for you to proceed. Do this!’ But really, putting the ownership on us felt like we were taking charge of it and taking charge of our own actions. So I liked that.”*

Multiple recommendations (High Comfort, Energy Saver and Super Energy Saver) allowed participants to negotiate their energy saving and comfort goals according to their motivations and priorities. Interestingly, changes in life situations affected participants' motivations and priorities, and offering multiple recommendations allowed them to negotiate their goals according to the change. For example, Amy chose a 'High Comfort' option for her husband since he had to stay home for the summer due to an injury. Another participant, Susan, lost her job during the study and she tried to use less cooling to save money.

Eco-feedforward supported decision-making by presenting quantified estimates of future savings.

ThermoCoach provides *eco-feedforward*, which presents quantified energy savings estimates (5%, 10% and 15%) coupled with comfort levels (Barely change, May decrease slightly and Will decrease). Finding a sweet spot between energy savings and comfort is a key consideration in setting a thermostat schedule. Therefore, participants found this combination useful when deciding which option to choose. In particular, quantified measures made energy savings more tangible and practical when justifying their decisions. Jim explained how eco-feedforward information helped him make decisions: *“It does give you an indication, and it certainly helps to quantify. ...It's a little bit harder to actually quantify the energy savings looking at the bill because you've got the other variable of your weather changes as well. But*

just basing it off of sheerly comfort and knowing that ...you should anticipate this type of savings. I think that's a very effective combination."

Admittedly, participants acknowledged that predicted energy savings were difficult to estimate due to many changing variables such as daily weather, physical conditions, and daily routines. Liz assumed that the estimation of energy savings was calculated based on the average temperature of her location. She further explained that while eco-feedforward information did not guarantee the energy reduction that was estimated, it was useful for deciding which schedule to choose: *"As a consumer, as long as I knew that it was just an estimate and there was no guarantee that I was gonna save X amount of dollars, then I think that might have some weight or bearing on what option that I choose."*

As suggested in (Yang et al., 2014), we observed in this study that *eco-feedforward* prompted participants' motivation for setting a new goal. For example, Mike was influenced to change his thermostat schedule to save more energy. He found it was comfortable and stayed with a more energy efficient schedule: *"Because, from the information provided for the choices that were made, with all things being equal, I would have expected to have less energy usage. And all things are never equal, obviously. But, what it did tell me is that ... had I made those slight changes, I should still be fairly comfortable and I should also notice some savings. And again, the difficulty is, because all things aren't equal, 'Did in fact that happen?' Well it's hard to say. But at the same time, I haven't gone back and modified the settings that were set up with that option."*

Customization allowed ways to accommodate preferences.

ThermoCoach creates schedule options to maximize setbacks and provides schedule options that increased setbacks. While prompting users to adopt new, more efficient setbacks, we found that ThermoCoach also supported user control by allowing users to accommodate their

particular preferences and needs through editing the schedule. This was found to be effective in increasing the adoption of recommendations. For example, four homes modified schedule options to better fit to their individual home and still made their schedule more efficient than before. One had a fish tank and changed the setback temperature for the daytime. Two homes chose a ‘High Comfort’ option, which made minimal changes to their schedule, but delayed the cooling start time a little bit. As a result, the changed schedule was still more efficient than their previous schedule.

Encouraging experiments with a setback strategy

As mentioned earlier, our previously reported results indicated that more participants in Group TC employed higher (i.e., more efficient) setback temperatures than other groups. As part of the present analysis, we performed additional data analysis and found a notable difference between Group TC and other groups. There were 15 homes across all groups that initially had a schedule without any setback and kept one single temperature all times. Interestingly, all six such homes that received ThermoCoach recommendations adopted setbacks. On the other hand, none of the other nine homes changed their schedules. They were in Groups P and N, and therefore did not receive ThermoCoach recommendations.

In this section, we explain how ThermoCoach was successful in encouraging participants to experiment with the schedule and adopt new higher setback temperatures.

Comparison of schedule options facilitates learning of different scheduling strategies and weighing pros and cons.

Placing participants’ existing schedules alongside different schedule recommendations initiated quick reflection and provided participants with learning opportunities. Many participants described recognizing similarities and differences between schedules easily as they looked through different schedule options. For example, they noticed how time and temperature settings varied among different options. This helped some participants to gain

insights into different ways to employ setback temperatures. For example, Amy had a night setting that she set to cool at 75°F throughout the night. When she found that recommended schedule options suggested raising the setting to 79°F in the middle of the night, she reflected on her existing schedule and contemplated an idea to create a new setback that she had not considered previously: *“We probably wouldn't be able to feel a difference during that sleeping pattern to feel the four-degree difference. So, you could save energy without affecting comfort, essentially. That was good to know.”*

Offering a more comfortable option lowers the barrier to acceptance.

Providing the ‘High Comfort’ option along with more aggressive plans helped to lower barriers to adopting a new setback that was higher than the existing schedule. Because the ‘High Comfort’ option did not make dramatic changes from participants’ existing schedule, having this option eased participants’ concerns or uncertainty during the process of adopting a new schedule. For example, some participants were not comfortable going for the more aggressive option at first. However, they were willing to try the “High Comfort” option since it was not *“that big of a hard shift.”* Also, participants who started with ‘High Comfort’ shared that such an option could help them gradually transition to a more aggressive option. For example, Jim described the process this way: *“One might say, ‘Okay, let me start at the lowest energy-savings.’ ‘Okay, well, that's fine.’ ‘Well, let me bump it up another level here to see how that works.’”*

Alternative options challenge users' beliefs and trigger users to experiment with a new schedule.

When participants were presented with new schedule recommendations, existing beliefs that might have hindered them from increasing the energy efficiency of their schedule were often challenged. One of the notable benefits of recommendations is supporting participants in correcting their misconceptions and encouraging them to experiment with a new schedule. For

example, Jim as noted earlier opted for ‘High Comfort,’ had grown up being taught, “*Leave the thermostat at one setting. That's the most efficient thing to do.*” He kept the schedule mostly at 76°F as shown in Figure 11.

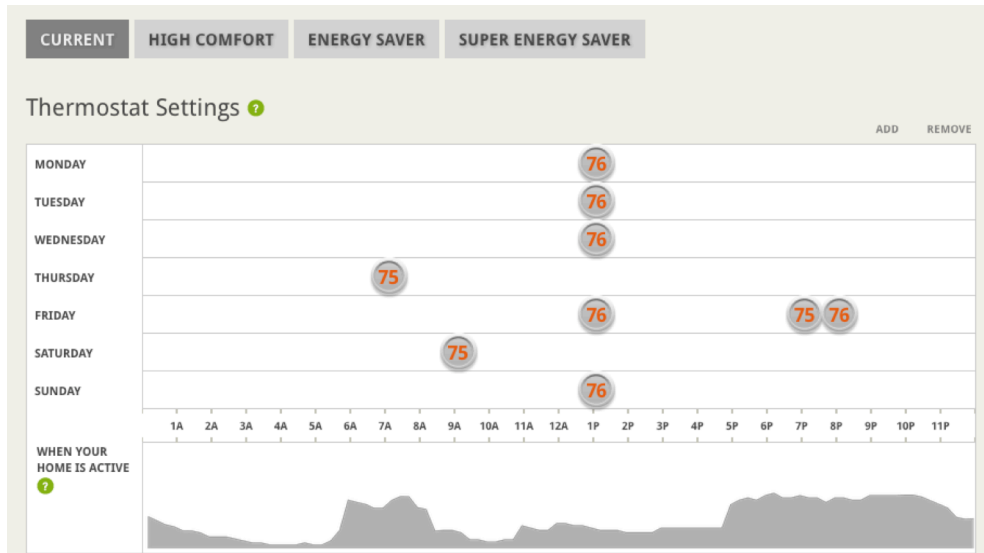


Figure 11. Jim’s previous schedule kept the temperature mostly at 76°F all times.

When Jim first received ThermoCoach recommendations, he found that the idea of a setback increasing energy savings was contrary to what he had always been taught. However, he decided to see how it would work and chose the ‘High Comfort’ option, which was the most conservative approach amongst the three options. His changed schedule is shown in Figure 12.

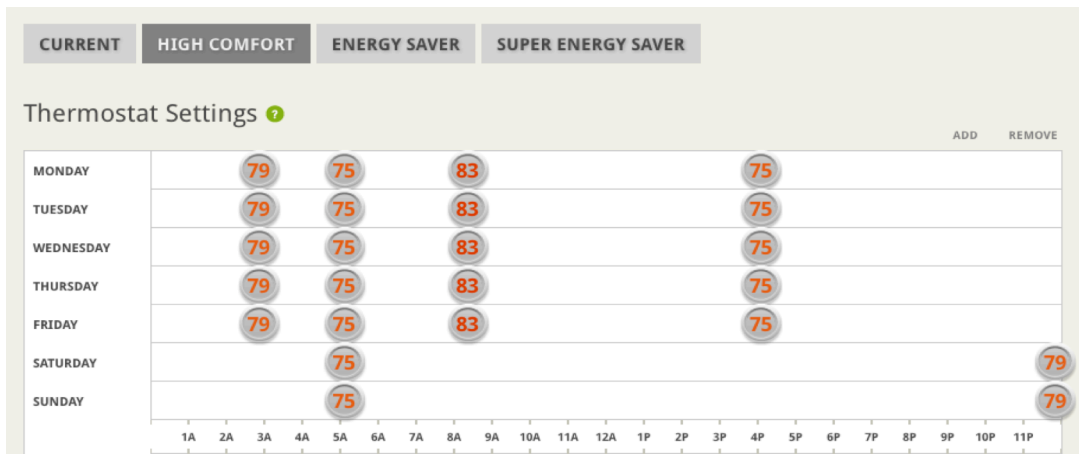


Figure 12. Jim adopted the ‘High Comfort’ option, which included two new setbacks; one was 8°F degrees higher and the other was 4°F higher than his normal temperature.

Jim also explained that ThermoCoach recommendations helped him overcome a certain reluctance to adjust the settings of his system: *“This was a very interesting way to be able to experience that change or experience a result of altering those settings without fear of putting in a completely inappropriate setting, if you will. I thought that was a very beneficial way of doing it, whereas if I was going to try and just do this on my own, I may not have been as well versed in terms of knowing what it would do. [...] It prompts people to think about changing settings when they, again, if you're like me, [are] more likely to leave it set.”*

Jim also added that, *“I would be more likely to experiment with that again to see if I could boost my savings and keeping my comfort level comparable to where it is.”* Like Jim, Tom found that ThermoCoach recommendations were good because *“otherwise, you would just leave it at what you had because that was comfortable, not realizing that a slight change can result in a real saving without a real impact to your comfort level.”*

On the other hand, Steven kept the temperature at 75°F all day. When he looked the projected energy savings that the ThermoCoach recommendations proposed, he was still uncertain about

the idea of adding a setback. Therefore, he accepted the ‘Super Energy Saver’ option, but adjusted the setback temperatures to stay between 76°F and 78°F. He explained why he changed the setback temperatures: *“Just kind of my limited knowledge. It seems like that takes more energy than just trying to keep a house at a steady state, but maybe I'm wrong about that.”*

Shortcomings of ThermoCoach

As we reported, ThermoCoach recommendations provided various benefits and assisted participants with improving the energy efficiency of their thermostat schedule. In this section, we report shortcomings and limitations of ThermoCoach recommendations based on participants’ insights and reflections.

Inability to assess the performance of recommendations lowers user trust in system and its recommendations.

Participants expressed their desire to check how effective the recommendation they chose was in delivering energy savings. For example, Patrick implemented the ‘Super Energy Saver’ option to save 10%, the greatest savings among all schedule options. However, Patrick found that it was not straightforward to know whether he indeed achieved the 10% energy savings that he anticipated. He explained that he could not trust the recommendations if he was not able to verify the actual energy savings after using the schedule: *“Because you can recommend, you can tell me I'm gonna get 10% savings if I choose option D. Or, a 1% savings if I choose option B. Or, no savings if I continue with option A. You can tell me that. But, I'm not gonna believe you until you actually give me some statistics that says, "For the entire month of June, you used 100 kilowatt hours. In the month July, you only used 87 kilowatt hours. Which represented actually a 13% savings, and we estimated it would be about a 10%." I need to be able to hear something as concrete as that before I have any confidence that I'm actually going to achieve something just because you tell me that I will.”*

As mentioned earlier, participants acknowledged that it is difficult to estimate energy savings when there are various factors that are dynamically changing. Thus, they did not expect the actual outcomes to match the estimates exactly. However, participants expected the system to provide concrete evidence to allow them to assess the effectiveness of implementing the recommendations. Without such evidence, they might not have trust or confidence that they would achieve the desired savings by following the system's recommendations.

Inability to detect a mismatch between the schedule and actual use misses opportunities for savings.

In addition to assessing the performance of the recommendation after it was implemented, it would also have been useful to assess how the recommendation was working while it was active. Liz accepted a recommended schedule and thought the schedule worked well for her home. She did not change the schedule after activation. However, during the final interview, she found that the schedule she activated had a setback temperature of 83°F, which was higher than she expected. After learning about this, Liz remembered the times when she noticed an 83°F setting on the Nest thermostat and she simply kept turning it down. It did not occur to her that she might need to check the schedule. Liz suggested that ThermoCoach should provide a new recommendation if users kept making overrides without realizing that their schedule was not working for them: “*I think some follow-up email would probably be nice to tell me that you're not really abiding by this recommendation. ...Then if the new email could perhaps say that, "You are consistently overriding the recommendations," and maybe suggest some new recommendations with maybe lower temperatures or an adjustment of the schedule somehow that still saves energy, but makes me more comfortable, like optimize the process. I think that would be helpful. I would probably respond really well to that also.*”

Failing to address preferences of decreases acceptance of recommendations.

Five out of 13 homes did not opt-in to any option and four homes edited recommendations before they activated them. Here, we explain reasons that those participants did not adopt schedules ThermoCoach recommended.

First, the setback temperatures recommended were too high for some. ThermoCoach suggested daytime setbacks that were 8 degrees higher than the regular temperature of each household. Some homes were already making efforts to save energy by having higher temperatures, such as 80°F, as setbacks or even as their regular temperature. Therefore, those homes thought an even greater setback temperature such as 88°F was too extreme. Second, many households had pets, but ThermoCoach did not take pets into consideration and when creating setbacks for times when there was no human occupancy. For example, Laura had rabbits and kept the temperature at 79°F during the day. When recommendations suggested a setback temperature of 87°F, she found that they were “so drastically” high. Third, participants prioritized comfort needs. Some homes mentioned that it was difficult for them to sleep if it was not cool when going to bed. Emma “vetoed” the recommendations since they suggested raising the nighttime temperature by four degrees. She explained that it would be uncomfortable for her family: “*It was the one that said that we should raise our temperature at night, and we said, ‘No way.’ [...] I’m already having hot flashes.*” Finally, two participants said that they were simply too busy. One commented that it would be rare for her to find a half an hour to set a new thermostat schedule. Another did not remember receiving the ThermoCoach email. One home may have had a system or networking error. This home chose a Super Energy Saver schedule, but the activation did not work due to an unknown error. They did not realize that they were using their old schedule until the final interview.

Two homes experienced discomfort after they accepted a recommendation. Nora found that her fish tank was looking unhappy. Her husband found that the setback temperature was too high

for the fish tank and adjusted the schedule. In Patrick's home, his wife felt it was quite warm when she returned home. Patrick lowered the setback temperature. Both homes fixed the problem by revising their schedule.

In summary, we have just described the benefits and shortcomings of ThermoCoach recommendations as they aimed to assist participants in improving their thermostat scheduling while supporting participants' agency. We found that ThermoCoach supported participants in employing new higher setbacks to increase energy savings and in managing the tension between energy savings and comfort goals.

Next, we discuss the effectiveness and shortcomings of the design features of ThermoCoach. Then we propose design implications for eco-coaching systems to better assist users with planning, executing and assessing their thermostat scheduling effectively.

DISCUSSION

ThermoCoach follows the core principle of eco-coaching: assisting users to take actions to save energy while supporting their agency to take ownership and make informed decisions. ThermoCoach employs several concrete design features to perform eco-coaching for thermostat scheduling. Our findings show that ThermoCoach assisted participants with their thermostat scheduling process, particularly with reflecting on and assessing existing schedule and alternative options, making decisions about balancing energy savings and comfort needs, and experimenting to improve a setback strategy.

In the following sections, we highlight challenges involved in designing recommendation-based eco-coaching systems like ThermoCoach.

Creating recommendations: Improving personalization

In personalizing schedule options for each household, ThermoCoach focused on two key aspects that greatly varied in individual households – *occupancy pattern* and *comfort preferences*. To generate a model for *occupancy patterns*, ThermoCoach collected data using various sensors for X weeks. Participants particularly liked that ThermoCoach collected data over time and because of this they considered the schedule options to be credible. Inferring *comfort preferences* was important because people might ignore a recommendation if it simply asked them to raise the temperature setting to save energy without considering their comfort. Thus, the system used existing temperature settings in users' schedules as a way to incorporate their comfort needs. These design choices were found to be effective in generating recommendations that fit the occupancy patterns and accommodated the comfort needs of individual households.

In our design choices, we decided to leave users in control of revising schedule options to meet any additional situations and comfort needs, such as maintaining cooler air for pets and at nighttime. However, it could be more effective to consider certain variables for personalization. For the successful adoption of recommendations based eco-coaching systems like ThermoCoach, it is critical to allow users to complement the system's lack of capability to understand or anticipate varied and changing needs and situations in individual homes. Doing so can also improve the performance of the system and increase satisfaction with the recommendations. Thus, we suggest that a system like ThermoCoach should ask users for information regarding, for example, pets and sleep preferences, to better understand their constraints and requirements for a thermostat schedule.

Having pets was a common factor to consider for scheduling. About 70% of participants in this study had pets. According to the 2015-2016 APPA National Pet Owners Survey, 65% of U.S. households own a pet¹⁰. Therefore, considering pets when generating schedule recommendations would be useful. For households with younger children, comfort needs were also prioritized over energy savings. Interestingly, some households commented that they would have opted-in to more aggressive options for the winter schedule because the energy bill tended to be much higher in winter than summer for participants in our study. One participant lost her job during the study and wanted to use less cooling to save money. Generating options to address changes in energy saving motivations should be considered to increase the benefit of eco-coaching.

Assessing recommendations: Multi-phase assessment

The ThermoCoach system provides estimated energy savings as a way to prompt users to adopt a schedule recommendation. However, it does not report how much following the recommendations actually saved. One participant mentioned that he would not trust the system unless it provided concrete evidence to show if or to what extent the recommendations worked as the system had proposed.

To build user trust with recommendation-based eco-coaching systems, systems should support users to assess the quality and performance of recommendations over time. We suggest that an eco-coaching system should provide not only projected energy savings estimation for the future, but also quick and easy assessment of how effective the recommendations were at delivering what they proposed to users after implementation.

¹⁰ 2015-2016 APPA National Pet Owners Survey. (2016). http://www.americanpetproducts.org/pubs_survey.asp

Assessing the actual performance of recommendations after use with consideration of real-world factors and conditions

While assessments would be useful, there is a challenge in evaluating the performance of recommendations—the actual energy saving outcomes—compared to the estimated savings. There are various factors that dynamically change in real environments that cannot be predicted in advance. For example, weather changes throughout the season and people’s daily schedule and activities vary; these in return affect how people heat/cool their house, and physiological and psychological factors affect their comfort preferences.

Indeed, many participants in our study reported this was a common reason that they were not able to assess their energy efficiency since they could not simply compare their energy bills month-by-month or year-by-year. Here, we note that several of our participants also mentioned that they understood that estimated savings were ‘estimated,’ and that actual savings would vary according to changing conditions and situations in the real environment, such as weather. What these participants wanted was information that helped them believe the system would bring the benefits it proposed as long as the conditions under which the estimates were made were maintained.

Thus, it is important for an eco-coaching system to indicate to what extent the actual energy savings out- or under-performs the initial estimation. Then, the system should also explain what factors affect the differences between actual and estimated savings. Providing information regarding to what extent and in what aspects the actual conditions and users behaviors in the real environment were different from the predicted conditions and user behavior pattern would be useful. The most obvious factor would be weather. The system could show how the weather differed from previous months and how the difference impacted the heating or cooling needs of the house. It could also show how occupancy patterns were different than the patterns that were used to generate the initial recommendations. As we

mentioned earlier, one household had one member who had to stay home due to injury, greatly increasing the amount of time during which that house was occupied.

Providing hindsight evidence with post-hoc simulation of alternative recommendations

We suggested that assessment of the performance of the recommendations after use would be beneficial. However, there would still be a lack of evidence that the recommendation was particularly good because there would be no way to compare the recommendations to an alternative. Users do not know what might have been if they had implemented other schedule options. It would be useful for users to be able to evaluate not only the recommendations they chose, but also alternatives that they did not implement.

We propose that eco-coaching systems should compare how alternatives might have worked compared to the recommendation they used. Eco-coaching systems could provide post-hoc simulations for the alternative options along with assessment of the chosen option to gauge the outcomes that might have been achieved. Understandably, users do not have all the necessary information to know which option would work mostly effectively for them at one time. However, it becomes much easier to gauge how different strategies would have worked afterwards.

Providing post-hoc assessment of the alternatives could be effective in providing opportunities to understand how different strategies would have worked. In particular, it would make it easier for users to correct their existing misconceptions and thus make more informed decisions for future scheduling. For example, Steven, who did not choose a schedule option with a higher setback temperature, might have been convinced if he had seen how alternative options could have worked under the same circumstances. Assessment of alternatives using post-hoc simulation would provide additional evidence to reinforce the performance of a chosen recommendation because it is easier to evaluate performances that are comparable to

each other. More importantly, this would provide an additional learning opportunity for users to discover pros and cons of different options in accommodating various everyday situations.

Hindsight is always 20/20. While projected estimation was useful for users in making decisions for their planning, reflecting on how their schedule worked as well as how alternatives might have worked could inform users about the impact their decisions had or might have had. Based on assessment of how effective alternatives might have been in bringing energy savings outcomes, users could also have opportunities to learn how different strategies would have worked for certain situations. We note that users would have reacted differently to alternatives if they were indeed were used; for example, users might have felt more uncomfortable and ended up overriding the temperature and consuming more energy. Simulations of alternatives would be still useful for users in reflecting on their decisions and possible outcomes.

Performing assessment for the schedule in use

So far, we have discussed the benefit of assessing recommendations after implementation. Lastly, we suggest that an eco-coaching system should conduct an ongoing assessment of the performance of the schedule that is in use. One of the functions that we described for an eco-coaching system includes monitoring user behaviors and their energy use and identifying discrepancies between them. In addition to monitoring user behavior and energy use to generate recommendations, an eco-coaching system should perform quick, ongoing checks to assess how the schedule in use is working. For example, as in the case of Liz, if users are making many overrides after accepting a recommendation, an eco-coaching system can follow up and ask if the users want to stay with their choice or if the system should provide new recommendations based on their adjustments to the schedule. This would provide an opportunity to discover time or temperature settings that do not work for them. If participants happen to be not committed to their decisions (e.g., a schedule they chose based on their

motivations for energy savings), follow-up checks could reinforce their decisions. The system would also need to monitor how the schedule is working in terms of occupancy patterns. When users' occupancy patterns change from those upon which the recommendations were based, the system should generate new recommendations to adjust to the modified needs for the home.

If provided, tools enabling users to assess the quality and performance of recommendations, including the schedule they used as well as alternatives, would increase the credibility of recommendations over time. This, in turn, would allow users to trust the system and further experiment with their scheduling to increase energy savings. In addition, this process could support users in exercising their discretion (knowledge and insights into different situations that were not sensed or interpreted by the system) to better evaluate the quality and performance of different strategies based on their particular situations. Users would be able to build and strengthen their understanding and ability to utilize and apply various strategies for scheduling to accommodate their interests and priorities.

CONCLUSION

Eco-coaching assists users by providing recommendations tailored to their behavior patterns and preferences and making it easier to take actions, but also leaves users to make decisions about whether or how they should follow such recommendations. From the user perspective, participants identified several benefits of ThermoCoach. It made it easier for them to generate a schedule, provided opportunities to reflect on their thermostat schedule by comparing alternative options and weighing pros and cons, and helped them to make informed decisions for individual homes' needs and situations. Further, it challenged their existing beliefs and encouraged experimenting with their scheduling.

CHAPTER 6.

CONCLUSION

In previous chapters, we have described a series of studies that investigated user interaction with and experience of intelligent domestic systems in the wild, particularly those that seek to learn about and adapt to users' behavior. We focus on informing and evaluating the design of intelligent systems that help users manage their home energy consumption more effectively.

In this section, we summarize three studies that we conducted and restate our contributions.

This dissertation research aims to better understand user interaction and experience of intelligent systems in the home, provide design guidelines and recommendations for intelligent systems in the home, and finally examine specific interaction concepts for supporting energy savings. To do so, this dissertation examines the following **research questions** (RQs) through a series of studies:

In the first study, in order to better understand the challenges of deploying an intelligent system in the home and to inform future design, we began with investigating the lived experience of an advanced thermostat, the Nest Learning Thermostat (Yang & Newman, 2013).

RQ 1: How do people understand and interact with intelligent systems in the home?

In the second study, we compared people's interactions with conventional thermostats with interactions their with the Nest, and observed how user relationships and experiences with an intelligent system changed over time (Yang, Newman, & Forlizzi, 2014).

RQ 2: How do the intelligent features impact users' interaction with their thermostat?

RQ 3: How does user interaction with an intelligent thermostat change over time, and how does it affect energy savings?

The first and second studies together described users' lived experience of intelligent technologies and demonstrated problems and challenges that users encountered with these technologies in their daily environments. Below, we state a set of **contributions representing empirical findings** based on the two studies.

- IV. Lack of support for intelligibility and user control in everyday intelligent technologies hinders users from understanding how the system interprets and adapts to users' behavior and situations, and thus deters them from intervening to guide or correct the system's behavior. (Chapter 3)
- V. Users' engagement with the system helps to address system shortcomings and improve performance. However, maintaining users' engagement over time is difficult when users have little motivation to go through the effort of understanding and assessing the system's behavior. (Chapter 3, Chapter 4)
- VI. Users' reliance on intelligent systems and diminished interactions results in missed opportunities for energy savings. Sustaining user interaction and engagement with intelligent system is critical to achieve the goal of energy savings. (Chapter 4)

Based on these findings, we developed design guidelines for end-user interaction with intelligent technologies. The second set of **contributions** consists of **design guidelines**.

- I. **Design guidelines for supporting user understanding and control.** We propose three avenues for future development of everyday intelligent technologies to support user understanding and control of intelligent systems for the home (Chapter 3).

Exception flagging: Our data supports the view that some amount of human behavior is unpredictable, some preferences change, some routines are unstable, and some contingencies are too rare to form a pattern. A key design challenge is to elicit input from users to help the system differentiate the data that represents regular, stable preferences or behavior from input that does not. Rather than have people give explanations about every intent they have, we can have people just note when a change is not something they want the system to remember. *Exception flagging* can allow people to provide additional information to the system without over-burdening them.

Incidental intelligibility: It is challenging to convey an understanding of how an intelligent system works given that users in the home are unlikely to pay a lot of attention to any individual system. One possibility might be to consider ways in which *incidental intelligibility*—interaction elements that increase users’ understanding of the system’s intelligent behavior that are embedded in the tasks users consciously seek to accomplish—could help users build understanding of a system behavior over the long term without asking their focused attention to learning how the system works as a discrete task.

Constrained engagement: Users are not likely to devote a great deal of effort to interacting with intelligent systems in the home. However, systems require some amount of engagement from the users to perform optimally. As we are evolving towards a world in which users engage with dozens if not hundreds of intelligent systems like the Nest, UbiComp researchers face the challenge of designing technologies that engage but do not overwhelm—a goal that we refer to as *constrained engagement*. Such engagement must be dramatically constrained, given

that the interaction between user and system is necessarily sparse and peripheral, yet continuous and long-lived.

- II. **Design guidelines for balancing system autonomy and user control.** Reflecting on the tensions between convenience and benefits of automation and user engagement necessary for energy savings, we propose a set of design implications that invite user participation and reflection with the goal of saving energy at home (Chapter 4).

Providing actionable recommendations: Designers should consider ways to generate concrete plans for increasing energy savings that leave users in control but are easy for users to implement. As an example, consider a recommendation for an improvement to the user's schedule that appears on the home screen of a thermostat control app or in an email. This recommendation could include an option that allows the user to implement the recommended change instantly.

Providing eco-feedforward: To help users decide whether such recommendations ought to be followed, systems could further provide *eco-feedforward* messages or visualizations to convey the projected impacts of implementing the recommended changes. Providing actionable recommendations along with information about the projected benefits of those recommendations would enable systems to suggest courses of action that align with system goals while allowing users to stay in control.

Stimulating reflection and reassessment: To sustain user engagement over time, eco-interaction technologies need to maintain lightweight engagement between users and the system on an ongoing basis. For example, a system could allow the current schedule to expire after a period of time, or ask users to choose between their existing schedule and a more efficient one proposed by the system as one way to stimulate evaluation and reassessment of the schedule.

In our final study, we evaluated the effectiveness of the guideline for balancing system autonomy and user control. To do so, we first developed a design approach that we call **eco-coaching**: giving personalized suggestions for specific actions that would reduce wasted energy. Then, we conducted a 12-week deployment study of the ThermoCoach system, which performs eco-coaching for thermostat scheduling to answer the following research questions:

RQ 4: How does the eco-coaching design approaches work for balancing system autonomy and user control?

RQ 5: How do schedule recommendations and eco-feedforward affect users' thermostat control practices and energy savings?

In the ThermoCoach study, we demonstrated and evaluated the design approach of eco-coaching to balance system autonomy and user control. The principles of eco-coaching are to provide personalized suggestions for specific actions that would reduce wasted energy, while supporting user agency in accomplishing their energy saving goals.

Here we provide **contributions to sustainable HCI**: insight into design challenges, a long-term field study, and findings and design recommendations regarding interaction strategies to support energy savings in the home.

- I. **Findings from a longitudinal study of eco-feedback technology.** We found that the combination of eco-feedback and machine learning-based personalization led to increased engagement with energy-saving features of the system in the short term, but that such engagement was not sustained over the long term. (Chapter 3 and Chapter 4)
- II. **Findings on effectiveness of the eco-coaching approach.** We found that the eco-coaching approach 1) made it easier for users to implement an effective thermostat schedule, 2) supported user agency in negotiating trade-offs between energy savings

and comfort, 3) facilitated learning different scheduling strategies as well as weighing pros and cons of different options, and 4) challenged users' beliefs about how well they were doing. These outcomes, in turn, were successful in getting users to employ and experiment with more efficient setback strategies. (Chapter 5)

III. Design guidelines for supporting users' assessment of system performance.

To build user trust with recommendation-based eco-coaching systems, systems should support users to assess the quality and performance of recommendations over time.

We suggest that an eco-coaching system should provide not only projected energy savings estimation for the future, but also quick and easy assessment of how effective the recommendations were at delivering what they proposed to users after implementation. (Chapter 5)

Assessing the actual performance of recommendations after use with consideration of real-world factors and conditions: Users wanted information that can help them trust the system would bring the benefits it proposed. Thus, it is important for an eco-coaching system to indicate to what extent the actual energy savings out- or under-performs the initial estimation. Then, the system should also explain what factors affect the differences between actual and estimated savings.

Providing hindsight evidence with post-hoc simulation of alternative recommendations: Eco-coaching systems should compare how alternatives might have worked compared to the recommendation they used. It would be useful for users to be able to evaluate not only the option they chose, but also alternatives that they did not implement. Post-hoc simulations for the alternative options can help users to gauge the outcomes that might have been achieved.

Performing assessment for the schedule in use: Eco-coaching systems should conduct an ongoing assessment of the performance of the schedule that is in use. When users' occupancy patterns change from those upon which the recommendations were based, the system should generate new recommendations to adjust to the modified needs for the home.

In this final chapter, we have provided a detailed list of contributions from previous chapters. This dissertation made contributions to the interdisciplinary fields of human-computer interaction (HCI) and ubiquitous computing (UbiComp).

LIMITATIONS AND CAVEATS

Here we clarify the scope of this thesis and thus the direction taken by the design implications for intelligent technologies for the home.

The scope of studies on providing design implications for sustainability

First, it is important to clarify the scope of this dissertation research on sustainability and thus the direction taken by the design implications, especially with regards to situating our recommendations within the larger space of technology to support energy savings in the home.

In this dissertation, we are not challenging the notion that people have relatively stable expectations for thermal comfort and that they expect indoor temperatures to be mechanically maintained at a level concordant with those expectations. More specifically, we are not engaging critiques of the cultural construction of thermal comfort (e.g., (Chappells & Shove, 2004)) or models of adaptive thermal comfort (e.g., (Clear, Morley, Hazas, Friday, & Bates, 2013)) that suggest that people can or should attain comfort through other means than mechanical heating and cooling.

While we find such alternative views compelling and highly deserving of consideration, we are focused here on investigating the bounds of energy efficiency that can be obtained within the commonly understood constraints of thermostat-controlled temperature regulation for obtaining personal comfort. From this perspective, we believe that finding a balance between automation and user engagement will be key to optimizing energy efficiency in the face of consumer expectations of comfort. We also believe that finding such a balance is a challenge that the field of HCI is particularly well suited to address. Even within the frame of improving thermostat control to achieve better energy efficiency in the face of a presumably stable comfort requirement, our study has limitations.

Limitations by the nature of the technology

The goal of this dissertation has been to illuminate the principles for designing intelligent systems for the home. This dissertation has demonstrated how users' understanding and control of intelligent systems are critical to the desired system performance. It also showed how users could successfully cooperate with systems to achieve their desired goals such as increasing energy savings in the home. However, studies in this dissertation have only studied a particular area where intelligent systems can support domestic life, namely energy savings. While we have argued that the commercial deployment of an intelligent thermostat, the Nest, and field deployment of the ThermoCoach have provided valuable opportunities for studying this issue, our studies are limited by the nature of the technology studied.

Different domestic technologies will vary in terms of complexity, distribution of labor, and relative importance to household members. It would be difficult to argue, for example, that findings from our studies could be blindly applied to adaptive systems that control lighting, security, or entertainment. While we think that some of our insights will apply (exception flagging is likely to be important for many machine learning-based systems, constrained

engagement could be a reasonable goal for mostly-disinterested stakeholders), further study will be needed to determine how and when to apply these principles.

Supporting user understanding, control and engagement is important for improving the system performance and increasing the benefit of the system. However, there would be different ways or levels of importance for different systems. For example, different systems require different levels of explicit user control. There might be less need for monitoring the system performance for sprinkler systems than for home security systems. The risk of inaccurate performance of home security system (i.e. not detecting when a burglar was entering the home) is greater than water sprinkler system (i.e. watering the plant the ground is moist after the rain). The system with higher risk would need to better support for user understanding of the system (e.g., easy understanding of how the sensors work) and assessment of system performance than with systems with lower risk.

In addition, different level of user engagement might be expected. People would feel more need for and be willing to check the performance of home security system and ensure the system alarms when somebody invades the home. On the other hand, it might not be the same when people think about water sprinkler system. The benefit of having water sprinkler system lies largely on the convenience of the systems automatically water the yard. Thus, assessing the system performance of water sprinkler might not be necessary to be frequent or even not necessary until people notice the plants are withering. As long as there is not a signal for issues with system performance, people might not even need to be engaged in the case like water sprinkler system beyond periodic assessment of water usage (e.g., to detect possible water leaks). However, with home security systems, it is more difficult for people to notice if the security system is working properly or not, there is not an obvious signal like withering plants. Because it would be too late or no use to improve or assess the system performance after an incident such as home invasion happened, trust-and-verify approach would not be acceptable

for a system like home security systems. Therefore, supporting user engagement should be designed with considerations for different factors such as user expectations and tolerance for system error, user motivations for engaging with the system (e.g., monitoring the system performance), and consequences of system failure. Different systems will need to have different ways of engaging users to support user understanding, control as well as to balance system autonomy and user agency and control.

The characteristics of our participants

As smart devices like the Nest achieve wider adoption, studies of different stakeholders within the home will be increasingly needed. As noted, our participants were disproportionately tech-savvy, affluent, and male. Though we focused on the 'primary' users of the Nest in our interviews and diary studies, we became aware of different levels of engagement among different house members, echoing patterns found in other studies of home automation (Brush et al., 2011; Mennicken & Huang, 2012). Primary users tended to be more engaged, meaning that they were willing to learn and employ advanced features of the Nest. Other house occupants often did not share the same interest, and in many cases used the Nest as they did their previous (conventional) thermostats. Other studies have identified the importance of gender roles with respect to technology configuration and use (Rode et al., 2004), as well as that of computer expertise and identity (Poole, Chetty, Morgan, Grinter, & Edwards, 2009). Further studies should strive to understand different perspectives within the home with respect to adaptive technologies, so as to provide a more balanced understanding of how such systems ought to be designed.

Our studies were restricted to the continental United States, and looked at a restricted set of people over a constrained period of time. The Nest users we studied were relatively affluent and technologically savvy compared to a more general population. All homes participated in

the ThermoCoach study were located within 30 miles of each other in Virginia, United States and were subject to similar weather conditions throughout the study.

Participants in our three studies were relatively highly educated, likely to have their own house, and were mostly married. Our conventional thermostat study participants were a bit more varied in these regards, but still not representative of the vast diversity of living situations, housing types, and individual differences found in US residences—to say nothing of differences across the globe. In addition, less motivated and educated users might be less inclined than our participants to monitor and improve the system operation, and so might benefit from a higher degree of automation and an even more constrained level of engagement. Finding the optimal level of engagement for different populations or even individuals remains a significant challenge.

In summary, understanding the issues with improving user control of intelligent systems in general, and with HVAC systems more specifically remains a significant challenge, and we look forward to further studies that will deepen and fill out the findings of the present work.

CONCLUDING REMARKS

The larger goal of this dissertation is to guide future efforts for designing effective interaction between users and intelligent systems for the home. This dissertation *enhances our understanding* of daily user interaction with and experience of intelligent systems, *identifies challenges* in deploying intelligent systems in the home, and *provides design approaches* to bridge the gap between system capability and user control.

This dissertation examines the various challenges and tensions that can arise in end-user interactions with everyday intelligent technologies as these systems seek to track and adapt to dynamic, nuanced day-to-day user behaviors, activities and situations. Based on findings from

field-based studies, this dissertation provides design considerations for desirable goals and design properties for end-user interaction with intelligent technologies for the home.

The intellectual merit of this dissertation derives from a series of field studies and user-centered design approach to investigate issues of designing intelligent systems in the home.

Our field studies with its accompanying diary study and interviews provide a valuable opportunity for us to observe how users understand, interact with and use intelligent systems in their everyday environments. To our knowledge, *this is the first study that looks at user interactions with and lived experience of an intelligent system in the home*, in particular one that monitors and “learns” user behavior and adapts its behavior accordingly: the Nest Learning Thermostat. There has been a lot of work in ubiquitous computing over the years looking at sensor-based intelligent systems in the home.

In very many cases, the evaluations of these systems have been conducted in laboratory-type settings, such as the AwareHome (Abowd, Bobick, Essa, Mynatt, & Rogers, 2002), House_n (Intille, 2002), or—when they are done in real homes such as the Adaptive Home (Mozer, 1998)—are installed for a short period of time, studied, and then removed. Therefore this dissertation provides valuable understanding of challenges to deploying an intelligent system in an everyday home environment, as it reflects end-users’ interaction with and lived experience of this type of technology over a long period of time.

Our comparative study, in which we compared an intelligent system to conventional systems, sheds light on how end-users interact with and respond to the interactions and functions that are unique to the intelligent system, and how this impacts users’ experience of controlling heating and cooling systems. It looks at how an alternative domestic heating and cooling system control (the Nest) evokes new types of interactions that differ from interactions with traditional thermostats. Alternative control interfaces for energy consuming appliances such as

heating and cooling systems are of increasing interest in HCI. Thus the findings from this study provide valuable insights for people working on designing controls for domestic appliance interfaces.

Our longitudinal study examines how interactions with intelligent systems change over time and what impact these interactions bring to energy savings outcomes. This study provides a *much-needed long-term field study* to observe how end-users adopt and use the intelligent system after the novelty has worn off: how the intelligent system is integrated into daily life over time, and how it affects daily practices of managing heating and cooling and thus its impacts on energy savings. We get beyond the novelty effects, and start to see how users adapt to the technology as they live with it, and indeed how they and the technology co-adapt together. Additionally, our study of users' interactions with the Nest over time provides insight into challenges for the design of systems to help users save energy. While there is an overabundance of eco-feedback technology systems that are being developed and evaluated, most of these are prototypes, tested in labs, and only evaluated with regard to their short-term impact. As such, this paper makes an important contribution to the community by providing a long-term field study of *an actually deployed system*.

Finally, this dissertation illuminates an important and notable interaction or interplay between human and machine. We argue that thoughtful, continuous involvement from users for control is critical to the success of both intelligent systems for the home and interventions that promote sustainable choices.

The findings from this dissertation have been valuable for demonstrating how users' understanding and control of intelligent systems are critical to the desired system performance. We argue for the importance of research into making systems intelligible, so that people can better adapt their communication with the system to achieve the desired effects, as well as

recover from errors. We also assert that designing more cooperative, collaborative and coordinated interactions between intelligent systems and their users and figuring out how to sustain those interactions over time is critical to the success of both digital home technology and interventions that promote sustainable choices.

As such, this dissertation contributes to the Human-Computer Interaction and UbiComp research fields by enhancing our understanding of how to design intelligent systems for the home, and we believe that the result of this dissertation research will generalize beyond HVAC systems to other adaptive, intelligent home systems.

Beyond contributing to the Human-Computer Interaction and UbiComp research fields, the broader impacts of our work stem from our ability to contribute solutions to the challenging problem of fostering energy-efficient operation of home heating and cooling equipment.

REFERENCES

- Abowd, G. D., Bobick, A. F., Essa, I. A., Mynatt, E. D., & Rogers, W. A. (2002, July). The aware home: A living laboratory for technologies for successful aging. In *Proceedings of the AAAI-02 Workshop "Automation as Caregiver"*(pp. 1-7).
- Abowd, G. D., & Sterbenz, J. P. (2000). Final report on the inter-agency workshop on research issues for smart environments. *Personal Communications, IEEE*, 7(5), 36-40.
- Al-Sanea, S. A., & Zedan, M. F. (2008). Optimized monthly-fixed thermostat-setting scheme for maximum energy-savings and thermal comfort in air-conditioned spaces. *Applied Energy*, 85(5), 326-346.
- Bellotti, V., & Edwards, K. (2001). Intelligibility and accountability: human considerations in context-aware systems. *Human-Computer Interaction*,16(2-4), 193-212.
- Bourgeois, J., Van Der Linden, J., Kortuem, G., Price, B. A., & Rimmer, C. (2014). Conversations with my washing machine: an in-the-wild study of demand shifting with self-generated energy. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 459-470). ACM.
- Brush, A. J., Lee, B., Mahajan, R., Agarwal, S., Saroiu, S., & Dixon, C. (2011, May). Home automation in the wild: challenges and opportunities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2115-2124). ACM.
- Chappells, H., & Shove, E. (2004). Comfort: a review of philosophies and paradigms. *University of Lancaster*.
- Clear, A. K., Morley, J., Hazas, M., Friday, A., & Bates, O. (2013, September). Understanding adaptive thermal comfort: new directions for UbiComp. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (pp. 113-122). ACM.
- Consolvo, S., Roessler, P., & Shelton, B. E. (2004). The CareNet display: lessons learned from an in home evaluation of an ambient display. In *UbiComp 2004: Ubiquitous Computing* (pp. 1-17). Springer Berlin Heidelberg.
- Cook, D. J., Youngblood, M., Heierman III, E. O., Gopalratnam, K., Rao, S., Litvin, A., & Khawaja, F. (2003). MavHome: an agent-based smart home. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on* (pp. 521-524). IEEE.
- Coskun, A., Zimmerman, J., & Erbug, C. (2015). Promoting sustainability through behavior change: A review. *Design Studies*, 41, 183-204.

- Costanza, E., Fischer, J. E., Colley, J. A., Rodden, T., Ramchurn, S. D., & Jennings, N. R. (2014). Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 813-822). ACM.
- Cowan, P. A., Field, D., Hansen, D. A., Skolnick, A., & Swanson, G. E. (2014). *Family, self, and society: Toward a new agenda for family research*. Routledge.
- Crabtree, A., & Rodden, T. (2004). Domestic routines and design for the home. *Computer Supported Cooperative Work (CSCW)*, 13(2), 191-220.
- Crabtree, A., Rodden, T., Hemmings, T., & Benford, S. (2003). Finding a Place for UbiComp in the Home. In *UbiComp 2003: Ubiquitous Computing* (pp. 208-226). Springer Berlin Heidelberg.
- Darby, S. (2006). The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*, 486, 2006.
- Davidoff, S., Lee, M. K., Zimmerman, J., & Dey, A. K. (2006). Socially-aware requirements for a smart home. In *Proceedings of the international symposium on intelligent environments* (pp. 41-44).
- Davidoff, S., Lee, M. K., Yiu, C., Zimmerman, J., & Dey, A. K. (2006). Principles of smart home control. In *UbiComp 2006: Ubiquitous Computing* (pp. 19-34). Springer Berlin Heidelberg.
- Davidoff, S., Zimmerman, J., & Dey, A. K. (2010). How routine learners can support family coordination. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2461-2470). ACM.
- Dey, A. K., Hamid, R., Beckmann, C., Li, I., & Hsu, D. (2004, April). a CAPpella: programming by demonstration of context-aware applications. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 33-40). ACM.
- Dey, A. K., Rosenthal, S., & Veloso, M. (2009). Using interaction to improve intelligence: how intelligent systems should ask users for input. In *Workshop on Intelligence and Interaction: IJCAI*.
- Dourish, P., & Bell, G. (2011). *Divining a digital future: Mess and mythology in ubiquitous computing*. MIT Press.
- Eagle, N., & Pentland, A. (2006). Reality mining: sensing complex social systems. *Personal and ubiquitous computing*, 10(4), 255-268.

Edwards, W. K., & Grinter, R. E. (2001, September). At home with ubiquitous computing: Seven challenges. In *UbiComp 2001: Ubiquitous Computing* (pp. 256-272). Springer Berlin Heidelberg.

Energy Information Administration (EIA) (2011). Table 2.1 Energy Consumption by Sector. Retrieved February 25, 2016, from http://www.eia.gov/totalenergy/data/monthly/pdf/sec2_3.pdf

Energy Savings from Nest White Paper Preview (2016). Retrieved February 25, 2016, from <https://nest.com/downloads/press/documents/efficiency-simulation-white-paper.pdf>

Fanger, P. O. (1970). Thermal comfort. Analysis and applications in environmental engineering. *Thermal comfort. Analysis and applications in environmental engineering*.

Fein, R. M., Olson, G. M., & Olson, J. S. (1993, April). A mental model can help with learning to operate a complex device. In *INTERACT'93 and CHI'93 conference companion on Human factors in computing systems* (pp. 157-158). ACM.

Feldmeier, M., & Paradiso, J. A. (2010, November). Personalized HVAC control system. In *Internet of Things (IOT), 2010* (pp. 1-8). IEEE.

Forlizzi, J. (2007). How robotic products become social products: an ethnographic study of cleaning in the home. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction* (pp. 129-136). ACM.

Forlizzi, J. (2008). The product ecology: Understanding social product use and supporting design culture. *International Journal of Design*, 2(1).

Friedewald, M., Da Costa, O., Punie, Y., Alahuhta, P., & Heinonen, S. (2005). Perspectives of ambient intelligence in the home environment. *Telematics and informatics*, 22(3), 221-238.

Froehlich, J., Findlater, L., & Landay, J. (2010). The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1999-2008). ACM.

Ganti, R. K., Pham, N., Ahmadi, H., Nangia, S., & Abdelzaher, T. F. (2010). GreenGPS: a participatory sensing fuel-efficient maps application. In *Proceedings of the 8th international conference on Mobile systems, applications, and services* (pp. 151-164). ACM.

Gupta, M., Intille, S. S., & Larson, K. (2009). Adding GPS-control to traditional thermostats: An exploration of potential energy savings and design challenges. In *Pervasive Computing* (pp. 95-114). Springer Berlin Heidelberg.

- Hazas, M., Brush, A. J., & Scott, J. (2012). Sustainability does not begin with the individual. *interactions*, 19(5), 14-17.
- Hekkert, P., Snelders, D., & Wieringen, P. C. (2003). 'Most advanced, yet acceptable': Typicality and novelty as joint predictors of aesthetic preference in industrial design. *British journal of psychology*, 94(1), 111-124.
- Herlocker, J. L., Konstan, J. A., & Riedl, J. (2000, December). Explaining collaborative filtering recommendations. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work* (pp. 241-250). ACM.
- Horvitz, E. (1999, May). Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 159-166). ACM.
- Huang, C. C. J., Yang, R., & Newman, M. W. (2015, September). The potential and challenges of inferring thermal comfort at home using commodity sensors. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 1089-1100). ACM.
- Intille, S. S. (2002). Designing a home of the future. *IEEE pervasive computing*, 1(2), 76-82.
- Jiang, X., Dawson-Haggerty, S., Dutta, P., & Culler, D. (2009, April). Design and implementation of a high-fidelity ac metering network. In *Information Processing in Sensor Networks, 2009. IPSN 2009. International Conference on* (pp. 253-264). IEEE.
- Kapoor, A., Lee, B., Tan, D., & Horvitz, E. (2010, April). Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1343-1352). ACM.
- Karapanos, E. (2013). User experience over time. In *Modeling Users' Experiences with Interactive Systems* (pp. 57-83). Springer Berlin Heidelberg.
- Kaushik, P., Intille, S. S., & Larson, K. (2008). Observations from a case study on user adaptive reminders for medication adherence. In *Pervasive Computing Technologies for Healthcare, 2008. PervasiveHealth 2008. Second International Conference on* (pp. 250-253). IEEE.
- Kempton, W. (1986). Two Theories of Home Heat Control*. *Cognitive Science*, 10(1), 75-90.
- Keyson, D. V., de Hoogh, M. P. A. J., Freudenthal, A., & Vermeeren, A. P. O. S. (2000). The intelligent thermostat: a mixed-initiative user interface. In *CHI'00 extended abstracts on Human factors in computing systems* (pp. 59-60). ACM.

Kidd, C. D., Orr, R., Abowd, G. D., Atkeson, C. G., Essa, I. A., MacIntyre, B., ... & Newstetter, W. (1999). The aware home: A living laboratory for ubiquitous computing research. In *Cooperative buildings. Integrating information, organizations, and architecture* (pp. 191-198). Springer Berlin Heidelberg.

Ko, A. J., & Myers, B. A. (2004). Designing the whyline: a debugging interface for asking questions about program behavior. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 151-158). ACM.

Koehler, C., Ziebart, B. D., Mankoff, J., & Dey, A. K. (2013, September). TherML: occupancy prediction for thermostat control. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (pp. 103-112). ACM.

Kujala, S., Roto, V., Väänänen-Vainio-Mattila, K., Karapanos, E., & Sinnelä, A. (2011). UX Curve: A method for evaluating long-term user experience. *Interacting with Computers*, 23(5), 473-483.

Kulesza, T. (2012). An explanation-centric approach for personalizing intelligent agents. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces* (pp. 375-378). ACM.

Kulesza, T., Burnett, M., Wong, W. K., & Stumpf, S. (2015). Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (pp. 126-137). ACM.

Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). Tell me more?: the effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1-10). ACM.

Kulesza, T., Wong, W. K., Stumpf, S., Perona, S., White, R., Burnett, M. M., ... & Ko, A. J. (2009). Fixing the program my computer learned: Barriers for end users, challenges for the machine. In *Proceedings of the 14th international conference on Intelligent user interfaces* (pp. 187-196). ACM.

Lim, B. Y., & Dey, A. K. (2009). Assessing demand for intelligibility in context-aware applications. In *Proceedings of the 11th international conference on Ubiquitous computing* (pp. 195-204). ACM.

Lim, B. Y., & Dey, A. K. (2010, September). Toolkit to support intelligibility in context-aware applications. In *Proceedings of the 12th ACM international conference on Ubiquitous computing* (pp. 13-22). ACM.

- Lim, B. Y., Dey, A. K., & Avrahami, D. (2009). Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2119-2128). ACM.
- Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., ... & Whitehouse, K. (2010). The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems* (pp. 211-224). ACM.
- Mackay, W. E. (2000). Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica*, 30(1), 177-193.
- McGee-Lennon, M. R., Wolters, M. K., & Brewster, S. (2011). User-centred multimodal reminders for assistive living. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2105-2114). ACM.
- McNee, S. M., Lam, S. K., Guetzlaff, C., Konstan, J. A., & Riedl, J. (2003, September). Confidence displays and training in recommender systems. In *Proc. INTERACT* (Vol. 3, pp. 176-183).
- Meier, A., Aragon, C., Hurwitz, B., Mujumdar, D., Perry, D., Peffer, T., & Pritoni, M. (2010). How people actually use thermostats. *ACEEE Summer Study on Energy Efficiency in Buildings. Pacific Grove, Calif.: American Council for an Energy Efficient Economy.*
- Meier, A., Aragon, C., Peffer, T., Perry, D., & Pritoni, M. (2011). Usability of residential thermostats: Preliminary investigations. *Building and Environment*, 46(10), 1891-1898.
- Meier, A. (2011). Thermostat interface and usability: a survey. *Lawrence Berkeley National Laboratory.*
- Mennicken, S., & Huang, E. M. (2012). Hacking the natural habitat: an in-the-wild study of smart homes, their development, and the people who live in them. In *Pervasive computing* (pp. 143-160). Springer Berlin Heidelberg.
- Monk, A., Blythe, M., & Kim, H. (2010). What do people want from the technology in their homes: Work at the Centre for Usable Home Technology (CUHTec). In *Information Technology Interfaces (ITI), 2010 32nd International Conference on* (pp. 313-318). IEEE.
- Mozer, M. C. (1998). The neural network house: An environment that adapts to its inhabitants. In *Proc. AAAI Spring Symp. Intelligent Environments* (pp. 110-114).
- Mozer, M. C. Lessons from an Adaptive Home. *Smart Environments: Technologies, Protocols, and Applications*, 271-294.

Nest | The Learning Thermostat. Retrieved September 24, 2012, from <http://www.nest.com/>

O'Brien, J., Rodden, T., Rouncefield, M., & Hughes, J. (1999). At home with the technology: an ethnographic study of a set-top-box trial. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 6(3), 282-308.

Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization science*, 3(3), 398-427.

Paepcke, S., & Takayama, L. (2010). Judging a bot by its cover: an experiment on expectation setting for personal robots. In *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on* (pp. 45-52). IEEE.

Peffer, T., Pritoni, M., Meier, A., Aragon, C., & Perry, D. (2011). How people use thermostats in homes: A review. *Building and Environment*, 46(12), 2529-2541.

Pierce, J., Fan, C., Lomas, D., Marcu, G., & Paulos, E. (2010). Some consideration on the (in) effectiveness of residential energy feedback systems. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems* (pp. 244-247). ACM.

Pierce, J., Schiano, D. J., & Paulos, E. (2010, April). Home, habits, and energy: examining domestic interactions and energy consumption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1985-1994). ACM.

Pisharoty, D., Yang, R., Newman, M. W., & Whitehouse, K. (2015, November). Thermocoach: Reducing home energy consumption with personalized thermostat recommendations. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments* (pp. 201-210). ACM.

Poole, E. S., Chetty, M., Morgan, T., Grinter, R. E., & Edwards, W. K. (2009, April). Computer help at home: methods and motivations for informal technical support. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 739-748). ACM.

Ranjan, J., Griffiths, E., & Whitehouse, K. (2014, November). Discerning electrical and water usage by individuals in homes. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings* (pp. 20-29). ACM.

Residential Energy Consumption Survey (RECS) - Data - U.S. Energy Information Administration (EIA). (2013). Retrieved September 15, 2013, from <http://www.eia.gov/consumption/residential/data/2009/>

- Rode, J. A., Toye, E. F., & Blackwell, A. F. (2004). The fuzzy felt ethnography—understanding the programming patterns of domestic appliances. *Personal and Ubiquitous Computing*, 8(3-4), 161-176.
- Rogers, Y. (2006). Moving on from Weiser's vision of calm computing: Engaging UbiComp experiences. In *UbiComp 2006: Ubiquitous Computing* (pp. 404-421). Springer Berlin Heidelberg.
- Rowan, J., & Mynatt, E. D. (2005). Digital family portrait field trial: Support for aging in place. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 521-530). ACM.
- Sadri, F. (2011). Ambient intelligence: A survey. *ACM Computing Surveys (CSUR)*, 43(4), 36.
- Scott, J., Bernheim Brush, A. J., Krumm, J., Meyers, B., Hazas, M., Hodges, S., & Villar, N. (2011). PreHeat: controlling home heating using occupancy prediction. In *Proceedings of the 13th international conference on Ubiquitous computing* (pp. 281-290). ACM.
- Shih, P. C., Han, K., Poole, E. S., Rosson, M. B., & Carroll, J. M. (2015). Use and adoption challenges of wearable activity trackers. *iConference 2015 Proceedings*.
- Simm, W., Ferrario, M. A., Friday, A., Newman, P., Forshaw, S., Hazas, M., & Dix, A. (2015). Tired energy pulse: exploring renewable energy forecasts on the edge of the grid. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 1965-1974). ACM.
- Strengers, Y. (2014). Smart energy in everyday life: Are you designing for resource man?. *interactions*, 21(4), 24-31.
- Strengers, Y. A. (2011). Designing eco-feedback systems for everyday life. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2135-2144). ACM.
- Stumpf, S., Rajaram, V., Li, L., Burnett, M., Dietterich, T., Sullivan, E., ... & Herlocker, J. (2007). Toward harnessing user feedback for machine learning. In *Proceedings of the 12th international conference on Intelligent user interfaces* (pp. 82-91). ACM.
- Stumpf, S., Rajaram, V., Li, L., Wong, W. K., Burnett, M., Dietterich, T., ... & Herlocker, J. (2009). Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies*, 67(8), 639-662.
- Suchman, L. (2006). *Human-machine reconfigurations: Plans and situated actions*. Cambridge University Press.

- Takayama, L., Pantofaru, C., Robson, D., Soto, B., & Barry, M. (2012). Making technology homey: finding sources of satisfaction and meaning in home automation. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (pp. 511-520). ACM.
- Tintarev, N., & Masthoff, J. (2007). Effective explanations of recommendations: user-centered design. In *Proceedings of the 2007 ACM conference on Recommender systems* (pp. 153-156). ACM.
- Tullio, J., Dey, A. K., Chalecki, J., & Fogarty, J. (2007). How it works: a field study of non-technical users interacting with an intelligent system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 31-40). ACM.
- Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management science*, 46(2), 186-204.
- Vermeulen, J., Luyten, K., van den Hoven, E., & Coninx, K. (2013). Crossing the bridge over Norman's Gulf of Execution: revealing feedforward's true identity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1931-1940). ACM.
- Weiser, M. (1991). The computer for the 21st century. *Scientific american*, 265(3), 94-104.
- Weiser, M. (1999). The computer for the 21st century. *Mobile Computing and Communications Review*, 3(3), 3-11.
- Weiser, M., & Brown, J. S. (1997). The coming age of calm technology. In *Beyond calculation* (pp. 75-85). Springer New York.
- Yang, R., & Newman, M. W. (2012). Living with an intelligent thermostat: advanced control for heating and cooling systems. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (pp. 1102-1107). ACM.
- Yang, R., & Newman, M. W. (2013). Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing* (pp. 93-102). ACM.
- Yang, R., Newman, M. W., & Forlizzi, J. (2014). Making sustainability sustainable: Challenges in the design of eco-interaction technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 823-832). ACM.
- Youngblood, G. M., & Cook, D. J. (2007). Data mining for hierarchical model creation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 37(4), 561-572.