

Learning from a Learning Thermostat: Lessons for Intelligent Systems for the Home

Rayoung Yang and Mark W. Newman

School of Information
University of Michigan
Ann Arbor, MI 48109 USA
{rayang, mwnewman}@umich.edu

ABSTRACT

Everyday systems and devices in the home are becoming smarter. In order to better understand the challenges of deploying an intelligent system in the home, we studied the experience of living with an advanced thermostat, the Nest. The Nest utilizes machine learning, sensing, and networking technology, as well as eco-feedback features. We conducted interviews with 23 participants, ten of whom also participated in a three-week diary study. Our findings show that while the Nest was well-received overall, the intelligent features of the Nest were not perceived to be as useful or intuitive as expected, in particular due to the system's inability to understand the intent behind sensed behavior and users' difficulty in understanding how the Nest works. A number of participants developed workarounds for the shortcomings they encountered. Based on our observations, we propose three avenues for future development of interactive intelligent technologies for the home: *exception flagging*, *incidental intelligibility*, and *constrained engagement*.

Author Keywords

Smart home; intelligent systems; sustainability.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

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 [28], 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 [30], 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., [7,8,22,23]) as well as commercial offerings [30]. In late 2011, the Nest thermostat [31] was introduced to the market and received a great deal of media attention [32]. 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., [9,10]), or with prototypes in experimental settings (e.g., [8,22]). 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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UbiComp'13, September 8–12, 2013, Zurich, Switzerland.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-1770-2/13/09 \$15.00.

<http://dx.doi.org/10.1145/2493432.2493489>

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., [16,20]). 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., [2,14,26]), 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., [3,9,10,15]), but few such projects have provided insight into the lived experience of occupants. A notable exception is Mozer's Adaptive House [15], 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 need 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 [5]. When considering adaptive home systems that utilize sensing and machine learning, issues of intelligibility and control become central to the concept of “manageability” [1,5]. 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 [27]. While extensive research has been done into how to design interfaces that render system behavior more *intelligible* [1,12,24], 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 [6], 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., [7]), which will, in turn, motivate people to save more energy. Strengers et al. [23], however, pointed out that obtaining information did not always cause people to take action or change their behaviors. Previous studies [18] 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 [8,22]. 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 [31].

The Nest features an attractive wall-mounted device, as well as smart phone and web-based control capabilities (Figure 1). 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.

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




Figure 1. 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).

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 METHOD 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.

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

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	Parti-cipant(s)	Adults (Children)	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

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 [33], 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 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 [32]. 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 [29]. 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.



Figure 2. 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.

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 2) 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."

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 3.

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



Figure 3. P2 posted screenshots of his schedule before and after he modified it.

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 [25] 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., [11,24]) or on eliciting more and higher quality input from the user (e.g., [4]). 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., [11,12,24]). 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 [13] and collaborative systems [17], 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 [28]. 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 [21]. 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.

Limitations

Our goal in this paper has been to illuminate the principles for designing intelligent systems for the home. While we have argued that the commercial deployment of the Nest has provided a valuable opportunity for studying this issue, our study is limited by the nature of the technology studied and the characteristics of our participants.

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 study 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.

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 [2,14]. 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 [20], as well as that of computer expertise and identity [19]. 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.

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.

ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation (0905460) and the University of Michigan Rackham Graduate Student Research Grant. We thank our participants and members of Michigan Interactive and Social Computing for their feedback and support. Mark Ackerman, Tom Finholt, and Tawanna Dillahunt provided helpful comments on earlier drafts of this paper.

REFERENCES

- Bellotti, V., and Edwards, W. Intelligibility and accountability: human considerations in context-aware systems. *Human-Computer Interaction* 16, 2-4 (2001), 193–212.
- Brush, A., et al. Home automation in the wild: challenges and opportunities. In *Proc. CHI 2011*, 2115–2124.
- Cook, D.J., et al. MavHome: An agent-based smart home. *PerCom 2003*, 521–524.
- Dey, A.K., Rosenthal, S., and Veloso, M. *Using Interaction to Improve Intelligence: How Intelligent Systems Should Ask Users for Input*. Presented at the *Workshop on Intelligence and Interaction: IJCAI 2009*.
- Edwards, W.K and Grinter, R.E. At home with ubiquitous computing: seven challenges. In *Proc. Ubicomp 2001*, 256-272.
- U.S. Energy Information Administration, Annual Energy Review, September 27, 2012
- Froehlich, J., Findlater, L., and Landay, J. The design of eco-feedback technology. In *Proc. CHI 2010*, 1999–2008.
- Gupta, M., Intille, S., and Larson, K. Adding gps-control to traditional thermostats: An exploration of potential energy savings and design challenges. In *Proc. Pervasive 2009*, 95–114.
- Intille, S.S. Designing a home of the future. *IEEE Pervasive Computing* 1, 2 (2002), 76–82.
- Kidd, C., et al. The aware home: A living laboratory for ubiquitous computing research. In *Proc. CoBuild 1999*, 191–198.
- Kulesza, T., et al. Fixing the program my computer learned: Barriers for end users, challenges for the machine. In *Proc. IUI 2009*, 187–196.
- Lim, B.Y., Dey, A.K., and Avrahami, D. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *Proc. CHI 2009*, 2119–2128.
- Mackay, W.E. Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica* 30, 1 (2000), 177–193.
- Mennicken, S. and Huang, E. Hacking the natural habitat: an in-the-wild study of smart homes, their development, and the people who live in them. In *Proc. Pervasive 2012*, 143–160.
- Mozer, M.C. Lessons from an Adaptive Home. In D.J. Cook and S.K. Das, eds., *Smart Environments*. John Wiley & Sons, Inc., 2005, 271–294.
- O'Brien, J., et al. At home with the technology: an ethnographic study of a set-top-box trial. *ACM ToCHI* 6, 3 (1999), 282–308.
- Orlikowski, W.J. The Duality of Technology. *Organization science* 3, 3 (1992), 398–427.
- Peffer, T., et al. How people use thermostats in homes: A review. *Building and Environment*, (2011).
- Poole, E.S., et al. Computer help at home: methods and motivations for informal technical support. In *Proc. CHI 2009*, 739–748.
- Rode, J.A., Toye, E.F., and Blackwell, A.F. The fuzzy felt ethnography—understanding the programming patterns of domestic appliances. *Personal and Ubiquitous Computing* 8, 3-4 (2004), 161–176.
- Rogers, Y. Moving on from Weiser's vision of calm computing : engaging UbiComp experiences. In *Proc. Ubicomp 2006*, 404–421.
- Scott, J., et al. PreHeat: controlling home heating using occupancy prediction. *Proc. Ubicomp 2011*, 281–290.
- Strengers, Y.A.A. Designing eco-feedback systems for everyday life. In *Proc. CHI 2011*, 2135–2144.
- Stumpf, S., et al. Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies* 67, 8 (2009), 639–662.
- Suchman, L. *Human-machine reconfigurations: Plans and situated actions*. Cambridge University Press, 2006.
- Takayama, L., et al. Making Technology Homey: Finding Sources of Satisfaction and Meaning in Home Automation. In *Proc. Ubicomp 2012*, 511-520
- Tullio, J., et al. How it works: a field study of non-technical users interacting with an intelligent system. In *Proc. CHI 2007*, 31–40.
- Weiser, M. The computer for the 21st century. *Scientific American* 265, 3 (1991), 94-104.
- Yang, R. and Newman, M.W. Living with an intelligent thermostat: Advanced control for heating and cooling systems. Presented at the *HomeSys 2012 Workshop: Ubicomp 2012*.
- 10 Smart Digital Appliances You Wish You Owned - CNBC. <http://www.cnbc.com/id/46807536>.
- Nest | Home. <http://www.nest.com/>.
- Nest | Reviews. <http://nest.com/reviews/>.
- Catch.com. <https://catch.com/>.