It is too Hot: An In-Situ Study of Three Designs for Heating

Alper T. Alan[1], Mike Shann[2], Enrico Costanza[1], Sarvapali D. Ramchurn[1], Sven Seuken[2]

[1]University of Southampton
Southampton, United Kingdom
{ata1g11, ec, sdr1}@ecs.soton.ac.uk

[2]University of Zurich
Zurich, Switzerland
{shannm, seuken}@ifi.uzh.ch

ABSTRACT
Smart energy systems that leverage machine learning techniques are increasingly integrated in all aspects of our lives. To better understand how to design user interaction with such systems, we implemented three different smart thermostats that automate heating based on users' heating preferences and real-time price variations. We evaluated our designs through a field study, where 30 UK households used our thermostats to heat their homes over a month. Our findings through thematic analysis show that the participants formed different understandings and expectations of our smart thermostat, and used it in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Based on the findings, we present a number of design and research implications, specifically for designing future smart thermostats that will assist us in controlling home heating with real-time pricing, and for future intelligent autonomous systems.

Author Keywords
Heating; field study; design; smart grid; real-time pricing; autonomous systems; internet of things.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
We live in a world where smart energy systems form an important part of our lives, not only within workplaces, but also at home [9]. These systems sense information about their users, learn from this information, and exploit what they have learned to make decisions on their users' behalf. Examples include smart thermostats that learn households' occupancy behaviour for automating heating of their house (e.g., Nest). Consistent with the improvements in these technologies, our ways of consuming energy change day by day.

Meanwhile, the energy market is undergoing significant changes with penetration of renewables, including wind turbines and solar farms [23]. Energy production from these renewables fluctuates depending on the weather. Such variable energy generation puts more strain on the energy market to balance supply with demand. One way to overcome this issue is to use real-time pricing schemes, where energy prices vary over short time intervals, typically hourly [3]. These real-time prices significantly change the role of consumers and offer lower energy bills. However, consumers are required to continuously monitor the changing prices to take the advantage of them. While this task might be daunting for people, it is well-suited for a smart system that can learn user preferences, monitor the prices, and respond to them autonomously. Prior HCI studies about renewable resources [4] and real-time pricing [8] have shown that autonomous computing systems can help people shift their loads by providing them with recommendations. However, in these studies the tasks were not fully automated by the systems, therefore user reactions towards agency was limited to the only suggestions received.

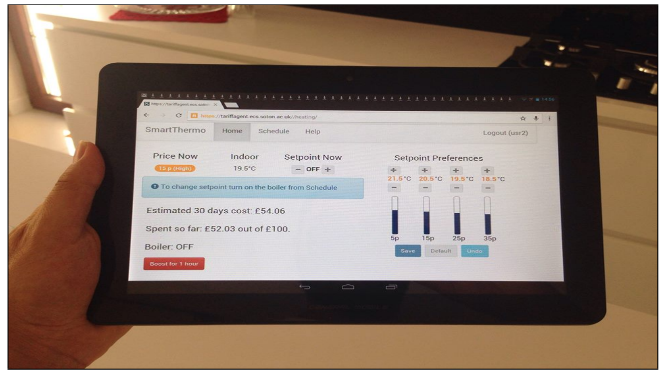

Figure 1. One of our smart thermostat designs running on a tablet.

Against this background, in this study we aim to explore how households adopt to and interact with a home smart thermostat designed to help them manage heating given real-time energy prices. To do so we developed and deployed three different smart thermostats: a *manual* one through which participants explicitly specify how the heating should respond to price changes, and two *learning-based* ones that employ an artificial intelligence (AI) algorithm to automate the temperature settings based on learned households' preferences. We conducted a field study with 30 UK households over a month. As the smart thermostat responded to the varying prices on the households' behalf, it caused a real impact on the comfort of its users. More specifically, in this study, we aim to observe people's feelings and expectations towards a smart thermostat that controls their home heating given real-time prices, and how they interact with such a thermostat in their everyday lives.

CHI'16, May 07-12, 2016, San Jose, CA, USA.
© 2016 ACM ISBN 978-1-4503-3362-7/16/05...$15.00.
http://dx.doi.org/10.1145/2858036.2858222

BACKGROUND

Interdisciplinary knowledge and research is vital to facilitate a better understanding of how user interactions can be designed for smart energy systems. The work we present in this paper lies at the boundary of AI and HCI. Hence, in this section, we briefly review prior literature on models and algorithms for smart heating, field studies of autonomous systems, and interaction with machine learning.

Models and Algorithms for Smart Heating

The combination of mathematical models and computer algorithms form the backbone of making systems smart. The literature includes numerous models and algorithms developed with different approaches for energy efficient heating. Some approaches focused on the models of the environment (e.g., the weather) to create an efficient heating schedule [30, 14]. Other approaches used motion sensing to detect people's presence, and control the heating based on the occupancy models of buildings [21, 11, 15].[1]

Shann and Seuken [22] presented a learning algorithm that elicits users' preferred temperatures for different energy prices and creates a comfort-cost trade-off model for each user. Lam et al. [10] introduced a thermal comfort model that updates based on the user's comfort feedbacks. However, all these studies have simulation based results. They do not have a real smart thermostat that is physically deployed to enable its users to interact with it to control their actual heating.[2]

Field Studies with Autonomous Systems

Conducting field studies has been considered a significant way of evaluating computing systems to better understand the social context and user experience [19]. Increasingly, researchers are taking AI technologies from labs to field deployments, where algorithms are used to bring autonomous systems to life in different forms, such as social robots, virtual agents and software agents. In this respect, Sauppé and Mutlu [20] conducted an ethnographic field study with robots in manufacturing sites, and showed that human workers approached the robot co-workers as social entities.

Yang and Newman [28] examined the real-world uptake of a smart thermostat with 23 participants. They highlighted how sub-optimal decisions taken by a smart thermostat are likely to cause frustration for users and may lead them to abandon the technology. Their follow-up study [29] has investigated users' long term interactions with the smart thermostat. Their findings suggest that users' interactions faded over time and resulted in unrealised energy saving opportunities. They also propose that an alternative design (i.e., a mixed-initiative system) might improve the sustainability of user engagement and the system's usefulness.

Bourgeois et al. [4] deployed energy-aware washing machines that provide users with suggestions on when to do their laundry based on the availability of green energy. They studied various intervention techniques with 18 households for 8 months, and showed that proactive suggestions sent by a software agent via text messages are more effective than the agent's email interventions. Similarly, Costanza et al. [8] proposed 'Agent B', a software agent that also helps users book their washing machine in a scenario where electricity prices change every 15 minutes. In a field experiment, 10 participants used Agent B for one month. The results suggest that Agent B helped users defer their laundry in response to real-time prices.

Alan et al. [1] proposed 'Tariff Agent', an agent that helps users select electricity tariffs on a daily basis. It provides three levels of autonomy (fully autonomous, semi-autonomous, and manual), and users can change the level at any time. In a field experiment with 10 users, Tariff Agent was used over a period of 12 days. The results indicate that people are willing to delegate some decisions to an agent but at the same time there is also a desire to stay in control. In a 6-week follow-up study with 12 users, they showed that the users took responsibility for undesired outcomes if the system's autonomy level could be adjusted flexibly [2].

Our study builds on this literature, but extends it in two key ways. First, while prior studies only made the autonomous operation of the systems tangible to participants only through financial rewards, in our study the system also directly influences the participants' thermal *comfort* by controlling the heating setpoint given real-time prices. Second, we evaluate three different UI designs for smart thermostats in the wild.

Interaction with Machine Learning

HCI researchers have long been using machine learning techniques in their studies [13]. The main area of these studies has focused on classification methods, and visualisation of machine learning algorithms. The underlying aim is to improve accuracy of classifiers through enabling users to explore data and interact with machine learning algorithms [26]. However, even though machine learning is becoming an important part of our everyday lives, and significantly changing the way we live (e.g., smart thermostats), there have been few studies that examine how users could interact with machine learning systems, beyond accuracy judgements [25].

In fact, to date no studies have examined how different designs of learning systems may impact people's perceptions. There is a significant gap in our understanding of how we should design interactions with machine learning systems, especially for the ones that might possibly intrude upon our daily activities. In this paper, we present a field study that aims to close this gap by focusing on households' perceptions of and interactions with two designs of a smart thermostat exploiting machine learning, and draw implications for the design of future smart energy systems.

SMARTTHERMO IN THE WILD

We aim to explore how to best design user interactions with a smart thermostat that is designed to automate home heating control when energy price varies in real-time. To do so we prototyped a smart thermostat based on envisioning [17], and evaluated it with a field study. To make the experience of the smart thermostat relevant to the participants of our

[1] Indeed, there are commercial thermostats that utilise occupancy detection (e.g., tado.com).)

[2] We note that Lam et al. deployed a mobile system but it was only for getting users' comfort feedbacks.

Table 1. Participants' profiles.

	Thermostat	Age	Occupation	Others
P1	Indirect L.	40	PhD Student	1 Child
P2	Manual	63	Maintenance Eng.	1 Adult
P3	Direct L.	37	Antiques Dealer	1 Ad., 2 Ch.
P4	Direct L.	43	PhD Student	1 Ad., 2 Ch.
P5	Manual	50	Estate Mng.	1 Ad., 1 Ch.
P6	Manual	36	Nanny	1 Ad., 1 Ch.
P7	Indirect L.	76	Retired	1 Adult
P8	Indirect L.	32	Radiographer	2 Ad., 1 Ch.
P9	Direct L.	44	Teacher	1 Ad., 2 Ch.
P10	Direct L.	62	Retired	2 Adults
P11	Manual	44	Teacher	1 Ad., 2 Ch.
P12	Indirect L.	50	Office Mng.	1 Adult
P13	Manual	40	Education	2 Ad., 1 Ch.
P14	Indirect L.	60	Lecturer	1 Adult
P15	Manual	53	Photographer	2 Ad., 1 Ch.
P16	Direct L.	60	Self Employed	1 Adult
P17	Indirect L.	58	Charity Mng.	1 Adult
P18	Manual	40	Accountant	1 Ad., 4 Ch.
P19	Direct L.	71	Retired	1 Ad.
P20	Indirect L.	56	Database Admin.	2 Adults
P21	Indirect L.	26	Contract Mng.	1 Ad., 1 Ch.
P22	Direct L.	22	Student	1 Adult
P23	Indirect L.	28	Sport Mng.	2 Adults
P24	Direct L.	69	Retired	1 Adult
P25	Direct L.	49	Gas Engineer	1 Adult
P26	Manual	64	Engineer	1 Adult
P27	Manual	91	Retired	1 Adult
P28	Direct L.	73	Retired	1 Adult
P29	Manual	75	Retired	Na
P30	Indirect L.	28	PhD Student	1 Ad. , 1 Ch.

study, we defined a real-time pricing scenario, based on actual historical spot prices in the UK electricity market in January 2014.[3] Following a study design proposed in prior work [8], this scenario was made tangible to participants through financial rewards, as detailed further below. For convenience, we removed extreme outliers from the historical pricing data making the prices range from 5 pence to 35 pence. During our field study the energy price was changed every 30 minutes, similarly to the UK market. In the following sections, we present the details of our study, including our study procedure and participants, descriptions of the thermostat designs, and the data collection and analysis methods we used.

Study Design

In order to gain an understanding of how to best design a smart thermostat for real-time prices, we decided to explore three thermostat designs:

1. Manual: This design aims to provide manual operation and involves no machine learning algorithm. Hence, in this design, users are required to manually specify how the temperature is going to be set at different prices through adjusting a number of setpoint sliders.

2. Direct Learning: This design uses the machine learning algorithm introduced in a prior work [22],[4] and aims to automate users' temperature decisions for different prices.

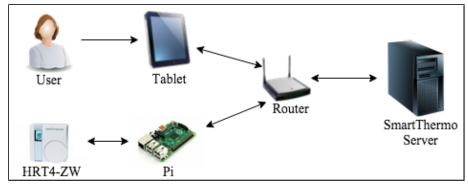

Figure 2. Overall system diagram.

When users make changes in the temperature, the learning algorithm correlates these changes with the prices and generates a user model. Thus, rather than requiring the user to manually specify the setpoint sliders, the learning algorithm automatically arranges them. Each time the user submits a temperature, the algorithm updates the user's model and the thermostat directly heats to the optimal temperature of the user model based on the current price. The aim of this design is to help users understand that the setpoints that they save are being learned by the smart thermostat for future use to determine the setpoint based on varying prices.

3. Indirect Learning: Similar to the direct learning thermostat, the same learning algorithm is used in this design. However, the rationale behind this design is to enable users to temporarily override the learning, and in this way hide from the users the complexity of the algorithm. Thus, in this design, each time the user submits a temperature, the algorithm updates the user model but the thermostat first heats to the inputted temperature - rather than heating to the optimal temperature of the model. Though, after one hour it goes back to auto-mode and sets the setpoint to the optimal temperature of the user model based on the then current price.

We conducted a study with 30 UK households (see Table 1) over a period of four weeks during February-March 2015. To recruit participants we distributed approximately 3000 study invitation letters around the city. We recruited households who had a broadband Internet connection and a central heating control, based on a first-come first-served basis. Participants were assigned to three groups, each corresponding to one thermostat design, one by one in the order 1, 2, 3. An online budget of £100 was then allocated to each household and participants started to use our system for heating their house. Their heating cost is calculated based on the number of hours their boiler was on, and subtracted from their online budget on each day. After four weeks, when the study ended, they received the amount left in their budget as experimental reward. By so doing, we aimed to encourage participants to respond to the prices, and make savings have a tangible impact. The idea of using monetary incentives to simulate real-time pricing is inspired by an early study where participants received payments of the value of electricity saved [24].

Technology

We equipped each household with a Horstmann HRT4-ZW thermostat, a Raspberry Pi (RPi) and an Android 4.4 tablet. Figure 2 shows the connections among different entities. The Horstmann thermostat is a standard room thermostat but can

[3]For practicality we recorded the prices about a year earlier than our study took place, that is in February-March 2015.

[4]There might be more advanced algorithms giving better results. We chose this algorithm due to its simplicity and robustness.

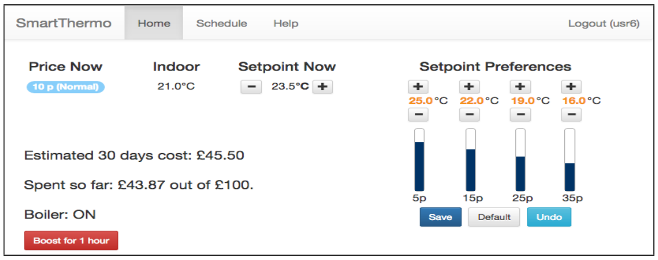

Figure 3. Manual - home page.

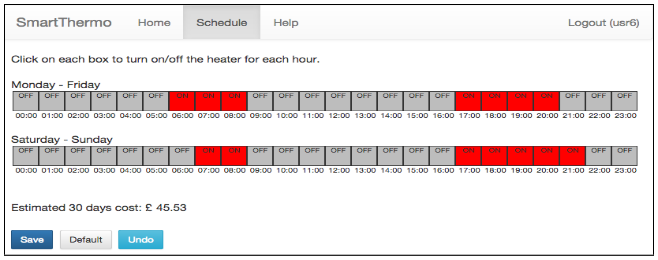

Figure 4. Manual - schedule page.

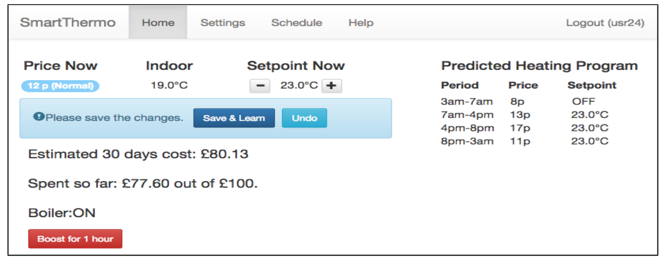

Figure 5. Direct learning - home page.

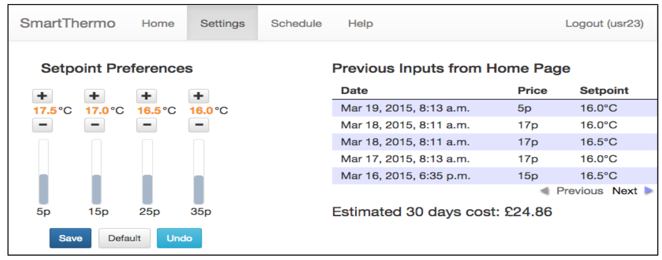

Figure 6. Direct learning - setting page.

be wirelessly controlled over the Z-Wave communication protocol from the RPi (through a RaZberry daughter card[5]). The RPi also connects through the home wireless broadband router to our web server, where the smart thermostat algorithm and UIs run. The RPi regularly pulls the indoor temperature from the thermostat (every 5 minutes), sends the temperature data to our server, and receives the latest individual heating plan. Based on the plan the RPi then controls the setpoint of the thermostat. The tablet allows participants to access our web application through the broadband connection, and to manipulate their own heating plan. Each tablet was installed with a software called Kiosk Browser Lockdown and our web application was set as the default one. We also added the application as a bookmark on the home screen of participants own devices (tablet or smart phone), if they wished.

Design Variations

Manual Thermostat

On the home page (Figure 3), users can see the current energy price and indoor temperature, and adjust the setpoint by pressing the *+/- buttons* next to it. Each press increases/decreases the setpoint by half a degree. To provide context for the current value, a label indicates whether the price is *normal* (bottom of the range), *high* (mid-range) or *very high* (top of the range). The price value and the label are color coded green, yellow or red for emphasis. At the bottom left of the home page users can find the *boost button*, which allows them to turn the heating on continuously for 1 hour, temporarily overriding the setpoint.

On the right side of the page, four *setpoint sliders* enable users to specify how the setpoint should be changed at different prices. In other words, these sliders allow users to directly specify how to trade off comfort and cost. These are positioned on the home page to make them easily visible and accesible, even at the risk of increasing the complexity of the page.

The schedule page (Figure 4) is another page that the manual thermostat users could access. This page allows the users to program the heating schedule that defines the boiler's on and off times. Due to the screen size of our tablets we decided to divide the schedule of a day into hourly-based time slots and group the days as weekdays and weekend. To change the boiler's status for a period of time the user only needs to touch on the time slots corresponding to the period. We provided this schedule page since we anticipated that users would expect such a functionality from a smart thermostat.

Both the home and the schedule pages display the 'Estimated 30 days cost' that reflects how the current settings on the setpoint sliders and the schedule impact the monthly cost of heating. When users make a change in the sliders or in the schedule, the estimated cost updates accordingly. Also it is important to note that users need to save any changes that they make in order to register that change into the system.

Direct Learning Thermostat

In this thermostat design (Figure 5), users directly interact with the machine learning algorithm. When the user presses the +/− buttons, the learning algorithm updates the user's model and displays the optimal setpoint based on the model. The algorithm uses Bayesian inference to update the model, which means it considers the user's individual temperature inputs as noisy data. Thus, the user might need to press the +/− buttons several times to change the setpoint a half degree, depending on the model's prior knowledge. However, to visualise the impact of each press, a pop-up message appears with two buttons, by which users can save or undo the setpoint change. Additionally, each press synchronously affects the 'Estimated 30 days cost' as well as a table called 'Predicted Heating Program'. This table shows the average temperatures that will be set by the thermostat, based on the

[5]http://razberry.z-wave.me/

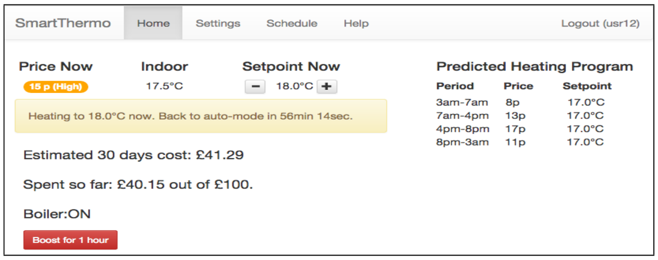

Figure 7. Indirect learning - home page.

predicted average prices at four time periods. The setpoint displayed for each time period changes dynamically according to the updates in the user model. Similar to the manual thermostat, the home page also contains the boost button, which turns the heating on continuously for 1 hour. The boost button does not influence the learning: it was designed as a way to define exceptions to the preferences.

With this thermostat design, users also have one additional page called settings. The settings page (Figure 6) aims to provide users an additional level of control and transparency on the learning algorithm. Similar to the home page of the manual design, there are four sliders representing the user's learnt temperature preferences for each price band. We moved the sliders into the settings page because the focus of the learning thermostat is simplicity of use. The user can see how these sliders are arranged by the thermostat by looking at a history table. The table lists the user's previous temperature inputs together with correlated prices. The user can adjust the sliders to specify his own preferences. By so doing, the user resets the learning algorithm and clears the table of previous inputs. Therefore, a confirmation pop-up is shown before the user saves any changes made in the sliders. The schedule page provided in the manual thermostat is also accessible by the users of the direct learning thermostat.

Indirect Learning Thermostat

In this design, the temperature input that users provide temporarily overrides the setpoint that the algorithm would set based on the user's model and the current price. Thus, the $+/-$ buttons work exactly in the same way as in the manual thermostat (each press increases/decreases the setpoint by half a degree). Once the user saves the new setpoint, the algorithm updates the user model based on the new input. Then, it waits for an hour to take the control back and change the setpoint to the learned one according to the then current price. Meanwhile, the thermostat heats to the inputted temperature. This process was explained with a pop-up message including a countdown timer, starting from 60 minutes, in the UI (Figure 7). Users can still change the setpoint and save it before the countdown finishes, which will restart the count down with the new setpoint. As in the direct learning thermostat, the home page also contains a boost button, which does not influence the learning. The settings and schedule pages are also provided for the indirect learning thermostats.

Method

During the course of the study, we recorded all users' interactions with the thermostat application. For instance, we recorded when participants changed the setpoint or the heating schedule, or when they used the boost button. Additionally, we collected detailed quantitative data about the heating habits of each household, including the temperatures the users set in response to real-time prices and how indoor temperatures varied over the course of the study. However, it was difficult to derive conclusions about the impact of the different thermostat designs on these data, since there are other factors affecting people's heating preferences (e.g., weather and home insulation) [16]. Therefore, in this paper, we focused on the qualitative data collected during the interviews.

Interviews

We conducted semi-structured exit interviews with family members at their homes. We interviewed 26 households.[6] The interviews were mostly held with the participant that signed the consent form at the beginning of the study, however some interviews also involved the participant's partner. In the interviews we asked participants open questions about their use, adoption and understanding of the thermostat. All interviews were audio-recorded, and lasted on average of 34 minutes (SD: 8 minutes, min: 18 minutes, max: 52 minutes).

Analysis

The interviews were fully transcribed and analysed through thematic analysis [5]. Four researchers were involved in this, while the coding was performed by two researchers. The analysis started by categorising the material at the sentence level through open codes. Initially 93 open codes were used, later grouped in broader categories that we discuss in the following section.

FINDINGS

In this section we first present an overview of the quantitative analysis we performed on the overall system usage, based on the automatic interaction logs. Secondly, we report the major findings of our thematic analysis. The analysis revealed six key themes: (1) orientation towards the thermostat's agency, (2) reactions to different UI features, (3) managing the home heating with real-time prices, (4) mental models of the thermostat's learning feature, (5) balancing cost and thermal comfort, and (6) limitations of the thermostat's learning model. We present the categories that revealed these themes in the following subsections. In excerpts, we use "F" for female and "M" for male to denote the gender of the household member.

Overview of Quantitative Analysis

Users of each thermostat design heated their home using our system for a month. They mostly interacted with the system via the tablet we provided or the tablet they already had. Some of these participants additionally used their mobile phones to access the system. Table 2 includes the overall data of system usage. We performed one-way ANOVA tests on the quantitative data. However, we could not find any significant differences or long-term effects in user interactions across the three deployed designs. This was also the case for the analysis of other quantitative data collected (i.e., setpoints and indoor temperatures). The analysis did not reveal

[6] Other 4 participants were not available for the interview.

Table 2. Overall quantitative data analysis.

	Manual		Direct L.		Indirect L	
	M	SD	M	SD	M	SD
Setpoint Changes from Home Page	36.7	52.8	13.3	14.4	18.5	22.8
Setpoint Changes from Settings Page	-	-	6.4	7.4	2.5	3
Schedule Changes	18.2	14.8	47.8	51.3	21.8	19.4
Boost Activations	17.5	14.1	7.3	4.3	14.2	11.3
Spent from Budget	£55	£22	£55	£33	£30	£15
Demand Response	36%	19%	34%	17%	47%	24%

Note. M = Mean. SD = Standard Deviation.

any significant differences in users' demand-responses or savings, which might be understandable given the interpersonal, contextual, and environmental differences of the users.

Orientation towards the Thermostat's Agency

All participants commented that they were happy with the thermostat autonomously responding to real-time prices on their behalf. The following is a typical response that we received in the interviews, when we asked participants about their feelings towards the agency of the system.

> **P18 (M):** "I'm happy with that if the thermostat understands that at this price I would rather avoid heating the house, and at this price I would like to heat the house, then I'm happy for it to take over that control, as long as it's very straightforward for me to override."

We observed that they felt in control overall and were also mostly confident with the way the thermostat was working.

> **Analyst:** "To what degree did you feel like the system worked for you, or it required you to do the work?"

> **P13 (F):** "So basically I just order the system to do the things for me and the system does the whole thing. "

Reactions to Different UI Features

In the interviews we observed that most participants understood well how to use the UI elements (e.g., the +/− buttons) of each thermostat design. Nearly all participants commented that the thermostat was easy to understand and use. All participants appeared to understand the functionality of the setpoint sliders and mostly appreciated their use.

The users of the direct learning thermostat were mostly aware of the fact that sometimes they were required to press the +/− buttons multiple times to achieve the desired setpoint value. However, they did not explicitly state that this was due to the learning feature.

> **P3 (F):** "The estimated cost would change before the degree thing changed. So, you press it. Sort of, like, Wow it needs four presses per half degree or something, and I was like, because I could see this number here was changed. It was doing something. Then I thought it must be incremental, must be in tenths rather than in halves."

The indirect learning thermostat users mostly explained the way the 60 minutes countdown works as though it goes back to the previously saved setpoint rather than the learned setpoint based on the price at the time after 60 minutes.

> **P23 (F):** "So then there's a countdown for 60 minutes and after that the temperature will resume to what was set previously. Occasionally I would reset the temperature again within those 60 minutes."

Managing the Home Heating with Real-time Prices

Here we detail how and why participants used the thermostat in different ways to heat their house with real-time prices.

Setpoint Preferences

Most participants of all three thermostat designs reported preferring to change 'setpoint now' from home page to control the indoor temperature with real-time prices. They also fiddled with the setpoint sliders, but the number of times was relatively low compared to the changes made in the 'setpoint now'. Interviews revealed that most participants were happy with the arrangement of the sliders and therefore they did not feel the need to alter them often. Also a few participants found the sliders complex, which led them to play with the 'setpoint now' more.

We had constrained the setpoint sliders to present a straight line corresponding to users' heating preferences. Therefore, changes made to one of the sliders affected the values of others. Overall the participants who interacted with the sliders found them easy to adjust and appreciated their use. In the interviews, only one participant griped about this linear relationship among the sliders. However, we noticed that this user only played with the first slider throughout the study, therefore they could make only parallel shift on the slider values without being able to change the slope of the line.

> **P15 (F):** "If we could've adjusted them differently and made our own decisions on these rather than they just go up automatically when you change one of the others, we would've preferred that."

Most participants kept the configuration of the sliders in descending order, starting with higher temperature at lower price and lowering the setpoint as the price increases. Two users calibrated the four sliders to have the same setpoint. In other words, they opted for a specific temperature setpoint to heat their house over the course of the study regardless of the heating cost.

> **P29 (M):** "You should be prepared to pay more, a higher rate, if you wanted to be more comfortable. I could change it as when I wanted it, but if I wanted to go to a higher temperature, it could cost me more. But, because I had set it at a flat rate, I wasn't bothered."

Also, participants reported that they did not need to change the sliders once they found their limit for how much comfort they could sacrifice to save money. The process of finding such a limit was generally a matter of trial and error:

> **P11 (F):** "It was freezing cold and it must've broken. I checked and that's when I saw it was on 35p and that's when I changed the lowest set point."

Some participants were more conscious of and certain about their tolerance limit for temperature even in the early stages of the field study. Therefore, once these participants arranged the sliders early on in the study, they stopped interacting with the sliders, and used other interface features, such as the schedule, to adjust the heating.

> **P16 (M):** "When we first got it, we looked at the pricing bands and made some decisions at that stage. We did it once and I don't think we revisited it. What we did visit, then, pretty regularly, probably every day, and maybe more often than once a day, we did revisit the schedule."

Schedule

Most participants told us how it was easier to access and change the heating program through our system compared to their previous programmable heating controls. Being able to easily turn on and off the heating by touching on the displayed hourly time slots, and being able to have different programs for weekdays and weekends, seemed to meet participant's favour.

> **P21 (F):** "We changed it most days because it was so easy to access. If I was going out and I knew that we wouldn't be home until five, I'd set it to come on at four. Whereas previously, we wouldn't even touch it on a normal one."

Here she refers to the "normal one" as a wall-mounted programmable thermostat that she had before taking part in our study. Her comment suggests that people may engage more with their heating systems when the systems are easy to control and access. Specifically, among the users of all UI designs, most participants liked being able to control the heating remotely from anywhere in the house or anywhere outside via their Internet-connected devices - rather than walking to a wall-mounted thermostat each time.

> **P23 (F):** "It was just useful to be able to change the temperature from wherever I was really. I could do it from work and quite often did or if I went from work to the supermarket and then came home, you could do it from the supermarket. So yes, it was really *clever*."

In the interviews, we observed that occupancy was the major factor affecting the way participants modified the schedule. They mostly tended to turn off the heating when no one was around, and turn it on if someone was at home. Participants who had regular lifestyles reported that they didn't need to change the schedule often, whereas some participants altered the schedule quite often due to their irregular lifestyles.

> **P28 (M):** "I did change this one [schedule] at the very beginning, but other than that I haven't touched it at all because, I've been working with this time limits for the boiler on and off for 25 years. It's suited my lifestyle. I'm a creature of habit really."

> **P10 (F):** "I'm retired so could be at home all day but at the last minute suddenly go off somewhere and we're three adults. So it's three people leading separate lives in a way rather than if we were a family with children and you'd know you would be in the house until half past eight go to school picking up. So our lifestyle is quite *erratic*."

We also observed that some participants used the schedule as a medium to respond to changing energy prices. For example, in the following quote, P9 (F) indicates that she moved the time that the thermostat normally comes on in the mornings nearly one hour earlier to benefit from lower prices.

> **Analyst:** "How do you feel about the real-time prices for heating energy?"

> **P9 (F):** "I noticed that it [price] was cheaper before 7am. Previously, I'd been putting the heating on like maybe 6:45 because we get up about 7:00, then leaving it on while we're getting ready for work and school and then turning it off. I changed that and started putting it on earlier, putting it on at 6:00 and then having it go off at 7:00, and it still kept the house warm enough until we went out sort of an hour or so later."

There were other factors that influenced participants' heating program, such as their daily activities or weather conditions. Participants mostly turned the heating on at times that they usually took showers, or turned it off when they used the oven. While on cold days participants arranged the schedule to make the thermostat come on for more time slots, they had fewer time slots on for milder days.

Boost

After deciding how to balance comfort and cost, participants tended to use the "boost" button for exceptional situations to turn the heating on instead of changing the setpoints on the sliders.

> **P2 (M):** "I tended to be comfortable with the settings that I had on it and sort of left it. The only time if it was really cold in the mornings when we got up, I'd press the boost to boost it and it probably only went on for an hour or two."

Some participants also commented that they used the boost button just to heat their home a bit more when the prices were lower.

> **P21 (F):** "I liked when it said £0.05 and I was like; yes, put the heating on, boost it!"

Mental Models of the Thermostat's Learning Feature

Only the users of direct and indirect thermostats were exposed to the machine learning algorithm. These users were required to click the save and learn button that appears every time they make a change in the setpoint from home page in order to register their preferred temperature into the thermostat. Including the text 'learn' in the save button seemed to be successful at conveying the fact that the thermostat was learning. However, when we asked the users' opinion about what the thermostat was learning in the interviews, three users reported that they had not thought about it before and therefore that they had no comment.

Among the participants who formed opinions about the learning feature, most participants appeared to have an understanding that is well-matched with the actual underpinnings of the thermostat's learning feature. Most participants were aware of the fact that the thermostat was trying to correlate their preferred setpoints to varying prices. It seemed that the display of previous inputs in the settings page supported their comprehension. Though, conceivably, no one seemed to be interested in how the thermostat was actually calculating the setpoint based on their previous inputs

> **Analyst:** "If you had to explain the learning feature to one of your friends, how would you explain it?"

> **P3 (F):** "It [thermostat] learns your tolerance for an increase in price. It learns your habits and your behaviours in terms of the price versus the temperature, and then it applies those, reapplies them for future events when the unit price goes up."

> **P23 (M):** "As you input your set point changes according to the prices and then the system starts to understand what your views are of that cost I suppose. That is what you think is expensive and that is what you think is cheap, and then make changes."

In these quotes, the participants are very clear about what the thermostat was trying to learn. They explain that the thermostat was learning their temperature preferences for different prices based on their previous inputs. Also, the participants express that the thermostat was learning in order to be able to autonomously respond to the changing prices on their behalf.

On the other hand, some participants had another interesting mental model description of the learning feature, which was neglecting the effect of the prices. These participants described the thermostat as though it was matching their preferred temperatures with the times and the days of the previous temperature inputs that they provided.

> **Analyst:** "So can you tell me what happens when you click to the save and learn button after you change the setpoint?"

> **P21 (M):** "Well, it [thermostat] updates and it changes the kind of the setpoint to what it is going to heat to, but it also learns what you have done. So, I am guessing that later on, if you are doing that at a certain point every day then it's going to learn that."

> **P30 (M):** "If I play a particular temperature as the setpoint and then click on save and learn, from what I understand is the system will take this reading to consideration for whether to turn the boiler on or off but at the same time try to see that at this particular time of the day, whether it's weekday or weekend and then try to replicate that during other days."

This misinterpretation of the learning feature was more prevalent among the indirect learning thermostat users than among the users of the Direct Learning Thermostat. Further exploration also revealed that none of the participants having the misinterpretation was familiar with the Nest thermostat, or in fact any other smart thermostats that exist in the energy market. We can therefore assume that they were not biased.

Balancing Cost and Thermal Comfort

Even though seeing the current price of energy had mostly impacted on how our participants heated their home over the course of the study, there were other significant factors that played key roles in the decisions of the participants for maintaining their thermal comfort at home. One of these key factors was occupancy. Most of our participants commented that they tended to turn off the heating for the times that no one was at home. Another important factor was outside weather as opposed to the indoor temperature: the colder the weather was, the longer the heating was on. Lastly, daily activities at home seemed to substantially influence the participants' heating preferences. While sitting still or having a shower caused participants to turn on the heating, cooking or other physical activities led them to keep the heating off.

> **P2 (M):** "I generally go out by nine o'clock I had the heating going off at eight o'clock in the morning. So it sort of warmed us up to have our showers and be comfortable in the morning, and weekends it depended whether we were in or out as to whether we left it on or knocked it off. So it revolved around our lifestyle and work patterns and things. And the temperature outside. If it was really cold outside then we would have it on longer."

> **P1 (M):** "Most of the time I tried to connect the schedule with my daily activities. For instance, I take shower in the morning, and sometimes I work here at home between 9 and 11. So these are the times that I turn on the heating. Most of the time between 12 and 3, I cook and turn off the heating, because it really doesn't feel cold."

Another interesting finding that emerged from the interviews was the ways participants attempted to maintain their thermal comfort at home without using our heating system. The most prevalent attempt was putting on one more layer of clothing (generally a jumper), or using a blanket when the energy prices are high. Also some participants took the advantage of their other heating sources such as wood-burning stoves, which is typical in small town houses in the UK.

> **P8 (M):** "I think we have probably spent less on our heating in general than we would have done normally. Normally we heat the house pretty much all the time in the winter. We did at times just put another jumper on."

Limitations of the Thermostat's Learning Model

As it is clear in the previous excerpts, the price was not the only factor affecting our participant's setpoint preferences. However, the learning algorithm used in both direct and indirect learning thermostats was only considering two inputs: the setpoint registered and the price at that current time. Therefore, the thermostat was automating the setpoint control only based on the price. This limited learning capability resulted in dissatisfaction among a few participants since the setpoint automatically set by the thermostat was not always the right temperature for its owner. The following quotes are the only ones from which we received such feedback from the participants.

> **P3 (F):** "There were times when I came in and I was like. Hang on a sec. My house is really warm and it must have been because it had learned something that. To do with the temperature. So, it must have said all the prices are this, so they like it warm when it's like this. It's like. Hell no. It's too hot!"

> **P9 (F):** "Well, if I understood the intention that it was trying to set my temperature according to the price, that didn't really work for me. I kind of wanted a combination. I kind of could see the point of that. But like I said, at night, I didn't want it so warm, though perhaps I quite sort of would like it to keep it a degree or two cooler when the temperature's high to save money or something like that. But I also wanted it to let me decide more and not decide for me all the time."

DISCUSSION

In this section, we revisit the major findings of our study, and discuss them in light of prior literature. We also present implications for interaction design of smart energy systems and for future research.

Designing a Thermostat for Real-Time Prices

Any thermostat designed for real-time prices will need to automate the heating at some level, as otherwise it would be a very difficult task for a human to monitor every price change and alter the heating accordingly. Prior research on autonomous systems suggests that these systems should allow their users to easily override the automated decisions at any point in time, without completely disabling the system's autonomy [1]. In this vein, in our study we observed that some participants used the boost button as a means to temporarily override their temperature preferences for exceptional situations, rather than resetting the learned preferences. These exceptional situations not only occurred when users felt cold and wanted to heat the house despite the high prices, but also happened when users wanted to heat the house a bit more than they would do normally in order to benefit from low prices (typically termed the rebound effect).

One of the most-liked features was the display of 'Estimated 30 days cost'. As P28-m said, "I've watched also my estimated cost each day, to see whether it varied at all. I had taken an interest in it, every day really, I've become almost fixated by it." We observed that the participants used it as a 'sandbox' area [12], by which they could view the consequences of different settings on the cost before approving them. Another well-liked feature was the ability to control the thermostat remotely. Participants commented that this feature affords them a high degree of convenience for heating their home. Most of them reported that they monitored their house (whether the heating was on or off) while they were away, or turned the heating on just before coming home. Furthermore, they found the use of it handy even within the house. For instance, one of our participants commented that she liked being able to take the tablet with her to bed so that she could turn the heating on in cold mornings without having to leave the bed.

Most participants found the thermostat's heating schedule easy to access and program. However, some participants perceived its hourly time slots as limiting their scheduling plan. This is understandable when one considers that today most heating controls provide finer resolutions (e.g., 10 to 30 minutes). Additionally, grouping the daily heating program by weekdays and weekend was not convenient for all participants to accommodate their occupancy patterns. As an example, one participant said her Saturdays and Sundays are totally different. We also had some participants who did not have any occupancy patterns at all and had to adjust the schedule quite a few times in a day. Therefore, further research is needed to address how to best design heating programs for people with unpredictable lifestyles.

Our field study showed that participants could use our thermostats to effectively manage their home heating and create temperature preferences based on real-time prices. As we expected these temperature preferences varied for different individuals. While most participants set lower temperatures at peak prices compared to lower price periods, two households kept the same temperature for all price bands. Furthermore, our participants adopted different strategies to respond to real-time prices. While most participants used the setpoint and the setpoint sliders for reacting to changing prices, some participants interestingly used the boost and the schedule features more than adjusting the setpoint for heating their home with real-time prices. This is in line with a previous study that examined people's use and mental model of their heating system [18], and revealed that setpoint adjustment was less prevalent among their participants compared to the adjustments of other devices, such as the programmer, override button and radiator valves.

Finally, we noted the several ways that our participants used to maintain their thermal comfort, especially when the prices were high, without using our heating systems, such as putting on one more layer of clothing or using a blanket. These observations show similarity to the findings of previous work [7, 6], which examine students' daily heating habits and report the similar activities without any financial benefits.

Expectations from Smart Home Heating Systems

While most participants perceived the thermostat as "smart" because of its learning capability of preferred temperatures and its ability to automate home heating based on changing prices, for some participants it was enough to describe the thermostat as smart just because of its remote control capability and its programmable schedule. This perception was due to the fact that these functionalities were mostly new to the participants. More importantly they experienced improvement in their quality of life as these functionalities assisted and facilitated their heating task. This finding is in line with prior research suggesting that computing technologies would be perceived to be "smart" if they offer an advantage for the users' daily tasks [12].

Regarding the learning feature of the thermostats, participants had different explanations and mental models. While some participants described the system as one that was trying to match their temperature preferences with changing prices, other participants thought that the system was learning the times of days that they set temperatures. Interestingly, participants who used the direct learning thermostat and had no technical background (e.g., P3, an antiques dealer) described more accurate mental models compared to the participants of the indirect learning thermostat with more technical background (e.g., P1-m and P30-m, both computer science PhD students). A previous study examining non-technical users' understandings of an intelligent system suggests that people's initial mental models and misconceptions stayed relatively constant over their study [27]. Therefore, we asked our participants, who had this misconception, if they were aware of any commercial smart thermostats, such as Nest that learns your schedule, in order to see if they had any initial knowledge that would have affected their mental models. However, they all reported that they had not heard of any smart thermostat before. This then may suggest that exposing users directly to the outcomes of learning algorithms may help users to create better mental models. Furthermore, while showing the correlation between previous temperature inputs and prices supported the users' understanding, a more useful method could be a notification system that periodically states what has been learned by the system. We believe these results highlight an important implication for future research in interaction with "smart" energy systems to try and discover the source of people's mental models and learning expectations.

Our system learns users' preferred temperatures at different prices to automate home heating. However, from the interviews, it was clear that the price was not only the factor that our users considered for heating their home. Other key factors were outside weather, occupancy and daily activities within the house. Some participants explicitly stated that the use of the thermostat could be more convenient if it could learn their occupancy patterns. Also, outside weather and the activities that they perform during a day within the house have a significant impact on how people feel the indoor temperature. For instance, most of our participants preferred to have the heating on when they shower and have the heating off when they use their oven or perform physical activities. Therefore, future design of learning thermostats should not only take into account occupancy patterns and outdoor temperatures [16], in addition to people's price preferences, but also people's daily routines (e.g., times that they shower and cook).

Studying Future Smart Energy Systems

In order to let participants experience a future scenario, we prototyped our system based on envisioning [17]. Our scenario depicts an energy market in which consumers can respond to real-time prices by using a smart thermostat that automatically controls heating on their behalf. Participants' statements about their perception and adoption of the smart thermostat indicate that combining experimental reward with a deployed prototype is an effective way to convey a future scenario to participants and allow them to obtain real life experience, echoing recent studies of future scenarios [1, 2, 8]. Extending these recent studies, our participants experienced the autonomous actions of the smart energy system not only through financial incentives, but also through the thermostat's automatic temperature changes. Such changes could directly influence our participants' comfort. Yet, similar to the results of previous studies, our participants mostly felt in control of their heating system and demonstrated a generally positive attitude towards the thermostat. Hence, we believe this finding reinforces those from those previous studies, revealing the potential of future autonomous smart energy systems.

One of the prerequisites for taking part in our study was to have a central heating system with a single boiler. However, we did not define any requirements on the type of thermostat previously installed, such as programmable or non-programmable, digital or analog. Our findings revealed that the type of thermostat familiar to our participants influenced their perception and use of our system. In particular, participants not used to a programmable thermostat focused mostly on the schedule feature of our system, since this was a new and significant feature for them. This circumstance turned out to steer attention away from our primary interest: the ability of our thermostats to automatically react to real-time prices. Hence, future research should take user fragmentation into account in the recruiting process of participants in order to improve the effectiveness of system designs and to obtain more focused results.

CONCLUSION

Smart energy systems that leverage machine learning techniques are increasingly integrated in all aspects of our lives, and they are changing the way that we perform our daily activities. The design of these systems plays a key role in how we adapt to and interact with them. To better understand how to design user interaction with such systems, we implemented three different smart thermostats that automate heating based on users' heating preferences and real-time prices. We evaluated our designs through a field study, where 30 UK households used our thermostats to heat their homes over a month.

Our findings through thematic analysis show that the participants formed different understandings and expectations of our smart thermostat, and used it in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Based on the findings, we provided a number of design and research implications, specifically for designing future smart thermostats that will assist us in controlling home heating with real-time pricing, and for future intelligent autonomous systems. These recommendations will assist designers in improving user experience with smart energy systems, which in return will help us to more smoothly integrate them into our everyday lives and actually benefit from them.

ACKNOWLEDGMENTS

Thanks to our participants for the courage of taking part in our study. Work partially supported by the EPSRC ORCHID project (EP/I011587/1), through a PhD scholarship from the Turkish Government, and by the SNSF (Swiss National Science Foundation) project #147211. Data URL: `http://doi.org/bbpr`. Study approved by University of Southampton FPSE Ethics Committee (ref: 13417).

REFERENCES

1. Alper T. Alan, Enrico Costanza, Joel Fischer, Sarvapali D. Ramchurn, Tom Rodden, and Nicholas R. Jennings. 2014. A Field Study of Human-agent Interaction for Electricity Tariff Switching. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS '14)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 965–972. http://dl.acm.org/citation.cfm?id=2615731.2617400

2. Alper T. Alan, Sarvapali D. Ramchurn, Tom Rodden, Enrico Costanza, Joel Fischer, and Nicholas R. Jennings. 2015. Managing Energy Tariffs with Agents: A Field Study of a Future Smart Energy System at Home. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers (UbiComp/ISWC'15 Adjunct)*. ACM, New York, NY, USA, 1551–1558. DOI: http://dx.doi.org/10.1145/2800835.2801626

3. Galen L. Barbose, Charles A. Goldman, and Bernie Neenan. 2004. A Survey of Utility Experience with Real Time Pricing. (12 2004), 127.

4. Jacky Bourgeois, Janet van der Linden, Gerd Kortuem, Blaine A. Price, and Christopher Rimmer. 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 (UbiComp '14)*. ACM, New York, NY, USA, 459–470. DOI: http://dx.doi.org/10.1145/2632048.2632106

5. Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.

6. Adrian Clear, Adrian Friday, Mike Hazas, and Carolynne Lord. 2014. Catch My Drift?: Achieving Comfort More Sustainably in Conventionally Heated Buildings. In *Proceedings of the 2014 Conference on Designing Interactive Systems (DIS '14)*. ACM, New York, NY, USA, 1015–1024. DOI: http://dx.doi.org/10.1145/2598510.2598529

7. Adrian K. Clear, Janine Morley, Mike Hazas, Adrian Friday, and Oliver Bates. 2013. Understanding Adaptive Thermal Comfort: New Directions for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 113–122. DOI: http://dx.doi.org/10.1145/2493432.2493451

8. Enrico Costanza, Joel E. Fischer, James A. Colley, Tom Rodden, Sarvapali D. Ramchurn, and Nicholas R. Jennings. 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 (CHI '14)*. ACM, New York, NY, USA, 813–822. DOI: http://dx.doi.org/10.1145/2556288.2557167

9. Nicholas R. Jennings, Luc Moreau, David Nicholson, Sarvapali Ramchurn, Stephen Roberts, Tom Rodden, and Alex Rogers. 2014. Human-agent Collectives. *Commun. ACM* 57, 12 (Nov. 2014), 80–88. DOI: http://dx.doi.org/10.1145/2629559

10. Abraham Hang-yat Lam, Yi Yuan, and Dan Wang. 2014. An Occupant-participatory Approach for Thermal Comfort Enhancement and Energy Conservation in Buildings. In *Proceedings of the 5th International Conference on Future Energy Systems (e-Energy '14)*. ACM, New York, NY, USA, 133–143. DOI: http://dx.doi.org/10.1145/2602044.2602067

11. Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. 2010. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys '10)*. ACM, New York, NY, USA, 211–224. DOI: http://dx.doi.org/10.1145/1869983.1870005

12. Sarah Mennicken, Jo Vermeulen, and Elaine M. Huang. 2014. From Today's Augmented Houses to Tomorrow's Smart Homes: New Directions for Home Automation Research. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 105–115. DOI: http://dx.doi.org/10.1145/2632048.2636076

13. Vassilis Moustakis. 1997. Do People in HCI Use Machine Learning?. In *Design of Computing Systems: Social and Ergonomic Considerations, Proceedings of the Seventh International Conference on Human-Computer Interaction, (HCI International '97), San Francisco, California, USA, August 24-29, 1997, Volume 2*. 95–98.

14. F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. 2010. Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions. In *American Control Conference*. Baltimore, USA. http://control.ee.ethz.ch/index.cgi?page=publications; action=details; id=3551

15. Athanasios Aris Panagopoulos, Muddasser Alam, Alex Rogers, and Nicholas R. Jennings. 2015. AdaHeat: A General Adaptive Intelligent Agent for Domestic Heating Control. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1295–1303. http://dl.acm.org/citation.cfm?id=2772879.2773317

16. Therese Peffer, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry. 2011. How people use thermostats in homes: A review. *Building and*

Environment 46, 12 (2011), 2529 – 2541. DOI: http://dx.doi.org/10.1016/j.buildenv.2011.06.002

17. Stuart Reeves. 2012. Envisioning Ubiquitous Computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1573–1582. DOI: http://dx.doi.org/10.1145/2207676.2208278
18. Kirsten M.A. Revell and Neville A. Stanton. 2014. Case studies of mental models in home heat control: Searching for feedback, valve, timer and switch theories. *Applied Ergonomics* 45, 3 (2014), 363 – 378. DOI: http://dx.doi.org/10.1016/j.apergo.2013.05.001
19. Yvonne Rogers, Helen Sharp, and Jenny Preece. 2011. *Interaction design: beyond human-computer interaction*. Wiley.
20. Allison Sauppé and Bilge Mutlu. 2015. The Social Impact of a Robot Co-Worker in Industrial Settings. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3613–3622. DOI: http://dx.doi.org/10.1145/2702123.2702181
21. James Scott, A.J. Bernheim Brush, John Krumm, Brian Meyers, Michael Hazas, Stephen Hodges, and Nicolas Villar. 2011. PreHeat: Controlling Home Heating Using Occupancy Prediction. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp '11)*. ACM, New York, NY, USA, 281–290. DOI: http://dx.doi.org/10.1145/2030112.2030151
22. Mike Shann and Sven Seuken. 2013. An Active Learning Approach to Home Heating in the Smart Grid. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13)*. AAAI Press, 2892–2899. http://dl.acm.org/citation.cfm?id=2540128.2540545
23. Will Simm, Maria Angela Ferrario, Adrian Friday, Peter Newman, Stephen Forshaw, Mike Hazas, and Alan Dix. 2015. Tiree 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 (CHI '15)*. ACM, New York, NY, USA, 1965–1974. DOI: http://dx.doi.org/10.1145/2702123.2702285
24. Robert E. Slavin, John S. Wodarski, and Bernard L. Blackburn. 1981. A Group Contingency for Electricity Conservation in Master-Metered Apartments. *Journal of Applied Behavior Analysis* 14, 3 (1981), 357–97. DOI: http://dx.doi.org/10.1901/jaba.1981.14-357
25. Simone Stumpf, Vidya Rajaram, Lida Li, Weng-Keen Wong, Margaret Burnett, Thomas Dietterich, Erin Sullivan, and Jonathan Herlocker. 2009. Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies* 67, 8 (2009), 639 – 662. DOI: http://dx.doi.org/10.1016/j.ijhcs.2009.03.004
26. Justin Talbot, Bongshin Lee, Ashish Kapoor, and Desney S. Tan. 2009. EnsembleMatrix: Interactive Visualization to Support Machine Learning with Multiple Classifiers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1283–1292. DOI: http://dx.doi.org/10.1145/1518701.1518895
27. Joe Tullio, Anind K. Dey, Jason Chalecki, and James Fogarty. 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 (CHI '07)*. ACM, New York, NY, USA, 31–40. DOI: http://dx.doi.org/10.1145/1240624.1240630
28. Rayoung Yang and Mark W. Newman. 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 (UbiComp '13)*. ACM, New York, NY, USA, 93–102. DOI: http://dx.doi.org/10.1145/2493432.2493489
29. Rayoung Yang, Mark W. Newman, and Jodi Forlizzi. 2014. Making Sustainability Sustainable: Challenges in the Design of Eco-interaction Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 823–832. DOI: http://dx.doi.org/10.1145/2556288.2557380
30. Zhe Yu, Liyan Jia, Mary C. Murphy-Hoye, Annabelle Pratt, and Lang Tong. 2013. Modeling and Stochastic Control for Home Energy Management. *IEEE Trans. Smart Grid* 4, 4 (2013), 2244–2255. DOI: http://dx.doi.org/10.1109/TSG.2013.2279171