

An exploratory study to design constrained engagement in smart heating systems

Simonas Skrebe
Centre for HCI Design, City, University of London
Northampton Square, London EC1V 0HB
Simonas.Skrebe@city.ac.uk

Simone Stumpf
Centre for HCI Design, City, University of London
Northampton Square, London EC1V 0HB
Simone.Stumpf.1@city.ac.uk

Smart heating systems that leverage complex models of user preferences and energy consumption within the home and the wider network in order to make intelligent heating decisions have started to be adopted in homes. While heating systems that allow the user to directly manipulate the heating schedule and temperature have been investigated in some detail, little is known about how to strike a balance between encouraging users to interact with the system but not to demand too much of their attention, what research has termed "constrained engagement" with calm technology. In this exploratory study, we investigated how participants responded to a number of scenarios involving a novel smart heating system in order to support controllability, intelligibility and user experience as part of a constrained engagement approach. We focused in particular on when participants wanted to engage with the smart heating system and how explanations from the system could influence user engagement. Our study contributes a better understanding of users' expectations towards smart heating systems that can form the basis of improved user interfaces.

Constrained engagement, smart heating systems, controllability, intelligibility, demand response.

1. INTRODUCTION

Smart heating systems, such as Nest, Ecobee and Lyric, offer to improve energy efficiency and maintain comfortable temperatures and have become increasingly popular in recent years, often working in conjunction with Internet of Things (IoT) solutions such as IFTTT, HomeKit and SmartThings to automate and connect different devices in the home. Heating systems that previously required users to manually set their schedule to switch the heating on and off to meet target temperatures are now being replaced by smart systems that take actions on the user's behalf to deliver comfort and reduce costs based on user preferences, house occupancy, environment characteristics, etc.

These smart heating systems differ substantially from typical domestic central heating systems (Figure 1) that users are familiar with: usually, traditional systems are built around a source, such as a boiler, which generates heat that is then distributed in some way to heat emitters, such as radiators. The heating can be controlled by the user in various ways: simple on/off boiler controls, thermostats that monitor the room temperature, thermostatic radiator valves (TRVs) on radiators, programmable thermostatic controls that can operate a heating schedule through to controls that can be accessed remotely through the Internet.

Modern smart heating systems instead are able to take actions by themselves to pre-heat the home to meet temperatures in a schedule or switch the heating off when it is sensed that nobody is present, making intelligent decisions to optimise comfort and savings on the user's behalf (Pierce and Paulos 2012). While traditional heating systems are essentially directly manipulated by users, smart heating systems make some heating decisions by themselves; hence, from the user's perspective, the system has temporarily some control over what decisions are executed and how. Of course, the user




Figure 1: Typical domestic heating system

can take back charge and enforce their own decisions at any point.

A new way of taking energy consumption into account that is particularly unfamiliar to users is *demand response*. This means that a smart heating system considers fluctuating consumption in the wider energy network as part of its decisions, for example, the system might take the initiative to reduce or shift the user's energy consumption during peak periods, incentivised through preferable energy tariffs.

Furthermore, new heat sources are also being introduced to homes, partly to make energy consumption more efficient or sustainable, or to leverage smart heating systems better. For example, a *heat pump* operates like a traditional boiler, however, instead of using fuel like gas or electricity to produce heat, a heat pump moves energy from outside (either air or the ground) and "pumps" it to a unit where it can be used to provide central heating. In a way it works similar to a refrigerator, but converts colder temperature into warmer temperature. It works unlike a gas boiler in that it typically runs at a lower temperature, and it takes longer to heat up a home. A *hybrid heat pump* works on the same principle but the crucial difference is that the system can switch between gas and electricity as a primary heat source. The system can, for example, determine to switch to run on gas if the external temperature is low and cannot economically support running as an electric heat pump, or the home needs a short temperature "boost". Therefore, hybrid heat pumps are ideally suited for demand response due to their ability to shift to different energy sources without a reduction in comfort.

Our study is set in a wider project to understand the overall value and user experience of wide scale hybrid heat pump deployment in demand response settings. Our project, FREEDOM (Flexible REsidential Energy Demand Optimisation and Management), led by Passiv Systems Ltd. and funded by Western Power Distribution and Wales and West Utilities, is the first of its kind to investigate consumer engagement with these new technologies as part of a smart heating system.

Previous work involving smart heating systems has called for *constrained engagement* between the system and user (Yang and Newman 2013), where intelligent, automated decision-making by a smart heating system and effective, engaging yet easy-to-understand user interactions are tightly integrated. Research in intelligent systems has shown that control and intelligibility are important aspects to consider in the design of smart systems (Parasuraman and Riley 1997, Lee and See 2004, Shneiderman 2007, Alan et al. 2016), especially to ensure that users understand the reasons behind the system's decisions and actions (Lim et al. 2009,

Kizilcec 2016). Yet little work has been carried out to understand how to best support constrained engagement with these very unfamiliar systems that combine demand response, hybrid heat pumps and intelligent behaviour, so that these systems are better aligned with user expectations and needs.

In order to investigate this, we conducted a study which asked users to engage with five scenarios of a fictitious system which combines all of these novel concepts. Participants were encouraged to sketch a potential user interface that envisioned constrained engagement. We used their verbalisations and discussions as a rich medium to analyse their interactions with a smart heating system. We focused on answering the following research questions:

- How can we support users to feel engaged with their smart heating system so that they feel in control and understand what the system is doing?
- How can we balance engagement with reducing the need for user input?
- What are the potential effects of constrained engagement on the user experience?

Our research contributes a better understanding of users' perceptions and expectations towards smart heating systems that can form the basis of improved smart heating user interfaces.

In this paper, we first outline previous research in smart systems that make intelligent decisions on the user's behalf. We then describe the design of our study. We present our findings related to the two research questions in turn. We follow this with a discussion of our findings and potential avenues for future work before concluding with a summary of our study.

2. RELATED WORK

Designing usable smart heating systems needs to take into account the extended, mixed-initiative interaction between a user and an intelligent system that make its own decisions (Horvitz 1999). Previous research has investigated these systems mainly in terms of user input, i.e. *controllability*, system outputs, i.e. *intelligibility*, and the effects this might have on the *user experience*. We will now provide an overview of related work for each of these aspects.

2.1 Controllability

A smart system takes charge of carrying out a user's tasks (Parasuraman et al. 2000), therefore, from the user's perspective, control over actions is, at least

temporarily, shifted from the user to the system in a *mixed-initiative* approach (Horvitz 1999). Previous research with smart heating systems has indicated that users find it difficult to manage the heating behaviour, especially in adapting the system to user preferences (Yang and Newman 2013). Much of the effectiveness of smart systems rests on them easing the cognitive load and optimising tasks for users (Shneiderman 2007) and there should be less need for the user to directly manipulate the system. This also fits with the notion of calm technology (Weiser and Brown 1997), in which the technology fades into the periphery. This balance between controllability and calm technology has been termed *constrained engagement* (Yang and Newman 2013), where user input is encouraged but also needs to be constrained such that it does not overwhelm the user. In smart heating systems, this means that we would like to make the most of the interaction channel between user and system, so that we keep the user informed of what decisions the system is making and we get input from the user to make the system more effective, yet ensure that the user does not need to give too much input. While research in smart systems outside of the heating domain indicates that user input is especially focused on points in the interaction when the system's behaviour runs counter to the user's expectations, or it makes wrong decisions in the eyes of the user (Alan et al. 2014, Hoff and Bashir 2015, Parasuraman and Riley 1997, Kizilcec 2016), we do not yet know how constrained engagement can be achieved in smart heating systems. We investigated this gap in our understanding involving these new smart heating systems which are very unfamiliar to users in our study through fictitious scenarios when the heating is not behaving as expected or violates user preferences.

2.2 Intelligibility

It has been shown that users need to understand how an intelligent system works in order to use it effectively, promote positive attitudes towards the system, increase their trust in the system's decisions, and to gain a sense of controllability (Gregor and Benbasat 1999, Lim et al. 2009, Cramer 2012; Kulesza et al. 2012, Holliday et al. 2016, Kizilcec 2016). Yang and Newman's (2013) study of the Nest thermostat found that despite overall positive perceptions of the thermostat, users did not find the intelligent features of the thermostat useful as they were unable to interpret them. Similarly, it was found that users often lacked sufficient know-how to make the best use of smart heating technology (Rodden et al. 2013).

It has been suggested that explanations about a smart system behaviour can increase intelligibility (Gregor and Benbasat 1999, Lim et al. 2009, Kulesza et al. 2010), for example, by describing why or why not a smart system behaved in a certain way,

what will happen next, the data the smart system uses to make decisions, etc. It has also been shown that explanations can increase the correctness of users' mental models (Kulesza et al. 2012) and might therefore address common misconceptions about how heating works. For example, people often think that "whacking the temperature up" on thermostats will speed up the rate at which the home temperature increases (Yang and Newman 2013, Fischer et al. 2016). However, current smart heating systems do not explain their intelligent decisions to a user, and thus do not encourage users to partake in constrained engagement. Our study investigates in detail how to communicate with users about the actions that a smart heating system is taking.

2.3 User Experience

Increased intelligibility positively affects the adoption and use of technology in general (Kizilcec 2016) and also smart heating systems specifically (Costanza et al. 2014). Inability to understand why smart systems perform in a certain way, especially when users are expecting different system actions, leads to frustration, mistrust, and can even deter people from using the technology altogether (Corritore et al. 2003, Parasuraman and Riley 1997). Our study investigates how perceived controllability and the intelligibility of the smart heating system can shape user experience.

3. STUDY DESIGN

We conducted an empirical study in which we asked participants to engage with five scenarios involving a fictitious smart heating system interface that made use of a hybrid heat pump in a demand response setting, and was able to provide some explanation of the actions it is taking. This kind of smart heating system is currently not in real-life use, hence we did not compare our approach with existing smart heating systems such as Nest, Ecobee, etc.

3.1 Participants

Using a convenience and snowball sampling approach, participants were recruited via email, social media, and posters displayed at City, University of London. We decided to involve mainly participants who did not already have experience with a smart heating systems because this is a representative sample of the wider population, given that smart heating controls are still relatively rare. In addition, this sample mirrors the target audience for the FREEDOM project.

To be eligible for the study, all participants had to at least contribute a share towards household energy bills and be over 18 years old. Fourteen participants were recruited for the study, equally split between 7 males and 7 females with a mean age of participants

of 33 (minimum 23, maximum 50) and a variety of occupations and income levels. Each participant received a £10 Amazon voucher for taking part.

All participants reported owning a smartphone device, 13 of 14 participants owned a laptop and 10 of 14 participants owned a tablet. Only one participant (P14) said they were somewhat confident at using technology; the remainder reported they were confident or very confident at using technology.

Eight participants owned a property and 6 were tenants. The most popular heating source was a gas boiler (11 participants), followed by community / district heating (2 participants) or a heat pump (1 participant). Only one participant used a smart heating control already; other participants controlled the heating by using the boiler controls, turning the radiators on or off manually through TRVs, or using an external thermostat. Nine participants (P2, P3, P4, P6, P8, P9, P11, P13, P14) ran some kind of schedule to heat their home at specified times, and then manually manipulated the heating based on whether they felt too cold or warm. The remainder of participants did not run a schedule, turning the heating on or off manually.

3.2 Procedure

Each session took approximately 1 hour and 15 minutes. At the beginning of the study, participants completed a background questionnaire asking about their demographics (age, gender, occupation, income), technology use (owned devices, confidence in using technology), property (rent or own, time lived in the property, number of occupants), and their energy provider and interaction with the system (company, type of energy fuel used, annual heating bill, type of heating system, interaction with the heating system). Then, we conducted brief semi-structured interviews to learn more about participants' heating system setup and understand how they used it.

The main part of the session started with a researcher giving an overview of hybrid heat pump systems and demand response and providing an example schedule to use. We then presented each scenario to participants in turn, and asked them how they would behave in each situation. The researcher and the participant drew a sketch on the fly (Figure 2), illustrating the user interface alongside explanations that described system actions. Throughout the study, participants were encouraged to draw on the sketches and verbalise their thoughts and decisions. The discussions around the sketches thus provided a rich medium to capture participant's perceptions and expectations (Tohidi et al. 2006).

After completing the main part, we conducted a semi-structured interview to further understand their choice of heating technology, what they regarded as comfortable temperatures in the home, the usability

of current heating systems, where they turn to for more information to get support for using their heating system, and their attitude to any financial rewards for using a smart heating system.




Figure 2: Example of sketch of the user interface, together with explanation of system behaviour.

3.3 The scenarios

The participants were first given a hypothetical schedule to ground their decision-making. The schedule was split into weekdays and weekends, each showing different occupancy periods when people were at home, out or asleep and the desired temperature levels. It represented a common occupancy routine pattern with an assumption that it would suit the majority of the study participants' lifestyles.

Five scenarios were created in consultation with a heating expert within the project team to explore participants' reactions and ensuing interactions with the system. These scenarios varied in how much the system behaviour deviated from common but possibly wrong mental models of how heating works, in the range of unexpected actions of the heating system, and in the content of the explanations that were provided for these actions.

Scenario 1 (over-ride schedule) explored the situation when the user wants to possibly override the current temperature as set in their schedule. The hypothetical situation was that they had returned home over the weekend feeling a bit cold from walk on a rainy day, with the current schedule set to Out. If they increased the temperature, we added a system explanation that the system is heating up, and that they could hear the heat pump humming but it did not get warm very quickly, corresponding to the behaviour of heat pumps.

Scenario 2 (system pre-heating ahead of schedule) explored the system taking the decision to heat in order to meet a scheduled occupancy period in the

morning. The hypothetical situation presented to the participants was that they woke up at 5am on a weekday morning in a warm room, with the heat pump humming and the radiators warm. The system explanation was that the system was pre-heating for the next schedule period and that it usually took 2 hours to heat up the home to the desired temperature.

Scenario 3 (system demand-response delay) investigated the system shifting the heating based on network demand, having delayed the heating but now heating up to the desired temperature. Note that in reality a hybrid heat pump would typically switch to a different, more economical fuel type instead of delaying the heating but we wanted to explore a case where the system's decision had some potentially significant impact on the user. We explained that the system was heating to meet the scheduled temperature but was delayed due to high energy demand, but their temporary discomfort was off-set by a heating bill rebate.

Scenario 4 (system demand-response switch-off) investigated the system switching off the heating due to network demand. This time we introduced an element in the explanation which discouraged repeatedly overriding the system's decision by losing their bill rebate (e.g. "If you increase the temperature 3 times a month during the demand response periods, you will lose your £5 rebate. You have already increased it: 1 time"). This case explored an artificially severe case of the system violating user expectations but we wanted to see how this differed to scenario 3.

Scenario 5 (system pre-heating ahead of demand-response) investigated the system pre-heating the home to avoid a peak network demand, slightly ahead of the set schedule.

3.4 Data Analysis

The qualitative data from the main session formed the heart of the analysis, and audio recordings of all participants' responses from contextual interviews were transcribed and analysed using MaxQDA v.12 qualitative analysis software.

We used a thematic analysis approach to analyse the transcripts (Braun and Clarke 2006) and we combined participant statements which provided similar feedback. For the analysis, we identified statements by participants within each of the five scenarios that related to key themes inspired by the literature review – control, comfort, cost, intelligibility, trust – and desired features (Table 1). To ensure the reliability of the data analysis, the researchers agreed on the key themes prior to the study. In addition, a second researcher reviewed the research findings in order to validate the qualitative data analysis.

We then further looked at the responses in relation to these themes and grouped them into particular topics that are discussed in the results part of this paper: controllability, intelligibility and user experience.

To support the data analysis, photographs of paper sketches were taken to provide contextual reference to specific scenario and participants' responses. In particular, we were looking for sketched elements that mapped back to the six themes we used in the analysis.

Theme	Statement relating to participants'...
Control	Perceived control or the lack of it when using a system.
Comfort	Perceived comfort when a heating system makes an independent decision influencing user's comfort levels.
Cost	Perceived impact of reduced or increased costs when using a smart heating system.
Intelligibility	Perceptions related to text and graphical explanations as well as understanding system features and decision making.
Trust	Participant's perceived trust towards a smart heating system.
UI features	Requirements for specific user interface features that they would like to see or interact with when using a smart heating system.

Table 1: Themes and their description used in the data analysis.

The background questionnaire data and the semi-structured interview data provided additional information as to any possible reasons for participants' statements or decision-making. The background questionnaire was analysed using Microsoft Excel to provide an overview picture about the participants, their use and interaction with technology, current heating system. Data from the semi-structured interviews was used as a reference of participants' overall view towards traditional and smart heating systems related to automation, control and demand response. Since the key part of the research was formed by the qualitative data and due to the low sample size, we did not break down the data from the main session by the questionnaire or interview responses.

4. RESULTS

4.1 Controllability

At the moment there is conflicting evidence about how users trade-off between comfort and cost when controlling smart heating systems. Previous research has indicated that most users are very cost-sensitive (Costanza et al. 2014), however, studies that motivated users to reduce energy use found that users prioritise convenience and comfort (Yang and Newman 2012). When asked explicitly during the study what was more important to them, 13 participants prioritised comfort over cost. However, 8 out of the 14 participants (P1, P2, P3, P6, P7, P8, P9, P11) were still mindful about energy costs and stated that they would not heat their homes excessively. This points to the importance of allowing the user a choice and the user should be engaged in making these decisions, for example by being able to provide input to the system how comfort and cost should be balanced. However, constrained engagement means that this should be a general setting, either learned or manually provided, instead of forcing the user to make a decision instance by instance.

Controllability was more important to our participants when the system took actions that had bigger perceived effects on participants. For example, in scenario 1 (over-ride schedule), eight participants stated they would slightly increase the temperature from the current set point (P3, P4, P5, P6, P7, P8, P9, P12). Two participants expressed that they would increase the temperature to as much as possible (P10, P11), indicating that they held an incorrect mental model that setting the thermostat very high makes it warmer quicker. However, small temperature differences or assuring them that the system was warming up meant that they did not perceive a need to interact with the system. For example, we found that participants were less concerned about possibly small fluctuations in the home temperature which they did not perceive to substantially affect their comfort. Seven participants did not care about small temperature differences in scenario 1 (P1, P9, P13) and in scenario 3 (P5, P8, P11, P13) and therefore did not take any corrective actions. Similarly, in scenario 4, four participants (P1, P2, P7, P9) stated that they were happy to devolve decisions to the system if the difference between current and scheduled temperature was very small. However, three participants (P5, P7, P12) reported feeling anxious about comfort in the long term, especially if their expectations might be significantly violated.

It was observed that participants tended to want to take charge of the smart system's actions when they thought the system was not meeting their preferences, or when the system's actions were unexpected. These instances provide starting points

for constrained engagement between the system and the user, and most of them can be easily enumerated. For example, unexpected actions by the system are that the temperature is below the set schedule or the temperature is above the set schedule. A user's mental model, from experience with traditional heating systems, is often that their schedule is directly linked with the heating switching on or off, hence, during delaying heating when the temperature was below their scheduled temperature, participants often sought to interact with the system to fix the perceived problem. In one instance in scenario 3, the system actions were even perceived as a fault:

"I would call the service, try to find the troubleshooting." (P2)

In scenario 4 (demand response switch-off), seven participants (P4, P5, P6, P7, P8, P13, P14) wanted to override the system by increasing the temperature, changing the schedule settings or even switching the smart system off completely, returning to manual control. However, in scenario 3 (demand response delay), with the promise of an incentive, six participants did not feel the need to interact (P1, P7, P12, P13, P14) even though the system actions were unexpected.

A particular problematic issue for participants was pre-heating (the temperature is above the current scheduled temperature) which was perceived to be outside of their control, with the fear that this might result in increased energy bills without being explicitly approved by the user. In fact, pre-heating the home is optimised to achieve comfort as outlined by the schedule. Participants attempted some workarounds to take charge of the pre-heating. Nine participants (P2, P3, P5, P6, P7, P10, P12, P13, P14) attempted to reduce the target temperature overall to stop the pre-heating so that the bedroom was cooler. Four participants (P7, P8, P9, P13) attempted to manage the pre-heating through adapting their schedule. For example, one participant (P8) wanted to reduce the overall temperature hoping the system would start heating later while two participants (P1, P7) wanted to force the system to heed their preferences by amending the scheduled occupancy periods to start the pre-heating times later:

"It's annoying as I will have to change all of the set points for the preheat time, rather than naturally setting the schedule i.e. 7-9, now I will have to set it for 8am to accommodate the need of the system. Maybe then it will be the temperature I want to. To take into account how the system works all of these things would annoy me, I want to set the schedule and for the system to work to make my schedule work." (P7)

Only one participant (P8) said that they would also attempt to control the temperature via TRVs to

ensure that temperatures in the bedroom would be kept low while the pre-heating of other rooms can continue.

In contrast, in scenario 5 (system pre-heating ahead of demand-response), once the participants appreciated the benefits of pre-heating in order to shift energy consumption, all participants were happy with the system actions. This indicates that system actions that are considered unusual are prime points during which user engagement takes place but that, given appropriate explanations or incentives, users might accept the unexpected system actions in the longer term.

Although users are able to step in at any point to change the temperature and they are fully in control of the schedule to optimise comfort, there are various ways that they could feel even greater confidence in the smart heating system. If they disagreed with the system's actions, participants mainly wanted the option to temporarily stop the heating, be able to delay heating, or set the heating times more directly.

Frequently throughout the study, participants voiced apprehension that system actions were happening without them being aware of them. For example, in scenario 3 where the system delayed switching on, six participants (P3, P4, P5, P6, P9, P10) were disappointed that the system took an action without having been notified at the time the decision was taken. In the absence of this notification during which point they could give explicit consent to the system's actions, five participants (P4, P5, P6, P9, P10) were interested in ways to opt out from the smart heating system in some way. This indicates that users would like to input some form of consent or notification of future system behaviour but we saw no evidence of users wanting to explore "what-if" situations in system behaviour in the scenarios we explored.

4.2 Intelligibility

In addition, participants requested much more information about why the system made these decisions, a point we will now cover as part of intelligibility.

We have already noted that participants wanted to be informed about system actions in advance of them taking effect, and they then wanted to consider what to do next, for example, have an option to override the action if they are not happy with what the system is doing or go along with it. In all scenarios, participants indicated that they would try to understand why the system was behaving in the way it did when they noticed unexpected behaviour or when they wanted the system to take some action.

First, participants wanted to be kept informed about the current system's state and actions (Figure 3 and

4). For example, participants were interested in learning about what the current schedule period is (P5, P7, P11), how long heating was delayed, projected time the target temperature will be reached, even to the point of which radiators were currently heating up, suggesting a need for detailed explanations about what was happening with the system at the time:

"Target temperature will be reached at 3:00pm." (P1, P4, P13)

"Target temperature will be reached in 3 hours." (P1, P2, P7, P8, P14)

"Heating will be delayed by 2 hours." (P1, P12)




Figure 3: P3's sketch for explaining scenario 1.

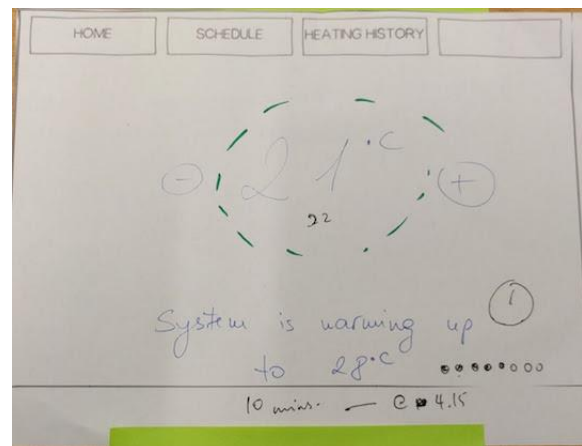


Figure 4: P7's sketch for progress indication (black dots) and an explanation how long it will take to reach a certain temperature.

As well as simply describing the current system actions, participants were also interested in the reasons for why the system is doing the current action, for example:

"Ok (...) so it was trying to make me pay less - saving me money, it was expecting me to come

at 6 o'clock and then allowed it to cool down.
Sounds like the machine is working well." (P5)

Participants also highlighted a need to receive more information about the system's future and past behaviour. This particularly occurred for actions that were perceived as occurring on a regular basis such as demand response or pre-heating. In these cases, participants wanted to know how often demand response occurs, projected times it will occur and who else is affected by demand response (P4, P10, P11), pre-heating start times and how long the temperature would be maintained for to ensure comfort levels (P3, P6, P9, P11).

Graphical explanations were found to be useful by a lot of participants (P1, P3, P4, P5, P6, P7, P8) to communicate the smart system's actions. These graphical explanations could include detailed information, such as scheduled temperatures, current temperature, and historic and projected future energy prices and costs, including peak demand, and energy type used (Figure 5). For example, having looked at the graph in scenario 5 (Figure 6), P5 thought the system's pre-heating shown in relation to price fluctuations demonstrated that the system worked well. This suggests that graphs might work well to explain complex information in order for users to quickly evaluate system's behaviour.

However, not all participants understood the explanations, especially if the actions ran counter to the participants' previous experiences with traditional heating systems. This occurred repeatedly across scenario 2, 3 and 4. Most participants were not familiar with pre-heating, demand response and hybrid heat pumps, and found these concepts hard to grasp from explanations in the UI alone. Participants stated that more needed to be done to educate system users about the way that smart heating systems worked more generally in order to interpret system actions, and this might need to be communicated in other forms, such as contracts, leaflets, brochures, etc.

Particularly with respect to system actions that might have significant effects on their individual comfort, participants needed more intelligibility:

"Would be very annoyed! It has delayed the heating, taken the control away and then offered £5." (P3)

The explanations for any system actions should stress the beneficial aspects of the system's actions, such as energy savings and impacts on other customers in the wider energy networks, for example:

"High demand. We will save you £5 if we delay your heating by 15 minutes" (P).




Figure 5: P1's sketch showing schedule indicating different fuel types.

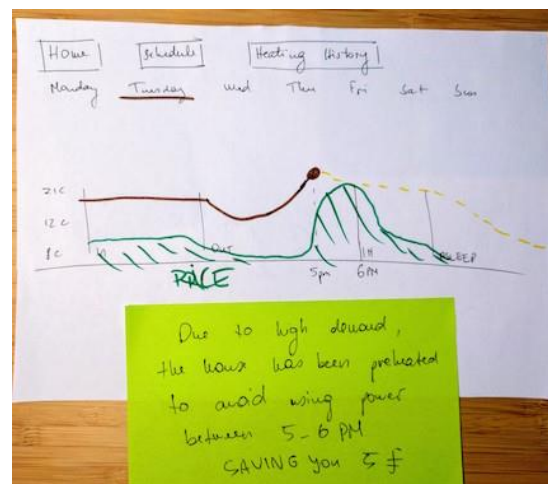


Figure 6: Sketch showing price fluctuation during the day (P5)

In these cases, participants wanted more information about how the system is going to achieve the required comfort levels, for example, how long it will take to heat up to target temperature and whether the system will be able to maintain the temperature without causing discomfort.

To make decisions and interact with the system effectively, these participants wanted detailed information about the link between comfort, energy use and system actions:

"Oh ok... now that it is good. It has recognised what it needs to do and it preheated the house to avoid us having that cold, that's really good. Again with things like that I would like to know if we can plan for". (P9)

4.3 User Experience

If explanations hit it right, they can make users feel more in charge, and improve the user experience of the system, even if its actions are unexpected:

"Everything can be forgiven if you have received an explanation." (P11)

"It gives you options." (P7)

Consequently, after having read the explanations, trust placed in the system might make participants' interactions with the smart heating system unnecessary:

"I feel positive that it is going to heat up. And the system is going to do what I decided what I want to do" (P14).

However, in other instances when there was not an explanation that was perceived to adequately justify the system's actions, or the system was felt to be taking too much control away from the user, study participants (P1, P3, P13, P14) were upset, annoyed, or even felt threatened and punished (P13). Participants stated that they were not being told a "full story" (P11) and as a result, some participants (P1, P3, P11, P12) would decide to either not interact with the system until they investigated further what to do next, or switch the energy provider. Often participants blamed energy companies for smart heating system actions, such as pre-heating homes or affecting comfort, similar to findings by Rodden et al. (2013). This indicates that eroding trust can have significant consequences, such as smart heating system disuse (Parasuraman and Riley 1997) but also transfer to distrusting energy companies in general.

5. DISCUSSION AND FUTURE WORK

Our study showed that participants thought that smart heating systems were an exciting idea and that they would be interested in trying it. However, participants also expressed a need to understand how the system works and why it takes some actions, so they can trust the decisions it makes, otherwise they might stop using the smart system.

We have investigated opportunities for the design of constrained engagement. Our results indicate that users might take unexpected behaviour by the smart heating system that significantly impacts their comfort as a point that might prompt further interactions. However, if given convincing explanations, they might not feel the need to intervene. How to carefully balance users feeling in control while also feeling engaged with the system is an open research question.

Our study underlined the importance of accurate mental models of the system for users to understand its actions and use it to its best ability (Kulesza et al. 2012). Smart heating systems use complex computer models to optimise user comfort. Participants' mental model were based around traditional combi-boiler central heating systems, and they frequently struggled to understand the decisions and actions the smart system was making. Therefore, it is important to explain novel heating

concepts but research into how to best achieve this is still warranted. It is questionable that all explanations to that level of depth could be achieved through the user interface of a heating control app but possibly other, more expansive forms of communications might be suitable to correct users' mental models. However, this might mean that these explanations might also overwhelm users, contrary to principles of constrained engagement (Yang and Newman 2013) and explanatory debugging (Kulesza et al. 2015).

We investigated in detail what might make convincing explanations of system actions. Textual explanations seemed to be enough to justify current actions but while participants consistently agreed that the explanations were useful they nonetheless should contain detailed information, for example, how long it will take to heat up the home, etc. Crucially, these explanations need to communicate the benefit to the user – why the system is performing in a certain way – rather than just describing system actions – what the system is doing.

We are heartened that complex graphical visualisations seem to be viable as a way to communicate heating decisions, both for past and future actions, and that they gave participants confidence about what the system was doing. Graphs seemed to be useful for two main reasons: first, participants were able to see the home heating patterns which explained how the system achieved the schedule, and second, how heating actions might correlate with other patterns in the wider network. However, further research is warranted about how textual and graphical explanations can work hand in hand, and how to craft the explanations so as not to overwhelm users.

Our work also points the way to better investigate these kinds of systems. Our approach encouraged participants to imagine concrete scenarios of potential future use, carefully constructed by the project team to explore dimensions such as control, intelligibility and trust. Participants then co-designed possible user interfaces, allowing us to gain an insight into their current mental models as well as their requirements for smart heating system design.

Our results clearly show that participants' interactions to manage a smart heating system coupled with increased intelligibility can positively impact trust. However, our study is not without its limitations. Although our approach aimed to create a realistic picture of how these smart heating systems might work, users typically depend on other types of feedback in addition to control user interfaces, such as heat emitting from radiators, the hum of a boiler, etc. For example, from our study it appears that small temperature differences might be tolerated well, however, we were not able to emulate changes in actual heat; what might look like a small difference

on paper might be felt as a much bigger temperature in real life.

Currently, our project is engaged in designing and implementing user interfaces for a smart heating system that explains its behaviour to users. In future work, we will be evaluating and comparing the effectiveness of textual and graphical explanations and its impact on trust and constrained engagement in a lab study. We are very excited about participating in a field study over the heating season in 2017/18 during which we will have the opportunity to gather feedback from users in a real world setting after the complete smart heating system has been installed.

6. CONCLUSIONS

We conducted a study with fourteen participants in which they were set fictitious scenarios and asked to help annotate a user interface sketch to investigate how users interact with smart heating systems. We found that

- Constrained engagement can be managed. Participants interacted with the smart heating system when their comfort was compromised or when the system actions were unexpected. Users expect to be notified when decisions are taken, however, they can be dissuaded from intervening to over-ride system behaviour through appropriate explanations;
- Designers of smart heating systems will need to take account of users trying to force the system to conform to their wrong mental models. Explanations can help but will need to go much deeper to explain important novel concepts in domestic heating;
- Users want detailed information about what the system is doing as well as the motivation for doing them. Graphical visualisations seem to be able to communicate actions past and future;
- Interactions with smart heating systems are important to user experience. If not supported appropriately this might lead to them abandoning the system as well as distrusting energy companies.

We have outlined future work to craft better explanations to increase users' sense of control, intelligibility, and user experience. Our work is the first step in a better understanding of users' expectations towards smart heating systems that can form the basis of improved user interfaces, with smart heating systems working hand-in-hand with users to optimise comfort and reduce costs.

7. ACKNOWLEDGEMENTS

We thank the participants of our study. This work was supported by the FREEDOM project, funded by Western Power Distribution and Wales and West Utilities. We also thank Edwin Carter from Passiv Systems Ltd. for his help with this research.

8. REFERENCES

- Alan, A.T., Shann, M., Costanza, E., Ramchurn, S. and Seuken, S. (2016) It is too hot: an in-situ study of three designs for heating. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
- Alan, A., Costanza, E., Fischer, J., Ramchurn, S.D., Rodden, T. and Jennings, N.R. (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*. International Foundation for Autonomous Agents and Multi Agent Systems. 965-972
- Braun, V. and Clarke, V. (2006) Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2). 77-101.
- Cramer, H. (2012) Personalized and autonomous are already everywhere; let's focus on awareness over trust. *Workshop on End-User Interactions with Intelligent and Autonomous Systems*, SIGCHI Conference on Human Factors in Computing Systems.
- Costanza, E., Fischer, J.E., Colley, J.A., Rodden, T., Ramchurn, S.D. and Jennings, N.R. (2014) Doing the laundry with agents: a field trial of a future smart energy system in the home. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 813-822.
- Corritore, C.L., Kracher, B. and Wiedenbeck, S. (2003) On-line trust: concepts, evolving themes, a model. *International Journal of Human-Computer Studies*, 58(6), 737-758.
- Fischer, J.E., Crabtree, Rodden, T., Colley, J.A., Costanza, E., Jewell, M.O. and Sarvapali D. Ramchurn. (2016) "Just whack it on until it gets hot": Working with IoT Data in the Home. In: *Proceedings of Conference on Human Factors in Computing Systems*. ACM, 5933-5944
- Gregor, S. and Benbasat, I. (1999) Explanations from intelligent systems: Theoretical foundations and implications for practice. *MIS Quarterly*. 497-530.
- Hoff, K.A. and Bashir, M. (2015) Trust in automation integrating empirical evidence on factors that influence trust. *Human Factors: The Journal of*

- the *Human Factors and Ergonomics Society*, 57(3). 407-434.
- Holliday, D., Wilson, S. and Stumpf, S. (2016) User Trust in Intelligent Systems: A Journey Over Time. In: *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 164-168.
- Horvitz, E. (1999) Principles of mixed-initiative user interfaces. In: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 159-166.
- Kizilcec, R.F. How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface. (2016) In: *Proceedings of the Conference on Human Factors in Computing Systems*. ACM, 2390-2395.
- Kulesza, T., Stumpf, S., Burnett, M., Wong, W.K., Riche, Y., Moore, T., Oberst, I., Shinsel, A. and McIntosh, K. (2010) Explanatory debugging: Supporting end-user debugging of machine-learned programs. In: *Proceedings of IEEE Symposium on Visual Languages and Human-Centric Computing*. IEEE, 41-48.
- Kulesza, T., Stumpf, S., Burnett, M. and Kwan, I. (2012) Tell me more?: The effects of mental model soundness on personalizing an intelligent agent. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1-10.
- Kulesza, T., Burnett, M., Wong, W.-K. and Stumpf, S. (2015) Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In: *Proceedings of the 20th International Conference on Intelligent User Interfaces*. ACM.
- Lee, J.D. and See, K.A. (2004) Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 46(1). 50-80.
- Lim, B.Y., Dey, A.K. and Avrahami, D. (2009) Why and why not explanations improve the intelligibility of context-aware intelligent systems. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2119-2128.
- Parasuraman, R. and Riley, V. (1997) Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 39(2). 230-253.
- Parasuraman, R., Sheridan, T.B. and Wickens, C.D. (2000) A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, 30(3). 286-297.
- Pierce, J. and Paulos, E. (2012) Beyond energy monitors: interaction, energy, and emerging energy systems. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 665-674.
- Rodden, T.A., Fischer, J.E., Pantidi, N., Bachour, K. and Moran, S. (2013) At home with agents: exploring attitudes towards future smart energy infrastructures. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 1173-1182.
- Shneiderman, B. (2007) Human responsibility for autonomous agents. *Intelligent Systems*, 22(2). IEEE, 60-61.
- Tohidi, M., Buxton, W., Baecker, R. and Sellen, A. (2006) User sketches: a quick, inexpensive, and effective way to elicit more reflective user feedback. In: *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*, ACM, 105-114.
- Weiser, Mark, and Brown, J.S. (1997) The coming age of calm technology. *Beyond calculation*. New York: Springer, 75-85.
- Yang, R. and Newman, M.W. (2012) Living with an intelligent thermostat: advanced control for heating and cooling systems. In: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 1102-1107.
- Yang, R. and Newman, M.W. (2013) Learning from a learning thermostat: lessons for intelligent systems for the home. In: *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 93-102.