

Log it While It's Hot: Designing Human Interaction with Smart Thermostats for Shared Work Environments

Stephen Snow
University of Southampton
Southampton, UK
s.snow@soton.ac.uk

Frederik Auffenberg
University of Southampton
Southampton, UK
fal12@ecs.soton.ac.uk

m.c. schraefel
University of Southampton
Southampton, UK
mc@ecs.soton.ac.uk

ABSTRACT

Smart thermostats offer impressive scope for adapting to users' thermal comfort preferences and saving energy in shared work environments. Yet human interactions with smart thermostats thus far amount to an assumption from designers that users are willing and able to provide unbiased data at regular intervals; which may be unrealistic. In this paper we highlight the variety of social factors which complicate users' relationships with smart thermostats in shared work environments. These include social dynamics, expectations, and contextually specific factors that influence motivations for interaction with the system. In response we outline our framework towards a *Smarter Thermostat*: one which better accounts for these messy social inevitabilities, is equipped for a decline in user feedback over time and one which augments rather than attempts to replace human intelligence- thereby ensuring a smarter thermostat does not create dumber humans.

Author Keywords

Smart thermostat; thermal comfort; shared work environments; office; reciprocity; participatory sensing

ACM Classification Keywords

H.5.2 User Interfaces: User-centered design

INTRODUCTION

Ongoing issues of climate change and peak demand provide a strong case for reducing building energy use. Given that Heating, Ventilation and Air Conditioning (HVAC) accounts for over 40% of a commercial building's energy consumption [8], improving the efficiency of HVAC systems through smart technology, offers advantages in terms of sustainability and running costs.

Smart thermostats which learn and adapt to users' thermal preferences, offer a promising solution to balancing a reduction of HVAC energy use while retaining thermal

satisfaction for users [1, 3, 10, 18, 21]. In order to achieve this, machine learning algorithms for smart thermostats tend to be grounded in a generalized model of thermal comfort (e.g. [7] or [11]), with an added capacity to learn peoples' preferences for certain parameters through user-input into the system. However, assumptions made about how humans can or should interact with such systems can limit their suitability in practice. Further, relatively less smart thermostat-related research concentrates on field deployments in shared work environments (e.g. offices) as opposed to in the home.

In shared work environments, participatory sensing mechanisms designed to assist smart thermostats in calculating peoples' thermal comfort thus far amount to an assumption that users are willing and able to engage with the system at regular intervals over extended time periods [10, 18, 23, 25, 27]. These assumptions contradict social research which finds everyday life to be messy and unpredictable [31] and that engagement with personal informatics tends to wane over time [15, 29, 33]. More fundamentally, if the notion of "smart" represents a system learning users' preferences and progressively absolving them from responsibility for regulating their thermal comfort, do we run a risk of smart thermostats creating humans who are less connected to the outdoor environment and less adaptable to thermal fluctuations?

In this paper, we describe the variety of social and interactional issues we encountered during our deployment of a novel participatory sensing system which allowed users to log their thermal comfort. Using a human-centred approach to understanding thermal comfort in work environments, our findings highlight a number of social and context-specific factors which can affect the use of such a system. We found that interaction with the system was motivated by user discomfort, was affected by pre-existing relationships with building managers and was characterized a decline in usage over time.

Based on these findings, we highlight the limits of relying on users as reliable and autonomous sensors. Participatory sensing systems for smart thermostats require careful consideration in design in order to account for the human factors which such systems will inevitably encounter in real-world deployments. Finally, we consider how these systems may reciprocate the input provided by their users. These factors shape our vision for a *Smarter Thermostat*,

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

CHI 2017, May 06-11, 2017, Denver, CO, USA
© 2017 ACM. ISBN 978-1-4503-4655-9/17/05...\$15.00
DOI: <http://dx.doi.org/10.1145/3025453.3025578>

one which: (1) Leverages the messy social inevitabilities that we highlight in our deployment and is equipped to handle changes in user engagement, and; (2) Augments, human intelligence through reciprocity and actionable information.

While these considerations for design have been developed in the specific context of smart thermostats in shared work environments, we expect they may potentially be generalizable to designers of smart systems more broadly.

Background: Refining smart thermostat operation through human input

A smart thermostat is a computer-based agent to control temperature via HVAC systems in the home or office. Two typical aims of smart thermostats are to: (1) reduce HVAC energy consumption by simplifying interactions between the human and the thermostat [3, 22] and (2) for the thermostat to provide agreeable thermal comfort to occupants by learning their heating schedules [32] or thermal comfort preferences [10, 18]. Machine learning approaches to thermal comfort modelling for smart thermostats involve extracting user-specific information from existing thermal comfort models, (e.g. the Static [11] or Adaptive [7] models) and learning these parameters through user input. User feedback is used to train and refine the thermal comfort model in order for the system to adjust for individual users' preferences. As such, designing for human involvement in smart thermostat systems is equally fundamental to the accuracy of the model to which it is grounded.

The necessity of human input into smart thermostat operations make it a problem well suited to HCI. Existing HCI work on thermal comfort extends to wearable sensors for personal thermostat control [12] and for sensing thermal comfort [5, 16]. Compact wrist-wearable pulse and skin temperature sensors can provide accurate real-time physiological thermal comfort information from multiple users [5, 16]. Such systems provide accurate per-individual information on thermal comfort at the individual level, however, given the transient occupation and higher numbers of users in shared work environments, their practicality for use here may be limited.

Literature on smart thermostat deployments in the home suggests that the value of human participation in smart thermostat systems has been compromised in the past by factors related to usability and the machine's inability to sense the intent of user input [32]. Smart thermostat deployments in the home thus far have had limited success in leveraging human input into their operations. In contrast to their common goal of energy saving, some have in fact increased energy use in the past [9]. Alan et al. [1] found users' expectations of the function of their smart thermostat did not always match the thermostats capabilities. The Nest home thermostat adapted heating schedules learned from users' interactions with the system. However the ability of the thermostat to learn preferences was complicated by too

much or too little input from users, which led to a mutual misunderstanding between the users and the system [32].

Beyond the home: Smart thermostats in shared work environments

Shared work environments such as offices represents a relatively under-explored context for smart thermostat research in HCI, with thermal comfort sensing and modelling work thus far constrained to the fields of artificial intelligence and machine learning [3, 10, 13, 18, 21, 23, 25, 26, 27]. Shared work environments provide a unique set of challenges for smart thermostats as distinct from the home. It is far easier for a system to learn the thermal comfort preferences of a family of 2-4 than it is an office of 20-40 people. Factors affecting use in the home such as cost/comfort trade-offs (e.g. [21]) are not applicable in shared work environments, given office workers or library visitors do not pay the buildings' utility bills.

Recent work suggests that the potential for participation in activities related to thermal comfort in offices can be influenced by social factors [28]. Workers in naturally ventilated offices were found to willingly tolerate thermal discomfort before requesting a window be opened or a heater adjusted, either out of consideration for other co-workers or simply not wanting to make a fuss [28]. In climate controlled buildings, achieving acceptable thermal comfort for a larger number of people can be equally complicated by the variety of personal preferences concerning how a thermostat "should" be controlled (i.e. set and forget versus adaptive control) [6]. An objective of this line of research is to understand comfort in offices in a broader context than individual thermal comfort alone, given that expectations and experiences of comfort are shaped in part by culture and conventions [4]: *"Rather than figuring out more efficient ways of maintaining 21-23° in the face of global warming, society should be embarking on a much more searching debate about the meaning of comfort and the ways of life associated with it"* [4 pp.39]. We re-visit this call-to-action later in the paper.

Mathur et al. [19, 20] visualized environmental variables in an office as part of a "Quantified Office" technology probe. Data on air quality, noise, self-reported mood and activity were collected and visualized on a screen in a break-out space. The project is somewhat unique, not only because it deployed shared informatics in a novel context, but that it represents one of the few studies of this nature to take a bottom-up approach, asking the question *"what information do employees find useful themselves?"* rather than assuming this first [19]. The authors suggest office-situated deployments should provide "actionable" quantifications-i.e. information upon which users can act [19].

Does it actually work? Validation of smart thermostats in-the-wild

Unlike the home, where smart thermostats have been observed in-the-wild [1, 32], the same is not true for smart thermostat deployments in shared work environments.

Various thermal comfort models and participatory sensing systems for office-based smart thermostats exist [3, 10, 13, 18, 23, 25, 26, 27], however the validation and practical implementation of these models does not always extend to in-the-wild user trials [10, 18, 25, 27].

Recent participatory sensing approaches for thermal comfort in offices include affording users the ability to influence the set-point of their offices' thermostat via a mobile [10, 18, 25] or PC application [27]. In each of these studies, users select their thermal comfort level between 'Cold' to 'Hot' on a scale. Initial user trials hint at the potential for considerable HVAC energy savings of 10% [10] to 18% [18]. Yet some potential difficulties remain regarding the applicability of such platforms to represent a human interface for smart thermostat over the long term. Lam et al. [18] for example requires input for the gender, height and weight for occupants; information some occupants may not feel comfortable disclosing. Some systems utilize human participation only in the training stage of their model and then rely solely on the model to provide thermal comfort from this baseline user data [13, 18]. This does not allow the system to adapt for a change of occupancy or user preferences over time. On the other hand participatory sensing systems such as [10, 23, 25, 27] all rely on regular user input over longer timescales, which may not be realistic given the tendency for user engagement in situated technology to evolve over time to less frequent interactions, or diminish all together [14, 15, 33].

Designing for "Resource Man"?

Notable in the design of several examples above, is the assumption that users will be sufficiently interested and engaged with the thermal comfort of their home or office to provide the system with regular and accurate data over extended timeframes. Such assumptions have been highlighted as problematic [30, 31]. Strengers [30] highlights the tendency of designers of smart technology to assume (wrongly) that every user is a "Resource Man"- i.e. a technically-minded rational consumer; who is aware of the factors affecting their personal informatics and willing and able to respond to fluctuations in price-signals with economically rational decisions. In reality everyday life is messy and unpredictable [31]; and many users are simply uninterested in the quantification of their personal informatics. Although the provocation of "*Are you Designing for Resource Man?*" is created for the purpose of highlighting the problematic assumptions made by energy utilities of their consumers [30, 31], the stereotype is applicable more broadly to human-machine interfaces such as smart thermostats. In particular, in the design of systems requiring human engagement or input, it is important to understand how social contexts and everyday life may influence or interrupt use [15, 29, 31]. In this paper, we outline our vision for how these assumptions of rationality may be avoided in the design of smart thermostats.

METHODS

Background

The method undertaken in this paper is usefully understood in the context of our previous work. In response to the practical limitations in applying existing thermal comfort models (i.e. Static [11], or Adaptive [7]) to smart thermostats, we proposed our own personalized thermal comfort model which only requires easily obtainable inputs from the user and utilizes Bayesian networks to adapt to individual users' preferences. This model, further described in [3] offers up to 23% more accurate predictions of users' actual comfort levels than previous models. Previously, however, the model had only been tested on training data from the ASHRAE RP-884 project [2] and requires further validation in the wild if it is to be incorporated into a smart thermostat. The motivation for this paper originated from our attempts to validate this model in the wild, using a novel means of user input which we explain here.

Experimental design and deployments

The original intent of this research was to determine people's definitions of acceptable temperatures and whether these fall within the values predicted by our model [3]. We developed a poster inviting people to log how they are feeling in relation to the temperature. We made 172 unique posters in total to be placed at various locations around (1) a university library and (2) several offices on one floor of a naturally ventilated office building in Southampton, UK.

The purpose of the deployment was two-fold. First, to validate the accuracy of our model in-the-wild with real users and in doing this, to gather an idea of the range of human factors that we anticipated such a deployment might encounter. Each poster featured the title "How's the temperature?" and a large QR code and a URL address unique to each poster (Figure 1 below).

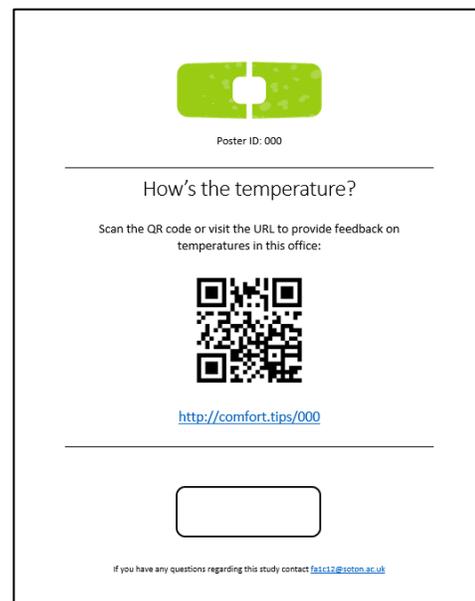


Figure 1: Poster

Scanning the QR code or entering the poster’s unique URL into a browser linked the user to a portal where thermal comfort could be logged by means of moving two sliders up or down and optionally selecting buttons representing different reasons for potentially sub-optimal thermal comfort (i.e. “Heat from radiators”, “Cold draft” etc). Refer Figure 2 below. The website was designed to be as simple as possible, for use from both PC’s and mobile devices. All votes to the system were anonymous and the website stored only a unique identifier for each login, simply to be able to identify peoples’ consecutive votes and to determine how many times and to which poster people had voted.



Figure 2: Web interface

The intention was to have the posters as physical reminders for people to register their thermal comfort with the system, in contrast to [10] and [18] who required users to open a mobile app and then select their location from a menu. On the bottom of each poster, we attached a small portable Joulo temperature logger which was pre-programmed to record the time, date and ambient temperature every 10 minutes for the university library deployment and every 4 minutes for the office deployment. This continuous logging allowed correlation of the ambient temperature at the time that each thermal comfort observation was recorded.

Library deployment: We positioned 143 posters over all 5 floors of a large university library. The library, originally built in 1914 and extended/modified in the 1930’s, 1950’s, 1990’s and 2000’s respectively, represents a patchwork of building extensions and HVAC arrangements. Heating is primarily by conventional radiators. Ventilation arrangements vary throughout the library; the majority is ventilated by forced mechanical ventilation, however some areas are naturally ventilated with user-openable windows and an area of the archives is air conditioned. In our deployment we aimed for a balance between representative geographic coverage for temperature measurements and choosing sufficiently public locations such that the posters achieved a presence in the library. Additionally 3 humidity sensors were attached to 3 of the posters in order to provide an indication of humidity across the library. Of the 143 posters deployed, 85% were placed in open spaces and 15% in bookable meeting rooms. The deployment of the posters lasted 5 weeks between May and June 2016. Adhering to ethical restrictions of the study and the wishes of the library, we did not approach students for qualitative interviews following the deployment and relied instead on

the comments box of the user interface for qualitative feedback (Figure 2).

Office deployment: In this deployment we positioned 29 posters and the three humidity sensors in 8 offices and 3 hallway locations around a single floor of a naturally ventilated office building. The offices where the posters were deployed ranged in occupancy from 2 people (one office) 3-10 people (4 offices) and three large open plan offices each with more than 20 occupants. The offices were populated with university administrative workers involved with finance, student services, human resources and legal. The posters remained in place for the two-week deployment between February and March 2016. All offices were naturally ventilated via occupant-controlled windows and were heated via conventional radiators. A glass-roofed atrium had been retrofitted to one side of the building some years prior. This connected two previously separate naturally ventilated buildings, compromising the ability of both buildings to ventilate. The smaller offices in our study with windows opening to the atrium featured portable, manually controlled air conditioning units which had been introduced owing to complaints about the heat in summer.

In each office we visited to deploy the posters we made a short address to the employees present, to explain the purpose of the posters and how to log their thermal comfort votes. We asked people to vote as often as they liked, but to ideally space votes at least one hour apart. We also instructed office members to pass on the information about the trial and extend our invitation to other office members who were not present during our visit. Wishing to replicate achievable real world conditions beyond the novelty phase, unlike [25] and [27], we did not remind people to vote. Instead we placed the posters in visible and accessible locations around the offices and corridors, so as to encourage the type of opportunistic interaction which may be sustainable over longer timescales.

Interviews

Following the office deployment, we invited participation in semi-structured interviews related to the deployment. Sampling for the interviews was restricted to those employees who worked in an office where one or more poster was deployed and who had logged their thermal comfort on one or more of the posters at least once throughout the deployment.

Semi-structured interviews averaging 30 minutes in length were carried out with 14 participants, which represented approximately 15% of the total workers in all of the offices where the posters had been deployed, and 23% of the total number of users who had logged at least one thermal comfort vote throughout the deployment. Participants comprised 13 females and 1 male. Ages ranged between 26 and 57 with a mean of 39. All participants worked for the university in administrative roles including human resources, financial services, student services and legal. The

gender bias of the interview participants is reflective of the majority female occupancy observed in the offices.

Discussions in the interviews related to the participants' work environment, the different factors affecting their thermal comfort, social dynamics around thermal comfort decisions such as opening or closing windows, symptoms they felt in the office and ways they improved their thermal comfort if it was poor. The second part of the interviews related to the usability of the system, motivations for accessing the system and potential improvements.

Analysis

Temperature data: At the end of the deployment, we collected all the Joulo temperature loggers and downloaded the logged temperature data. In the library deployment, 17 of the 142 loggers had either disappeared (12), stopped working (2) or had been configured wrongly (3). For feedback provided to these loggers, we interpolated temperature data from the nearest three loggers. We used thin plate interpolation to calculate an approximation of the temperature between the loggers for the visualization of the heat map created from the data (Figures 3A and 3B).

Feedback from user interface: Data from the “Comments Box” of the thermal comfort interface was downloaded to a spreadsheet. These were read individually and assigned into four categories and later quantified: (1) heat-related (2) cold-related (3) ventilation/stuffy/humidity-related (4) other.

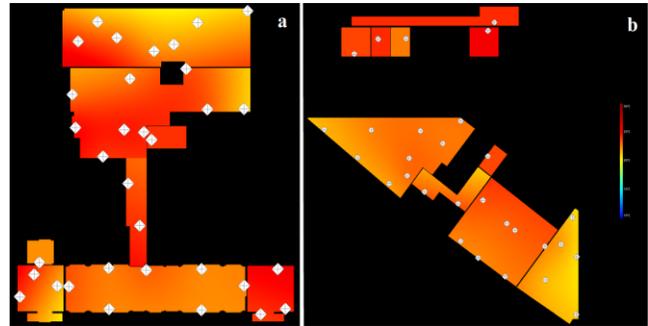
Interviews: All interviews were transcribed verbatim and read by the research team. Initially, responses were aligned to the different key questions asked. Following this, a second read of the transcripts was used in a thematic analysis where transcripts were read independently of the question asked and responses were manually coded according to themes emergent from the data.

FINDINGS

The findings from both the library and office deployment are presented in terms of (1) an overview of the quantitative findings from both deployments, (2) a discussion of the key motivations for users to record their thermal comfort, (3) reflections on usability aspects of the system from both deployments and (4) social factors and office politics observed.

Overview: Temperature

The average temperature across all Joulos over the study periods of both deployments were 24.67°C in the Library and 22.13°C in the Office. As can be seen in Figures 3A and 3B, however, this temperature was not evenly distributed geographically.



Figures 3A and 3B: Averaged geographic temperature distribution across the library (3A) and office deployment (3B)

Figures 3A and 3B show the temperature distribution across the top floor of the Library (3A) and across the entire deployment across the single floor of the Office (3B).

In the Library, we offer no explanation for the distribution of the temperature, as the building represents a patchwork of retrofits and HVAC situations. The relatively high average temperature during the library deployment was due to one unusually warm week in early May coinciding with the winter heating not yet turned off. The temperatures varied slightly less in the office deployment, with the warmest offices being those with windows that opened out to the enclosed atrium (Figure 3A and 3B).

Overview: Votes logged

Both deployments received a substantial number of votes throughout the respective deployments (refer Table 1). Votes were relatively well distributed between posters, however some posters located in quieter areas of the library received none. Table 1 (below) shows the distribution of the votes between the two deployments.

	<u>Library</u>	<u>Office</u>
Total votes logged	990	167
Total users who voted	688	57
Av no. votes per user	1.44	2.93
Av no. votes per poster	6.92	5.76

Table 1: Distribution of votes

Finally, Figures 4 and 5 below show the temporal distribution of votes over the course of the two deployments

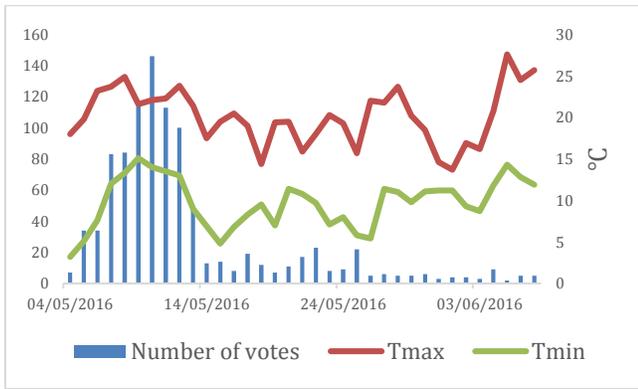


Figure 4: Number of votes plotted against min and max outdoor temperature: Library deployment

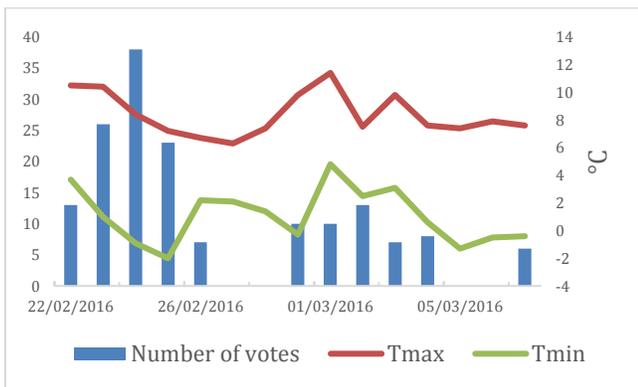


Figure 5: Number of votes plotted against min and max outdoor temperature: Office deployment

Figures 4 and 5 both show a pronounced drop in the number of votes recorded per day following an initial peak. The pronounced novelty affect described above was validated in the qualitative interviews, where five participants recalled being too busy or forgetting to log their thermal comfort on the system in the second week of the deployment: “[I was] busy and kept forgetting. It was only when I felt like, ‘Oh god! It’s too hot!’” (P7). In the Library (Figure 4), as mentioned, a warm week in May occurred while the heating was still active, overheating areas of the library and prompting a large number of votes. As such it is unclear in Figure 4 (library) as to whether the drop in votes correspond to a novelty effect or simply a return to more reasonable temperatures in the library once the weather cooled and the heating was turned off.

Far more significant than the timing of votes, however, was the nature of the votes recorded in terms of thermal comfort. Figure 6 shows the distribution of all temperatures logged by the Joulos throughout the office deployment (blue) and the normalized distribution of votes per temperature logged (red).

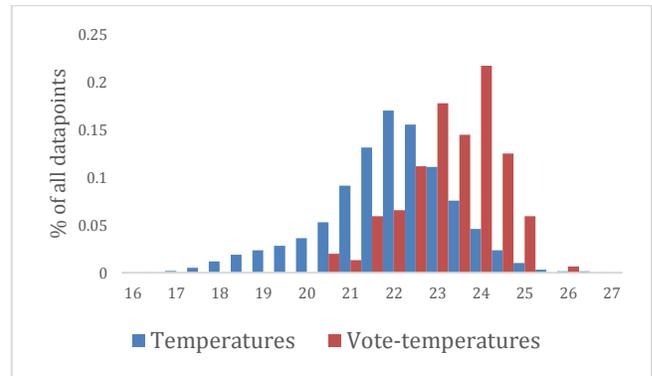


Figure 6: Distribution of indoor temperatures and votes per temperature

What can be seen here is the marked tendency for people to log it while it’s hot- i.e. the vast majority of votes were placed when the temperature was higher.

In the library, of the 990 total votes cast over the four-week deployment. 88.4% corresponded to people feeling warm or hot with only 7.96% registering a ‘neutral’ vote and 3.64% wanting it to be cooler. This pattern was repeated in the office, with only 15% of all the votes correlating to employees wanting it to be cooler and less than 10% of all the 167 votes cast corresponded to people feeling neutral. In this respect, our system was essentially used as a digital complaints box.

Due to the single-sided nature of the votes recorded in our study as described above, we were unable to validate our comfort model on this data. To accurately learn user preferences, the model requires votes in a range of different conditions. Having only ‘too hot’ votes in a steadily hot environment causes the model to be biased towards predicting ‘too hot’ in any case, regardless of the actual thermal conditions.

Motivations for use- Heat

It emerged that motivation for recording thermal comfort votes in the systems was primarily related to the temperature in the office. Of the 303 comments entered into the comments box of the user interface from the library deployment, 251 were concerned with the heat or being too hot, compared to only 5 relating to being too cold. “Too hot on Level 1”, I’M MELTING!!!”, “Can’t concentrate”. Lack of ventilation, high humidity and stuffiness were also commonly reported alongside temperature complaints.

In the interviews following the office deployment, eight of the 14 participants recalled being reminded to vote only when they felt uncomfortable, i.e. too hot or cold and did not, or simply forgot, to interact with the system when they felt comfortable:

“I think you’re more likely to remember to vote if you are feeling cold or hot or warm. It was just a reminder in your mind, ‘Ooh, we need to log these things,’ rather than when you were feeling just comfortable with anything” (P4)

"[Voted mainly]... near the beginning of the trial, mainly because it was that was when it was hotter and I wanted to comment on it and then later in the week it cooled down. [laughter]" (P11)

This shows how discomfort acted as a key reminder to interact with the system and users were less likely to remember to interact with such a system if they were feeling comfortable. Despite our purposeful placement of certain posters in thoroughfares, while several of our interviewees reported opportunistic interactions with the system, (i.e. being reminded by walking past a poster) this reminder appeared ultimately insufficient to retain peoples' engagement with the system into the second week of the deployment (Figures 4 and 5). Contrary to our expectations, social motivations did not feature prominently. Only one participant mentioned being socially motivated to vote. In this case, P12 would vote together with another colleague, typically on the way to or from a tea break.

Motivations- Fighting for a cause

What the quantitative results for temperature and thermal comfort in the office deployment do not illuminate, is the long history of negotiations between occupants and building managers ("Estates and Facilities") around the issue of overheating and ventilation. The thermal comfort of the building was an on-going cause of frustration among participants and resulted in a persistent tension between building occupants and the Estates and Facilities department. Interviewees related the different work-arounds employed over the recent years to attempt to improve thermal comfort in the building:

"They've installed some extra ventilation things to try and help with the ventilation and the temperatures in those rooms, but I don't know if it works (P1)

"During the summer, they issued everyone with a little bottle of water so that they could spray themselves because it was that hot. That was the only way they could think of trying to cool people down" (P8)

Despite this, it was felt these work-arounds were ultimately unhelpful and participants continued to feel unsatisfied with their office environment, yet powerless to improve their situation.

Unhappiness with the thermal performance of the building (at least in summer) represented an exceptionally strong theme in the first interview (P1). Based on this, we then asked P1 and each subsequent participant specifically about whether they believe they would have engaged with the system less if they knew the data would not be shared with anyone who could make a difference. Nine of the 14 participants agreed this to be true for them. Participant 14 summed this up: *"I would probably just say, 'Oh, well what's the point then?'"*. These comments were made in full knowledge of the fact that we were researchers and were not connected to the Estates and Facilities department.

Despite this, when also asked specifically whether they might have exaggerated their responses (i.e. voted "hot" when feeling neutral) in pursuit of generating a more extreme outcome, none of the participants agreed they had exaggerated their votes. This shows that despite participation with the system being influenced by the wish to have their voices heard, that participants still voted truthfully about their thermal comfort and did not exaggerate their responses.

Usability of the system

Being HCI researchers, we were particularly interested in the usability of the poster and the web interface, and asked a number of questions to these ends in the qualitative interviews. Interestingly however, usability did not emerge as a specific issue in any of the interviews and none of the 303 comments written in the Comments Box of the library deployment related to usability. Interview participants reported the interface easy to access through both the QR code and the URL, despite one participant noting she needed to be taught how to download and use a QR scanner on her phone by another colleague. Two participants mentioned that they would like to have been able to list the adaptive measures they had already taken before voting, e.g. Removed cardigan (✓) removed shoes (✓), turned on desk fan (✓) and *still* feeling "hot".

Although usability was not flagged as a problem, when asked, several participants agreed they would like to have more quantification of the conditions in the office. Six of the 14 participants mentioned they would have liked a quantification of how stuffy it was in their office. The sentiments of the majority of participants was summed up by P4: *"We feel like the air quality is low, but nothing to show any measurements for that" (P4)*. Towards this, two participants, had already brought their own thermometers in from home and we found that one of the digital humidity sensors we had originally deployed in a corridor had been moved into an office. P11 explained that "one of the guys" from the office next door had done this, so they could press the button on the logger and read the temperature and humidity in the office.

Social factors around air quality

A key part of the interview process was to gather qualitative data on the social context of thermal comfort and air quality in offices. We observed that negotiations around thermal comfort were mostly polite and amiable. However, we found ample evidence suggesting that social factors such as politeness and empathy can serve to complicate otherwise rational actions towards improving one's satisfaction with thermal comfort conditions in the office. People would not always speak up if they were unhappy with the air quality in the office, out of not wanting to cause a fuss, or discomfort to their co-workers. The fact that the adequate ventilation of certain offices relied on opening windows was problematic in the winter:

“I might want the window open which will then cause a cold draught to somebody... Yeah, I'd like to open a window but that might not be right... So I'll turn my fan on [instead]” (P11)

Participant 4 believed a number of people in her office elected to put up with air quality they were unhappy with, purely out of consideration for others.

Interviewer to P4: *“Do you think that there's a lot of people who suffer in silence, who'd like the window to be open but don't want to disturb others?”*

P4: *“Oh, definitely at this point of the year, yeah, you will find that in the office.... [Later]: I think people will sort of, before they give their own opinions, will look at everybody else first, and then they'll go along with the consensus”*

This was the case for P7 who put consideration for others before her own thermal (dis)comfort in relation to desk fan use. Although her fan caused her discomfort, she turned it on anyway out of consideration for others:

“I wear contact lenses, so if I have the fan on directly in my face, then it dries them [my eyes] out. I do struggle with that, but I tend to keep quiet in the summer because I know how bad it is in there” (P7)

It can be seen here how consideration for others can be placed before personal comfort, resulting in participants not taking advantage of very simple measures to make themselves more comfortable.

P3 and P5 who worked in one of the three large offices both noticed a marked discrepancy in their office between people who feel the cold versus people who feel the heat. P3 related this as: “two factions” in the office- *“people who are permanently hot no matter what the temperature is and some who are really, really cold”* (P3). This did not affect social relations in the office and both participants mentioned the negotiations around thermal comfort in the office were almost always agreeable:

“I'm sitting there watching them and layers come off and then the fan goes on and the window goes open. And it is generally warm. The other faction are permanently walking around with neck scarves and jackets on, so they make it fairly... [obvious]” (P3)

These findings highlights that the agreeable and polite negotiations around thermal comfort do not necessarily equate to satisfaction with thermal comfort, nor satisfaction with ways that others achieve their own personal comfort. What also becomes particularly clear in these findings is the variety of ways that social factors in offices can impact thermal comfort and air quality and that access to means of adjusting thermal comfort (i.e. windows, radiators, fans),

does not mean they will be used in ways that might be rationally expected.

DISCUSSION: TOWARDS A SMARTER THERMOSTAT

In this paper we have argued that designing for human interactions with smart thermostats is just as important as the accuracy of the thermal comfort model to which they are grounded. Results show that contrary to our expectations that interactions with the system would be largely opportunistic, motivations for interacting were in fact much more complex. People tended to vote only when they were uncomfortable and many interactions in the office deployment took place in the hope that the resultant data might provoke actions by building managers to improve conditions in the building. In this way the system was used as something of a digital complaints box and the lack of neutral or cold votes negated our ability to validate our model on this data.

Given evidence that females are more likely to express dissatisfaction than males in equal thermal conditions [17], it might be tempting to conclude that the gender bias in our interviews is indicative of females' exhibiting a higher thermal sensitivity than males. While we acknowledge the gender bias, we do not wish to speculate here on the possible influence of gender, given: (1) the existing female-weighted gender balance of the offices; (2) the user-feedback system not recording input for the gender of contributors; and (3) the possibility that the females who volunteered for our interview were simply more pro-active in volunteering for the interviews than their male counterparts, rather than being more thermally sensitive.

Overall, participation in the study was characterized by a decline in engagement over the course of the deployment (Figures 4 and 5) with participants reporting being busy or simply forgetting. This finding is consistent with literature on personal informatics and eco-feedback, where engagement with situated systems has also been found to decline over time or transition to less frequent forms [14, 15, 29, 33]. Additionally, we found that polite and amiable negotiations around adjustments of windows, radiators or desk fans did not necessarily equate to satisfaction with thermal comfort. These findings echo those of previous research [28]. These various and bespoke social factors affecting use are not yet well represented in the literature on smart thermostats in shared work environments. We have also demonstrated the limits of assuming humans to be reliable sensors for thermal comfort, willing and able to provide accurate data at regular intervals. In this way we have shown how the falsity of the “Resource Man” stereotype [30] extends from the home to shared work environments.

Our findings suggest that we need to eschew the assumption that humans will act as regular and reliable sensors for thermal comfort. A move is necessary towards design seeking to take advantage of when a person actually wants to interact with a thermostat- or any such technology-

rather than expecting people will interact (and continue to interact) with a system out of altruism or obligation. From these findings we have synthesized design considerations towards a *Smarter Thermostat* for shared work environments; one which: (1) Leverages the messy social inevitabilities that we have highlighted in our deployment so they are, by design, equipped to handle changes in user engagement, and; (2) Augments human intelligence through reciprocity and actionable information. We discuss these in turn below.

Leveraging the Digital Complaints Box: Equipped to handle diminishing user engagement

Existing models for participatory sensing of thermal comfort in shared work environments operate either by: (a) learning user input only during an initial training phase; thereby being less adaptable to changes in occupancy [13, 18], or (b) requiring relatively high and sustained levels of human input into the system in order to function [23, 25, 27] which may not be sustainable in the long term. In our study, despite aiming for opportunistic engagement, the average number of votes per user was only 1.44 and 2.93 in the library and office deployments respectively, and overall engagement with the system declined over time.

For this reason, we believe a *Smarter Thermostat* is one which is equipped to handle diminishing user engagement; and which is capable of operating within the constraints of users' tendency to only log their thermal comfort when they are uncomfortable. As such, we envisage a *Smarter Thermostat* utilizing user input to instead determine users' *limits of comfort*. Precisely, the thermostat determines to what extent it can mimic the outdoor temperature before complaints begin to be registered. In the training phase, this would involve the thermostat operating within a purposefully large temperature range and narrowing this range based on users' feedback if and when the environment is uncomfortable.

In this way, the thermostat operates within the largest acceptable range for a given office. Given a key goal of smart thermostats is saving energy [1, 9, 21], the energy saving implications for an HVAC system maintaining the temperature in an office between (for example) 20 and 25°C are far greater than the same system attempting to keep the temperature within a narrower range of 21-22°C as conventional thermostats often attempt to. This in turn encourages human adaptability to fluctuations in temperature [7] and maintains a better connection to outdoor conditions. Encouraging adaptation in this way holds the potential to broaden office occupants' understandings and expectations of comfort beyond ubiquitous mechanical intervention [4]. Pragmatically, a *Smarter Thermostat* is not affected by the likely decline in user engagement over time attributed to other situated interfaces [14, 15, 29, 33], but it still provides users with the ability to log their thermal (dis)comfort, should they become dissatisfied with it. This simply means maintaining

an interface for human input as we did throughout our deployments, such that the system may continue to respond to changes in occupancy, such as employee turn-over or visitors with different thermal comfort preferences.

Augmenting human intelligence: Ensuring smarter thermostats = smarter humans

A common feature of smart thermostat design is their ability to learn occupants' temperature preferences and then autonomously provide user-optimal temperatures with minimal further input [18, 27, 32]; gradually reducing the need for users to interact with the system. However we are not convinced such a rationale for design is necessarily in the best interests of users. We share an unease expressed by Rogers [24] with the tendency of "smart" technology to take responsibility away from humans; reducing the need for them to think for themselves or engage with the issue at hand. We argue that such a rationale runs the risk of creating users who are less adaptable to changes in temperature [7] and less engaged with the smart technology designed to support them.

Indeed, although engagement in our system was variable, we found ample evidence that the quality of the indoor environment was of keen interest to many of our participants. Two had brought thermometers in from home to measure the temperature, four had engaged directly with Estates and Facilities and six of the 14 participants noted how they would have liked a quantification for how stuffy it was in the office. Thus we see it as important that smart technology aims to foster this engagement [24]

We argue a *Smarter Thermostat* is one which aims to augment intelligence, not absolve responsibility. Rather than users entering data into a black box which is supposed to improve their thermal comfort- as is currently the case- we draw from Mathur [19] and suggest a *Smarter Thermostat* could reciprocate user input, by offering users information they can act upon themselves. An example of what this could look like in practice is Figures 3A and 3B in our findings, where we have provided a thermal-map from both the office and library deployment. Users would be able to use such maps to choose (or even be recommended) warmer, cooler, fresher, i.e. more personally suitable locations to work in a library or open-plan office. Such a system supports the various personal preferences of users that we identified (i.e. preferences for cool or warmth) and how these preferences may change over time, for example new arrivals to a country becoming acclimatized, or due to personal circumstances such as menopause (P12), pregnancy (P10) or other health conditions (P14).

Yet we have demonstrated that maintaining thermal comfort in the workplace is not an isolated personal adjustment. It is a negotiated, social activity. Smarter Thermostats incorporating a form of situated display as we have advocated above, allows reflection on how individual and group actions impact upon others. For example being able to see how the opening of a certain window causes a

cold spot over an adjacent colleague. We envisage Smarter Thermostats contributing to shared understandings of thermal cause-and-effect within offices; contributing to comfort being managed and negotiated as a team, rather than only by individuals taking individual actions, as is currently the case. We echo Chappels and Shove's call to promote debate around definitions of comfort itself, rather than concentrating only on how to maintain a narrow comfort range more efficiently [4]. Thermostat design of the nature we describe here may be a starting point for this.

Beyond the office, there is the potential to augment the intelligence of external parties such as building managers or health and safety officers. This is important given the propensity for these parties to act at times as "gate keepers" for ceiling vents, heating boilers or external windows [28]. Providing these groups with more granular and immediate information on conditions within different offices allow for a quicker response time to open vents or adjust settings if conditions become sub-optimal. This type of information could also provide useful feedback on the thermal performance of recent retrofits, space planning, or identifying malfunctions in HVAC performance.

Considerations for future work

Finally, we briefly outline privacy as a concern for future work in this space. Although privacy was not a primary focus in our interviews, we anticipate that privacy concerns have the potential affect the use of smart thermostats in shared work environments. Our interview findings suggest an inclination for users to put up with a considerable level of thermal discomfort rather than speak up, so as not to cause a fuss. Thus despite many potential benefits of thermal visualizations as we have described above, we do not believe that users' votes should feature in this information. Even if displayed in an aggregated form (for example thermal maps based on user votes rather than ambient temperature), this may still affect use. As an example, aggregated data can easily become attributable if surrounding colleagues are absent. Certain people may elect not to vote in order to maintain the status quo, in precisely the same way that some participants did not vocalize their opinions in our study. Further research is warranted into how a *Smarter Thermostat* might best reciprocate user participation without affecting their inclination to participate in the first place.

CONCLUSIONS

Considering the performance of smart thermostats in shared work environments is contingent upon human input into the system, it is surprising the current deficit in human-centred approaches to design. The research described above represents a contribution to the literature to address this gap, in offering a tested, novel means of gathering thermal comfort information from users in shared work environments. Based on our results, we have argued that the design of the human interaction with smart thermostats is just as important as the accuracy of the thermal comfort

model to which it is grounded. Towards this, we outlined two key considerations for the design of a *Smarter Thermostat* concerned with; (1) coping with the messy reality of user engagement; and (2) augmenting human intelligence rather than absolving responsibility.

This work offers several contributions to the literature on smart thermostats in shared work environments:

First, using a bottom-up approach to understanding thermal comfort, our findings demonstrate a multitude of individual, social and contextual factors affecting the use of participatory sensing systems. These include use of the system being influenced by discomfort, pre-existing relationships with building managers and a novelty affect with usage over time. These aren't yet well accounted for in the existing literature on smart thermostats in shared work environments.

Second, we highlight the limits of relying on users as reliable and autonomous sensors. Our findings foreground the need for smart thermostats to be capable of providing thermal comfort and energy efficiency within the realistic constraints of the type and frequency of human data they are likely to receive.

Third, we have highlighted the possibility and value of reciprocity in smart technology; opportunities for smart systems to augment human intelligence- rather than absolve responsibility- by providing users with actionable information. The value of such information may extend beyond individual users to foster a more shared understanding of comfort in offices and the creation of thermal datasets for use by building managers, energy auditors and potentially policymakers. Yet further research is warranted into users' expectations and preferences for privacy disclosure.

Although much further work is necessary, we anticipate this rationale for smart thermostats has relevance to the design of smart systems more generally. Namely, that any smart technology might benefit from design which equips it to deal with changes in use that arise from the messy nature of everyday life; and that might enable a more empowering experience by reciprocating its user input with useful information.

ACKNOWLEDGMENTS

This work is funded by the Engineering and Physical Sciences Research Council, UK, under the Challenging Engineering scheme. The ReFresh Project: Remodeling Building Design and Sustainability from a Human Centered Perspective (grant number EP/K021907/1). The Joulo temperature loggers were developed with the assistance of: University of Southampton Carbon Management Plan project - 'CMP165 - Joulo Temperature & Light Logging and Analysis of Offices and Laboratories' - awarded to Prof. Alex Rogers and Dr. Reuben Wilcock.

REFERENCES

1. Alan, A. T., Shann, M., Costanza, E., Ramchurn, S. D., & Seuken, S. 2016. It is too Hot: An In-Situ Study of Three Designs for Heating. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 5262-5273. ACM.
2. ASHRAE 55, 2010. Thermal Environmental Conditions for Human Occupancy (ANSI Approved), 2010.
3. Auffenberg, F., Stein, S., Rogers, A. 2015 A personalised thermal comfort model using a Bayesian network, *Proceedings of the 24th International Conference on Artificial Intelligence*, p.2547-2553, July 25-31, 2015, Buenos Aires, Argentina.
4. Chappells, H., & Shove, E. 2005. Debating the future of comfort: environmental sustainability, energy consumption and the indoor environment. *Building Research & Information*, 33(1), 32-40.
5. Chin, J. 2015. Design and implementation of an adaptive wearable thermal comfort data acquisition prototype. 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* (pp. 585-590). ACM.
6. Clear, A., Friday, A., Hazas, M., & Lord, C. 2014. Catch my drift?: Achieving comfort more sustainably in conventionally heated buildings. In *Proceedings of the 2014 conference on Designing Interactive Systems (DIS)*, pp. 1015-1024. ACM.
7. de Dear, R.J., and Brager, G.S. 1998. Developing an adaptive model of thermal comfort and preference. In *ASHRAE Transactions*, volume 104, pp. 145–167. ASHRAE.
8. Department of Industry, Innovation and Science. (2016) Heating, Ventilation and Air Conditioning High Efficiency Systems Strategy. Available at: <http://industry.gov.au/Energy/EnergyEfficiency/Non-residentialBuildings/HVAC/FactSheets/Documents/HVACFSEnergyBreakdown.pdf>. Accessed: 18/9/16
9. EPA. Summary of research: Findings from the programmable thermostat market. Washington, DC: 2004.
10. Erickson, V.L. and Cerpa, A.E. 2012. Thermovote: Participatory sensing for efficient building HVAC conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pp. 9–16.
11. Fanger, P.O. 1970 Thermal comfort. Analysis and applications in environmental engineering. In *Thermal comfort. Analysis and applications in environmental engineering*, pp.244.
12. Feldmeier, M., & Paradiso, J. A. 2010. Personalized HVAC control system. In *Internet of Things (IOT), 2010*, pp. 1-8. IEEE.
13. Ghahramani, A., Jazizadeh, F., & Becerik-Gerber, B. 2014. A knowledge based approach for selecting energy-aware and comfort-driven HVAC temperature set points. *Energy and Buildings*, 85, pp.536–548.
14. Gulotta, R., Forlizzi, J., Yang, R., & Newman, M. W. 2016. Fostering Engagement with Personal Informatics Systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS)* (pp. 286-300). ACM.
15. Hargreaves, T., Nye, M., Burgess, J. 2013. Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term. *Energy Policy* 52, pp. 126-134.
16. Huang, C.C., Yang, R., Newman, W.W. 2015. The potential and challenges of inferring thermal comfort at home using commodity sensors. In *Proceedings of UbiComp '15*. ACM, pp.1089-1100.
17. Karjalainen, S. 2012. Thermal comfort and gender: a literature review. *Indoor air*, 22(2), 96-109.
18. Lam, A.H., Yuan, Y., and Wang, D. 2014. An Occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings: Categories and subject descriptors. In *Proceedings of the fifth International Conference on Future Energy Systems (ICFS)*, pp. 133–143.
19. Mathur, A., Broeck, M. Van Den, Vanderhulst, G., Mashhadi, A., Kawsar, F., & Laboratories, B. 2015. Tiny habits in the giant enterprise : Understanding the dynamics of a quantified workplace. *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*, pp.577–588.
20. Mathur, A., Broeck, M. V. D., Vanderhulst, G., Mashhadi, A., & Kawsar, F. 2015. Quantified workplace: Opportunities and challenges. In *Proceedings of the 2nd workshop on Workshop on Physical Analytics*, pp. 37-41. ACM.
21. Panagopoulos, A. A., Alam, M., Rogers, A., & Jennings, N. R. 2015. AdaHeat: A General Adaptive Intelligent Agent for Domestic Heating Control. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems* (pp. 1295-1303)
22. Pisharoty, D., Yang, R., Newman, M. W., & Whitehouse, K. 2015. Thermocoach: Reducing home energy consumption with personalized thermostat recommendations. In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, pp. 201-210. ACM.

23. Purdon, S., Kusy, B., Jurdak, R., & Challen, G. 2013. Model-free HVAC control using occupant feedback. In *Proceedings - Conference on Local Computer Networks, LCN*, pp. 84–92.
24. Rogers, Y. 2006. Moving on from weiser's vision of calm computing: engaging ubicomp experiences. In *Proceedings of the 8th international conference on Ubiquitous Computing (UbiComp'06)*, Paul Dourish and Springer-Verlag, Berlin, Heidelberg, pp.404-421.
25. Sarkar, C., Nambi, A. U., & Prasad, R. V. 2016. iLTC: Achieving Individual Comfort in Shared Spaces. In *International Conference on Embedded Wireless Systems and Networks (EWSN 2016)*, pp. 65–76.
26. Shann, M., and Seuken, S. 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, pp. 2892–2899
27. Shetty, S. S., Chinh, H. D., & Panda, S. K. 2015. Strategies for Thermal Comfort Improvement and Energy Savings in existing Office Buildings using Occupant Feedback. In *2015 IEEE International Conference on Building Efficiency and Sustainable Technologies*, pp. 23–27.
28. Snow, S., Soska, A., Chatterjee, S. K., schraefel, m.c. 2016. Keep calm and carry on: Exploring the social determinants of indoor environment quality. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 1476-1482.
29. Snow S., Buys L., Roe P., Brereton, M. 2013. Curiosity to cupboard: self-reported disengagement with energy use feedback over time. In: *Proceedings of the Australian Conference on Human-Computer Interaction (OzCHI 2013)*. ACM, pp. 245–254
30. Strengers, Y. 2014. Smart energy in everyday life: Are you designing for resource man? *Interactions* 21.4 pp. 24-31.
31. Strengers, Y. 2013. Smart energy technologies in everyday life: Smart Utopia? Springer.
32. Yang, R and Newman, M. W. 2013. Learning from a learning thermostat: Lessons for intelligent systems for the home. *Proceedings of the 2013 ACM international joint conference on Pervasive and Ubiquitous Computing (UbiComp '13)*, pp. 93–102.
33. Yang, R., Newman, M. W., & Forlizzi, J. 2014. Making sustainability sustainable: challenges in the design of eco-interaction technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM pp. 823-832.