

The Potential and Challenges of Inferring Thermal Comfort at Home Using Commodity Sensors

Chuan-Che (Jeff) Huang, Rayoung Yang, Mark W. Newman

School of Information, University of Michigan

Ann Arbor, Michigan, USA

{chuanche,rayang,mwnewman}@umich.edu

ABSTRACT

For decades, researchers have investigated ways to infer human thermal comfort. Studies have usually required cumbersome sensors and human observers, making them inappropriate for use in naturalistic settings such as the home. Emerging wearable and smart home sensing devices offer the opportunity to develop new models of thermal comfort based on data collected *in-situ*. To explore this opportunity, we deployed a sensing system in seven homes and collected self-report data from 11 participants for four weeks. Our system captures many factors employed in previous thermal comfort research, as well as new factors (e.g., activity level, sweat level). Machine learning-based models derived from the collected data show improvement over previous techniques, however significant prediction errors remain. In analyzing these errors we identify six problems that pose challenges for inferring comfort in the wild. Based on our findings, we suggest techniques to improve future in-situ thermal comfort modeling efforts.

Author Keywords

Smart home; thermal comfort; wearable sensors.

ACM Classification Keywords

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

INTRODUCTION

Researchers have been exploring ways to model human thermal comfort for more than 40 years. Such exploration is driven by the potential for increasing people's quality of life (e.g., improving building comfort quality); reducing energy consumption (e.g., intelligent thermostats that react to people's comfort level [9]); and advancing knowledge of the connection between physiological and psychological factors regarding comfort. While prior work on modeling thermal comfort such as Predicted Mean Vote (PMV) [8],

and Adaptive Thermal Comfort [24] provide insights into the major factors affecting people's comfort, there remain several challenges for inferring thermal comfort in real homes and offices. First, devices such as near-body temperature sensors traditionally used to collect measurement data are bulky and cumbersome for people to carry or wear on their body. Second, trained human observers or extensive questionnaires have generally been needed to record data that are difficult to detect with available sensors such as clothing insulation and activity level. Various activities at home such as cooking, dining or cleaning are known to affect one's thermal comfort, but these factors have been difficult to capture without the presence of human observers. Finally, most of the previous models were designed to serve large groups of people (e.g., occupants of an office building) rather than individuals or small groups, and thus are not ideal for settings that contain only a few people, such as the home.

These limitations have prevented researchers from developing techniques that could continuously infer one's comfort in naturalistic settings, especially for places where people conduct varied activities and exhibit adaptive behaviors. Furthermore, these limitations make such models unsuitable for UbiComp applications such as intelligent thermostats that intend to improve occupants' comfort by responding to present conditions [7].

Recently, wearable fitness trackers and smart home sensors have become widely available (e.g., Fitbit [31], SmartThings [32]). Although wearable fitness trackers are designed to monitor physical activities or health status, we observe that they could also capture several factors that are influential to human thermal comfort, such as activity level, ambient temperature, and sweat secretion. These emerging sensing devices provide the opportunity to address several challenges with prior approaches and permit the instrumentation of everyday home environments to gather data for inferring thermal comfort. By integrating commodity wearable and in-home sensors, we envision a system that is able to infer thermal comfort in-situ and at home with minimal setup and disruption.

While we are not the first to consider comfort sensing in naturalistic settings, our work offers three novel contributions: First, we present a new technique to infer thermal comfort using off-the-shelf wearable and in-home sensors in a domestic environment. Using wearable sensors

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

UbiComp '15, September 7–11, 2015, Osaka, Japan .

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

ACM 978-1-4503-3574-4/15/09...\$15.00.

DOI: <http://dx.doi.org/10.1145/2750858.2805831>

allows us to automatically obtain data important to inferring individual's thermal comfort that was previously difficult to obtain, as well as new sources of data that could prove valuable including physical movement (which can be converted into Metabolic Equivalent of Task—a factor in PMV models), sweat level (inferred from galvanic skin response (GSR)), and skin temperature.

Second, we demonstrate the feasibility of our technique by conducting a 4-week sensor deployment and ESM study [14] in seven households. The findings from our feasibility study validate the potential of this new technique in inferring personal thermal comfort at home under naturalistic conditions.

Our third contribution identifies six situations that pose challenges for inferring people's thermal comfort at home. We believe that these situations will be challenging for any system that aims to automatically infer thermal comfort in naturalistic settings.

RELATED WORK

Models of Thermal Comfort

Numerous studies have been sought to improve our understanding of human thermal comfort. Such studies have covered the influence of physiological factors, acclimatization, and culture on thermal comfort [1–3,8,16,17,26]. Researchers have investigated the effects of these factors in steady-state and in transition (e.g., changing from cold to warm) [12]; applied heat transfer theory to derive formulas for thermal comfort [8]; and used machine-learning to infer comfort models [23]. While most of the early research regarded people as passive dwellers with no control of the environment, more recent studies have demonstrated the importance of viewing people as active agents who actively configure the environment to maintain their comfort level [24]. While it is impossible to discuss all the thermal comfort models explored in prior literature, we briefly describe two that are the most widely used and relevant to our work—Predicted Mean Vote (PMV) [8] and Adaptive Thermal Comfort [24].

Developed by Fanger [8], PMV operates on the assumption that human thermal comfort is achieved when thermal load, skin temperature and sweat secretion are within a comfortable boundary, given an activity level. Based on these assumptions, the Predicted Mean Vote provides an index that combines six parameters deemed essential to thermal comfort—air and radiant temperatures, humidity, wind velocity, clothing level and metabolic rate. Given these parameters, the model can produce an index ranging from -3 (cold) to +3 (hot), indicating the thermal comfort quality of a building environment. Although this index has been shown to be accurate for buildings with central heating, ventilating, and air conditioning (HVAC) systems and to obtain the mean thermal comfort of large groups of people, it has not worked well for buildings without centrally-controlled systems or for individuals [17,24].

On the other hand, Adaptive Thermal Comfort [24] provides an explanation for why such deviation from PMV exists—people are actually more tolerant to warm and cold conditions than PMV predicts. The primary reason is that people adapt by various means, such as opening windows or changing clothes. This suggests that buildings that support adaptive behaviors, like the home, could allow a wider thermal comfort zone.

Thermal Comfort Sensing at Home: Challenges and Opportunities

In this work, we focus on inferring personal thermal comfort at home continuously while allowing occupants to behave naturally. This focus is important for three reasons. First, having a better technique to infer personal comfort at home could make a significant impact, as research has shown that people spend most of their time at home—around 15.6 hours per day in the U. S. [20]. A better way to infer personal comfort at home could potentially improve the comfort of home residents. For example, intelligent thermostats could potentially be developed that react to occupants' comfort level in real-time [9]. Second, to the best of our knowledge, there is no existing technique that allows inferring of thermal comfort at home while occupants are conducting their natural routines and exhibiting adaptive behaviors. Third, models such as PMV have been shown to be inaccurate for inferring comfort for a small group of people [17]. The average U. S. household is only 2-3 people [33], a group too small for models such as PMV to make accurate inference.

Additionally, there are two major barriers that limit previous research from reaching into everyday households. First, many previous studies relied on cumbersome equipment such as bulky near body temperature and humidity sensors [1,2]. These devices were inconvenient for users to carry continuously. Second, trained observers or extensive questionnaire were often required to collect data that are influential to inferring thermal comfort, but challenging to detect via sensors, such as activity level and clothing insulation [1]. Due to these barriers, it would be infeasible to apply previous sensing methods in UbiComp applications. For example, Clear et al. [7] have outlined several possible applications that use thermal comfort as a system input, including, for example, a thermal comfort portal that allows people to reflect on their practice of maintaining thermal comfort.

Comfort Sensing Systems for Naturalistic Settings

To infer human thermal comfort in naturalistic settings, sensors and other tools used for data collection must be minimally disruptive, blending into people's everyday routines. Although a less explored area, some research has investigated various approaches to reach this goal. For example, Feldmeier and Paradiso [9] developed a system that continuously infers comfort in naturalistic settings. While their studies were conducted in an office setting, it would be feasible to apply their approach in the home, inasmuch as the sensors used are amenable to the home

environment. They developed a machine learning-based model that predicts thermal comfort based on input from wearable and embedded indoor temperature as well as humidity sensors. The predictions were then used to control the HVAC system in a large (zoned) office building. However, their models incorporated limited physiological information since they collected only room temperature and humidity with the worn sensors. None of the other essential physiological factors such as metabolic equivalent were included in the models, nor were those factors captured by the sensors used in their study. Additionally, their system was deployed in an office building, thus how to adapt it to perform effectively in home is still a question.

SPOT [13] and the system developed by Nouvel and Alessi [25] also aim at inferring personal comfort in naturalistic settings. SPOT uses the Microsoft Kinect's skeleton tracking capability [30] for inferring metabolism. The system categorizes the skeleton information into four types of posture, namely reclining, seated and relaxed, sedentary activity, and standing. A predefined metabolic rate is then mapped to each of the postures. However, one's metabolic rate might vary significantly in certain postures (e.g., "standing"), meaning that estimates can be wildly inaccurate. The system developed by Nouvel and Alessi expects people to provide explicit comfort feedback whenever their metabolic rate or clothing level is changed. Because people shift between different activities and clothing levels relatively often, asking people to report changes in their metabolic rate or clothing level would seem impractical for a sustained deployment. Finally, these systems have only been deployed in offices, and for only one or two participants. Therefore, it is unclear whether such systems would work in the home.

In the following sections, we will first introduce our experimental system for investigating the potential of thermal comfort sensing in naturalistic settings using commodity wearable devices. Then, we will present our study method and the findings from our analyses.

EXPERIMENTAL COMFORT SENSING SYSTEM

To explore the potential of our approach, we developed an experimental system for collecting the required sensor data and user comfort feedback, as well as for performing the thermal comfort inference. Our experimental comfort sensing system contains five components: (1) HomeHub, (2) wearable sensors, (3) in-home sensors (4) a mobile ESM tool and (5) a web-based diary tool. After the introduction of these five components, the rationale for our sensor selection is then provided.

HomeHub: The HomeHub is the central component for collecting data from the wearable and the in-home sensors. A notebook PC is used for the hardware of the HomeHub (we used an ASUS X200MA), and we built the software of the HomeHub on top of the Lab of Things framework [4].

Wearable Sensors: We employ the Basis B1 [34], a wrist-worn fitness tracker, for collecting data useful for comfort prediction. Several of the data sources provided by the Basis B1 (i.e., activity level, skin temperature, and galvanic skin response) have not been employed in previous efforts to predict thermal comfort in real time, and so represent new potential sources of information. Basis B1 measures skin and near-body air temperatures, galvanic skin response (GSR), heart rate, step count, and estimated calorie consumption, calculated once per minute. We use GSR to approximate sweat level. We further use the per minute calorie consumption offered by Basis B1 and the weight of the individual to approximate a person's metabolic equivalent of task (MET). One study has shown that Basis B1 is able to estimate energy expenditure with 76.5% accuracy [21]. However, this study was conducted before the introduction of BodyIQ technology, which was available at the time of our study. In a forum post [35], the lead researcher of Basis claimed that with BodyIQ, the accuracy outperforms other major competitors, some of which claim around 90% accuracy.

Basis B1 synchronizes with the user's smartphone, which uploads the data to the Basis cloud storage. We automatically fetch the sensor data from Basis's web service every 15 minutes, which is the maximum rate attainable via Basis' API.

In-home Sensors: We use the Aeotec MultiSensor [36] to track room-level air temperature and humidity. The HomeHub samples each MultiSensor every 3 minutes via Z-Wave.

Mobile ESM Tool: To collect an individual's thermal comfort feedback given different environmental and bodily conditions, we employed an Android-based Experience Sampling Method (ESM) tool called Minuku [5]. We configured Minuku to send brief questionnaires to participants based on time and location (e.g., "at home"). Each questionnaire asked for the person's thermal sensation, comfort sensation, current activity, indoor location, clothing level, and brief notes that might help them recall the reasons for their sensation and comfort report when completing the end-of-day diary entry. We



Figure 1: (Left) The ESM interface; (Right) The Web-based Diary Tool.

developed our ESM strategy to minimize interruption, while collecting enough information related to people's thermal comfort to help them to recall more in-depth information later in the day. Our implementation of ESM allows us to collect not just sensor data and user comfort feedback, but also users' explanation of their perception, such as the reasons that might have caused their discomfort. This extra information allows us to investigate the feasibility of our technique in real households and identify potential challenges and opportunities for new approaches.

Web-based Diary Tool: At the end of each day, participants were asked to provide more information to explain the thermal reports they submitted throughout the day. A web-based diary tool (see Figure 1, right) was designed to facilitate these diaries. This tool displays an individual's thermal reports created throughout the day, as well as visualizations of the sensor streams that could help the person recall his indoor location and activity.

Sensor Selection: We chose the types of sensors based on the widely-used PMV model, which states that the primary factors influencing a person's thermal comfort include air temperature, radiant temperature, wind velocity, humidity, metabolic rate and clothing level. In addition to these six parameters, we also track skin temperature and sweat level, which are implicit in the PMV model: when calculating PMV, these two additional factors can be approximated by the six basic parameters using other models [8] developed in human thermal comfort, rather than directly measured.

While some factors such as radiant temperature and wind velocity are challenging to track precisely, especially in large buildings or outdoors, our setting of a single family home allows us to make a few reasonable simplifications. First, we assume the radiant temperature is the same as the air temperature, since rooms in houses are relatively small. Second, we designate a fixed wind velocity (0.2 m/s) based on the average winds speed of indoor ventilation (between 0.05 to 0.4 m/s) [27]. Finally, clothing level is obtained through participants' self-reports, collected with each ESM response. Four options, ranging from nearly naked to heavy clothing level, were provided, along with examples.

While PMV is known to have limited predictive power for individuals [24], we feel that the factors in the model are comprehensive and well-studied, thus serving as a good basis for our work. We will further discuss how we incorporate Adaptive Thermal Comfort and dynamic transitions by using a machine-learning based approach, along with a person's previous body and environmental states in the data pre-processing section.

STUDY METHOD

The goals of our study were to explore the feasibility of our approach and to investigate the potential challenges of inferring thermal comfort at home in naturalistic settings. We recruited 14 participants from 9 households in Michigan. 11 of the participants from 7 households were

able to complete the study. Three people dropped because of important family events that reduced the time they could stay at home. The recruitment was done through Craigslist, mailing lists and the snowball sampling method. We recruited participants who have an Android phone and who stay at home during waking hours for at least 5 hours a day on average. In addition, we recruited half of the households to have two participants in order to explore individual differences with respect to thermal comfort. Participants were compensated depending on the number of reports they provided during the ESM study. The amount of each person varied, with the average compensation being US\$44.90. For households with multiple participants, we compensated each of them with US\$10 extra. Table 1 provides an overview of participants and their household information.

Parti.	Gen-der	Valid Report s	House-hold	House Size (sqft)	Adult (Child)	Type of thermostat
P1	F	187	H1	1400-2000	4	Manual
P2	F	98	H1	1400-2000	4	Manual
P3	M	138	H2	< 800	2	Manual
P4	F	91	H2	< 800	2	Manual
P5	M	143	H3	< 800	2*	Manual
P6	M	131	H4	1400-2000	2*	Nest
P7	F	113	H5	800-1400	2	Prog.
P8	F	10	H6	800-1400	2 (1)	Manual
P9	M	2	H6	800-1400	2 (1)	Manual
P10	M	107	H7	800-1400	2	Prog.
P11	F	112	H7	800-1400	2	Prog.

Table 1: Participant Information (*: one member left in the middle of the study; Prog: Programmable Thermostats)

Our study took place between 30th August and 4th October, 2014, and consisted of a semi-structured initial interview, followed by a four-week sensor deployment and ESM study. We then conducted exit interviews. During the initial interview, we collected information on participants' daily schedules and how satisfied they were with regard to their thermal comfort in the different rooms in their houses. We used participants' daily schedules to help determine the rooms in which to place the multi-sensors. Additionally, we also collected information on how they used their thermostats and other comfort-related appliances such as fans and dehumidifiers, as well as their attitudes towards the trade-offs between saving energy and remaining comfortable.

After the interview, we installed 4 multi-sensors at different locations in each participating house and provided each participant with a Basis B1. Participants were asked to wear the Basis B1 whenever they were awake and at home. In addition, the ESM tool was installed on participants' phones. To ensure every participant understood how to use the tools, we guided each of them to create one thermal comfort report using the ESM tool and to provide a detailed comment using the web-based diary tool.

Immediately after the sensors were deployed, the ESM study was started. During the four-week ESM study period, participants self-reported their thermal sensation, comfort sensation, location within the home, and their activity information. Sensor data detected by the multi-sensors and Basis B1 were stored in our database. Participants were not exposed to any predictions we made, and no prediction results were used to make any changes to their home HVAC system during the study.

For self-reports, we used both the 7-level thermal sensation index introduced in PMV, and also a standard 4-level comfort sensation index used in several thermal comfort studies (e.g., [12]). Using two indices allowed us to resolve the ambiguity of the labels in the thermal sensation index: for many people, the “cool” label might actually represent a comfortable and preferred feeling.

Our ESM tool prompted each participant approximately every 30 minutes when s/he was at home and during a pre-specified awake time window. Participants could ignore individual prompts, but they were expected to answer at least 6 reports per day. Participants could also actively report whenever they like, although we encouraged them to only initiate a report when they felt uncomfortable. At the end of the study, we found out that some participants had deactivated the GPS tracking of the ESM tool because it drained too much battery from the phone. Some of them therefore only initiated thermal reports, rather than responding to prompts. It’s possible that participants created different number of reports related to uncomfortable situations because of this issue, but we expect no effect on the validity of their answers.

A 30-minute exit interview was conducted following the ESM study. Before the exit interview, we calculated the PMV index of each thermal report a participant created using Fanger’s approach [8]. We used Fanger’s PMV rather than our own model’s prediction because our model was not finalized until after the exit interviews were conducted. While PMV is inaccurate for predicting individual’s thermal comfort, the calculations allowed us to investigate situations in which there were large differences between the PMV and participants’ reports. During the exit interviews, we asked participants to recall what happened at the moment of a report, and the potential reasons for the wrong PMV prediction. To facilitate recall, we asked the participant to review their diary entries and comments.

DATA PROCESSING

From the 11 participants who completed the 4-week study, we collected 1,431 thermal comfort reports (details are provided in Table 1). However, only 1,132 thermal comfort reports were considered to be complete— i.e., containing all the associated information, including thermal sensation, comfort sensation, indoor location and activity, as well as data detected by sensors, including air temperature, humidity, skin temperature, near body air temperature, GSR, and metabolic rate at the moment of report. In

addition, only 9 of the 11 participants created more than 90 valid reports; the other two participants were not able to properly maintain the synchronization between their Basis B1 devices and smartphones due to software configuration issues, thus rendering many of their reports useless for training our models.

The “comfortable” sensation dominated the reports: within the 1,132 reports, 50.6% of them are labeled as “neutral” on the thermal sensation index, and 76% of them are labeled as “comfortable” on the comfort sensation index. This may be due to the fact that the weather in the study area was unusually mild during the study period, with average temperatures at 16°C, (max 33°C; min 0°C). However, there was one week (13 Sep. to 20 Sep.) that the temperature dropped below normal. The average temperature in that week was 11°C (max 26°C; min 0°C).

The percentage difference between thermal and comfort sensation also confirms the usefulness of having two indices. For example, some people interpreted “slightly cool” or “cool” as a comfortable and preferred temperature.

Feature extraction

For each of the reports, we further obtained features related to an individual’s previous state, including the metabolic rate, skin temperature, near body air temperature, and sweat level (GSR) 30 and 10 minutes before the report. For each of the features, we smoothed the data by taking an average over a five-minute window. These features were inspired by [12], which demonstrates that, in addition to the factors modeled by steady-state thermal comfort models like PMV, dynamic transitions between warm and cold environments also affect people’s thermal sensation. However, one problem we faced when extracting these features was that some reports had no data from the previous state. For example, when a participant had just woken up in the morning and worn his or her wearable sensors for just a few minutes, there would not be any sensor data available for the previous 15 minutes. To handle this problem, we filled in missing values with the average sensor reading for that participant. Finally, we extracted the temperature and humidity data by referencing the room information participants provided in the thermal report. We smoothed the room temperature and humidity data by taking the average over a 30-minute window. If the room did not contain a multi-sensor, we then averaged the sensor readings from the two adjacent rooms in which we had multi-sensors deployed.

We further categorized the activity provided by the participants using their free text descriptions reported through the diary tool. As we already captured activity-level through the measurement of metabolic equivalent of task, activity type information was collected primarily to investigate the psychological influence of activities. We generated our activity categories by combining the American Time Use Survey Activity Lexicon [29] with our own heuristics of which activities might effect an

individual’s thermal comfort sensation, such as “cooking” (see Table 2).

relaxing and leisure, working at home, eating, housekeeping, cooking, waking up, grooming, arriving home, socializing, exercising, sleeping, feeling sick, other

Table 2: Activity Categories

Thermal comfort index

As no standard scale exists in the previous literature for training machine-learning based comfort models, we created a 5-level thermal comfort index for this work. Our 5-level index combines the 4-level comfort sensation index suggested by Gagge et al. [12] and the 7-level thermal sensation index used in PMV. Table 3 illustrates the mapping from the thermal sensation and comfort sensation indices into our 5-level index. This mapping resulted in a scale ranging from, *Uncomfortably-Cold* (UC-Cold) to *Uncomfortably-Warm* (UC-Warm) at the extremes, with *Slightly-Uncomfortably-Cold* (S-Cold), *Comfortable* (COM), *Slightly-Uncomfortably-Warm* (S-Warm), as intermediate levels.

The intuition behind this transformation is that thermal sensation alone cannot represent an individual’s complete thermal comfort assessment. For example, some participants actually interpreted “cold” as a comfortable feeling at times, and most of them enjoyed being “slightly cool”. Comfort sensation, on the other hand, only provides the intensity of discomfort, but not the direction of discomfort. However, from our exit interviews we found out that participants tended to have a common interpretation of the comfort sensation index. They interpreted uncomfortable as situations in which they would definitely take action to change the thermal environment, while slightly uncomfortable meant they were uncomfortable but did not experience an urgent need to take action. Therefore, by combining these two indices, we could obtain an index that captures both the intensity of discomfort and the warm-cold direction of the discomfort.

Besides the advantage of integrating the intensity and the warm-cold direction of discomfort, we believe our index is an improvement upon the 3-level scale used by Feldmeier and Paradiso, and the 7-level index used by Nouvel and Alessi [25]. First, compared to the 3-level scale, we allowed participants to report discomfort that was relatively tolerable (e.g., S-Cold, S-Warm), thus taking into account people’s tolerance to warm and cold situations, as suggested by Adaptive Thermal Comfort. Second,

compared to the 7-level scale, we maintain a smaller set of thermal classes, thus facilitating the training of machine-learning based models.

ANALYSES & FINDINGS

Two analyses were conducted in this research. In the first analysis, we developed our comfort model using a machine-learning approach. We compared the accuracy of different feature sets, as well as the accuracy of our approach compared to other baseline models inspired by previous research [8,9]. In the second analysis, we explored the challenging situations for inferring thermal comfort at home in naturalistic settings by examining cases where our model failed to make an accurate prediction.

Analysis 1: Accuracy of Comfort-Aware Model

To develop our comfort model, one step is required before the training to handle our imbalanced dataset. In our dataset, 76% of the data are labeled as COM on our 5-level thermal comfort index as shown in Table 3. The remaining classes each cover between 4-8% of the data. To prevent training a classifier that always predicts comfortable, we adjusted the weight of each datum given its class to create equally balanced class labels (following [18]).

Baseline models

Random Forest and Gaussian-kernel SVM were selected to train our models, as these two methods are most likely to perform the best on various data sets [10]. Three models were chosen for comparison: the first is a ZeroR model that always predicts COM; the second is a decision tree model with the estimated PMV index as the only feature, chosen for testing the efficacy of PMV in predicting individual’s thermal comfort as inspired by [25]; the third model is a Gaussian-kernel SVM with humidity and near body air temperature (SVM-H/T) as the only two features. This final model is inspired by Feldmeier and Paradiso’s work [9], which is the only other work that uses wearable sensors for comfort sensing in naturalistic settings. In their work a simpler model, Fisher’s discriminant was adopted due to the requirement of monotonicity. This constraint was imposed because the end-goal of their system was to automatically control a HVAC system—monotonicity was required in their model to prevent hysteresis. Removing the monotonicity constraint allows us to employ more sophisticated and presumably more accurate models. We abandoned monotonicity because automatic temperature adjustment is not the only application that could be enabled by comfort sensing. Other research that is more interested in measuring the comfort quality of homes does not require

	Comfort Sensation	Thermal Sensation	Thermal Comfort Index	#reports	%reports
1	Uncomfortable/Very uncomfortable	Cooler than neutral	Uncomfortably cold (UC-Cold)	76	6.7%
2	Slightly uncomfortable	Cooler than neutral	Slightly uncomfortably cold (S-Cold)	43	3.8%
3	Comfortable		Comfortable (COM)	862	76.1%
4	Slightly uncomfortable	Warmer than neutral	Slightly uncomfortably warm (S-Warm)	61	5.4%
5	Uncomfortable/Very uncomfortable	Warmer than neutral	Uncomfortably warm (UC-Warm)	90	8%

Table 3: Frequency of reports for each level of our combined 5-level thermal comfort index.

	Machine Learning Models (RF: Random Forest)								Baselines		
	RF +BASE	RF +NO-CLO	RF +ACT	RF +BASIS	SVM +BASE	SVM +NO-CLO	SVM +ACT	SVM +BASIS	ZeroR	SVM- H/T	DT-PMV
MAE	0.76(0.15)	0.80(0.17)	0.75(0.15)	0.96(0.20)	0.73(0.12)*	0.77(0.12)	0.75(0.13)	0.91(0.11)	1.2(0.0)	1.20(0.03)	1.10(0.19)
MSE	1.20(0.41)	1.31(0.48)	1.24(0.45)	1.84(0.73)	0.98(0.30)*	1.08(0.30)	1.07(0.33)	1.33(0.28)	2.0(0.0)	2.00(0.09)	1.92(0.78)

Table 4: Mean Absolute Error (MAE) and Mean Squared Error (MSE) for Random Forest (RF) and Support Vector Machine (SVM) models with different feature sets. Three baselines are provided for comparison, as described in the text. The best performing model (SVM+BASE) is highlighted with a *.

an automatic agent to take actions. Even for applications like intelligent thermostats, mixed-initiative [15] approaches might be taken. Such approaches could benefit from a more accurate prediction while remaining tolerant to non-monotonicity.

Since our class labels are ordinal in nature—there is a natural ordering of the classes from uncomfortably warm to uncomfortably cold—we therefore trained our multi-class classifiers on top of a simple ordinal classifier developed by Frank and Hall [11]. This ordinal classifier allows us to transform any multi-class classifier into a classifier for ordinal variables. Note that the class weight adjustment is applied on the 5-level thermal comfort index, rather than the multiple binary classification problems generated by the simple ordinal classifier.

In addition to the prediction models, four different feature sets were investigated: (1) BASE; (2) BASE without clothing level (NO-CLO); (3) BASE with activity information (ACT); and (4) Basis B1 feature set (BASIS). The BASE feature set contains all the major parameters believed to influence a person’s thermal comfort. These features are provided by the multi-sensor, Basis B1, and participant self-reports. The NO-CLO contains all the features in BASE except for the self-reported clothing level. This was selected to test if using only sensor data is adequate for comfort prediction. ACT feature set contains all the features in BASE plus activity labels (again, see Table 2). ACT was included as we conjectured that the type of activity a person is conducting might change their expectation about the ideal temperature. Finally, the BASIS feature set was selected to test if using only wearable sensors would be adequate for comfort inference.

Ten-fold cross-validation was performed on the whole dataset ten times, yielding 100 trials of model accuracy. We use Mean Absolute Error (MAE) and Mean Squared Error (MSE) as our evaluation metrics. While MAE is a standard way to evaluate ordinal variable classifiers, we included MSE to further penalize classifiers with large error, as we did not want our classifiers to predict cool while a person is feeling warm. Finally, the Wilcoxon signed-rank test is used to compare model accuracy.

Result: SVM with BASE performs the best

Table 4 shows the prediction results for each model. We found that SVM and Random Forest with BASE feature set perform very similarly to each other with respect to MAE, although SVM has significantly smaller MSE. Both of these

models outperform the three baselines ($p < 0.01$). Compared to the baselines, SVM with BASE reduces the MSE by 51% compared to ZeroR, 49% compared to DT-PMV, and 51% compared to SVM-H/T. SVM-H/T’s prediction result is very close to ZeroR, indicating that having only near body air temperature and humidity is not sufficient for predicting thermal-sensation at home, perhaps because people exhibit adaptive behaviors and conduct various activities. Overall, the relatively small MAE and MSE of our model indicates that it is able to control the prediction error to within one ordinal class distance most of the time—if a person is feeling comfortable, then most of the time the prediction is bounded between slightly uncomfortably warm and slightly uncomfortably cold.

In addition, we also found that adding activity labels (e.g., “cooking,” “exercising”) directly as features did not improve the model. This could be because there are only a few thermal reports for each type of activity. Furthermore, particular activities may have different effects on different individuals. More advanced graphical modeling might be needed to utilize activity information for improving the models.

Although the SVM model with BASE feature set performs significantly better ($p < 0.01$) than the SVM model without clothing level information (NO-CLO), their errors are very close to each other—the MSE for BASE is 0.98 (SD=0.3), whereas for NO-CLO it is 1.08 (SD=0.3). This similar performance may have resulted from relatively low diversity in clothing levels at home, or from participants’ inaccurate estimation of their clothing level. We will discuss the latter cause in the second analysis.

Finally, we found that by using only features provided by Basis B1, the MSE is 1.35 times higher than SVM with BASE. It is surprising that even though Basis B1 is able to detect near body air temperature, having no room temperature and humidity information increases the error considerably.

Analysis 2: Challenging Situations

In our second analysis, we further investigated challenging situations for comfort inference. Specifically, we were interested in knowing *the situations and factors behind inaccurate predictions*.

To conduct such an analysis, we selected the best model obtained from our previous study, that is, SVM with the BASE feature set. To generate predictions for all the

reports, we conducted two rounds of model training. First, we trained the SVM model by using 50% of the thermal comfort reports (training set), and then performed the prediction on the rest of the reports (testing set). We then trained another model based on the testing set and applied it to generate the predictions for the training set. Table 5 shows the resulting confusion matrix. Note that with only 50% of the data used for training, it is likely that our model would generate more errors than normal (e.g., if we had trained on 90% of the data as in our first Analysis), however, as our goal was to qualitatively analyze the potential situations for inaccurate prediction, we feel that the higher number of error cases is acceptable, albeit slightly less efficient.

To investigate the challenging situations for comfort inference and identify potential solutions, we examined information from participants' comments in their exit interviews, activity information in comfort reports, and raw sensor data.

Challenging situations for prediction

We identified 79 error cases where the prediction was more than two classes away from the true label, such as when the model generated UC-Warm while the true label was COM, S-Cold, or UC-Cold. The cells highlighted in Table 5 indicate errors we investigated. Our qualitative analysis of the 79 errors allowed us to identify six situations that lead to inaccurate prediction: (1) *short-term effect or local heat source*; (2) *dynamic transitions*; (3) *extra cover or wind effect*; (4) *light-weight exercising and housekeeping*; (5) *problems with data collection tool*; and (6) *individual differences*.

Short-term effect or local heat source refers to the situation where the indoor temperature is different from the current thermal sensation of the participant. For example, it might be because the participant was drinking a cold or warm beverage, close to a hot stove, or had just taken a shower. Under this situation, a participant's near body air temperature and skin temperature might be within a comfortable zone, while the ambient temperature was low or high. For example, P3 commented "*I felt warmer because I was reading the news and checking email with my laptop on my lap. Even though the room was still cool from earlier, the laptop made me feel warm and kept me comfortable.*" He reported COM while the air temperature was 19.8 °C and his skin temperature was only 28 °C, thus making the prediction to be S-Cold.

The second source of the prediction error is a result of *dynamic transitions* between cold and warm situations. We did not have enough thermal reports related to such transitions to train the model effectively for these cases. Furthermore, the sampling rate supported by Basis B1 (1 minute) and our smoothing window (5 minutes) cannot capture quick transitions well. For example, one type of transition occurs when participants enter their warmer homes from colder outdoor environment. Since we

conducted the study during late summer/early fall, there were only a few days that the outdoor temperature was low, thus such events are not well-represented in our training dataset. Another type of transition occurs when people move from a warm bed into a colder room. In this situation, participants usually had a high skin temperature from their cozy bed, while the room temperature at the moment captured by the multi-sensor was relatively low. This inconsistency confuses the model. For example, P4 commented that "the room was [at] a comfortable temperature" with "waking up" as her activity. The room temperature at the time was only 18.9 °C, while her skin temperature 15 minutes before the report was 31 °C, which is an unusually high skin temperature. Therefore the model predicted she felt UC-Cold, while she actually felt COM.

A third cause of prediction error is when a person might have *extra cover* that was not included in their self-reported clothing level, or might have been affected by un-captured *wind effect*. For example, P11 commented that "*The puppy was in my lap, which warmed me up*" in one case, and noted "*Was still in bed under heavy blankets*" in another report. Additionally, certain locations where participants spent time could have an effect, as an upholstered sofa or carpet can effectively increase a person's clothing level [22]. Additionally, our simplified assumption about constant indoor wind velocity occasionally led to prediction problems. For example, P1 commented that she had her fan on while her skin temperature was 33.7 °C and the air temperature was 27.8 °C. The model predicted she felt UC-Warm due to the high temperature while she reported comfortable due to the additional wind effect.

The fourth type of error occurs when an individual is doing *light exercise or housework* during cool days, including activities like moving objects around the house, performing small household repairs, or simply walking around the house. In such cases, an individual might have higher clothing level and slightly higher approximated metabolic equivalent while the other metrics are still low, resulting in a prediction that is cooler than what the individual actually perceives (e.g., predicting COM when the participant reports UC-Warm). Here, it could be that the high clothing level amplifies the effect of the slightly heightened metabolic equivalent. Although the PMV index, a feature included in our machine-learning based model, is supposed to capture such interaction between clothing and

Prediction \ True	UC-COLD	S-COLD	COM	S-WARM	UC-WARM
UC-COLD	8	17	0	0	0
S-COLD	7	39	15	8	0
COM	22	186	410	271	10
S-WARM	3	8	17	64	7
UC-WARM	2	1	2	26	9

Table 5: Confusion Matrix (Cells highlighted in gray represent the cases used for our fault analysis.)

metabolism, the misprediction may be generated either because of the inaccurate estimation of the clothing level from the participants, or because of the low skin and air temperatures.

There were *problems with the data collection tool* and data handling. For example, there were several instances that participants did not wear their Basis B1 or had just started to wear it, thus the inaccurate prediction is due to our naive way of handling missing values (i.e., fitting the mean value obtained from all the reports of a participant). In addition, participants may have interpreted the rating scales in different ways at different times. For example, P11 reported her comfort sensation as “slightly uncomfortable” due to the fact that she felt warm. However, she also labeled thermal sensation as “slightly cool” due to the fact that she had just drunk a bottle of cold water. The correct labels to train our model should be “slightly uncomfortable” for comfort sensation and “slightly warm” for thermal sensation, as the “uncomfortable” feeling was caused by the warm temperature. Therefore, the ambiguity in the interpretation of the data collection instructions caused the inaccurate prediction.

Individual differences include the tolerance of heat, cold, sweat and humid environments and the different interpretation of the comfort and thermal sensation scales. For example, P10 reported “comfortable” with comment “*At the desk, my hands were getting cold. I am used to my hands getting cold, though, so it wasn't uncomfortable.*” As his estimated skin temperature was 26.7 °C and the room temperature was 16.5 °C, which is below the typical comfort zone for most people, the model incorrectly predicted UC-Cold. In another case, P1 reported that she felt comfortable after she had exercised at home. The model wrongly predicted her comfort level as UC-Warm as her skin temperature and the air temperature were high—32.6 °C and 28.56 °C respectively—as was her metabolic equivalent. Finally, another source of individual difference is sickness. P11 was sick for more than two weeks during the study, thus her perception and body conditions were different than normal. In one thermal report with mild air temperature at 21 °C, she noted “[*I felt the body] temperature went up because I was feeling sick due to a bad headache.*”

DISCUSSION

Our work shows that sensing thermal comfort in the wild is promising, but challenges remain. Widely used comfort models and prior techniques that aim at inferring comfort in naturalistic settings are insufficient for inferring individual’s thermal comfort at home. In fact, our result shows that these previous techniques—mostly designed for large buildings and offices—did not perform better than a naive ZeroR baseline. The dynamic nature of home activities and people’s adaptive behaviors make comfort inference much more difficult in the home than in climate chambers and in offices.

In an effort to infer thermal comfort at home that allows domestic residents to live naturally, we proposed a new technique that uses wearable fitness trackers and in-home sensors to capture various factors that are underexplored in previous in-situ sensing research. By incorporating metabolic equivalent, sweat, and skin temperature sensors, we saw that our technique can reduce prediction error by 50% compared to the baseline models. In addition, our result shows the benefits of having both wearable and in-home sensors. Having only one type of sensor is not sufficient to predict comfort accurately. While wearable sensors are useful to obtain continuous activity-level, sweat and temperature measurements, they are ineffective at capturing the larger thermal comfort picture of the room. Moreover, it would be challenging for the wearable device to understand a transition such as getting out from under a warm comforter in a cold winter while one wakes up. On the other hand, having only in-home sensors is insufficient for capturing discomfort caused by high activity-level and transitions, such as coming back home from exercising.

While we are interested in understanding more specifically which temperature sensors (i.e., skin, near-body and room temperatures) are most critical, our analysis is inconclusive. However, our work is the first step towards understanding the benefit of utilizing these various sensors in reaching better prediction accuracy. We would argue, however, that future work on in-situ sensing should have all three types of sensors, as the combination of the three could potentially help identify various situations. For example, whether users are close to any local heat source could potentially be identified by the combination of near-body and room temperature. Clothing level might potentially be inferred by the difference between skin and near-body temperature, as explored by SPOT [13].

Implications for Improving Thermal Comfort Inference

While there are still some challenging situations for inference, a closer look at the reports reveals that many of these situations could be easily resolved, and some of them could potentially be tolerated depending on the intended application.

First, several of the errors resulted from the way we handle missing data and the way we collect thermal comfort feedback. Note that while we did throw out a report if the user was not wearing the Basis B1 at the time, missing data was filled in for reports that are associated with users just starting to wear Basis B1 (e.g., when they woke up in the morning), as we didn’t want to overlook these important moments. When deploying an application based on such a sensing technique, we could easily remove such errors by checking whether the users have been wearing their wearable tracker for a certain period of time. Alternatively, sensor data from previous days that share similar activity patterns could be used to fill in the missing sensor data. For the problem with inconsistent labeling, we could revise the ESM interface to insure the warm-cold direction of the

comfort sensation is the same as thermal sensation. For example, we could ask a clarifying question such as, “Do you want the temperature to be cooler or warmer?”

The errors related to rare or quick transitions could potentially be resolved by increasing the sampling rate of wearable sensors, along with better handling of time series data. Currently the Basis B1 has a sampling rate of one time per minute—note that this is the sampling rate accessible by 3rd party developers rather than the true sampling rate of the sensors—which is due to its focus on fitness related applications. On top of this relatively low sampling rate, our five-minute smoothing window makes quick transitions harder to observe. To better identify transitions between cold and warm environments or the presence of additional heat sources, the variance of the sensor reading along with the mean we use in this study could be used. In addition, activity recognition techniques could be applied to identify some transitions such as home/away status [6]. Furthermore, if comfort sensing techniques such as ours become useful, it is possible that wearable sensor manufacturers would broaden their services to include comfort sensing, thus tailoring their device to serve such a purpose—e.g., by allowing 3rd party developers to access sensor data in real-time. Finally, a larger deployment with more data to train the models might help to model relatively rare transitions.

Two other challenges we identified include extra “clothing” and wind effects that were not captured by the sensors or by users’ self-reports. Through additional insights gained from the exit interviews, we found that whether or not people have extra covering is sometimes related to their locations in the room. For example, P11 usually had her blanket on when she sat on the sofa in her TV room. This suggests it might be possible to infer people’s extra clothing level via part-of-the-room indoor positioning, although it would require additional training data from each of the individuals. However, we also found that when such extra clothing exists, the difference between the near body temperature detected by Basis B1 and the air temperature detected by the multi-sensor is larger. We could potentially use this information to help us identify if extra clothing exists.

Finally, individual differences might be resolved by taking a personalization approach. However, the standard ways of training personalized models—such as generating a model for each individual based on one’s own data—may not be feasible for our approach since the number of reports required to cover all five thermal comfort labels at a sufficient level may be excessive. While such isolated models may not work for our approach, “Groupization” could be a better solution as it seeks to build personalized models for an individual by using data from other similar people. This approach has been used in personalized information retrieval [28] and activity recognition [19] systems, and is particularly suitable for our application

scenario where the training data provided by each individual is minimal. Groupization reduces the amount of feedback an individual needs to provide in order to train the model, while improving upon the prediction accuracy based on a population-based model.

Limitations

There are a few limitations of our study. First, study was conducted at the end of the summer/beginning of fall when the temperature was relatively mild. Further deployment is needed to inspect the feasibility of our approach in more extreme weather, as seasonal difference might be an important factor to include in the model. In addition to the relatively short deployment conducted in the summer, our thermal comfort reports are only provided by 9 participants, a small population. More research is needed to consider the role of individual difference and to validate if this approach could work for a larger group of people.

Although extensive collection of thermal comfort feedback from each individual is required in our experimental system, we envision the future system in which online inference could rely on a small group of users who provide feedback for training a population-based model. If derivation from the population-based model is found, the system could then prompt a user for more comfort feedback in order to create a personalized model or reassign them to a more appropriate group.

CONCLUSION

In this paper, we present a new technique for inferring people’s thermal comfort at home under naturalistic settings using in-home sensors as well as off-the-shelf wearable devices equipped with sensors that allow estimations of metabolic equivalent, sweat and skin temperature. A sensor deployment and experience sampling study was conducted in 7 households with 11 participants to validate the potential of such approach. Our study results reveal the advantages of this approach, challenging situations for prediction, and potential directions for improve in-situ comfort sensing at home.

ACKNOWLEDGMENTS

This work was funded in part by NSF award IIS-1149601 and a Microsoft Research grant. We would like to thank our participants, as well as Kevyn Collins-Thompson, Tawanna Dillahunt, members of the Interaction Ecologies group, and our anonymous reviewers for helpful comments.

REFERENCES

1. Nick Baker and Mark Standeven. 1996. Thermal comfort for free-running buildings. *Energy and Buildings* 23, 3, 175–182. [http://doi.org/10.1016/0378-7788\(95\)00942-6](http://doi.org/10.1016/0378-7788(95)00942-6)
2. Arash Beizaee and Steven K. Firth. 2011. A comparison of calculated and subjective thermal comfort sensation in home and office environment. Retrieved June 7, 2015 from

- <https://dspace.lboro.ac.uk/dspace-jspui/handle/2134/14021>
3. Gail S. Brager and Richard J. de Dear. 1998. Thermal adaptation in the built environment: a literature review. *Energy and Buildings* 27, 1, 83–96. [http://doi.org/10.1016/S0378-7788\(97\)00053-4](http://doi.org/10.1016/S0378-7788(97)00053-4)
 4. A. J. Bernheim Brush, Jaeyeon Jung, Ratul Mahajan, and James Scott. 2012. HomeLab: shared infrastructure for home technology field studies. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, ACM, 1108–1113. <http://doi.org/10.1145/2370216.2370450>
 5. Yung-Ju Chang, Gaurav Paruthi, and Mark Newman. 2015. A Field Study of Comparing Three Crowd-Based Annotation Approaches to Collecting Transportation Data.
 6. Liming Chen and Ismail Khalil. 2011. Activity Recognition: Approaches, Practices and Trends. In *Activity Recognition in Pervasive Intelligent Environments*, Liming Chen, Chris D. Nugent, Jit Biswas and Jesse Hoey (eds.). Atlantis Press, 1–31. Retrieved April 30, 2015 from http://link.springer.com/chapter/10.2991/978-94-91216-05-3_1
 7. Adrian K. Clear, Janine Morley, Mike Hazas, Adrian Friday, and Oliver Bates. 2013. Understanding Adaptive Thermal Comfort: New Directions for UbiComp. *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, 113–122. <http://doi.org/10.1145/2493432.2493451>
 8. P. O. FANGER. 1970. Thermal comfort. Analysis and applications in environmental engineering. 244 pp.
 9. Mark Feldmeier and Joseph A. Paradiso. 2010. Personalized HVAC control system. *Internet of Things (IOT), 2010*, IEEE, 1–8. Retrieved March 22, 2014 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5678444
 10. Manuel Fernández-Delgado, Eva Cernadas, Senén Barro, and Dinani Amorim. 2014. Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research* 15, 1, 3133–3181.
 11. Eibe Frank and Mark Hall. 2001. A Simple Approach to Ordinal Classification. *Proceedings of the 12th European Conference on Machine Learning*, Springer-Verlag, 145–156. Retrieved December 16, 2014 from <http://dl.acm.org/citation.cfm?id=645328.649997>
 12. A. P. Gagge, J. A. J. Stolwijk, and J. D. Hardy. 1967. Comfort and thermal sensations and associated physiological responses at various ambient temperatures. *Environmental Research* 1, 1, 1–20. [http://doi.org/10.1016/0013-9351\(67\)90002-3](http://doi.org/10.1016/0013-9351(67)90002-3)
 13. Peter Xiang Gao and S. Keshav. 2013. SPOT: A Smart Personalized Office Thermal Control System. *Proceedings of the Fourth International Conference on Future Energy Systems*, ACM, 237–246. <http://doi.org/10.1145/2487166.2487193>
 14. Joel M. Hektner, Jennifer A. Schmidt, and Mihaly Csikszentmihalyi. 2007. *Experience Sampling Method: Measuring the Quality of Everyday Life*. SAGE.
 15. Eric Horvitz. 1999. Principles of Mixed-initiative User Interfaces. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 159–166. <http://doi.org/10.1145/302979.303030>
 16. Michael A Humphreys and J Fergus Nicol. 2002. The validity of ISO-PMV for predicting comfort votes in every-day thermal environments. *Energy and Buildings* 34, 6, 667–684. [http://doi.org/10.1016/S0378-7788\(02\)00018-X](http://doi.org/10.1016/S0378-7788(02)00018-X)
 17. Byron W. Jones. 2002. Capabilities and limitations of thermal models for use in thermal comfort standards. *Energy and Buildings* 34, 6, 653–659. [http://doi.org/10.1016/S0378-7788\(02\)00016-6](http://doi.org/10.1016/S0378-7788(02)00016-6)
 18. Sotiris Kotsiantis, Dimitris Kanellopoulos, and Panayiotis Pintelas. *Handling imbalanced datasets: A review*. GESTS International Transactions on Computer Science and Engineering. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.96.9248>
 19. Nicholas D. Lane, Ye Xu, Hong Lu, et al. 2014. Community Similarity Networks. *Personal and Ubiquitous Computing* 18, 2, 355–368. <http://doi.org/10.1007/s00779-013-0655-1>
 20. Judith A Leech, William C Nelson, Richard T Burnett, Shawn Aaron, and Mark E Raizenne. 2002. It’s about time: A comparison of Canadian and American time–activity patterns. *Journal of Exposure Analysis and Environmental Epidemiology* 12, 6, 427–432. <http://doi.org/10.1038/sj.jea.7500244>
 21. Jung-Min Lee, Youngwon Kim, and Gregory J. Welk. 2014. Validity of consumer-based physical activity monitors. *Medicine and Science in Sports and Exercise* 46, 9, 1840–1848. <http://doi.org/10.1249/MSS.0000000000000287>
 22. E. A. McCullough, B. W. Olesen, and S. Hong. Thermal insulation provided by chairs. *ASHRAE Transactions-American Society of Heating Refrigerating Airconditioning Engin* 100(1), 795-804.
 23. Ahmed Chérif Megri, Issam El Naqa, and Fariborz Haghghat. 2005. A Learning Machine Approach for Predicting Thermal Comfort Indices. *International*

- Journal of Ventilation* 3, 4, 363–376.
<http://doi.org/10.5555/ijov.2005.3.4.363>
24. J. F. Nicol and M. A. Humphreys. 2002. Adaptive thermal comfort and sustainable thermal standards for buildings. *Energy and Buildings* 34, 6, 563–572.
[http://doi.org/10.1016/S0378-7788\(02\)00006-3](http://doi.org/10.1016/S0378-7788(02)00006-3)
 25. Romain Nouvel and Franck Alessi. 2012. A novel personalized thermal comfort control, responding to user sensation feedbacks. *Building Simulation* 5, 3, 191–202. <http://doi.org/10.1007/s12273-012-0076-5>
 26. Elizabeth Shove. 2003. *Comfort, Cleanliness and Convenience: The Social Organization of Normality*. Berg Publishers, New York, NY, USA. Retrieved April 15, 2014 from <http://site.ebrary.com/lib/alltitles/docDetail.action?docID=10060563>
 27. J. Thorshauge. 1982. Air-velocity fluctuations in the occupied zone of ventilated spaces. *ASHRAE transactions* 88, 753–764.
 28. Thanh Tien Vu, Dawei Song, Alistair Willis, Son Ngoc Tran, and Jingfei Li. 2014. Improving search personalisation with dynamic group formation. ACM Press, 951–954.
<http://doi.org/10.1145/2600428.2609482>
 29. 2012. ATUS Activity Coding Lexicons. Retrieved February 10, 2015 from <http://www.bls.gov/tus/lexicons.htm>
 30. 2015. Kinect. *Wikipedia, the free encyclopedia*. Retrieved July 17, 2015 from <https://en.wikipedia.org/w/index.php?title=Kinect&oldid=670483329>
 31. Fitbit. Retrieved from <https://www.fitbit.com/>
 32. SmartThings. *SmartThings*. Retrieved June 8, 2015 from <http://www.smartthings.com/>
 33. U. S. Census Bureau, American Community Survey, 5-Year Estimates. Retrieved from http://quickfacts.census.gov/qfd/meta/long_HSD310213.htm
 34. Basis B1. Retrieved from <http://www.mybasis.com/>
 35. Forum post that discussed the accuracy of Basis B1 on calories burned (energy expenditure). Retrieved from <https://support.mybasis.com/hc/communities/public/questions/201390174-Accuracy-of-the-calories-burned>
 36. Aeotec MultiSensor. *Aeotec*. Retrieved February 10, 2015 from <http://aeotec.com/z-wave-sensor>