

ThermalSense: Determining Dynamic Thermal Comfort Preferences using Thermographic Imaging

Juhi Ranjan*[†] and James Scott*

* Microsoft Research, Cambridge, UK

[†] University of Virginia, Charlottesville, USA

juhi@virginia.edu, jws@microsoft.com

ABSTRACT

We present *ThermalSense*, a method for dynamically detecting and predicting thermal comfort by using thermographic imaging to look for the physiological markers of vasodilation or vasoconstriction. We describe how *ThermalSense* can be used to infer how to control heating and cooling systems and reduce energy use while maintaining comfort.

We evaluate *ThermalSense* using a study involving thirty individuals over five weeks in an office building. Our study shows that, on around 40% of occasions, the HVAC system could have expended less energy to achieve comfort. It further demonstrates that thermographic imaging can be used to infer whether heating or cooling must be activated to maintain comfort, with an accuracy of 94-95%.

ACM Classification Keywords

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

Author Keywords

Thermal comfort; thermographic imaging; HVAC; energy

INTRODUCTION

Thermal regulation of buildings is important for safety and comfort, and at the same time represents a significant proportion of worldwide energy use. For example, in residential buildings in the US, it accounts for 48% of total energy use [1], or 61% in Canada and the UK [3, 29]. In nearly all thermal regulation systems, human comfort is represented in the control loop by using a “set point” air temperature, which the system tries to match by heating or cooling as necessary.

However, it has long been known that human thermal comfort relies on many other factors. Fanger’s seminal work on the Predicted Mean Vote (PMV) metric [8] utilized six quantities: air temperature, relative humidity, air velocity, mean radiant temperature, clothing level and metabolism level. Radiant




Figure 1: Thermographic images of the same person when (a) they want it to be warmer, and (b) they want it to be cooler, showing detectable differences. Mean temperatures: (a) palm: 30.9°C, forehead: 34.8°C, cheek: 32.5°C. (b) palm: 34.3°C, forehead: 35.1°C, cheek: 34.9°C

heat can, for example, cause people to feel warmer on a sunny day when the walls and furniture warm up and re-radiate heat, even if the air temperature is the same. Metabolism levels can also cause people’s thermal sensation to differ [14].

It is well-known that people are comfortable in a range of temperatures at any given time. The commonly-used ANSI/ASHRAE-55 standard allows for air temperature differences of up to 3°C in office spaces while maintaining occupant comfort [6]. In our own study (reported later), the

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 '16, September 12-16, 2016, Heidelberg, Germany Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4461-6/16/09\$15.00

DOI: <http://dx.doi.org/10.1145/2971648.2971659>

median range of air temperatures reported as “comfortable” was 1.4 °C.

The motivation for this work is to find a new method by which heating and cooling systems can be controlled, which minimizes energy use while maintaining comfort. To do this, we need to sense whether the occupant would be comfortable at a less energy-intensive temperature (i.e., colder if the heating is on, or warmer if cooling is on), at any given time.

To achieve this, we decided to explore using thermographic imaging of people to look for physiological signs of feeling warm or cool. The human body uses cutaneous vasodilation and vasoconstriction to help regulate internal temperature, bringing more blood to the skin surface to cool the body down, and vice versa to preserve warmth. This can be detected by using long-wave infra-red (thermographic) imaging cameras, as illustrated in Figure 1. We call this approach “ThermalSense.”

The main contributions of this paper are:

1. Algorithms for dynamically estimating occupants’ real-time thermal preferences using thermographic imaging and machine learning.
2. A study (with 698 data samples over 5 weeks spanning 30 individuals) and analysis into how occupants reported their dynamic thermal sensations and thermal comfort preferences, which also showed that, on around 40% of occasions, energy could be saved if realtime thermal preferences were used rather than using standard air temperature based control.
3. A further analysis (based on the same study) that demonstrated that ThermoSense (e.g. based on Fanger’s PMV factors). ThermalSense achieved an accuracy of 94-95% at predicting whether heating/cooling should be actuated to maintain comfort while minimizing energy use.

RELATED WORK

Early, seminal work in predicting thermal preferences was done by P.O. Fanger [8], whose Predicted Mean Vote (PMV) uses six primary factors known to affect thermal sensation - air temperature, relative humidity, radiant temperature, air speed, clothing insulation and metabolic rate. This was not designed for a realtime sensing scenario, but instead for offline analysis.

PMV uses a 7-point thermal sensation scale ranging from -3 (cold) to +3 (hot) where 0 (neutral) is assumed to be the most desirable level. However, other work [19] has identified that people often do not interpret “neutral” as most comfortable/desirable, and often prefer non-neutral temperatures.

To add further complexity, more recent work has highlighted how the relationship between thermal sensation and thermal comfort is by no means immutable. People adapt to their environment, both using immediate remedies such as wearing a sweater, and also through changing their expectations of what constitutes “comfortable” [2, 26, 5].

In our study (to be described later), we ask participants both about their current thermal sensation and their comfort at that temperature as well as other temperatures. This allows us to compare with findings from the literature, and also to ensure

that our analysis addresses the differences between thermal sensation and comfort.

There are existing systems that aim to adapt temperature dynamically. Neurothermostat [25] is early work in this area, which tries to learn an occupant’s behavior in changing thermostats, and then automate this control on the user’s behalf. Thermovote [7] brings humans more directly into the loop, using participant feedback on whether people are feeling hot/cold/neutral to decide how they want to adjust temperatures to improve occupant comfort.

SPOT [13] and its successor SPOT+ [12] (with room occupancy prediction ability) try to customize the PMV equation for a person by determining an offset from the comfort which would be experienced by the mean population in a building. It uses multiple sensors deployed in a room to determine all six parameters of PMV, and thereby develop a personalized equation for PMV.

Huang et. al explored the idea of using off the shelf sensors in wearable devices to monitor the PMV factors, as well as new factors such as sweat and activity level [17]. They postulated that ASHRAE scale of thermal sensation did not consider the comfort of a person at non-neutral sensations, and therefore defined their own scale which considered comfort as well as sensation.

A study conducted in Iran in summer and winter seasons, predicted the thermal sensation experienced by visitors in an outdoor location [20]. They measured factors such as air temperature, radiant heat, air velocity and relative humidity. Other than microclimatic factors, they also captured demographic factors, such as gender, reason to visit the place, clothing insulation etc.

More research in this area, as well as some recent commercial products, focuses on the problem of ensuring that the HVAC system is only active when people are around. Some work in this area uses occupancy sensors in the home, e.g. Lu et. al. [23], PreHeat [31] and Koehler et. al. [21]. Other work by Gupta et. al. [15] uses GPS from a user’s smartphone to predict future occupancy. Google’s Nest and similar products have brought occupancy-reactive heating to the commercial market.

THERMALSENSE APPROACH

ThermalSense aims to automatically and dynamically optimize for both energy saving and comfort through dynamically sensing thermal comfort preferences. Figure 2 illustrates how such dynamic sensing can allow an HVAC system to both avoid comfort “mistakes” as people’s thermal comfort goals change over the course of a day, and save energy by not heating or cooling beyond what is necessary for comfort. In contrast, a constant-air-temperature system must target a conservative temperature close to the average of the comfortable range, if changes in this range are not sensed.

We use thermographic imaging of skin temperatures to infer not only how comfortable people are currently, but also predict how comfortable they would be if it were slightly warmer or slightly cooler.




Figure 2: If an HVAC system can keep the temperature at the edge of the comfortable range at all times, energy can be saved compared to a constant air temperature approach, while also maintaining comfort.

We now describe the physiology behind how skin temperatures (as illustrated in Figure 1) are indicative of thermal comfort. The human body uses thermal regulation techniques to heat and cool itself, in order to maintain a constant core temperature. In addition to other methods such as sweating and thermogenesis (e.g., shivering), a key method employed by the body is widening or narrowing blood vessels in the skin [16].

Vasodilation, i.e. widening of blood vessels, causes increased blood flow in the skin, allowing the body to radiate heat. Conversely, vasoconstriction decreases blood flow in the skin, and reduces heat loss. The state of vasodilation and vasoconstriction in the extremities of the human body are considered to be indicative of the thermal state of the body [27]. Furthermore, previous work has reported that a person’s thermal comfort can be related to their skin temperatures [11, 10, 32].

DATA COLLECTION

To explore the feasibility of using thermographic images in this way, we conducted a data collection-based study with offline analysis. The data gathered included sensor measurements and also captured realtime thermal preferences through surveys.

We conducted the data collection with participants working in an office building in the UK for 5 weeks during summer. Using an office building rather than a home environment, allowed us to get data from 30 participants over this period. While some studies of thermal comfort are done in highly controlled and varied climate chambers, our focus was on in-situ and realtime thermal preferences that an occupant experiences as part of a normal temperature-regulated indoor environment.

The office was equipped with a HVAC system which was operational for the full period, i.e. maintaining a controlled air temperature. This environment was a deliberately difficult one for our approach — it would have been much easier to infer that people want cooling when they are in a hot room rather than in one kept at a comfortable standard temperature. However, what ThermalSense aims to do is to enable small

1. Please consider for five seconds, and then indicate how you feel: Hot, Warm, Slightly Warm, Neutral, Slightly Cool, Cool, Cold
2. On a scale of 1-5, how comfortable are you
 - a. At the current temperature
 - b. If it were a bit warmer
 - c. If it were much warmer
 - d. If it were a bit cooler
 - e. If it were much cooler
3. Please list all clothes that you are wearing

Figure 3: Survey conducted each time thermographic readings were taken

changes to air temperature in an HVAC-controlled environment, that maintain comfort but reduce energy use. This motivates studying an environment which is already under HVAC control.

We recruited participants who worked in the same building, through email and direct contact, offering a 25 GBP gift card as gratuity. There were 30 participants, 24 male and 6 female, from three different floors of a five-story office building, 9 with individual offices who had the ability to adjust the office temperature to their preference, and 21 with shared office spaces which used a fixed shared setpoint. One of the authors of the paper was a participant. All participants were aware of our overall idea to use thermographic imaging to estimate comfort, but we did not show the participants their own thermographic images.

We visited participant’s workspaces twice per working day, once in the morning, once in the afternoon, over a period of 5 weeks. At each visit, if the participant was available, we verbally conducted a short survey about thermal sensation and thermal comfort, and we recorded sensor data concerning ambient conditions as well as thermographic images - details below.

We collected a total of 700 responses across all participants. The median number of responses per participant was 24, the minimum was 8 and the maximum was 33. At the end of the study, we conducted exit interviews in which we asked the participants to reflect more deeply upon their perception of thermal comfort.

The survey, shown in Figure 3, was designed and refined through pilot deployments to take under 1 minute to complete, to avoid interrupting the participants’ work too much. The first question concerned thermal sensation using the ANSI/ASHRAE 55 standard, which is a seven point scale ranging from Cold to Hot. The second question of the survey looked at thermal comfort, both at the current temperature and also warmer or cooler temperatures. The third question of the survey allowed us to measure the clothing insulation level, based on a database [24], as used in Fanger’s PMV equation.

1. Please explain briefly what “how you feel” means to you and how do you determine your response? (For example, do you focus internally on your body on specific body parts (if so which?), or externally on your perception of the environment, or both?)
2. Specifically, with the word “neutral”, do you interpret this as the most comfortable temperature for you, or do you interpret this to mean that the environment is at a neutral temperature.
3. If your response to the previous question is the latter, which point on the scale of (Hot, Warm, Slightly Warm, Neutral, Slightly Cool, Cool, Cold), do you think you are typically most comfortable at?

Figure 4: Post-study survey

Since participants were sitting when we visited them, we used the standard metabolic level of 1.0 MET.

We also measured air temperature and humidity (with a Silicon Labs Si7020 sensor), and radiant temperature (with the thermographic camera discussed later), which are again factors well known to influence thermal comfort and form part of Fanger’s PMV equation. We did not measure air velocity with every sample since initial experiments showed that the air velocity was low and constant - we used a constant 0.1 m/s^2 for air velocity.

Finally, during each visit we recorded various body temperatures for participants. We used a FLIR A655sc thermographic camera, which has a 640x480 resolution and a sensitivity of 0.05°C , with the T198065 80° field-of-view lens. We took two thermographic images of the participant — one showing the backs of their hands and another showing the front of their hands (see Figure 5).

In future, in an integrated solution, we envisage a computer vision system would determine when someone was present and then to sample automatically from the right body location, either make use of the thermographic image alone, or an RGB or depth image from an additional camera. However, for this study, since we did not know which body part temperatures would be most informative for our application, we used a manual approach.

To choose which body part temperatures to measure, we conducted a pilot study with 4 participants where we controlled the ambient temperature and got them to change clothing and their metabolic level as well (with an exercise bike). Based on this preliminary study and also on the regions of the body that are visible when wearing normal indoor clothing, we chose to capture the body regions shown in Figure 5, specifically 7 facial regions: forehead, cheeks (x 2), lips, jaw (includes lips), upper neck and lower neck, and 6 hand regions: center of




Figure 5: Body regions where thermographic data was collected: forehead, cheeks, lips, jaw (including lips), upper neck, lower neck, palm core, palm (includes palm core), and back-of-hand (not shown in figure).

palm (x 2), palm (which includes center of palm) (x 2), and back of hand (x 2). The initial experiments showed us that the nose region was not useful.

Another potential region for thermal annotation were the fingertips. This was based on the ideas proposed by Humphrey et al. [18], in which they suggested that fingertip temperatures along with radiant temperature was highly correlated to the thermal sensation reported by participants. However, according to a report by Wang et al. [32], these temperatures vary substantially among people expressing the same thermal sensation, exceeding 10°C among people feeling neutral or cooler than neutral. They suggested that this variability might limit the accuracy of any method that uses fingertip temperature as a predictor of thermal sensation or comfort. Therefore, we decided to leave out the fingertip regions.

For the captured thermographic images, which comprise a temperature per pixel, we used FLIR’s Research IR software offline to manually annotate each region as a rectangular area and obtain the min, max, mean, range, and standard deviation of temperatures across the rectangle. These were later used as features input into our machine learning classifications, as discussed in the next section.

Finally, we conducted a short interview at the end of our study, as shown in Figure 4. This allowed us to find out how the participants related their responses to the thermal sensation and thermal comfort questions; as others have found [19], a “Neutral” thermal sensation is not always regarded as most comfortable. We also asked how people evaluated their thermal comfort, as we wanted to see if we could correlate this with the body part temperatures that we measured.

THERMAL COMFORT DYNAMICS

The first step in our investigation is to explore how thermal comfort changes over time and how it relates to thermal sensation. As previously stated, thermal sensation and thermal

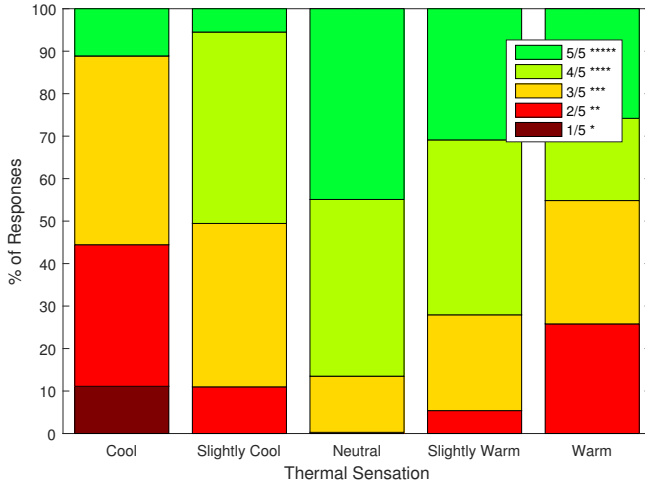


Figure 6: Thermal comfort ratings at current thermal sensation. A 5/5 rating indicates maximum comfort, and 1/5 rating indicates least comfort. Participants did not always equate a neutral sensation with high comfort, or a non-neutral sensation with low comfort.

comfort have a complex relationship and it is important to validate that our participants’ responses are reasonable, and in line with previous findings, so we can rely on them in training and evaluating our thermographic imaging approach.

We began by analyzing the survey responses. Table 1 shows the number of responses received per thermal sensation. There were no reports of either ‘Hot’ or ‘Cold’ thermal sensation, which was no surprise given that the HVAC system was operating normally, and under 6% of responses were “Cool” or “Warm”.

Relating thermal comfort (survey question 2) to thermal sensation (survey question 1), Figure 6 shows the ratings received for the current thermal sensation on a scale of 1 to 5, where 1 being least comfortable, and 5 being most comfortable. Our participants showed differences in interpretation between sensation and comfort, consistent with past studies in this area [19].

Thermal Sensation	# of Responses
Cool	9
Slightly Cool	91
Neutral	363
Slightly Warm	204
Warm	31

Table 1: Number of times each thermal sensation was reported.

In a further analysis, we combined the data from the comfort and sensation scores in order to identify the range of thermal sensations which led to maximum comfort. This range often consisted of more than one sensation level because the maximum comfort score could be given to both e.g. the current temperature and also to a slightly cooler temperature. For the




Figure 7: Participants often preferred cooler-than-neutral or warmer-than-neutral thermal sensations, and also varied their preference over time

purposes of this extrapolation we assumed that “a bit warmer” on question 2 equated to one point on the thermal sensation scale (e.g. moving upward from “Neutral” to “Slightly warm”). Note that this assumption is not relied on by other analyses in this paper.

Factor	% of participants
Whole Body	40
Head	30
Torso	17
Arms	40
Legs/Feet	17
Least Comfortable Part	7
Environment	7

Table 2: Factors considered by participants in determining their thermal sensation

This analysis is illustrated in Figure 7, which shows that participants often associated maximum comfort with a non-neutral thermal sensation. Furthermore, their preferences were not consistent; most participants varied between desiring warmer-than-neutral or cooler-than-neutral, though many were biased towards one over the other.

These observations were further backed up in the exit survey (Figure 4). Question 1 of the survey asks the participant to elaborate on what thermal sensation meant to them and how they determined their response - whether they focused on their body or the perceived environment (air temperature). About 93% of the participants referred to the thermal perception of their individual body parts, while 7% of the participants referred to their perception of the environment, to determine their thermal sensation. Table 2 shows a summary of what percentage of participants refer to each of the factors for determining thermal sensation. The total exceeds 100% because some participants referred to more than one factor to make this determination.

Thermal Sensation	# of Participants
Neutral	19
Cooler than neutral	4
Warmer than neutral	6
Depends on activity	1

Table 3: Preferred thermal sensation reported during exit interview.

Many of the participants referred to thermal discomfort felt at very specific parts of the body, such as arms and legs. P1 says, *“I would use my own temperature, not the environment. It probably was slightly disproportionately focused on my arms and head, since that’s where I notice heat the most, but it was more of a general sense of how I was hot.”* P12 says, *“I focus on my arms and legs. If they are feeling chilled, then I feel ‘slightly cool’, if they are feeling warm, then I feel ‘slightly warm’. If neither, then I feel neutral”*.

Some participants generally based their responses on whether they felt extreme thermal discomfort in any part of the body. P15 says, *“There’s always variations in how different parts of my body feel, so usually when I think of how I feel overall, I consider what the least-comfortable part of my body feels like, so my answer tends to be a bit biased in that direction (i.e.: if my hands are really cold, my answer will likely be that I’m slightly cool for example)”*. P23 says, *“I think I would say ‘cold’ if I feel cold on any part of my body (e.g. if either feet or head are feeling cold, I would say I am feeling cold). Analogously, for feeling warm. If nothing is out of the ordinary, (I’d say) neutral”*.

Some participants used discomfort felt at different parts of the body for determining cold sensation vs. warm sensation. P28 says, *“I guess (I used) my body to decide. At this time of your survey since I wear shorts and sandals my legs and feet feel the cold first (so that’s where I focus for cold) and my face for heat. Not sure what happens if I have a hot face and cold feet”*.

While, some participants focused on individual parts of the body, others took a more holistic view in making this determination. P3 says, *“it involved considering how my body, on the whole, felt temperature wise. It was a somewhat internal sense, though it did not rely on any body part in particular. I somewhat found myself thinking about my ‘core’ body, if anything. It was almost entirely focused on my body, though not the environment around it”*. P17 says, *“I mostly focus on how warm my body feels, from my waist upwards specifically. When I am too warm I usually feel it around my shoulders, chest, back, face etc, so that is where I pay most attention to first, to gauge how I actually feel”*. P2 says, *“I am generally trying to think how I feel in my body but then often cold hands/arms will sway me from neutral to slightly cool. Otherwise it’s generally body temperature”*.

A small minority of the participants also considered their perception of the room temperature. P6 says, *“I think I focus on body parts and sort of the external environment as well. I think about how my hands and feet feel I guess because they are the extremities and then how my skin feels in general, especially on my arms and face as they are usually uncovered.*

I also think about what the temperature in the room is but to a lesser extent than the body parts mentioned”. P14 says, *“It is always about the perceived temperature differential between me and the environment (I am aware that this is not objectively true but this is the perception I use)”*.

As can be seen in Table 3, 11 out of 30 participants stated that they preferred non-neutral thermal sensations. Participant P1 reported, *“I definitely prefer being slightly cool... I interpreted question 1 as asking my temperature, not my comfort level.”* Participant P3 said, *“I typically think of myself as most comfortable at Warm. In context of the last question, this means that I am more comfortable if my body feels ‘warm’ than if it feels ‘neither warm nor cold.’”*. Another participant preferred different thermal sensations depending on what they doing. P15 says, *“If I’m coding or trying to get something done, I’ll often like it to be a bit cooler, so I don’t fall asleep and stay alert... Whereas if I’m more in the mood to waste some time online, etc., I’ll like it to be warmer”*.

In conclusion for the above analysis, we have validated that our participants’ responses seem to be in line with recent work (e.g. [19]) in differing between their approach to thermal sensation and thermal comfort. Through the rest of this paper, we focus on the participants’ responses concerning thermal comfort - both currently and if cooler or hotter - and regard them as the ground truth that we need to predict.

Variability in comfortable temperatures

We analyzed how thermal comfort changes with air temperature. Figure 8 shows box plots for each participant of the air temperature distributions when the current temperature was reported as delivering maximum comfort (green), and when it was not (yellow).

For some participants such as P2, P4 and P27, a single constant air temperature could be found which results in maximum comfort nearly all or all of the time. However, in most cases, air temperature-based control is insufficient. In other words, the same air temperature can be reported at different times to be maximum-comfort, or to be less-than-maximum-comfort.

POTENTIAL FOR ENERGY SAVING

If we refer back to Figure 2, we can see various kinds of performance states of an HVAC system. One simple performance indicator is whether the air temperature is inside the green band, i.e. comfort is achieved. Comfort mistakes can take two forms: one where too little energy was used (e.g., not heating enough in winter), and where too much energy was used (e.g. over-heating in winter). Even when comfortable, any energy used beyond that required to get to the edge of the green band (i.e., the red and blue lines on the figure) can be regarded as wasted energy.

With this categorization in mind, Figure 9 shows the thermal preferences exhibited during the study, through the viewpoint of both a cooling requirement (e.g. in a hot country/summer), and a heating requirement (e.g. in a cold country/winter).

The top segment refers to situations where the HVAC system got it right — the participant was comfortable, and would not




Figure 8: Box-plots of air temperatures where the participant reported maximum comfort versus air temperatures where the participant reported that comfort would be improved at a lower or higher temperature.

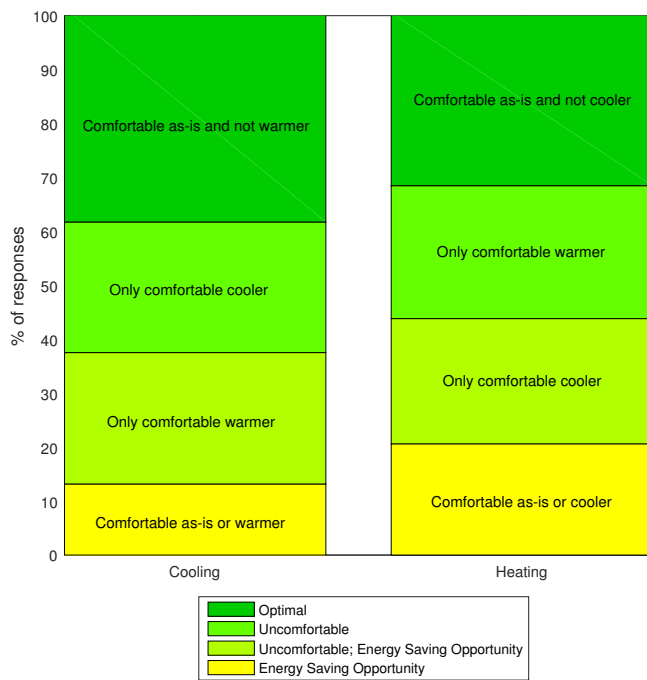


Figure 9: Potential energy savings based on the participant responses in the study

be as comfortable if less energy would have been expended. The second-from-top segment refers to situations where the system spent too little energy, i.e. not enough to achieve comfort. The second-from-bottom segment refers to situations where over-cooling or over-heating was done, i.e. too much energy was spent and comfort was not maintained. The bottom segment refers to conditions where the participant was comfortable, but they would also have been comfortable if less energy were spent.

The bottom two segments combined represent instances where energy could have been saved. These sum to just under 40% for cooling, and just above 40% for heating. This shows that there is a huge opportunity to do better than existing approaches, if real-time thermal comfort preferences could be measured.

Because of the nature of the data collection study, we do not attempt to translate these identified saving opportunities into actual energy numbers. This would depend not only on whether the system was heating or cooling, but also on how extreme the outside temperatures are.

ESTIMATION OF THERMAL PREFERENCES

We now evaluate the ability of ThermalSense to determine real-time thermal preference. As previously described with reference to Figure 5, we collected thermal images covering 13 body regions and for each region extracted 5 features: max, min, mean, range and standard deviation. We also include air temperature as a feature in all analyses, since this is easy to measure and already measured by existing HVAC systems.

To evaluate ThermalSense, we focus on predicting the necessity of energy expenditure: should the HVAC system now perform energy consumption, or can it avoid doing so. As such, the classification tasks are binary classifications where the responses that fall into the second-from-top segment in Figure 9 form one class (where energy must be used or discomfort will occur) and all other response types form the other class (where energy does not currently need to be used to maintain comfort). As the classes are unbalanced, we oversample the minority class [4].

For the machine learning algorithm we selected *Rotation Forests* [30], which is a classifier ensemble method based on *Random Forest*. We chose this as ensemble methods based on Random Forests have been shown to achieve high classification accuracy in a variety of datasets [9]. We also tried using random forests, kNN, decision trees and SVM classifiers, and these all performed worse; we do not report detailed comparisons due to space constraints. Results are obtained using 10-fold cross-validation.

We compared six approaches:

1. The modal class. This provides a baseline that a “best guess” would achieve. When participant ID is not provided, this is 50% since we balance the classes. When participant ID is provided, this is not 50% since individ-




Figure 10: Comparison of ThermalSense (AT+BodyTemp) against alternative approaches, for classifying whether to increase energy use or not. Analysis is done for both cooling (e.g. summer) and heating, and both with the participant ID and without. AT = air temperature, RH = relative humidity, Clo = clothing insulation, RT = radiant temperature.

ual participants’ responses are not balanced, and so a “best guess” can be more accurate.

2. Fanger’s equation [8], which estimates thermal sensation, given air temperature, radiant temperature, relative humidity, clothing level, metabolic level and air velocity (the latter two being regarded as constant in this study).
3. Machine learning with air temperature and relative humidity only - i.e., using the same sensors as might be used in a standard thermostat.
4. Machine learning with Fanger’s equation factors - air temperature, relative humidity, clothing level and radiant temperature.
5. Machine learning with body temperatures as described above and air temperature (ThermalSense).
6. Machine learning with everything: body temperatures and Fanger’s equation factors.

We perform the above comparison for four scenarios — both with and without the participant ID being provided to the machine learning algorithm during training and testing, and modeling both a cooling-oriented situation (e.g., in summer) and a heating-oriented situation.

Figure 10 shows the accuracy obtained comparing the various methods and scenarios described. Fanger’s PMV equation performs similar to or worse than the modal value. Using air and humidity data alone performs badly (<65% accuracy) when the participant ID is not known, but much better (>90%) when it is known. In other words, simply knowing the ID of the person present as well as having per-person training results a significant improvement.

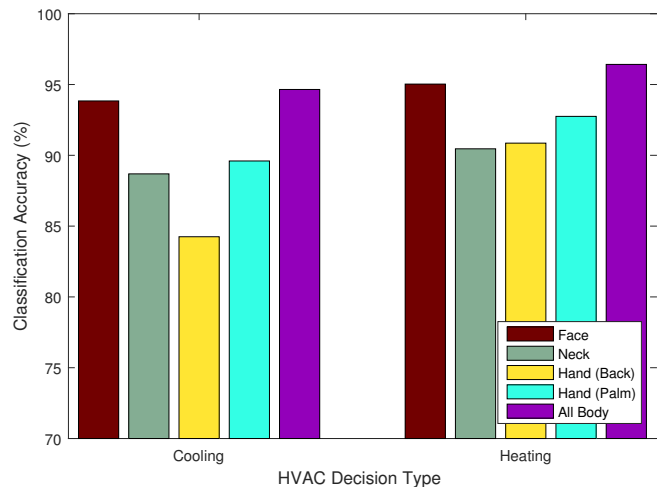


Figure 11: Accuracy of classifying whether to use energy or not based on ThermalSense with data from various parts of the body (and air temperature).

Machine learning based on factors used in Fanger’s PMV equation can do better still, with accuracy of 72%-74% without person ID and person-specific training and 91%-92% with person ID/training.

ThermalSense performs well. Without the participant ID, it achieves 94%-96% accuracy in predicting whether energy consumption is necessary to maximize comfort. With participant ID, the accuracy improves only by a tiny amount if at all. This is very heartening as it implies that our approach can work without person identification while the system is operating. This can also have privacy advantages (though the use of thermal imagery is still something that requires a careful consideration of privacy and data security). The use of relative humidity, clothing level and radiant temperatures in addition to body temperatures also does not improve the accuracy much further, only by 0.5%-1%. This is again good, because clothing insulation level in particular is difficult to infer.

However, ThermalSense does require training data for the individual being sensed. While our cross-validation did not mix training and test sets, data from each person was present in the training set and the test set. When using different people for training sets and test sets, ThermalSense’s performance dropped to 63% (cooling) and 68% (heating) as compared to the 94%-96% achieved with prior training on each individual — but without realtime person ID.

Body Regions

One barrier to using thermographic imaging for HVAC control is in the need to capture thermal information about specific body parts. In our study, we explored the most commonly visible body parts in office environments, namely regions of the face and hands, and achieved an overall accuracy of around 95%.

To further understand if these regions are all necessary to achieve this accuracy, we re-ran the analysis using subsets of the thermographic data. As Figure 11 shows, when split

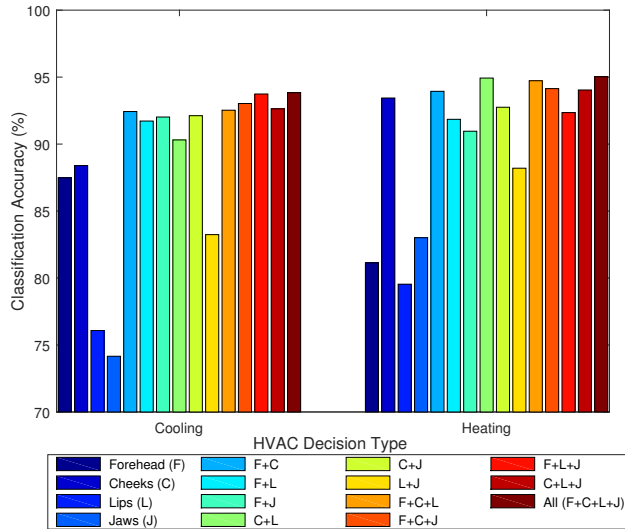


Figure 12: Accuracy of classifying whether to use energy or not based on ThermalSense using data from facial regions (individual and combinations), and air temperature.

into regions, the face (5 regions) outperforms the neck (2 regions), the rear of the hand (1 region) or the palm of the hand (2 regions). Accuracy of using facial features alone stands at 94%-95%, compared to 94-96% for all regions combined.

Figure 12 further breaks down the regions of the face and looks at them individually and in combination. Among individual sections of the face, we observe that the use of temperature features from cheek performs much better than the other three regions. However, no subset of the regions performs within 1% of all facial regions combined, for both heating and cooling conditions.

In summary, we have shown a 94-95% classification accuracy for predicting whether an HVAC system must expend energy or not to maintain comfort, and that this can be achieved using thermographic data from only an individual’s face, and without knowing that individual’s identity at runtime. This is very encouraging.

DISCUSSION

While encouraging, our investigation is also limited in many ways. We had only 30 participants, who were all in only one geographic location and using one HVAC system (which was a commercial one in a five-storey office building, rather than a smaller office or domestic system). The data collection period was only 5 weeks, during a specific season of the year (summer). The study would clearly benefit from being repeated with variation in the above aspects.

Furthermore, we did not build a closed-loop system, so we cannot demonstrate or quantify actual energy or comfort improvements at this stage. In the rest of this section we will elaborate further on the challenges of turning the ThermalSense approach and offline evaluation into a closed-loop system.

Building out a real system

One stage of a real system which we did not explore in this study is automatic segmentation of the thermographic images - we relied on manual segmentation during this investigatory phase. However, now that we know that we only need to track the head of a person to determine their thermal comfort preference, we can use known vision techniques that can track faces using only thermographic imaging [22, 28, 34].

ThermalSense’s inferences are based on occupants’ predictions of their thermal comfort level if it were “slightly” warmer or cooler. This raises a few questions. Firstly, how accurate are our occupants’ predictions — they may occasionally be wrong. Second, how much is “slightly” and does this differ per-participant? If a closed-loop system were built, the system could learn this during a training period, by observing how reported preferences change as the temperature is changed.

Cooling and heating systems have actuation latency, and this can be significant (e.g. up to hours for under-floor heating). What we have shown is the ability of ThermalSense to make accurate instantaneous recommendations, but the performance under realistic latencies must be further evaluated.

One of the key questions in implementing such a system would be — where should thermographic camera(s) be placed. If the system is used to control the HVAC for a office room, then deploying the camera in a way that can track the face of a person when they are at their desk might be sufficient. However, in other spaces such as homes, placement is not as easy. It is unclear whether one would need enough cameras to afford a continuous view of each user [33] or whether it would be sufficient to place a smaller number of cameras in key areas, and rely on extrapolation between times that users were sensed. Furthermore, spaces are often shared among multiple occupants, and this raises difficulties for both thermographic measurement and temperature control.

As far as the cost of individual thermal cameras is concerned, the prices are currently prohibitive - many thousands of dollars. This has recently been decreasing though they are still expensive compared to e.g. webcams. For example, the Melexis MLX90621 is a 16x4 pixel thermographic sensor costing under USD60, while the FLIR One is a 160x120 thermographic camera accessory for a phone costing USD250. It is reasonable to extrapolate that in future thermographic cameras may be affordable enough to deploy widely.

Privacy is a concern with ThermalSense, as is the case with any other camera based system. However since ThermalSense only needs simple temperature features from the facial regions of a person, it may be appropriate to immediately extract this data from each frame and discard the original image, which may go some way to alleviating privacy concerns.

While the performance of ThermalSense currently relies on per-participant training (though not realtime participant identification), further work can look into how this can be minimized or eliminated, e.g. by using a large training set based on hundreds or thousands of individuals.

User Interface

The proposed system removes the need to have manually-controlled setpoint temperatures, since thermographic image sensing is used instead. This is in contrast to work such as PreHeat [31] which removes the need to have manually controlled HVAC operation timings, using occupancy sensing. The combination of the two means that an HVAC system with no user controls could be implemented. Of course, the user should be in ultimate control and overriding UI mechanisms need to be provided to allow the user to handle exceptional situations or cases where predictions are wrong.

By using thermographic imaging, ThermalSense somewhat changes the expectation with regards to energy saving behaviors. For example, with our proposed system, in winter you would never need to put on a sweater because the room would warm up sufficiently if you were not wearing one. This could cause people to be less aware of their energy usage, and cause higher energy use. The user interface must be carefully designed to avoid such effects.

On the other hand, for the conscious consumer who does form the habit of putting on a sweater at home in winter, our proposed system would automatically adjust, without requiring manual control. In this sense, actions such as exercising or wearing more/less clothing are treated as “natural” user interface inputs, without needing to separately inform the HVAC system of how to respond.

CONCLUSION

ThermalSense is a new approach for determining thermal comfort using thermographic imaging, and therefore to enable the control of HVAC systems in order to minimize energy use while maintaining comfort, and also removing the need to manually control the preferred air temperature. People feel comfortable in a range of temperatures, and the aim of this paper has been to discover if the ThermalSense approach can not just determine current thermal comfort, but also predict whether energy could be saved because the user would still be comfortable at a less energy-intensive temperature.

Using data gathered in a 30-participant, 5-week trial in an office building, we found that around 40% of the time there were opportunities to save energy compared to a static air-temperature-based setpoint. We demonstrated that ThermalSense can predict with high accuracy whether it is necessary to actuate the HVAC system to maintain comfort. Specifically, for a cooling-oriented environment our accuracy was 94%, while in a heating-oriented environment it was 96%, using temperatures measured on hands, neck and face. Using just facial features, which may be easier to sense in real deployments, its accuracy drops only to 94-95%. While ThermalSense currently requires per-person training, it does not require realtime identification of the occupant. ThermalSense outperformed alternative approaches based on air temperature, relative humidity, radiant temperature, occupant identity, and occupant clothing level.

Future work in this area includes validating this result in other scenarios (e.g. different climates, building types), and turning this proof-of-concept into a closed-loop system, which

will require computer vision to locate and segment faces in thermal images and integration into an HVAC control system. With further effort in these areas, we hope the ThermalSense approach can reduce HVAC energy consumption while maintaining or improving thermal comfort.

ACKNOWLEDGMENTS

We thank Mike Hazas, Kamin Whitehouse and our anonymous reviewers for their valuable feedback in the preparation of this paper, John Helmes for assistance with building the measurement rig, Jakub Dostal, and all our study participants. This work is supported in part by NSF Graduate Research Fellowship Program under Grant 0809128.

REFERENCES

1. Energy Information Administration. 2009 Residential Energy Consumption Survey. <http://www.eia.gov/consumption/residential/>.
2. Brager, G. S., and de Dear, R. J. Thermal Adaptation in the Built Environment: A Literature Review. *Energy and Buildings* 27, 1 (1998), 83 – 96.
3. Energy Policy Branch Energy Sector Energy Forecasting Division. In *Canada Energy Outlook*, Natural Resources Canada (1997), 1–2.
4. Chawla, N. V. Data Mining for Imbalanced Datasets: An Overview. In *Data mining and knowledge discovery handbook*. Springer, 2005, 853–867.
5. Clear, A. K., Morley, J., Hazas, M., Friday, A., and Bates, O. Understanding Adaptive Thermal Comfort: New Directions for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, ACM (2013), 113–122.
6. CBE Thermal Comfort Tool, Center for the Built Environment, University of California Berkeley. <http://cbe.berkeley.edu/comforttool/>.
7. Erickson, V. L., and Cerpa, A. E. Thermovote: Participatory Sensing for Efficient Building HVAC Conditioning. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, ACM (2012), 9–16.
8. Fanger, P. O. *Thermal Comfort. Analysis and Applications in Environmental Engineering*. (1970).
9. Fernández-Delgado, M., Cernadas, E., Barro, S., and Amorim, D. Do We Need Hundreds of Classifiers to Solve Real World Classification Problems? *The Journal of Machine Learning Research* 15, 1 (2014), 3133–3181.
10. Frank, S. M., Raja, S. N., Bulcao, C. F., and Goldstein, D. S. Relative Contribution of Core and Cutaneous Temperatures to Thermal Comfort and Autonomic Responses in Humans. *Journal of Applied physiology*, 5 (1999), 1588–1593.
11. Gagge, A., Stolwijk, J., and Hardy, J. Comfort and Thermal Sensations and Associated Physiological Responses at Various Ambient Temperatures. *Environmental Research* 1, 1 (1967), 1 – 20.

12. Gao, P. X., and Keshav, S. Optimal Personal Comfort Management Using SPOT. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, ACM (2013), 1–8.
13. Gao, P. X., and Keshav, S. SPOT: A Smart Personalized Office Thermal Control System. In *Proceedings of the fourth international conference on Future energy systems*, ACM (2013), 237–246.
14. Goto, T., Toftum, J., de Dear, R., and Fanger, P. O. Thermal Sensation and Thermophysiological Responses to Metabolic Step-Changes. *International Journal of Biometeorology* 50, 5 (2006), 323–332.
15. Gupta, M., Intille, S. S., and Larson, K. Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, Springer-Verlag (Berlin, Heidelberg, 2009), 95–114.
16. Houdas, Y., and Ring, E. *Human Body Temperature: Its Measurement and Regulation*. Springer Science & Business Media, 2013.
17. Huang, C.-C. J., Yang, R., and Newman, M. W. The Potential and Challenges of Inferring Thermal Comfort at Home using Commodity Sensors. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ACM (2015), 1089–1100.
18. Humphrey, M., McCartney, K., Nicol, K., and Raja, I. An Analysis of Some Observations of the Finger Temperature and Thermal Comfort of Office Worker. *Indoor Air* 99 (1999), 1.
19. Humphreys, M. A., and Hancock, M. Do People Like to Feel Neutral?: Exploring the Variation of the Desired Thermal Sensation on the ASHRAE Scale. *Energy and Buildings* 39, 7 (2007), 867–874.
20. Kariminia, S., Shamshirband, S., Motamedi, S., Hashim, R., and Roy, C. A Systematic Extreme Learning Machine Approach to Analyze Visitors Thermal Comfort at a Public Urban Space. *Renewable and Sustainable Energy Reviews* 58 (2016), 751–760.
21. Koehler, C., Ziebart, B. D., Mankoff, J., and Dey, A. K. TherML: Occupancy Prediction for Thermostat Control. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, ACM (2013), 103–112.
22. Kong, S. G., Heo, J., Abidi, B. R., Paik, J., and Abidi, M. A. Recent Advances in Visual and Infrared Face Recognition-A Review. *Computer Vision and Image Understanding* 97, 1 (2005), 103–135.
23. Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., Field, E., and Whitehouse, K. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, ACM (2010), 211–224.
24. McCullough, E. A., Jones, B. W., and Huck, J. A *Comprehensive Database for Estimating Clothing Insulation*. ASHRAE, 1985.
25. Mozer, M. C., Vidmar, L., Dodier, R. H., et al. The Neurothermostat: Predictive optimal control of residential heating systems. *Advances in Neural Information Processing Systems* (1997), 953–959.
26. Nicol, F., Humphreys, M., and Roaf, S. *Adaptive Thermal Comfort: Principles and Practice*. Taylor & Francis, 2012.
27. O'Brien, C., Castellani, J. W., and Sawka, M. N. Thermal Face Protection Delays Finger Cooling and Improves Thermal Comfort during Cold Air Exposure. *European journal of applied physiology* 111, 12 (2011), 3097–3105.
28. Pavlidis, I., and Dowdall, J. B. Near-Infrared Method and System for Use in Face Detection, Apr. 11 2006. US Patent 7,027,619.
29. Rathouse, K., and Young, B. Domestic Heating: Use of Controls. Technical Report RPDH 15. In *Building Research Establishment, UK* (2004).
30. Rodriguez, J. J., Kuncheva, L. I., and Alonso, C. J. Rotation Forest: A New Classifier Ensemble Method. *IEEE Trans. Pattern Anal. Mach. Intell.* 28, 10 (Oct. 2006), 1619–1630.
31. Scott, J., Brush, A. J., Krumm, J., Meyers, B., Hazas, M., Hodges, S., and Villar, N. PreHeat: Controlling Home Heating using Occupancy Prediction. In *Proceedings of the 13th international conference on Ubiquitous computing*, ACM (2011), 281–290.
32. Wang, D., Zhang, H., Arens, E., and Huizenga, C. Observations of Upper-Extremity Skin Temperature and Corresponding Overall-Body Thermal Sensations and Comfort. *Building and Environment* 42, 12 (2007), 3933–3943.
33. Wang, Y., and Cao, G. Barrier Coverage in Camera Sensor Networks. In *Proceedings of the Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ACM (2011), 12.
34. Wang, Y.-T., and Ming-Shi, C. Human Face Recognition using Thermal Image. *Journal of Medical and Biological Engineering* 22, 2 (2002), 97–102.