Energy & Buildings 312 (2024) 114063

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enb

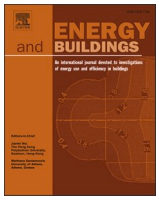

Hot, cold, or just right? An infrared biometric sensor to improve occupant comfort and reduce overcooling in buildings via closed-loop control

Ronnen Levinson [a,*], Donghun Kim [a], Howdy Goudey [a], Sharon Chen [a], Hui Zhang [b], Ali Ghahramani [b], Charlie Huizenga [b], Yingdong He [b], Akihisa Nomoto [b], Edward Arens [b], Ana Álvarez Suárez [c], David Ritter [c], Markus Tarin [c], Robert Prickett [d]

[a] Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[b] Center for the Built Environment, University of California, Berkeley, CA, USA
[c] MoviTHERM, Irvine, CA, USA
[d] Daikin Silicon Valley, Santa Clara, CA, USA

ARTICLE INFO

Keywords:
Air conditioning
Overcooling
Thermal comfort
Skin temperature
Infrared thermography
Machine vision
Closed-loop control
Energy savings

ABSTRACT

To improve occupant comfort and save energy in buildings, we have developed a closed-loop air conditioning (AC) sensor-controller that predicts occupant thermal sensation from the thermographic measurement of skin temperature distribution, then uses this information to reduce overcooling (cooling-energy overuse that discomforts occupants) by regulating AC output. Taking measures to protect privacy, it combines thermal-infrared (TIR) and color (visible spectrum) cameras with machine vision to measure the skin-surface temperature profile. Since the human thermoregulation system uses skin blood flow to maintain thermoneutrality, the distribution of skin temperature can be used to predict warm, neutral, and cool thermal states. We conducted a series of human-subject thermal-sensation trials in cold-to-hot environments, measuring skin temperatures and recording thermal sensation votes. We then trained random-forest classification machine-learning models (classifiers) to estimate thermal sensation from skin temperatures or skin-temperature differences. The estimated thermal sensation was input to a proportional integral (PI) control algorithm for the AC, targeting a sensation level between neutral and warm. Our sensor-controller includes a sensor assembly, server software, and client software. The server software orients the cameras and transmits images to the client software, which in turn assesses occupant skin temperature distribution, estimates occupant thermal sensation, and controls AC operation. A demonstration conducted in a conference room in an office building near Houston, TX showed that our system reduced overcooling, decreasing AC load by 42% when the room was occupied while improving occupant comfort (fraction of "comfortable" votes) by 15 percentage points.

1. Introduction

Delivering a thermally comfortable indoor environment is a primary goal of building heating, ventilation, and air conditioning (HVAC) systems. Most HVAC control systems regulate indoor air temperature and humidity to setpoints obtained from industry standards or operators' empirical judgments. However, environmental parameters, such as air movement and radiation, and personal factors, such as clothing, activity level, and thermal adaptation over various timescales substantially influence the optimum thermostat setpoint for an occupant. Controlling to a fixed temperature setpoint rather than to a measurement of occupants' actual comfort is a longstanding shortcoming that causes large fractions of commercial building occupants to find their thermal environment uncomfortable [1,2].

HVAC operation often produces overheating in winter and overcooling in summer [3–6]; our current study pertains to the latter. Overcooling is highly energy-intensive, and correcting it offers substantial savings. Across a range of climates in the United States, raising an excessively low thermostat cooling setpoint by 1 °C reduces annual HVAC energy use by an average of 10 % [7]. Overcooling is also a worldwide phenomenon, now widespread in hot and humid climates of Asia [8,9].

The U.S. Energy Information Administration projects that by 2030, the United States will consume 1.4 Quad/y [1.5 EJ/y] of primary energy

* Corresponding author.
E-mail address: RML27@cornell.edu (R. Levinson).

https://doi.org/10.1016/j.enbuild.2024.114063
Received 15 August 2023; Received in revised form 30 January 2024; Accepted 5 March 2024
Available online 6 March 2024

to cool commercial buildings and another 2.4 Quad/y [2.5 EJ/y] to cool homes [10]. Eliminating overcooling could reduce cooling energy use in the U.S. commercial building sector by 35 %, saving 0.5 Quad/y [0.53 EJ/y] [11], and may save a comparable amount of energy in U.S. homes.

There are several causes for such overcooling, all based on HVAC design and operation assumptions that are more conservative than necessary [5–9]. Overcooling could be caught and corrected by the HVAC control system if it could detect the occupants' thermal comfort, as well as the number of occupants in each thermal zone. Knowing these would allow the building to be controlled responsibly at lower supply air flows and seasonally appropriate indoor temperatures.

Thermal comfort is traditionally assessed by surveying occupants, but it can be predicted from measured physiological variables such as a person's skin temperature [12,13]. Commercially available or prototype solutions for addressing personal comfort include smartphone applications that ask occupants to rate their comfort [14]; networked personal hardware, such as wearable wireless sensors that measure core and skin temperatures [15]; or these techniques combined with occupant counting [16]. However, these solutions require occupant participation, may not be effective in transiently populated public spaces, and do not detail the spatial temperature distribution needed for holistic assessment of room heat balance [17]. Available passive (non-participatory) systems use occupant *count* to control the HVAC equipment, yielding *open-loop* comfort control—that is, HVAC regulation without assessment of occupant comfort. Currently, no passive systems provide the occupant comfort feedback required for *closed-loop* comfort control.

Closed-loop control systems with application-appropriate sensors have been demonstrated to substantially reduce cooling loads in buildings. For example, Bell et al. [18] found that the energy consumed to cool computer server rooms could be reduced by 30–40 % if air conditioning (AC) operation was regulated using air temperatures measured at server fan inlets—the locations that matter—rather than at the ceiling return.

We seek to improve occupant comfort and save cooling energy by implementing a closed-loop air conditioning sensor/controller that radiatively detects occupants and their thermal sensations, then uses this information to reduce overcooling (cooling-energy overuse that discomforts occupants) by regulating AC output. Note that we use the term "air conditioning" to refer only to space cooling, by which we mean mechanical removal of heat from the occupied space. This scheme will supplement conventional wall-mounted room air temperature sensors with a wall- or ceiling-mounted, wide-angle sensor that views the occupants. It combines thermal-infrared (TIR) and color (visible spectrum) cameras with machine vision to determine the location, skin-surface (hereinafter, simply "skin") temperature profile, and thermal sensation of each observable occupant and can evaluate the collective sensation of the occupants.

To maintain privacy, the sensor does not identify individuals or share images, and discards images after using them to help locate heads and hands; no personally identifiable information (PII) is used or generated. Fully anonymous image analysis can also provide real-time occupancy counting that can be used by heating, lighting, and ventilation controls; that can provide data for scheduling services with machine learning; and that can support building security. This sensor could be installed in either commercial or residential buildings and would be especially valuable for public spaces with transient populations, such as meeting rooms, auditoriums, restaurants, and stores, in which occupants have little control over their comfort.

Our approach combines two recent innovations to assess occupant skin temperature distribution: (1) low-cost (≤US$165 as of July 2023), medium resolution (160 × 120 pixel), uncooled TIR camera detectors (microbolometers) which can be coupled with very-low cost (≤US$10), megapixel-resolution color camera detectors; and (2) open-source computer vision image recognition software that can quickly locate body parts of interest, such as the face, nose, and hands.

This article proceeds in seven steps:

- **Physiology.** We review how skin blood flow and temperature respond to thermal stress.
- **Sensing.** We describe the development and testing of the radiometric sensing hardware and software needed for skin temperature distribution thermography, including color and thermal cameras, temperature references, and machine-vision code.
- **Sensation.** We summarize the relationship between thermographic skin temperature distribution and thermal sensation developed in a series of human-subject trials.
- **Control.** We elaborate the incorporation of thermal sensation estimated from skin-temperature distribution as input to an AC control algorithm that seeks to minimize cooling energy use while maintaining occupant comfort.
- **Integration.** We describe the development of our sensor-controller that incorporates the radiometric sensor, thermal sensation model, and control algorithm to minimize overcooling.
- **Demonstration.** We summarize the performance of our sensor-controller assessed through trials in an office space.
- **Future development.** We present plans to upgrade our device.

2. Skin blood flow and temperature responses to thermal stress

Controlling skin blood flow is one of the primary mechanisms by which the human thermoregulatory system maintains thermoneutrality (a body core temperature of approximately 37 °C). Under warm conditions, thermoregulatory vasodilation can boost skin blood flow to 6−8 L/min [19], representing as much as 60 % of total cardiac output [20]. In cold conditions, thermoregulatory vasoconstriction can limit the skin blood flow to nearly zero. Slight changes in skin blood flow can result in large changes in heat transfer to the environment; an 8 % increase in skin blood flow over the entire body doubles the body's heat transfer to the environment. During cold stress the vasoconstrictor system quickly activates and reduces blood flow. After removal of the cold stress, the skin blood flow quickly returns to the normothermia conditions. Vasodilation begins when the body core temperature rises above a temperature threshold. Blood flow rates to different areas of the skin differ under vasoconstriction or vasodilation. Such behavior can be explained by the cardiovascular territories (i.e., regions supplied by specific arteries) supplying blood. In addition, the distribution of cutaneous vessels is not uniform across the body.

Changes in skin blood flow cause skin temperature to vary, especially in exposed parts of the body. Areas with higher density of vessels enable higher blood circulation, resulting in larger temporal variation in skin temperature. Capturing accurate skin temperatures from different cardiovascular territories allows us to characterize the thermoregulatory response to heat and cold stresses. We then use the time-series of temperature measurements to infer the thermoregulation state and estimate thermal comfort. Past research relating thermal sensations to skin temperatures and skin temperature differences indicates that head (forehead, cheek, or nose) and hand skin temperatures can be used to predict both "warm" and "cool" thermal states with over 90 % accuracy [12,21,22].

3. Development of the radiometric sensing hardware and software

We began by assessing the accuracy with which TIR cameras can measure occupant surface temperatures within a narrow field of view, then evaluated techniques to measure surface temperature distribution over a wide field of view with such cameras. We explored options for machine-vision software to identify body parts of interest—e.g., face, nose, and hands—in a color image, then developed and tested a series of hardware-software systems for the radiometric measurement of face, nose, and/or hand skin temperatures.

3.1. Thermographic measurement accuracy

To assess the accuracy of absolute and differential TIR skin temperature measurements, we performed a suite of experiments in an office setting that examined major possible sources of error, including distance, viewing angle, temporal drift in camera response, and ambient thermal reflections. We assessed accuracy primarily by comparing radiatively sensed temperatures to an active temperature reference (ATR—a surface of known thermal emittance whose temperature is measured with a contact thermometer, and controlled with an electric heating circuit). We also performed some measurements with a passive temperature reference (PTR—a surface of known thermal emittance whose temperature is measured with a contact thermometer, but is not regulated).

Trials were conducted with two uncooled-microbolometer thermal cameras: a high-resolution FLIR SC660 with a 45° wide-angle lens (640 by 480 pixels; spectrum 7.5–13 μm; absolute accuracy ± 1 °C or 1 % of range; differential accuracy 45 mK @ 30 °C; field of view 45° × 34°) and a medium-resolution FLIR A315 with a 45° lens adapter (320 by 240 pixels; spectrum 7.5–13 μm; absolute accuracy ± 2 °C or 2 % of range; differential accuracy 50 mK @ 30 °C; field of view 45° × 34°). The FLIR SC660 was used to radiometrically measure the temperature of a CI Systems SR-80-4A Infrared Blackbody ATR (10 cm × 10 cm; thermal emittance 0.95 ± 0.02; temperature accuracy ± 0.05 °C; spatial uniformity ± 0.01 °C) set to 35 °C, while the FLIR A315 was used to radiometrically measure the temperature of a Thermoworks IR-500 Portable Infrared Calibrator/Blackbody ATR (diameter 58 mm; thermal emittance 0.95; temperature accuracy ± 0.8 °C; temperature stability ± 0.1 °C) set to 55 °C.

3.1.1. Distance error

Distance trials with the high-resolution camera viewing the 100 cm^2, 35 °C ATR at distances 1–34 m and with the medium-resolution camera viewing the 26 cm^2, 55 °C ATR at distances 1–4 m, both at normal incidence to the ATR surface, found that increasing distance induced an almost linear decrease in the thermographic measurement of ATR temperature, reaching an underreporting of 1 °C at 10 m for the former and at 3 m for the latter. We attribute this to reduced atmospheric transmittance (scattering and absorption of the thermal infrared radiation) and effective loss of resolution at large distances.

3.1.2. Angle error

Angle trials with the high-resolution camera viewing its ATR at 1 m and with the medium-resolution camera viewing its ATR at 2 m, both conducted at incidence angles 0–75°, showed a decrease in thermographic measurement of ATR temperature when the incidence angle (that between the ATR surface normal and the camera-surface line) increased. Underestimates were minor (i.e., within a few tenths of a Kelvin) at angles < 40° but grew rapidly at larger angles. For example, at an extreme incidence angle of 75°, the underestimate reached 1 °C for the high-resolution camera and 11 °C for the medium-resolution camera. We attribute these discrepancies to the directional nature of emittance and the effective loss of resolution at oblique angles.

3.1.3. Temporal drift

Stray heat within an uncooled microbolometer (an array of pixels whose resistances are altered by absorption of incident thermal radiation) induces drift in radiometric temperature measurement. To compensate, the camera periodically initiates a non-uniformity correction (NUC) operation in which a metal shutter or "flag" of known temperature and thermal emittance is moved in front of the sensor. To explore the effects of sensor drift on TIR measurement accuracy, we used the high-resolution camera to collect time series of images of the 35 °C ATR temperature with NUC suppressed and with NUC active. Immediately after the camera was turned on, TIR measurements of the ATR drifted 0.6 °C without NUC and 0.4 °C with NUC over 15 min. After warming up for 1 h, drift was <0.13 °C without or with NUC. We found the most effective tool for mitigating the effect of sensor drift to be a post-processing image-wide offset correction via an external temperature reference. The key takeaway from this suite of experiments, however, was that regardless of how much the absolute TIR measurements were influenced by uncontrollable sensor drift, differential TIR measurements remained steadfastly stable (<0.1 °C) and largely immune to drift.

3.1.4. Motion blur

The planned TIR camera application requires wide-view real-time assessment of temperature signals in an indoor environment. The field-of-view (FOV) angle desired could be well north of the 40–50° offered today by affordable TIR camera packages. Several approaches could yield FOVs exceeding 90° (semi-hemispherical) or perhaps even 180° (hemispherical):

1. A single stationary TIR sensor outfitted with wide-angle optics that provide the desired FOV
2. A stationary system of multiple sensors and optics that collectively cover the desired FOV
3. A single TIR camera (of modest FOV) capable of rotation over the desired FOV via a pan-and-tilt mount

Geometric distortion would make Approach 1 difficult and component cost rules out Approach 2, at least in the current market where a low-resolution TIR camera (160 by 120 pixels) costs about US$165. To evaluate Approach 3, we assessed the effects of camera motion on thermography to determine whether a thermal camera must be stationary to take an accurate image.

The long time constant of an uncooled microbolometer sensor—on the order of tens of milliseconds [23,24]—makes its output subject to blurring when either the camera or the subject moves. This, combined with the fact that individual pixel outputs from the microbolometer focal plane array are read out sequentially rather than instantaneously [23], results in TIR motion blur presenting as multiple layers of distortion: first, spreading and blending of temperature signals over adjacent pixels; and second, geometric skewing of the subject based on the intersection of pixel readout direction and frame movement. Also, the United States limits the export of thermal cameras with frame rates exceeding 9 Hz [25].

We evaluated motion-based distortion by placing the medium-resolution camera on a rotating tripod roughly 1 m away from the 10 cm × 10 cm CI Systems SR-80-4A Infrared Blackbody ATR set to 35 °C. With the TIR camera recording at a frame rate of 9 Hz, we rotated the assembly from right to left around the pivot point of the tripod mount at five different rotational speeds between 1 and 16 revolutions per minute (RPM). We extracted the frames centered on the ATR plate and compared their likenesses to that of a still TIR image. The average pixel readouts within equivalently sized square selections were calculated to serve as a measure of TIR temperature errors introduced by motion blur.

As expected, the degree of blur increases with increasing rotational speed. At 1 RPM, the extracted frame shows no geometric distortions, insignificant motion blur, and the apparent TIR reference plate temperature is equivalent to that of the still image. At the highest rotational speed tested (16 RPM), the extracted frame instead shows significant signal blurring along the direction of movement as well as noticeable geometric distortion in the shape of the subject, making the square reference emitter plate resemble a rhombus. The apparent TIR temperature of the plate under these circumstances falls nearly 1.5 °C below the known setpoint value.

3.2. Machine vision

It is difficult to outline the face, nose, and hands—hereinafter, "body parts"—in a thermal image without supporting data from a

contemporaneous color image because (a) images from thermal cameras typically have far fewer pixels than those from color cameras; (b) the temperature-based edges in a thermal image are typically fuzzier than the color-based edges in a color image; and (c) it can be difficult to distinguish warm body surfaces from adjacent non-human warm surfaces in a thermal image.

A typical complementary metal–oxide–semiconductor (CMOS) color image sensor (say, 2,500 × 2,000 pixels) offers about 250 times the spatial resolution of a low-resolution microbolometer thermal image sensor (160 × 120 pixels) at about one-tenth the cost. Therefore, we sought to (a) recognize body parts of interest such as the face, nose, and hands, in a color image of the human subject; (b) register (align) the color image with a thermal image of the same subject; and (c) overlay the color-image feature outlines onto the thermal image to locate the body parts in the latter.

3.2.1. Feature recognition in color image

In a companion study that considered only facial feature recognition [26], some of the current authors compared the accuracy and speed of three open-source machine-vision algorithms—Bulat's Face Alignment [27,28], InsightFace [29,30], and FaceNet [31,32]—to find in a color image the facial landmarks needed to infer the locations of the subject's nose, cheeks, and forehead. They selected Bulat's Face Alignment because it was twice as fast as InsightFace and much more accurate than FaceNet.

To locate facial features and hands in a color image, we tested three machine-vision algorithms in the current study: the 2D Face Alignment Network (2D-FAN) [33], an open-source, face-only algorithm scheme closely related to Bulat's Face Alignment [27,28] that we would augment with a hand-segmentation (recognition) model trained with the EgoHands dataset [34,35]; OpenPose (face, hands, and body pose; 70 facial, 2 × 21 hand, and 25 pose keypoints; one or more subjects; free for non-commercial use only) [36,37]; and MediaPipe Holistic (face, hands, and pose; 468 facial, 2 × 21 hand, and 33 pose keypoints; one subject; open source) [38,39].

2D-FAN was fast and accurate (as expected from our experience with Bulat's Face Alignment) but our hand recognition model based on EgoHands was not successful. OpenPose was slower than 2D-FAN but could accurately locate facial, hand, and pose keypoints even with multiple subjects in the image. MediaPipe Holistic quickly and accurately located facial, hand, and pose keypoints. We did not quantitively compare the performance of MediaPipe to that of OpenPose because our goal was to identify and explore, rather than score, machine-vision tools that could locate a subject's nose, face, and hands in a color image. However, other researchers have quantitatively compared the speed and accuracy of these two algorithms [40–42].

Using keypoints returned by MediaPipe Holistic, we connect 36 facial keypoints to outline the face, 26 facial keypoints to outline the nose, 12 hand keypoints (each) to outline the left and right hands, and 7 hand keypoints (each) to outline the left and right palms. We also connect facial keypoints to outline other regions that might be of interest, including the left and right eyes (17 keypoints each) and the portion of the face that could be obscured if the subject wears eyeglasses (36 keypoints).

3.2.2. Transformation of feature shapes from color image to thermal image

Substantial attention has been paid to aligning color and thermal images [43–46] within the mature field of image registration [47–49]. We explored two techniques for mapping the body feature shapes, such as face, nose, and hand outlines, from the color image into the thermal image.

Edge alignment. Prior work by some of the authors [26] cropped and downscaled the color image to the resolution of the thermal image (160 × 120 pixels), used Canny edge detection [50] to find edges in the color and thermal images, applied phase correlation [51] to identify the pixel shift (characteristic x and y offsets) between these edges, and then used this shift to translate feature keypoints from the color image to the thermal image. When in the current study we applied this technique to color and thermal images of a human subject in an office, we found that the phase correlation would occasionally yield poor registration (an inaccurate shift) if the dominant edges were located in the subject's background rather than on or around the subject's body. This was most likely to occur in spaces with poor lighting or strong background reflections, such as those from glass walls. It may be possible to detect and ignore instances of inaccurate edge-based registration by eliminating shifts that lie outside of an acceptable range determined by manual calibration.

Pose keypoint mapping. We determined that MediaPipe Holistic could locate pose keypoints in a *thermal* image. After lengthy experimentation with different sets of pose keypoints, we chose to create two least-squares homographies (planar transformation matrices) with the OpenCV [52] function findHomography():

- An eye-ear-shoulder homography based on six eye keypoints (left_eye_inner, left_eye, left_eye_outer, right_eye_inner, right_eye, right_eye_outer), two ear keypoints (left_ear, right_ear), and two shoulder keypoints (left_shoulder, right_shoulder) in the color image and in a thermal image upscaled to the height of the color image
- A hand homography based on eight hand keypoints (left_wrist, left_pinky, left_index, left_thumb, right_wrist, right_pinky, right_index, right_thumb) in the image pair

We then use the OpenCV function perspectiveTransform() to apply the eye-ear-shoulder homography to the face, nose, eye, and eyeglasses outlines, and to apply the hand homography to the left and right hand and palm outlines.

We found that this scheme will occasionally fail (inaccurately transform the outlines of the face, nose, or hands) if MediaPipe has low confidence in its detection of pose landmarks in the thermal image, or if the transformation matrix is ill-conditioned—i.e., exhibits a high condition number indicating a large change in output for a small change in input. Therefore, we manually reviewed a large set of color-thermal image pairs to identify for each homography (eye-ear-shoulder, hand) thresholds for (a) the minimum detection confidence of the thermal image pose and (b) the maximum transformation condition number, that together yield high (e.g., 90 % or 95 %) confidence that the transformation is accurate. Our quality-assurance process rejects any transformation if the detection confidence of the thermal image pose is too low or if the condition number of the transformation matrix is too high.

3.3. Radiometric sensing hardware

Our sensor assembly includes a pair of "tracking" narrow-view camera sensors, one color and one TIR, on a movable platform; a wide-view stationary color camera; an onboard passive temperature reference; a laser pointer; and a server computer (Fig. 1). This small unit (25 cm L × 10 cm W × 12 cm H) can be mounted on a tripod to view a single occupant, or on a wall or ceiling to view multiple occupants (one at a time).

- **Platform**—a two-axis computer-controlled motorized pan-tilt platform (Pimoroni Pan-Tilt HAT) whose surface normal can span at least a full hemisphere.
- **Tracking cameras**—a pair of platform-mounted, minimally separated narrow-view cameras, comprising a CMOS color camera sensor (OmniVision OV5647; 2,592 × 1,944 pixels; 54° × 41° FOV)[1] and a

[1] A lower-resolution color camera would be fine since we operate it at only 960 × 720 pixels.

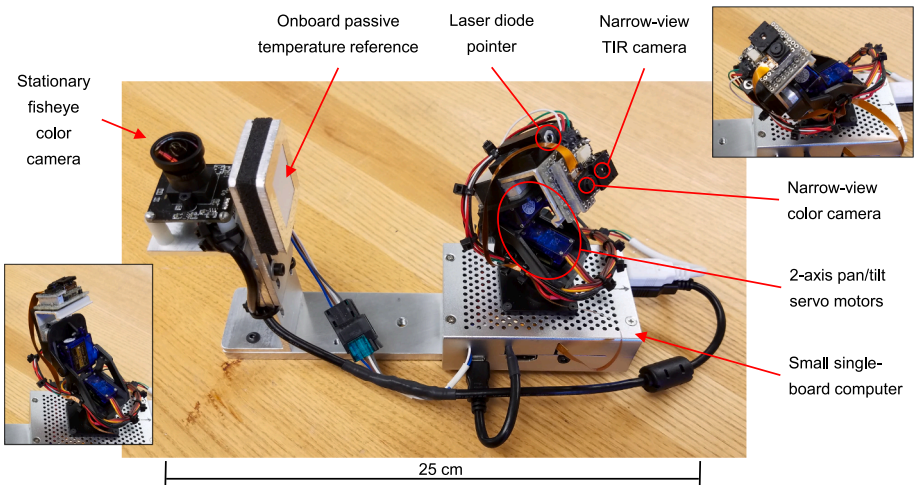

Fig. 1. Sensor hardware assembly, including (clockwise from left) the stationary camera; the onboard PTR; the laser pointer, narrow-view TIR camera, and narrow-view color camera on the pan-tilt platform; and the server computer. Images inset at left and right help illustrate the hemispherical range of angles attainable for the cameras on the pan-tilt platform.

microbolometer TIR sensor (FLIR Lepton 3.5; 160 × 120 pixels; 51° × 49° FOV).

- **Stationary camera**—a stationary fisheye color camera (Sony IMX291; 1,945 × 1,109 pixels, 175° FOV) near the platform that provides a wide overhead view of the room and its occupants when the assembly is installed high in the room, facing downward.
- **Onboard passive temperature reference (PTR)**—an insulated high-thermal-conductance metal plate near the platform with both high thermal emittance (TE) and low TE surfaces whose contact temperature is continuously measured with a rear-mounted digital thermistor (Microchip Technology MCP9808). The PTR is used to calibrate temperatures measured with the TIR camera.
- **Laser diode pointer**—a low-power color laser mounted on the platform, adjacent to the tracking cameras, that when activated shows in the stationary camera's image the approximate center of the tracking cameras' field of view.
- **Server computer**—a single-board computer (Raspberry Pi 3) running the server software detailed in Section 3.4.1.

Note that our sensor assembly neither requires nor incorporates the FLIR SC660 and FLIR A315 thermal cameras used in our initial exploration of thermographic measurement techniques.

We also place a second, larger PTR (Fig. 2) near the subject rather than on the sensor assembly. It includes a heated wire around its perimeter to make it easier to find in the thermal image. This near-subject PTR may be removed in the next generation of our device.

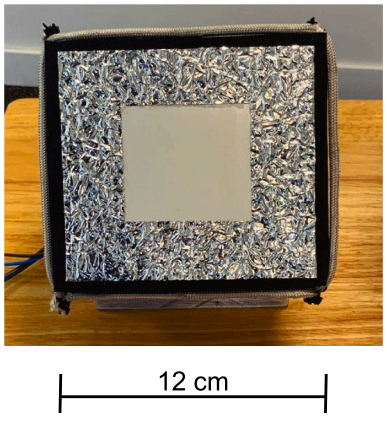

Fig. 2. Near-subject PTR with heated wire around its perimeter.

Our sensing hardware expands on that developed by some of the authors in a previous study [26] by adding the pan-tilt platform, onboard PTR, and laser diode pointer. It employs the same microbolometer (FLIR Lepton 3.5) but different color sensors.

3.4. Radiometric sensing software

3.4.1. Server

Our server software[2] is a set of Python 3 scripts that orients the platform, tracking cameras, and laser pointer; captures images from the three cameras and contact temperatures from one or more PTRs; and transmits the platform orientation, images, and PTR temperatures to the client computer. On client request, it

- points the tracking cameras in a specified direction using the pantilt-hat library;
- captures a narrow-view color image from the tracking color camera using the imutils library;
- captures a narrow-view thermal image from the tracking thermal camera using the purethermal1-uvc-capture library;
- measures the temperatures of digital thermistors in the on-assembly PTRs and near-subject PTR using the adafruit_mcp9808 library; and
- transmits the images and temperatures to the client using the imageZMQ library, employing socket communications to receive requests and confirm delivery.

3.4.2. Client

Our entirely new client software includes radiometric-sensing, thermal-sensation, and control elements. Its sensing component

- requests from the server tracking-cameras images and near-subject PTR temperature for a single subject (occupant);
- uses the PTR temperature to calibrate absolute temperatures in the thermal image;
- employs MediaPipe Holistic to locate body parts of interest (e.g., face, nose, hands) in the color and thermal images as described in Section 3.2;

[2] Our server software descends from what was called "client" software in the "iEye" system developed by one of the authors [26,53]. It has been refactored to act as the server rather than the client in the transmission of images and temperatures.

- computes rolling time-median skin temperature statistics within each body region of interest; and
- calculates skin temperature differences of interest, such as (face_p90 – nose_median) and (face_p90 – hand_max), where p90 denotes 90th percentile.

4. Development of the thermal sensation model

We conducted many human-subject thermal-sensation trials in cool, neutral, or warm environments, measuring skin temperatures and recording thermal sensation votes. We then trained random-forest classification machine-learning models—hereinafter, "classifiers"—to estimate thermal sensation from skin temperatures or skin-temperature differences.

In a series of trials detailed by some of the authors in Ref. [54], we exposed a total of 34 human subjects to cool, neutral, or warm environments while measuring their skin temperatures with a color-thermal camera system comprising the FLIR A315 medium-resolution TIR camera detailed in Section 3.1, a FLIR Blackfly S BFS-PGE-50S5C color camera (2,448 × 2,048 pixels), and a custom thermal imaging capture & analysis tool (TI-CAT) developed by the authors from MoviTHERM. We simplified thermal sensation votes reported on a nine-point scale (−4 = very cold, −3 = cold, −2 = cool, −1 = slightly cool, 0 = neutral, +1 = slightly warm, +2 = warm, +3 = hot, +4 = very hot) by classifying sensations lower than −1 as cool (simplified −1), between −1 and +1 (inclusive) as neutral (simplified 0), and higher than +1 as warm (simplified +1). We then trained each classifier to predict the most-probable value of simplified sensation (−1, 0, or +1) based on one of the 16 temperatures or temperature differences involving nose, cheek_max, cheek_median, hand_max, and/or hand_median temperatures.

We subsequently trained additional classifiers using skin temperatures and skin-temperature differences measured with the current system described in Section 3.3 and Section 3.4, most of which were recorded during the same human-subject trials reported in Ref. [54]. The new classifiers use time-median temperatures or temperature differences based on face_p90, face_median, nose_median, hand_max, and/or hand_median.

The sensing component of our client software applies these new classifiers to predict both most-probable simplified sensation (−1, 0, or +1) and probability-weighted simplified sensation (a continuous value between −1 and +1 equal to −1 × cool_probability + 0 × neutral_probability + 1 × warm_probability). Since a discrete (integer) value could result in undesired cycling, the probability-weighted simplified sensation is used to guide our AC-control algorithm.

5. Development of the AC control algorithm

We implemented a standard discrete-time proportional-integral (PI) control logic, using a 1 s sampling time, to close the loop between sensation detection and AC operation. The PI controller receives a tracking error defined as the absolute difference between the 10-second moving average of the predicted thermal sensation (probability-weighted simplified sensation) and the target sensation value, and calculates the thermostat setpoint using the PI control logic. After conducting some tests, the proportional gain and integral time constant were set at 4 (°F per unit change on the simplified sensation scale) and 15 s, respectively, with a 1 s sampling time and a 10 s moving-average window. The controller also incorporates a standard anti-windup algorithm, which resets the integral error when the calculated thermostat setpoint, serving as the control variable in our case, exceeds a certain threshold specified below.

The adjustable target sensation value plays a crucial role in balancing energy consumption and thermal comfort. For our investigation, a target sensation value of 0.5 was selected to strike a balance between neutrality and warmth on a simplified 3-point sensation scale, as determined by a parametric study.

To tailor the PI control logic, adjustments were made to increase the penalty for tracking errors and expedite responses in cases of room discomfort, such as excessive warmth or coolth. This adjustment involved amplifying the tracking error using a specified function. The output of the AC control, which is the room air temperature setpoint, is bounded by lower and upper thresholds, specifically 70 °F [21.1 °C] and 78 °F [25.6 °C], respectively. If unable to detect skin temperature after 10 min, we assume the room is unoccupied and revert to a predefined schedule of room temperature setpoints (either 76 °F [24.4 °C] or 80 °F [26.7 °C]) until the skin temperature is again detected. An anti-windup logic handles control saturation and mode transitions.

Once the thermostat setpoint is established, whether from the AC control or the predefined schedule based on occupancy detection, it is transmitted to the AC system via the AC adapter. This adapter serves as an intermediary, conveying setpoint instructions from our client software to the AC system. The communication was established using the AC adapter's Application Programming Interface (API), enabling the modification of the setpoint within the AC system via a "PUT" request.

6. System integration

Our complete system, named "Goldilocks", integrates the sensing hardware and software (Section 3.3 and Section 3.4) with the thermal sensation classifiers (Section 4) and control algorithm and hardware (Section 5). It comprises the

- sensor assembly (pan-tilt platform, color and thermal cameras, laser pointer, onboard PTR, and server computer);
- external hardware (client computer, AC adapter, near-subject PTR, network switch for client–server communication, survey device for human subject trials, and air temperature & relative humidity sensors for system performance trials);
- server software (Python 3) that orients the tracking cameras and transmits images and PTR contact temperatures to the client; and
- client software (Python 3) that assesses occupant skin temperature distribution, estimates occupant thermal sensation, and controls AC operation.

The client software currently runs on a Windows 10 laptop PC (HP OMEN 15-dh1054nr) but could be ported to a single-board Linux computer with GPU, such as the NVIDIA Jetson Nano.

Fig. 3 illustrates Goldilocks in operation. In panel A (top left), MediaPipe Holistic locates 468 facial, 2 × 21 hand, and 33 pose keypoints in a color image containing the subject (one of the authors). In Panel B (bottom left), a subset of these keypoints is used to outline the subject's face, nose, eyes, eyeglasses region, hands, and palms in that color image. In Panel C (top right), MediaPipe Holistic locates 33 pose keypoint in the corresponding thermal images, along with 21 keypoints in the subject's left hand that we ignore. In Panel D (bottom right), the outlines from the color images are mapped into the thermal image using the planar transformations described in Section 3.2.2. Panel D also shows skin-temperature statistics derived from the thermal image, along with values of probability-weighted simplified thermal sensation predicted from various skin-temperature statistics (e.g., FN = face and nose temperatures; FNH = face, nose, and hand temperatures; FN_dif = differences between face temperature and hand temperature; FNH_dif = differences between face, nose, and hand temperatures).

7. Demonstration

7.1. Experiment

We tested our system in a conference room in an office building near Houston, Texas from August 2022 to January 2023 [55]. Since the room was in the building's interior zone (no exterior windows or walls) and the space outside the room was conditioned to 72 °F [22.2 °C], we

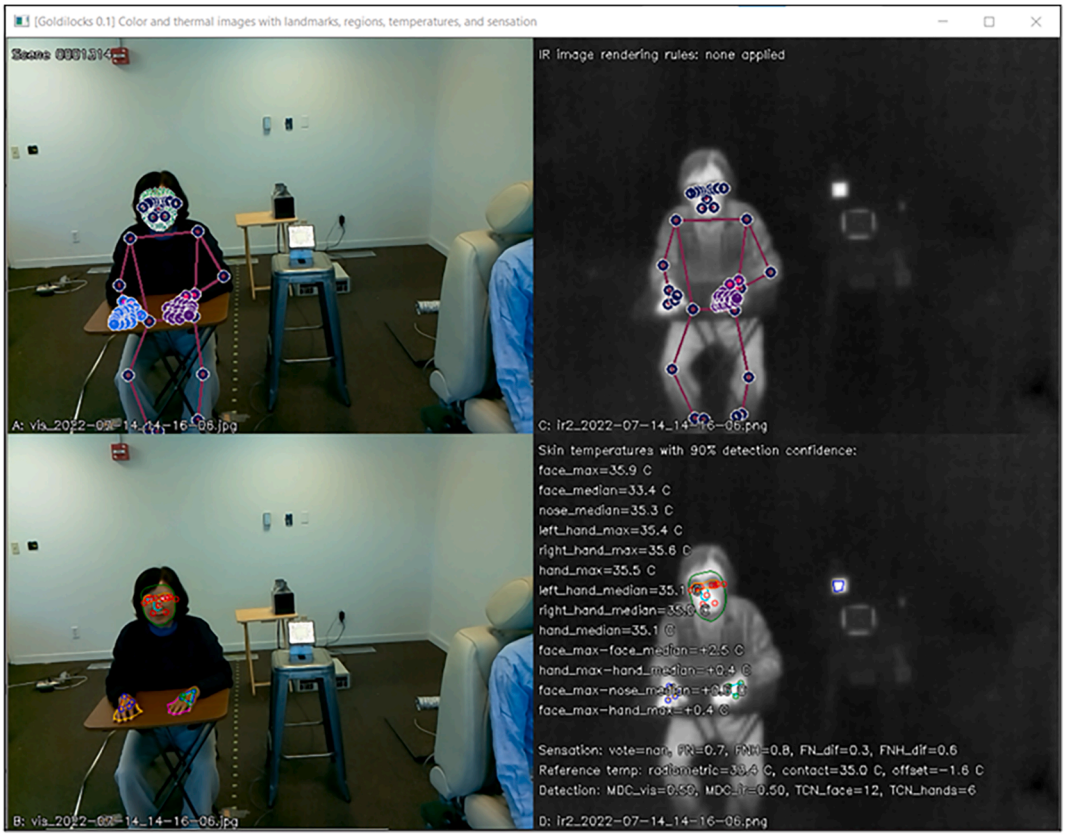

Fig. 3. Operation of the "Goldilocks" client software showing (A) face, hand, and pose landmarks in the color image; (B) facial pose landmarks, face outline, nose outline, eye outlines, eyeglasses outline, hand outlines, and palm outlines in the color image; (C) pose landmarks in the thermal image; and (D) face landmarks, hand landmarks, and transformed face, nose, eye, eyeglasses, and hand outlines in the thermal image. The near-subject PTR can be seen in panels A and B, and the resistance-heated border of the PTR can be seen in panels C and D. The black box in panels A and B and the corresponding bright patch in panels C and D is an ATR (high-emittance surface fixed at 35 °C) that is not integral to our system. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

operated a 1,500 W electric resistance heater in the room during trial hours (weekdays 09:00–17:00) to simulate a summer heat load from conduction, infiltration, and solar heat gain through the building envelope.

We implemented three control strategies in the room during trial hours:

- **Conventional Control:** fixed setpoint of 72.0 °F [22.2 °C], equal to the setpoint throughout the office building during occupied hours
- **TIR A Control:** thermostatic setpoint regulated by our system with an unoccupied (empty room) setpoint of 80 °F [26.7 °C]
- **TIR B Control:** same as TIR A Control but with an unoccupied setpoint of 76 °F [24.4 °C]

The room was occupied by at most one person at a time. Over the course of the five-month trial we had about 20 unique subjects, predominately male and mostly 20–50 years old. We did not choose the test subjects. The participants were the occupants of the office, and they used the room as a workspace or conference room as usual. Following a research protocol approved by the University of California at Berkeley Committee for the Protection of Human Subjects (IRB-2020–12-13922), each participant self-reported at will (a) time spent in the room (<5 min, 5 min–1 h, >1 h), (b) thermal sensation (seven-point scale), and (c) comfort vote (yes/no). Participants were uncompensated and fully informed about the nature of the study.

We collected occupation times, thermal sensation votes, and comfort votes with a survey tablet, and measured the AC's air-cooling rate (rate of heat removal) as the product of the air mass flow rate (known from the unit's speed setting) and the air's enthalpy drop from inlet to outlet (based on inlet and outlet measurements of air temperature and humidity).

To ensure that the camera correctly detected the face and hands of the occupant, the subject was asked to sit in a designated chair and work as usual. The camera was positioned so that the horizontal distance between the camera and the designated chair was about 2.0 m and the height of the camera was about 1.5 m above floor level (Fig. 4b). Fig. 5 shows a representative color image and the feature outlines of a trial subject as captured by the sensor-controller.

7.2. Results

On the seven-point scale, 35 % of the sensation votes under Conventional Control that set the room temperature to 72.0 °F (22.2 °C) were "Neutral" and the rest were either "Slightly cool", "Cool", or "Cold", with no "Warm" or "Hot" votes (Fig. 6a). This indicates that the Conventional Control overcooled the occupants.

TIR A Control made the occupants more thermally neutral than Conventional Control, reducing the mean air-cooling rate (rate of heat removal by the AC) by 18 % when the room was occupied (Fig. 7b) and by 62 % over the entire trial period (room occupied or unoccupied). The latter result indicates that our system worked as an occupant sensor and saved energy by raising the setpoint when the room appeared empty. However, some occupants complained that the room was hot and uncomfortable, especially immediately after entering the room, because the conference room was substantially warmer than the rest of the office. Although the percentage of "Neutral" votes under TIR A Control was

Fig. 4. Demonstration configuration, showing photos of (a) the electric space heater used to simulate a summer heat load; (b) the conference room (about 4.3 m L × 4.3 m W × 3.7 m H), in which the subject's chair is marked with an asterisk; (c) a close-up of the sensor-controller mounted on a tripod; (d) the survey tablet used to collect thermal sensation and comfort votes; and (e) a temperature/relative humidity logger mounted above the room's conventional thermostat.

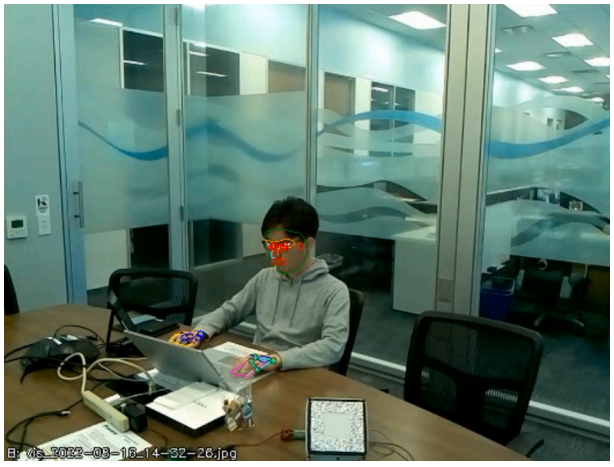

Fig. 5. Color image and feature outlines of a trial subject (one of the authors) in the conference room as captured by our sensor-controller. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

higher than that under Conventional Control, those of "Warm" and "Hot" votes, which are considered outside the thermally comfortable range, were 6 % and 12 % respectively (Fig. 6a). There was no significant difference in thermal comfort votes between Conventional Control and TIR A Control (Fig. 6b). The higher unoccupied setpoint temperature increased the time needed to cool the air to a comfortable temperature, which could reduce thermal satisfaction.

TIR B Control made the occupants more thermally neutral than Conventional Control, lowering the mean air-cooling rate by 42 % when the room was occupied (Fig. 7c) and by 18 % over the entire trial period. On a seven-point scale, 56 % of the occupants, the highest fractions among the three control strategies, responded "Neutral", and over 90 % of the sensation votes were either "Slightly cool", "Neutral", or "Slightly warm". TIR B Control also yielded 76 % "Comfortable" votes, higher than TIR A Control (62 %) and Conventional Control (61 %).

7.3. Discussion

The machine-learning model that we used in this field study to predict occupant thermal sensation was trained using skin temperatures and thermal sensation data obtained in cold environments (ambient air temperature ≤ 20 °C [68.0 °F]), neutral environments (around 24 °C [75.2 °F]), and hot environments (≥30 °C [86.0 °F]). However, the actual room air temperature was 22–26 °C [71.6–78.8 °F] and did not enter the cold or hot regimes (≤20 °C or ≥30 °C) employed in model training. This led to overprediction of the "Neutral" thermal sensation and resulted in a suboptimal prediction accuracy. Therefore, it is necessary to conduct additional sensation trials in warm and cool (rather than hot and cold) environments to retrain and improve the model. The sensation votes and skin-temperature measurements collected in the conference-room trial will provide the initial data for this effort, to be supplemented by the participation of a wider demographic.

There may be cultural expectations for overcooling from building designers, operators, and occupants. Without access to real-time feedback, designers and engineers assume that a low ambient temperature is comfortable for occupants. For occupants, long-term exposure to overcooling may create expectations for overcooling, because people adapt to their environments physiologically and psychologically [56]. With feedback-providing devices like that described in our study, building designers and operators can set the ambient temperature based on

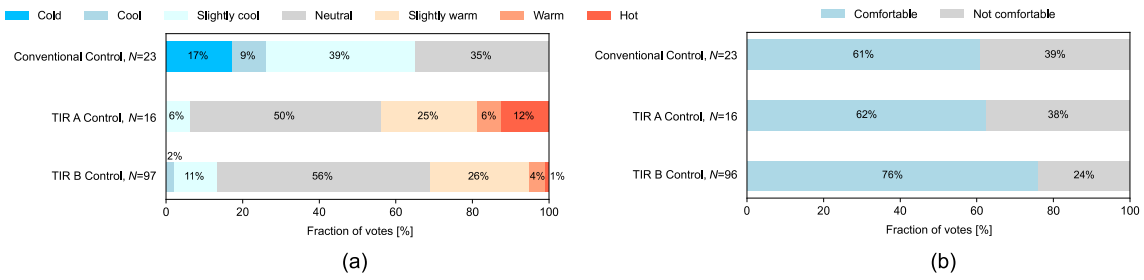

Fig. 6. Distributions of (a) thermal sensation votes and (b) thermal comfort votes under each control algorithm. *N* represents the total number of votes.

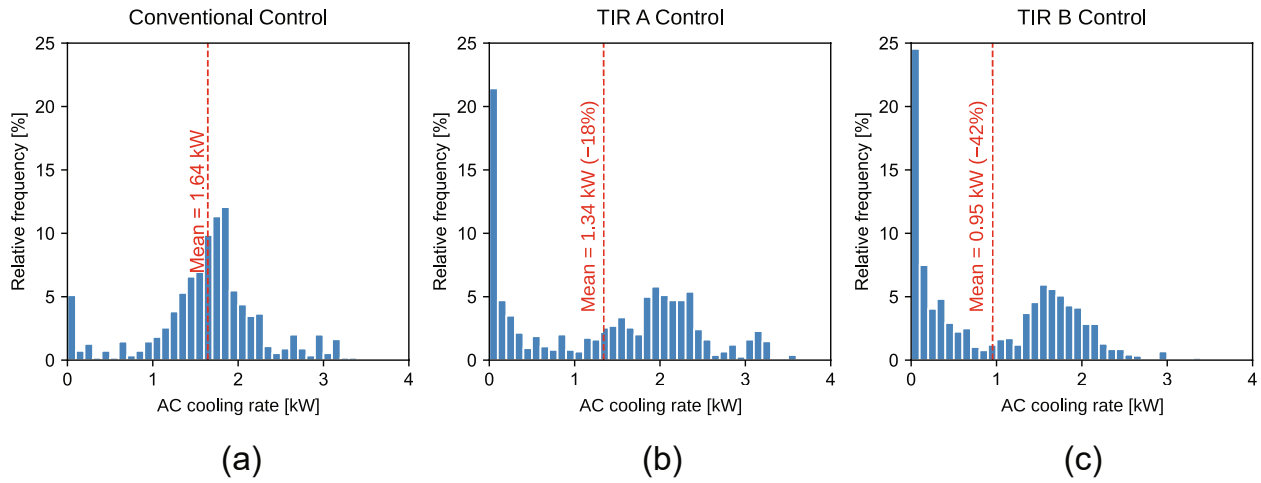

Fig. 7. Distributions of the air-cooling rate (rate of heat removal by the AC) when the conference room was occupied under (a) Conventional Control, (b) TIR A Control, or (c) TIR B Control. TIR A Control and TIR B Control reduced the mean air-cooling rate by 18% and 42%, respectively compared to Conventional Control.

occupant response rather than assumptions. We posit that occupant satisfaction will be improved, and that expectations for overcooling can be reduced. For example, a large-scale study by two of the authors cut occupant overcooling discomfort complaints in office buildings by 50 % through feedback-enabled control of supply airflow rates [5,6].

8. Future development

We plan to upgrade our system to detect and act upon the thermal sensations of multiple occupants, and to improve the assessment of the thermal sensations of individual occupants.

The upgraded server software will locate anonymous occupants in the wide-view image and aim the tracking cameras at each occupant in turn. The upgraded client software will improve skin-temperature sensing, thermal sensation prediction, and AC control. Specifically, it will

- calculate the rolling time-median time derivatives of skin temperatures and skin-temperature differences to assess how skin temperatures are evolving;
- train new classifiers or dynamic models using both rolling time-median temperatures or temperature differences and their time derivatives;
- train new classifiers or dynamic models for a large and more diverse set of human subjects under milder thermal conditions more representative of offices or other indoor public spaces;
- provide an option to generate locally trained classifiers by using the survey device to record occupant thermal sensation votes while the invention is operating using pre-trained classifiers;
- direct the server software to locate multiple occupants in the space and apply a modified sensing algorithm (same as above but using the onboard PTR for thermographic calibration) to assess the thermal sensation of each occupant;
- calculate ensemble (multi-occupant) thermal sensation metrics based on the thermal sensations of individual occupants—e.g., minimum, maximum, mean, and median;
- use one or more of the ensemble sensation metrics to regulate the AC control;
- automatically learn the site-specific occupancy schedule via a statistical learning technique and incorporate it into the AC control for better savings and comfort; and
- investigate sensitivities of the optimal control parameters (such as P and I gains, default thermostat set point for the non-presence mode) for other sites.

We may also explore the application of our system to the prevention of overheating in winter.

9. Summary

The human thermoregulation system uses skin blood flow to adjust its heat balance with the thermal environment to maintain thermoneutrality at the brain and the body core. This causes skin temperature to vary, especially in exposed parts of the body. The distribution of skin temperature can be used to predict warm, neutral, and cool thermal states.

To improve occupant comfort and save energy in buildings, we have developed “Goldilocks”, a closed-loop AC sensor-controller that predicts occupant thermal sensation from the thermographic measurement of skin temperature distribution, then uses this information to reduce overcooling (cooling-energy overuse that discomforts occupants) by regulating AC output. Taking measures to protect privacy, it combines TIR and color cameras with machine vision to measure the skin-surface temperature profile.

We began by assessing the accuracy with which TIR cameras can measure surface temperatures within a narrow field of view, then evaluated techniques to measure surface temperature distribution over a wide field of view with such cameras. We explored options for machine-vision software to identify body parts of interest—e.g., face, nose, and hands—in a color image, then developed and tested a series of hardware-software systems for the radiometric measurement of face, nose, and/or hand skin temperatures.

We conducted a series of human-subject thermal-sensation trials in cold-to-hot environments, measuring skin temperatures and recording thermal sensation votes. We then trained random-forest classification machine-learning models to estimate thermal sensation from skin temperatures or skin-temperature differences. To avoid overcooling and save energy, we created a PI control algorithm for the air conditioner that targets an estimated thermal sensation between neutral and warm on a simplified three-point scale (cool, neutral, or warm).

Our sensor-controller includes a sensor assembly (pan-tilt platform, color and thermal cameras, laser pointer, integrated PTR, and server computer); external hardware (client computer, network router, AC adapter, near-subject passive temperature reference, survey device, and air temperature & relative humidity sensors); server software; and client software. The server software orients the tracking cameras and transmits images and PTR contact temperatures to the client. The client software assesses occupant skin temperature distribution, estimates occupant thermal sensation, and controls AC operation.

A demonstration in an office building near Houston, TX showed that our system reduced overcooling. It improved occupant comfort (fraction of comfortable votes) by 15 percentage points and decreased the need for air conditioning, lowering the cooling load by 42 % when the room was occupied.

We plan to upgrade our system to detect and act upon the thermal sensations of multiple occupants, and to improve assessment of the thermal sensation of each individual occupant.

CRediT authorship contribution statement

Ronnen Levinson: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Software, Supervision, Writing – original draft, Writing – review & editing. **Donghun Kim:** Formal analysis, Investigation, Methodology, Software, Supervision, Writing – original draft. **Howdy Goudey:** Investigation, Methodology. **Sharon Chen:** Investigation, Writing – original draft. **Hui Zhang:** Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Writing – original draft, Writing – review & editing. **Ali Ghahramani:** Conceptualization, Formal analysis, Investigation, Methodology, Software, Writing – original draft, Writing – review & editing. **Charlie Huizenga:** Formal analysis, Investigation, Methodology, Software, Writing – review & editing. **Yingdong He:** Formal analysis, Investigation, Methodology, Software. **Akihisa Nomoto:** Formal analysis, Investigation, Software, Writing – original draft, Writing – review & editing. **Edward Arens:** Conceptualization, Funding acquisition, Investigation, Resources, Writing – review & editing. **Ana Álvarez Suárez:** Formal analysis, Investigation, Methodology, Software, Writing – original draft. **David Ritter:** Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Markus Tarin:** Conceptualization, Funding acquisition, Project administration, Supervision, Writing – original draft. **Robert Prickett:** Funding acquisition, Project administration, Resources, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ronnen Levinson reports financial support was provided by Daikin U.S. Ali Ghahramani and Ronnen Levinson have patent #WO2021050369A1: Autonomous comfort systems issued to University of California. Corresponding author Ronnen Levinson is an Associate Editor of *Energy & Buildings*.

Data availability

Data will be made available on request.

Acknowledgements

This study was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, with cost share provided by Daikin U.S.; MoviTHERM; and the Center for the Built Environment (CBE) at the University of California, Berkeley. We thank Haley Gilbert (LBNL subcontractor) for helping coordinate project execution; Alexander Merritt (UC Berkeley) and Sang woo Ham (LBNL) for technical support in trials; Sri Swaminathan, Norman Pennant, and Bevnoty Attia of Daikin U.S. for facilitating the Texas field trial; Chuncheng Piao and Kevin Ninomiya of Daikin U.S. for supporting project execution; and Marina Sofos, Erika Gupta, and Brian Walker of the Building Technologies Office, U.S. Department of Energy for their support and guidance.

References

[1] C. Huizenga, S. Abbaszadeh, L. Zagreus, E.A. Arens, Air quality and thermal comfort in office buildings: Results of a large indoor environmental quality survey, in: Healthy Buildings 2006, Lisbon, Portugal, 2006. pp. 393–397. https://escholarship.org/uc/item/7897g2f8 (accessed August 12, 2023).

[2] L.T. Graham, T. Parkinson, S. Schiavon, Lessons learned from 20 years of CBE's occupant surveys, Build. Cities 2 (2021) 166–184, https://doi.org/10.5334/bc.76.

[3] M.J. Mendell, A.G. Mirer, Indoor thermal factors and symptoms in office workers: findings from the US EPA BASE study, Indoor Air 19 (2009) 291–302, https://doi.org/10.1111/j.1600-0668.2009.00592.x.

[4] T. Parkinson, S. Schiavon, R. De Dear, G. Brager, Overcooling of offices reveals gender inequity in thermal comfort, Sci. Rep. 11 (2021) 23684, https://doi.org/10.1038/s41598-021-03121-1.

[5] E. Arens, H. Zhang, T. Hoyt, S. Kaam, F. Bauman, Y. Zhai, G. Paliaga, J. Stein, R. Seidl, B. Tully, J. Rimmer, J. Toftum, Effects of diffuser airflow minima on occupant comfort, air mixing, and building energy use (RP-1515), Sci. Technol. Built Environ. 21 (2015) 1075–1090, https://doi.org/10.1080/23744731.2015.1060104.

[6] G. Paliaga, H. Zhang, T. Hoyt, E. Arens, Eliminating overcooling discomfort while saving energy, ASHRAE J. (2019) 14–28, https://escholarship.org/uc/item/5t665086 (accessed August 12, 2023).

[7] T. Hoyt, E. Arens, H. Zhang, Extending air temperature setpoints: simulated energy savings and design considerations for new and retrofit buildings, Build. Environ. 88 (2015) 89–96, https://doi.org/10.1016/j.buildenv.2014.09.010.

[8] A. Chen, V.-W.-C. Chang, Human health and thermal comfort of office workers in Singapore, Build. Environ. 58 (2012) 172–178, https://doi.org/10.1016/j.buildenv.2012.07.004.

[9] S.C. Sekhar, Thermal comfort in air-conditioned buildings in hot and humid climates–why are we not getting it right? Indoor Air 26 (2016) 138–152, https://doi.org/10.1111/ina.12184.

[10] US Department of Energy Building Technologies Office, Scout Baseline Energy Calculator, (n.d.). https://scout.energy.gov/baseline-energy-calculator.html (accessed July 30, 2023).

[11] S. Derrible, M. Reeder, The cost of over-cooling commercial buildings in the United States, Energ. Build. 108 (2015) 304–306, https://doi.org/10.1016/J.ENBUILD.2015.09.022.

[12] D. Wang, H. Zhang, E. Arens, C. Huizenga, Observations of upper-extremity skin temperature and corresponding overall-body thermal sensations and comfort, Build. Environ. 42 (2007) 3933–3943, https://doi.org/10.1016/j.buildenv.2006.06.035.

[13] H. Zhang, E. Arens, C. Huizenga, T. Han, Thermal sensation and comfort models for non-uniform and transient environments: Part I: Local sensation of individual body parts, Build. Environ. 45 (2010) 380–388, https://doi.org/10.1016/j.buildenv.2009.06.018.

[14] K. Tweed, Building Robotics Launches New Ecosystem to Scale Human Sensors for Comfier Buildings, (2015). https://www.greentechmedia.com/articles/read/Building-Robotics-Launches-New-Ecosystem-for-Happier-Office-Workers (accessed July 30, 2023).

[15] C. Sugimoto, Human sensing using wearable wireless sensors for smart environments, in: 2013 Seventh International Conference on Sensing Technology (ICST), IEEE, Wellington, New Zealand, 2013, pp. 188–192, https://doi.org/10.1109/ICSensT.2013.6727640.

[16] A. Rowe, Human-in-the-loop Sensing and Control for Commercial Building Energy Efficiency and Occupant Comfort, (2018). https://www.energy.gov/sites/prod/files/2018/06/f52/31691_Rowe_050218-1330.pdf.

[17] K. Chen, Q. Xu, B. Leow, A. Ghahramani, Personal thermal comfort models based on physiological measurements – a design of experiments based review, Build. Environ. 228 (2023) 109919, https://doi.org/10.1016/j.buildenv.2022.109919.

[18] G.C. Bell, B. Storey, M.K. Patterson. Control of Computer Room Air Conditioning using IT Equipment Sensors, Lawrence Berkeley National Laboratory, Berkeley, CA, 2009, https://doi.org/10.2172/982931.

[19] J.M. Johnson, D.W. Proppe, Cardiovascular Adjustments to Heat Stress, in: R. Terjung (Ed.), Comprehensive Physiology, 1st ed., Wiley, 1996: pp. 215–243. https://doi.org/10.1002/cphy.cp040111.

[20] L.B. Rowell, Cardiovascular adjustments to thermal stress, in: R. Terjung (Ed.), Comprehensive Physiology, first ed., Wiley, 2011, pp. 967–1023, https://doi.org/10.1002/cphy.cp020327.

[21] C. Dai, H. Zhang, E. Arens, Z. Lian, Machine learning approaches to predict thermal demands using skin temperatures: steady-state conditions, Build. Environ. 114 (2017) 1–10, https://doi.org/10.1016/j.buildenv.2016.12.005.

[22] A. Ghahramani, G. Castro, B. Becerik-Gerber, X. Yu, Infrared thermography of human face for monitoring thermoregulation performance and estimating personal thermal comfort, Build. Environ. 109 (2016) 1–11, https://doi.org/10.1016/j.buildenv.2016.09.005.

[23] B. Oswald-Tranta, M. Sorger, P. O'Leary, Motion deblurring of infrared images from a microbolometer camera, Infrared Phys. Technol. 53 (2010) 274–279, https://doi.org/10.1016/j.infrared.2010.04.003.

[24] M. Vollmer, K.-P. Möllmann, Infrared Thermal Imaging: Fundamentals, Research and Applications, first ed., John Wiley & Sons, Ltd, Weinheim, 2010. https://doi.org/10.1002/9783527630868.

[25] Infrared Cameras Inc., Thermal Camera Export Restrictions: Exportable and Embargoed Countries for Thermal Infrared Imaging Camera Sales, (n.d.). https://infraredcameras.com/thermal-camera-export-restrictions (accessed July 31, 2023).

[26] A. Ghahramani, Q. Xu, S. Min, A. Wang, H. Zhang, Y. He, A. Merritt, R. Levinson, Infrared-fused vision-based thermoregulation performance estimation for personal thermal comfort-driven HVAC system controls, Buildings 12 (2022), https://doi.org/10.3390/BUILDINGS12081241.
[27] A. Bulat, Face Recognition, (2023). https://github.com/1adrianb/face-alignment (accessed August 5, 2023).
[28] A. Bulat, G. Tzimiropoulos, How Far are We from Solving the 2D & 3D Face Alignment Problem? (and a Dataset of 230,000 3D Facial Landmarks), in: 2017 IEEE International Conference on Computer Vision (ICCV), 2017: pp. 1021–1030. https://doi.org/10.1109/ICCV.2017.116.
[29] J. Guo, J. Deng, InsightFace: 2D and 3D Face Analysis Project, (2023). https://github.com/deepinsight/insightface (accessed August 5, 2023).
[30] J. Deng, Y. Zhou, S. Cheng, S. Zaferiou, Cascade Multi-View Hourglass Model for Robust 3D Face Alignment, in: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), 2018: pp. 399–403. https://doi.org/10.1109/FG.2018.00064.
[31] I. de P. Centeno, MTCNN, (2023). https://github.com/ipazc/mtcnn (accessed August 5, 2023).
[32] K. Zhang, Z. Zhang, Z. Li, Y. Qiao, Joint face detection and alignment using multitask cascaded convolutional networks, IEEE Signal Process. Lett. 23 (2016) 1499–1503, https://doi.org/10.1109/LSP.2016.2603342.
[33] A. Bulat, How far are we from solving the 2D & 3D Face Alignment problem? (and a dataset of 230,000 3D facial landmarks), (2023). https://github.com/1adrianb/2D-and-3D-face-alignment (accessed August 5, 2023).
[34] S. Bambach, EgoHands: A Dataset for Hands in Complex Egocentric Interactions, (n.d.). http://vision.soic.indiana.edu/projects/egohands/ (accessed August 5, 2023).
[35] S. Bambach, S. Lee, D.J. Crandall, C. Yu, Lending A Hand: Detecting Hands and Recognizing Activities in Complex Egocentric Interactions, in: 2015 IEEE International Conference on Computer Vision (ICCV), 2015: pp. 1949–1957. https://doi.org/10.1109/ICCV.2015.226.
[36] G. Hidalgo, Y. Raaj, OpenPose, (2023). https://github.com/CMU-Perceptual-Computing-Lab/openpose (accessed August 5, 2023).
[37] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, Y. Sheikh, OpenPose: realtime multi-person 2D pose estimation using part affinity fields, IEEE Trans. Pattern Anal. Mach. Intell. 43 (2021) 172–186, https://doi.org/10.1109/TPAMI.2019.2929257.
[38] Google, MediaPipe, (2023). https://github.com/google/mediapipe (accessed August 5, 2023).
[39] Google, MediaPipe, Google for Developers (n.d.). https://developers.google.com/mediapipe (accessed August 5, 2023).
[40] A. Saiwa, Openpose vs Mediapipe for Dynamic Vision | Beyond Poses, (2023). https://saiwa.ai/blog/openpose-vs-mediapipe/ (accessed August 5, 2023).
[41] P. Radzki, Detection of human body landmarks - MediaPipe and OpenPose comparison, HearAI (2022). https://www.hearai.pl/post/14-openpose/ (accessed August 5, 2023).
[42] J.-L. Chung, L.-Y. Ong, M.-C. Leow, Comparative analysis of skeleton-based human pose estimation, Future Internet 14 (2022) 380, https://doi.org/10.3390/fi14120380.
[43] J. Klimaszewski, M. Kondej, M. Kawecki, B. Putz, Registration of infrared and visible images based on edge extraction and phase correlation approaches, in: R. S. Choraś (Ed.), Image Processing and Communications Challenges 4, Springer, Berlin, Heidelberg, 2013, pp. 153–162, https://doi.org/10.1007/978-3-642-32384-3_19.
[44] H. Li, W. Ding, X. Cao, C. Liu, Image registration and fusion of visible and infrared integrated camera for medium-altitude unmanned aerial vehicle remote sensing, Remote Sens. (Basel) 9 (2017) 441, https://doi.org/10.3390/rs9050441.
[45] Y. Dong, C. Fei, G. Zhao, L. Wang, Y. Liu, J. Liu, S. Fan, Y. Li, X. Zhao, Registration method for infrared and visible image of sea surface vessels based on contour feature, Heliyon 9 (2023) e14166, https://doi.org/10.1016/j.heliyon.2023.e14166.
[46] J. Han, E.J. Pauwels, P. de Zeeuw, Visible and infrared image registration in man-made environments employing hybrid visual features, Pattern Recogn. Lett. 34 (2013) 42–51, https://doi.org/10.1016/j.patrec.2012.03.022.
[47] L.G. Brown, A survey of image registration techniques, ACM Comput. Surv. 24 (1992) 325–376, https://doi.org/10.1145/146370.146374.
[48] B. Zitová, J. Flusser, Image registration methods: a survey, Image Vis. Comput. 21 (2003) 977–1000, https://doi.org/10.1016/S0262-8856(03)00137-9.
[49] Z. Xiong, Y. Zhang, A critical review of image registration methods, Int. J. Image Data Fusion 1 (2010) 137–158, https://doi.org/10.1080/19479831003802790.
[50] Canny edge detector, Wikipedia (2023). https://en.wikipedia.org/wiki/Canny_edge_detector (accessed August 6, 2023).
[51] Phase correlation, Wikipedia (2022). https://en.wikipedia.org/wiki/Phase_correlation (accessed August 5, 2023).
[52] OpenCV, Open Source Computer Vision Library, OpenCV (n.d.). https://opencv.org/ (accessed August 6, 2023).
[53] A. Ghahramani, R.M. Levinson, S.K. Min, K. Chen, A.Y. Wang, Autonomous comfort systems, WO2021050369A1, 2021. https://patents.google.com/patent/WO2021050369A1/en (accessed July 30, 2023).
[54] Y. He, H. Zhang, E. Arens, A. Merritt, C. Huizenga, R. Levinson, A. Wang, A. Ghahramani, A. Alvarez-Suarez, Smart detection of indoor occupant thermal state via infrared thermography, computer vision, and machine learning, Build. Environ. 228 (2023) 109811, https://doi.org/10.1016/j.buildenv.2022.109811.
[55] A. Nomoto, D. Kim, H. Zhang, Y. He, C. Huizenga, E. Arens, R. Prickett, S. Swaminathan, R. Levinson, Field study of thermal infrared sensing for office temperature control. 2023 ASHRAE Annual Conference (2023), https://escholarship.org/uc/item/69r9q3kg.
[56] R.J. de Dear, G.S. Brager, Developing an adaptive model of thermal comfort and preference, ASHRAE Transactions 104 (1998) 145–167. https://escholarship.org/uc/item/4qq2p9c6.