

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

Hot, Cold, or Just Right? An Infrared Biometric Sensor to Improve Occupant Comfort and Reduce Overcooling in Buildings via Closed-loop Control

### Permalink

<https://escholarship.org/uc/item/8j72j627>

### Authors

Levinson, Ronnen

Kim, Donghun

Goudey, Howdy

et al.

### Publication Date

2025-02-01

Peer reviewed



Building Technologies Department

Building Technologies and Urban Systems Division

Lawrence Berkeley National Laboratory

# Hot, Cold, or Just Right? An Infrared Biometric Sensor to Improve Occupant Comfort and Reduce Overcooling in Buildings via Closed-loop Control

## Project Report

Ronnen Levinson, Donghun Kim, Howdy Goudey, Sharon Chen, Hui Zhang<sup>1</sup>, Ali Ghahramani<sup>1</sup>, Charlie Huizenga<sup>1</sup>, Yingdong He<sup>1</sup>, Akihisa Nomoto<sup>1</sup>, Alexander Merritt<sup>1</sup>, Edward Arens<sup>1</sup>, Ana Álvarez Suárez<sup>2</sup>, David Ritter<sup>2</sup>, Markus Tarin<sup>2</sup>, and Robert Prickett<sup>3</sup>

<sup>1</sup> Center for the Built Environment, University of California, Berkeley, CA

<sup>2</sup> MoviTHERM, Irvine, CA

<sup>3</sup> Daikin Silicon Valley, Santa Clara, CA

February 2025

doi:10.20357/B7161Q





## **Disclaimer**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or the Regents of the University of California.

Lawrence Berkeley National Laboratory is an equal opportunity employer.

## **Copyright Notice**

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript or allow others to do so, for U.S. Government purposes.

## Acknowledgments

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 (Lawrence Berkeley National Laboratory subcontractor) for helping coordinate project execution; Alexander Merritt (UC Berkeley) and Sang woo Ham (Lawrence Berkeley National Laboratory) for technical support in trials; Sri Swaminathan, Norman Pennant, and Bevnoty Attia of Daikin U.S. for facilitating the Texas field trial; Chun-cheng Piao and Kevin Ninomiya of Daikin U.S. for supporting project execution; and Marina Sofos, Erika Gupta, Brian Walker, and Cecilia Johnson of the Building Technologies Office, U.S. Department of Energy for their support and guidance.

The authors thank the following experts for reviewing this report (affiliations do not imply that those organizations support or endorse this work):

- Alan Meier, Lawrence Berkeley National Laboratory
- Stefano Schiavon, University of California, Berkeley

## Email addresses of authors

Ronnen Levinson\*: [RMLevinson@LBL.gov](mailto:RMLevinson@LBL.gov)

Donghun Kim: [DonghunKim@LBL.gov](mailto:DonghunKim@LBL.gov)

Howdy Goudey: [CWGoudey@LBL.gov](mailto:CWGoudey@LBL.gov)

Sharon Chen: [SSChen@LBL.gov](mailto:SSChen@LBL.gov)

Hui Zhang: [zhanghui@berkeley.edu](mailto:zhanghui@berkeley.edu)

Ali Ghahramani: [ghahramani@nus.edu.sg](mailto:ghahramani@nus.edu.sg)

Charlie Huizenga: [huizenga@berkeley.edu](mailto:huizenga@berkeley.edu)

Yingdong He: [heyindingdong2022@hnu.edu.cn](mailto:heyindingdong2022@hnu.edu.cn)

Akihisa Nomoto: [monyo323232@gmail.com](mailto:monyo323232@gmail.com)

Alexander Merritt: [merrittalexandere@gmail.com](mailto:merrittalexandere@gmail.com)

Edward Arens: [earens@berkeley.edu](mailto:earens@berkeley.edu)

Ana Álvarez Suárez: [a.a.alvarez.suarez@gmail.com](mailto:a.a.alvarez.suarez@gmail.com)

David Ritter: [d.ritter@movitherm.com](mailto:d.ritter@movitherm.com)

Markus Tarin: [m.tarin@movitherm.com](mailto:m.tarin@movitherm.com)

Robert Prickett: [Robert.Prickett@daikincomfort.com](mailto:Robert.Prickett@daikincomfort.com)

\* Corresponding author

## Suggested citation

Levinson, R., Kim, D., Goudey, H., Chen, S., Zhang, H., Ghahramani, A., Huizenga, C., He, Y., Nomoto, A., Merritt, A., Arens, E., Suárez, A. Á., Ritter, D., Tarin, M., & Prickett, R. (2025). *Hot, Cold, or Just Right? An Infrared Biometric Sensor to Improve Occupant Comfort and Reduce Overcooling in Buildings via Closed-loop Control: Project Report*. Lawrence Berkeley National Laboratory, Berkeley, CA, USA.  
<https://doi.org/10.20357/B7161Q>

# 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 and 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.



# Table of Contents

Abstract.....	iii
Table of Figures .....	vi
Abbreviations .....	vii
1 Introduction .....	1
1.1 The challenge of overcooling .....	1
1.2 Proposed solution.....	2
1.3 Project overview .....	2
1.4 Report structure.....	4
2 Skin blood flow and temperature responses to thermal stress.....	4
3 Development of the radiometric sensing hardware and software .....	5
3.1 Thermographic measurement accuracy.....	5
3.1.1 Distance error .....	6
3.1.2 Angle error.....	6
3.1.3 Temporal drift .....	6
3.1.4 Motion blur.....	7
3.2 Machine vision.....	8
3.2.1 Feature recognition in color image.....	8
3.2.2 Transformation of feature shapes from color image to thermal image.....	9
3.3 Radiometric sensing hardware .....	10
3.4 Radiometric sensing software.....	12
3.4.1 Server.....	12
3.4.2 Client .....	13
4 Development of the thermal sensation model .....	13
5 Development of the AC control algorithm .....	14

6	System integration .....	15
7	Demonstration .....	17
7.1	Experiment .....	17
7.2	Results .....	19
7.3	Discussion .....	21
8	Future development .....	22
9	Summary .....	22
	References .....	23
	Appendices .....	30

# Table of Figures

Figure 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..... 11

Figure 2. Near-subject PTR with heated wire around its perimeter..... 12

Figure 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. .... 17

Figure 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. ... 19

Figure 5. Color image and feature outlines of a trial subject (one of the authors) in the conference room as captured by our sensor-controller. .... 19

Figure 6. Distribution of (a) thermal sensation votes and (b) thermal comfort votes under each control algorithm. *N* represents the total number of votes. ....20

Figure 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 Controls and TIR B Control reduced the mean air-cooling rate by 18% and 42%, respectively compared to Conventional Control. ....21

# Abbreviations

AC	air conditioning
ATR	active temperature reference
CMOS	complementary metal–oxide–semiconductor
FN	face and nose
FNH	face, nose, and hand
FOV	field of view
HVAC	heating, ventilation, and air conditioning
NUC	non-uniformity correction
PI	proportional-integral
PTR	passive temperature reference
RPM	revolutions per minute
TIR	thermal infrared

# 1 Introduction

## 1.1 The challenge of overcooling

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].

Current 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 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.

People make distinctive poses or movements when they are thermally uncomfortable [19]. Tracking these gestures has been proposed as a means of quantitatively assessing thermal comfort [20] and ultimately controlling an HVAC system [21].

## 1.2 Proposed solution

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 and 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 the 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 ( $\leq$ US\$165 as of July 2023), medium resolution ( $160 \times 120$  pixel), uncooled TIR camera detectors (microbolometers) which can be coupled with very-low cost ( $\leq$ \$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.

## 1.3 Project overview

Our three-year project encompassed four tasks:

1. **Sensing:** Develop a wide-view, ceiling-mounted sensor that uses a visual camera, a TIR camera, and machine vision to map occupant skin surface temperatures and non-human room surface temperatures (Years 1 & 2)
2. **Comfort:** Develop an experimentally validated logic that predicts occupant comfort from skin surface temperature distribution (Years 1 & 2)
3. **Control:** Develop a closed-loop feedback control algorithm based on sensor outputs (surface temperature maps and occupant comfort statistics) to govern HVAC operation (Years 1 & 2)
4. **Integration & Demonstration:** Assemble and field-test in a commercial building a prototype sensor/controller that incorporates the sensor, comfort logic, and HVAC control algorithm (Years 2 - 3)

We began Task 1 (Sensing) by assessing the accuracy with which high- and medium-resolution TIR cameras measure occupant and room surface temperatures within a narrow field of view (Subtask 1.1), then evaluating techniques to measure surface temperature distribution over a wide field of view with such cameras (Subtask 1.2). We also employed machine-vision software to identify occupant surfaces from high-resolution color images (Subtask 1.3). We completed this task by combining TIR imaging and machine vision to measure the body-part (e.g., head and hand) temperatures of each occupant (Subtask 1.4).

We began Task 2 (Comfort) by evaluating skin temperatures from different body parts and selecting the parts that strongly correlate with thermal comfort for both steady-state and transient conditions (Subtask 2.1). We then created algorithms that use occupant skin temperatures to predict thermal comfort (Subtask 2.2). We completed this task by predicting the environmental requirements for providing and maintaining occupants' comfort (Subtask 2.3), and tuning the comfort prediction models for our prototype sensor/controller (Subtask 2.4).

We began Task 3 (Control) by developing a closed-loop algorithm that uses the Task 2 thermal comfort model and room surface temperature measurements to control the HVAC system (Subtask 3.1). We then verified the correct operation of the controller and assessed its closed-loop performance through emulation prior to full-scale testing (Subtask 3.2).

We began Task 4 (Integration & Demonstration) by building a prototype sensor/controller suitable for field trials (Subtask 4.1). We then assessed the performance of the initial prototype sensor/controller in occupied commercial buildings, and refined the prototype accordingly (Subtask 4.2). Performance was quantified by measuring overcooling (overuse of cooling energy in a manner that discomforts the occupants) with and without our sensor/controller connected to the HVAC system.

## 1.4 Report structure

This report expands upon our open-access journal article summarizing the project (*Appendix M: Hot, Cold, or Just Right? An infrared biometric sensor to improve occupant comfort and reduce overcooling in buildings via closed-loop control*; also Ref. [22]) by incorporating the project's detailed task reports as cited appendices. It 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 [23], representing as much as 60% of total cardiac output [24]. In cold conditions, thermoregulatory vasoconstriction can limit the skin's 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 the removal of the cold stress, the skin blood flow quickly returns to normothermia conditions. Vasodilation begins when the body's 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 [12,25,26].

The physiology of human thermal comfort is explored in *Appendix F: Selection of body parts that represent comfort for steady-state and transient thermal conditions*.

## 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  $\mu\text{m}$ ; absolute accuracy  $\pm 1$  °C or 1% of range; differential accuracy 45 mK @ 30 °C; field of view 45°  $\times$  34°) and a medium-resolution FLIR A315 with a 45° lens adapter (320 by 240 pixels; spectrum 7.5–13  $\mu\text{m}$ ; absolute accuracy  $\pm 2$  °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  $0.95 \pm 0.02$ ; temperature accuracy  $\pm 0.05$  °C; spatial uniformity  $\pm 0.01$  °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  $\pm 0.8$  °C; temperature stability  $\pm 0.1$  °C) set to 55 °C.

The distance-error, angle-error, and temporal-drift experiments summarized in Sections 3.1.1 - 3.1.3 are detailed in *Appendix A: Quantification of TIR temperature measurement accuracy versus angle and distance to subject*. The motion-blur experiments summarized in Section 3.1.4 are elaborated in *Appendix B: Selection of wide-view TIR imaging temperature measurement technique(s)*. Options for thermographic assessment of temperature wide-view head cameras are further explored in *Appendix D: Assess occupant facet temperatures and non-human surface temperature using wide-view overhead camera*.

### 3.1.1 Distance error

Distance trials with the high-resolution camera viewing the 100 cm<sup>2</sup>, 35 °C ATR at distances 1 – 34 m and with the medium-resolution camera viewing the 26 cm<sup>2</sup>, 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 the 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 [27,28]—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 [27], 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 [29].

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 distortion or significant 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 ATR 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,  $2,500 \times 2,000$  pixels) offers about 250 times the spatial resolution of a low-resolution microbolometer thermal image sensor ( $160 \times 120$  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 [30], some of the current authors compared the accuracy and speed of three open-source machine-vision algorithms—Bulat’s Face Alignment [31,32], InsightFace [33,34], and FaceNet [35,36]—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) [37], an open-source, face-only algorithm scheme closely related to Bulat’s Face Alignment [31,32] that we would augment with a hand-segmentation (recognition) model trained with the EgoHands dataset [38,39]; OpenPose (face, hands, and body pose; 70 facial,  $2 \times 21$  hand, and 25 pose keypoints; one or more subjects;

free for non-commercial use only) [40,41]; and MediaPipe Holistic (face, hands, and pose; 468 facial,  $2 \times 21$  hand, and 33 pose keypoints; one subject; open source) [42,43].

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. Our evaluation of these machine-vision tools is detailed in *Appendix C: Machine-vision software to identify subject areas of interest for thermal comfort*.

MediaPipe Holistic quickly and accurately located facial, hand, and pose keypoints. We did not quantitatively 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 [44–46].

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 [47–50] within the mature field of image registration [51–53]. 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 [30] cropped and downscaled the color image to the resolution of the thermal image ( $160 \times 120$  pixels), used Canny edge detection [54] to find edges in the color and thermal images, applied phase correlation [55] 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 [56] 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 (Figure 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)<sup>1</sup> and a microbolometer TIR sensor (FLIR Lepton 3.5; 160 × 120 pixels; 51° × 49° FOV).

---

<sup>1</sup> A lower-resolution color camera would be fine since we operate it at only 960 × 720 pixels.



- **Stationary camera**—a stationary fisheye color camera (Sony IMX291;  $1,945 \times 1,109$  pixels,  $175^\circ$  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.




Figure 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

We also place a second, larger PTR (Figure 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.




Figure 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 [30] by adding the pan-tilt platform, on-board 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<sup>2</sup> 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

---

<sup>2</sup> Our server software descends from what was called “client” software in the “iEye” system developed by one of the authors [30,57]. It has been refactored to act as the server rather than the client in the transmission of images and temperatures. The sensing performance of “iEye” is detailed in *Appendix E: Prototype sensor/controller can measure occupant facet temperature skin temperature differences to within 0.2 K*, while the performance of the iEye control algorithm is described in *Appendix I: Development and evaluation of algorithms to predict ambient requirements as constraints for HVAC optimization using prototype sensor/controller*.



- 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-uvic-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;
- 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 90<sup>th</sup> 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. [58], 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 x 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 Sections 3.3 and 3.4, most of which were recorded during the same human-subject trials reported in Ref. [58]. 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  $-1 \times \text{cool\_probability} + 0 \times \text{neutral\_probability} + 1 \times \text{warm\_probability}$ ). Since a discrete (integer) value could result in undesired cycling, the probability-weighted simplified sensation is used to guide our AC-control algorithm.

The human-subject thermal-sensation trials and the development of algorithms to predict thermal sensation from skin-temperature distribution are detailed in *Appendix G: Develop algorithms to predict thermal comfort using occupant skin temperature measurements* and *Appendix H: Smart detection of indoor occupant thermal state via infrared thermography, computer vision, and machine learning*.

## 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.<sup>3</sup> 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

---

<sup>3</sup> We initially designed and simulated a model predictive control (MPC) algorithm before switching to a simpler PI control logic. The design of the MPC algorithm is detailed in *Appendix J: Design of MPC for comfort sensing technologies for controlling occupant thermal comfort*.

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 minutes, 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 (Sections 3.3 and 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.

Figure 3 illustrates Goldilocks in operation. In panel A (top left), MediaPipe Holistic locates 468 facial,  $2 \times 21$  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).

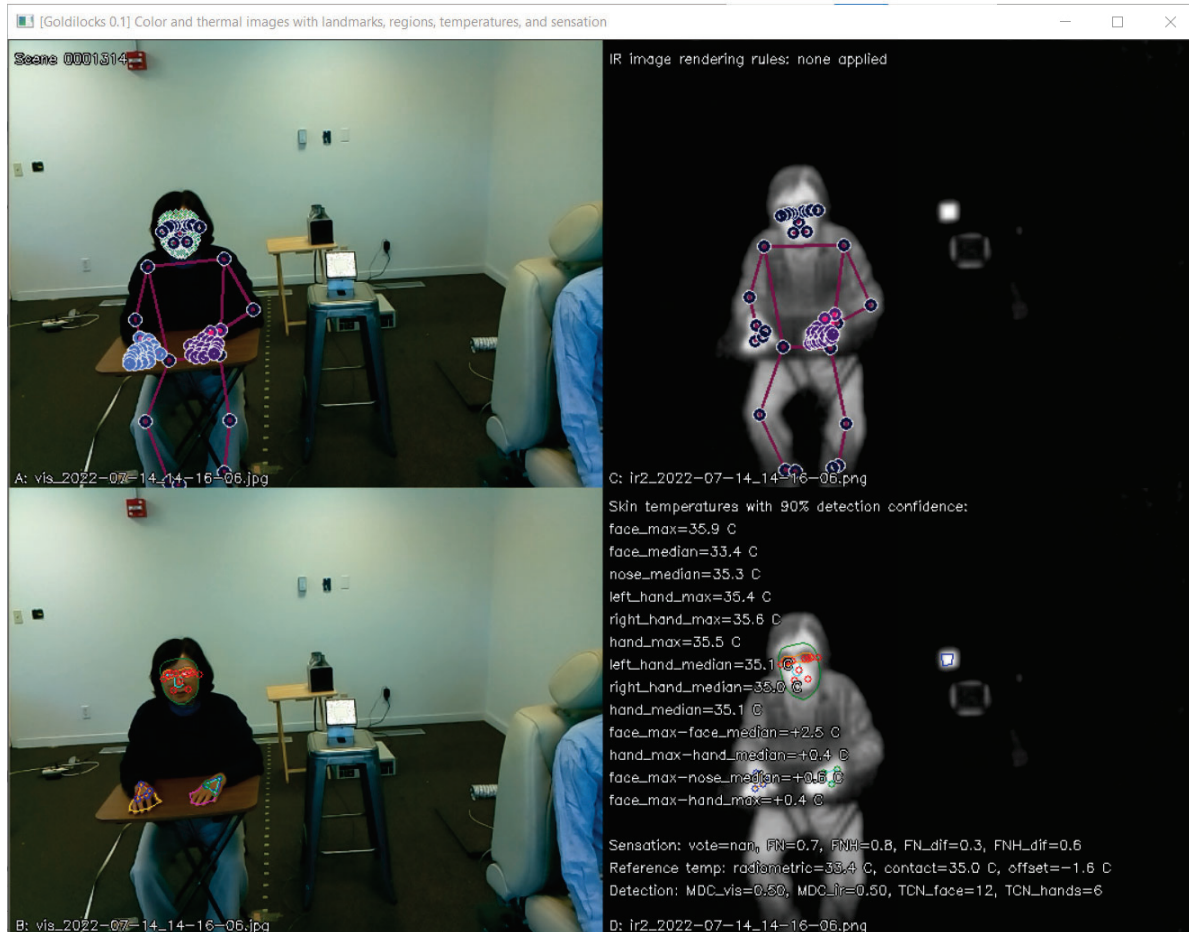


Figure 3. Operation of the “Goldilocks” client software showing (A) face, hand, and pose landmarks in the color image; (B) facial pose landmarks, face outline, nose outline, eye outlines, eyeglasses outline, hand outlines, and palm outlines in the color image; (C) pose landmarks in the thermal image; 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.

## 7 Demonstration

The following demonstration is detailed in *Appendix L: Field study of thermal infrared sensing for office temperature control*.

### 7.1 Experiment

We tested our system in a conference room in an office building near Houston, Texas from August 2022 to January 2023 [59]. 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 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 hour, >1 hour), (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 (Figure 4b). Figure 5 shows a representative color image and the feature outlines of a trial subject as captured by the sensor-controller.





Figure 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.



Figure 5. Color image and feature outlines of a trial subject (one of the authors) in the conference room as captured by our 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 (Figure 6a). This indicates that the Conventional Control overcooled the occupants.




Figure 6. Distribution of (a) thermal sensation votes and (b) thermal comfort votes under each control algorithm.  $N$  represents the total number of votes.

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 (Figure 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 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 (Figure 6a). There was no significant difference in thermal comfort votes between Conventional Control and TIR A Control (Figure 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 (Figure 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%).



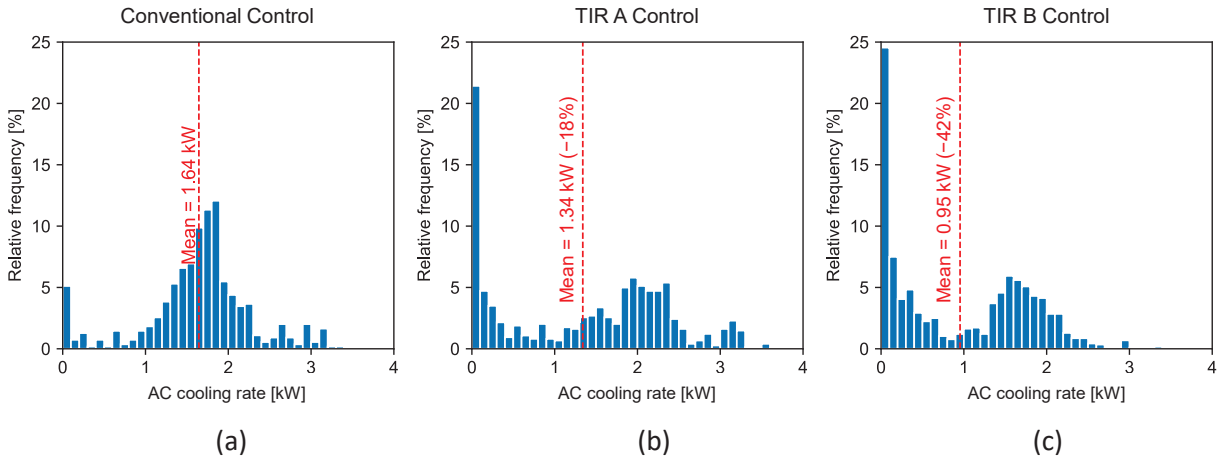


Figure 7. Distributions of the air-cooling rate (rate of heat removal by the AC) when the conference room was occupied under (a) Conventional Control, (b) TIR A Control, or (c) TIR B Control. TIR A Controls and TIR B Control reduced the mean air-cooling rate by 18% and 42%, respectively compared to Conventional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

## 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  $\leq 20$  °C [68.0 °F]), neutral environments (around 24 °C [75.2 °F]), and hot environments ( $\geq 30$  °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 ( $\leq 20$  °C or  $\geq 30$  °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 [60]. With feedback-providing devices like that described in our study, building designers and operators can set the ambient temperature based on 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

Berkeley Lab filed a provision patent application for our technology in October 2023.

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.

## 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, *Buildings and 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), *Science and Technology for the Built Environment* 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 Journal* (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, *Building and Environment* 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, *Building and Environment* 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, *Energy and Buildings* 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, *Building and Environment* 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, *Building and Environment* 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](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, *Building and Environment* 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 Lab, Berkeley, CA, 2009. <https://doi.org/10.2172/982931>.
- [19] A. Meier, X. Cheng, W. Dyer, G. Chris, T. Olofsson, B. Yang, Non-invasive assessments of thermal discomfort in real time, in: CATE 2019 – Comfort at the Extremes: Energy, Economy and Climate, Ecohouse Initiative Ltd., 2019. <https://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-158145> (accessed June 15, 2024).
- [20] B. Yang, X. Li, Y. Hou, A. Meier, X. Cheng, J.-H. Choi, F. Wang, H. Wang, A. Wagner, D. Yan, A. Li, T. Olofsson, H. Li, Non-invasive (non-contact) measurements of human thermal physiology signals and thermal comfort/discomfort poses -A review, *Energy and Buildings* 224 (2020) 110261. <https://doi.org/10.1016/j.enbuild.2020.110261>.
- [21] A. Meier, W. Dyer, C. Graham, Using Human Gestures to Control a Building’s Heating and Cooling System, in: *Energy Efficiency in Domestic Appliances and Lighting (EEDAL)*, 2017: pp. 627–635. <https://doi.org/10.2760/264880>.
- [22] R. Levinson, D. Kim, H. Goudey, S. Chen, H. Zhang, A. Ghahramani, C. Huizenga, Y. He,

- A. Nomoto, E. Arens, A.Á. Suárez, D. Ritter, M. Tarin, R. Prickett, Hot, cold, or just right? An infrared biometric sensor to improve occupant comfort and reduce overcooling in buildings via closed-loop control, *Energy and Buildings* (2024) 114063. <https://doi.org/10.1016/j.enbuild.2024.114063>.
- [23] 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>.
- [24] L.B. Rowell, Cardiovascular Adjustments to Thermal Stress, in: R. Terjung (Ed.), *Comprehensive Physiology*, 1st ed., Wiley, 2011: pp. 967–1023. <https://doi.org/10.1002/cphy.cp020327>.
- [25] C. Dai, H. Zhang, E. Arens, Z. Lian, Machine learning approaches to predict thermal demands using skin temperatures: Steady-state conditions, *Building and Environment* 114 (2017) 1–10. <https://doi.org/10.1016/j.buildenv.2016.12.005>.
- [26] A. Ghahramani, G. Castro, B. Becerik-Gerber, X. Yu, Infrared thermography of human face for monitoring thermoregulation performance and estimating personal thermal comfort, *Building and Environment* 109 (2016) 1–11. <https://doi.org/10.1016/j.buildenv.2016.09.005>.
- [27] B. Oswald-Tranta, M. Sorger, P. O’Leary, Motion deblurring of infrared images from a microbolometer camera, *Infrared Physics & Technology* 53 (2010) 274–279. <https://doi.org/10.1016/j.infrared.2010.04.003>.
- [28] M. Vollmer, K.-P. Möllmann, *Infrared Thermal Imaging: Fundamentals, Research and Applications*, 1st ed., John Wiley & Sons, Ltd, Weinheim, 2010. <https://doi.org/10.1002/9783527630868>.
- [29] 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).
- [30] 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) 1241. <https://doi.org/10.3390/buildings12081241>.
- [31] A. Bulat, *Face Recognition*, (2023). <https://github.com/1adrianb/face-alignment> (accessed August 5, 2023).
- [32] 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>.
- [33] J. Guo, J. Deng, InsightFace: 2D and 3D Face Analysis Project, (2023).  
<https://github.com/deepinsight/insightface> (accessed August 5, 2023).
- [34] 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>.
- [35] I. de P. Centeno, MTCNN, (2023). <https://github.com/ipazc/mtcnn> (accessed August 5, 2023).
- [36] K. Zhang, Z. Zhang, Z. Li, Y. Qiao, Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks, IEEE Signal Processing Letters 23 (2016) 1499–1503.  
<https://doi.org/10.1109/LSP.2016.2603342>.
- [37] 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/ladrianb/2D-and-3D-face-alignment> (accessed August 5, 2023).
- [38] S. Bambach, EgoHands: A Dataset for Hands in Complex Egocentric Interactions, (n.d.).  
<http://vision.soic.indiana.edu/projects/egohands/> (accessed August 5, 2023).
- [39] 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>.
- [40] G. Hidalgo, Y. Raaj, OpenPose, (2023). <https://github.com/CMU-Perceptual-Computing-Lab/openpose> (accessed August 5, 2023).
- [41] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, Y. Sheikh, OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields, IEEE Transactions on Pattern Analysis and Machine Intelligence 43 (2021) 172–186. <https://doi.org/10.1109/TPAMI.2019.2929257>.
- [42] Google, MediaPipe, (2023). <https://github.com/google/mediapipe> (accessed August 5, 2023).
- [43] Google, MediaPipe, Google for Developers (n.d.). <https://developers.google.com/mediapipe> (accessed August 5, 2023).
- [44] A. Saiwa, Openpose vs Mediapipe for Dynamic Vision | Beyond Poses, (2023).



<https://saiwa.ai/blog/openpose-vs-mediapipe/> (accessed August 5, 2023).

- [45] P. Radzki, Detection of human body landmarks - MediaPipe and OpenPose comparison, HearAI (2022). <https://www.hearai.pl/post/14-openpose/> (accessed August 5, 2023).
- [46] 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>.
- [47] 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](https://doi.org/10.1007/978-3-642-32384-3_19).
- [48] 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 Sensing* 9 (2017) 441. <https://doi.org/10.3390/rs9050441>.
- [49] 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>.
- [50] J. Han, E.J. Pauwels, P. de Zeeuw, Visible and infrared image registration in man-made environments employing hybrid visual features, *Pattern Recognition Letters* 34 (2013) 42–51. <https://doi.org/10.1016/j.patrec.2012.03.022>.
- [51] L.G. Brown, A survey of image registration techniques, *ACM Comput. Surv.* 24 (1992) 325–376. <https://doi.org/10.1145/146370.146374>.
- [52] B. Zitová, J. Flusser, Image registration methods: a survey, *Image and Vision Computing* 21 (2003) 977–1000. [https://doi.org/10.1016/S0262-8856\(03\)00137-9](https://doi.org/10.1016/S0262-8856(03)00137-9).
- [53] Z. Xiong, Y. Zhang, A critical review of image registration methods, *International Journal of Image and Data Fusion* 1 (2010) 137–158. <https://doi.org/10.1080/19479831003802790>.
- [54] Canny edge detector, Wikipedia (2023). [https://en.wikipedia.org/wiki/Canny\\_edge\\_detector](https://en.wikipedia.org/wiki/Canny_edge_detector) (accessed August 6, 2023).
- [55] Phase correlation, Wikipedia (2022). [https://en.wikipedia.org/wiki/Phase\\_correlation](https://en.wikipedia.org/wiki/Phase_correlation) (accessed August 5, 2023).
- [56] OpenCV, Open Source Computer Vision Library, OpenCV (n.d.). <https://opencv.org/> (accessed August 6, 2023).
- [57] 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).

- [58] 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, *Building and Environment* 228 (2023) 109811. <https://doi.org/10.1016/j.buildenv.2022.109811>.
- [59] 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, in: *2023 ASHRAE Annual Conference*, Tampa, FL, 2023. <https://escholarship.org/uc/item/69r9q3kg>.
- [60] 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>.

## Appendices

- A. Quantification of TIR temperature measurement accuracy versus angle and distance to subject [Task 1.1 / Milestone 1.1]
- B. Selection of wide-view TIR imaging temperature measurement technique(s) [Task 1.2 / Milestone 1.2]
- C. Machine-vision software to identify subject areas of interest for thermal comfort [Task 1.3 / Milestone 1.3]
- D. Assess occupant facet temperatures and non-human surface temperature using wide-view overhead camera [Task 1.4 / Milestone 1.4]
- E. Prototype sensor/controller can measure occupant facet temperature skin temperature differences to within 0.2 K [Task 1 / Go/No-Go Decision Point 2]
- F. Selection of body parts that represent comfort for steady-state and transient thermal conditions [Task 2.1 / Milestone 2.1]
- G. Develop algorithms to predict thermal comfort using occupant skin temperature measurements [Subtask 2.2 / Milestone 2.2]
- H. Smart detection of indoor occupant thermal state via infrared thermography, computer vision, and machine learning [Subtask 2.2 / Milestone 2.2]
- I. Development and evaluation of algorithms to predict ambient requirements as constraints for HVAC optimization using prototype sensor/controller [Tasks 2.3 + 2.4 / Milestones 2.3.1 + 2.3.2 + 2.4]
- J. Design of MPC for comfort sensing technologies for controlling occupant thermal comfort [Task 3.2 / Milestone 3.2]
- K. Identification of field test site [Task 4.2 / Milestone 4.2.1]
- L. Field study of thermal infrared sensing for office temperature control [Task 4.2 / Milestones 4.2.2 + 4.2.3]
- M. Hot, Cold, or Just Right? An infrared biometric sensor to improve occupant comfort and reduce overcooling in buildings via closed-loop control [Task 4.2 / Milestones 4.2.2 + 4.2.3]

# Appendix A: Quantification of TIR temperature measurement accuracy versus angle and distance to subject

## HCJR Task 1.1 / Milestone 1.1

Sharon Chen<sup>a</sup>, Ana Alvarez<sup>b</sup>, Howdy Goudey<sup>a</sup>, Dave Ritter<sup>b</sup>, Markus Tarin<sup>b</sup>, and Ronnen Levinson<sup>a</sup>

<sup>a</sup> Heat Island Group, Lawrence Berkeley National Laboratory, Berkeley, CA

<sup>b</sup> MoviTHERM, Irvine, CA

2020-01-21

## Abstract

To assess the accuracy of absolute and differential thermal infrared temperature measurements achievable under the constraints posed by the *Hot, Cold, or Just Right* application, we performed a suite of experiments that examined each possible source of error—first in isolation, then combined—in test scenes composed of human and non-human subjects in office rooms. We collected thermal image sets that varied the camera-to-subject distance and viewing angle, represented camera performance in both low-drift and high-drift conditions, and included a variety of both human and non-human subjects. We assessed accuracy by comparing radiatively sensed temperatures to both contact-sensed temperatures and external temperature references. We found that, in isolation, increasing camera-to-subject distance and viewing angle did degrade performance as we expected. We observed that while temporal drift was a significant and uncontrollable source of error in absolute temperature measurements, it had little effect on differential measurements. Poor resolution of the target can contribute error to both absolute and differential temperatures; this topic merits further study. Finally, unexpected results observed in our test images of human and non-human subjects suggest that there may be other uncontrollable environmentally induced sources of error yet to be identified.

## 1 Introduction

The eventual use case for thermal infrared (TIR) temperature sensing proposed in the *Hot, Cold, or Just Right* (HCJR) project involves characterizing temperatures of small human features under non-ideal circumstances. The TIR camera must be able to accurately measure differential temperatures between centimeter-scale human facial features located far from the camera and at oblique angles. The objective of this report, therefore, is to provide an assessment of the accuracy limits of a TIR camera operating within these environmental constraints. We aim to identify the sources of TIR measurement error (both those inherent in the measurement principle and caused by hardware limitations), experimentally isolate and explore the effect of those sources of error (in

both trend and magnitude), and to directly compare the results of TIR test images of human and non-human objects inside typical indoor environment to temperatures measured by contact sensors.

## 2 Principles of TIR temperature measurement

The microbolometer sensor in many infrared cameras consists of a two-dimensional matrix of infrared-sensing elements that collectively generate a map of incident infrared radiation. When these sensing elements absorb long-wave infrared radiation (LWIR) from a scene, their temperatures (and thus electrical resistances) change; these resistance changes are measured and then translated by back-end electronics into scene temperature values. The transformation considers the following points.

**1. Only a portion of the radiation detected by the sensor originates from thermal emission by the target of interest.** Every object emits radiation of a magnitude proportional to the fourth power of its absolute temperature. The radiant flux received by the camera sensor is a combination of thermal emission from the target, thermal reflection by the target, and thermal emission by the atmosphere occupying the line-of-sight path between the target and camera (FLIR Systems, Inc., 2015). Solving for the object temperature requires isolating the portion directly contributed by thermal emission from the target.

**2. A non-zero amount of radiation traveling from the target to the camera is “lost” to absorption and scattering by the intervening air.** This attenuation of radiation by the atmosphere is characterized by an atmospheric transmittance term that varies with scene composition and local meteorological conditions (Holst, 2000). Atmospheric transmittance decreases with increasing concentration of LWIR-absorbing gases such as water vapor. Its magnitude is also affected by the presence of atmospheric aerosols, like fog or pollution, that may absorb and scatter radiation. Finally, the length of the air path determines the scaling of the attenuation (Holst, 2000). Many infrared imaging software packages perform atmospheric transmittance corrections based on user inputs for ambient relative humidity and target distance.

**3. Some portion of the radiation detected by the camera sensor consists of thermal emission by the air path between the target and the camera.** The magnitude of this term depends on the temperature of the air path between the camera and target. It is often assumed that temperature is constant along this path length.

**4. Some portion of the radiation detected by the sensor consists of thermal reflection by the target.** A target of non-zero reflectance will, by definition, reflect some ambient radiation towards the camera sensor. For simplicity, we can describe the environmental radiation incident upon a target’s surface as emission by a blackbody of temperature  $T_{\text{refl}}$  (Holst, 2000).  $T_{\text{refl}}$  (often termed “reflected apparent temperature”, or simply “reflected temperature”) can be determined experimentally via either of two methods described in ASTM E1862-14, the *Standard Practice for Measuring and Compensating for Reflected Temperature Using Infrared Imaging Radiometers* (ASTM International, 2014). In the “Reflector Method”,  $T_{\text{refl}}$  is the measured apparent temperature

of a diffuse infrared reflector that is placed at the same location and angle as the subject. In the “Direct Method”, the camera itself is moved to the location of the target and measures directly the average scene temperature at various angles.

If all surfaces in the vicinity of a scene are at the same temperature, the measured local ambient (or atmospheric) temperature  $T_{atm}$  can sometimes be an acceptable proxy for  $T_{refl}$ . However, consider a situation where the camera operator (a heat source) stands in the field of view of the nearby target surface. A temperature sensor suspended some short distance away monitoring room temperature would not properly capture the radiation emitted by the human and incident upon the target. The  $T_{atm}$  measurement would in this case underestimate  $T_{refl}$ .

Table 1 presents a summary of the user-supplied terms an infrared camera must consider in order to transform detected radiation into a temperature value. They must either be supplied by the user or calculated from user inputs.

Table 1. Accurate TIR measurement requires knowledge of the following variables.

Variable	Description	Method
$\epsilon_{obj}$	Thermal emittance of target	Direct measurement; look-up tables
$T_{refl}$	Reflected apparent temperature	Direct measurement via ASTM E1862
$T_{atm}$	Local atmospheric temperature	Direct measurement
$D$	Distance between target and camera aperture	Direct measurement
$RH$	Local relative humidity	Direct measurement

### 3 Equipment specification, selection, and acquisition

Berkeley Lab and MoviTHERM each performed their own explorations of TIR measurement accuracy using their native equipment and capabilities. Doing so ensured that our conclusions were not biased to any one particular operator or lab set-up and expanded the pool of experiments that we could collectively perform. The instruments and equipment available to each institution are described in this section.

### 3.1 Thermal cameras




Figure 1. Two photographs showing (a) the tripod-mounted the FLIR SC660 at rest, and (b) in use.

Infrared images at **Berkeley Lab** were captured with a FLIR SC660 handheld microbolometer-based infrared camera outfitted with a 45° wide angle lens (see Figure 1). It is a high-resolution LWIR thermographic imager with an accuracy of roughly  $\pm 1$  °C or 1% of range. Select manufacturer-reported specifications are reproduced in Table 2 below.

Table 2. Selected manufacturer-reported specifications of the FLIR SC660 with 45° lens used at Berkeley Lab.

Item	Specification
Resolution	640 × 480 pixels
Detector type	Focal Plane Array (FPA), uncooled microbolometer
Spectral range	7.5–13 $\mu\text{m}$
Object temperature range	–40 °C to +120 °C
Accuracy	$\pm 1$ °C or 1% of range
Noise Equivalent Temperature Difference (NETD) <sup>a</sup>	<45 millikelvin (mK) @ +30 °C
Field of view (FOV)	45° × 34°
FOV of individual sensing element	1.3 milliradians (mrad)
Minimum focus distance	0.2 m
Focal length	19 mm
F-number	1.1

<sup>a</sup> Refer to Section 5.1.1 for a description of this term.

Unless otherwise specified, the camera was powered on and allowed 1 hour for its output to stabilize before use. Because the camera has no option allowing the user to control timing of non-uniformity corrections (NUCs; refer to Section 5.1.1 for an in-depth discussion of spatial non-

uniformity noise and methods for non-uniformity correction), we usually allowed it to trigger NUC events on its manufacturer-programmed (native) schedule. Due to the volume of images captured, timing of NUC events was not recorded. Anecdotally, we observed the elapsed time between NUC events to range from approximately 3 s to 30 min, and that NUC events often occurred very frequently immediately after turn-on and slower the longer the camera remained powered on.

Inspection and analysis of thermal images collected at Berkeley Lab was performed in the free version of FLIR Systems' FLIR Tools software.



Figure 2. FLIR A315 TIR Camera with 45° lens adapter used at MoviTHERM.

For all testing described in this document, **MoviTHERM** used a FLIR A315 thermal imaging camera with the 45° FOV (field of view) lens adapter installed (Figure 2). This camera—one of two FLIR A315 cameras acquired for HCJR project testing—has an infrared sensor resolution of  $320 \times 240$  pixels. Other pertinent camera specifications are detailed in Table 3.

Table 3. Details of the FLIR A315 camera set used at MoviTHERM.

Item	Specification
Sensor type	Uncooled microbolometer
Spectral range	7.5 to 13 $\mu\text{m}$
Image capture frequency	15 Hz/30 Hz/60 Hz
Target temperature measurement ranges	
Low temperature mode	-20 to 120 $^{\circ}\text{C}$
High temperature mode	0 to 350 $^{\circ}\text{C}$
Motorized focus	Automatic and manual focus modes
Non-uniformity correction (sensor drift compensation)	Automatic or manual operation
Accuracy	$\pm 2^{\circ}\text{C}$ or $\pm 2\%$ of the reading
Thermal sensitivity	$< 0.05^{\circ}\text{C}$ @ $+30^{\circ}\text{C}$
NETD (“thermal contrast”)	50 mK
On-board emissivity correction	Variable from $\epsilon = 0.01$ to 1.0
Optics transmission correction	Automatic, based on internal sensor feedback
Native camera optics	
Field of view (FOV)	$25^{\circ} \times 18.8^{\circ}$
Focal length	18 mm
Minimum focal distance	0.4 m
45° lens adapter attachment	
Field of view (FOV)	$45^{\circ} \times 33.8^{\circ}$
Focal length	9.66 mm
Minimum focal distance	0.2 m



Figure 3. FLIR Blackfly S BFS-PGE-50S5C Color Camera w/ Tamron 23FM08L Lens used at MoviTHERM.

In addition to the FLIR A315 TIR camera, MoviTHERM used a FLIR Blackfly S color camera (Figure 3) to capture color reference images. These provided timestamped visual records of the blackbody emitter and contact temperature measurements while the thermal images were collected.



## 3.2 Blackbody reference emitters




Figure 4. Photograph of the CI Systems SR-80-4A Infrared Blackbody used at Berkeley Lab. The emitting surface is a 10 cm x 10 cm black-painted metal plate in the center of the blackbody head.

At **Berkeley Lab**, a stable, controllable, and calibrated infrared illumination source, or “reference emitter”, was provided by a CI Systems SR-80-4A Infrared Blackbody (Figure 4; Table 4). Its emitting surface is a 10 cm × 10 cm high-emissivity coated metal plate whose temperature can be controlled to a 0.01 °C resolution within the temperature range of 5 °C to 100 °C.

In all Berkeley Lab trials, the reference emitter setpoint temperature was programmed to 35.00 °C; this value was chosen because it was the highest human forehead temperature recorded in previous human subject trials. To ensure a stable output, the system was always allowed more than 30 min of warm up time before use.

Since the reference emitter provides a uniform 10 cm square area of known emissivity and stable temperature, it was more often utilized as a target for our experiments rather than as an external temperature reference. Given that the broad goal of our experiments was to assess the errors resulting from non-ideal experimental setups, it can be assumed that the radiatively sensed temperatures reported in this document were not corrected by an external reference (unless otherwise specified).

Table 4. Selected manufacturer-reported specifications of the CI Systems SR-80-4A reference emitter used at Berkeley Lab.

Item	Specification
Temperature range	5 °C to 100 °C
Temperature accuracy	±0.05 °C
Resolution (setpoint, readout)	±0.01 °C
Uniformity	±0.01 °C
Emissivity	0.95 ± 0.02




Figure 5. Thermoworks IR-500 Portable Infrared Calibrator/Blackbody Reference Emitter used at MoviTHERM.

All **MoviTHERM** IR imaging tests incorporated a Thermoworks IR-500 Portable Infrared Calibrator/Blackbody reference emitter (Figure 5; Table 5) in the field of view (Figure 6). The IR-500 unit combines a 58 mm diameter thermal target, a class-A PT100 temperature sensor with a microcomputer temperature controller, and a device for rapid heating and cooling.




Figure 6. Thermal image of the Thermoworks IR-500 Portable Infrared Calibrator/Blackbody reference emitter unit, captured using the FLIR A315 at 1 m.

Table 5. Details of the Thermoworks IR-500 Portable Infrared Calibrator/Blackbody used at MoviTHERM.

Item	Specification
Temperature range	50 – 500 °C
Accuracy in test range	±0.8 °C at setpoint points ≤ 100 °C
Stability in test range	±0.1 °C at setpoint points ≤ 100 °C
Emitter surface thermal emittance	0.95
Resolution	0.1 °C

### 3.3 Contact temperature sensors

Two types of contact temperature sensors were available for use at **Berkeley Lab**: fine-gauge T-type thermocouples (from Omega Engineering) and standard-precision negative temperature coefficient 10 kΩ interchangeable epoxy bead thermistors (Littelfuse KS103G3). The T-type thermocouples have a measurement range of 0 °C to 200 °C, with accuracy reported as the greater of either ±1.0 °C or 0.75% of reading (above 0 °C). The thermistors have a measurement range of 0 °C to 70 °C with interchangeability accuracy of ±0.1 °C. Loggers continually recorded sensor temperature outputs at intervals of 10 s, and all temperature measurements reported in this document represent the average of at least three consecutive readings.

**MoviTHERM** captured contact temperature measurements with an Extech 421508 portable digital thermometer (Figure 7; Table 6) with a K-type thermocouple. The K-type thermocouple was mounted to a low thermal emittance (0.04—0.07) aluminum plate using thermal paste. Standard blue masking tape was used to secure the thermocouple lead and prevent movement during testing.



Figure 7. Image of the Extech 421508 Portable Thermometer w/ K-type thermocouple used at MoviTHERM.

Table 6. Specifications of Extech 421508 Portable Thermometer w/ K-type thermocouple used at MoviTHERM.

Item	Specification
Sample rate	1 Hz
K-type thermocouple measurement range/resolution	-200 to 1,372 °C / 0.1 °C (in measurement range)
Extech thermometer accuracy	±(0.05% full scale + 0.3 °C) from -50 to 1,370 °C
Probe accuracy	±2.2 °C or ±0.75% full scale (whichever is greater) from 0 to 800 °C

While thermal image capture was underway, live temperature readings from the Extech 421508 Liquid Crystal Display (LCD) were captured concurrently using the FLIR Blackfly Color (visible spectrum) camera referenced earlier. All contact temperature measurement values noted in the MoviTHERM test results were extracted from the color images.

## 4 Berkeley Lab experiments

### 4.1 Test environment

Most tests were performed inside a closed conference room with no windows, at locations where both camera and subject would not be exposed to drafts. Preliminary location scouting showed the air temperature in this room to be fairly stable, so no changes were made to the conference room’s normal temperature control scheme.

One test was performed inside a long interior hallway with no foot traffic. Several vents were found along its 2.4 m (8 ft) high ceiling, which was a cause for concern even though no drafts were identified at the time of the experiment. Nevertheless, we ensured that over the course of the trials both camera and subject were never placed closer than 2 m from the nearest vent.

### 4.2 Reflected temperature corrections

Reflected apparent temperature was measured, as much as was convenient, following the “Reflector Method” described in ASTM E1862-14. We used as our infrared reflectors crumpled and flattened sheets of aluminum foil (shiny side forward). These infrared reflectors were placed adjacent to—and photographed alongside—the target so that a unique  $T_{\text{refl}}$  correction could be applied to each individual photograph.

All radiatively sensed temperature measurements in this document, unless otherwise noted, incorporate  $T_{\text{refl}}$  corrections.

### 4.3 Berkeley Lab experiment 1: Varying viewing angle

The goal of this experiment was to investigate how the apparent radiatively sensed temperature of the reference emitter hot spot (set to 35.00 °C) varied with viewing angle. We define here the viewing angle as the angle between the camera-to-target line of sight and the surface normal of the target. To achieve this, the reference emitter head was mounted on a turntable with its emitter aperture positioned roughly 1 m from the camera. A total of six different viewing angles, ranging from 0° to 80°, were assessed; no angles past 80° were measured as the horizontal dimension of the emitter plate footprint appeared to dip below 4 pixels. At each of the six angles, 15 thermal images (taken at 20 s intervals) were captured and averaged (Figure 8).




Figure 8. Photographs of the reference emitter taken at viewing angles of 0°, 45°, 75°, and 80°.

### 4.4 Berkeley Lab experiment 2: Investigating the magnitude of measurement drift

The goal of this experiment was to assess the magnitude of temporal noise and drift. We attempted to do so by comparing how the camera performed during periods of low drift and high drift, both with and without its manufacturer-programmed NUC schedule. Based on anecdotal observations that camera native NUC frequency tended to decrease the longer the camera stayed powered on, we assumed measurement drift to be high immediately after the camera was turned on, and low when it was allowed to operate continuously for upwards of an hour. Again, the 35.00 °C reference emitter served as the test subject in these trials (Figure 9).

In the first part of the experiment, the camera was powered on and allowed more than one hour for its output to stabilize. The reference emitter head was positioned roughly 1 m from the camera at a viewing angle of 0°. Two sets of images were captured: one set with the camera programmed to record a photograph every 20 s over 5 min, and another with the camera programmed to record a photograph every 8 s over 20 min. In the former, the camera was allowed to operate on its native NUC schedule; in the latter, the NUC scheme was suppressed. Note that while the NUC routine was allowed to run during the warm-up period, it suppressed 30 min before the beginning of the experiment.

The second part of the experiment employed a similar setup to the first, but here the image sets were captured immediately after powering on the camera. In other words, the camera was not allowed to “warm up” before the measurements began. Two sets of images were collected: one with the camera allowed to operate on its native NUC scheme, and another with the camera NUC

scheme suppressed. In the former, three 15-min image time series were captured at 20 s intervals. In the latter, several 3-min to 15-min image time series were captured at intervals ranging from 1 s to 10 s. During some of these trials, contact sensors mounted on a neighboring wall surface in the background simultaneously captured time series (at 10 s intervals) of the wall temperature.

We assessed the stability of absolute temperature by comparing the camera output to the reference emitter setpoint temperature. Since a second reference emitter was not available, we assessed the stability of differential temperature by tracking temporal changes in the temperature difference between separate areas within the same scene. (While the reference emitter plate temperature does not appear to be completely uniform across its entire surface, its temperature distribution pattern does seem to remain stable over time.) We drew four equal-sized quadrants (Q1 to Q4) across the surface of the emitter plate and found their spatially averaged temperatures to indeed differ from each other by a few tenths of a degree Celsius. We therefore chose these quadrant temperature differences to serve as our stable references of differential temperature.

In the unlikely event that the reference emitter could not provide stable emission over the course of the experiment, we decided to concurrently monitor the temperature of a 10 cm square area of wall located 2 m behind the reference emitter (Figure 9). This area was chosen such that it appeared clearly within the camera frame, but was physically located far enough away from the reference emitter that we could be confident it would not be affected by heat from the temperature control unit. We confirmed this assumption via measurements from two contact temperature sensors (one thermistor and one thermocouple).






Figure 9. The reference emitter plate was divided into four quadrants, the differences of which served as a reference for assessing the stability of differential temperature.

#### 4.5 Berkeley Lab experiment 3: Varying distance

To achieve the long distances necessary for this experiment, the camera and reference emitter were moved into a long interior hallway. The camera and reference emitter head were positioned such that the viewing angle to the emitter plate was  $0^\circ$ . Image sets were first captured with the camera positioned 1 m and 34 m—the latter being the distance at which the 10 cm square emitter plate appeared to be the size of 1 image pixel—away from the reference emitter. Additional image sets were captured with the camera moved to intermediate distances until a total of 8 different distances were assessed (Figure 10). Each set consisted of a series of 15 images captured at 20 s intervals. Relative humidity was not measured in this experiment and was assumed to be 50%.

At distances past 4 m it became increasingly difficult to verify that the camera viewing angle was indeed  $0^\circ$ , so some amount of camera-target alignment error should be expected in those corresponding data points. Furthermore, because distances were measured with an approximately 7 m measuring tape, some amount of experimental error should be expected in distance measurements past 7 m. Given these uncertainties, the pixel footprint of the reference emitter plate reported in the following table are approximate and largely based on visual inspection of the infrared images.



When the camera-to-target distance is large, the target (in this case, the 10 cm × 10 cm reference emitter plate) appears very small in the recorded image (Table 7). This poses challenges to accurate determination of the target’s temperature from the thermal image.

Let us use the term “footprint” to describe the collection of pixels within a thermal image that map onto some target of interest whose temperature we would like to determine. The most straightforward method of calculating target temperature from a thermal image is to average the temperatures of the pixels inside its footprint. We know that each pixel within a thermal image represents the output of an individual sensing element in the camera’s focal plane array. In order for the temperature reported by that pixel to be considered an accurate representation of the target, the target must completely fill the sensing element’s FOV. That is to say, if the sensing element “sees” not just the target but also its background, its output correspondingly represents a blend of both target and background and should therefore not be used to calculate the target’s temperature. Pixels at the footprint perimeter are understandably the most questionable. When the target footprint is large, the area-to-perimeter ratio is also large and any inaccuracies in pixel choice at the footprint border has negligible effect on the average temperatures. When the target footprint is small, however, the area-to-perimeter ratio is also small and pixel choice at the footprint border becomes a significant source of bias.

Of course, a different calculation method would be subject to other concerns. If the target is expected to be significantly warmer or cooler than its background, we can bypass any bias from perimeter pixel choice by assigning target temperature to be, respectively, the maximum or minimum temperature within the target footprint. This method, however, is instead vulnerable to outlier results from bad pixels.

Therefore, we analyzed these images using both methods. In our “area mean” method estimating average temperature  $T_{\text{mean}}$ , we first identified the pixels in an image we believed to represent the reference emitter and simply calculated the average pixel value. In our “area max” method estimating maximum temperature  $T_{\text{max}}$  or minimum temperature  $T_{\text{min}}$ , we simply extracted the maximum or minimum pixel value instead.




Figure 10. Photographs of the reference emitter taken at distances of (a) 34.1 m, (b) 6.4 m, (c) 2.2 m, and (d) 1.0 m.

Table 7. Reference emitter pixel footprint for each distance assessed.

Distance (m)	Reference emitter footprint in image (estimated)	
	Pixel length $d$	Pixels in area ( $d^2$ )
1.0	75	5,625
2.2	35	1,225
4.0	19	361
6.4	11	121
10.6	7	49
18.1	4	16
25.9	2	4
34.1	1	1

#### 4.6 Berkeley Lab experiment 4: Trials with a human subject

In this experiment, a human subject was seated next to the reference emitter, head roughly 0.5-1.5 m from the camera. Contact sensors were attached with medical tape to the following two points on the subject’s face: (1) middle of the forehead, and (2) tip of the nose (Figure 11). Each set of contact sensors included one bare epoxy bead thermistor (hereafter referred to as “TH”), one bare wire thermocouple (hereafter referred to as “TC1”), and one thermocouple with the bare sensing end attached to a thin foil plate (3 mm square; hereafter referred to as “TC2”). While effort was made to ensure good contact between the sensors and skin at the start of a measurement event, this could not be verified. Good contact between the skin and sensors was complicated by interference from skin oils as well as the curvature of the features, particularly the nose.

Several sets of measurements were captured, each with the camera positioned at a different viewing angle, distance, or human subject comfort rating (Figure 12). Each set consisted of a 30 s thermal image time series (1 s intervals) of the subject’s bare face, followed by a 30 s contact temperature time series with the sensor placement described above. Due to limited camera programming options, the camera’s NUC scheme was suppressed during the 30 s of thermal image measurements. In order to minimize errors, this experiment was only performed after the camera’s native NUC frequency had slowed to every 20 min.

Forehead temperatures were extracted from the thermal images via the “area max” method (hereafter referred to as “ $T_{IR,max}$ ”). Nose tip temperatures were similarly extracted via an “area min” method (hereafter referred to as “ $T_{IR,min}$ ”). All radiatively assessed skin temperatures were further offset-corrected (“ $T_{IR,corr}$ ”) via the reference emitter (set point 35.00 °C).




Figure 11. Contact sensors measurements (a) and thermographic measurements (b) of the subject's forehead and nose.




Figure 12. Six images of the human test subject were assessed, encompassing different viewing angles in both the vertical ("V") and horizontal ("H") plane.

## 5 MoviTHERM experiments

MoviTHERM collected narrow-view data using three experiments, broadly defined as follows:

- Blackbody reference tests
- Human subject tests
- Inanimate object and surfaces tests

## 5.1 Sources of uncertainty in infrared thermography

Any experiment description must include an understanding of the sources of measurement uncertainty. Minor discrepancies will be observed between the theoretical expectations and the actual measurement results collected during testing. This measurement uncertainty section tries to account for these deviations.

### 5.1.1 Measurement uncertainty due to camera limitations

The first source of measurement uncertainty can be attributed to the FLIR A315 camera itself—or more specifically, to the camera’s microbolometer. A microbolometer is an electronic device that measures the intensity of the TIR energy striking its surface. The A315’s microbolometer is an array of  $320 \times 240$  thermally isolated cells or “pixels”. Each pixel absorbs incident radiation; as this radiation is absorbed, the pixel temperatures increase, resulting in a change of electrical resistance. A thermal picture or thermograph is constructed by mapping the corresponding resistance measurements for each pixel in the two-dimensional sensor array.

One unavoidable byproduct of the camera’s electronic design is an unintended signal degradation due to stray heat within the camera. Like all electronic circuits, heat is generated inside the camera (and in the sensor itself) as a natural consequence of electrical resistance. This internal resistance causes the sensor, and the electronic components near the sensor, to warm the sensor. The pixels are still influenced primarily by the incident radiation, however the ratios between the background temperature and the pixel temperatures in the scene begin to shift. Background temperature changes degrade the signal-to-noise ratio, and measurement errors are introduced as the resistance measured against a moving baseline is no longer reflective of just the incident radiation alone. Errors tend to increase over time as electrical resistance causes the background temperature at a certain time to increase from the initial background temperature. The effect of this error appears as a gradual drift in temperature values.

To compensate for this drift, the camera periodically initiates a NUC operation. During a NUC operation, a metal shutter or “flag” with a known temperature is moved in front of the sensor. A correction is applied to the output value of each pixel to eliminate or “zero out” the accumulated background errors. Each NUC operation essentially resets the sensor to the nominal, calibrated output values. However, between NUC cycles, the temperature will drift for reasons discussed in the previous paragraph. The instantaneous error and rate of temperature drift depends on several factors. If the target temperature is substantially higher than the sensor temperature, the signal-to-noise ratio will be high, and the impact of sensor temperature changes will be minimal. However, if the sensor temperature and target temperature are close in value, more measurement uncertainty will be seen. In the context of the HCJR project, we can expect the human target temperature to be similar to the sensor temperature, so the impact of sensor drift must be taken into consideration.

According to FLIR, the camera manufacturer, the temperature accuracy of the A315 camera is  $\pm 2\text{ }^\circ\text{C}$  or  $\pm 2\%$  of the full scale. The relative pixel-to-pixel temperature accuracy specification is much better than this; the specified NETD (Noise Equivalent Temperature Difference, a.k.a. “Thermal Contrast”) is 50 mK. This means that a temperature difference as small as 50 mK can

be resolved by the sensor. However, due to the issues discussed previously, the temperature readings reported by the camera are only accurate to within  $\pm 2$  °C (or  $\pm 2\%$  of the reading). To improve the accuracy of the readings, a target of known thermal emittance and temperature can be located in the image frame. For this function, a blackbody reference emitter provides a useful reference. Using the difference between the actual and measured temperature of the blackbody target, a temperature correction can be applied to every pixel in the image. While the actual temperature is only accurate to  $\pm 2$  °C, the in-frame, relative temperature accuracy is specified at 50 mK or 0.05 °C. If the absolute temperature of one pixel is known, the absolute temperature of every other pixel can be derived using the relative temperature difference.

To achieve this level of accuracy, the temperature and thermal emittance of the reference target must be known with a high degree of accuracy. Furthermore, a reliable measurement sample comprised of a statistically significant number of pixels must be captured in the image frame. It follows that any uncertainty in the blackbody temperature measurement will create uncertainty, even in corrected measurements.

### 5.1.2 Measurement uncertainty due to the blackbody temperature control limitations

A second source of measurement uncertainty in the testing at MoviTHERM was due to the Thermoworks IR-500 blackbody temperature reference (see Section 3.2). When evaluating and/or correcting for camera drift, we need to assume a stable point of reference. For analysis, the blackbody setpoint was assumed to be an accurate and stable point of reference; however, no device is perfect.

According to the manufacturer's specifications, the Thermoworks IR-500 blackbody has a range accuracy of  $\pm 0.8$  °C and a stability limit of  $\pm 0.1$  °C for temperature setpoints  $\leq 100$  °C. Since all testing was done with a temperature setting of 55 °C, these are the accuracy specifications applicable to the MoviTHERM tests.

It's important to note that the uncertainty from the blackbody controller must be considered additive to the uncertainty from the camera, and that these are random, rather than systemic errors. At any given point in time, the blackbody PID temperature controller offers  $\pm 0.1$  °C. Any temperature deviation from the user-entered setpoint could be a systemic error, however a more accurate measurement reference would be needed to characterize this error. In the absence of a more accurate reference, we are left with  $\pm 0.8$  °C of uncertainty.

As we have seen, the camera contributes some measurement instability. The blackbody also contributes some instability. A third reference may help to characterize the proportions contributed by the blackbody and camera respectively, but any additional device will also introduce some uncertainty that must be characterized.



## 5.2 MoviTHERM experiment 1: Blackbody reference tests

The blackbody emitter is a temperature-controlled test surface of known emittance. The blackbody unit provides a digital readout of the current temperature of the emitter element in real-time while an integrated controller maintains the user-defined temperature setpoint. Due to the greater potential for errors inherent in ad hoc contact thermocouple testing, the blackbody controller temperature readout was believed to be the most reliable reference available for testing.




Figure 13. The Thermoworks IR-500 blackbody reference emitter was imaged at normal ( $0^\circ$ ) view angle at distances of 1 to 4 m.

Test images of the blackbody emitter were collected at varying distances and angles, as depicted in Figure 13 and Figure 14. Results relating to IR temperature accuracy and variability were largely derived from the blackbody reference testing.




Figure 14. The Thermoworks IR-500 blackbody reference emitter was imaged at view angles of  $0^\circ$  to  $75^\circ$ .



### 5.3 MoviTHERM experiment 2: Human subject testing




Figure 15. Sample of TIR of Human Subject Testing image.

The HCJR project seeks to extract temperature data from one or more human subjects and utilize this temperature data to determine the relative comfort level of the various human subjects in an imaged scene. In order to effectively assess comfort, the comfort model must obtain accurate human temperature measurements. The MoviTHERM human subject testing focused primarily on temperature measurements as distance changed (Figure 16), as well as the potential effect of sensor resolution on human subject temperatures.

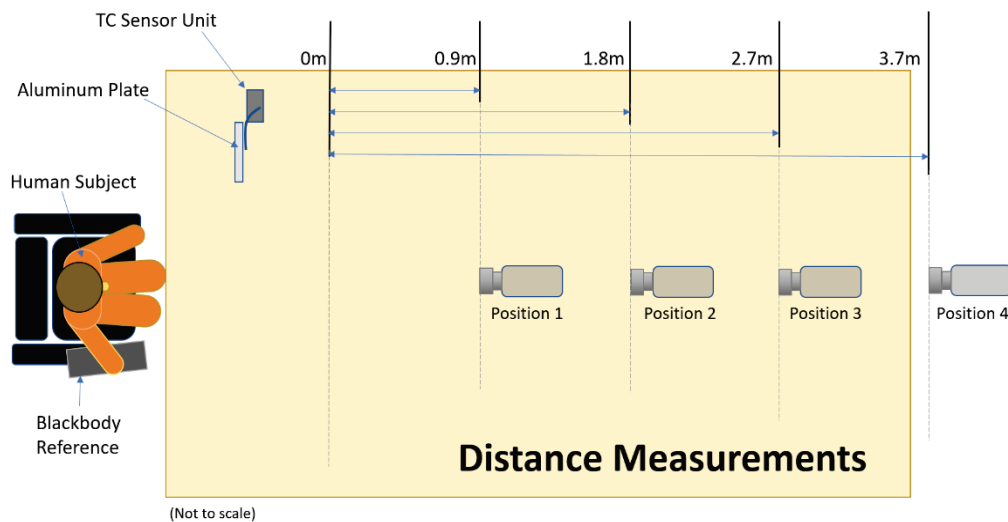


Figure 16. The human subject and reference emitter were imaged at a view angle of  $0^\circ$  and distances up to 4 m.

Angular measurements were intentionally excluded from the MoviTHERM human subject sample set because human surfaces are typically curved by nature. It was presumed that a high number of variables (uncontrolled body position, head and/or arm rotation)—combined with no suitable human reference sensor in our lab—would make the human subject angular results less definitive than the blackbody angular testing outlined in the previous section.

The thermal emittance of human skin is known to be stable and predictable. Therefore, it is believed that angular test results derived from the more tightly controlled blackbody angular tests can be successfully extrapolated to human skin surfaces, and will be more numerically accurate than samples collected from human subjects at arbitrary angles. By limiting the angular measurements to the flat, blackbody surface, we don't have to worry about the effect of the angle between the normal of the curved surface with respect to the angle of view from the camera.

Temperature measurements derived from the outer extremities of curved skin surfaces (sides of the face, arms, hands, etc.) are altered as the camera capture angle deviates from the surface normal. Furthermore, the size and location of selected sample regions can produce significantly different results at different test distances using the medium resolution sensor. We believe it would have been difficult to isolate the effect of angle versus the impact of the other variables.

#### 5.4 MoviTHERM experiment 3: Inanimate objects and surface testing

The HCJR sensor will see both human and inanimate surfaces. Inanimate objects will be composed of various materials and therefore can be expected to exhibit varying degrees of emittance, thermal reflectance, and in some cases, thermal transmission. In most cases, actively heated or cooled objects may also be present, such as electronic devices, heating vents, or cooling vents. A variety of common office items were assembled for testing as shown in Figure 17.




Figure 17. Inanimate objects and surfaces tests subjects—Color (a) and TIR (b) images of the scene.

The purpose of the inanimate object testing was to capture temperature measurements at various distances and angles. In this way, the inanimate object testing was similar to the blackbody distance and angle testing introduced earlier (refer to Figure 13 and Figure 14). In addition, test data was captured over time as the temperature of the room evolved over several hours. The objective of this test was mainly observational—tracking the temperature differences of various common office objects composed of different materials, and with surfaces of unspecified thermal emittance and angular position. While this test was less rigorous in terms of deriving useful numerical data, it

was intended to illuminate the degree to which undefined, inanimate objects can be used to derive useful background temperature data from the scene.

Absolute, instantaneous temperature values from the TIR camera may be of limited use if little is known about the physical properties of various object in the scene. This test was intended to explore the practicality of using temperature changes of fixed objects (i.e. furniture, walls) to track temperature changes over time.

## 6 Berkeley Lab results

### 6.1 Emittance

To achieve any degree of accuracy in radiatively sensed temperature measurements, the emittance of the target must be known. It can either be directly measured or looked up in a reference table.

To explore the magnitude of errors caused by assigning an incorrect value to as a target's emissivity, we arbitrarily chose an image of the reference emitter (with known emittance  $\epsilon_{\text{true}} = 0.95$ ) from Berkeley Lab experiment 2 (viewed at an angle of  $0^\circ$  and distance of 1 m) and varied its FLIR Tools emittance input by up to  $\pm 0.05$  (or 5 points). Figure 18 plots the changes in apparent temperature corresponding to differences in assigned emittance.  $T_{\text{true}}$  is the radiatively sensed temperature output when the emittance is set to the manufacturer-specified value ( $\epsilon_{\text{true}}$ ), and  $T_{\text{apparent}}$  is the radiatively sensed temperature output when the emittance is set to an incorrect value ( $\epsilon_{\text{guess}}$ ).




Figure 18. The difference between apparent radiatively sensed temperature and true reference emitter temperature induced by errors in assumed emittance.

Recall that the manufacturer-reported accuracy of the camera is roughly  $\pm 1$  °C. In these examples of high-emittance surfaces at temperatures upwards of 30 °C, errors in assumed emittance on the order of several points correspond to errors in reported temperature of roughly several tenths of a degree.

For the HCJR use case, we expect errors resulting from incorrect assignment of emittance to be small, given that the emittance of human skin is known to be 0.98. However, some examples of common situations where this may not necessarily hold true include when the skin is covered by another material, such as water or sweat, or when the subject is viewed at an oblique angle, where directional emissivity becomes a concern.

## 6.2 Distance

Distance becomes an important variable to measurement accuracy when the atmospheric transmittance cannot be assumed to be unity. Some scenarios where this is the case include when the ambient relative humidity is very high, or when the subject is very far from the camera.

In the following figures, we report the results of Berkeley Lab experiment 3, which imaged the 35.00 °C reference emitter at various distances.  $T_{\text{mean}}$  data points were determined via our “area mean” method;  $T_{\text{max}}$  data points were determined via our “area max” method. In the  $T_{\text{mean,corr}}$  and  $T_{\text{max,corr}}$  data sets we supplied the camera software with distance information; in the  $T_{\text{mean}}$  and  $T_{\text{max}}$

data sets we left this input at shortest distance in our data set (1 m). Figure 19 plots the radiatively sensed temperatures versus distance; Figure 20 focuses on the effect of applying distance-based atmospheric transmittance corrections; and Figure 21 compares the radiatively sensed temperature results to the apparent size (in terms of pixel footprint) of the reference emitter in the thermal image. Because the emitter plate is square, we report its pixel footprint in units of pixel side lengths (“characteristic length”) rather than the total pixel area.

We observe broadly from Figure 19 that, as expected, as the distance between the subject and camera increases, the apparent radiatively sensed temperature of the subject deviates from the known true value. The “area mean” processing method in all cases produced a lower apparent temperature than the “area max” method, the degree of which increases significantly past 10 m. This is unsurprising, as this marks the point where the pixel footprint of the reference emitter falls to roughly a  $4 \times 4$  grid. At distances between 1 m and 5 m, there is no significant difference in the temperature result when the camera is not supplied an accurate distance for its atmospheric transmittance correction; past 5 m, however, discrepancies of up to a few tenths of a degree are observed. Finally, regardless of which processing method was used or whether a proper atmospheric transmittance correction was applied, distances past 10 m resulted in errors greater than  $1^\circ\text{C}$ .

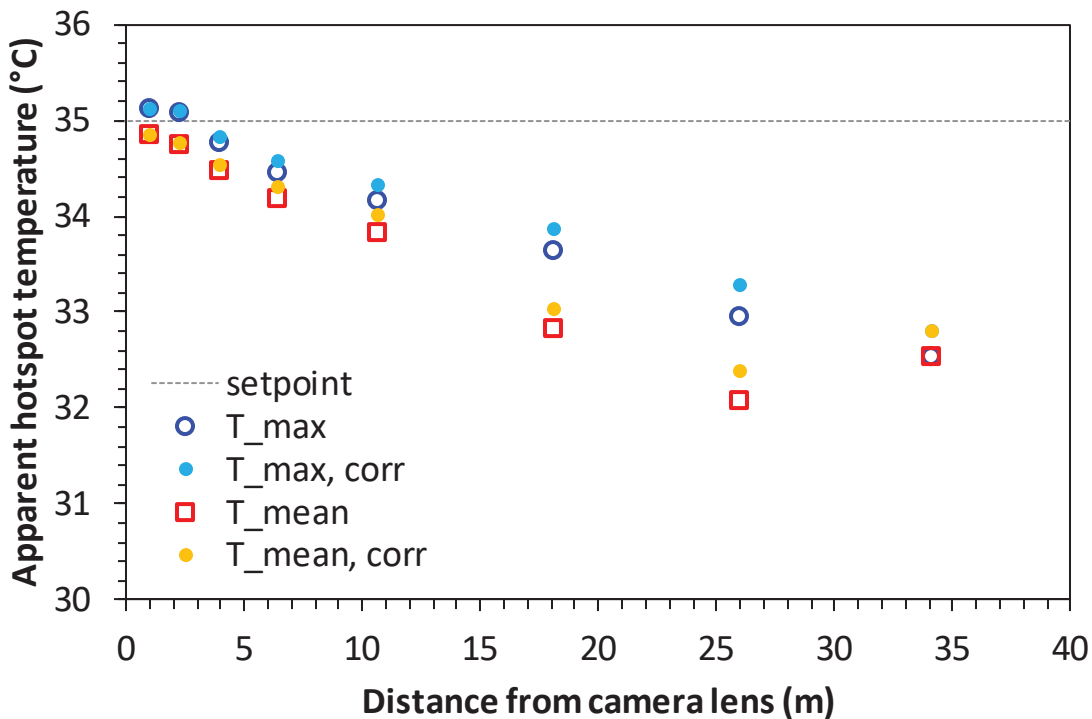


Figure 19. The apparent radiatively sensed temperatures determined by photographs at varying distances.

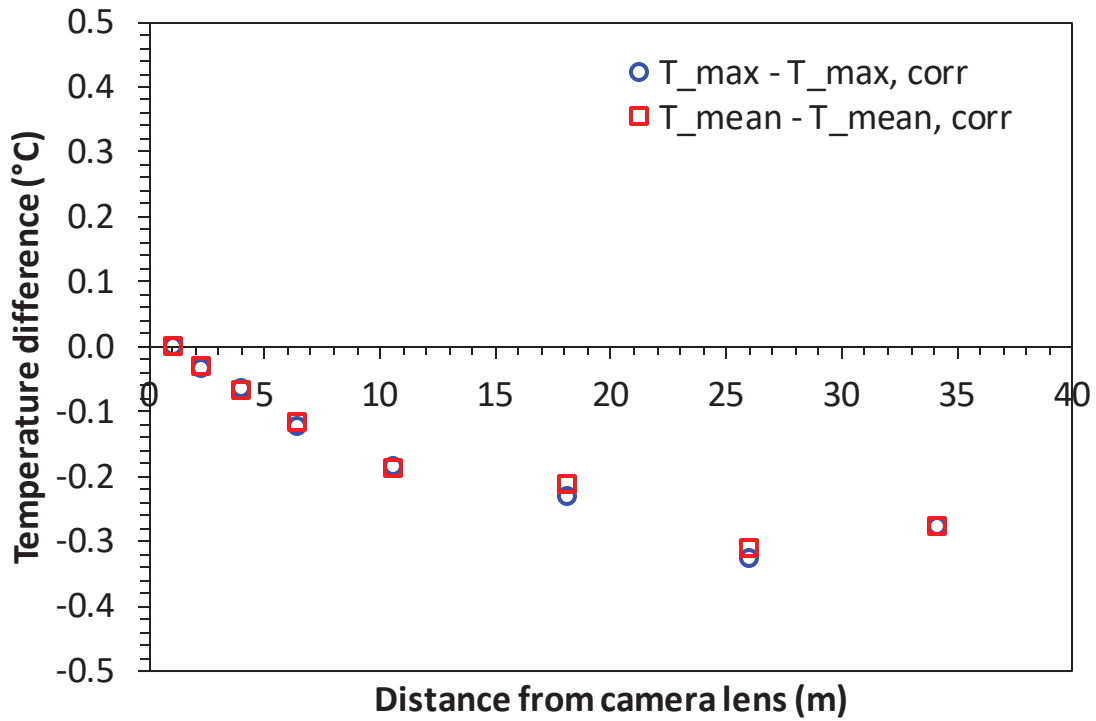


Figure 20. Change in the apparent temperature ( $T - T_{\text{corr}}$ ) when a distance (atmospheric transmittance) correction is not applied.

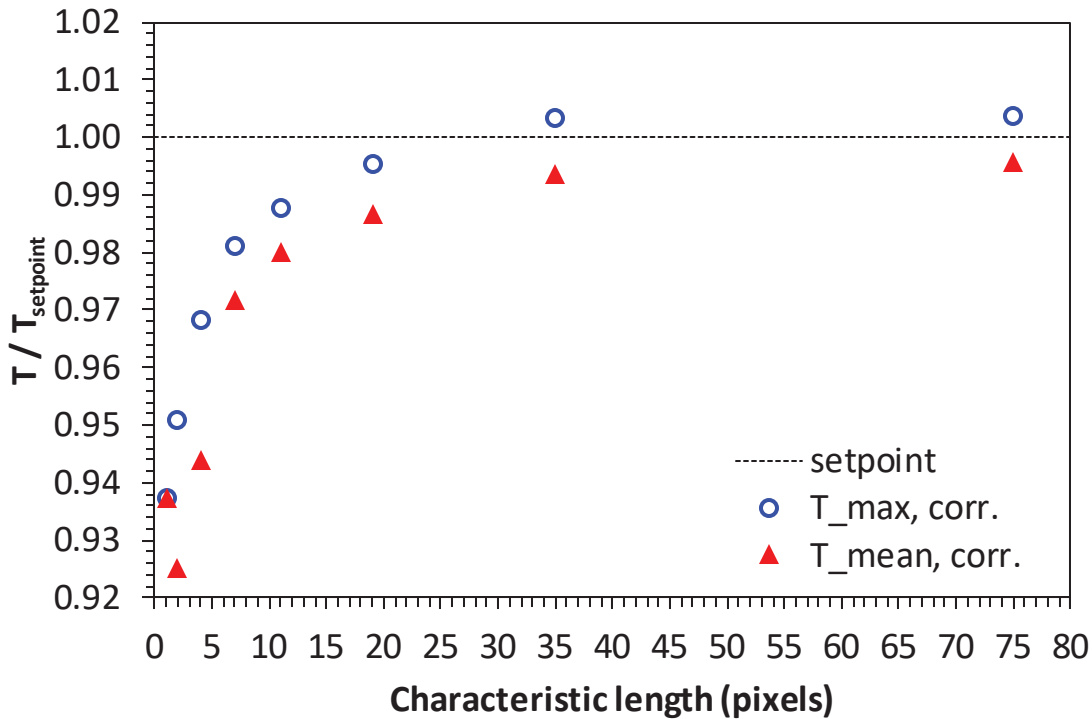


Figure 21. Apparent radiatively sensed temperature (expressed as a percentage of the known true value) versus image pixel footprint of the reference emitter.

The camera manufacturer recommends that accurate measurements require the footprint of the target in the image to be no smaller than  $3 \times 3$  pixels. In this trial, we find that an image footprint of  $4 \times 4$  pixels produces a temperature result nearly  $1^\circ\text{C}$  below the that of the largest footprint, and that the influence of pixel footprint size only effectively disappears when it exceeds  $35 \times 35$  pixels.

It should be noted that the targets that will be of interest in the HCJR use case—the tip of the nose, the center of the forehead, and the hand—are much smaller than the reference emitter used in this experiment. If the  $10\text{ cm} \times 10\text{ cm}$  reference emitter appears to be  $3 \times 3$  pixels large at a distance of approximately 10 m, then a  $1\text{ cm} \times 1\text{ cm}$  nose tip would meet the same threshold at approximately 1 m. A camera with poorer resolution would require an even shorter distance to meet the same threshold.

### 6.3 Viewing angle

In Berkeley Lab experiment 1, we hoped to investigate the magnitude of errors in TIR measurement as the camera viewed the target at increasingly oblique angles. These errors arise as a result of the following:

- Effective emittance is unknown because emittance varies with viewing angle.



- As viewing angle increases, the image footprint of the target decreases in size. At oblique viewing angles, the target may not even be visible.
- For very large objects, there may exist a non-negligible difference in camera-to-target distance between the its near and far edges.

Of the items listed above, only the first two are of concern in the HCJR use case. Since the reference emitter has a known emittance and temperature, we simply calculated the “area mean” temperature of the emitter plate in each image and plotted it as a function of viewing angle (Figure 22).

Because emissivity decreases with increasing viewing angle, we expect radiatively sensed temperatures to vary similarly with viewing angle. Indeed that is the case; we observe that the apparent temperature of the reference emitter decreases slowly with viewing angle up to approximately 60°, after which it changes rapidly.

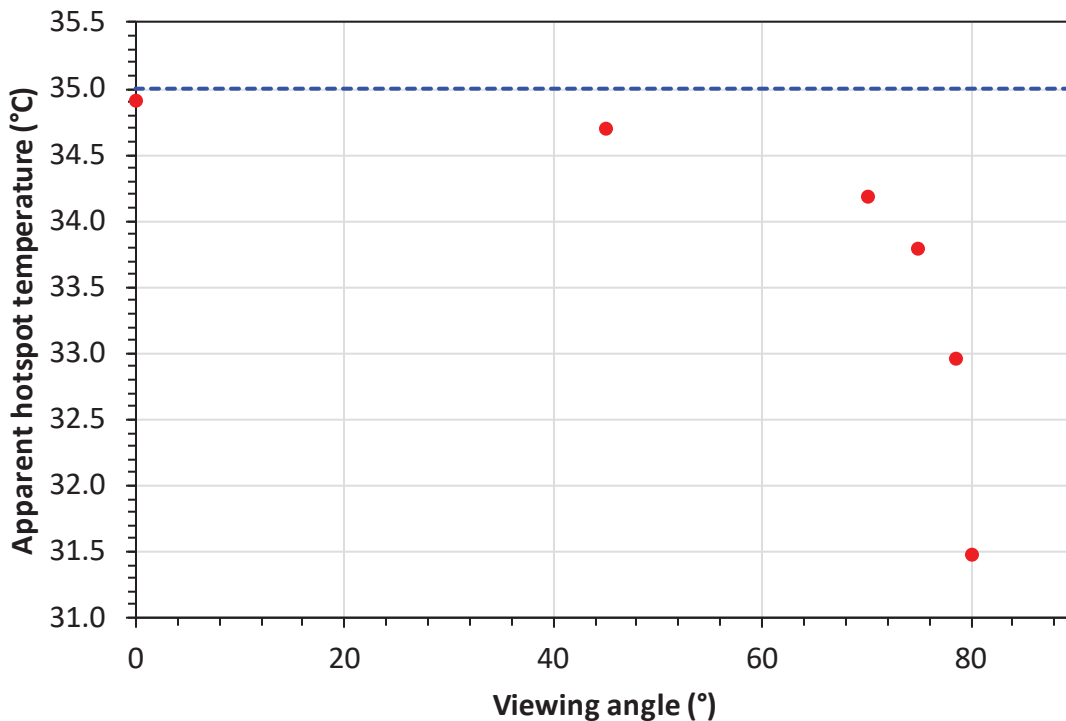


Figure 22. Apparent radiatively sensed temperature as function of viewing angle.

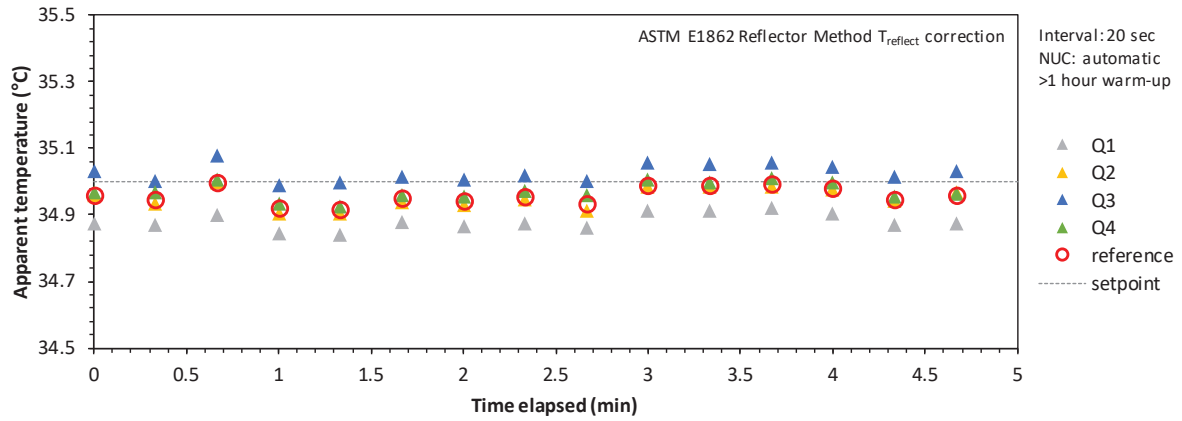
We expect the practical impact of directional emissivity to be mitigated by the fact that the body parts of interest—the nose, the forehead, and the hands—are largely curved in nature. However, we do expect viewing angle to significantly impact the pixel footprint size.

## 6.4 Measurement drift


Temporal drift in camera output caused by temperature instability of the focal plane array is unavoidable, and many thermal cameras employ a shutter-based periodic NUC routine to address this source of measurement error. In an ideal setup, the NUC event would trigger before each photograph, but we acknowledge that such a strategy might not be practical in the HCJR use case. Each NUC event takes the camera “offline” for at least a few seconds, making on-demand photo capture more difficult. Additionally, triggering the mechanical shutter too often may significantly degrade its useful life.

The goal of experiment 2 was to explore the magnitude of (uncorrected) temporal drift as well as the effectiveness of the NUC routine in maintaining measurement accuracy. The next four figures show time series of absolute and differential radiatively sensed temperatures under four different camera conditions:

- (1) stable (low drift) condition with active NUC routine (Figure 23)
- (2) stable (low drift) condition without NUC routine (Figure 24)
- (3) unstable (high drift) condition with active NUC routine (Figure 25)
- (4) unstable (high drift) condition without NUC routine (Figure 26)

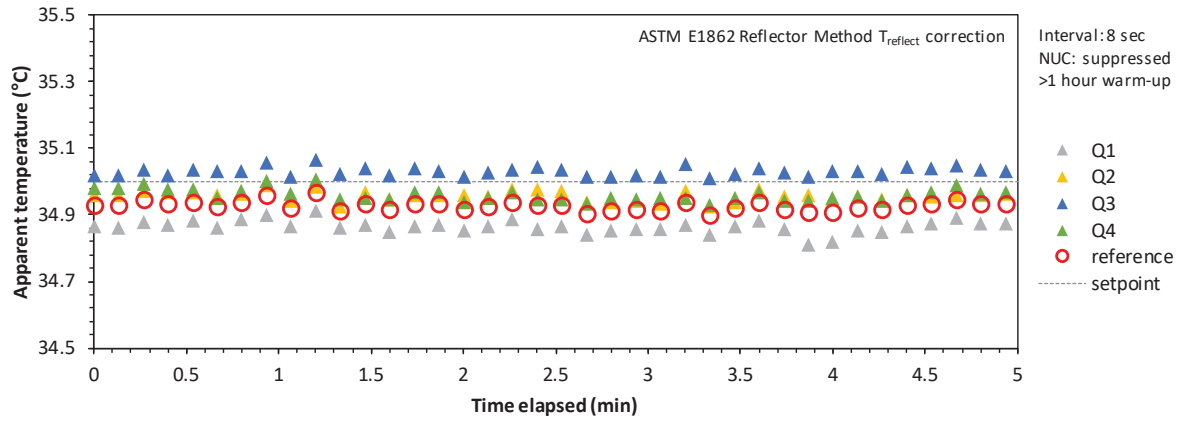


(a)




(b)

Figure 23. Time series of (a) absolute and (b) differential temperatures captured when the camera was allowed a 1 h warm up period. The automatic NUC routine was allowed to run.




(a)




(b)

Figure 24. Time series of (a) absolute and (b) differential temperatures captured when the camera was allowed a 1 h warm up period. The automatic NUC routine was suppressed.

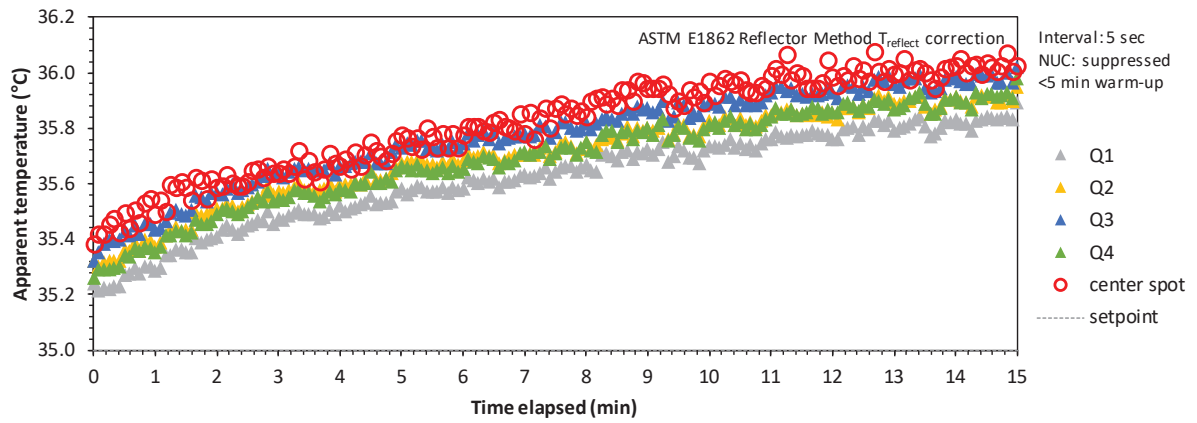


(a)




(b)

Figure 25. Time series of (a) absolute and (b) differential temperatures captured when the camera was not allowed to warm up. The automatic NUC routine was allowed to run.



(a)



(b)

Figure 26. Time series of (a) absolute and (b) differential temperatures captured when the camera was not allowed to warm up. The automatic NUC routine was suppressed.




Figure 27. The radiatively sensed ( $T_{\text{wall,IR}}$ ) and contact-sensed ( $T_{\text{wall,TH}}$ ;  $T_{\text{wall,TC}}$ ) temperatures of the background wall in (a) condition 3 and (b) condition 4.  $T_{\text{wall,IR,corr}}$  is simply  $T_{\text{wall,IR}}$  corrected with a stable external reference.

From these four figures, we can observe the following:

1. When the camera was allowed to warm up, the time series of both absolute and differential temperatures showed temporal stability both with and without the NUC routine active. Furthermore, all reported absolute temperatures remained within 0.2 °C of the reference emitter setpoint.
2. When the camera was not allowed to warm up, the time series of absolute temperatures showed temporal drift regardless of whether the NUC routine was active. Without NUC, camera output was observed to drift approximately 0.6 °C over 15 min. With NUC, the camera output was observed to drift approximately 0.4 °C over 15 min. Although each NUC event did indeed shift camera output closer to the known true value (35.00 °C), no measurements in either trial reached within 0.2 °C of the reference emitter setpoint in the 15 min of sampling. On the other hand, regardless of whether the NUC routine was active, the differential temperatures remained remarkably stable.



3. Comparing the radiatively sensed and contact-sensed temperatures of the background wall confirm that the drift and fluctuation patterns in the thermal image time series (Figure 25; Figure 26) are not real. The temperature time series recorded by both contact sensors are stable over time.
4. When we correct the radiatively sensed temperature of the background wall ( $T_{\text{wall,IR}}$ ) using the reference emitter, the drift and fluctuation patterns largely disappear (Figure 27) and the temporal trend appears as stable as the contact-sensed temperatures. Furthermore, the reference-corrected radiatively sensed wall temperature ( $T_{\text{wall,IR,corr}}$ ) differs not more than half a degree Celsius from the highest-accuracy contact sensor (thermistor “TH”).
5. Regardless of the condition of the camera, and whether or not a NUC routine was active, the differential temperatures remained remarkably consistent. Note, however, that these differential temperatures were calculated from relatively large pixel areas and are thus resistant to the influence of “bad” pixels; we would expect to see more variability if each quadrant were orders of magnitude smaller.

This short experiment allows us to draw two conclusions:

1. The NUC routine alone does not completely correct for measurement drift (at least over time periods of 15 min).
2. Differential temperatures appear to be largely immune to drift.
3. Applying an external reference correction is necessary for accurate absolute TIR measurements regardless of whether an NUC routine is active.

The second conclusion is comforting, given that the most important measure in the HCJR use case is the differential temperature between body parts.

## 6.5 TIR and contact sensor measurements of human facial features

It is difficult to draw any conclusions from the results of this experiment because contact sensor measurements of skin are, in general, subject to their own myriad problems and cannot be considered a reference against which to compare the performance of the radiatively sensed results. The sensor shape, contact area, and attachment mechanism (e.g., tape, weights) are several of many setup variables that can perturb the skin in different ways and potentially induce an unrepresentative result (Macrae et al., 2018).

Combined with the fact that in this specific experiment good contact between the sensors and skin could not be conclusively confirmed, and that the human subject showed a slight degree of skin irritation to the medical tape adhesive, we caution against using the following results as a gauge of IR sensor performance until additional trials are performed.

Despite the aforementioned sources of experimental error, in four out of six image sets the radiatively sensed absolute temperatures of the forehead and nose were within 1 °C of at least one of the contact-sensor values. In three of the six image sets, the radiatively sensed differential temperatures differed from the contact sensor results by only a few tenths of a degree Celsius.







Figure 28. For each of the image sets assessed (panel *a*; labeled A through F), the output of the contact sensors (TC1, TC2, and TH) were compared to the radiatively sensed temperatures ( $T_{\text{IR}}$  and  $T_{\text{IR,corr}}$ ). Results for the forehead are shown in panel *b*; results for the nose are shown in panel *c*; and differential temperature results are shown in panel *d*.

## 7 MoviTHERM results

### 7.1 A note on sources of uncertainty related to the blackbody

For our testing, the blackbody internal reference was more accurate and stable than any other measurement device we had on hand. However, there is also the issue of temperature uniformity. Due to the nature of the temperature distribution across the emitter surface, and the rapid temperature falloff at the edges of the emitter, the sample region and the particular pixels sampled will impact the results, as shown in Figure 29.

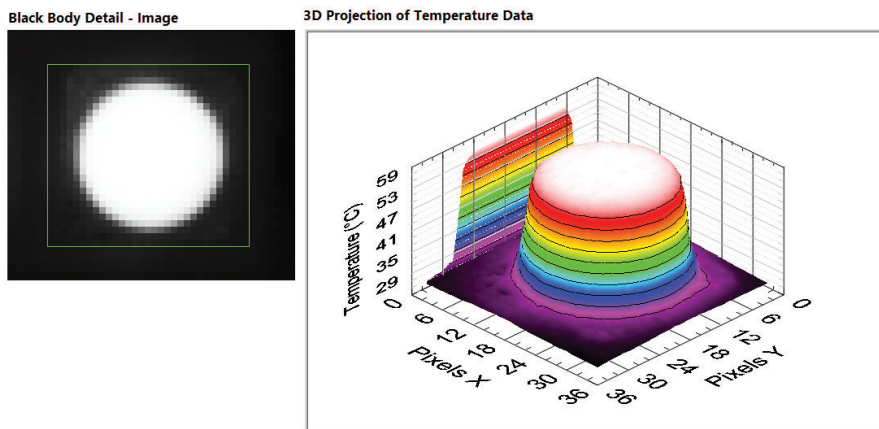


Figure 29. Blackbody temperature uniformity. Temperature data 3D projection as a function of the square region number of pixels per side, starting at 36 × 36, which corresponds to the green square depicted in the black body detail image.

It is important to understand that the blackbody temperature values are impacted by the region selected. Differences to the average value of the region, and the standard deviation of the measurement will be different if a different set of pixels is selected, as shown in Figure 30.




Figure 30. Temperature variation due to sample size and location variability.

It is further noted that lower resolution appeared to affect the blackbody temperature measurement accuracy (as noted later in the results section). So while it is possible to correct for a blackbody measurement, the corrected value can still be affected by uncertainty caused by physical characteristics of the blackbody, or distance and resolution limitations.

## 7.2 Blackbody reference test results

Figure 31 shows the results obtained for both the distance and angle of view datasets. Different distances go from 1 m to 4 m in steps of 2 m, while angles go from 0° to 75° in steps of 15°. For the measurement of the temperatures of the blackbody working surface, rectangular regions were chosen by hand trying to obtain the maximum number of pixels without interfering with edges or complicated areas. As explained in Section 7, this is not a trivial task. Table 8 details the size of this regions, so it can be considered in the discussion of the results. The region pixel values are obtained averaging 5 s of images taken at 1 frame per second (fps) in order to further denoise the data.

Table 8. Size of the largest possible blackbody regions selected for the temperature study.

Size of regions in distance dataset		Size of regions in angle dataset	
Distance (m)	Number of pixels	Angle	Number of pixels
1	120	0°	132
2	25	15°	99
3	6	30°	70
4	2	45°	50
		60°	16
		75°	8

The overall standard deviation of distances for all data collected is 0.48 m. Temperature measurements of the blackbody emitter set to 55 °C appear to drop approximately 0.5 °C per meter as the distance between the camera and the blackbody increases.

The overall standard deviation of angles for the angle view dataset is 4.1°. Temperature measurements seem to remain unchanged at angles less than about 40°. At angles greater than 40°, temperature decrease accelerates as angles increase. At an angle of 75°, temperature loss may be greater than 10 °C.

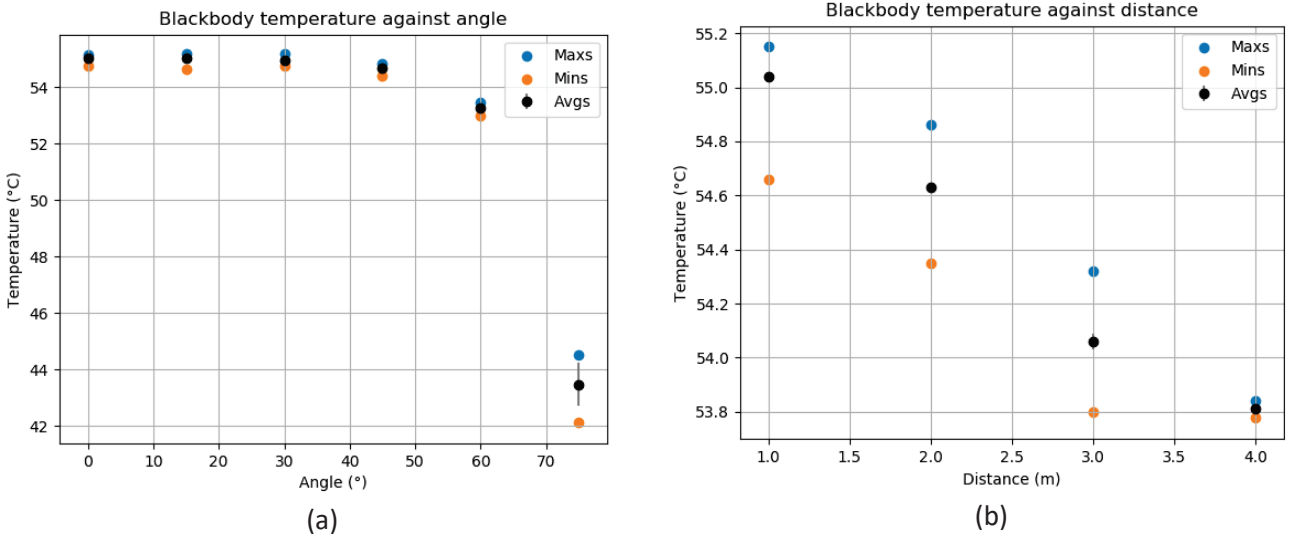


Figure 31. Blackbody reference test results. Temperatures of the region of interest characterized as the maximum, average and minimum values of the region as a function of (a) angle of view from the camera center to the normal of the blackbody surface and (b) distance from the camera center to the blackbody center.

### 7.3 Blackbody reference test discussion

The results show a descending trend for both distance and angle independent variables, which implies there is a loss of temperature measurement accuracy. In this section, we are going to try to explain the possible causes and factors of this loss.

Distance affects temperature in what looks like an almost linear loss, of around half degree Celsius or around 1% of the absolute temperature per meter. The most obvious effect of distance is the fact that the further an object is, the smaller the region of interest becomes. With the loss of resolution, the boundary of warm objects gets averaged with the colder background, and that, together with the human error in the choosing of the region, leads to an important decrease in temperature. In fact, some works (Neto et al., 2006) conclude that in practice distance only factors as a loss of resolution issue.

However, besides the resolution, there's another possible source of error: the atmospheric transmission. The atmosphere, particularly water vapor and carbon dioxide, between the thermal camera and the target object attenuates infrared radiation emitted by the object. This attenuation is noticeable, even though it is strongly dependent on the wavelength of the camera, it is still significant even for long wave cameras such as the one we are using. For example, in normal outside conditions, at 10 m the transmission is 96%. A 4% loss of temperature with distance is significant, and the camera manufacturer doesn't recommend its use with distances further than 10 m (Modest, 2013).

To illustrate this, we show the transmission coefficients of a simplified model of atmospheric transmission for our wavelength and conditions in Figure 32. At 4 meters, there is a 99% transmission, which explains part of our error, but not all of it. We believe the addition of the loss of resolution, atmospheric attenuation and systematic errors does.

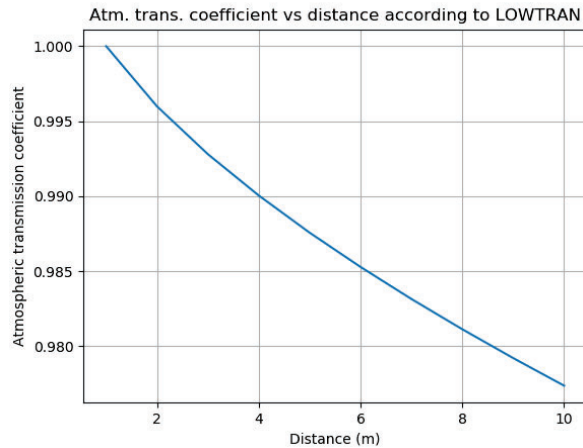


Figure 32. Decay with distance of the atmospheric transmission coefficient according to the simplified transmission model LOWTRAN (Minkina et al., 2016) in normal conditions (ambient temperature 15 °C and 50% humidity), calibrated at 1 m.

In the case of the angle of view or of deflection, temperatures decrease non-linearly as the angle of camera deflection increases above 40°, which the expected loss of resolution cannot alone explain.

Nunak et al. (2015) and FLIR Systems, Inc. (2016) show that temperature measurements can be expected to remain relatively constant for angles of less than about 40°. Reflectivity changes with the angle of view as shown in the Fresnel equations. Fresnel obtained the equations, whose formulation is outside of the scope of this report, that are used to predict the reflectance of non-rough surfaces, which depends only on the refractive index and the angle of incidence. The effect of Fresnel reflectance is not very significant for conductive materials, but for non-conductive or dielectrics is very important, and it goes up non-linearly, as Figure 33 depicts.

Measuring the temperature of very reflective surfaces challenging because of background reflections and because of the corresponding low thermal emittance. Thermal reflectivity ( $r$ ) is related to thermal emittance ( $\varepsilon$ ) by

$$r = 1 - \varepsilon \tag{1}$$

according to Kirchhoff's law of thermal radiation for opaque materials, which implies that for a camera with normal settings—that is, calibrated for a particular thermal emittance that is usually in the 0.85-0.95 range—temperatures of surfaces viewed with higher angle of view will appear much lower.






Figure 33. Relationship between the angle of view and the reflectance coefficient of a surface with real refractive index  $n=1.5$  from air according to the Fresnel equations.

### 7.4 Human test results

In the human subject test temperature measurements were obtained as distance between the camera center and the chair changes for distances between 0.9 m (3 ft) and 3.7 m (12 ft). Figure 34 shows the temperatures obtained for the blackbody reference as well as the standard deviation of all the different features through distance. As we can see, the standard deviation in the hand measurements is the highest, followed by the nose. Cheek measurements are below the threshold of the blackbody.




Figure 34. Human test results. Panel *a* shows maximum, average and minimum values of the blackbody region as a function of distance to the camera center, used as a correction for the temperatures of the rest of features studied, whose standard deviations are shown in panel *b*.

Furthermore, Figure 35 shows the results for the human features or regions measured: nose, cheeks and right hand. Each feature shows different trends and variability, although all of them are within the expected error of the camera.

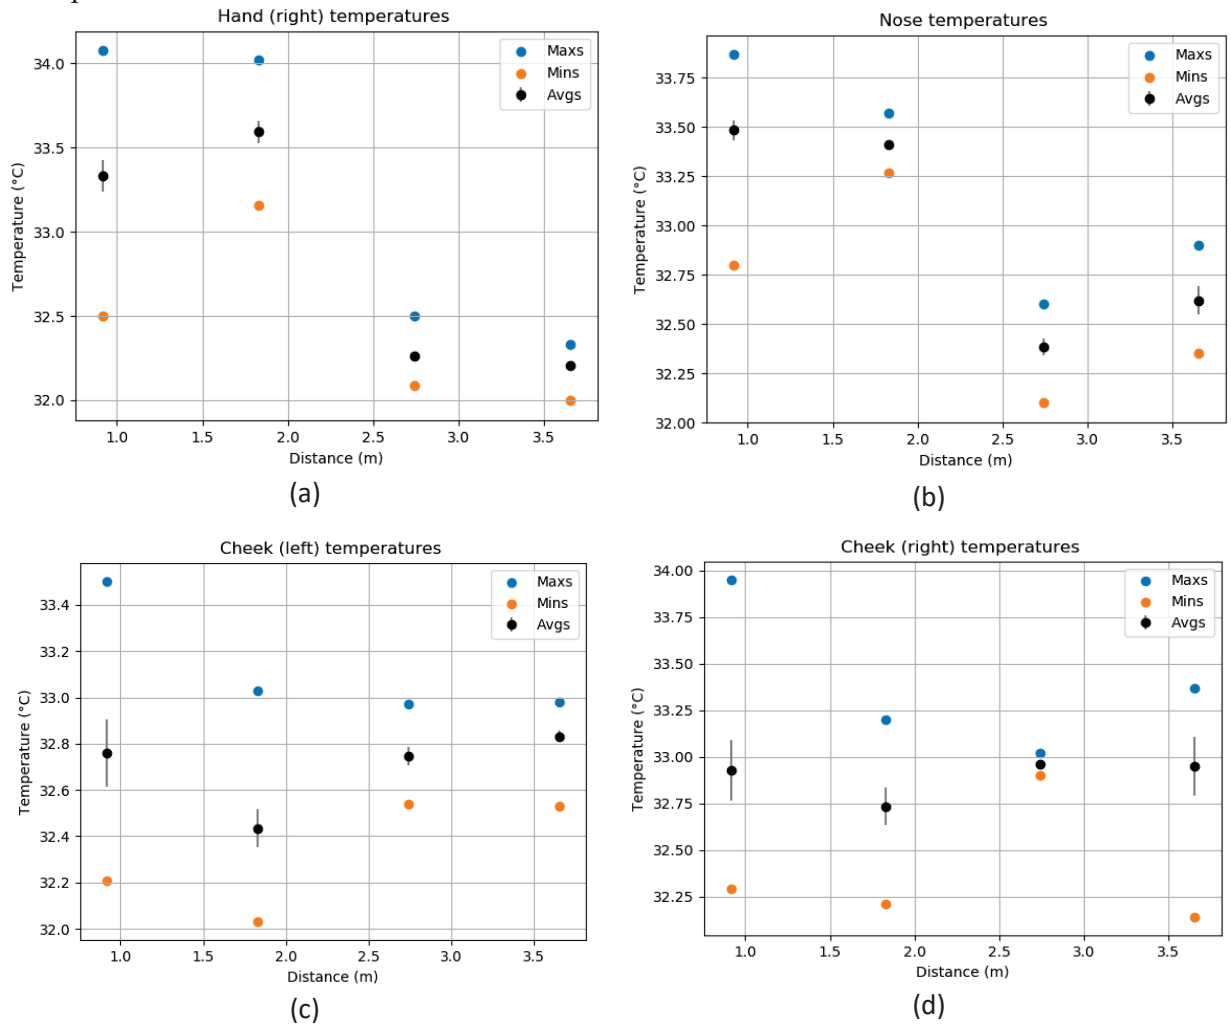


Figure 35. Maximum, average and minimum values of the following regions of the human subject as a function of distance to the camera center: (a) right hand, (b) nose, (c) and (d) both cheeks.

#### 7.4.1 Human test discussion

Human skin is a near ideal radiator in the thermal infrared spectrum (4–40  $\mu\text{m}$ ) with a thermal emittance between 0.95 and 0.98. With the camera calibrated to an emittance of 0.95, little variability was predicted; however, the human features studied did not behave as expected. With distance, as the blackbody reference test showed, we expected a smooth descending trend. However, hand and nose temperatures drop down only at approximately 2 m and cheeks give approximately stable quantities.

Together with the instrumental errors associated to the camera and the blackbody reference emitter, we find possible repercussions of unavoidable systematic errors, such as the background and subject temperature changes in the time it takes to move the setting. Measuring with only one camera, the chair and human subject had to be moved for each of the different thermal images, and it is possible that the subject got slightly warmer during that process. That would compensate the distance attenuation and thus explain the stability of the cheek measurements.

The human error in the selection of the non-trivial regions is also an important factor in these results. Establishing specific points on the hand, nose and cheeks is difficult, especially after a distance change which makes the area smaller, as illustrated in Figure 36. Different region choices might appear correct from the perspective of the human interpreter. However, even a few pixels of positional difference can give very different numerical results.

The fact that hand measurements have the highest variability is consistent with the fact that the hand is a thermally complex region, so precisely choosing of the points to capture a measurement becomes more important.

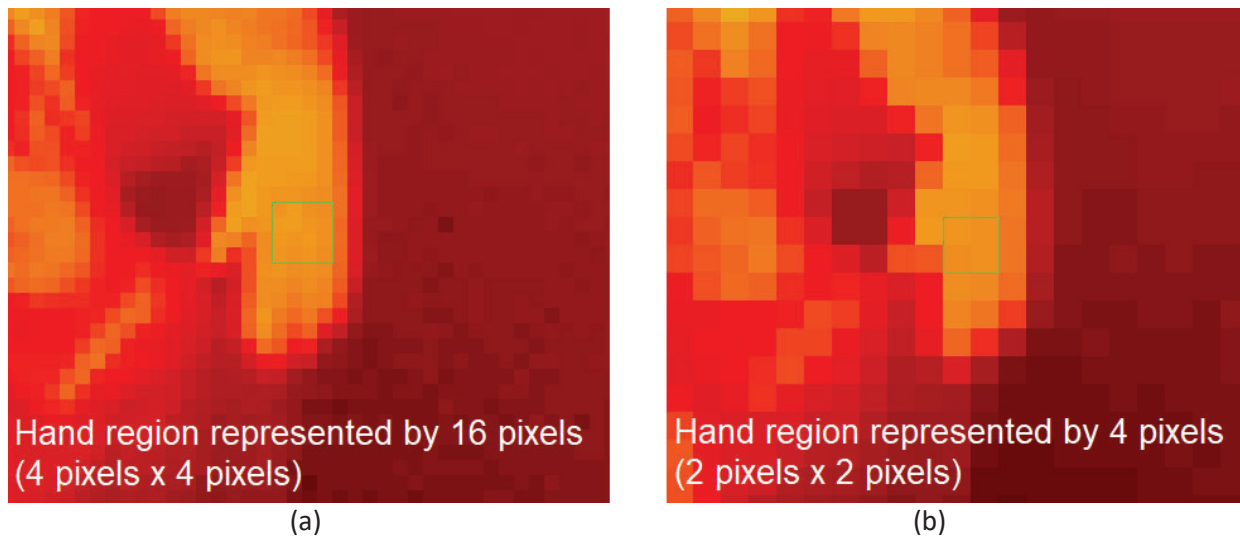


Figure 36. As camera distance increases, spatial resolution decreases. A smaller number of relevant pixels are available, and the placement of the region becomes increasingly more critical. In the figure above, the same hand region is shown in both (a) an image with higher resolution and (b) an image with lower resolution.

## 7.5 Inanimate object and surface test results

We can see in Figure 37 how this test has the highest standard deviations of all the tests performed, especially the aluminum plate, which shows a standard deviation of 0.8 °C. The blackbody emitter temperature displayed against distances from 1 to 4 m has a descending trend similar to the ones seen before.

The aluminum plate temperature values obtained with a contact sensor versus its corresponding thermal image measure are shown in Figure 38. For 1 m, we find the smallest difference in between these two measurements: 2.5 °C. That gap becomes higher as TIR temperatures go higher with distance.




Figure 37. Panel *a* shows maximum, average, and minimum values of the blackbody region as a function of distance to the camera center during Inanimate Objects and Surface Testing. Blackbody temperatures were used as a correction for the temperatures of the rest of features studied. Standard deviations are shown in panel *b*.

Finally, Figure 39 shows the aluminum plate and the rest of surfaces variations with distance. A very clear ascending trend is shown for the wall, table and chair measurements.


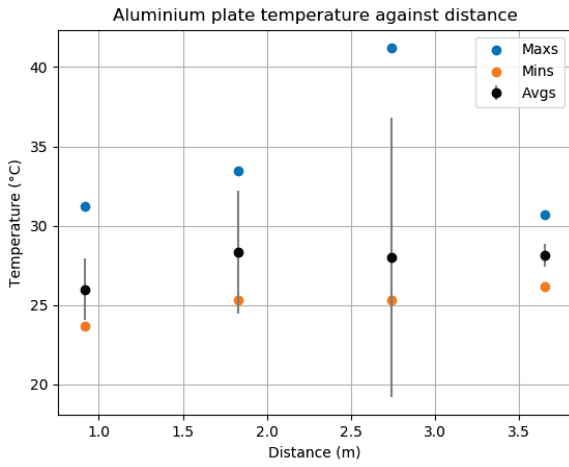



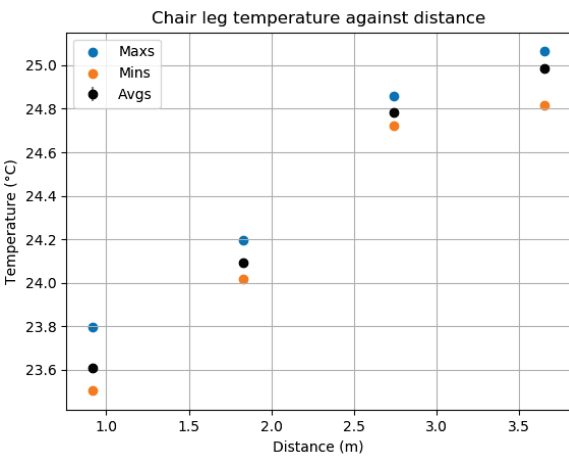
Figure 38. Comparison of average (with error bars) values of the temperature of the aluminum plate as a function of the distance measured with the TIR camera and with a contact sensor.



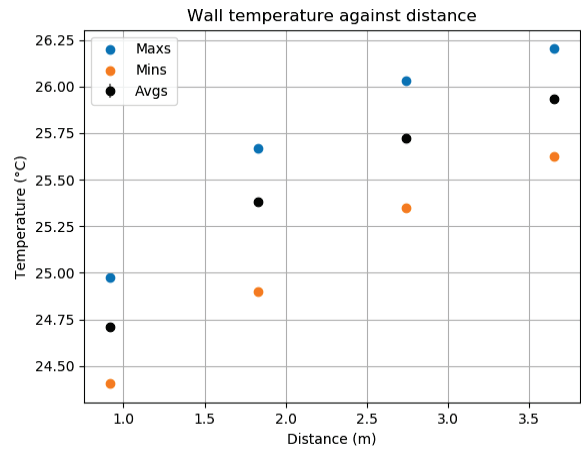
(a)



(b)



(c)



(d)

Figure 39. Maximum, average and minimum values of the following regions of the human subject as a function of distance to the camera center: (a) aluminum plate, (b) table, (c) chair leg, and (d) wall.

## 7.6 Inanimate object and surface test discussion

The objective of the inanimate object and surface testing was to capture temperatures from a typical, centrally heated and cooled office environment. To more accurately mirror a typical internal environment, window blinds and office doors were intentionally left open, facilitating wider ambient temperature variations than the previous tests allowed.

Whereas the blackbody and human testing discussed in the previous sections involved subjects with regulated internal temperatures and known thermal emittance, the inanimate object and surfaces tests targeted objects with unspecified thermal characteristics, and with no active

temperature control. As such, these objects could be impacted by heating and cooling of the environment more substantially than the previous subjects. In addition, the effects of thermal reflectance, while not specifically characterized for the objects under test, could add an additional vector of uncertainty.

While these choices introduced uncontrolled (and in some cases, difficult to explain) variables to this round of tests, we believe the results may be more indicative of real-world conditions that one may encounter when capturing ad hoc temperature measurements of inanimate objects in this type of environment.

Despite the fact that we attempted to eliminate camera drift by forcing a NUC operation before each measurement cycle, and further tried to optimize regions and timing, the uncontrolled room temperature variables appear to have much more impact on the data in this round of testing than the effects of distance in our previous tests. Based on the preliminary results collected, it appears that environmental considerations can overshadow the impact of distance, and in some cases, even the effect of angles in all but the most extreme cases (i.e. greater than 60° deflection from the orthogonal plane).

The variability in the uncontrolled parameters was intended to mirror real-life office conditions. In testing, it was determined that these factors are difficult to quantify and eliminate from the results. Therefore, only general observations can be made about the temperature accuracy of these objects in a real-world office environment.

Additional testing in a more tightly controlled lab environment may lead to more definitive conclusions with regard to inanimate object and surface temperature accuracy. However, given that the real-world target environments are not tightly controlled lab environments, additional testing in a controlled environment may be of dubious value unless the target environments can be tightly controlled.

While it is difficult to draw any definitive conclusions from this round of testing, it appears that uncontrolled environmental factors can be expected to impact TIR camera temperature accuracy on these objects and surfaces. Without a comprehensive understanding of the physical properties of the surfaces (specifically the thermal emittance, transmissivity, and thermal reflectance—in conjunction with other potential heat sources in the vicinity of an object), it is unlikely to expect the ad hoc thermal camera measurements to provide as reliable an environmental reference as other types of temperature measuring devices (i.e. thermostats, contact TC sensors, thermistors, etc.)

However, given that the proposed comfort model is primarily based on the measurement of human body temperatures, and the thermal emittance of skin is known to be almost ideal at 0.98%, it is believed that the results noted here should have minimal impact on the viability of the comfort model. If additional environmental parameters are necessary inputs to the comfort model, conventional sensors, or TIR camera sensors sampling large surface areas of known thermal emittance and reflectance (such as walls) may be useful measurement sources, provided windows

and active heat sources are shielded to prevent erroneous stray radiation from impacting measurements.

## 8 Summary

The following factors, if not properly accounted for in the experimental setup, all contribute in varying magnitude to error in thermal infrared measurements:

- Object emittance
- Distance between camera and object
- Viewing angle between the camera and object
- Sensor drift caused by environmentally induced temperature instability
- Operator-dependent methodology for calculating object temperature from the thermal image
- Ambient thermal reflections

Berkeley Lab and MoviTHERM each performed a suite of experiments with the objective of assessing the impact each of these factors have on overall measurement accuracy.

An initial set of controlled blackbody reference emitter tests were performed to examine in isolation the effect of camera-target distance and viewing angle on TIR measurement accuracy. In Berkeley Lab's experiments, a high-resolution camera viewed a 35.00 °C reference emitter at distances ranging from 1 m to 34 m; in MoviTHERM's experiments, a medium-resolution camera viewed a 55 °C reference emitter at distances ranging from 1 m to 4 m. Both experiments showed that, as expected, increasing distance induced an almost linear decrease in the TIR measurements of reference emitter temperature. In the Berkeley Lab trials, a 1 °C underreporting of temperature was observed when distance had increased to 10 m; the MoviTHERM trials reached the same threshold at a distance of 3 m. The main sources of this distance-based loss of accuracy are assumed to be the combined effect of decreasing atmospheric transmittance and effective loss of resolution at large distances.

Both Berkeley Lab's and MoviTHERM's camera viewing angle trials showed a decrease in TIR measurement of reference emitter temperature when camera-target viewing angle increased. This decrease was slight (i.e., within a few tenths of a degree Celsius) when viewing angles remained below 40°, only to deteriorate rapidly at larger angles. For example, at an extreme angle of 75 °C MoviTHERM's camera underreported the reference emitter temperature by approximately 11 °C and Berkeley Lab's camera underreported the reference emitter temperature by approximately 1 °C. We attribute these discrepancies to the directional nature of emittance as well as the effective loss of resolution at oblique viewing angles.

To explore the effects of sensor drift on TIR measurement accuracy, Berkeley Lab collected a suite of TIR image time series of its stable 35.00 °C reference emitter while the camera's NUC routine was both active and suppressed. The results showed that when environmental conditions favored drift in the sensor response (such as in the time period immediately after the camera is turned on)



an active NUC routine did not completely mitigate temporal instability in absolute TIR measurements. When the NUC routine was suppressed, TIR measurements of the stable reference emitter were observed to drift up to 0.8 °C within a 15 min period. When the NUC routine was active, drift was still observed at up to 0.5 °C over the same length of time. 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 most significant takeaway from this suite of experiments, however, was that regardless of how much the absolute TIR measurements were impacted by uncontrollable sensor drift, differential TIR measurements remained steadfastly stable and largely immune to drift.

In the last set of trials, Berkeley Lab and MoviTHERM each sought to explore TIR measurement accuracy in images more closely resembling the HCJR application. Berkeley Lab imaged human facial features at different viewing angles and attempted to validate the radiatively sensed forehead and nose temperatures against contact-sensed measurements of the same. MoviTHERM performed additional distance-varying trials using both a human subject and a collection of inanimate objects as the targets of interest.

It was determined during Berkeley Lab's human subject trials that contact sensor measurements of human skin are subject to many unavoidable sources of error and therefore should not serve as references for validating TIR measurements. While it is promising to observe that in four out of six images the absolute and differential TIR temperatures agreed to within 1 °C of at least one of the contact sensors, we cannot make any definitive conclusions about the agreement between TIR and contact-sensed measurements before repeating the experiment using a non-human thermal manikin as the subject.

The results of MoviTHERM's distance-varying trials with both a human subject and inanimate objects were both curious. Instead of a smoothly decreasing relationship between TIR measurement and distance, the human hand and nose showed a stepwise decrease in TIR measurement with increasing distance, while both cheeks showed no change at all. These results lent credence to our suspicion that viewing angle would have a less of an impact on human targets than on the flat surface of a reference emitter. On the other hand, the radiatively sensed temperatures of the tested inanimate objects actually appeared to *increase* with increasing distance. These results suggest that uncontrolled environmental factors outside of those previously identified may have a nontrivial impact on TIR measurement accuracy.

From these investigations we expect that parameters such as emittance, viewing angle, and thermal reflections should have a relatively small degrading effect on measurement accuracy due to the nature of our targets of interest. Human skin is of known high thermal emittance (0.98) and low thermal reflectance, and the forehead and nose are both features with curvature. Additionally, the camera-object distances and relative humidity levels expected in an indoor environment also maintain atmospheric transmittance levels at a relatively high floor. It is difficult, for example, to imagine use cases involving distances exceeding 10 m or relative humidity levels approaching 100%. Finally, while sensor drift does have a significant impact on the accuracy of absolute temperature measurements, it has no discernible effect on differential temperature measurements.

This shows that our measure of interest—differential temperatures between body parts—is not subject to this uncontrollable source of error. And in cases where accuracy in absolute temperatures might be necessary, we identified external referencing to be a reliable correction technique.

Ultimately, the results suggest that poor target resolution (which depends on the interplay between camera-target distance, camera-target viewing angle, choice of camera resolution, and temperature calculation method) may remain a nontrivial contributor to absolute TIR measurement error and possible contributor to differential TIR measurement error. The specific constraints this imposes on absolute and differential TIR measurement accuracy will be investigated in-depth in our next subtask, where we will perform experiments with a ceiling-height camera imaging a controlled-temperature thermal manikin at different distances and positions.

## References

- ASTM International, 2014. ASTM E1862-14: Standard Practice for Measuring and Compensating for Reflected Temperature Using Infrared Imaging Radiometers. West Conshohocken, PA. <https://doi.org/10.1520/E1862-14R18>
- FLIR Systems, Inc., 2015. The Ultimate Infrared Handbook for R&D Professionals. Retrieved from <https://www.flir.com/discover/rd-science/the-ultimate-infrared-handbook-for-rnd-professionals/>
- FLIR Systems, Inc., 2016. Technical Note: UAS Radiometric Temperature Measurements. Retrieved from <https://www.flir.com/discover/suas/uas-radiometric-temperature-measurements/>
- Holst, G. C., 2000. *Common Sense Approach to Thermal Imaging*. Bellingham, WA: SPIE Optical Engineering Press
- Macrae, B. A., Annaheim, S., Spengler, C. M., & Rossi, R. M., 2018. Skin temperature measurement using contact thermometry: A systematic review of setup variables and their effects on measured values. *Frontiers in Physiology*, 9. <https://doi.org/10.3389/fphys.2018.00029>
- Minkina, W., & Klecha, D., 2016. Atmospheric transmission coefficient modelling in the infrared for thermovision measurements. *Journal of Sensors and Sensor Systems*, 5(1), 17–23. <https://doi.org/10.5194/jsss-5-17-2016>
- Modest, M., 2013. *Radiative Heat Transfer, 3rd Edition*. Elsevier Academic Press. <https://doi.org/10.1016/C2010-0-65874-3>
- Neto, E. W., Costa, E. D., & Maia, M. A., 2006. Influence of Emissivity and Distance in High Voltage Equipments Thermal Imaging. In: 2006 IEEE/PES Transmission & Distribution Conference and Exposition: Latin America, 1-4. <https://doi.org/10.1109/TDCLA.2006.311553>

Nunak, T., Rakrueangdet, K., Nunak, N., & Suesut, T., 2015. Thermal Image Resolution on Angular Emissivity Measurements using Infrared Thermography. Lecture Notes in Engineering and Computer Science: Proceedings of the International MultiConference of Engineers and Computer Scientists, 1, 323-327. Retrieved from [http://www.iaeng.org/publication/IMECS2015/IMECS2015\\_pp323-327.pdf](http://www.iaeng.org/publication/IMECS2015/IMECS2015_pp323-327.pdf)

# Appendix B: Selection of wide-view TIR imaging temperature measurement technique(s)

## HCJR Task 1.2 / Milestone 1.2

Sharon Chen<sup>a</sup>, Dave Ritter<sup>b</sup>, Howdy Goudey<sup>a</sup>, Markus Tarin<sup>b</sup>, and Ronnen Levinson<sup>a</sup>

<sup>a</sup> Heat Island Group, Lawrence Berkeley National Laboratory, Berkeley, CA

<sup>b</sup> MoviTHERM, Irvine, CA

2020-05-25

## Abstract

The *Hot, Cold, or Just Right* application requires thermal infrared (TIR) images of most occupants within a room. Many low-cost TIR cameras have a field of view (FOV) less than or equal to 60°, which may be insufficient to image an entire room. In this report we identify and assess three strategies for wide-view TIR imaging and three options for camera placement. If cost and image resolution are a concern, we recommend installing a single camera on a rotating pan-and-tilt mount. Options for camera placement have different advantages and disadvantages and are suited to different use cases.

We supplemented our prior assessments of TIR temperature measurement accuracy with trials that imaged from near and far pre-defined areas (“targets”) of different sizes on a thermal manikin. We considered a TIR temperature measurement to be accurate when it agreed closely with a corresponding contact temperature measurement. In the short-range trial, we imaged the thermal manikin from various angles at distances less than 2.5 m. Here, we observed that the accuracy of TIR-determined target temperatures was poor when the targets were too small to register on a  $3 \times 3$  array pixel array, but were otherwise within the nominal accuracy specification ( $\pm 2$  °C). Within the subset of well-resolved measurements, we observed no clear relationship between accuracy and camera angle, distance, or reference correction. In the long-range trial, we imaged the manikin from 3.7 m and 5.8 m away, distances at which both targets were poorly resolved. Contrary to what we saw in the short-range trial, the accuracy of nearly all measurements lay within the nominal specification of the camera. The conflicting trends observed in these two trials suggest that we need to perform further trials to clarify the requirements for accurate TIR imaging. We will supplement this analysis with results from upcoming human subject studies.

## 1 Introduction

The HCJR application will require monitoring of human occupants who could be located almost anywhere inside a room. Therefore, the camera system that we develop must be able to accurately record the thermal infrared (TIR) temperatures of targets spanning a very wide field of view. Many affordable camera packages today such as the FLIR Lepton 3.5 have low resolution (e.g.,  $160 \times 120$  pixels) and modest fields-of view (e.g., below 60°). Given constraints like these, we investigated (a) strategies that can be employed to achieve

wide-view temperature characterization of an indoor space and (b) how their implementation affects TIR measurement accuracy.

## 2 Short-range contact-TIR measurement comparisons using a thermal manikin subject

We began by using trials with a manikin to elaborate our previously reported analysis comparing the accuracy of TIR temperature measurement to contact temperature measurements.

### 2.1 Experimental set-up

All thermal infrared (TIR) and color images generated in this quarter’s experiments were captured with either (a) the FLIR SC660 thermal camera, which captures both TIR and color images; or (b) the MoviTHERM-developed dual-camera assembly that pairs the FLIR A315 thermal camera and the FLIR Blackfly S BFS-PGE-50S5C color camera. Our TIR image analyses utilized both the free version of FLIR Tools (for FLIR SC660 TIR images) and a suite of custom programs developed by MoviTHERM (for FLIR A315 TIR images). A screenshot from one of the MoviTHERM-developed programs is shown in Figure 1. All contact temperature measurements were measured using an epoxy-bead thermistor. A discussion of the specifications of the equipment described here can be found in the Milestone 1.1 report [1].

Our team obtained a thermal manikin from Berkeley Lab’s FLEXLAB facility to serve as a target subject in our second round of contact versus TIR temperature comparison trials. This manikin is comprised of a clothed full-body female-shaped clothing display dummy (“doll”) with lengths of adhesive-backed heating wire (85 W) wound around its surface (Figure 2a). When viewed through a TIR camera, the temperature distribution across the thermal manikin is not uniform; areas corresponding to the manikin’s fabric “skin” are near ambient temperature, while areas including or immediately adjacent to the heating wire (or immediately adjacent) are significantly warmer.

To give the TIR camera two simple and pre-defined areas on which to focus, we taped two targets of known emittance to the “face” of the manikin (Figure 2b):

- a 5 cm × 10 cm target covering the forehead (hereafter described as the “forehead target”), and
- a 1.3 cm × 1.3 cm target covering the nose (hereafter described as the “nose target”).

We did not include a target for the hand, which was identified in Subtask 1.1 [1] as a promising candidate for use in our human comfort model, because its complex shape was difficult to identify and manually extract during our TIR image analyses.

These targets were composed of four layers—two layers of aluminum tape (high thermal conductivity) below two layers of masking tape (high thermal emittance<sup>1</sup>)—and were placed to contact both the heating tape strips and the manikin’s fabric skin. When viewed through a TIR camera, the targets present as rectangular areas of uniform temperature that are warmer than the manikin’s skin (about 21 °C) but cooler than the maximum temperature observed on the heating tape strips (about 45 °C) (Figure 2c). Initial

---

<sup>1</sup> The thermal emittance of the masking tape surface was measured to be 0.89 using a Devices & Services AE-1 portable emissometer following ASTM Standard C1371 [2].

assessments of these targets showed that they stabilized at slightly different temperatures; the nose target was consistently about 2 °C warmer than the forehead target. To assess the contact temperature of each target, a bead thermistor was placed in the center of each target, between the aluminum and masking tape layers. For the purposes of this experiment, we considered these contact measurements the true measure of manikin temperature against which the accuracy of the corresponding TIR measurements would be judged.

The intention of this experiment was to test the agreement between TIR- and contact-determined temperatures using a non-human subject with human-shaped and human-sized features. Unlike the reference emitter plate, the manikin is three-dimensional, with curved surfaces that are more visible (relative to a flat surface) across a wider range of camera angles. Unlike a human subject, the surface temperature of the inanimate manikin can be measured by contact sensors without the probe effect that occurs with attempting the same on living skin [3].

However, there are still many ways in which the manikin is an inadequate substitute for a live human subject, and these manikin tests are meant to supplement, rather than supplant, trials with real humans. The temperature signatures of the forehead and nose targets cannot be adjusted and are not meant to mimic the size, shape, or pattern of temperature distributions across a human face. Also, the manikin is bald, so this experiment does not address forehead visibility (or lack thereof) when the subject has hair, a head covering, or a hat.

In previous experiments we assessed how camera angle and distance independently affected TIR measurement accuracy when viewing a relatively large and flat uniform surface. In this experiment we aimed to observe whether the same trends hold when viewing a human-shaped and human-sized subject. In total, we collected a set of 14 color and TIR image pairs that viewed the face of the manikin at different angles and distances, which were subject to limitations imposed by the size of the room in which the experiment was performed. Panels *a*, *b*, and *c* in Figure 3 show, respectively, how we measured the camera-to-subject distance ( $D$ ; varying from 0.6 m to 2.4 m), horizontal camera angle ( $H$ ; varying from 0° to 90°), and vertical camera angle ( $V$ ; varying from 0° to 90°). Figure 4 displays the color and TIR image pairs from six different camera views captured at 0.6 m; Figure 5 displays the TIR images from six different camera views captured at 1.5 m; and Figure 6 displays the TIR images from two different camera views captured at 2.4 m. The distances (assessed by a measuring tape) and camera angles (adjusted manually via the 15°-interval angle markings on the tripod mount) reported here are approximate and subject to human error. These errors are most pronounced in the reported camera angles because the camera tripod's pivot point is offset a small distance behind and below the camera lens.

## 2.2 Results and discussion

The image compilation shown in Figure 4 shows the six different camera views assessed at a camera-target distance of 0.6 m. TIR images are shown the left column and corresponding color images on the right. The image compilation in Figure 5 shows TIR images captured from the six different views assessed at a camera-target distance of 1.5 m. The image set in Figure 6 shows the two TIR images from the two views assessed at a camera-target distance of 2.4 m.

Quantitative results of the 14-view survey are summarized in Table 1. All TIR measurement results reported in this table were extracted with MoviTHERM's HCJR-specialized "IR-CAT" software (Figure 1). In each TIR image, regions corresponding to each target were identified using visual inspection based on context

clues informed by the corresponding color image. TIR temperatures of each target were subsequently calculated as the average of all temperature pixels within the selected region ( $T_{\text{IR,avg}}$ ). For poorly resolved entries failing the industry “rule-of-thumb” regarding sampling area—which recommends that it encompass a pixel array no smaller than  $3 \times 3$ —the calculated TIR temperatures are highlighted in red and should not be considered valid results. Contact temperatures measured by the thermistors ( $T_{\text{TH}}$ ) corresponding to each camera view were determined by averaging the 30-second time series centered on the TIR image capture timestamp.

Table 1 reports experimental results in terms of TIR-contact temperature differences instead of absolute temperature results. The first three columns of Table 1 show the combination of camera distance ( $D$ ), vertical camera angle ( $V$ ), and horizontal camera angle ( $H$ ) that describe each distinct camera view. The reference correction (i.e., an offset correction based on the difference between the TIR measurement and known setpoint of the reference emitter) associated with each view is reported in the next column. The next four columns show both the uncorrected and reference-corrected TIR-contact temperature differences for both the forehead and nose targets. Finally, the last two columns report the nose-forehead temperature differences observed in the thermistor and TIR data sets, respectively.

The suite of camera views we tested is not comprehensive enough to support anything more than qualitative analyses. Broadly speaking, in all entries of well-resolved TIR temperatures—both raw and reference-corrected—we observed agreement with contact temperatures within  $\pm 2$  °C, which is the stated accuracy of the FLIR A315 camera. Beyond that, the trends in these results are challenging to understand. We had expected to see close agreement between TIR and contact results at near-ideal camera views—i.e., short distances and near-zero camera angles—and poorer agreement with increasing distance and angle. We expected that reference corrections would generally improve agreement between TIR and contact temperatures. We expected this temperature agreement to be closer in the forehead target results than in the nose target results due to size differences between the two targets.

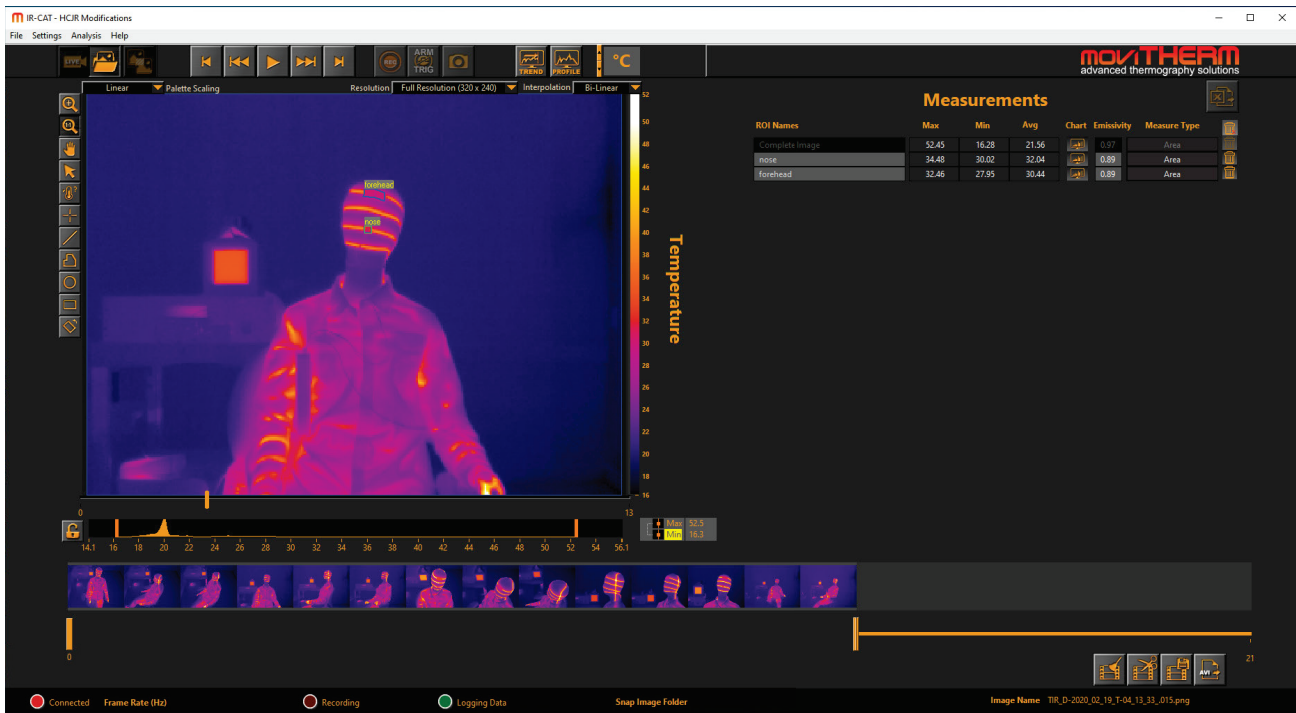
However, none of these expectations held true across all results. In general, we did observe that as expected, absolute TIR temperatures decreased as distance and angle increased. But the *absolute difference* between TIR and contact measurements did not actually significantly change with camera distance or angle. When camera views were near ideal, TIR temperatures were slightly higher than contact temperatures, and when camera views were less-than-ideal, TIR temperatures were slightly lower than contact temperatures. The reference correction did not significantly improve TIR-to-contact agreement, either; in roughly half of the measurements, the difference between the two increased when the correction was applied. Finally, we did observe that the size discrepancies between the two targets resulted in more invalid nose target TIR measurements (eight) than forehead target TIR measurements (one). However, among the camera views in which both targets yielded valid TIR measurements, the nose target TIR-to-contact agreement was generally closer than the forehead target TIR-to-contact agreement.

We do not currently have a framework for understanding the unexpected trends detailed above. However, we can still draw several useful takeaways from this experiment. First, we observed that the unobscured forehead target of the manikin was clearly visible in all views where the camera looked downward at a vertical angle of  $45^\circ$ . From this observation we can extrapolate that the same would hold true for the three-dimensional nose of a real human. This supports a strategy of mounting HCJR cameras on the ceiling.

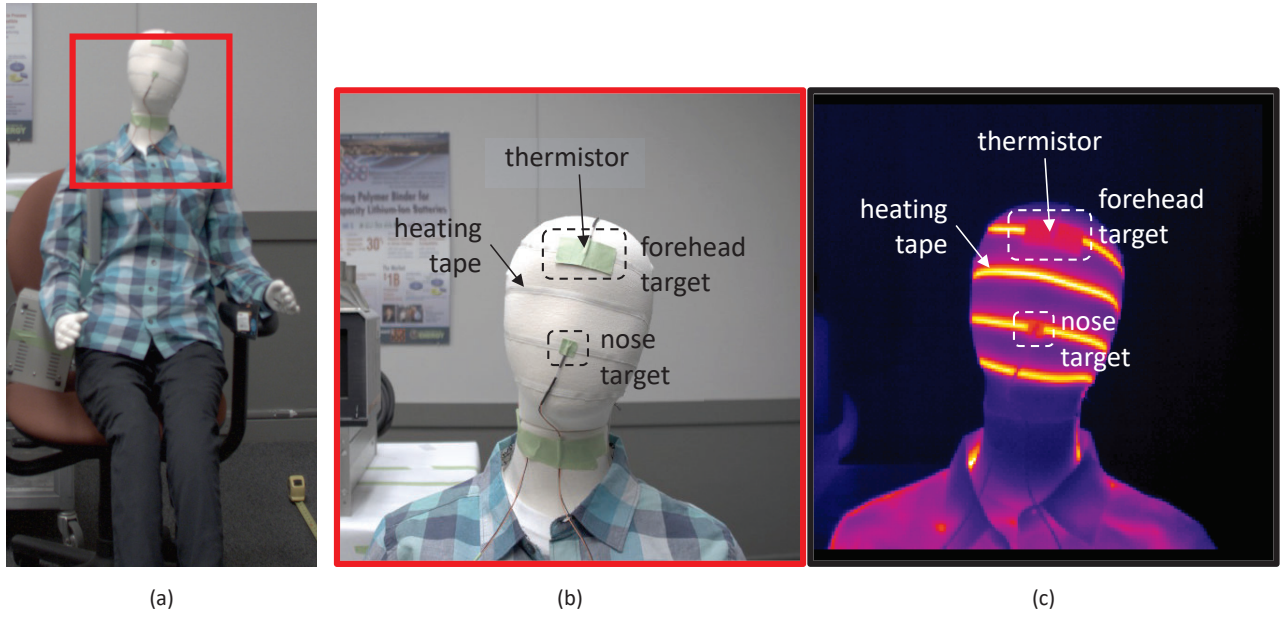


Second, we observed that when the camera distance was 2.4 m, no view of the nose target provided a sampling area that met the minimum required for accuracy per the industry rule-of-thumb (i.e., a  $3 \times 3$  array of image pixels); therefore, we consider these TIR measurements to have suspect accuracy. On the other hand, the forehead target—more than 30 times larger than the nose target—did fulfill the minimum sampling area requirement in both views at 2.4 m, so we consider these TIR measurements to be accurate. As shown in the last two rows of Table 1, the nose target-to-forehead target temperature differences calculated from the TIR images differ significantly from those calculated from the thermistor timeseries. In this situation where size differences result in both targets being unequally resolved by the camera—to such a significant degree—we see that it is the smaller of the two (the nose target) that determines the range of camera-to-manikin distances where accurate nose-to-forehead determinations are possible.

We are interested in knowing how vantage (i.e., distance and angle) affects the accuracy of TIR-calculated temperature differences because we believe that a likely model for human comfort involves comparing the warm and cold spots on a person's face and/or hands [1]. In this experiment, we saw that the large size discrepancy between the nose and forehead targets made accurate TIR-determined temperature differences difficult to measure accurately at a camera-to-manikin distance exceeding 1.5 m. That being said, it is unclear at this time whether the lessons learned here are transferrable to cases where the subject is a living human being. Our assessment of a person's thermal comfort will ultimately involve selecting the warmest or coldest spots warmest or coldest spots on the face or hand; without testing with human subjects, we don't have a good idea of how these features will manifest in TIR images. What are the relative sizes of these hot or cold spots? What shape do they take? These are questions we hope to answer in upcoming human subject trials.




1  
 2 Figure 1. Screen capture showing region selection and extraction of feature temperatures in a TIR image using MoviTHERM-developed software.  
 3

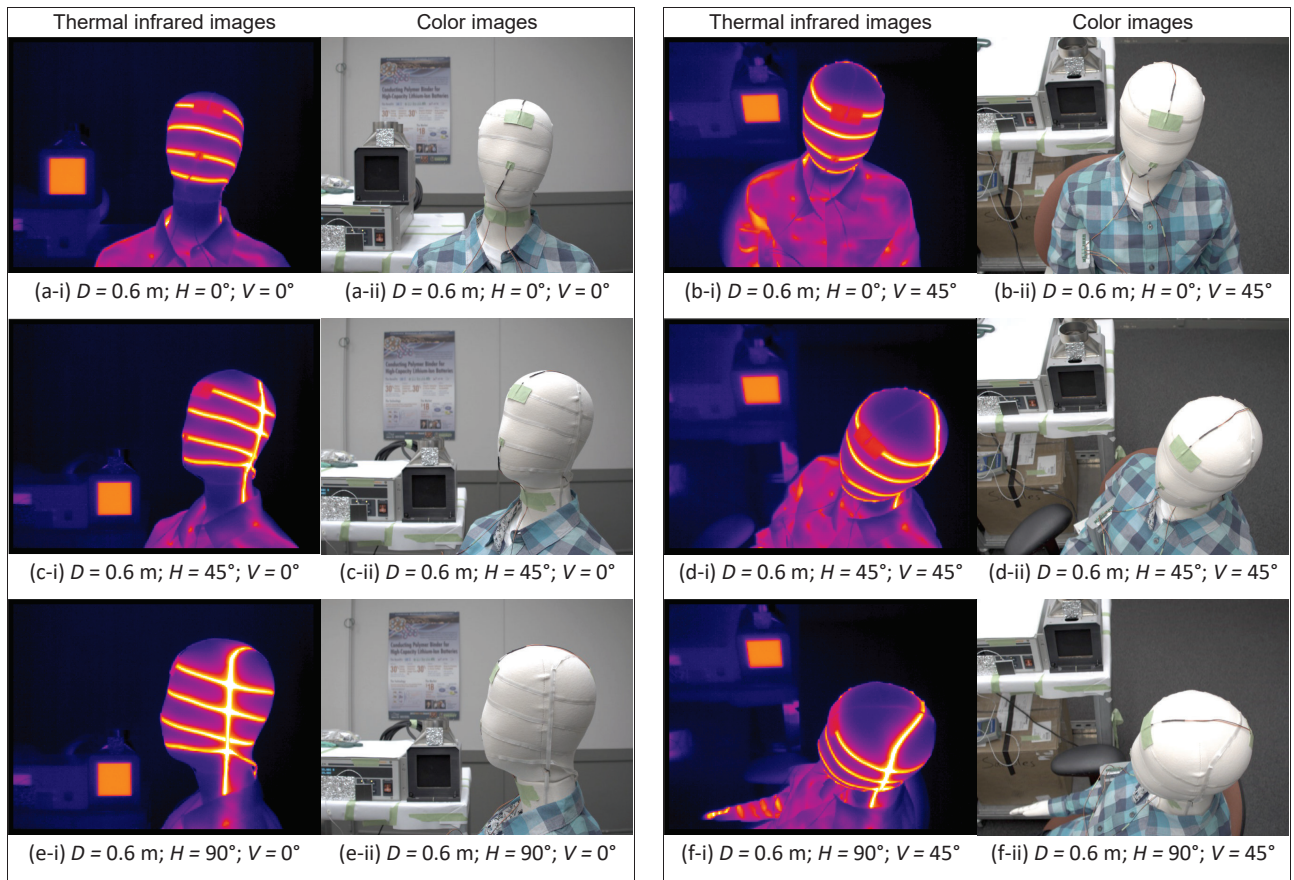


4 Figure 2. Images showing (a) the manikin, in color; (b) its facial targets, in color; and (c) its facial targets, in thermal infrared.

5



6 Figure 3. In this experiment we described each camera-to-subject viewpoint using three parameters: (a) the distance between camera and  
 7 subject,  $D$ ; (b) the angle from which the camera views the subject in the horizontal plane,  $H$ ; and (c) the angle from which the camera views the  
 8 subject in the vertical plane,  $V$ .  
 9



10 Figure 4. Six different camera-to-subject views (a – f) were assessed from a distance of 0.6 m. For each view, (i) thermal (left) and (ii) color (right)  
 11 images are shown side-by-side.

	$V = 0^\circ$	$V = 45^\circ$
$H = 0^\circ$		
$H = 45^\circ$		
$H = 90^\circ$		

12 Figure 5. Thermal images from the six different camera-to-subject views  
 13 assessed at a distance of 1.5 m.  
 14

15

	$V = 0^\circ$
$H = 0^\circ$	
$H = 45^\circ$	

16 Figure 6. Thermal images of the two different camera-to-  
 17 subject views assessed at a distance of 2.4 m.