Poster Abstract: Low-Cost, Perspective Invariant and Personalized Thermal Comfort Estimation

Peter Wei
pw2428@columbia.edu
Columbia University
New York, New York, USA

Yanchen Liu
yl4189@columbia.edu
Columbia University
New York, New York, USA

Xiaofan Jiang
jiang@ee.columbia.edu
Columbia University
New York, New York, USA

ABSTRACT

In this poster abstract, we present a thermal comfort estimation system using low-cost thermal camera based sensor nodes. This system extracts perspective invariant, non-intrusive thermal measurements, is easily deployable and low-cost, and can incorporate individual thermal feedback for more personalized thermal comfort estimates. In comparison with baseline methods, our system is able to improve thermal comfort estimates on the ASHRAE 7-point thermal sensation scale by up to 64% over baseline methods.

ACM Reference Format:

Peter Wei, Yanchen Liu, and Xiaofan Jiang. 2021. Poster Abstract: Low-Cost, Perspective Invariant and Personalized Thermal Comfort Estimation. In *The 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021) (IPSN '21), May 18–21, 2021, Nashville, TN, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3412382.3458781

1 INTRODUCTION

In commercial buildings, research has largely focused on energy optimization [1, 9]; however, occupant thermal comfort is also a critical factor that impacts both health and productivity [2]. In [3], authors found that only 2% of surveyed buildings meet the 80% occupant thermal satisfaction requirement defined by ASHRAE Standard 55 [8], suggesting that for most buildings, both measurement and optimization of thermal comfort can still be improved.

Recent studies [4, 6] have used thermal cameras to measure facial temperature statistics to estimate individual thermal comfort; however, these studies have left unaddressed challenges such as model personalization and robustness against occupant behavior. In this work, we develop a novel thermal comfort estimation system that is easily scalable, learns general as well as personal comfort models, and is robust to different facial perspectives.

2 SYSTEM DESIGN AND EVALUATION

2.1 Thermal and Environmental Sensing

To predict thermal comfort, we utilize thermal and environmental sensing features. The thermal sensing node is composed of a FLIR One Pro thermal camera (lower cost than most commercial thermal cameras, often 1,000$+), and an NVIDIA Jetson Nano board, to minimize cost. To reduce the overhead of deployment, a software library was developed for the Jetson Nano to continuously receive thermal and RGB images from the FLIR One Pro, encrypt and transmit the images securely to a cloud server. Future works can process images directly on the Jetson Nano, to preserve privacy.

We also collect ambient temperature and humidity features, which are noted to be important in various studies and are used in the ASHRAE 55 Standard for estimating thermal comfort. We deploy environmental sensors consisting of a TMP35 sensor and a Huzzah Feather board to measure and transmit ambient temperature and humidity time series data to the server every 15 seconds.

2.2 Comfort Feedback

In our deployment at Columbia University, we provide occupants with a web interface for submitting thermal comfort feedback on the ASHRAE 7-point thermal sensation scale. This feedback is submitted along with an ID for later association with identified thermal measurement data. To increase the amount of training data that can be collected with less ground truth data (which requires human interaction), each comfort feedback response is used to label data in a small time window around the feedback response.

2.3 Personalized Thermal Comfort

The thermal sensing node provides RGB and thermal images to our feature extraction pipeline (Figure 1), which is composed of five components: head detection, head orientation estimation, distance estimation, facial identification, and model mapping.

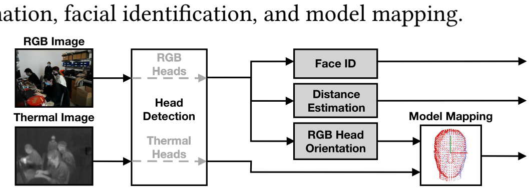

Figure 1: The thermal feature extraction pipeline outputs ID, distance estimate, and facial thermal features for each detected occupant in the RGB and thermal images.

2.3.1 Head Detection and Orientation Estimation. To detect occupants in the RGB and thermal images, we procured a custom training dataset of 100,000 labeled head bounding boxes from the South China University of Technology (SCUT) [5], and a hand labeled dataset of 1,000 RGB and thermal images from our deployment in a commercial building, which provides data similar to our anticipated conditions such as image quality and facial coverings. We retrained YOLO-v3 [7] on this dataset and can achieve over 98.0% and 91.0% precision and recall respectively. To estimate head

IPSN '21, May 18–21, 2021, Nashville, TN, USA

ACM ISBN 978-1-4503-8098-0/21/05…$15.00
https://doi.org/10.1145/3412382.3458781

orientation, we retrain a state-of-the-art head orientation model, FSA-Net [10] using a hand labeled dataset consisting of over 1,000 head images with pitch, roll, and yaw angles. We found improvement of 48.8% and 56.1% in average yaw and pitch error over the pre-trained model after retraining with the hand labeled dataset.

2.3.2 Distance Estimation. As temperature measurements from the thermal camera vary with distance, we utilize distance as a feature in estimating comfort. Our method to estimate the distance from the FLIR camera is based on a relationship between the area of the detected bounding box in the image (A), and the distance to the view plane containing the projected bounding box (D), as: AD^2 = constant. We fit a numerical constant based on multiple ground truth measurements. From distance, we can also derive the horizontal and vertical displacement (X_b and Y_b) in the view plane from the height and width of the view plane. Finally, the distance from the camera (R) can be estimated as $R^2 = X_b^2 + Y_b^2 + D^2$.

2.3.3 Face ID. To associate comfort estimates to specific occupants, our network uses a VGG-16 backbone which feeds into a convolutional layer and dense layers for classifying occupants. To train the model to differentiate between occupants, head bounding boxes are hand labeled using over 100 examples for each occupant.

2.3.4 Model Mapping. We perform mapping of thermal image measurements to a head model using the C++ Point Cloud Library. This involves selecting the points that are visible to the camera based on the orientation, flattening the points onto a plane, mapping the 2D thermal measurements to the flattened points, and reprojecting to recover the thermal measurement mapping. For comfort estimation, we extract median, maximum, and mean temperature features from the forehead, left and right temple, left and right cheek, and nose areas by computing these statistics over the temperature values of points on the head model that correspond to each region.

2.3.5 Personal Thermal Comfort Estimation. Once data has been collected, we train a general comfort model, using all of the labeled data from all occupants. For more personalized comfort estimates, personal comfort models are trained on labeled data for each occupant based on submitted comfort feedback. As thermal comfort preferences may vary between different people, it is important that each model is tuned to the specific preferences of the occupant.

2.3.6 Thermal Feature Extraction Pipeline Applications. The thermal feature extraction pipeline can also be used as a backbone in other applications such as fever screening, which can utilize temperature estimation to screen for febrile humans. Different architecture and prediction models are required, but the pipeline can serve as a critical component of future application specific systems.

2.4 Preliminary Results

We obtained approval from the Columbia University IRB (AAAS9589) to conduct a two week study in a commercial building on 10 occupants. We trained a random forest regressor (RF), and a linear regressor (LR) general comfort model; personal comfort models with random forest regressors, separated by occupant ID; and two baseline models: the ASHRAE 55 Standard [8], and Li et al. [4] which is a trained linear regression model on maximum and average thermal measurement statistics. Evaluation error on the second week data is shown in Figure 2. This study indicates that the ASHRAE 55 Standard, which doesn't use realtime measurements, has the highest error, while personalized comfort models achieve the lowest error, suggesting that occupant variations in thermal preferences are important and can be learned with personalized models.

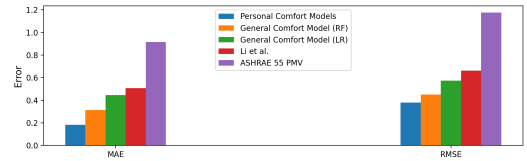

Figure 2: Comfort estimation error for the general and personal comfort models, Li et al., and ASHRAE 55 PMV.

3 CONCLUSION

Preliminary results show that personal comfort estimation has the potential to capture differences in thermal preferences and improve thermal comfort estimates over general models. We present a low-cost thermal comfort estimation system that is invariant to occupant perspectives, and demonstrate a 64% improvement in thermal comfort estimation error over baseline methods.

ACKNOWLEDGMENTS

This research was partially supported by the National Science Foundation under Grant Numbers CNS-1704899, CNS-1815274, CNS-11943396, and CNS-1837022. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Columbia University, NSF, or the U.S. Government or any of its agencies.

REFERENCES

[1] Yun Cheng, Kaifei Chen, Ben Zhang, Chieh-Jan Mike Liang, Xiaofan Jiang, and Feng Zhao. 2012. Accurate real-time occupant energy-footprinting in commercial buildings. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. 115–122.
[2] Ji Jia, Chengtian Xu, Shijia Pan, Stephen Xia, Peter Wei, Hae Young Noh, Pei Zhang, and Xiaofan Jiang. 2018. Conductive thread-based textile sensor for continuous perspiration level monitoring. *Sensors* 18, 11 (2018), 3775.
[3] Caroline Karmann, Stefano Schiavon, and Edward Arens. 2018. Percentage of commercial buildings showing at least 80% occupant satisfied with their thermal comfort. (2018).
[4] Da Li, Carol C Menassa, and Vineet R Kamat. 2019. Robust non-intrusive interpretation of occupant thermal comfort in built environments with low-cost networked thermal cameras. *Applied energy* 251 (2019), 113336.
[5] Dezhi Peng, Zikai Sun, Zirong Chen, Zirui Cai, Lele Xie, and Lianwen Jin. 2018. Detecting heads using feature refine net and cascaded multi-scale architecture. In *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2528–2533.
[6] Juhi Ranjan and James Scott. 2016. ThermalSense: determining dynamic thermal comfort preferences using thermographic imaging. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 1212–1222.
[7] Joseph Redmon and Ali Farhadi. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018).
[8] ASHRAE Standard. 2010. Standard 55-2010, Thermal environmental conditions for human occupancy. *American Society of Heating, Refrigerating and Air Conditioning Engineers* (2010).
[9] Peter Wei, Stephen Xia, Runfeng Chen, Jingyi Qian, Chong Li, and Xiaofan Jiang. 2020. A Deep-Reinforcement-Learning-Based Recommender System for Occupant-Driven Energy Optimization in Commercial Buildings. *IEEE Internet of Things Journal* 7, 7 (2020), 6402–6413.
[10] Tsun-Yi Yang, Yi-Ting Chen, Yen-Yu Lin, and Yung-Yu Chuang. 2019. Fsa-net: Learning fine-grained structure aggregation for head pose estimation from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1087–1096.