Human Factors and Wearable Technologies, Vol. 29, 2022, 63–72
https://doi.org/10.54941/ahfe1001475

Sensor Fusion for Remote Multi-Body Temperature Monitoring

Sean Hackett, Florian Alber, Haocheng Zheng, and Yang Cai

Carnegie Mellon University, 4720 Forbes Ave. Pittsburgh, PA 15213, USA

ABSTRACT

We present a prototype of the remote multi-body temperature screening system that uses an RGB-D camera for face tracking and distance measurement and a thermal imaging camera for temperature detection. An automated calibration algorithm is implemented based on the distance to the target, the ambient temperature, and a reference object. Our field tests include outdoor and in-vehicle body temperature monitoring. We found that face tracking and distance measurements help to improve the dynamic automatic remote body temperature monitoring. Wearing a face mask would impair face tracking, sensor fusion of thermal, RGB, and depth data, we have reduced the false detection of faces. In addition to the indoor environment, we evaluated the system in extreme conditions, including multiple types of face masks, outdoor, vehicle checkpoints, and under a canopy. We found that the vehicle drive-through and canopy offer improved performance over outdoors.

Keywords: Thermal, Thermography, Fever detection, Face detection, Face tracking, Face tracking, 3D vision, Human factors, Wearable, Sensor fusion, Sensors, Augmented reality

INTRODUCTION

Thermography imaging has been widely used as a screening device in checkpoints during past infectious disease epidemics, including SARS and COVID-19. A recent study in hospital labs shows that automated thermal image processing of the back yielded two risk scores that demonstrated higher sensitivity in the detection of COVID-19 (CDC, 2014; Brzezinski 2021). On the other hand, a flood of thermal scanners such as the temperature tablet kiosks (iOmniscient, 2022; Wilson, 2005; Healthline, 2022) that hit the market often failed to yield enough accuracy (Harwell, 2021). In this study, we aim to improve temperature measurement with sensor fusion of face tracking, distance measurement, and self-calibration. We would also like to explore extreme conditions such as wearing different kinds of masks, multiple bodies in motion, in-car drive-through, under a canopy, and outdoor.

THE RGB-D THERMAL IMAGING SYSTEM

The RGBD camera, thermal cameras, and microprocessor are contained in an enclosure with a tripod mount so they can be easily set up in any space. Figure 1 shows the system setup. The enclosure will point toward the forehead of the subjects and detect their IR reflections and distance. The system

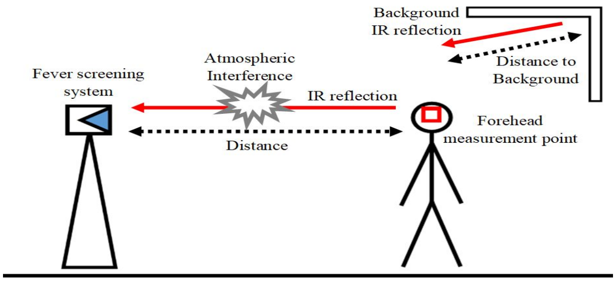

Figure 1: Fever screening system setup.

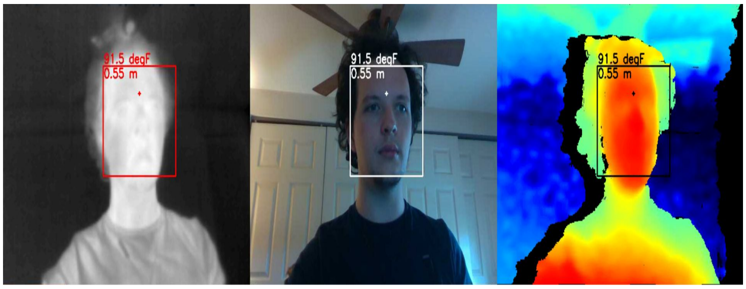

Figure 2: Distance versus temperature in thermography.

also has to compensate for the background temperature of the frame and atmospheric interference between the thermal camera and the subject.

The low-cost thermal camera FLIR Lepton 3.5 (FLIR, 2022) has a stated accuracy of +/- 9F which is obviously not suitable for body temperature tracking where a difference of just 1F could be critical. Measurements from the thermal camera were recorded for multiple water sources and compared with a digital thermometer. It was found that the thermal camera had an offset value that would change with each Flat Field Correction (FFC) which is an offset calibration to compensate for errors that build up over time in the camera. After correcting this offset value by manual or automated methods, the accuracy of the thermal chip is much improved, with an interquartile range of 0.5F, and a standard deviation of 0.44F.

As the ambient temperature of the room changes, this causes the output of the thermal camera changes, even if we are measuring a fixed point with a constant temperature. Also, the thermal camera performing an FCC will change the output temperature from the thermal camera. By measuring the temperature of a point in the background manually, in this case with a thermal measurement gun, we can then compare the temperature of the point from the gun and the thermal camera and generate an offset value to apply to the frame (Fhkam, 2005; Medline, 2022), as shown in Figure 2.

DISTANCE VS TEMPERATURE

Initial prototyping of the system was done at a lower resolution to optimize for speed. At 480p for the RGB frame, the face detection could operate at a

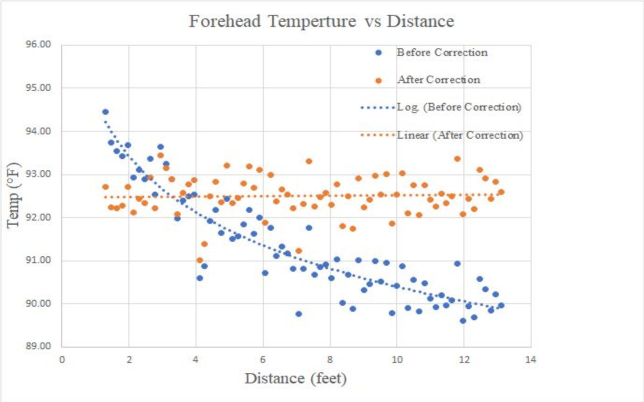

Figure 3: Temperature vs distance from the thermal camera.

maximum distance of 8 feet After optimizing the software the fever screening system could run at high speed at a higher resolution of 720p, extending the distance of face detection to 13 feet.

The distance between the subject and the thermal camera has an effect on measured temperature. This is due to the atmosphere which is between the thermal imaging camera and the object attenuating infrared radiation emitted by the object (Minkina and Klecha, 2016). Tests were carried out to measure the change in measured temperature of a participant as they walked towards the thermal camera, with their distance to the camera calculated by tracking their forehead in the depth image. The results of this are shown in Figure 3. For the distance of 1 to 13 feet, the measured temperature appears to be logarithmic, shown in blue. By applying the correction to the measured temperature, the output temperature more accurately described the actual temperature of the measured subject. The corrected temperature, shown in orange, is now flat. Variations around this corrected temperature could be due to the tracked point on the forehead changing slightly when walking towards the thermal camera and small inaccuracies of the camera.

DETECTING MASKED FACES

With the goal to use this system in public places, a challenge for face detection is detecting masked faces. It is more difficult to detect masked faces using traditional methods without a large training set. With the method of face detection, the quality of detected faces can be adjusted. With unobscured faces, this can be set high to reduce false positives. However, with a mask, the face will not be detected. By reducing the quality the face can be detected but introduces false positives as shown in Figure 4.

With sensor fusion, this problem can be overcome by combining the coordinates of the detected face in the RGB frame with the temp in the thermal frame. We know the typical temp range of the human face and can use this to reject false faces which have a temp outside of this. The maximal distance to the current face can be detected due to the resolution being approx 4m so any face detected further than this will also be rejected. This method has

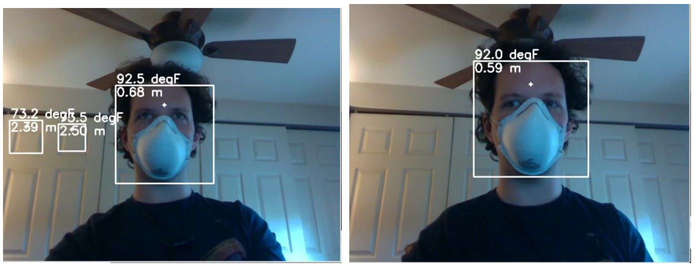

Figure 4: False Positives (left) and eliminated (right).

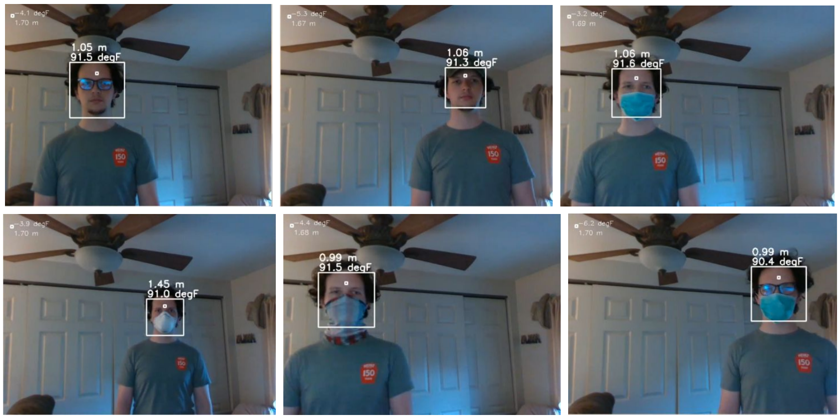

Figure 5: From top left to bottom right: glasses, hat, surgical mask, N95 mask, bandana, glasses + surgical mask.

reduced false positives to almost 0. Multiple masks have also been tested as shown in Figure 5.

Testing was carried out for these face obstructions to create a confusion matrix and determine their accuracy. A 30-second clip with 120 samples of each was recorded as the participant moved around the space and followed the same movement in each. The face was always in the frame so no true negatives occurred, and any false negative was an error, but occurrences of this were very low. See Table 1.

As expected the highest true positive accuracy was observed with no mask with over 90%, as the face is completely unobstructed. There is then a small reduction in accuracy with glasses. When wearing a baseball hat the temperature on the forehead was unable to be detected if the peak of the hat covered the forehead, or if so much of the hat covered the face so that the face could not be detected. If the hat only covers the forehead then future work could choose a different point on the face for temperature detection.

With each of the 3 types of masks (surgical, N95, and bandana) they performed according to how much of the face each covered, and how they distorted the shape of the face. The N95 mask changed the shape of the face and therefore showed the lowest accuracy. The surgical mask performed best as it covered the face the least and didn't distort the face shape, and the bandana performed between the two other masks, as although it covers a large

Table 1. Confusion matrices of face detection indoors (120 samples over 30 seconds).

No Mask

		Actual Results	
		P	N
Predicted Results	P	90.83%	4.17%
	N	9.17%	-

Glasses

		Actual Results	
		P	N
Predicted Results	P	85.83%	2.50%
	N	14.17%	-

Baseball Hat

		Actual Results	
		P	N
Predicted Results	P	65.83%	0.83%
	N	34.17%	-

Surgical Mask

		Actual Results	
		P	N
Predicted Results	P	81.67%	0.00%
	N	18.33%	-

N95 Mask

		Actual Results	
		P	N
Predicted Results	P	68.33%	1.67%
	N	31.67%	-

Bandana

		Actual Results	
		P	N
Predicted Results	P	73.33%	5.00%
	N	26.67%	-

Mask and Glasses

		Actual Results	
		P	N
Predicted Results	P	63.33%	0.83%
	N	36.67%	-

area, it doesn't distort the face shape as much as the N95 mask. When wearing glasses along with the surgical mask the accuracy was the lowest of all the face obstructions tested. Even with the lowest accuracy of true positives at 63% the face and temperature only need to be tracked in a small number of sequential frames to establish if a fever is present, so all the results provided a suitable result.

SIMULATED FEVER TEST

To simulate a face a hot towel was placed on the forehead to elevate the temperature above the normal range, as shown in Figure 6. A temperature in excess of 96.1F indicates a fever so the box and around the face switches to red so that an individual in a crowd can be easily identified. When using the hot towel to simulate the fever the forehead temperature decays very rapidly as shown in Figure 7. With a towel of 100F, the simulated fever above 96F is only maintained for approximately 5 seconds and returns to normal after approx 30 seconds.

OUTDOOR PERFORMANCE

The fever screening system will be more useful if it can operate in a variety of environments, especially if this can be done before a potentially sick

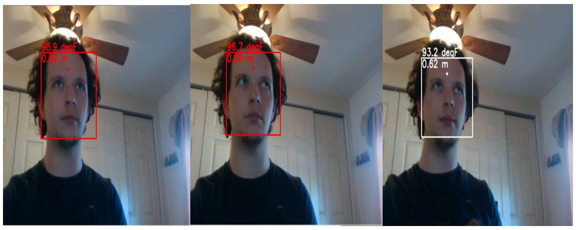

Figure 6: Simulated fever as it reduces (left to right).

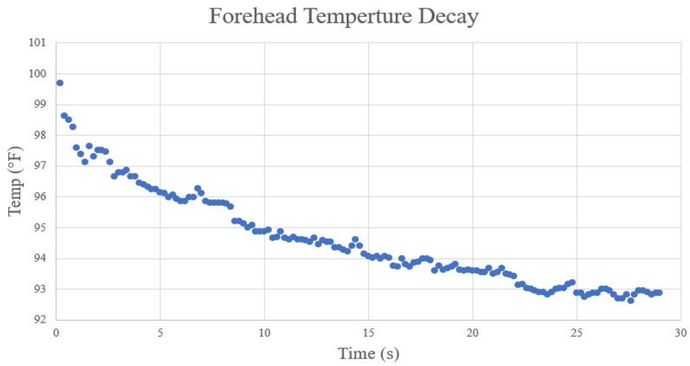

Figure 7: Forehead temperature decay for simulated fever.

person enters an indoor place. Here the system is tested outdoors, a car drives through and under a canopy with the temperature accuracy of each calculated. An outdoor environment offers more challenges than a relatively controlled indoor environment. First of all, inside there is a fairly constant ambient temperature of the background but outside there is a very large range, especially on a sunny day where dark surfaces heat up, and the sky which has a much lower temperature. This large range can reduce the temperature accuracy. Also trying to set a fixed temperature point to improve this accuracy is more difficult as the outdoor temperature can change quickly. Finally, lighting and shadows outside can make face detection more challenging.

Testing of the system needed to be carried out quickly to avoid changes in outdoor conditions, so the temperature point in the frame was measured and quickly set in the fever screening system. The face could then be accurately detected and false positives outside the human temperature range rejected as shown in Figure 12. After some time the fixed temperature point could change (e.g. due to the sun, clouds changing the temperature of its surface) and the output temperature of a face could be incorrect and the face rejected.

CAR DRIVE-THROUGH

A potential use case for a fever screening system could be at drive-through checkpoints such as a gated entrance to a facility or a border crossing point. This is desirable as it would prevent a person with a potential fever from

Figure 8: Outdoor Testing.

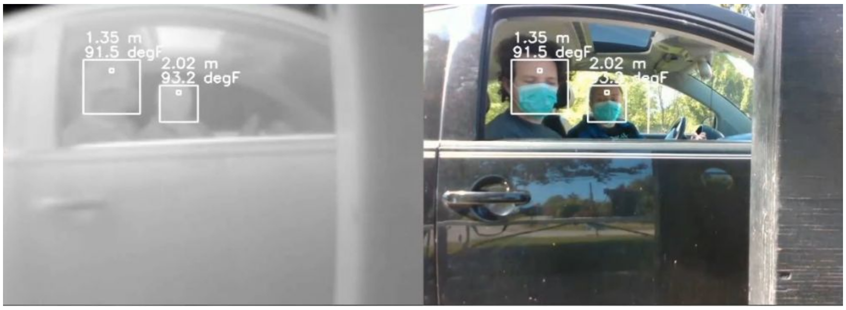

Figure 9: Car Drive-Through Testing.

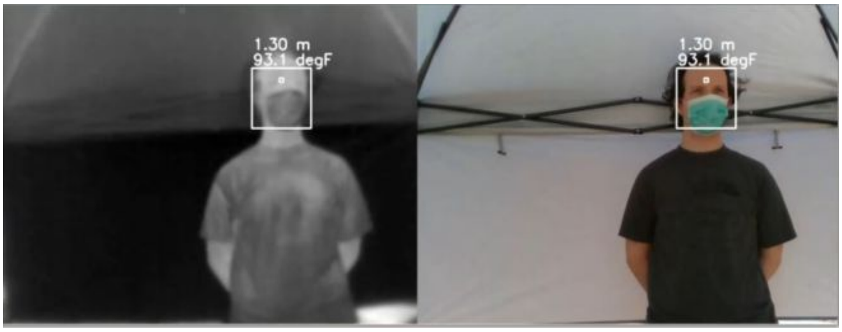

Figure 10: Canopy Testing.

interacting with anyone before they left their vehicle. For this reason, the fever screening system was set up outdoors next to a parked vehicle facing in through the window.

The system was tested with two participants in the vehicle wearing masks and the face and temperature of each were detected as shown in Figure 9. This was carried out during the day and external lighting would need to be provided to detect the face in lower light conditions. Since the testing was carried out in an outdoor environment the temperature measurement was still challenging, especially with the heat from the black paint on the car. A fixed temperature measurement point was placed near the camera and manually measured to improve the temperature accuracy. To improve this a sheltered environment could provide more control.

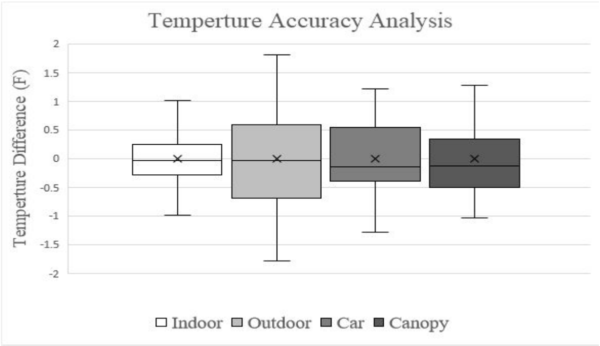

Figure 11: Temperature Accuracy Analysis.

THERMOGRAPHY UNDER A CANOPY

To overcome this challenge of the dynamic temperature outdoors the system was also tested under a canopy. Here the conditions are much more constant and the fixed temperature point needed to be updated much less often as it was covered from the sun and only ambient temperature would affect it. The face and temperature were detected more accurately than without the canopy or in a vehicle as shown in Fig. 10.

TEMPERATURE ANALYSIS

To determine the accuracy of the temperature measurement from the system analysis was carried out in the four different environments; indoors, outdoors, in the car, and under the canopy. A 30-second clip with 120 samples was carried out for each and the forehead temperature was recorded. Before the test, the forehead temperature was recorded with a handheld IR thermometer to act as the true value. The box plot in Figure 11 shows the results for each environment, showing the interquartile range and the maximum and minimum values. The recorded forehead temperature was corrected to the true value. The standard deviation for each is; indoors - 0.44F, outdoors - 0.88F, car - 0.65F, canopy - 0.68F.

The temperature accuracy mostly relies on setting a robust ambient temperature point. Indoors this can be set as a background and will remain fixed for a long duration. Outdoors will change much more rapidly due to temperatures changing quickly, shadows, etc. In the car and indoors offer a more stable environment but less so than indoors. Therefore indoor temperature accuracy performs the best, and outdoors performs the worst, as shown in Figure 11. The car and canopy offer improved performance over outdoors.

FACE DETECTION OUTDOORS

Our test was carried out for face obstruction in the outdoor scenarios to create a confusion matrix and determine their accuracy. A 30-second clip with 120 samples was recorded. The face was always in the frame so no true negatives occurred, and any false negative was an error, but occurrences of this were very low. The results for all are shown in Table 2, with all scenarios

Table 2. Confusion matrices of face detection outdoors (120 samples over 30 seconds).

Outdoors - No Mask

		Actual Results	
		P	N
Predicted Results	P	92.50%	1.25%
	N	7.50%	-

Outdoors - Mask

		Actual Results	
		P	N
Predicted Results	P	82.50%	5.00%
	N	17.50%	-

Car - Mask - Driver

		Actual Results	
		P	N
Predicted Results	P	100.00%	11.25%
	N	0	-

Car - Mask - Passenger

		Actual Results	
		P	N
Predicted Results	P	95.00%	1.25%
	N	5.00%	-

Canopy - No Mask

		Actual Results	
		P	N
Predicted Results	P	97.50%	1.67%
	N	2.50%	-

Canopy - Mask

		Actual Results	
		P	N
Predicted Results	P	86.25%	0.00%
	N	13.75%	-

proving to have good face detection results. In the outdoor test, the subject walked around the frame in front of the camera and followed the same path with, and without a mask. The outdoor results for face detection were very similar to indoors.

During the car drive-through test, the subjects were stationary in the car seat. So without any movement or disturbances, the results were very high as face detection is not challenging in this scenario. As in the outdoor test, during the canopy test, the subject walked around the frame but was closer to the camera due to the space of the canopy. For this reason, the canopy results are higher than outdoors.

CONCLUSION

In this study, we developed a sensor fusion method to combine face detection, distance measurement, and thermography to remote sense multiple body temperature automatically. We found that face tracking and distance measurements help to improve the dynamic automatic remote body temperature monitoring. Wearing a face mask would impair face tracking, with sensor fusion of thermal, RGB, and depth data, we have reduced the false detection of faces. In addition to the indoor environment, we evaluated the system in extreme conditions, including multiple types of face masks, outdoor, vehicle checkpoints, and under a canopy. We found that the monitoring accuracy for indoor can reach +/- 0.25F. The vehicle drive-through and canopy offer improved performance over outdoors. In the future, we would like to add additional features that are available through the RGBD camera such as distance tracking of people, skeleton tracking, or detecting people in low light environments e.g. bars, nightclubs. Also, the use of a known heat source in the frame could be used as a standard to prevent any offset in future use.

ACKNOWLEDGEMENT

The study is sponsored by NIST PSCR / PSIA Programs and UTC Mobility21 program. The author is grateful to Scott Ledgewood for his discussions and Brian Kokkila for his support on firefighter helmets and insightful comments.

REFERENCES

Brzezinski, R. et al. (2021) Automated processing of thermal imaging to detect COVID-19, Scientific Reports, Vol. 11, No. 17489, 2021: https://www.nature.com/articles/s41598-021-96900-9

CDC (2014). Non-Contact Temperature Measurement Devices: Considerations for Use in Port of Entry Screening Activity. 2014.

Fhkam, D.K.K. et al (2005). A brief report on the normal range of forehead temperature as determined by the noncontact, handheld, infrared thermometer, American Infection Control, Vol. 33, No. 4, 2005, pp. 227–229: https://www.sciencedirect.com/science/article/pii/S0196655305000052

FLIR (2022), UAS Radiometric Temperature Measurements: https://www.flir.com/discover/suas/uas-radiometric-temperature-measurements/

Harwell, D. (2021). Those fever scanners that everyone is using to fight COVID can be wildly inaccurate, researchers find. Washington Post, March 4, 2021

Healthline. (2022). As Many as 80 Percent of People with COVID-19 Aren't Aware They Have the Virus: https://www.healthline.com/health-news/50-percent-of-people-with-covid19-not-aware-have-virus

iOmniscient (2022), iQ-FeverCheck helps contain COVID-19: https://iomni.ai/wp-content/themes/clikthot/images/iQ%20FeverCheck%20White%20Paper%20March%2026%202020.pdf

MedlinePlus (2022), Body Temperature Norms, MedlinePlus: https://medlineplus.gov/ency/article/001982.htm#:~:text=The%20average%20normal%20body%20temperature,by%20an%20infection%20or%20illness.

Minkina, W. and Klecha, D. (2016). Atmospheric transmission coefficient modeling in the infrared for thermovision measurements, Journal of Sensors and Sensor Systems, Vol. 5, No. 1, 2016 https://www.researchgate.net/publication/303470908_Atmospheric_transmission_coefficient_modelling_in_the_infrared_for_thermovision_measurements

Wilson, M.E. et al. (2005). Airport Fever Screening: Seeking SARS, Finding Dengue: https://www.jwatch.org/id200503110000004/2005/03/11/airport-fever-screening-seeking-sars-finding