2025 Smart Systems Integration Conference and Exhibition (SSI) | 979-8-3315-1244-6/25/$31.00 ©2025 IEEE | DOI: 10.1109/SSI65953.2025.11107216

Bio-Remote Sensing in Real-Time Thermographic Face Detection and Respiratory Rate Measurement

Kianoush Rassels
BioMechanical Engineering
TU-Delft
Delft, The Netherlands
K.Rassel@tudelft.nl

David M.J. Tax
Pattern Recognition Laboratory
TU-Delft
Delft, The Netherlands
D.M.J.Tax@tudelft.nl

Paddy French
Bioelectronics Research Group
TU-Delft
Delft, The Netherlands
P.J.French@tudelft.nl

***Abstract*—Measuring respiratory rates for different age groups during** monitoring **and patient treatment at the hospital is extremely important. Monitoring respiratory rate for a long time provides physicians and nurses valuable information about the patient's health condition. Incorrect respiratory rate information of adults or infants can result in incorrect diagnosing and treatment of the patient. The traditional respiratory rate measurement and monitoring is contact based. However, these are quite obtrusive since the patient needs to be connected to the monitoring apparatus with wires. These methods could cause damage to vulnerable skin like preterm infants and create stress or pain.**

This paper introduces a novel thermographic Bio-Remote sensing approach that enables real-time face detection and respiratory rate measurement of subjects using a single thermal camera system. The algorithm achieves this without requiring nostril location, instead utilising thermal images and minimum temperature profiles for accurate detections and measurement. Furthermore, this paper discusses the significance of combining Deep Learning (DL) with the Thermal Imaging technique to provide a safer, faster, and more practical solution for hospitals by accurately measuring the respiratory rate compared to a monitoring device as the golden standard.

Keywords—Thermographic face detection, respiratory rate measurement, Bio-Remote Sensing, Infrared thermal imaging, thermography technology, Deep Learning.

I. INTRODUCTION

The accurate monitoring of respiratory rate is essential for assessing patient health and diagnosing various conditions. Traditional methods involving physical probes with wires, respiration belts, or devices can be invasive and discomforting, particularly for vulnerable populations such as preterm infants [1]. These techniques may also introduce errors due to skin sensitivity or physical stress.

A contactless digital image sensor-based Remote-Photoplethysmography (RPPG) offers a non-invasive approach for respiratory rate measurement. However, its reliance on ambient light poses significant technical challenges in environments with varying lighting conditions, such as low or dark settings that may introduce noise and affect measurement accuracy.

However, thermographic technology presents a promising solution, providing a safe and harmless method for respiratory rate measurement without dependency on environmental light conditions. This technique allows for contactless monitoring, enabling physicians and nurses to focus on patient treatment and providing care rather than managing devices and handling wires and connecting to the patient.

The application of thermographic technology is particularly significant in clinical settings where remote monitoring is essential, such as with COVID-19 patients [2]. Integrating real-time thermographic detection with tracking systems offers a robust, non-invasive method for measuring vital signs and monitoring physiological changes across diverse age groups.

This study introduces a novel, reliable, and accurate single-camera thermographic Bio-Remote Sensing approach that enables real-time effective face detection and respiratory rate measurement. By integrating DL techniques with Thermal Imaging, the method presents a fast, practical, and safe solution comparable to established devices like the Philips VM6 patient monitor, which serves as a golden standard for measuring vital signs including heart rate, respiratory rate, and blood pressure.

II. MEASUREMENT SETUP

The measurement setup consists of a FLIR thermal camera; an Ubuntu 22.04; a Philips portable patient monitoring, recording, and alarming apparatus; and dedicated software running on the laptop for real-time pre- and post-processing of the thermal images.

A. Used Hardware and Software

In this study, an FLIR A700 infrared camera was connected to a laptop using an RJ45 Cat.5 unshielded ethernet cable containing an adapter/power supply.

The laptop specs are as follows: Ryzen 9 5900HX 8-core/16-thread, NVIDIA GeForce RTX 3050 Ti 4GB DDR6 GPU, 32 GB 3200 MHz DDR4, with Ubuntu 22.04 Windows Subsystem for Linux (WSL). The thermal measurements can be done at 7 frames per second (fps) or 9 fps with a resolution of 640x480 pixels.

The Philips SureSigns VM6 portable patient monitoring apparatus has an accuracy of ±1 rpm in the range of 0 to 120 rpm and ±2 rpm in the range of >120 rpm [3], was connected to the subject to validate the respiratory rate measured by the FLIR A700 thermal camera, see Fig. 1.

Fig. 1. Philips SureSigns VM6 patient monitoring apparatus (left picture*)*, FLIR A700 infrared camera (right picture).

B. Experimental Setup

The experimental setup involves detecting the face, measuring the temperature, and extracting the respiratory rate in a controlled indoor environment, and can be modelled as demonstrated in Fig. 2. Our proprietary software running on the laptop will measure the temperature with an accuracy of ±0.01 °C by compensating for the radiation sources and thermal noises that can affect the temperature measurements [4].

For accurate facial temperature measurement, the thermal camera must be positioned in front of the human subject at an appropriate distance. The lens of the infrared camera should be aligned with the subject's face such that the Field of View (FoV) is maintained within a 55° angle, as illustrated in Fig. 2. This setup ensures consistent directional emittance and reflectance during temperature measurement [5], which are essential for accurate readings. Additionally, this positioning guarantees that the infrared camera's FOV path remains unobstructed, enabling precise thermal data capture. This approach not only enhances the accuracy of temperature data but also guarantees the constant directional emittance and reflectance during the temperature measurement [5]. By adhering to this positioning standard, the subject's face can be detected based on facial temperature analysis and subsequently the respiratory rate can be measured.

In case, there is a deviation from the required measurement conditions, it leads to a decrease in measured temperature accuracy. This reduction in accuracy subsequently results in inaccuracies during respiratory rate measurement. To address these deviations, the system automatically provides visual feedback on the display about the reduced confidence level of the measurements. Additionally, an alarm is triggered and sent to the caregiver, indicating potential inaccuracies.

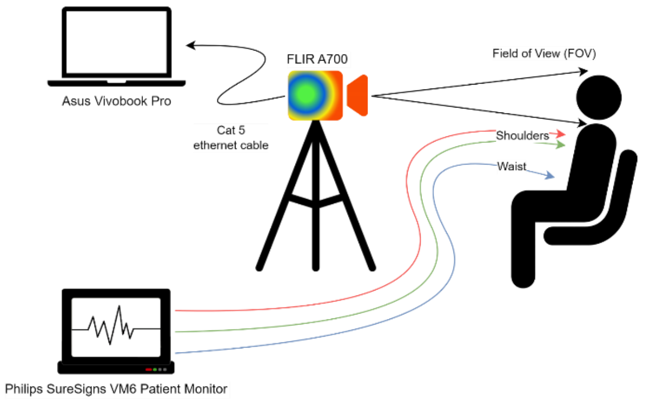

Fig. 2. The laptop (upper-left picture) receive real-time thermal images for analysis, the FOV of the FLIR A700 infrared camera (middle picture) is toward the human face (right picture) detects the face and measure the facial temperature profile, and The Philips SureSigns VM6 patient monitoring apparatus (lower-left picture) conventionally measures respiratory rate via attached wires to the human subject as an golden-standard in order to compare it with the extracted thermal respirator rate.

III. Face Detection and Respiratory Rate Extraction

This section provides a comprehensive understanding of the filter chain, enabling real-time face detection and respiratory rate extraction using only a single thermal camera without requiring nostril localisation [2][6][7], as demonstrated in Fig. 3.

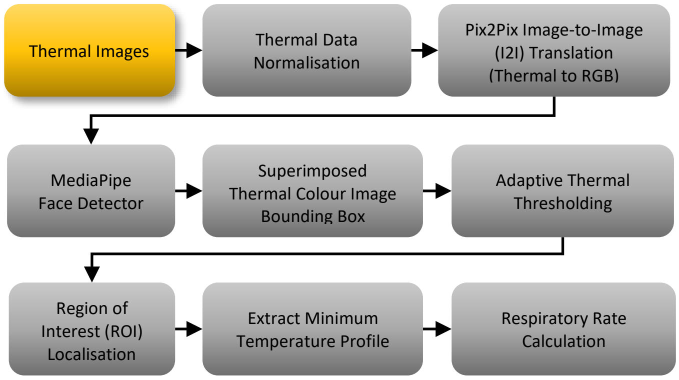

Fig. 3. Thermographic filter chain in detecting the face and extracting the respiratory rate.

A. DL Model and Thermal Face Detection

As depicted in Fig. 4, FLIR thermal raw data is fed through a series of steps to generate an RGB image suitable for face detection using a single thermal camera. By normalisation of raw thermal data, this step adjusts intensity values within a standard range, enhancing consistency in subsequent processing.

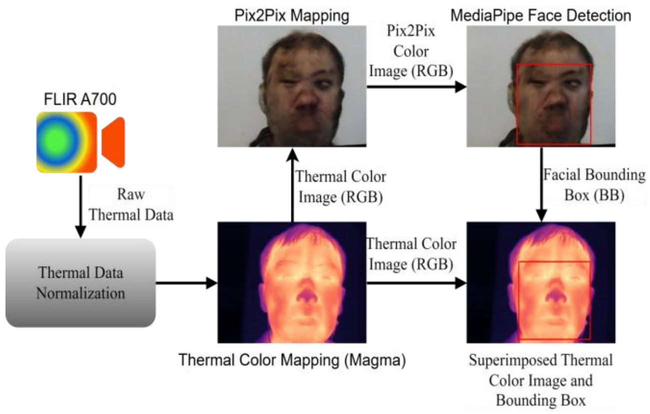

Fig. 4. Subject's face in the FoV of the thermal camera for continuous temperature measurement for real-time thermal face detection.

Next, the Magma colour map is applied to normalised thermal data, transforming it into a visible RGB image. The Magma colour map is chosen for its ability to balance warm and cool colours, effectively highlighting both subtle temperature differences and spatial information, which in turn helps in accurate facial feature detection. This RGB thermal image is then used as input into the Pix2Pix model [8][9], which performs image-to-image translation. The model generates an RGB image that closely mirrors the original thermal data, maintaining high fidelity. Finally, the output from Pix2Pix is used by MediaPipe Face Detection [10] model developed by Google Research in real-time to locate facial bounding boxes (BBs). These BBs are superimposed onto the thermal colour image, providing precise localisation of the subject's face.

This overlay facilitates accurate measurements and analysis, enabling robust facial detection and respiratory rate extraction without reliance on specific anatomical landmarks like nostrils.

B. Detecting Thermal Face with DL Training Model

Detecting faces in thermal images presents unique challenges due to the nature of thermal imaging, which captures heat signatures rather than visible features.

Conventional face detection algorithms are typically designed for RGB images, limiting their utility in thermal scenarios without additional hardware. In this approach the Pix2Pix conditional generative adversarial network (cGAN) converts thermal images into RGB format, enabling efficient face detection with existing methods.

The primary objective is to achieve face detection without reliance on additional RGB camera hardware, thereby optimising hardware requirements and operational efficiency. The Pix2Pix model is trained using resources from the Delft AI Cluster (DAIC) [11], a High-Performance Computing (HPC) cluster at TU-Delft. DAIC comprises Linux servers equipped with substantial processing power and memory, capable of managing large, computationally intensive tasks, including GPU-enabled jobs. This advanced infrastructure supports the training process, enabling the model to learn effective mappings from the thermal domain to the visual domain.

The Annotated Thermal Faces in the Wild (TFW) dataset [12] serves as the primary training data for this study. TFW provides pairs of images captured using both a thermal camera and a regular camera, encompassing indoor and outdoor settings. However, our analysis is limited to 7200 indoor images.

C. *Adaptive Thermal Thresholding and Localisation of the Region of Interest (ROI)*

In this study, after detecting and tracking a face in real-time using thermal imaging technology combined with DL, the next step will focus on localising specific facial features, particularly the nostrils.

This process is performed in both time and spatial domains to enhance computational efficiency and reduce the complexity of data processing. The detection and localisation of facial features are achieved through an *Adaptive Thermal Thresholding* method, see lower figure, (Fig. 5).

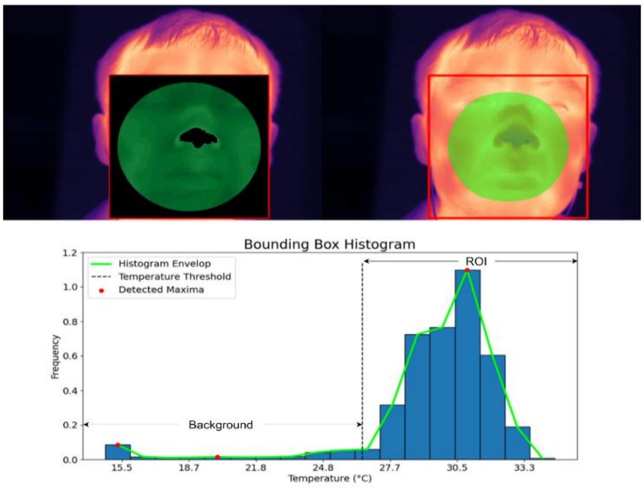

Fig. 5. Adaptive Thermal Thresholding per frame to isolate and exclude subject's face centre (ROI) from the background and face contours, respectively.

This approach utilises the histogram of BBs to dynamically adjust thresholds per frame, thereby minimising background noise and improving accuracy by eliminating the influence of the background and face contour effects (temperature gradients).

By selecting a median temperature value within each BB, we effectively eliminate background temperatures and isolate the foreground regions corresponding to the facial area. The algorithm identifies the maximum peak in the temperature profile histogram of the BB, which represents the most common temperature within that region, and applies a 2σ range to the left of this peak to further isolate the face from surrounding areas with different thermal characteristics, see lower figure, Fig. 5.

To achieve accurate temperature measurements in thermal face analysis, a circular mask is superimposed over the detected face within its BB. This mask defines the ROI, including colder regions around the nose as shown in the upper-left figure (Fig. 5). To ensure a complete and unbroken mask circle without any gaps, an adaptive morphological operation is used to close any holes within the mask, as depicted in the upper-right figure (Fig. 5). Finally, morphological erosion refines the mask's edge to focus on the central facial area, enhancing both accuracy and reliability for subsequent thermal analysis per frame by minimising lower background temperature interference.

D. *Extracting Minimum Temperature Profile and respiratory Rate*

By analysing the temperature profile characteristics and its variations from the thermographic signal, we can measure the temperature changes associated with breathing. The key step involves identifying the global minima in the temperature variation, using a cubic-spline interpolation to denoise the data and ensure a reliable detection of these minima, see Fig. 6.

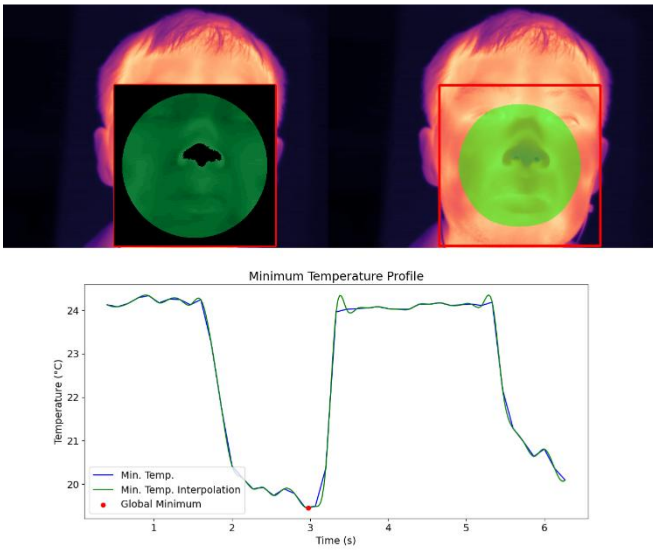

Fig. 6. Reparatory rate extraction by focusing on the center of the face (ROI) to detect minimum temperature profile.

To achieve this, a 5-second FIFO buffer for initialisation and to identify a clear respiratory cycle is used. The buffer is filled with the measured minimum temperature in ROI until at least one complete period of the respiratory rate is detected. In cases where the breath cycle or global minimum are not detected in the buffer, a second buffer is employed until a reliable detection is made. This process is repeated until two consecutive global minima are detected, allowing us to measure the time difference between these minima, which represents one complete breathing cycle. This time difference is then converted into measured respiratory rate per minute and updated every second for real-time display. In case the subject holds his/her breath, the system updates the displayed breath rate by reducing the respiratory rate based on the last detected cycle or global minimum.

This approach ensures accurate and reliable respiratory rate measurement in real-time while maintaining robustness against periods of irregular breathing or apnea.

IV. Conclusions

In this study, we present a novel, real-time, and accurate Bio-Remote Sensing method for measuring the respiratory rate using a single thermographic camera. To enhance the face detection process, we eliminated the need for an RGB-camera, thereby reducing complexity and dependency on environmental light conditions and noise. Additionally, this method offers a solution for thermal face detection that enhances operational efficiency while minimising dependencies through DL techniques.

Our approach is distinguished by its safety, precision, real-time functionality, and ability to operate without requiring nostril localisation. It utilises a hybrid model combining model-based techniques and AI-based DL with thermal imaging to detect facial features and measure respiratory rates, outperforming current techniques and moreover it is compared to traditional golden-standard methods. The system operates efficiently, providing updates every second, making it suitable for real-world applications in clinical environments. This research not only reduces hardware requirements and integrates seamlessly with conventional RGB-based detection algorithms but also addresses challenges such as varying lighting conditions and noise artifacts inherent in thermal imaging.

Acknowledgements

The authors would like to thank Caliang Xie for supporting this project. We are also grateful to Tom Dries, founder of Thermal Focus BV in Belgium, and TELEDYNE Technologies company (FLIR), which supported this project by providing state-of-the-art thermal camera hardware.

We also express our gratitude for the full support and help of the Delft AI Cluster (DAIC) at TU-Delft, including the provision of the necessary computational resources for training the Pix2Pix model.

In addition to the dedicated efforts of our research team and the crucial contributions from various individuals and entities. This work has been approved by the TU-Delft Human Research Ethics Committee (HREC) in the Netherlands and overseen by the Dutch Ministry of Health.

References

[1] C. B. Pereira, X. Yu, T. Goos, I. Reiss, T. Orlikowsky, K. Heimann, B. Venema, V. Blazek, S. Leonhardt, and D. Teichmann, "Noncontact monitoring of respiratory rate in newborn infants using thermal imaging," *IEEE transactions on Biomedical Engineering*, vol. 66, no. 4, pp. 1105–1114, 2018.

[2] Z. Jiang, M. Hu, L. Fan, Y. Pan, W. Tang, G. Zhai, and Y. Lu, "Combining visible light and infrared imaging for efficient detection of respiratory infections such as covid-19 on portable device," *arXiv preprint arXiv:2004.06912*, 2020.

[3] Philips suresigns vm6. Accessed: 2025-03-07. [Online]. Available: https://www.philips.com.au/healthcare/product/HC863065/suresigns-vm6-portable-bedside-monitor

[4] K. Rassels and P. French, "Accurate body temperature measurement of a neonate using thermography technology," in *2021 Smart Systems Integration (SSI)*. IEEE, 2021, pp. 1–5.

[5] A. F. Mills, *Basic Heat and Mass Transfer*, 2nd ed. Pearson New International Edition: Pearson Education Limited, 2014.

[6] L. Chen, N. Liu, M. Hu, and G. Zhai, "Rgb-thermal imaging system collaborated with marker tracking for remote breathing rate measurement," in *2019 IEEE Visual Communications and Image Processing (VCIP)*. IEEE, 2019, pp. 1–4.

[7] J. Fei and I. Pavlidis, "Thermistor at a distance: unobtrusive measurement of breathing," *IEEE transactions on biomedical engineering*, vol. 57, no. 4, pp. 988–998, 2009.

[8] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," *CVPR*, 2017.

[9] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.

[10] V. Bazarevsky, Y. Kartynnik, A. Vakunov, K. Raveendran, and M. Grundmann, "Blazeface: Sub-millisecond neural face detection on mobile gpus," *arXiv preprint arXiv:1907.05047*, 2019.

[11] Delft AI Cluster (DAIC), "The delft ai cluster (daic), rrid:scr_025091," 2024. [Online]. Available: https://doc.daic.tudelft.nl/

[12] A. Kuzdeuov, D. Aubakirova, D. Koishigarina, and H. A. Varol, "Tfw: Annotated thermal faces in the wild dataset," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 2084–2094, 2022.