medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

AutoTriage - An Open Source Edge Computing Raspberry Pi-based Clinical Screening System

Chaitra Hegde[1*], Zifan Jiang[2*], Jacob Zelko[2], Pradyumna Byappanahalli Suresha[1], Rishikesan Kamaleswaran[3], Matt A. Reyna[3], and Gari D. Clifford[2,3]

[1]School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA
[2]Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, GA
[3]Department of Biomedical Informatics, Emory University, Atlanta, GA
[*]Both authors contributed equally to this work

Abstract: With the recent COVID-19 pandemic, healthcare systems all over the world are struggling to manage the massive increase in emergency department (ED) visits. This has put an enormous demand on medical professionals. Increased wait times in the ED increases the risk of infection transmission. In this work we present an open-source, low cost, off-body system to assist in the automatic triage of patients in the ED based on widely available hardware. The system initially focuses on two symptoms of the infection - fever and cyanosis. The use of visible and far-infrared cameras allows for rapid assessment at a 1m distance, thus reducing the load on medical staff and lowering the risk of spreading the infection within hospitals. Its utility can be extended to a general clinical setting in non-emergency times as well to reduce wait time, channel the time and effort of healthcare professionals to more critical tasks and also prioritize severe cases.

Our system consists of a Raspberry Pi 4, a Google Coral USB accelerator, a Raspberry Pi Camera v2 and a FLIR Lepton 3.5 Radiometry Long-Wave Infrared Camera with an associated IO module. Algorithms running in real-time detect the presence and body parts of individual(s) in view, and segments out the forehead and lip regions using PoseNet. The temperature of the forehead-eye area is estimated from the infrared camera image and cyanosis is assessed from the image of the lips in the visible spectrum. In our preliminary experiments, an accuracy of 97% was achieved for detecting fever and 77% for the detection of cyanosis, with a sensitivity of 91% and area under the receiver operating characteristic curve of 0.91.

Although preliminary results are promising, we note that the entire system needs to be optimized before use and assessed for efficacy. The use of low-cost instrumentation will not produce temperature readings and identification of cyanosis that is acceptable in many situations. For this reason, we are releasing the full code stack and system design to allow others to rapidly iterate and improve the system. This may be of particular benefit in low-resource settings, and low-to-middle income countries in particular, which are just beginning to be affected by COVID-19.

Coronavirus | COVID-19 | Cyanosis | Edge Computing | Far Infra-red Camera | Fever | Object Detection | Triage

Correspondence: ***gari@gatech.edu***

Introduction

With a dramatic increase in emergency department (ED) visit rates over the last four decades, both in the United States and around the world [24, 19, 20, 29], accurate and timely triage is essential for assuring patients' safety and optimal resource allocation. Crowding in the ED can affect the triage process, leading to longer waiting times for triage, longer ED length of stays and potentially poorer outcomes [14]. More acutely, crowding in ED during a pandemic such as COVID-19 could increase the risk for health professionals as well as patients.

Computer-aided triage systems have been proposed over the years with the help of browser-based applications that exchange information with existing medical records [1], wearable sensors [11] and automatic initial interpretation of CT scans [12]. However, most existing methods either require significant interaction between the patients and the healthcare workers or need active input from the patients. Hence, there is an emerging need for an automatic triage system that works passively and requires minimal attention and interaction from both patients and health professionals. In this work, we focus on real-time identification of febrile status and cyanosis in patients, since these are two critical symptoms in the recent COVID-19 crisis.

The Emergency Severity Index (ESI), used by most EDs in the United States [17], records the febrile state of young children and manifestation of cyanosis in all age groups. While core temperature is difficult to measure non-invasively, there is some evidence that infrared cameras are able to do so to some level of acceptable accuracy [8]. In particular, we considered temperature in the forehead area and color distribution of the lip as indicators for the febrile state and cyanosis.

Many systems that detect a febrile state via infrared imaging [22, 26, 27, 13] have been reported and evaluated over the years and a deep learning method was introduced for face segmentation in [13]. Cyanosis is a bluish discoloration of the skin or other areas of the peripheral body resulting from poor circulation or inadequate oxygenation of the blood. More specifically, it is due to increased concentration of reduced hemoglobin (Hb) in the circulation and is clinically evident at an oxygen saturation of 85% or less. Mild cyanosis is more challenging to detect. Cyanosis can be observed in the lips, ears, trunk, nailbed, hands and conjunctiva. Circumoral areas (around the mouth) have been compared in detecting cyanosis resulting from arterial hypoxemia. It has been noted that while the tongue is the most sensitive area, the lips are more specific ([6], chapter 45). For this work, we therefore focused on the lips, since they are easier to observe than the tongue and have been identified in ED emergency response

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

assessment systems.

Our proposed system consists of a low cost minicomputer - a Raspberry Pi (RasPi), a Google Coral USB accelerator tensor processing unit (TPU), a visible camera and a thermal camera, which are all portable and relatively inexpensive. By leveraging object detection and machine learning classification techniques, the system was designed to be capable of segmenting out regions of interest and classify febrile state and cyanosis in real-time.

Methods

Hardware configuration. Fig. 1 shows the proposed system. The various parts and their approximate costs are: a RasPi 4 with 4 GB of RAM ($55) and a 16 GB microSD card ($7), a Google Coral USB accelerator ($74.99), a visible light (Red, Green and Blue; RGB) RasPi Camera v2 ($10.99) and a FLIR Lepton 3.5 Radiometry Long-Wave Infrared Camera ($200) with its associated IO module ($100). The total off-the-shelf cost of the system is around $450. The visible light and thermal videos are captured at 25hz and 9hz, respectively. The RasPi was used as the main processor which takes in the video streams, processes them and can display the resulting video in real-time on an external monitor (via the HDMI connector). The forehead and lip detection were computed on the Google Coral USB accelerator, since this part of the algorithm required a deep neural network and hence significantly more computational power. The full setup can be seen Fig. 2.

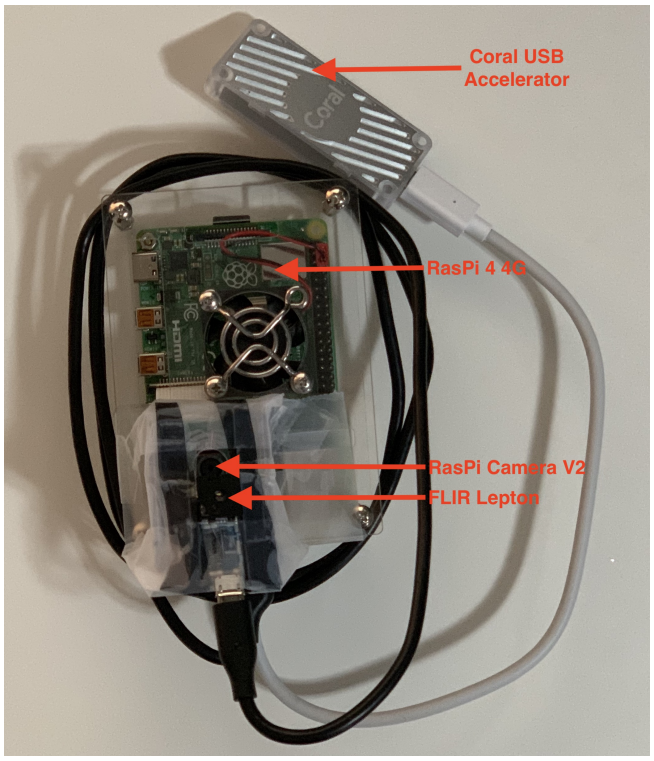

Fig. 1. Hardware configuration consisting of a Raspberry Pi 4 (4 GB RAM), a Google Coral USB accelerator (top), a RasPi Camera v2 and a FLIR Lepton 3.5 Radiometry Long-Wave Infrared Camera.

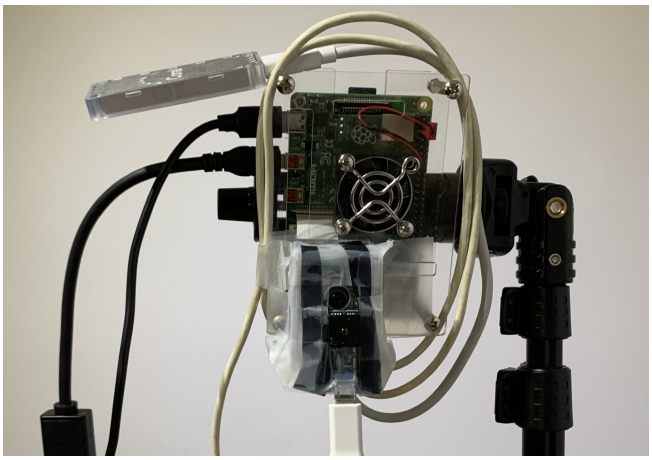

Fig. 2. Installation of system with all cabling including power and external monitor for visualization (not show, or needed for detection).

Overview of algorithms. Our algorithm performs two functions: fever estimation and cyanosis classification. For fever estimation, we detect the forehead region, and for cyanosis classification we detect the lip region. These regions are then analyzed to identify color distributions indicative of fever or cyanosis using a post processing machine learning algorithm. The overall flow is shown in Fig. 3 and a more detailed description of the algorithms now follows.

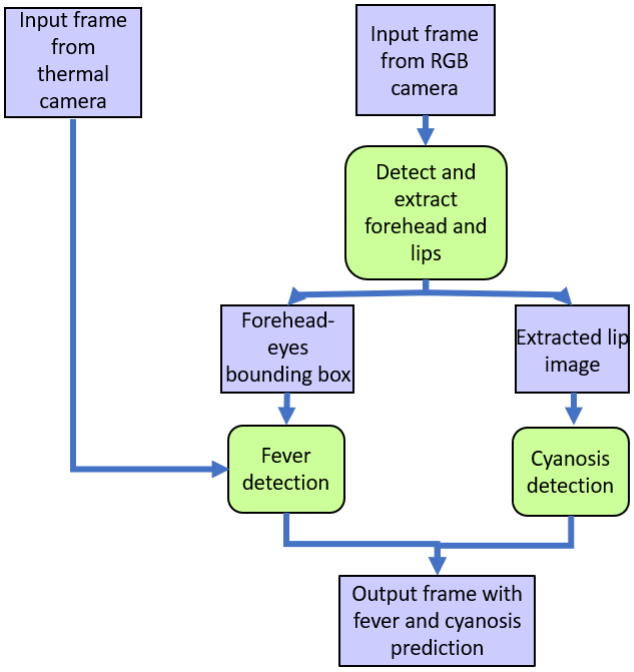

Fig. 3. AutoTriage pipeline. The purple boxes denote inputs and outputs at different stages. The green boxes denote computations.

Forehead and lip detection. To detect and segment out the forehead and lips for further analysis, a computer vision-based human pose detection algorithm knows as PoseNet is used [23]. This algorithm uses a convolutional neural network to regress the six degrees of freedom camera pose from a single RGB image in an end-to-end manner with no need for additional engineering or optimisation. PoseNet detects

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

"keypoints" - the image coordinates of key parts of the body, such as the elbow, knee, eyes and nose. We use the keypoints of the left and right eye as references to identify a bounding box around the forehead and lips. If we define the distance between the eyes to be D pixels, for the forehead, we use a bounding box that has width 2D and a height 1.2D. The base of the bounding box is 0.2D below the eyes. This ensures that the forehead-eyes area is captured. This is an important site, since the inner canthus of the eye is consistently the warmest area on the head and the most suitable area for fever detection [28]. For the lips, we move a distance D below the eyes and create a bounding box with width D and height 0.5D. These values were set empirically (through trial and error using the authors) to best capture the lips. In future experiments these values can be optimized on larger datasets. Note that the coordinates obtained by applying this heuristic are rounded to the closest integer value. Fig. 4 shows an example of the forehead and lip detection on one of the authors using this approach in the visible spectrum (upper plot) and the corresponding FLIR image (lower plot). Note that the images have slight FOV, image angle and translational differences since the cameras cannot take images from the same location in space and operate at different (non-synchronous) sampling rates.

It is important to note that there are other state-of-the-art object detection algorithms that are capable of detecting faces, such as YOLO [25]. YOLO can also be used to segment out the forehead and lips. However, using PoseNet in this implementation allows for rapid deployment on the Google Coral Accelerator TPU and is therefore simple to implement in real time. Another benefit of using PoseNet is that the keypoints generated by it for the rest of the body could potentially be used in the future to further enhance assessment of other symptoms such as posture or balance abnormalities, or behaviors such as vomiting, falls, and tonic-clonic seizures.

Febrile state detection. Unlike the previous studies, our proposed system detects key points using visible light video and then transforms the coordinates of the bounding boxes of the region of interest (ROI) to coordinates in the thermal video. After finding the ROI in the thermal video, the average of the ten pixels exhibiting the highest temperatures were averaged to produce a final temperature estimate. Lastly, a threshold was set to determine the febrile state.

Thermal output calibration. To achieve a more accurate measurement of the body temperature, we followed the guidelines from the FLIR Lepton 3.5 datasheet [10] and methods detailed in previous studies [13] that used a similar infrared camera and performed the radiometry calibration on the FLIR Lepton camera. The conversion from the 14 bit pixel values from the FLIR Lepton camera to a temperature in Celsius can be written as:

$$T = \frac{B}{exp(R/S - O) + F} - 273.15 \quad \textbf{(1)}$$

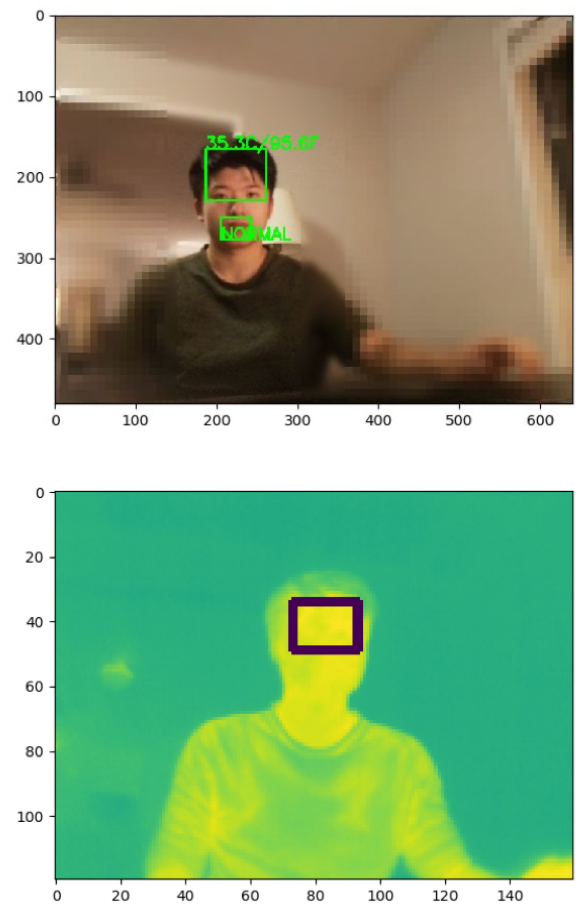

Fig. 4. Bounding boxes detected for the forehead and lips from a 1m range in visible light video (upper image) and thermal video (lower image) using PoseNet.

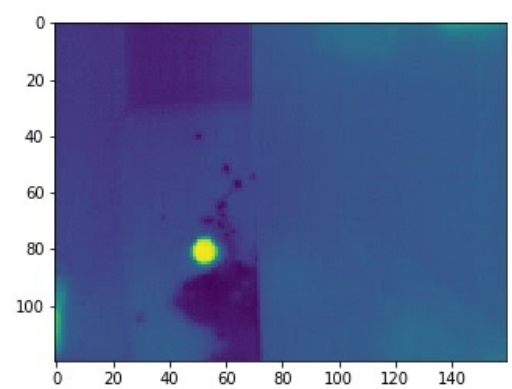

Fig. 5. Illustration of the water calibration experiment. The bright spot in the lower left quadrant of the FLIR output represents the water heated to varying known temperatures.

where S denotes the output pixel value from the FLIR Lepton, T denotes the actual temperature and R, B, F, O are parameters required for the conversion. Within the FLIR Lepton's operating temperature range, the typical values of F and B are 1 and 1428, respectively. We used bottles (with open lids) of heated water with temperature ranging from

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

$35-40°$C as a heat source and located them at one meter to the camera and approximately in the center of the field of view (FOV). (See Fig. 5.) The reference temperature of the water was measured three times using a Braun IRT6500 thermometer and averaged. The reference has an accuracy of $\pm 0.2°$C within the $35-42°$C measurement range. Also, the top ten pixels in the heat source were selected and averaged as the final output of the FLIR Lepton camera. R (camera responsivity) and O (the offset) in Eq. 1 are then fitted using the above-described experiment via a nonlinear least square method with soft $l1$ loss as follows:

$$L_{1;soft}(z) = 2 \cdot (\frac{1}{(1+z)^{0.5}} - 1). \quad \textbf{(2)}$$

The Root Mean Square Error (RMSE) metric was used to evaluate the fitting error. The fitted curve was then implemented to convert pixels values to temperature output.

ROI registration in thermal video. Since the forehead-eyes area is detected in the visible light image sequences, a coordinates transformation is needed to find the forehead-eyes location in the thermal video. The transformation depends on the resolutions and FOVs of the two cameras and the relative physical displacement between them. The resolution of the RasPi Camera was set at 1640×1232 pixels and the corresponding FOV was $62.2°$ horizontally and $48.8°$ vertically. The resolution of the FLIR Lepton camera was set to be 160×120 pixels and the corresponding FOV was $57°$ horizontally and $71°$ diagonally. Since there only exist less than 2 cm distance and consequently a small angle difference between the cameras, the transformation can be expressed as:

$$\begin{aligned} xL_{thermal} &= \frac{xL_{visible}+x_{bias}}{x_{ratio}} \\ yL_{thermal} &= \frac{yL_{visible}+y_{bias}}{y_{ratio}} \\ xR_{thermal} &= \frac{xR_{visible}+x_{bias}}{x_{ratio}} \\ yR_{thermal} &= \frac{yR_{visible}+x_{bias}}{y_{ratio}} \end{aligned} \quad \textbf{(3)}$$

Where xL, yL are the coordinates for the left vertex of the bounding box and xR, yR denote the right vertex. x_{ratio}, y_{ratio} are the resolution ratios between RasPi Camera and FLIR Lepton Camera. And x_{bias}, y_{bias} are the view difference caused by different FOVs and can be calculated as:

$$\begin{aligned} h &= \tan\frac{FOV_{horizontal-thermal}}{2} / \tan\frac{FOV_{horizontal-visible}}{2} \\ v &= \tan\frac{FOV_{vertical-thermal}}{2} / \tan\frac{FOV_{vertical-visible}}{2} \\ x_{bias} &= Resolution_{visible-horizontal} \cdot \frac{h-1}{2} \\ y_{bias} &= Resolution_{visible-vertical} \cdot \frac{v-1}{2} \end{aligned} \quad \textbf{(4)}$$

In practice, because of the angle difference and distance between the cameras, which varies with different mounting schemes, empirical offsets were added to ensure accurate transformation.

Threshold selection. The Merck Manual [3] defines fever as an elevated body temperature that is higher than $37.8°$C orally [16] and Cleveland clinic advised patients with a fever higher than $100.4°$F / $38°$C to isolate themselves as of April 2020 [7]. Forehead (temporal) temperature is usually 0.5 °F (0.3°C) to 1°F (0.6°C) lower than an oral temperature measurement [9].

Experimental setup. The preliminary test was conducted on a healthy male subject. A heated cloth was put on the subject's forehead to raise the temperature of the subject's forehead. The estimated temperature from the proposed system was recorded immediately after a measurement from the thermometer in the center of the subject's forehead. The subject was sitting one meter away from the cameras. The RMSE between the estimated value and temperature from the thermometer was used to evaluate the accuracy of our proposed system.

Cyanosis detection. Once the lip region is segmented from the face, an algorithm to detect cyanosis is applied to the region. Various algorithms were tested on a small dataset. The algorithm that performed the best is used to classify images from the camera as cyanotic or non-cyanotic. Leave-one-out cross validation method was used to report the results.

Dataset. We created a small dataset consisting of images of cyanotic and non-cyanotic lips from the internet. It consists of 35 images of cyanotic lips and 35 images of non-cyanotic lips. The images belong to various races, ages and genders, although this is not balanced. The number of lighter skin colored images are higher than darker skin images, more so for cyanotic images. Also, the number of images of young to middle-aged people is higher than the number of images of other age groups. This is due to the limited number and types of cyanotic images available on the internet. The dataset is included in our Github repository [5].

Classification algorithm. We implemented three classifiers: K nearest neighbors (KNN), logistic regression (LR) and support vector classifier (SVC). The input to these classifiers were frequency of pixel intensities from each channel (R, G, B channels) using a simple histogram with eighteen equally spaced bins (six for each of the three color channels). The rationale behind this was that the color distribution would be different in cyanotic and non-cyanotic lips, but there would be some colors in common. In other words, not all of the lip would be cyanotic, and some areas outside of the lip would be included. The bins representing these colors could then be regularized out. For the KNN algorithm, we used k=3 neighbors, uniform weights and the Euclidean distance metric. (We chose three clusters for the KNN to capture the two classes plus noise.) L2 regularization was used for the logistic regression implementation. For the SVC a regularization parameter $C = 2$ was chosen via grid search. Radial basis function kernel was used to make SVC a non-linear classifier. We used the *sklearn* package in Python3.7.3 for each of these implementations and used the default values for parameters not stated above.

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Results

Febrile state detection. Fig. 6 shows the measurements used in the calibration process and the fitted results. The calibrated values of R and O were determined to be 580061 and 25113, respectively, and the resultant RMSE was $0.57°C^2$. Fig. 7 shows the measured data points in the preliminary experiment, in which the proposed system achieved an RMSE of $0.41°C$ and a Pearson correlation coefficient of 0.96. When applying $37.4°C$ as the threshold for febrile state detection, an accuracy of 96.7% and an area under receiver operating characteristic curve (AUC) of 0.97 were achieved.

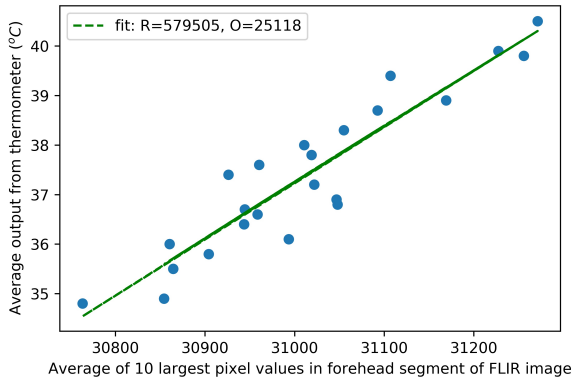

Fig. 6. Calibration of FLIR images using a bottle of water: The least square curve of best fit between the average temperature from the thermometer and the FLIR pixel values had a root mean square error (RMSE) of 0.57 °C.

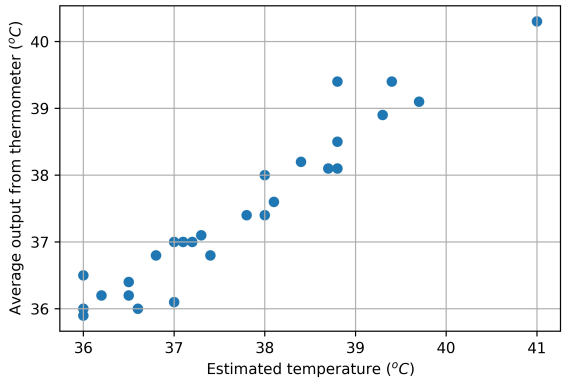

Fig. 7. Estimated temperature vs. temperature measured from thermometer. The correlation between the parameters was found to be 0.96 and the RMSE difference between them was 0.41 °C.

Table 1. Confusion Matrix for KNN for cyanosis detection.

		Actual Values	
		Cyanotic	Non-cyanotic
Predicted Values	Cyanotic	28	13
	Non-cyanotic	7	22

Cyanosis detection. Tables 1, 2, 3 show the confusion matrices for each classifier. Fig. 8 shows the receiver operating characteristic curve and AUC of each classifier. Table 4 summarizes the accuracies, AUC, sensitivity and specificity of the three classification models.

Table 2. Confusion Matrix for LR for cyanosis detection.

		Actual Values	
		Cyanotic	Non-cyanotic
Predicted Values	Cyanotic	27	10
	Non-cyanotic	8	25

Table 3. Confusion Matrix for SVC for cyanosis detection.

		Actual Values	
		Cyanotic	Non-cyanotic
Predicted Values	Cyanotic	32	13
	Non-cyanotic	3	22

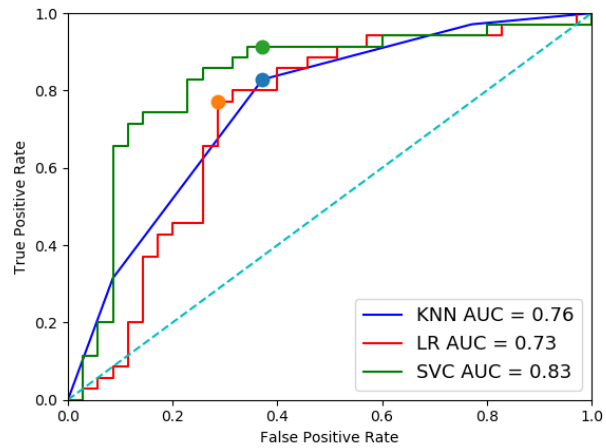

Fig. 8. Receiver Operating Characteristic Curves for the three classifiers evaluated in this work for cyanosis detection. The filled circles represent the operating points resulting in the other performance statistics. (Small differences exist due to the LOOCV approach.)

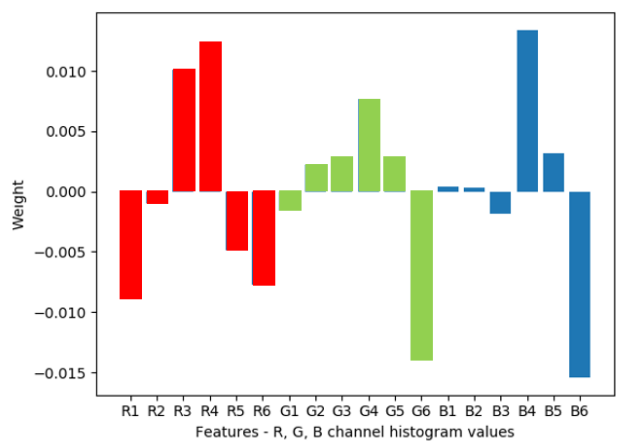

Fig. 9. Weighting of features from logistic regression. The features are the histogram values the R, G and B channels. Positive coefficients refer to cyanotic condition and negative refer to non-cyanotic.

To assess the relative importance of the features used for classification, we visualized the weights assigned to different

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Table 4. Performance of assessed cyanosis detection classifiers.

Algorithm	Accuracy (%)	AUC	Se	Sp
KNN	71.4	0.76	0.83	0.63
LR	74.3	0.73	0.77	0.71
SVC	77.2	0.83	0.91	0.63

features by the logistic regression classifier, since this is easier to interpret that the parameters of the SVR or KNN. Fig. 9 shows the weights for the eighteen features used (six bins each for each of the three R, G and B channels). The first six features correspond to the red channel (R1 - R6). The next six features correspond to the green channel (G1 - G6) and the last six features correspond to the blue channel (B1 - B6).

Discussion

Forehead and lip segmentation. Detecting and segmenting out the forehead and lips is the first step in our pipeline. The accuracy of this stage can determine the accuracy of the remaining stages. This step is dependent on the performance of PoseNet, which sometimes has false positive detection of people. In this implementation, it is limited to detecting 10 people at a time.

Fever detection. Previous meta-reviews suggest that peripheral temperature may not be sufficient to determine fever [21, 18]. This suggests that our proposed system, along with all traditional methods that measure peripheral temperature, like an ear thermometer, are not suitable to be used to perform an accurate diagnosis of febrile state. However, the proposed system is useful to perform mass early screening of the febrile state as a triage tool. The body temperature varies throughout the day in accordance with the circadian rhythm ([6], chapter 218). Taking this into account and having a dynamic threshold can reduce the number of false positive and false negative fever detection. Besides, body temperature could vary based on the ambient temperature. Having a reference temperature can help solve this issue. Additionally, the thermal camera in selection does not innately have the level of accuracy required for this task. Thus, a calibration for the targeted temperature range needs to be performed. However, the temperature calibration of the thermal camera is not a trivial task and can be inaccurate depending various aspects, like environmental temperature and surface condition. Also, a previous study suggests that improper use and interpretation of the infrared camera can lead to inaccurate triage [15]. Hence, it is important to understand that the proposed system is only reliable for the designated task under limited conditions. For example, the presence of common cosmetics or clothing such as a turban or hijab can affect the accuracy of the estimated temperature from thermal camera [30]. With a higher budget, the use of a thermal camera with higher accuracy and, if possible, one which is pre-calibrated can lead to a more reliable system. But a higher cost will inevitably lower the accessibility of the system.

Cyanosis detection. For the detection of cyanosis, out of the three classification approaches evaluated, the SVC exhibited the highest accuracy (77.1%) and KNN has the highest AUC (0.93), although this is not significantly higher than the AUC of 0.91 for the SVR. The SVR also produces the fewest false negatives (missed cynaosis), which at triage, is probably the most important feature of this system. For the open source implementation, we therefore chose the SVR, although we note this is somewhat arbitrary at this point given the size of the data set we used.

To visualize the effect of the features on the overall classification we plot the LR coefficients for each features (see Fig. 9). It can be observed that the red channel exhibits higher importance for cyanosis detection. When creating a histogram of the pixel values with six equally spaced bins for each of the three color channels, it can be seen that every channel (R, G and B) contributes to the classification.

Limitations. We note that the lip cyanosis dataset we used is relatively small and contains relatively good quality images. In the wild, the quality of images is not guaranteed to be high. This may be due to variations in lighting, occlusions, movement, angle of presentation, distances much greater than 1m, among other issues. Consistency of ambient lighting is an important factor to ensure correct classification of cyanosis [6]. Camera parameters such as field of view and shutter speed can also influence the absolute color detected by the camera. Applying a color correction by using a color reference in the frame can solve this issue [2]. Cosmetics applied to and around the lip can also interfere with the classification. In practice, we observed that if the mouth is open or teeth are visible, the cyanosis classifications tend to be inaccurate. Lighting conditions also play a major role in the algorithm's output.

Perhaps the most important issue to consider is that of skin color and the variation of presentation of cyanosis across the human race. This is an under-explored area of research, but research into racial bias in facial recognition algorithms [4] has highlighted just how dangerous it can be to use these algorithms out-of-the-box, without tuning to a population or thought about the bias it can create.

Conclusion

In this work, we have proposed a system that can detect fever and cyanosis using a combination of visible light and thermal camera operating on an edge computation platform that is running state-of-the-art deep learning. The system does not require any direct interaction between the device and either patients or healthcare workers. The source code needed to replicate our proposed system can be found on Github [5]. It is important to note that PoseNet is image size and rotationally invariant (at least for most behaviors), and although we optimize the analysis to work at a 1m distance from the camera, this invariance should create a robustness to movement to and from the camera, as well as within the frame. Many improvements can be made to this system to increase the classification performance and stability, including larger

medRxiv preprint doi: https://doi.org/10.1101/2020.04.09.20059840; this version posted April 11, 2020. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

population studies and end-to-end deep learning. However, the need for *something* is acute and will be increasingly so in low resource areas. We therefore feel it is appropriate to release this work prior to peer-review to solicit feedback and encourage others to improve the system.

Through this work, put together as a rapid response under a few days under lockdown, we hope to provide a starting point for automatic triage in clinical settings. Improving on this work could lead to novel implementations that can help streamline triage in clinics and hospitals, potentially during the current pandemic, where non-contact and rapid screening has distinct advantages for infection reduction.

ACKNOWLEDGEMENTS

This work was spun out from work funded by gifts from the Cox Foundation and the Center for Discovery.

References

[1] Dominik Aronsky et al. "An Integrated Computerized Triage System in the Emergency Department". In: *AMIA Annu Symp Proc* 2008 (2008). ZSCC: 0000046 https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2656061/, pp. 16–20. ISSN: 1942-597X. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2656061/ (visited on 04/02/2020).

[2] Nur Fatihah Binti Azmi et al. "Color Correction of Baby Images for Cyanosis Detection". In: *Medical Image Understanding and Analysis*. Ed. by Mark Nixon, Sasan Mahmoodi, and Reyer Zwiggelaar. Cham: Springer International Publishing, 2018, pp. 354–370. ISBN: 978-3-319-95921-4.

[3] Krystal Bullers. *Merck Manuals*. https://web.archive.org/web/20191001153051/https://www.merckmanuals.com/professional/infectious-diseases/biology-of-infectious-disease.

[4] Joy Buolamwini and Timnit Gebru. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification". In: *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Ed. by Sorelle A. Friedler and Christo Wilson. Vol. 81. Proceedings of Machine Learning Research. New York, NY, USA: PMLR, 23–24 Feb 2018, pp. 77–91. URL: http://proceedings.mlr.press/v81/buolamwini18a.html.

[5] Chaitra Hegde and Zifan Jiang, Pradyumna Byappanahalli Suresha, and Gari D. Clifford. *AutoTriage - An Open Source Raspberry Pi-based Clinical Screening System*. Version 0.1. Apr. 2020. DOI: 10.5281/zenodo.3740776. URL: https://github.com/cliffordlab/AutoTriage_release/releases.

[6] Vivian L Clark and James A Kruse. "Clinical methods: the history, physical, and laboratory examinations". In: *Jama* 264.21 (1990), pp. 2808–2809.

[7] *Coronavirus (COVID-19) Media Information*. https://newsroom.clevelandclinic.org/2020/04/02/coronavirus-covid-19-media-information/. Accessed: 2020-4-4.

[8] Canadian Agency for Drugs and Technologies in Health. *Non-contact thermometers for detecting fever: a review of clinical effectiveness*. 2014.

[9] *Fever Temperatures: Accuracy and Comparison*. https://www.cigna.com/individuals-families/health-wellness/hw/medical-topics/fever-temperatures-tw9223. Accessed: 2020-4-4.

[10] *FLIR LEPTON® Engineering Datasheet*. Lepton Engineering Datasheet, Rev: 200. FLIR Systems. Mar. 2018.

[11] Tia Gao and David White. "A next generation electronic triage to aid mass casualty emergency medical response". In: *2006 International Conference of the IEEE Engineering in Medicine and Biology Society*. Vol. Supplement. ISSN: 1557-170X. Aug. 2006, pp. 6501–6504. DOI: 10.1109/IEMBS.2006.260881. URL: 10.1109/IEMBS.2006.260881.

[12] Roman Goldenberg et al. "Computer-aided simple triage (CAST) for coronary CT angiography (CCTA)". In: *International journal of computer assisted radiology and surgery* 7.6 (2012), pp. 819–827.

[13] Jia-Wei Lin, Ming-Hung Lu, and Yuan-Hsiang Lin. "A Thermal Camera Based Continuous Body Temperature Measurement System". en. In: *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*. ZSCC: 0000000. Seoul, Korea (South): IEEE, Oct. 2019, pp. 1681–1687. ISBN: 978-1-72815-023-9. DOI: 10.1109/ICCVW.2019.00208. URL: https://ieeexplore.ieee.org/document/9022540/ (visited on 04/03/2020).

[14] M. Christien van der Linden, Barbara E. A. M. Meester, and Naomi van der Linden. "Emergency department crowding affects triage processes". en. In: *International Emergency Nursing*. Special Issue: Triage 29 (Nov. 2016). /https://www.sciencedirect.com/science/article/abs/pii/S1755599X16300052, pp. 27–31. ISSN: 1755-599X. DOI: 10.1016/j.ienj.2016.02.003. URL: 10.1016/j.ienj.2016.02.003 (visited on 04/02/2020).

[15] James B Mercer and E Francis J Ring. "Fever screening and infrared thermal imaging: concerns and guidelines". In: *Thermology International* 19.3 (2009), pp. 67–69.

[16] *Merck Manual. Fever*. https://web.archive.org/web/20191001153051/https://www.merckmanuals.com/professional/infectious-diseases/biology-of-infectious-disease/fever. Accessed: 2019-10-01.

[17] Amir Mirhaghi. "Most patients are triaged using the emergency severity index". en. In: *J. Nucl. Cardiol.* 24.2 (Apr. 2017). ZSCC: 0000000, pp. 738–738. ISSN: 1071-3581, 1532-6551. DOI: 10.1007/s12350-016-0704-z. URL: http://link.springer.com/10.1007/s12350-016-0704-z (visited on 04/03/2020).

[18] Binoy Mistry et al. "Accuracy and Reliability of Emergency Department Triage Using the Emergency Severity Index: An International Multicenter Assessment". en. In: *Annals of Emergency Medicine* 71.5 (May 2018), 581–587.e3. ISSN: 01960644. DOI: 10.1016/j.annemergmed.2017.09.036. URL: https://linkinghub.elsevier.com/retrieve/pii/S0196064417317456 (visited on 04/02/2020).

[19] Claire Morley et al. "Emergency department crowding: a systematic review of causes, consequences and solutions". In: *PloS one* 13.8 (2018).

[20] John C Moskop et al. "Another look at the persistent moral problem of emergency department crowding". In: *Annals of emergency medicine* 74.3 (2019), pp. 357–364.

[21] Bruce K Muma et al. "Comparison of rectal, axillary, and tympanic membrane temperatures in infants and young children". en. In: *Annals of Emergency Medicine* 20.1 (Jan. 1991). ZSCC: 0000121 https://www.sciencedirect.com/science/article/abs/pii/S0196064405811163, pp. 41–44. ISSN: 0196-0644. DOI: 10.1016/S0196-0644(05)81116-3. URL: 10.1016/S0196-0644(05)81116-3 (visited on 04/04/2020).

[22] Eddie Y.K Ng, G.J.L Kawb, and W.M Chang. "Analysis of IR thermal imager for mass blind fever screening". en. In: *Microvascular Research* 68.2 (Sept. 2004). ZSCC: 0000158, pp. 104–109. ISSN: 00262862. DOI: 10.1016/j.mvr.2004.05.003. URL: https://linkinghub.elsevier.com/retrieve/pii/S0026286204000548 (visited on 04/03/2020).

[23] George Papandreou et al. "PersonLab: Person Pose Estimation and Instance Segmentation with a Bottom-Up, Part-Based, Geometric Embedding Model". In: *The European Conference on Computer Vision (ECCV)*. Sept. 2018.

[24] Jesse M. Pines et al. "International Perspectives on Emergency Department Crowding". en. In: *Academic Emergency Medicine* 18.12 (2011). ZSCC: NoCitationData[s0] https://onlinelibrary.wiley.com/doi/full/10.1111/j.1553-2712.2011.01235.x, pp. 1358–1370. ISSN: 1553-2712. DOI: 10.1111/j.1553-2712.2011.01235.x. URL: 10.1111/j.1553-2712.2011.01235.x (visited on 04/02/2020).

[25] Joseph Redmon et al. "You Only Look Once: Unified, Real-Time Object Detection". In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016.

[26] E.F.J. Ring et al. "Detecting Fever in Polish Children by Infrared Thermography". en. In: *Proceedings of the 2008 International Conference on Quantitative InfraRed Thermography*. ZSCC: NoCitationData[s0]. QIRT Council, 2008. DOI: 10.21611/qirt.2008.03_07_17. URL: http://qirt.gel.ulaval.ca/archives/qirt2008/papers/03_07_17.pdf (visited on 04/03/2020).

[27] Armote Somboonkaew et al. "Mobile-platform for automatic fever screening system based on infrared forehead temperature". en. In: *2017 Opto-Electronics and Communications Conference (OECC) and Photonics Global Conference (PGC)*. ZSCC: NoCitationData[s0]. Singapore: IEEE, July 2017, pp. 1–4. ISBN: 978-1-5090-6293-5. DOI: 10.1109/OECC.2017.8114910. URL: http://ieeexplore.ieee.org/document/8114910/ (visited on 04/03/2020).

[28] ISOISO TR. "13154: Medical electrical equipment—Deployment, implementation and operational guidelines for identifying febrile humans using a screening thermograph". In: *International Organization for Standardization* (2009).

[29] Lindsey Woodworth. "Swamped: Emergency Department Crowding and Patient Mortality". In: *Journal of Health Economics* 70 (2020), p. 102279.

[30] Kaikai Zheng et al. *Infrared assessment of human facial temperature in the presence and absence of common cosmetics*. en. preprint. ZSCC: 0000000. Epidemiology, Mar. 2020. DOI: 10.1101/2020.03.12.20034793. URL: http://medrxiv.org/lookup/doi/10.1101/2020.03.12.20034793 (visited on 04/03/2020).