Multispectral Video Fusion for Non-contact Monitoring of Respiratory Rate and Apnea

Journal Article**Author(s):**

Scebba, Gaetano; Da Poian, Giulia; [Karlen, Walter](#)

Publication date:

2021-01

Permanent link:

<https://doi.org/10.3929/ethz-b-000414662>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IEEE Transactions on Biomedical Engineering 68(1), <https://doi.org/10.1109/tbme.2020.2993649>

Funding acknowledgement:

150640 - Intelligent Point-of-Care Monitoring: A Swiss Army Knife Approach to mHealth (SNF)

Multispectral Video Fusion for Non-contact Monitoring of Respiratory Rate and Apnea

Gaetano Scebba, Giulia Da Poian, and Walter Karlen, *Senior Member, IEEE*

Abstract—Continuous monitoring of respiratory activity is desirable in many clinical applications to detect respiratory events. Non-contact monitoring of respiration can be achieved with near- and far-infrared spectrum cameras. However, current technologies are not sufficiently robust to be used in clinical applications. For example, they fail to estimate an accurate respiratory rate (RR) during apnea. We present a novel algorithm based on multispectral data fusion that aims at estimating RR also during apnea. The algorithm independently addresses the RR estimation and apnea detection tasks. Respiratory information is extracted from multiple sources and fed into an RR estimator and an apnea detector whose results are fused into a final respiratory activity estimation. We evaluated the system retrospectively using data from 30 healthy adults who performed diverse controlled breathing tasks while lying supine in a dark room and reproduced central and obstructive apneic events. Combining multiple respiratory information from multispectral cameras improved the root mean square error (RMSE) accuracy of the RR estimation from up to 4.64 monospectral data down to 1.60 breaths/min. The median F1 scores for classifying obstructive (0.75 to 0.86) and central apnea (0.75 to 0.93) also improved. Furthermore, the independent consideration of apnea detection led to a more robust system (RMSE of 4.44 vs. 7.96 breaths/min). Our findings may represent a step towards the use of cameras for vital sign monitoring in medical applications.

Index Terms—Respiratory rate, apnea, non-contact physiological monitoring, nearables, multispectral fusion, thermal imaging.

I. INTRODUCTION

MONITORING respiratory activity is critical for assessing the state of health in humans. Respiratory rate (RR) is an important vital sign and a strong predictor of severe illness [1]. While normal RR ranges between 12-20 breaths/min in healthy adults at rest, RR outside of this range strongly correlates with specific adverse events, such as congestive heart failure [2]. Furthermore, monitoring respiratory activity is used to diagnose breathing disorders and lung diseases, such as pneumonia and central or obstructive sleep apnea. Despite its clinical relevance, the continuous and accurate monitoring of RR is highly undervalued [3].

This work was supported by the Swiss National Science Foundation (SNSF) under Grant 150640.

G. Scebba, G. Da Poian and W. Karlen are with the Mobile Health Systems Lab, Institute of Robotics and Intelligent Systems, Department of Health Sciences and Technology, ETH Zurich, Switzerland (e-mails: gaetano.scebba@hest.ethz.ch, walter.karlen@ieee.org)

Fig. 1: Quantifying respiratory activity using far-infrared (FIR) and near-infrared (NIR) cameras. Dedicated fusion models exploit the respiratory information from the multispectral cameras and estimate respiratory rate (RR) and detect apnea.

Current approaches to objectively assess respiratory activity require direct contact with the patient. The obtrusiveness of these approaches reduces their acceptability in the clinical and telemedicine settings [4]. Therefore, there has been an increased interest in developing nearables for non-contact, unobtrusive monitoring of vital signs. One of these approaches includes the estimation of RR using cameras [5]. The ubiquitousness of cameras and the advancement in computer vision enable physiological monitoring in more natural settings at home or in emergency settings for triage of potentially contagious patients [6]. Furthermore, the absence of direct contact with the skin allows for the monitoring of pre-term infants and other patients who suffer from skin irritations due to prolonged use of electrodes [7]. Another advantage of using cameras is the direct recognition of motion artifacts that reduce the vital sign quality and trigger false alarms [8].

Currently, cameras to monitor RR are not widely adopted in clinical settings as they lack accuracy during continuous operation. The robustness of the measurement method is impacted in many ways. RGB cameras operating in the visible spectrum require moderate illumination, preventing their application in dark environments [9]. Most far-infrared (FIR) cameras are high-end products and excessively expensive, whereas consumer-accessible FIR cameras are challenged by low pixel resolution and sampling rates, which makes the localization of a specific region of interest (ROI) difficult [10]. An important limitation of most current algorithms for camera-based respiration monitoring is that they have only been tested under optimal conditions and with a low number of subjects. Furthermore, many of these algorithms only implement a

trivial apnea detection that fails under the presence of noise.

We propose a novel non-contact system to monitor respiratory activity by combining near-infrared (NIR) and FIR cameras (Fig. 1). Our innovative approach extracts multiple streams of respiratory information from independent video modalities and then applies two dedicated data fusion models to determine whether the subject is breathing (apnea detection), and at which frequency (RR estimation). We designed and conducted a series of experiments to demonstrate that fusing multispectral videos leads to a significant improvement in the estimation of RR when compared to state-of-the-art methods, and increases the robustness of RR estimates when apnea events are present.

II. BACKGROUND

A. Multisensory fusion

Data fusion has been successfully implemented in biomedical signal processing for RR estimation with conventional contact sensors [11]. Numerous approaches fuse either the RR from photoplethysmographic (PPG) [12]–[14] or electrocardiographic (ECG) waveform modulations [15]–[17]. Others propose combining multiple sensor modalities, such as ECG and PPG [18]–[20], ECG and thoracic impedance [21], and ECG and accelerometry [22]. Fusing multiple vital signs is also advantageous in detecting apneic events. Furthermore, models exist that combine oxygen saturation with other modalities, such as ECG [23], [24], electroencephalography [25], respiratory effort and ECG [26], nasal airflow [27] or tracheal sounds [28].

In this work, we build on several of these approaches that have demonstrated the benefit of fusing multisensory information to increase accuracy and robustness of the RR estimation and apnea detection. However, while existing literature focused on vital signs obtained from contact sensors, our data fusion algorithms leverage respiratory signs that can be solely obtained from video data, thus enabling non-contact estimation of respiratory activity.

B. Respiratory activity detectable with digital cameras

Respiratory activity and its absence are characterized by specific patterns that can be extracted through appropriate analysis of videos recorded with RGB, NIR or FIR cameras. Four respiratory signs are directly observable: 1) respiratory induced motion, 2) thermal airflow variation, 3) respiratory plethysmographic modulation, and 4) apneic events.

1) *Respiratory induced motion (RM)*: RM of the torso is the most noticeable respiratory sign and is visible to the naked eye. It is produced by the volume variation of the lungs that generate the inflow and outflow of air. In particular, the lungs can be expanded and contracted generating the periodic motion of the torso in two distinct sections [29]. The motion of the abdominal wall is caused by the diaphragm that vertically lengthens and shortens the lungs, and the elevation and depression of the rib cage increases and decreases the anteroposterior diameter of the lungs [30]. When an obstructive apneic event occurs, the collapse of the upper airways prevents air exchange, but a persistent respiratory movement

of the torso is still present [31]. However, during a central apneic event, the RM is completely absent. Several algorithms have been proposed to quantify the RM. Differential image processing [32], [33] or motion tracking of salient points extracted from the chest ROI with optical flow [34] and dense optical flow [35] enable the extraction of respiratory activity.

2) *Thermal airflow (TA) signal*: The temperature fluctuations of the airways due to airflow are a respiratory sign that can be detected with FIR cameras [36]. The physical phenomenon is based on the radiative and convective heat transfer component during the breathing cycle [37], which results in a periodic increase and decrease of the temperature at the tissues around the nasal cavity. These observable temperature fluctuations are quantifiable in a FIR video as pixel intensity variations of the nostrils ROI [37]–[39]. During obstructive and central apnea there is a suspension of the air exchange between lungs and atmosphere and consequently, no temperature variation is observable at the nostrils.

3) *Respiratory plethysmographic (RP) modulation*: The various respiratory modulations of the PPG waveform are well described in literature [12]. The most prominent respiratory sign observable with cameras is the PPG baseline modulation. Each respiratory cycle causes variations of the intra-thoracic pressure that induce changes in the blood exchange between the pulmonary and systemic circulation. This phenomenon results in a periodic variation of the baseline of the PPG waveform. In the event of obstructive apnea, the persistent motion of the torso continues to induce the intrathoracic pressure variations. During a central apnea event, the RP modulation is drastically reduced in amplitude due to the absence of any RM. The PPG baseline modulation can be quantified through the linear combination of the RGB components in a video of skin [5]. Other studies apply blind-source separation [40], [41], continuous wavelet filtering [42], or auto-regressive models [4] to extract the RP modulation.

4) *Apneic events*: Apneic events are defined as the suspension of respiratory activity lasting for more than 10 seconds [43]. Apnea detection algorithms frequently focus on detecting an absence of RR [5], [44]. However, this is not reliable as frequency derived RR does not tend towards zero if there is noise present. Furthermore, this simplistic view on apnea detection prevents the distinction between the absence of breathing and the absence of a respiratory signal, as well as between obstructive and central apnea.

Pereira *et al.* have exploited the potential of combining multiple respiratory sources of information obtained from cameras [45]. They propose a black-box fusion approach, in which the TA fluctuations and the RM obtained from high-end FIR cameras are considered as independent data sources [44].

In contrast to the above described work, we present the first camera-based approach to monitor a broader range of respiratory activity by combining multiple types of respiratory modulations extracted from multiple modalities. Our approach simultaneously extracts multiple respiratory signals and with specialized data fusion models estimates the RR and classifies apnea. With a third fusion stage we combine that information to produce an apnea-aware RR quantification.

Fig. 2: Multispectral camera-based respiratory monitoring system. (a) The videos from far-infrared (FIR) and near-infrared (NIR) cameras are processed with a (b) Multispectral Region of Interest (ROI) localization algorithm [10]. The localized ROIs are used to extract (c) the Thermal Airflow (TA) signal from the nostril ROI, and (d) the Respiratory Motion signal both from the FIR and NIR chest ROIs. (e) In the Signal Quality based fusion (SQb Fusion), (e1) Respiratory Rate (RR) and Signal to Noise Ratio (SNR) are computed by the frequency analysis of the TA, and the Respiratory Motion from the FIR (RM_{FIR}) and the NIR (RM_{NIR}). (e2) The weighted median combines all RR estimates using their signal-to-noise-ratio (SNR) as weights and calculates a RR estimate. (f) The apnea detector (f1) extracts time and frequency features from the TA, RM_{NIR} , and RM_{FIR} signals and (f2) classifies the signals into apnea or breathing with an ensemble of support vector machines. (g) The smart signal quality based fusion (S^2 Fusion) fuses the results from the SQb Fusion and the apnea classifier (h) to obtain an apnea sensitive estimation of RR.

III. ALGORITHM DEVELOPMENT

We developed a system for estimating respiratory activity from synchronized NIR and FIR cameras (Fig. 2). We chose this pair as they can be operated independently of visible light. Our algorithmic approach consisted of three main steps. Firstly, we localized the nostril and chest ROIs in both the NIR and FIR spectra. Secondly, three respiratory signals were extracted from these ROIs. Finally, using our novel Smart Signal quality based Fusion (S^2 Fusion), a RR estimate obtained from a modulation fusion model and an apnea classification model were combined to obtain a more accurate quantification of the respiratory activity.

A. Multispectral ROI localization

The detection and tracking of the ROIs was built upon a multispectral ROI detector that we previously proposed in [10].

1) *Image registration*: In order to ensure the pixel-to-pixel correlation between the FIR and NIR frames, we applied a spatial image registration model, consisting of an affine transformation T , such as

$$T = \begin{bmatrix} s_x & 0 & t_x \\ 0 & s_y & t_y \end{bmatrix}, \quad (1)$$

where s_x and s_y specified the scaling factors and t_x and t_y the translation factors. The definition of these factors was dependent on the distance between the object of interest and the camera [10]. Therefore, we defined the transformation model during the calibration phase of the cameras and the

same factors were applied for all the recordings. The spatial alignment between FIR and NIR frames enabled the projection of the ROIs identified in the NIR frame to the FIR frame.

2) *Detection and tracking*: Five facial landmarks were detected in the NIR frame by applying a cascade convolutional neural network (CCNN) [46]. We obtained the nostril ROI as a rectangle centered over the nose landmark, with 20 pixel height and 15 pixel width. To locate the chest ROI, we used the coordinates of the chin landmark, which were derived as the bottom ordinate of the face box obtained from the CCNN model. The final chest ROI was defined using frame size dependent geometrical relations.

After successful ROI detection, an object tracker compensated for the movements of the subjects during the recording. As the chest ROI was defined by geometrical dependencies between the nostrils and face position, we tracked the nostril ROI on the NIR video only. The tracker first extracted feature points from the ROI using the minimum eigenvalue algorithm [47], then tracked these points with the Kanade-Lucas-Tomasi single points tracker [48], followed by correcting potential tracking failures with the method described by Kalal *et al.* [49]. To further mitigate potential tracking errors, we re-triggered the detection of the nostril ROI every 10 seconds.

B. Signal extraction

We extracted the RM from NIR and FIR and the TA from FIR videos. As a preprocessing step, we applied median filtering with a 3×3 mask to each NIR frame, and we enhanced

the contrast in each FIR frame by normalizing the pixel value to the interval $[0, 1]$. Following the preprocessing, we extracted the respiratory signals from each of the ROIs. The signals were then filtered and the RR was extracted from the maximal power spectral density.

1) **Thermal airflow:** The thermal fluctuations at the nostril ROI resulted in a pixel intensity variation in the FIR video. Thus, given a correctly identified ROI at time t , we computed the TA_{FIR} signal as the spatial average of the pixel intensity within the nostril ROI (Fig. 2 c) such as,

$$TA_{FIR}(t) = \frac{1}{WH} \sum_x \sum_y^H I_{ROI}(x, y, t), \quad (2)$$

where W and H are the width and height of the ROI, and $I_{ROI}(x, y, t)$ is the pixel intensity at pixel x, y at time t .

2) **Respiratory induced motion:** The periodic motion of the torso was extracted using Farnebäck dense optical flow algorithm [50]. In particular, we used the vertical velocity profiles of each tracked pixel, which were obtained as the ratio between the pixel displacement derived by the optical flow algorithm and the time between two consecutive frames. In contrast to estimating the optical flow for all the pixels of each frame [34], we restricted the optical flow calculation to the pixels within the chest ROI only. For this, we divided the ROI into a 5×7 cell grid and averaged the velocity profiles within each cell, obtaining 35 velocity profiles (Fig. 2 d). To reduce the number of the velocity profiles for further processing, we calculated the standard deviation σ_p of each profile over a window of 12 s and then excluded the profiles that did not meet the criterion, such as

$$\sigma_p < \text{median}(\Sigma_p) + 2 \cdot \text{IQR}(\Sigma_p), \quad (3)$$

where Σ_p is a vector containing the standard deviation of the 35 velocity profiles and IQR was the interquartile range. The RM signal was obtained by averaging the remaining velocity profiles. As the multispectral ROI localization enabled the identification of the chest ROI in the NIR and FIR frames, we extracted a RM signal from each spectral video, thus obtaining the RM_{NIR} and RM_{FIR} signals (Fig. 2 d).

3) **RR estimation:** The TA_{FIR} , RM_{NIR} , and RM_{FIR} signals were processed to estimate the respective RR values. For each signal, we applied the regularized least squares detrending algorithm introduced by Tarvainen *et al.* [51], with the smoothing parameter $\lambda = 300$. In the case of the TA signals, we also applied a Butterworth band pass filter, with filter order set to 2 and frequency range to 0.015 – 0.75 Hz. We did not apply any filter to the RM_{NIR} and RM_{FIR} signals to preserve their information content for the apneic detection task. The signals were then windowed and the Lomb-Scargle power spectral density (PSD) [52], [53] computed. The frequency with the highest power density in the range of 0.015 – 0.75 Hz was selected and converted to breaths/min. The Lomb-Scargle PSD estimator was chosen because cameras often have unstable frame rates and it did not require an evenly sampled time series.

C. Smart signal quality based fusion (S^2 Fusion)

The S^2 Fusion was designed as a multilevel data fusion algorithm that processes respiratory signals extracted from independent multispectral videos in order to quantify the respiratory activity. Practically, it consisted of three components: 1) a signal quality based fusion (SQb Fusion) to merge the RR estimates computed from each respiratory signal, 2) an apnea detector to extract temporal and spectral features from the respiratory signals and classify them as either an apnea or respiratory epoch, and 3) a final fusion to combine both models into a respiratory activity estimation (Fig. 2 g).

1) **Signal quality based fusion:** The SQb Fusion algorithm estimated the RR considering the signal quality from each of the three respiratory signals (TA_{FIR} , RM_{NIR} , and RM_{FIR}). The signal-to-noise ratio was computed such as

$$SNR_i = \left(\frac{\sum_{f_{peak}-\frac{k}{2}}^{f_{peak}+\frac{k}{2}} P(f)}{\sum_0^{f_{peak}-\frac{k}{2}} P(f) + \sum_{f_{peak}+\frac{k}{2}}^{f_s} P(f)} \right)_i, \quad (4)$$

where $P(f)$ was the Lomb-Scargle PSD, f_{peak} was the frequency of the peak with the highest power density, $k = 2$ breaths/min was the margin parameter [54], and f_s the sampling frequency. RR_{SQb} was obtained from the weighted median using the RR_i from each signal i and the weights w_i such as

$$w_i = \frac{SNR_i}{\sum_i SNR_i} \cdot M, \quad (5)$$

where $M = 24$ was an empirically determined scaling factor.

2) **Apnea detector:** We separated respiratory epochs from apneic epochs with an ensemble of support vector machines (SVMs, Fig. S1, Supplementary Material). In addition to the SNR used for the SQb fusion, we extracted time domain features such as the number of mean crossings, variance, and standard deviation from the 12 s windows. These features highlighted the morphological changes of the signals during the apneic events. In contrast, the frequency domain feature SNR described the periodicity of the breathing activity. All features were standardized to have zero mean and unit variance before they were fed into the SVM ensemble classifier. Test features were standardized based on training data statistics. The SVM consisted of one layer of regression models R_i and one classification model C (Fig. S1, supplementary material). Each R_i extracted the posterior probability $Pr(Y = apnea|X_i)$ for the set features X_i to belong to an apneic epoch. This way, each R_i consisted of an expert model for each respiratory signal. The purpose of C was to aggregate the Pr from each expert R_i and to compute the final output y^{apnea} . The design of the SVM ensemble was inspired by the work of Schwab *et al.* on multitasking networks [55].

IV. METHODS AND MATERIALS

We conducted a series of experiments to test whether 1) the SQb Fusion algorithm is more accurate in estimating the RR compared to a monospectral approach, 2) the combination of respiratory sources extracted from independent multispectral cameras increases the detection rate of apneic events, 3) the S^2 Fusion algorithm produces robust RR during apneic events

and can substantially improve the accuracy during recordings that include central or obstructive apneic events, and 4) there is a performance bias towards a subject's sex.

A. Experimental protocol

To evaluate the performance of our video-based respiratory monitoring algorithm, we designed an experimental protocol consisting of one task where participants were breathing spontaneously and four tasks where they followed different breathing patterns that included apneic events. Participants lay on a bed in a supine position and for each task breathed through the nose. During the tasks with apneic events, they breathed following the rhythm of a metronome. We only considered the supine body posture because it is the most common position leading to apneic events during sleep [56]. To challenge our algorithms, all videos were recorded in a dark environment (< 5 Lux).

Spontaneous Breathing Participants breathed spontaneously for 4.5 min. During this task, participants turned their head 45° left and right for 2 min each (Fig. 3 a).

Central Apnea Participants breathed for 3 min with a constant RR of 10 breaths/min and performed 3 central apneic events, each lasting 20 s (Fig. 3 b).

Obstructive Apnea Participants breathed for 3 mins with a constant RR of 10 breaths/min and performed 3 obstructive apneic events, each lasting 20 s. For each apneic event occurred, participants were instructed to simulate a blockage of the airways, and to keep the thorax moving to mimic an ongoing respiratory effort (Fig. 3 c).

Central Apnea - Blanket Participants breathed for 3 min with a constant RR of 10 breaths/min and performed 3 central apneic events, each lasting 20 s. During the entire duration of this task, participants were covered to the chin with a blanket to hide the chest contours (Fig. 3 d).

Central Apnea - Arbitrary duration Participants breathed with a constant RR of 10 breaths/min and 2 central apneic events whose duration was based on the participant's breath holding capacity (Fig. 3 e).

The study was conducted according to the ethical guidelines of Helsinki. Institutional research ethics board approval was obtained (ETH Zurich EK 2017-N-60). After informed written consent, healthy volunteers with no history of cardio-respiratory disease were enrolled in the study.

B. Experimental setup

The setup consisted of two cameras interfaced with a Raspberry Pi 3 B microcomputer (Raspberry Pi Foundation, Cambridge, UK). Custom software was developed to ensure simultaneous recording of the cameras, video compression, and data storing on the SD card. NIR videos were recorded with the See3cam_CU40 (econ-Systems, Chennai, India) with a frame rate of 15 Hz and a resolution of 336×190 pixels. FIR videos were recorded with a FLIR Lepton version 3.5 (FLIR Systems Inc., Wilsonville, United States) with an average frame rate of 8.7 Hz and a resolution of 160×120 pixels. We used a NIR LED array to enable recordings in insufficient lighting.

Fig. 3: The experimental protocol consisted of 5 different tasks: (a) spontaneous breathing for 4.5 minutes; constant breathing at 10 breaths/min interrupted by apneic events that simulated (b) central, (c) obstructive, (d) central while the body was covered with a blanket, and (e) central with arbitrary duration based on subject breath holding capacity.

Reference respiratory effort was recorded using a certified ezRIP module (Philips Respironics, Pennsylvania, USA) at a sampling rate of 50 Hz. The device used two elastic bands equipped with piezo-resistive sensors placed over the thorax and abdomen, which is one of the recommended setups for measuring respiratory effort in clinical polysomnography [43].

Collected videos and reference respiratory signals were transferred to a PC, where additional processing and analysis were performed with Matlab R2018b (MathWorks Inc., Natick, USA). The sliding window size was set to 12 s for all processing steps and was recomputed at the lowest camera sampling rate (FIR, ~ 8.7 Hz).

C. SVM ensemble training and testing

The adoption of the SVM ensemble for apnea detection implied a split of the dataset that could ensure a fair evaluation of the model. For this reason, we applied the leave-one-subject-out validation scheme, which is a particular case of the K-fold validation with K equal to the number of subjects. This validation scheme allowed us to maximize the amount of data used for training while accounting for potential high variance problems [57]. At each run, the training and test sets included the data of N-1 subjects and 1 subject respectively, with N equal to the number of recruited subjects. The training set consisted of the recordings obtained from *Central Apnea* and *Obstructive Apnea*, whereas the test set included the recordings obtained from all the tasks with apneic events. We did not include the recordings obtained from the *Spontaneous Breathing*, *Central Apnea – arbitrary duration*, and *Central Apnea – covered with blanket* tasks in the training set because they did not provide additional information regarding the signal morphology of apneic events and would have increased the class imbalance (Fig. S2, Supplementary Material).

D. Evaluation

The performance of the SQb Fusion algorithm for RR estimation was compared against algorithms that estimate the RR from single respiratory signals before fusion. Additionally, we implemented two baseline fusion algorithms based on mean and median of the RR estimates obtained from RM_{NIR} , RM_{FIR} , and TA_{FIR} signals. The RR estimates of all the evaluated algorithms were reported independently of their underlying estimation quality.

The performance of the SVM ensemble for apnea detection, which was trained with the features extracted from all signals ($RM_{NIR}+RM_{FIR}+TA_{FIR}$), was compared to the performance of three baseline SVM binary classifiers, trained with the features extracted from the individual RM_{NIR} , RM_{FIR} , and TA_{FIR} signals (Fig. S3, Supplementary Materials).

We evaluated the performance of the S^2 Fusion to quantify the respiratory activity on a combined experiment that includes both the RR estimation and apnea detection tasks. In particular, we compared it against the SQb Fusion using all the recordings that included apneic events. To guarantee an objective evaluation without overlap between training and test set in the S^2 Fusion evaluation experiment, we applied the same leave-one-subject-out validation scheme introduced in section IV-C.

E. Performance metrics

In order to describe the performance of the evaluated algorithms, we evaluated the accuracy (agreement between the RR estimates and a reference) and robustness (accuracy distribution in challenging or altering situations, specifically apneic events or differences in demographics).

To obtain accuracy, we calculated Bland-Altman plots, Pearson's correlation coefficient, and the root mean square error (RMSE). As reference RR, we used the RR estimates obtained from the thorax respiratory effort sensor. For comparison, we pooled all the estimated RRs within non-overlapping segments of 15 s and computed the median thereof. Bias and limits of agreement (LoA) for Bland-Altman analysis were calculated using the formulas for repeated measurements [58]. Pearson's correlation coefficient r was computed using a 95% confidence interval. The RMSE was computed between RR_{ref} and each fusion algorithm across segments for each subject and overall. F1 score, sensitivity, and specificity were calculated to compare the apnea classification.

To describe robustness, across-subject RMSE distributions were split by sex and displayed as boxplots. Also, we compared the across-subject RMSE distributions between the RR estimates obtained from the S^2 Fusion and the SQb Fusion algorithms to evaluate each apneic event task. Middle, bottom, and top horizontal lines of boxes depicted the median, lower, and upper quartile, respectively, and crosses depicted outliers.

We tested the normality of the RMSE distributions using the Shapiro-Wilk test and compared them with the Wilcoxon Rank Sum test. Multiple distribution comparisons were Bonferroni corrected. Significance levels were set to $p < 0.001$ (***) or $p < 0.05$ (*) unless stated otherwise.

Fig. 4: Comparison of respiratory signals from one subject during a central apnea task for (from top) reference respiratory effort from the thorax, the abdomen, TA_{FIR} , RM_{FIR} , RM_{NIR} signals, and respiratory activity (RA) estimations. SQb Fusion and S^2 Fusion algorithm were only available after the first 12 s window was collected.

V. RESULTS

We obtained videos of 30 healthy participants (17 females and 13 males, mean age: 27 ± 3 y). Subjects varied in facial appearance with a Fitzpatrick skin tone between type I and VI (I:2, II:7, III:19, IV:1, V:1, VI:0). Seven males had facial hair. A total of 492 minutes of video were recorded and available for analysis. An illustrative example of the signals obtained from the reference contact sensors and cameras, as well as estimation output, is depicted in Fig. 4. Motion artifacts substantially corrupted the recording of one subject.

A. Respiratory rate estimation performance

A total of 135 minutes (540 segments) of *Spontaneous Breathing* task were evaluated for the RR estimation accuracy. The SQb fusion showed the highest accuracy (Fig. 5). This was supported by the highest correlation ($r=0.92$), lowest RMSE (1.6 breaths/min), and the best agreement from the Bland-Altman analysis. The bias was 0.47 breaths/min and the LoA were 3.46 and -2.51 breaths/min, outperforming all

Fig. 5: Bland-Altman (a-d) and scatter plots (e-h) of the estimated RR using (a, e) the Signal Quality based (SQb) fusion and the individual respiratory signals (b, f) RM_{NIR} , (c, g) RM_{FIR} , and (d, h) TA_{FIR} on the Spontaneous Breathing data divided into 15 s long segments. The dashed blue lines depict the limits of agreement, the solid red line the bias. The dotted yellow line depicts the linear relation that best fits the data. r is the Pearson's correlation coefficient and RMSE the root mean square error across all segments.

TABLE I: Apnea detection performance for all apnea tasks on the test set. Median (mean absolute deviation) for F1 score, sensitivity, and specificity were obtained for the SVM ensemble trained with all the features, and features from the individual respiratory signals RM_{NIR} , RM_{FIR} and TA_{FIR} .

Obstructive Apnea			
Source	F1	Sensitivity	Specificity
$RM_{NIR}+RM_{FIR}+TA_{FIR}$	0.86 (0.12)	0.91 (0.16)	0.78 (0.06)
RM_{NIR}	0.76 (0.12)	0.75 (0.17)	0.82 (0.05)
RM_{FIR}	0.76 (0.15)	0.80 (0.20)	0.69 (0.05)
TA_{FIR}	0.75 (0.14)	0.74 (0.17)	0.72 (0.08)
Central Apnea			
$RM_{NIR}+RM_{FIR}+TA_{FIR}$	0.93 (0.02)	0.98 (0.03)	0.86 (0.04)
RM_{NIR}	0.93 (0.08)	0.95 (0.13)	0.92 (0.06)
RM_{FIR}	0.88 (0.04)	0.98 (0.06)	0.81 (0.06)
TA_{FIR}	0.75 (0.10)	0.77 (0.13)	0.70 (0.07)
Central Apnea - Arbitrary duration			
$RM_{NIR}+RM_{FIR}+TA_{FIR}$	0.87 (0.07)	0.96 (0.10)	0.82 (0.08)
RM_{NIR}	0.89 (0.08)	0.91 (0.09)	0.87 (0.08)
RM_{FIR}	0.87 (0.08)	0.89 (0.11)	0.86 (0.07)
TA_{FIR}	0.72 (0.12)	0.69 (0.14)	0.78 (0.08)
Central Apnea - Blanket			
$RM_{NIR}+RM_{FIR}+TA_{FIR}$	0.92 (0.03)	1.00 (0.09)	0.86 (0.05)
RM_{NIR}	0.94 (0.04)	0.98 (0.10)	0.92 (0.07)
RM_{FIR}	0.90 (0.05)	0.99 (0.11)	0.81 (0.08)
TA_{FIR}	0.74 (0.15)	0.79 (0.19)	0.72 (0.07)

the algorithms based on individual respiratory signals from single video modalities (Fig. 5 e-h). The SQb Fusion obtained a median RMSE of 1.17 breaths/min across subjects, which was significantly lower than the baselines (median fusion 1.60 breaths/min, $p < 0.05$ and mean fusion 1.96 breaths/min, $p < 0.01$, Table S1, Fig. S4, Supplementary Materials). The RMSE of TA_{FIR} showed a significantly different distribution between female and male subjects (median RMSE 2.65 vs. 1.58 breaths/min, Fig. 6).

B. Apnea detection

A total of 357 minutes of video recordings containing apnea were evaluated for apnea and respiratory activity detection. Fusing the features extracted from multispectral videos led to the best performance in detecting apneic events in the *Obstructive Apnea* task (F1 score = 0.86, Table I). The performance of the SVM ensemble trained with all the features or using only those extracted from the RM_{NIR} signal showed similar results on the *Central Apnea* task (F1 score = 0.93 for both sources).

C. Respiratory activity detection

The S^2 Fusion algorithm showed a statistically significant reduction of RMSE compared to the SQb Fusion in all apnea tasks (Fig. 7). The RMSE over all tasks was 4.44 breaths/min for S^2 Fusion and 7.96 for SQb Fusion.

Fig. 6: Boxplot comparison of the root mean square error (RMSE) from the Spontaneous Breathing data between subjects of different sex for the signal quality based (SQb) Fusion algorithm and the algorithms based on the individual respiratory signals.

Fig. 7: Comparison of the root mean square error (RMSE) for the S²Fusion and SQb Fusion algorithms by apneic event type: (a) Central Apnea, (b) Central Apnea – Arbitrary duration, (c) Obstructive Apnea, (d) Central Apnea – Blanket.

VI. DISCUSSION

We presented a novel fusion approach for respiratory monitoring from multispectral videos. Our algorithm features two dedicated data fusion models that combine multichannel respiratory information to estimate the RR, and also to detect the absence of respiratory activity. For the first time, these two components are combined in a single system together to be studied extensively. Our evaluation of videos from 30 subjects performing diverse breathing patterns showed that multisensory respiratory information combined with apneic events detection estimates RR with higher accuracy than applying single camera-based approaches, therefore demonstrating robustness during apnea events. The accuracy of the fusion approach did not show a statistical difference between male and female subjects. Moreover, in contrast to several state-of-the-art methods developed for high-end thermal cameras, our multispectral video fusion pipeline enabled low-end

mobile thermal cameras to satisfactorily perform the task of respiratory monitoring.

A. Respiratory rate estimation

The SQb Fusion combined the RR estimates obtained from the RM and TA signals. The key element of the proposed fusion approach was the use of the SNR of each respiratory signal within the fusion model. Redundancy was important for improving the estimation of the RR, illustrated by a significant improvement of the RMSE using fusion when compared to the results from a single spectral camera.

We demonstrated that adopting multispectral sources provides an additional gain in accuracy. This expands the work of Pereira *et al.*, who combined multiple respiratory signals from a single high-end video source to improve RR estimation [44]. In combination, this enables new applications where robustness against noise and movement artifacts is essential. For example, head or torso movements can cause the temporary loss of one of the regions of interest and affect the quality of the extracted respiratory signal leading to an inaccurate estimation of the RR. We included such challenging situations into our experimental data by asking the participants in the *Spontaneous Breathing* protocol to rotate their heads. Our analysis clearly showed that the SQb Fusion discarded the deteriorated, motion artifact affected estimations and emphasized those with higher SNR. However, when the movement involved both head and torso regions, as involuntarily occurred in one subject, the respiratory signals from all sources were deteriorated, resulting in a RMSE higher than 3 breaths/min. To cope with motion artifacts, dedicated computer vision methods could be further developed to quantify the motion of the whole body and use these insights as a quality assessment of the respiratory signals (e.g. by directly eliminating fusion estimates [12]).

We analyzed the performance of the investigated RR algorithms for potential differences linked to the subjects' sex (Fig. 6). We found that estimating the RR using only the TA signals lead to a significantly higher RMSE for female subjects. A deeper analysis revealed that the higher estimation error on the female group was caused by an incorrect ROI localization, likely related to long hair. No significant difference was observed for the performance of the RR algorithms based on motion signals from both NIR and FIR videos. Most importantly, no significant difference was observed in the SQb Fusion, suggesting that the fusion algorithm was reducing the bias seen in the TA signal effectively and robust to sex differences. The small size of each group (17 female and 13 male) does not allow a conclusive statistical evaluation on all algorithm components or demographic sub-features (such as facial hair or skin colour), but we believe that computer vision algorithms for physiological monitoring should systematically be tested for potential sex and other biases.

B. Apnea detection

Our apnea detector utilized a SVM ensemble classifier with a set of features extracted from multiple respiratory signals as input. This enabled the detection of apneic events

independently from the RR estimation and without requiring any pre-calibration steps. We demonstrated that detecting apneic events by fusing multiple respiratory signals extracted from multispectral independent videos provided higher F1 scores than adopting single spectral cameras. The importance of fusing multiple respiratory signals was underlined by the higher performance in detecting obstructive apneic events. During obstructive apneic events when the occlusion of the upper airways blocks air exchange, the respiratory movements of the torso were still present and could lead to a misdetection. However, the combination of the features extracted from the RM_{NIR} , RM_{FIR} , and TA_{FIR} signals enabled the SVM model to accurately discriminate breathing from obstructive apneic events. Fusing the features from all the respiratory signals showed a marginal impact on the central apnea detection tasks. Interestingly, the RM_{NIR} signal showed similar results to the fusion of all the features which was in contrast to its low performance in the RR estimation task (Fig. 5 b, f). A possible explanation could be that the two tasks had different requirements for signal quality. While the RR estimation task depended on high quality signals, the detection of apnea may only require processed features that were also available from low quality RM_{NIR} signals.

Our research and algorithm design focused on nighttime applications with monitoring of respiration in bed. This real-world scenario could easily be challenged by a body partially covered by a blanket. In Li *et al.* RM is reliably extracted from NIR videos without being negatively influenced by the presence of a blanket texture [34]. Our experiments confirmed these findings also for the S²Fusion and expanded the application to the task of apnea detection. In fact, we demonstrated that motion tracking applied to both NIR and FIR videos is a valid methodology to detect apneic events, even if one channel is unreliable.

C. Respiratory monitoring system

To the best of our knowledge, the S²Fusion is the first published fusion model that addressed the task of estimating the RR and detecting apnea in a single system. By combining the RR estimates obtained by the SQb Fusion and the apnea detection score obtained by the apnea detector, we forced the S²Fusion to account for abnormal breathing events, therefore increasing the robustness of the RR estimation from video recordings during apneic events. The modular architecture of the fusion could enable a more granular detection of apnea events, such as the distinction between the type of apnea with a more sophisticated, non-binary classifier.

While the presented work demonstrated the feasibility of the approach and is a significant step forward for the design of reliable non-contact respiratory monitors, it is also evident that further validation in clinical settings with patients is needed. While we have put great effort into diversifying data collection with a comparatively large number of subjects and a broadly defined protocol including spontaneous breathing, various apneic events and confounding factors such as blanket coverage and head movements, the dataset is not fully representative of future applications. First, our subjects were healthy

and breathing in the expected range of 6–22 breaths/min. Therefore, we did not investigate the performance of our algorithm for higher RR ranges as would occur during disease or in a pediatric population. We believe that guiding healthy subjects to breathe with a RR in the range of 25–35 breaths/min can result in accentuated respiratory signs and could make the validation of any video-based algorithm for RR estimation highly biased. Furthermore, our recordings only contained simulated apneic events. The respiratory motion signals obtained by the reference measurement systems visually showed a lot of similarity to those of real patients suffering from central and obstructive sleep apnea. Nevertheless, full overnight recordings from healthy subjects and patients suffering from sleep apnea are needed to fully assess the value of non-contact respiratory monitoring for use as a telemedicine apnea assessment tool. There is significant potential for a system like ours to be implemented in overnight polysomnography screenings where cameras are already part of the standard equipment, but play a minor role in the overall assessment.

VII. CONCLUSION

We presented a unique approach to monitor respiratory activity based on the fusion of multispectral videos. It explored the idea of extracting multiple respiratory signals from independent multispectral videos and strategically combining them to address RR estimation and apnea detection. Our findings demonstrated that fusing multiple respiratory signals from multispectral cameras increases the accuracy of respiratory activity when compared to single camera modalities. In addition, our S²Fusion further highlighted the advantage of addressing the tasks of estimating RR and detecting apnea independently. We experimentally demonstrated that low-cost multispectral mobile cameras could potentially be used to monitor respiration. We, therefore, are confident that our results constitute a step towards the implementation of camera sensors to unobtrusively and robustly monitor respiratory activity both in the clinic and in remote locations.

VIII. ACKNOWLEDGMENTS

We thank T. Meyer and R. Moser for contributing to the technical development, C. Maschio for assisting with the data collection, and J. Lim for helping to revise this manuscript. We thank the reviewers for their feedback on the manuscript. We are very grateful to all subjects for participating in our experiments.

REFERENCES

- [1] M. A. Cretikos *et al.*, “Respiratory rate: The neglected vital sign,” *Med. J. Aust.*, vol. 188, no. 11, pp. 657–9, 2008.
- [2] T. Flenady *et al.*, “Accurate respiratory rates count: So should you!” *Australas. Emerg. Nurs. J.*, vol. 20, no. 1, pp. 45–7, 2017.
- [3] M. Elliott, “Why is Respiratory Rate the Neglected Vital Sign? A Narrative Review,” *Int Arch Nurs Heal. Care*, vol. 2, no. 3, pp. 2–5, 2016.
- [4] L. Tarassenko *et al.*, “Non-contact video-based vital sign monitoring using ambient light and auto-regressive models,” *Physiol. Meas.*, vol. 35, no. 5, pp. 807–31, 2014.
- [5] M. van Gastel *et al.*, “Robust respiration detection from remote photoplethysmography,” *Biomed. Opt. Express*, vol. 7, no. 12, pp. 4941–57, 2016.

- [6] C. Hegde *et al.*, "Autotriage - an open source edge computing raspberry pi-based clinical screening system," *medRxiv*, 2020, medRxiv: 10.1101/2020.04.09.20059840.
- [7] M. M. Baharestani, "An overview of neonatal and pediatric wound care knowledge and considerations," *Ostomy. Wound Manag.*, vol. 53, no. 6, pp. 34–6, 2007.
- [8] C. Muroi *et al.*, "Automated false alarm reduction in a real-life intensive care setting using motion detection," *Neurocrit. Care*, p. 419–26, 2019.
- [9] G. Scebbba *et al.*, "Improving ROI detection in photoplethysmographic imaging with thermal cameras," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* IEEE, 2017, pp. 4285–8.
- [10] G. Scebbba *et al.*, "Multispectral camera fusion increases robustness of ROI detection for biosignal estimation with nearables in real-world scenarios," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* IEEE, 2018, pp. 5672–5.
- [11] P. H. Charlton *et al.*, "Breathing rate estimation from the electrocardiogram and photoplethysmogram: A review," *IEEE Rev. Biomed. Eng.*, vol. 11, pp. 2–20, 2017.
- [12] W. Karlen *et al.*, "Multiparameter respiratory rate estimation from the photoplethysmogram," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 7, pp. 1946–53, 2013.
- [13] M. D. Peláez-Coca *et al.*, "Cross time-frequency analysis for combining information of several sources: application to estimation of spontaneous respiratory rate from photoplethysmography," *Comput. Math. Methods Med.*, vol. 2013, 2013.
- [14] J. Lázaro *et al.*, "Deriving respiration from photoplethysmographic pulse width," *Med. Biol. Eng. Comput.*, vol. 51, no. 1-2, pp. 233–42, 2013.
- [15] J. Lázaro *et al.*, "Electrocardiogram derived respiratory rate from QRS slopes and R-wave angle," *Ann. Biomed. Eng.*, vol. 42, no. 10, pp. 2072–83, 2014.
- [16] G. D. Clifford *et al.*, "Data fusion for improved respiration rate estimation," *EURASIP J. Adv. Signal Process.*, vol. 2010, 2010.
- [17] C. Orphanidou *et al.*, "Data fusion for estimating respiratory rate from a single-lead ECG," *Biomed. Signal Process. Control*, vol. 8, no. 1, pp. 98–105, 2013.
- [18] J. Lee *et al.*, "Respiratory rate extraction from pulse oximeter and electrocardiographic recordings," *Physiol. Meas.*, vol. 32, no. 11, pp. 1763–73, 2011.
- [19] C. Orphanidou, "Derivation of respiration rate from ambulatory ECG and PPG using Ensemble Empirical Mode Decomposition: Comparison and fusion," *Comput. Biol. Med.*, vol. 81, pp. 45–54, 2017.
- [20] X. R. Ding *et al.*, "A pulse transit time based fusion method for the noninvasive and continuous monitoring of respiratory rate," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* IEEE, 2016, pp. 4240–3.
- [21] C. L. Shen *et al.*, "Respiratory Rate Estimation by Using ECG, Impedance, and Motion Sensing in Smart Clothing," *J. Med. Biol. Eng.*, vol. 37, no. 6, pp. 826–42, 2017.
- [22] N. N. Lepine *et al.*, "Robust respiration rate estimation using adaptive Kalman filtering with textile ECG sensor and accelerometer," in *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* IEEE, 2016, pp. 3797–800.
- [23] L. Poupard *et al.*, "Novel mathematical processing method of nocturnal oximetry for screening patients with suspected sleep apnoea syndrome," *Sleep Breath.*, vol. 16, no. 2, pp. 419–25, 2012.
- [24] G. Memis and M. Sert, "Multimodal Classification of Obstructive Sleep Apnea Using Feature Level Fusion," in *Proc. - IEEE 11th Int. Conf. Semant. Comput. ICSC 2017*. Institute of Electrical and Electronics Engineers Inc., 2017, pp. 85–8.
- [25] D. Alvarez *et al.*, "Spectral analysis of electroencephalogram and oximetric signals in obstructive sleep apnea diagnosis," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* IEEE, 2009, pp. 400–3.
- [26] H. M. Al-Angari and A. V. Sahakian, "Automated recognition of obstructive sleep apnea syndrome using support vector machine classifier," *IEEE Trans. Inf. Technol. Biomed.*, vol. 16, no. 3, pp. 463–8, 2012.
- [27] D. Sommermeyer *et al.*, "Detection of sleep disordered breathing and its central/obstructive character using nasal cannula and finger pulse oximeter," *J. Clin. Sleep Med.*, vol. 8, no. 5, pp. 527–33, 2012.
- [28] F. Espinoza-Cuadros *et al.*, "Speech Signal and Facial Image Processing for Obstructive Sleep Apnea Assessment," *Comput. Math. Methods Med.*, vol. 2015, 2015.
- [29] A. C. Guyton and J. E. Hall, *Textbook of Medical Physiology*, 2006.
- [30] V. B. Zordan *et al.*, "Breathe easy: Model and control of human respiration for computer animation," *Graph. Models*, vol. 68, no. 2, pp. 113–32, 2006.
- [31] L. L. Patrick Lévy, Malcolm Kohler, Walter T. McNicholas, Ferran Barbé, R. Doug McEvoy, Virend K. Somers and J.-L. Pépin, "Obstructive sleep apnoea syndrome," *Sleep Med. Rev.*, vol. 1, pp. 1–21, 2015.
- [32] M. Bartula *et al.*, "Camera-based System for Contactless Monitoring of Respiration," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2013, pp. 2672–5.
- [33] C.-W. Wang *et al.*, "Unconstrained video monitoring of breathing behavior and application to diagnosis of sleep apnea," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 2, pp. 396–404, 2013.
- [34] M. H. Li *et al.*, "Non-Contact Vision-Based Cardiopulmonary Monitoring in Different Sleeping Positions," *IEEE J. Biomed. Health Inform.*, vol. 21, no. 5, pp. 1367–75, 2017.
- [35] R. Janssen *et al.*, "Video-based respiration monitoring with automatic region of interest detection," *Physiol. Meas.*, vol. 37, no. 1, pp. 100–14, 2016.
- [36] Z. Zhu *et al.*, "Tracking human breath in infrared imaging," in *Int. Conf. Bioinform. Biomed. Eng.* IEEE, 2005, pp. 227–31.
- [37] A. K. Abbas *et al.*, "Neonatal non-contact respiratory monitoring based on real-time infrared thermography," vol. 10, no. 1, p. 93, 2011.
- [38] C. B. Pereira *et al.*, "Remote monitoring of breathing dynamics using infrared thermography," *Biomed. Opt. Express*, vol. 6, no. 11, pp. 4378–94, 2015.
- [39] Y. Cho *et al.*, "Robust tracking of respiratory rate in high- dynamic range scenes using mobile thermal imaging," *Biomed. Opt. Express*, vol. 8, no. 10, pp. 4480–503, 2017.
- [40] M.-Z. Poh *et al.*, "Advancements in noncontact, multiparameter physiological measurements using a webcam," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 1, pp. 7–11, 2011.
- [41] J. Jorge *et al.*, "Data fusion for improved camera-based detection of respiration in neonates," in *Opt. Diagnostics Sens. XVIII Towar. Point-of-Care Diagnostics*, vol. 10501. SPIE, 2018.
- [42] F. Bousefsaf *et al.*, "Continuous wavelet filtering on webcam photoplethysmographic signals to remotely assess the instantaneous heart rate," *Biomed. Signal Process. Control*, vol. 8, no. 6, pp. 568–74, 2013.
- [43] R. B. Berry *et al.*, "The AASM Manual for the Scoring of Sleep and Associated Events," *Am. Acad. Sleep Med.*, vol. 53, no. 9, pp. 1689–99, 2013.
- [44] C. B. Pereira *et al.*, "Noncontact monitoring of respiratory rate in newborn infants using thermal imaging," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 4, pp. 1105–14, 2018.
- [45] C. B. Pereira *et al.*, "Multisensor data fusion for enhanced respiratory rate estimation in thermal videos," in *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2016, pp. 1381–4.
- [46] K. Zhang *et al.*, "Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–503, 2016.
- [47] Jianbo Shi and Tomasi, "Good features to track," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. CVPR-94*. IEEE Comput. Soc. Press, 1994, pp. 593–600.
- [48] C. Tomasi *et al.*, "Detection and Tracking of Point Features," *Int. J. Comput. Vis.*, vol. 9, 1991.
- [49] Z. Kalal *et al.*, "Forward-Backward Error: Automatic Detection of Tracking Failures," in *2010 20th Int. Conf. Pattern Recognit.* IEEE, 2010, pp. 2756–9.
- [50] G. Farneback, "Two-frame motion estimation based on polynomial expansion," in *Lect. Notes Comput. Sci.* Springer, 2003, pp. 363–70.
- [51] M. P. Tarvainen *et al.*, "An advanced detrending method with application to hrv analysis," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 2, pp. 172–75, 2002.
- [52] N. R. Lomb, "Least-squares frequency analysis of unequally spaced data," *Astrophys. Space Sci.*, vol. 39, no. 2, pp. 447–62, 1976.
- [53] J. D. Scargle, "Studies in astronomical time series analysis. ii-statistical aspects of spectral analysis of unevenly spaced data," *Astrophys. J.*, vol. 263, pp. 835–53, 1982.
- [54] G. De Haan and V. Jeanne, "Robust pulse rate from chrominance-based rPPG," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 10, pp. 2878–86, 2013.
- [55] P. Schwab *et al.*, "Not to Cry Wolf: Distantly Supervised Multitask Learning in Critical Care," in *Proc. 35th Int. Conf. Mach. Learn. 2018*, no. 80, 2018, pp. 4518–27.
- [56] D. P. White, "Pathogenesis of obstructive and central sleep apnea," *Am. J. Respir. Crit. Care Med.*, vol. 172, no. 11, pp. 1363–70, 2005.
- [57] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *IJCAI*, vol. 14, no. 2. Montreal, Canada, 1995, pp. 1137–45.
- [58] J. M. Bland and D. G. Altman, "Measuring agreement in method comparison studies," *Stat. Methods Med. Res.*, vol. 8, no. 2, pp. 135–60, 1999.