



OPEN

Robust in-vehicle respiratory rate detection using multimodal signal fusion

Joana M. Warnecke^{1,2}✉, Joan Lasenby² & Thomas M. Deserno¹

Continuous health monitoring in private spaces such as the car is not yet fully exploited to detect diseases in an early stage. Therefore, we develop a redundant health monitoring sensor system and signal fusion approaches to determine the respiratory rate during driving. To recognise the breathing movements, we use a piezoelectric sensor, two accelerometers attached to the seat and the seat belt, and a camera behind the windscreen. We record data from 15 subjects during three driving scenarios (15 min each) city, highway, and countryside. An additional chest belt provides the ground truth. We compare the four convolutional neural network (CNN)-based fusion approaches: early, sensor-based late, signal-based late, and hybrid fusion. We evaluate the performance of fusing for all four signals to determine the portion of driving time and the signal combination. The hybrid algorithm fusing all four signals is most effective in detecting respiratory rates in the city ($P = 62.42$), highway ($P = 62.67$), and countryside ($P = 60.94$). In summary, 60% of the total driving time can be used to measure the respiratory rate. The number of signals used in the multi-signal fusion improves reliability and enables continuous health monitoring in a driving vehicle.

The respiratory rate is one out of four primary vital signs to derive the health status of individuals^{1,2}. It is significant to predict adverse cardiac events³, emotional stress⁴, and cognitive load⁵. Nicolo et al. defined thirteen goals for respiratory monitoring¹, ranging from the presence of breathing to cardiac events as well as environmental stress¹. Moreover, the respiratory rate provides information for airways and other lung structures, which are affected by chronic respiratory diseases (CRDs)². Asthma, chronic obstructive pulmonary disease (COPD), occupational lung disease, and pulmonary hypertension are some of the most prevalent⁶. According to the World Health Organisation (WHO), more than 262 million people suffer from asthma⁶. In most cases, asthma can be treated with inhaled medications⁷. However, COPD is mostly irreversible and particularly affects adults⁸. It is characterized by breathlessness, sputum production, and chronic coughing⁶. Furthermore, COPD patients develop more severe COVID-19 symptoms⁹. Continuous health monitoring and regular medical check-ups yield early detection of symptoms, improve therapeutic outcomes for CRDs, and prevent cardiac events¹⁰. In Western countries, people spend, on average, 35 min per day in a car¹¹. Therefore, unobtrusive in-vehicle sensors can potentially integrate continuous health monitoring into our daily lives¹².

Currently, in-vehicle monitoring focuses on the tiredness tracking of the driver and uses eye- and face-tracking or steering wheel, pedal, and lane movements^{13,14}. More recently, state-of-the-art research also focuses on vital signs recorded in automotive environments, in particular the respiratory rate^{15,16}. For instance, Ju et al. integrated a piezoelectric sensor in the seat belt to derive the drivers' stress level under laboratory conditions¹⁷. Baeck et al. applied similar sensors while driving¹⁸. Vavrinsky et al. used ballistocardiography (BCG) in a bucket seat¹⁹. Vinci et al. suggested radar sensors in the seat backrest at the height of the lunge to measure the respiratory rate²⁰. Guo et al.²¹ used a near-infrared camera for respiratory rate detection. However, current research is limited mostly to only one particular sensor concept and aims at supporting driving assistance systems rather than monitoring the long-term health status of vehicle occupants¹⁵. In previous work, we used an accelerometer attached to the seat belt under different driving conditions to determine the best position²² and a de-noising method for respiratory rate monitoring²³.

In this work, we present a redundant sensor system and signal fusion approaches for in-vehicle respiratory rate detection to answer the following research question: (1) *Which portion of the driving time is utilisable to measure the respiratory rate robustly?*, (2) *Which sensor combination delivers the most reliable results?*, and (3) *What is the performance difference for respiratory rate detection between the driving scenarios?*

¹Peter L. Reichertz Institute for Medical Informatics of TU Braunschweig and Hannover Medical School, 38106 Braunschweig, Germany. ²Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK. ✉email: joana.warnecke@plri.de

Material

Sensor system

Based on our previous review¹⁶ and the performance assessment by Leonhardt et al.¹⁵, we choose two accelerometer sensors (BiosignalPlux Explorer, Plux Wireless Biosignals, Lisboa, Portugal), a piezoelectric sensor (BiosignalPlux Explorer, Plux Wireless Biosignals, Lisboa, Portugal), and an RGB camera (Raspberry Pi Foundation, Cambridge, United Kingdom). The camera has a resolution of 1280×720 pixels and records with ten frames per second (Fig. 1a). The channel hub (BiosignalPlux Explorer, Plux Wireless Biosignals, Lisboa, Portugal) connects the accelerometer, piezoelectric sensor, and chest belt and sends the recordings via Bluetooth, and the camera has a wired connection to the Raspberry Pi (Fig. 1b).

For ground truth measurements, we use a chest belt (BiosignalPlux Explorer, Plux Wireless Biosignals, Lisboa, Portugal) (Fig. 1c). We place the belt around the upper thorax of the subject and measure the peaks of the movement. According to our previous work²³, we attach the accelerometers to the seat belt and the seat to measure noisy breathing movements and noise only, respectively (Fig. 1a). We integrate the piezoelectric sensor into the seat belt and attach the camera behind the steering wheel (Fig. 1c). We install all sensors into a street-legal vehicle (VW Tiguan, Volkswagen AG, Wolfsburg, Germany). To determine the respiratory rate optically, we average the green channel for specific regions of interest (ROI): (1) belt and (2) chest.

Experimental setup

We selected 15 test persons with different ages (20–67 years), genders (female: 6, male: 9), and body mass index (21–44). For each subject, we record 15 min of driving: (1) in the *city*, (2) on the *highway*, and (3) in the *countryside*. Following the navigation system, all subjects drove the same predetermined route.

Existing studies often show the technical feasibility in a simulated environment with a low number of subjects^{17,19,20}. Guo et al. conducted one experiment with five subjects under real driving conditions with a near-infrared camera²¹. The advantage of our study is the recording of various sensors for respiratory rate detection, three different driving scenarios, a higher number of subjects, and a publicly available data set.

Driving scenarios, such as highway, city, and countryside, are different due to a variety of factors, including road types, traffic conditions, speed limits, environmental variables, and the types of challenges. Therefore, we recorded data in three scenarios. Different driving scenarios feature distinct road conditions. Highways typically have well-defined lanes and relatively stable traffic patterns, while city roads often involve frequent lane changes, pedestrians, and complex traffic dynamics. The countryside has more uneven roads because the road conditions are often not as good as on the highway or in the city center. Collecting data in all scenarios enables a comparison between these different scenarios to analyse the data recorded with various road conditions. Each scenario has various traffic patterns. The highway typically has high-speed limits, while city driving involves frequent stops and starts and congestion. Countryside includes fewer vehicles and lower speed limits but can pose challenges like winding roads and variable road quality. The speed limit for the highway was 130 km/h, for the city 50 km/h, and for the countryside between 50 and 100 km/h.

Ethics approval

We record all data anonymously following the Helsinki Declaration. The ethics committee at TU Braunschweig (Braunschweig, Germany) approved the study's procedures (internal process number: D_2022-13). Informed consent was obtained from all subjects. Specific consent was obtained for identifying images in an online open-access publication.

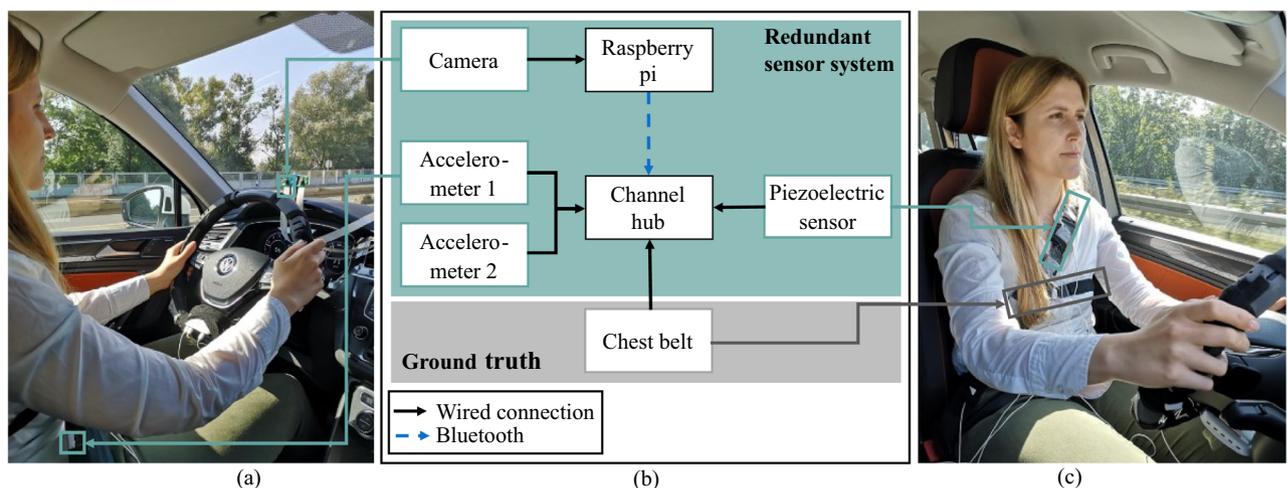


Figure 1. Sensor system. (a) The camera is behind the steering wheel, and the accelerometers are in the seat belt as well as on the right side of the driver's seat. (b) Schematic sensor system. (c) The piezoelectric sensor sits in the seat belt, and the chest belt for reference measurements is strapped around the subject's chest.

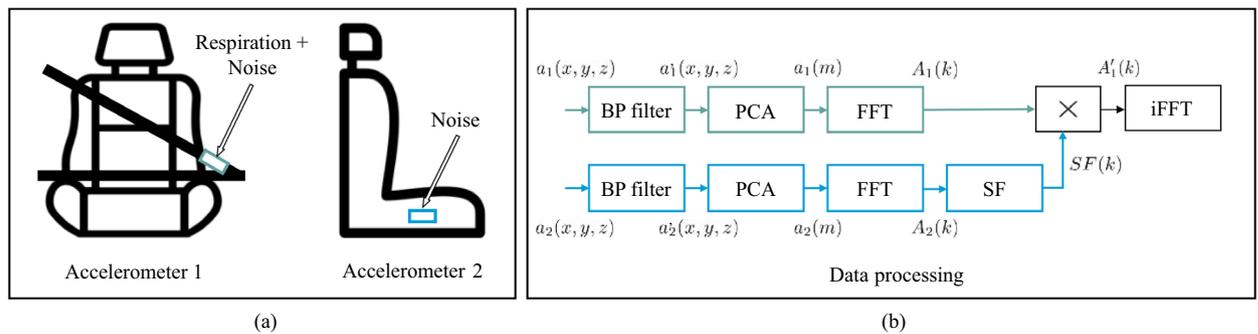


Figure 2. Pre-processing of the accelerometer signals. (a) Sensor positions. (b) Processing pipeline.

Signal pre-processing Piezoelectrical sensor

The piezoelectrical sensor determines the pressure that is generated when the breathing chest expands against the seat belt. As the sensor directly turns the pressure into an electrical signal (voltage), we do not apply additional preprocessing but use the raw data directly.

Accelerometers

For noise reduction, we attached *accelerometer 1* to the seat belt at the position side-waist, which measures the respiration and the noise (Fig. 2a). We selected the position side waist based on a previous publication²², which evaluated the positions of the shoulder, chest, side waist, and waist for respiratory rate detection with an accelerometer. The *accelerometer 2* is attached to the seat of the driver on the right side and measures the environmental noise (Fig. 2a). The required pre-processing of the signals uses bandpass (BP) filtering and principal component analysis (PCA). The noise canceling is then computed in the Fourier domain, using the fast Fourier transform (FFT) (Fig. 2b).

The magnitude in X_1 of the *accelerometer 1* includes the frequency distribution of respiration plus noise and X_2 just the noise. In X_2 , a higher magnitude of the frequency component implies higher noise. Specifically, $|X_2(k)|$ denotes the k th amplitude in an FFT series of *accelerometer 2*. Using Eq. (1), we calculated the suppression factor (SF) that suppresses the frequency components in X_1 based on X_2 magnitudes and $\mu(|X_1|)$ is the mean value of the X_1 amplitudes²³:

$$SF(k) = e^{-\frac{|X_2(k)|}{\mu(|X_1|)}} \quad (1)$$

Using the suppression factor (SF) in Eq. (2), we calculated the frequency distribution of the suppressed signal (supp):

$$X_{supp}(k) = SF(k) \cdot X_1(k) \quad (2)$$

The output is a de-noised signal for sensor $A'_1(k)$. In contrast to the previous publication, we applied the inverse FFT (iFFT) because the signal fusion approach needs the signal in the time domain as an input²³. For further details, we refer to our previous work²³.

Video

We attached a black-and-white chessboard pattern on the seat belt to increase the contrast between the clothes and the seat belt. This makes our vision system unresponsive to the color of the clothing the subject is wearing. Using the static setting in the car, we extract two ROIs covering the belt ($video_{Belt}$) and chest ($video_{Chest}$) (Fig. 3). The ROIs are a rectangle and determined by the static positions of the camera as well as the driver's seat: $video_{Belt}$ (x: 701, y: 550, w: 10, h: 190) (Fig. 3a) and $video_{Chest}$ (x: 401, y: 550, w: 400, h: 190) (Fig. 3b).

This means that the ROI belt has a width of 10 pixels and a height of 190 pixels at positions x: 701 and y: 550. To ensure the accuracy of ROI extraction, we verify their positions during the pre-processing stage and make necessary position adjustments. Using Eq. (3), The average (avg) green colour values of each cut-out pixel in the ROI are calculated by adding up all the color values and dividing by the total number of values:

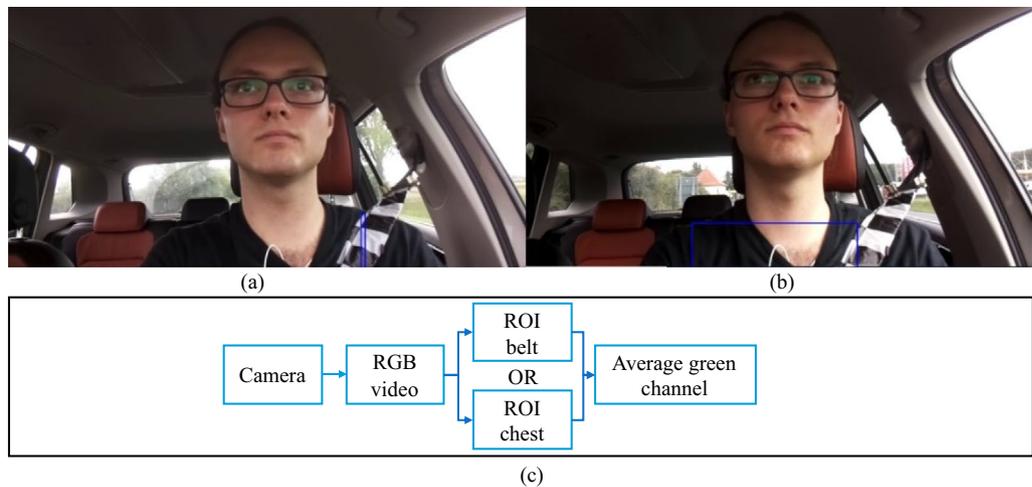


Figure 3. Video signal. (a) ROI belt. (b) ROI chest. (c) Data processing.

$$\text{color}_{\text{avg}} = \frac{1}{w \cdot h} \sum_{w=1}^n \sum_{h=1}^m \text{Pixel}_{w,h} \quad (3)$$

The average colour values of the ROIs change with every in- and exhaling movement because the position of the belt changes. Therefore, the respiratory rate can be calculated based on the movement of the belt.

Implementation for statistical analysis

To derive the signal from the ROI of $video_{\text{Belt}}$ and $video_{\text{Chest}}$, we use the library OpenCV (version 4.5). We implement the CNN-based approach in Python (version 3.8.5) using the libraries TensorFlow (version 2.3.1) and Keras (version 2.4.3). The high-performance computer cluster *Phoenix* calculates all training and testing²⁴. For the evaluation, we use the library Numpy (version 1.19.0).

Input data

In total, our sensors deliver four input channels and one reference signal as ground truth. The recordings differ in signal-to-noise ratio (SNR). Figure 4 gives examples of lower (Fig. 4a) and higher SNRs (Fig. 4b), respectively. The arbitrary unit (au) represents the unit for the accelerometer and the video signal.

According to Chandra et al.²⁵, we apply additional pre-processing for signal fusion:

1. Upsampling all data to a unified sampling rate of 200 Hz,
2. Median filtering, and 3. Amplitude normalization to the interval $[-1, 1]$.

We split the signal into snippets of 201 overlapping samples. Following Chandra et al.²⁵, test and training snippets overlap by 200 and 190 data points, respectively. We use the leave-one-subject-out cross-validation scheme: we train the CNN on 14 subjects and use the data of the remaining subjects for testing. We repeat this procedure 15 times and average the results.

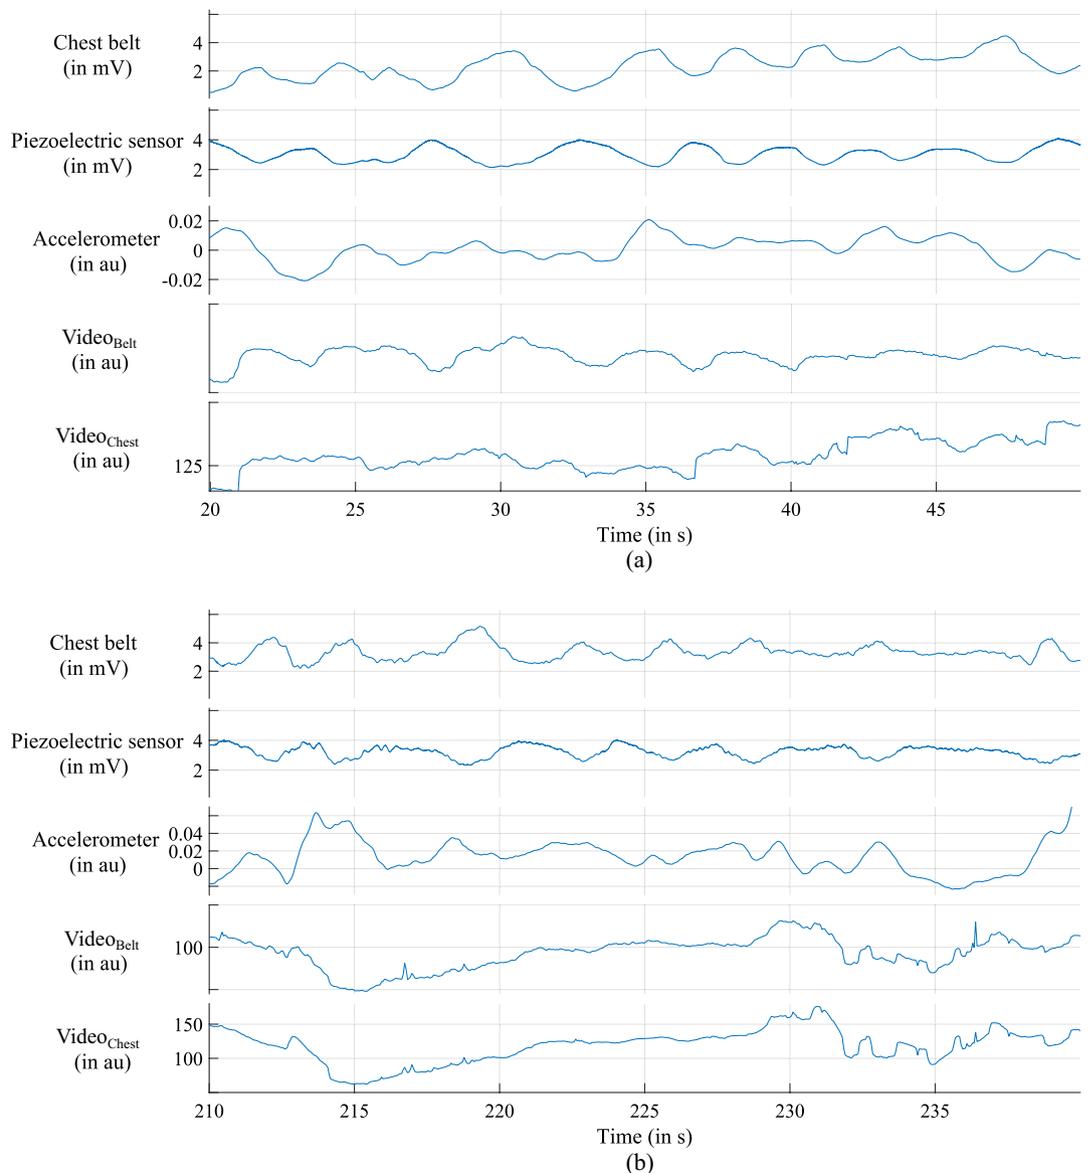


Figure 4. Input data from subject 5678 during the scenario *city* (a) Input data with low SNR from 20 to 50 s. (b) Input data with high SNR from second 210 to 240. We created the visualisation with MATLAB (MATLAB version R2021a, The MathWorks, Natick, United States).

Signal fusion

We calculate the performance of (1) early fusion, (2) signal-based late fusion, (3) sensor-based late fusion, and (4) hybrid fusion (Fig. 5). The latter combines the other three with a majority voting²⁶. Many authors reported that one sensor is insufficient to measure vital signs in a medical setting and suggested redundancy with respect to both the number of sensors and the physical base of sensor systems^{15,16}. However, signals from multiple sensors need a fusion strategy. The algorithm from Chandra et al. estimates the heartbeat location of multiple signals based on a convolutional neural network (CNN)²⁵. This general approach can be applied to other data. Münzner et al. compared early fusion with several late fusions²⁷.

Early fusion merges in a first convolutional layer. Sensor-based late fusion merges in the dense layer, and signal-based late fusion with an increased number of extracted features. For respiratory rate detection, we combine these approaches and extend them with a majority vote.

The first layer receives our five input signals. The second layer is a convolutional layer, which generates a feature map. To extract the features, we use four filters with a length of 20 samples for each of the sensor signals²⁵. To prevent over-fitting, we use a dropout rate of 0.5 in the dropout layer. The pooling layer reduces the dimension of the feature map. The dense layer classifies snippets into binary classes 1 and 0, which means that the snippet contains a heartbeat or not, respectively. The activation function is a sigmoid function. The output layer generates a vector Y of multiple labels that are either 0 or 1 for a specific snippet. The majority vote takes place after the output layer.

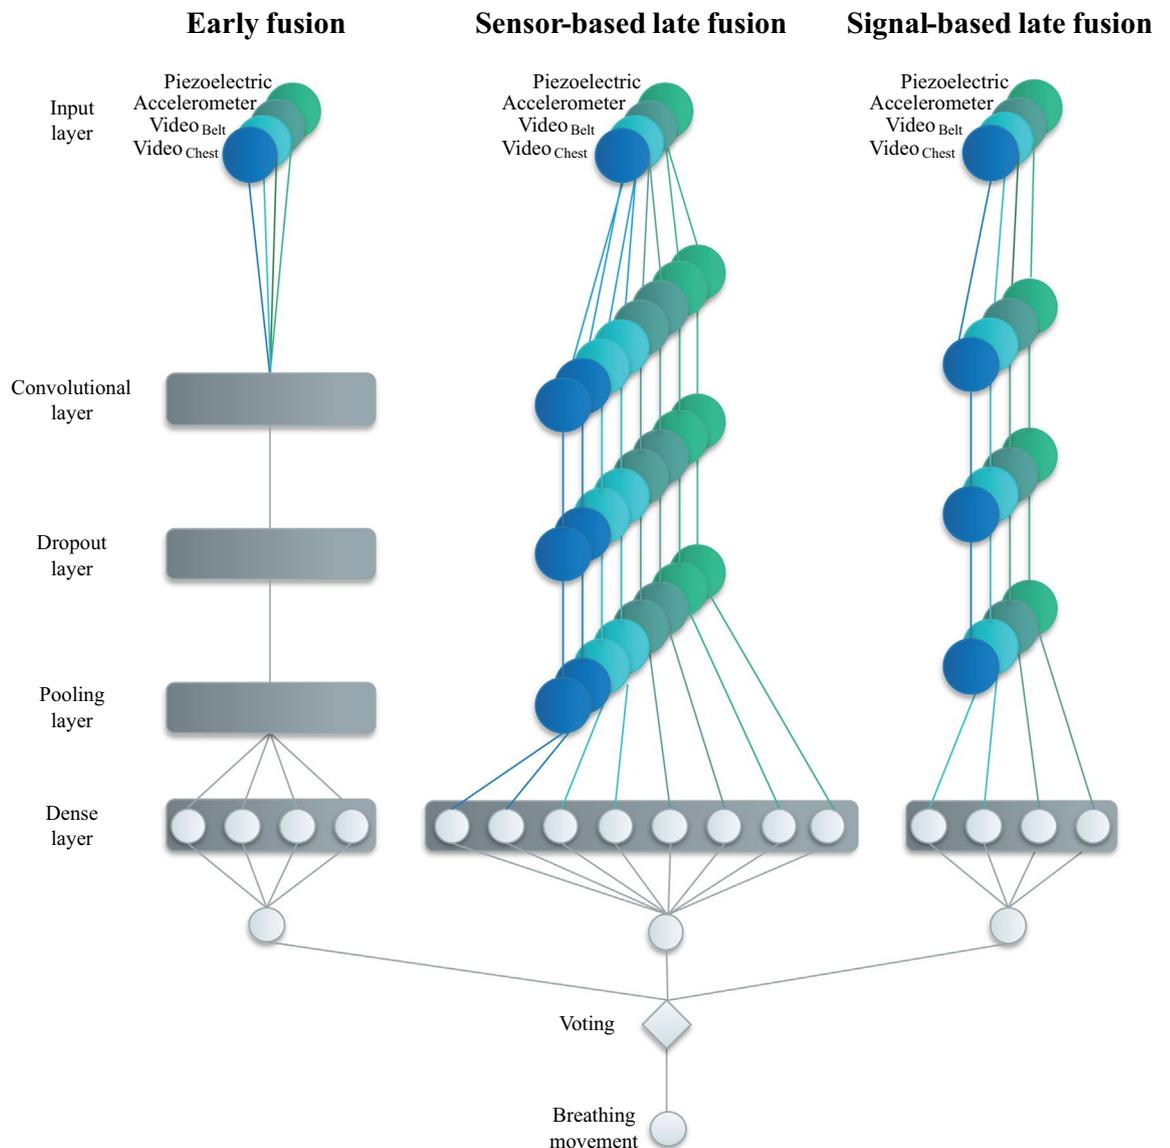


Figure 5. Hybrid signal fusion for multiple sensors.

Approach	Piezo [%]	Acc [%]	Video _{Belt} [%]	Video _{Chest} [%]
Early fusion	49.08	49.16	49.66	53.38
Signal-based late fusion	49.10	55.79	51.80	49.51
Sensor-based late fusion	49.10	49.19	53.58	51.40
Hybrid fusion	48.82	49.17	53.31	52.21

Table 1. Performance in percent for one signal during the three driving scenarios. Significant values are in bold.

Approach	Piezo+ Acc [%]	Piezo+ Video _{Belt} [%]	Piezo+ Video _{Chest} [%]	Acc+ Video _{Belt} [%]	Acc+ Video _{Chest} [%]	Video _{Belt} + Video _{Chest} [%]
Early fusion	52.35	46.58	54.25	50.32	46.58	42.21
Signal-based late fusion	53.18	56.08	47.36	48.44	50.89	52.08
Sensor-based late fusion	48.77	48.87	50.63	47.93	55.15	46.73
Hybrid fusion	50.32	50.91	50.71	50.13	51.29	46.89

Table 2. Performance in percent for two signals during the driving scenario. Significant values are in bold.

Approach	Piezo+Acc+Video _{Belt} [%]	Piezo+ Acc+ Video _{Chest} [%]	Piezo+ Video _{Belt} + Video _{Chest} [%]	Acc+ Video _{Belt} + Video _{Chest} [%]
Early fusion	52.39	50.83	53.63	53.59
Signal-based late fusion	55.64	54.36	48.34	46.29
Sensor-based late fusion	48.49	46.96	51.24	50.48
Hybrid fusion	52.35	49.51	50.67	49.06

Table 3. Performance in percent for three signals during the scenario *city*. Significant values are in bold.

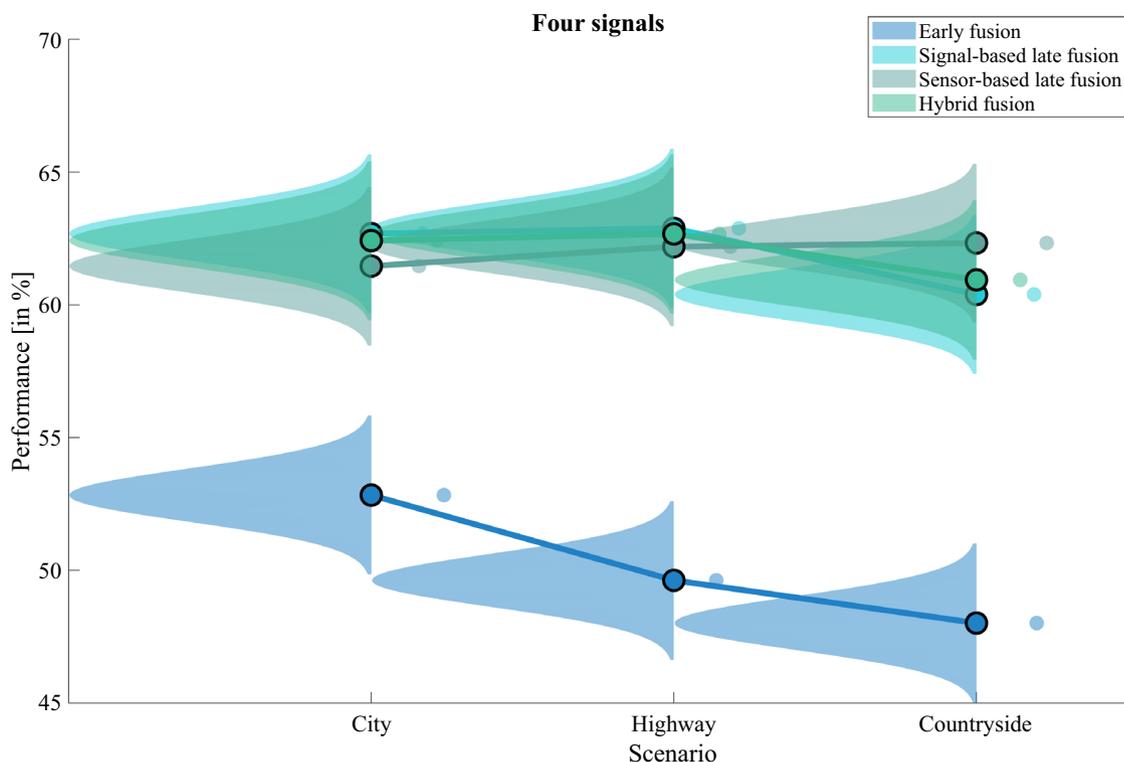


Figure 6. Performance of four signals for different driving scenarios. This plot was created with the MATLAB package from Allen et al.²⁸ (MATLAB version R2021a, The MathWorks, Natick, United States).

An early fusion approach uses a single CNN and fuses signals within the convolutional layer (Fig. 5). By using two CNNs per signal, the sensor-based late fusion method extracts more features per signal. The signal-based late fusion uses one CNN per signal. Both of these approaches merge in the dense layer.

Evaluation

Comparing snippets with ground truth, we obtain true-positive (TP), true-negative (TN), false-positive (FP), and false-negative (FN) classifications. We calculate performance $P = (PPV + S)/2$ based on the positive predictive value $PPV = TP/(TP + FP)$ and the sensitivity $S = TP/(TP + FN)$, as these evaluation metrics reflect the FP and FN numbers²⁵.

We first evaluate the different sensors and their combinations, disregarding the fusion algorithms and the driving scenarios. We then determine the best type of fusion and analyse the impact of the different driving scenarios. Finally, we determine the portion of driving time that can be used for reliable respiratory rate measurements.

Results

Comparison between one signal

The tables compare the average performance during the three driving scenarios for the four sensors to identify the signals with the best performance (Table 1). The signal-based late fusion yields the highest performance ($P_{max} = 55.79$) with the signal accelerometer (Table 1). In summary, the video_{Belt} has the highest score twice, and the accelerometer and video_{Chest} have the highest score once for a signal fusion approach. We will describe the de-noised accelerometer signal as *Acc* and the piezoelectric signal as *Piezo* in the following text to simplify the notation.

Approach	City [%]	Highway [%]	Countryside [%]	Mean _P [%]
Early fusion	52.83	49.62	48.01	50.15
Signal-based late fusion	62.68	62.88	60.39	61.98
Sensor-based late fusion	61.46	62.19	62.33	61.99
Hybrid fusion	62.42	62.67	60.94	62.01

Table 4. Performance in percent for four signals during driving. Significant values are in bold.

Comparison between two signals

The comparison of two sensor pairs shows that signal-based late fusion has the highest performance ($P_{max} = 56.08$) with the pair piezoelectric sensor and video_{Belt} (Table 2). The sensor pair video_{Belt} and video_{Chest} has the lowest performance for the early fusion approach ($P = 541.21$). In summary, the combination of accelerometer and video_{Chest} achieves the best performance twice, and the piezoelectric and video_{Belt} once.

Comparison between three signals

The three combined signals, piezoelectric signal, accelerometer, and video_{Belt} have the highest score ($P_{max} = 55.64$) with the approach signal-based late fusion (Table 3). On the other hand, this fusion approach also has the worst performance with the combination of piezoelectric sensor, video_{Belt}, and video_{Chest} ($P = 46.29$). However, the combination of a piezoelectric sensor, accelerometer, and video_{Belt} as well as a piezoelectric sensor, video_{Belt}, and video_{Chest} yields two times the highest performance per signal fusion approach.

Comparison between four signals

The raincloud plot visualises the distribution of the performance during the driving scenarios to show the performance differences (Fig. 6). The early fusion approach consistently performs poorly, with the lowest scores in all scenarios and the lowest average performance overall (Fig. 6). However, the distribution for the other fusion algorithms is comparable (Fig. 6).

With the combination of all four signals, the signal-based late fusion approach achieves the highest score ($P_{max} = 62.88$) in the scenario highway (Table 4). However, the hybrid fusion achieves the highest average performance ($Mean_P = 62.01$) overall driving scenarios. The early fusion technique yields the lowest average performance ($Mean_P = 50.15$).

Discussion

The major challenge for health monitoring during driving is the changing data quality due to artifacts, which are caused by movements of the car or driver due to driving activities or talking^{15,16}. Moreover, the different physical attributes, e.g., height and weight, of the test subjects lead to a high standard deviation in performance¹⁵. The iPPG signal is especially sensitive to light changes and movements²⁹. Moreover, other methods, such as PCA and FFT, could lead to better signal extraction. Furthermore, it is possible to integrate more sensors into the redundant sensor system. The system could be extended by a magnetic induction sensor integrated into the seat backrest³⁰. We excluded the BCG sensor because the pretest in the car showed that the SNR was low. For the pretest, we placed the sensor at the backrest to measure ballistic forces generated by the lungs. We also excluded a radar sensor because the API did not allow sufficient time synchronization. For our pretest, we used the radar sensor from Acconeer A111. The Raspberry Pi and camera are low-cost devices. The reliability of the Raspberry Pi is proven in other publications^{31,32}. There is also the possibility of optimizing the structure of signal fusion models. The performance of models with adjusted parameters could be investigated in further studies, e.g., the number of convolutional layers³³, decreasing or varying the learning rate of the Adam optimizer³⁴, or other activation functions³⁵. A long short-term memory (LSTM) or autoencoder-based structure could also increase the performance³⁶.

The early fusion has the lowest performance with $Mean_P = 50.15$. This is attributed to the extraction of fewer features compared to sensor-based fusion, which achieves the highest score after the hybrid fusion process. In sensor-based fusion, a more extensive feature set is obtained by employing two CNNs to extract features in the convolutional layer for a single signal, resulting in enhanced performance.

Moreover, it is also important to compare the results from this paper with the literature. Ju et al.¹⁷ showed the technical feasibility of the pressure sensor integrated into the seatbelt, and Vavrinsky et al.¹⁹ of a pressure sensor integrated into the seat to measure the respiratory rate. However, they did not collect data from various subjects and presented an evaluation. Beak et al. measured the electrocardiogram (ECG), galvanic skin response (GSR), and respiration to detect the driver's stress in a study with four male subjects under real driving conditions¹⁸. For the respiration rate measuring, they integrated a piezoelectric sensor into the seat belt. They did not use a reference sensor such as a chest belt to collect a ground truth for the respiratory rate to evaluate the respiratory rate monitoring because the focus of this study was on the detection of the driver's stress level. Vinci et al. used a microwave interferometer radar and conducted three measurements for 30 seconds to show the technical feasibility and waveform of the signal in a simulated environment²⁰. Guo et al. recorded data with a near-infrared time-of-flight camera to derive the respiratory rate of five subjects under real driving conditions and compared the results with the measurements of a chest belt as a reference²¹. The authors calculated the RR per minute. In contrast to our evaluation, they did not compare the position of the breathing movement with the reference. For 43 % of the driving time on the highway, they had a difference of 0 to 3 breaths per minute (BPM) compared to

the reference²¹. The novelty of our work lies in the fusion of the different signals for respiratory rate detection. Moreover, we collected data from 15 subjects under three different driving scenarios to enable further algorithm development with these data. The accuracy of the sensor system and fusion with respect to the other evaluation metrics and recordings is not directly comparable.

For future work, we will record the controller area network (CAN) Bus data to detect artifacts that are caused by environmental disturbances. As suggested by Fu et al.³⁷, we will also integrate movement detection using depth cameras. Other important vital signs are the heart rate and temperature. Future systems should record all primary vital signs, i.e., body temperature, pulse rate, and blood pressure. Additionally, we will publish a paper for in-vehicle heartbeat detection. The heartbeat is another important vital sign to detect cardiovascular diseases, such as stroke. The study encompasses 19 healthy subjects and adheres to the same experimental design as described in this paper³⁸. We publish two papers because the description of the different sensor systems, pre-processing, and analysis would exceed the scope of one journal paper.

Conclusion

To detect respiratory rate while driving, we developed a redundant sensor system and signal fusion approaches. Based on our results, the hybrid fusion and all four sensors have the highest performance for in-vehicle respiratory rate detection: *city* ($P = 62.42$), *highway* ($P = 62.67$), and *countryside* ($P = 60.94$). The result also shows that the fusion of multiple signals improves robustness. Furthermore, the voting system of the hybrid algorithm not only outperforms the other fusion algorithms but also presents several distinct advantages. This approach to algorithm fusion leverages the strengths of different techniques and combines them in a way that maximizes overall performance. By allowing each component algorithm to contribute its information through a voting mechanism, the hybrid algorithm achieves a higher accuracy and robustness. As a take-home message and to answer the initial research questions: *We can monitor the respiratory rate continuously for over 60 percent of our driving time with a low variance between the driving scenarios, and the combination of all sensors delivers the most reliable performance.* In summary, the results show the potential to detect CRD symptoms in an early stage.

Data availability

We publish the anonymous data in text (.txt) file format over the data storage system of TU Braunschweig under CC BY 4.0 to allow researchers to reproduce and apply further algorithms (link: <https://doi.org/10.24355/dbbs.084-202210201440-0>, accessed on November 16th, 2023): chest belt (reference signal), accelerometer 1 and 2, piezoelectric signal, video_{Belt}, and video_{Chest}, and subject information (subject ID, age, height, weight, gender, and known diseases).

Received: 19 April 2023; Accepted: 14 November 2023

Published online: 22 November 2023

References

- Nicolò, A., Massaroni, C., Schena, E. & Sacchetti, M. The importance of respiratory rate monitoring: From healthcare to sport and exercise. *Sensors (Basel)* **20**(21), 6396. <https://doi.org/10.3390/s20216396> (2020).
- Steichen, O., Grateau, G. & Bouvard, E. Respiratory rate: The neglected vital sign. *Med. J. Aust.* **189**(9), 531–2. <https://doi.org/10.5694/j.1326-5377.2008.tb02164.x> (2008).
- Churpek, M. et al. Derivation of a cardiac arrest prediction model using ward vital signs. *Crit. Care Med.* **40**(7), 2102–8. <https://doi.org/10.1097/CCM.0b013e318250aa5a> (2012).
- Tipton, M., Harper, A., Paton, J. & Costello, J. The human ventilatory response to stress: Rate or depth?. *J. Physiol.* **595**(17), 5729–52. <https://doi.org/10.1113/JP274596> (2017).
- Grassmann, M., Vlemincx, E., von Leupoldt, A., Mittelstädt, J. & Van den Bergh, O. Respiratory changes in response to cognitive load: A systematic review. *Neural Plast.* **2016**, 1–16. <https://doi.org/10.1155/2016/8146809> (2016).
- Chronic respiratory diseases [Internet]. Genève: World Health Organization [cited 2022]; [about screen 1]. Available from: <https://www.who.int/health-topics/chronic-respiratory-diseases>.
- Reddel, H. et al. Global initiative for asthma strategy 2021: Executive summary and rationale for key changes. *Am. J. Respir. Crit. Care Med.* **205**(1), 17–35. <https://doi.org/10.1164/rccm.202109-2205PP> (2022).
- Jiexing, L. et al. Role of pulmonary microorganisms in the development of chronic obstructive pulmonary disease. *Crit. Rev. Microbiol.* **47**, 1–12. <https://doi.org/10.1080/1040841X.2020.1830748> (2021).
- Beltramo, G. et al. Chronic respiratory diseases are predictors of severe outcome in covid-19 hospitalised patients: A nationwide study. *Eur. Respir. J.* **58**(6), 2004474. <https://doi.org/10.1183/13993003.04474-2020> (2021).
- Von Hertzen, L. & Haahela, T. Signs of reversing trends in prevalence of asthma. *Allergy* **60**(3), 283–92. <https://doi.org/10.1111/j.1398-9995.2005.00769.x> (2005).
- National travel survey: England 2018 [Internet]. London: Department of Transportation [cited 2022]; [about screen 1]. Available from: <https://www.gov.uk/government/statistics/national-travel-survey-2018>.
- Deserno, T. Transforming smart vehicles and smart homes into private diagnostic spaces. *Conf Proc APIT* 165–171, <https://doi.org/10.1145/3379310.3379325> (2020).
- Sahayadhas, A., Sundaraj, K. & Murugappan, M. Detecting driver drowsiness based on sensors: A review. *Sensors (Basel)* **12**(12), 16937–53. <https://doi.org/10.3390/s121216937> (2012).
- Ahmed, M., Hussein, H., Omar, B., Hameed, Q. & Hamdi, S. A review of recent advancements in the detection of driver drowsiness. *J. Algebr. Stat.* **16**(2), 2657–63 (2022).
- Leonhardt, S., Leicht, L. & Teichmann, D. Unobtrusive vital sign monitoring in automotive environments—a review. *Sensors (Basel)* **18**(9), 3080. <https://doi.org/10.3390/s18093080> (2018).
- Wang, J., Warnecke, J., Haghi, M. & Deserno, T. Unobtrusive health monitoring in private spaces: The smart vehicle. *Sensors (Basel)* **20**(9), 2442. <https://doi.org/10.3390/s20092442> (2020).
- Ju, J. et al. Real-time driver's biological signal monitoring system. *Sens. Mater.* **27**(1), 51–59 (2015).
- Baek, H., Lee, H., Kim, J., Choi, J. & Kim, K. Nonintrusive biological signal monitoring in a car to evaluate a driver's stress and health state. *Telemed. e-Health* **15**, 182–9. <https://doi.org/10.1089/tmj.2008.0090> (2009).
- Vavrinsky, E., Tvarozek, V., Stopjakova, V., Solarikova, P. & Brezina, I. Monitoring of car driver physiological parameters. *Conf. Proc. Adv. Semicond. Dev. Microsyst.* **2010**, 227–30. <https://doi.org/10.1109/ASDAM.2010.5667021> (2010).

20. Vinci, G., Lenhard, T., Will, C. & Koelpin, A. Microwave interferometer radar-based vital sign detection for driver monitoring systems. *Conf. Proc. IEEE Microw. Intell. Mob.* **2015**, 1–4. <https://doi.org/10.1109/ICMM.2015.7117940> (2015).
21. Guo, K. *et al.* Contactless vital sign monitoring system for in-vehicle driver monitoring using a near-infrared time-of-flight camera. *Appl. Sci.* **12**, 4416. <https://doi.org/10.3390/app12094416> (2022).
22. Wang, J., Warnecke, J. & Deserno, T. The vehicle as a diagnostic space: efficient placement of accelerometers for respiration monitoring during driving. *Stud. Health Technol. Inform.* **258**, 206–210 (2019).
23. Warnecke, J., Wang, J. & Deserno, T. Noise reduction for efficient in-vehicle respiration monitoring with accelerometers. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* **2019**, 3257–61. <https://doi.org/10.1109/EMBC.2019.8857654> (2019).
24. TU Braunschweig, Hardware [Internet]. [cited 2023]; [about 1 screen]. Available from: <https://www.tu-braunschweig.de/it/hpc>.
25. Chandra, B., Sastry, C. & Jana, S. Robust heartbeat detection from multimodal data via CNN-based generalizable information fusion. *IEEE Trans. Biomed. Eng.* **66**(3), 710–717. <https://doi.org/10.1109/TBME.2018.2854899> (2018).
26. Warnecke, J. *et al.* Sensor fusion for robust heartbeat detection during driving. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* **2021**, 447–450. <https://doi.org/10.1109/EMBC46164.2021.9630935> (2021).
27. Münzner, S. *et al.* CNN-based sensor fusion techniques for multimodal human activity recognition. *Conf. Proc. Symp. Wear Comp.* **2017**, 158–165. <https://doi.org/10.1145/3123021.3123046> (2017).
28. Allen, M., Poggiali, D., Whitaker, K., Marshall, T. & Kievit, R. Raincloud plots: A multi-platform tool for robust data visualization. *Wellcome Open Res.* **63**(4), 1–4. <https://doi.org/10.12688/wellcomeopenres.15191.2> (2019).
29. Chen, M., Zhu, Q., Zhang, H., Wu, M. & Wang, Q. Respiratory rate estimation from face videos. *Conf. Proc. IEEE Biomed. Health Inform.* **2019**, 1–4. <https://doi.org/10.1109/BHI.2019.8834499> (2019).
30. Vetter, P., Leicht, L., Leonhardt, S. & Teichmann, D. Integration of an electromagnetic coupled sensor into a driver seat for vital sign monitoring: Initial insight. *Conf. Proc. IEEE Veh. Electron. Saf.* **2017**, 185–90. <https://doi.org/10.1109/ICVES.2017.7991923> (2017).
31. Islam, M., Nooruddin, S., Karray, F. & Muhammad, G. Internet of things: Device capabilities, architectures, protocols, and smart applications in healthcare domain. *IEEE Internet Things* **10**, 3611–3641. <https://doi.org/10.1109/JIOT.2022.3228795> (2023).
32. Pardeshi, V., Sagar, S., Murmurwar, S. & Hage, P. Health monitoring systems using IoT and raspberry pi—A review. *Conf. Proc. Innov. Mech. Ind. Appl.* **2017**, 134–137. <https://doi.org/10.1109/ICIMIA.2017.7975587> (2017).
33. Li, F. *et al.* Feature extraction and classification of heart sound using 1d convolutional neural networks. *J. Adv. Signal Process.* **59**, 1–11. <https://doi.org/10.1186/s13634-019-0651-3> (2019).
34. Smith, L. N. Cyclical learning rates for training neural networks. *Conf. Proc. Appl. Comput. Vis.* **2017**, 464–72. <https://doi.org/10.1109/WACV.2017.58> (2017).
35. Isaac, B., Kinjo, H., Nakazono, K. & Oshiro, N. Suitable activity function of neural networks for data enlargement. *Conf. Proc. Control Autom. Syst.* **2018**, 392–397 (2018).
36. Park, S., Choi, S. & Park, K. Advance continuous monitoring of blood pressure and respiration rate using denoising auto encoder and lstm. *Microsyst. Technol.* **28**, 2181–2190. <https://doi.org/10.1007/s00542-022-05249-0> (2022).
37. Fu, C., Mertz, C. & Dolan, J. Lidar and monocular camera fusion: On-road depth completion for autonomous driving. *Conf. Proc. IEEE Intell. Trans. Syst. (ITSC)* **2019**, 273–278. <https://doi.org/10.1109/ITSC.2019.8917201> (2019).
38. Robust in-vehicle heartbeat detection using multimodal signal fusion. <https://doi.org/10.1038/s41598-023-47484-z>.

Acknowledgements

We would like to thank all the volunteers for their time and effort who supported the data acquisition and recording process.

Author contributions

Conceptualization, J.M.W.; methodology, J.L., T.M.D., and J.M.W.; validation, J.L., T.M.D., and J.M.W.; formal analysis, J.M.W.; investigation, T.M.D., J.L., and J.M.W.; resources, J.L., and T.M.D.; writing—original draft preparation, J.M.W.; writing—review and editing, J.L., and T.M.D.; visualization, J.M.W.; supervision, T.M.D. and J.L.; project administration, T.M.D.; funding acquisition, T.M.D. All authors reviewed the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL. This work is funded by the Lower Saxony Ministry of Science and Culture under Grant Number ZN3491 within the Lower Saxony “Vorab” of the Volkswagen Foundation and Center for Digital Innovations.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to J.M.W.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023