APOLLO 2.0: NON-INTRUSIVE VITAL SIGN MONITORING WITH WEARABLE TECHNOLOGY

Kenneth Koltermann
Advisor: Dr. Gang Zhou
William & Mary

April 7, 2023

Abstract

Portable, non-intrusive, and accurate vital sign detection systems are imperative for early emergency intervention where mere seconds can be the difference between life and death. Monitoring respiration rate in real-time can identify early signs of respiratory distress, but respiration rate is the most difficult of the three main vital signs to monitor in uncontrolled environments. To solve this problem, we introduce the Apollo V2, the 2nd iteration of our wearable vital sign monitoring shirt which is suitable for real-world environments. The shirt has no rigid components except for a small non-intrusive control patch and uses textile sensors to monitor respiration rate. We tested our system in a small user study consisting of both a sleep study and several awake activities. We demonstrate that Apollo V2 can accurately monitor respiration rate with an average error of less than 1 breaths per minute for sedentary and light exercise, and an average error of less than 1.4 breaths per minute for select medium to high intensity exercises.

Introduction

Accurate, real-time vital sign monitoring is important for diagnosing cardiovascular disease and preventing death by allowing faster intervention in life or death scenarios by medical professionals. The three most important vital signs to monitor are heart rate, body temperature, and respiration rate. There are many accurate, real-time vital sign monitoring devices such as the Holter monitor,[10] modern smart-watches,[7,9,17] or the Polar H10 heart rate monitor.[5] These systems are capable of accurately measuring heart rate and temperature, but are unable to accurately quantify respiration rate.

Respiratory arrest due to choking, asthma attack, infection, etc. can occur while a subject has a normal heart rate, and existing clinical respiratory monitoring systems are not suitable for real-world (IE: outside of a medical clinic) due to intrusiveness or lack of portability. There is a clear need for a respiration detection system that is portable, non-intrusive, and accurate that is suitable for daily-living and vigorous physical activity. Such a need leads us to the following research questions:

RQ1: *What wearable textile design allows for accurate respiration monitoring regardless of physical activity level?*

RQ2: *What features and models allow for the accurate estimation of respiration rate?*

In this paper, we answer **RQ1** by designing a wearable shirt that is capable of accurately measuring respiration signals by placing the respiration sensor in a manner that will be robust against different body types. We use 3 conductive fabric strips placed vertically and horizontally across the chest and abdomen so that respiration events can be identified regardless of body position. To answer **RQ2**, we compare both filtering and peak detection methods against a trained convolutional neural network (CNN) in order to demonstrate that machine learning methods are more suited for respiration detection. We demonstrate that our CNN based approach is capable of detecting respiration with a breaths-per-minute-error (bpme) of 0.47 for standing positions and 0.95 bpme for light exercises such as walking. Additionally, our preliminary results show that for medium to high intensity exercise, Apollo is capable of detecting respiration with an average bpme of 1.38.

Wearable Vital Sign Detection System Design

In this section, we describe the design of our Apollo V2 shirt that is capable of accurately monitoring the respiration of the wearer even during different wearer activity levels. We also describe in detail the differences between the original Apollo V1 shirt and the Apollo V2 shirt.

Apollo V2 Shirt Template

The underlying shirt template that our flexible textile sensors are attached to is made out of a spandex-blended compression material that accommodates multiple body types, is breathable for temperature regulation, and can be worn comfortably during a wide range of activities from sleeping to vigorous exercise. The Apollo V2 shirt is sleeveless unlike the Apollo V1 version that had short sleeves. The V2 version is sleeveless because we found that arm and shoulder movement in the V1 shirt would create noisy artifacts in the collected data. This is because arm movement in a shirt with sleeves will cause movement of the fabric of the rest of the shirt. For example, arm swinging while walking when testing Apollo V1 reduced the breathes per minute error (bpme). We explain this fact in more detail in the next section. Without sleeves, arm and shoulder movement has minimal impact on the collected data, and reduces the bpme.

Compression material is chosen because it form fits the wearer. Standard shirt sizes (US) are a simple function of height and circumference of the wearer that allows people of different body shapes to fit into the shirt size. Figure 1 shows several different body shapes. A compression shirt ensures that there is minimal "bagginess" of the shirt which could impact sensor performance. Finally, this compression material will have less overall movement during vigorous exercise since it will maintain skin contact with the wearer. Respiration belts used in a clinical setting[3,6] need to be manually fitted to the wearer which reduces ease-of-use while our system can accommodate multiple body types without requiring adjustment.

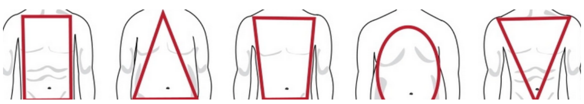

Figure 1: Common Body Shapes

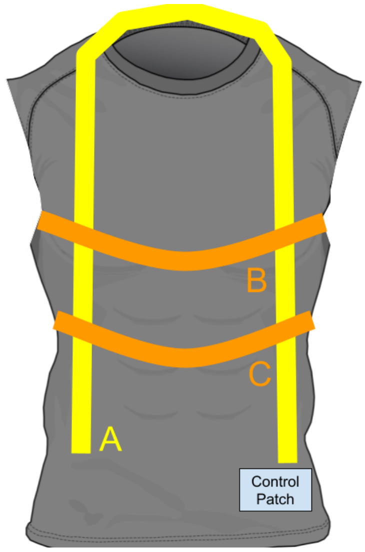

(a) Apollo V2 Shirt Design

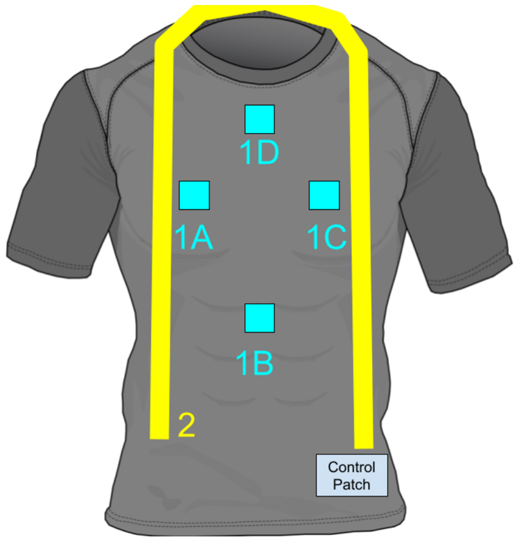

(b) Apollo V1 Shirt Design

Figure 2: Comparison Between 1st and 2nd Year Apollo Designs

Respiration Sensing Textile Sensors

To answer **RQ1**, we've tested many different shirt and sensor designs over the past year. To accurately determine the respiration rate of the user, we sewed three different conductive, textile-based, stretch sensors into the base shirts as shown in Figure 2a. These textile stretch sensors[14,15] change resistance depending upon the stretch force applied. The more the fabric is stretched, the lower the resistance across the fabric is reduced making an ideal material for flexible stretch sensors. As a wearer inhales and expands their diaphragm/abdomen, these stretch sensors will show a marked difference in resistance than when a wearer exhales. In Figure 2a, sensor A consists of a long strip of conductive fabric that extends from the bottom left of the front of the shirt vertically around the neck of the shirt to the bottom right of the front of the shirt. This sensor is placed in such a fashion as to be robust against changing variables such as different body types that fit into the same shirt or different types of breathing. For example, someone who "belly" breathes will experience less expansion of their chest than someone who diaphragm breathes.[2]

Traditionally, respiration belts designed to monitor respiration in a clinical setting are placed around the circumference of the chest and abdomen. These respiration belts are then manually tightened to suit the user's body type. However, in a shirt that multiple body types can wear, manually tightening of the shirt is not easily done. Sensors B and C in Figure 2a are conductive fabric strips that are sewn horizontally across the chest and abdomen in a similar way to how respiration belts function. As described in Section 4, sensors B and C are not as accurate as sensor A under most circumstances. However, certain postures such as the fetal position cause the sensor A to have slack which means that no stretch force is applied to sensor A during respiration. Sensors B and C are still useful in these positions.

Figure 2a and Figure 2 show the main sensor placement differences between Apollo V1 and V2. Unlike Apollo V1, Apollo V2 does not have pressure sensors (1A-D) on the chest and abdomen. Past works[16] have shown that pressure sensors can be used to accurately detect heart beat ballistic signals, but this was shown under sedentary conditions. With Apollo V1, we found that the compression material of the shirt was not sufficient to ensure that these pressure sensors maintained enough physical contact with the wearer during light exercise. Unfortunately, significant noise was present in the collected data which prevented accurate measurement of the user's heart rate. Furthermore, while the pressure sensors could be used for respiration monitoring, we found that the stretch sensors were much more robust against repeated physical events such as the vibration of the heel striking the floor during gait. For real-world deployment, Apollo V2 can be coupled with smartwatches[9,17] or other fitness trackers[7] that have difficulty recording respiration rate, but accurately estimate heart rate.

Apollo V2 Control Patch

Similarly to Apollo V1, Apollo V2 requires a control patch on the front lower-left of the shirt. The control patch is fully detachable using conductive metal snap buttons so that the shirt can be machine washed. The patch consists of a Arduino Uno microcontroller,[13] a battery with a voltage converter, 3 resistors, a 9 degrees of freedom (DOF) BNO055 IMU sensor,[12] and a Bluetooth module. Conductive thread[4] is used to connect these components as well as connect the control patch to the textile sensors on the shirt. The Apollo V1 shirt used the CurieNano[1] microcontroller which made the microcontroller smaller, but the IMU contained on the CurieNano does not perform sensor fusion like the BNO055 IMU sensor does. Since different body types may influence the orientation of the control patch on the shirt, an IMU with sensor fusion was important for "normalization" between individuals. The control patch samples data at a rate of 100Hz and transmits the data to an Android smartphone via Bluetooth. The 3 resistors contained in the control patch are necessary to construct 3 voltage dividers which enables our system to record the change in resistance of each textile sensor. Each textile sensor is connected via snap buttons and conductive thread to different analog pins on the microcontroller.

Vital Sign Modeling

The data collected by the control patch consists of 3 analog pin readings pertaining to stretch sensors A, B, and C, as well as 3 accelerometer values (AX, AY, AZ) after sensor fusion is performed by the IMU. As noted in the previous iteration of this work, the textile sensor data can be influenced by multiple factors including body movement not related to respiration, body type, and different postures. In this section, we detail the digital signal processing techniques used to clean the data as well as the two forms of respiration rate estimation used.

Signal Processing of Textile Sensor Data

Stretch sensors A, B, and C are each connected to different analog pins on the microcontroller which are capable of reporting a value in the range of 0-1023. This raw value is converted to voltage using the following equation where A_i is an analog reading pertaining to one of the textile sensors, V_{in} is the voltage supplied by the microcontroller, and V_i is the resultant voltage measured by the analog pin:

$$V_i = A_i \times \frac{V_{in}}{1023} \quad (1)$$

Using this recorded voltage, we can then calculate the resistance of each textile sensor using Equation 2 derived from the standard voltage divider formula. R_{S_i} is the resistance of sensor i and R_{r_i} is the resistance of static resistor on the control patch pertaining to sensor i.

$$R_{S_i} = R_{r_i} \times \left(\frac{V_{in}}{V_{out}} - 1 \right) \quad (2)$$

Given the properties of the conductive textile material, we generally expect the resistance of each fabric strip to decrease as the diaphragm expands stretching each sensor. An important note about these textile sensors is that each sensor will have a maximum resistance, R_{max}, and a minimum resistance, R_{min}, which are defined by the properties of each sensor (conductive material type and textile dimensions). Typically R_{max} is reached before the textile sensor is maximally stretched which has implications for determining the size of an Apollo V2 shirt for the wearer.

The finer details about how someone breathes may vary from person to person, but one constant is that the diaphragm or abdomen will expand and contract during standard respiration. The stretch force applied during respiration by two different people that fit into the same Apollo shirt may be different. As such, it is important that we normalize the data from each subject individually to avoid data from subjects with a wider range of recorded resistances compressing the data from subjects with a smaller range. The data for each subject is normalized to a range of 0 to 1 which is more suitable for our modeling purposes.

Respiration Rate Estimation Using Non-Machine Learning Methods

As shown in the previous iteration of Apollo, peak detection algorithms[18, 19] can be a reliable means of extracting the respiration rate from textile based stretch sensors. Inhalation stretches these sensors and reduces the resistance measured. As a result, the peaks in the data will denote where the chest is deflated and the valleys will denote where the chest is inflated.

The time-series resistance data exhibits significant high-frequency noise arising from measurement error by the microcontroller. Normal human respiration is usually 12-40 breaths per minute, or 0.2-0.67Hz. Higher frequencies that are dominant in the sensor data should not be due to human respiration and can be ignored. To do this, we use a 5th order Butterworth[8] bandpass filter between 0.2-0.67Hz to remove high frequency noise. Lower frequency signals are also removed, though these are less impactful than high frequency noise. Figure 3a and Figure 3b show this transformation from raw data into a denoised, filtered form. We use an adaptive peak detection algorithm[18] to mark the peaks in the filtered data which denote where exhalation has occurred. Figure 3b shows a comparison of the ground truth respiration markers against the detected peaks in the data for a sample of time where a subject was walking. As can be seen, the number of peaks matches the number of respiration events marked in the ground truth. There is some time offset between these values as a result of the ground truth labeling method. We discuss this in a later subsection.

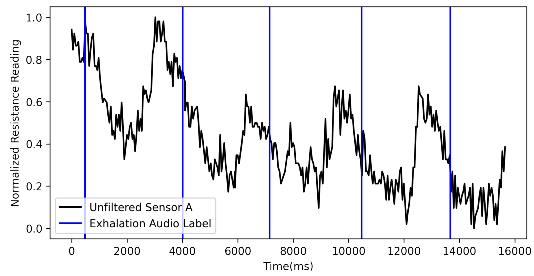

(a) Unfiltered Sensor A Data

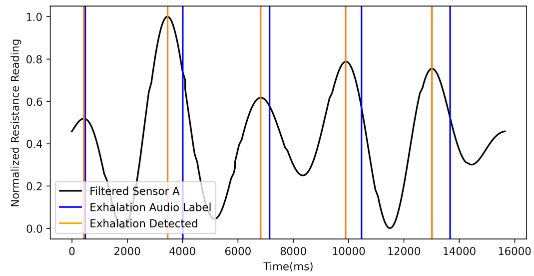

(b) Filtered Sensor A Data With Peak Detection

Figure 3: Comparison Between Unfiltered and Filtered Sensor A Data

Respiration Rate Estimation With Neural Networks

In the previous version of Apollo, we elected to not use machine learning due to the limited amount of data that we were able to collect due to Covid-19 in order to answer **RQ2**. However, we were able to collect more subject data this year which allowed us to train a convolutional neural network (CNN) for respiration rate estimation. CNNs are proven to be effective at extracting patterns that exist in time-series data, but are less computationally expensive than long-short-term-memory recurrent neural networks. Additionally, in theory, previous respiration events should have minimal impact on future respiration events for healthy subjects. This may not be the case for someone with a respiratory illness, but the subjects in our study are all healthy. As such, a CNN is suitable because respiration events are less spatially linked as the time between events is increased. For respiratory rate detection and diagnosing respiratory illnesses, a transformer-based neural network capable of associating many individual respiration events may be more useful. Figure 4 shows the structure of our CNN model. Our model takes in two inputs: a 1D time-series window consisting of the normalized, unfiltered sensor data, and a 1D time-series window consisting of normalized and filtered data discussed in the previous section. The size of these windows is 5 seconds, or 500 values since we record data at a frequency of 100Hz, and the step size for these windows is 0.25s. The These inputs enter two different branches of our CNN model consisting of three 1D convolutional layers which allows the model to learn patterns that exist in the raw data separate from the filtered data. Three convolutional layers per branch was chosen for two

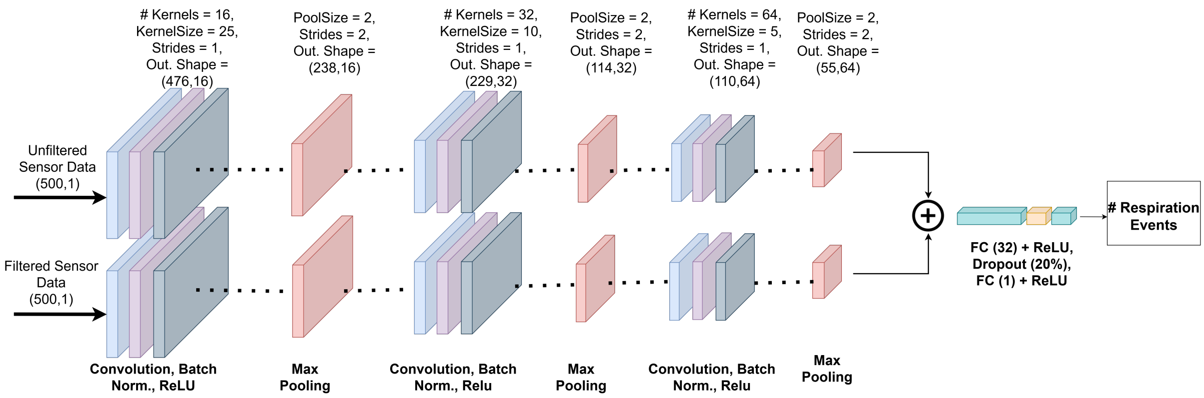

Figure 4: Apollo V2 CNN Model

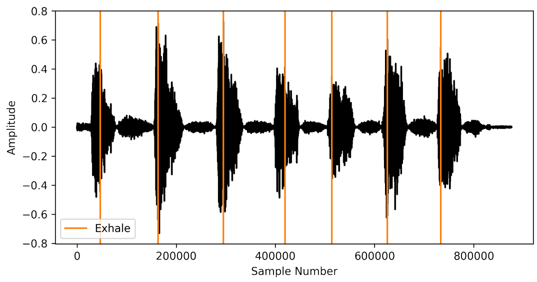

Figure 5: Apollo V2 CNN Model

SciPy peak detection was performed with a minimum distance determined by visual inspection of the audio signal in order to automatically mark inhalation and exhalation. Visual inspection of the labels was performed to ensure accuracy, and any incorrect labels were corrected through video and audio inspection. Figure 5 shows a typical sample of audio recorded from one of the subjects in our study. Exhalation is marked by a larger sound amplitude than inhalation.

reasons. 1. While we have sufficient data to train our model, deeper neural networks typically require more data for adequate training to occur. 2. A shallower CNN will avoid the vanishing gradient problem since typical measures for alleviating this problem with deeper neural networks such as ResNet[11] based architectures will require more data than we have collected. The output of the model is an integer denoting the number of respiration events, specifically inhalations, exist in the input data. To get the number of breaths per-minute, we perform a moving average over all the windows contained within a certain time frame.

Respiration Ground Truth

The ground truth for respiration rate was derived from the recorded respiration audio of each subject. During the tests performed by the subjects in our study, each subject wore an N95 mask with a microphone inserted into the mask. The mask served to muffle and reduce noise from the environment which helped us isolate the breathing signal of the subject. Originally, we explored existing machine learning models for respiration analysis, but found that by controlling the user study environment, we could eliminate the majority of external noise that could potentially cause issues. We detail this in the next section.

Experimental Results

In this section, we provide details about our user study, and evaluate the ability of Apollo to monitor respiration rate under different user activities. Furthermore, we show the experimental performance of Apollo when tested with medium to high intensity activities.

User Study

In Apollo V1, we collected small amounts of data from three users for sitting, standing, and walking activities. For Apollo V2, we collected significantly more data for an additional three users for sleeping, standing, and for various walking speeds. Additionally, we collected data for medium to high intensity activities from a single user. All data collection was conducted in strict accordance with an accepted IRB protocol. During each test, users in our study wore a N95 mask with a microphone inside connected to an Android smartphone which recorded respiration audio to serve as the ground truth for our study.

For the sleeping portion of our study, each user was asked to wear the Apollo V2 system while taking a nap for 1 hour. No restriction was placed on sleeping position or user comfort (pillow size, blankets, etc). For both the awake and sleeping portions of our study, no restrictions were placed on the user for adjusting the shirt as they might adjust their clothing at anytime in the real-world. For the awake portion of our study, users were asked to repeatedly stand still for one minute, then continuously walk for two minutes. Users were asked to repeat this this sequence as many times as they felt comfortable, and users were given the option to take a break whenever they wished to. For the medium to high intensity activities, a single user performed jumping jacks, pull ups, and ran at a fast pace. The three users in our Apollo V2 study all fit into the same XL shirt, and each fit into a different body type category shown in Figure 1. This fact is essential for us to determine the effectiveness of Apollo for different body shapes that fit into the same sized shirt.

Respiration Evaluation for Sensor A

The most common and important metric other respiration detection works use is breaths per minute error (bpme). Too large a bpme and a system becomes unreliable. In literature, a bpme of ≤ 1 is considered highly accurate, though the degree of acceptable error may depend upon the specific applications of the system. In our case, we are looking to achieve two goals: a low bpme of ≤ 1 and a stable bpme across multiple subjects, akin to accuracy and precision.

To evaluate the peak detection based approach for respiration analysis, bpme is a simple function of comparing the number of breaths detected vs. the ground truth for the entire dataset since this method does not "fit" to the data like a machine learning model does. Given the limited number of subjects we have, we've opted to perform our evaluation of our CNN model under subject-dependent conditions since it is highly likely that our model will not generalize to unknown subjects with vastly different characteristics than the subjects present in the training data. As such, we've elected to perform 10-fold cross validation for each subject where the test set for each fold is a random slice of 20% of a subjects data. 10-fold cross validation is performed for each activity. An important note is that we ensure that the 20% of the data used as the test set for each fold is not present in the training data to ensure that no data leakage artificially boosts the performance of our model.

Table 1 shows the performance for the peak detection algorithm using data from Sensor A (vertical textile sensor) in Figure 2a and Sensor 2 in Figure 2b for standing and walking for all 6 subjects. As expected, bpme increases during walking (1.34) when compared to standing (0.48). However, we believe this increase to be acceptable given the simplicity of this approach. Table 2 shows the 10-fold cross validation performed for each subject's standing and walking tasks. As can be seen, the CNN model slightly outperformed the peak detection algorithm, achieving a bpme of (0.47) for standing and a bpme of (0.95) for walking. We hypothesize this is the case because the CNN model was able to learn from both the raw and filtered signal data which may have helped the CNN model learn to ignore noise in the signal stemming from foot falls. The minimal difference between standing bpme for peak detection and our CNN model may suggest that we are approaching the minimum bpme that Apollo is capable of achieving, possibly due to limitations in the materials and fabrication of the shirts themselves. Sensors B and C proved to be less accurate than sensor A for subjects 4, 5, and 6. This is likely due to body shape influencing the maximum amount of stretch applied during respiration. For example, sensor C was at near maximum stretch while sensor B was relatively slack for participant 5.

Participant 4 completed medium to high intensity activities (jumping jacks, pull ups, running). The average bpme for our CNN model for this participant was 1.38 bpme. The error for the peak detection algorithm was 7.10 bpme. These preliminary results reveal that our CNN model is capable of ignoring noise in the data generated from body movement.

Conclusion and Future Work

In this work, we demonstrated Apollo V2, the 2nd iteration of our wearable vital sign monitoring shirt. We demonstrated that our Apollo shirt and CNN based approach is capable of detecting respiration with a breaths-per-minute-error (bpme) of 0.47 for standing positions and 0.95 bpme for light exercises such as walking. Additionally, our preliminary results showed that for medium to high intensity exercise, Apollo is capable of detecting respiration with an average bpme of 1.38. We believe that further improvements can be made by segmenting the vertical respiration sensor in to sections to even better accommodate different body shapes, as well as further improve movement noise reduction.

Acknowledgements

This work was funded by the Virginia Space Grant Consortium (VSGC) to whom we are most grateful for their support. Additionally, we wish to acknowledge the subjects who agreed to participate in our user study for giving their time to aid us in our research.

Activity	Sub. 1	Sub. 2	Sub. 3	Sub. 4	Sub. 5	Sub. 6	Avg.
Standing	0.49	0.32	0.77	0.67	0.59	0.41	0.48
Walking	1.12	1.52	2.51	0.90	0.84	1.21	1.34

Table 1: Peak Detection Error (BPME) For Vertical Stretch Sensor

Activity	Sub. 1	Sub. 2	Sub. 3	Sub. 4	Sub. 5	Sub. 6	Avg.
Standing	0.45	0.37	0.62	0.21	0.76	0.38	0.47
Walking	0.99	0.87	1.26	0.74	0.95	0.88	0.95

Table 2: CNN Detection Error (BPME) For Vertical Stretch Sensor

References

[1] Dfrobot curienano. https://www.dfrobot.com/product-1537.html.

[2] Diaphragmatic breathing exercises amp; benefits. https://my.clevelandclinic.org/health/articles/9445-diaphragmatic-breathing.

[3] Go direct respiration belt. https://www.vernier.com/product/godirect-respirationbelt/.

[4] Pbo fiber: Zylon. http://www.toyobo-global.com/seihin/kc/pbo/.

[5] Polar h10 heart rate monitor. https://www.polar.com/us-en/products/accessories/h10_heart_rate_sensor.

[6] Respiration sensor. https://bio-medical.com/respirationsensor.html.

[7] Benedetto, S., Caldato, C., Bazzan, E., Greenwood, D., Pensabene, V., and Actis, P. Assessment of the fitbit charge 2 for monitoring heart rate. *PLoS ONE 13* (01 2018).

[8] Butterworth, S. On the theory of filter amplifiers. *Experimental Wireless and the Wireless Engineer 7* (1930).

[9] Falter, M., Budts, W., Goetschalckx, K., Cornelissen, V., and Buys, R. Accuracy of apple watch measurements of heart rate and energy expenditure in patients with cardiovascular disease (preprint), 08 2018.

[10] Galli, A., Ambrosini, F., and Lombardi, F. Holter monitoring and loop recorders: From research to clinical practice. *Arrhythmia Electrophysiology Review 5* (08 2016).

[11] He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *CoRR abs/1512.03385* (2015).

[12] Industries, A. Adafruit 9-dof absolute orientation imu fusion breakout - bno055. https://www.adafruit.com/product/2472.

[13] Industries, A. Arduino uno r3. https://www.adafruit.com/product/4806.

[14] Industries, A. Knit jersey conductive fabric - 20cm square. https://www.adafruit.com/product/1364.

[15] Industries, S. Eeontex conductive stretchable fabric. https://www.sparkfun.com/products/retired/14112.

[16] Kiaghadi, A., Homayounfar, S. Z., Gummeson, J., Andrew, T., and Ganesan, D. Phyjama: Physiological sensing via fiber-enhanced pyjamas. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. 3*, 3 (sep 2019).

[17] Sarhaddi, F., Kazemi, K., Azimi, I., Cao, R., Niela-Vilen, H., Axelin, A., Liljeberg, P., and Rahmani, A. M. A comprehensive accuracy assessment of samsung smartwatch heart rate and heart rate variability. *PLOS ONE 17* (12 2022), e0268361.

[18] Scholkmann, F., Boss, J., and Wolf, M. An efficient algorithm for automatic peak detection in noisy periodic and quasi-periodic signals. *Algorithms 5* (11 2012), 588–603.

[19] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods 17* (2020), 261–272.