2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) | 979-8-3503-2447-1/23/$31.00 ©2023 IEEE | DOI: 10.1109/EMBC40787.2023.10340949

Towards a Lightweight Classifier to Detect Hypovolemic Shock

Leena Pramanik, Christopher L. Felton, Robert W. Techentin,
David R. Holmes, III, Ph.D., Timothy B. Curry, M.D., Ph.D., Michael J. Joyner, M.D.,
Victor A. Convertino, Ph.D., and Clifton R. Haider, Ph.D.
Mayo Clinic, Rochester, MN, USA; {felton.christopher, haider.clifton}@mayo.edu

***Abstract*—Predicting the ability of an individual to compensate for blood loss during hemorrhage and detect the likely onset of hypovolemic shock is necessary to permit early clinical intervention. Towards this end, the compensatory reserve metric (CRM) has been demonstrated to directly correlate with an individual's ability to maintain compensatory mechanisms during loss of blood volume from onset (one-hundred percent health) to exsanguination (zero percent health). This effort describes a lightweight, three-class predictor (good, fair, poor) of an individual's compensatory reserve using a linear support-vector machine (SVM) classifier. A moving mean filter of the predictions demonstrates a feasible model for implementation of real-time hypovolemia monitoring on a wearable device, requiring only 408 bytes to store the models' coefficients and minimal processor cycles to complete the computations.**

I. Introduction

The ability to measure a patient's progression of blood loss from point of injury to exsanguination permits for therapeutic intervention at an earlier, more efficacious point. Specifically, in order to survive pronounced blood loss, the body must maintain delivery of oxygen by blood to the tissues across the body through many compensatory mechanisms, e.g. vasoconstriction. In other words, the delivery of oxygen to the tissue is physiologically compensated up to the point of decompensation [3] at which point so much blood has been lost that the body can no longer maintain delivery of oxygen to the tissues at which point vitals such as heart rate increase and blood pressure decreases until death or resuscitation efforts are successful. Unfortunately, the standard vital signs clinically monitored, e.g., heart rate and blood pressure, fail to provide early indication for the triage of hypovolemia [1]. The primary limitation of standard clinical vitals are two-fold: 1) it is well known that vitals, e.g. blood pressure, often exhibit a marked change only when the individual is in late stages of hypovolemia and the body is no longer able to compensate for further blood loss, and 2) every individual has a unique ability to compensate for blood loss so a given change in standard vitals does not reflect the same change in compensatory reserve across individuals limiting their utility.

To address the need to improve detection of a patient's ability to compensate for further blood loss, a compensatory reserve (CR) physiologic model has been proposed [1]. The CR model has demonstrated in human subject studies [1], [2], [3], how an individual physiologically compensates to progressive blood loss to exsanguination of decompensation, i.e., has a personalized reserve tank. The standard vitals are sustained by the compensation, enough so, that the onset and progress towards exsanguination is not detected until decompensation, vitals crash, i.e., until the reserve tank runs out in the late stages of hypovolemia.

The compensatory reserve physiologic model provides a mechanism for monitoring the progression of hypovolemia to the point of decompensation. The CR is described with a CR metric (CRM) which ranges from 100-0%, where 100% represents a healthy individual, no reserve used, and 0% indicates an individual's point of decompensation, all reserve consumed. CR has been demonstrated in various human subject experiments [4], [2], [5], [3], and extraction of an estimate of CRM from invasive arterial measurements has been presented in [1]. A machine-learning (ML) convolutional neural-network regression model has been created to generate CRM [6] from volumetric waveforms such as arterial blood pressure (ABP), reconstructed brachial arterial blood pressure [7], [8], and photoplethysmography (PPG).

The work herein presents an alternative lightweight classifier to evaluate a subject's CR as either good, fair, or poor from PPG signals measured with a noninvasive optical finger-clamp, like a finger-clamp used with pulse-oximeters to measure SpO2. The lightweight ML model is suitable to monitor real-time progression towards a decompensation state in an ambulatory field scenario. A three-class characterization of the CRM can achieve a lightweight implementation capable of accurately predicting useful information about the patient's status, where the "fair" (moderate compromise) and "poor" (unstable) classes are particularly of interest as these determine the need for intervention.

For ambulatory monitoring, the lightweight prediction algorithm must execute on a small, low-energy wearable device, and such devices often comprised of a 100-mAh to 200-mAh battery with limited computational processing power. Nonetheless, most modern low-energy microcontrollers include a single precision floating-point unit (FPU) and can operate at close to 100 MHz; however, to conserve energy the average execution frequency is maintained as low as possible to extend runtimes. The CRM ML classifier provides an approach that uses minimal resources including memory storage and cycles to execute the algorithm. A lightweight implementation on a small wearable device enables ubiquitous noninvasive and ambulatory monitoring of hypovolemia, internal hemorrhage, and other life-threatening conditions. If undiagnosed [9] or only detected in late stages of hypovolemia [10], these conditions can benefit from earlier detection and intervention that can be provided by the

979-8-3503-2447-1/23/$31.00 ©2023 IEEE

lightweight algorithm described herein and implemented in a low-energy wearable device for extended monitoring times.

II. Methods

Noninvasive physiologic signals related to arterial waveforms (PPG, FinoMeter, etc.), captured during prior human subject lower body negative pressure (LBNP) studies, [1] were analyzed and used to train the CRM classifier. The data used to train and evaluate the ML classifier developed and reviewed herein was created with a set of lower body negative pressure (LBNP) human subject tests completed by the United States Army Institute of Surgical Research (USAISR) [1]. The data was available to the Mayo Clinic researchers through a Cooperative Research And Development Agreement (CRADA) and the conducting human subject Institution's Review Board (IRB) approved all experimental procedures involving human subjects. The USAISR LBNP protocol is reviewed in [1]. The USAISR LBNP dataset contains 201 subjects total, and 59 subjects with PPG waveforms used for this effort. The PPG was collected with a finger-clamp commonly used to sense SpO2, however the full-waveform PPG signal was recorded for evaluation.

The CR is estimated by extracting features from the volumetric signal and using a support vector machine (SVM) classifier to map features to a predicted class of the CRM estimate of good, fair, or poor. Although various sensing of the arterial volumetric changes are available, the work herein focuses on photoplethysmography (PPG). In addition, lightweight classifiers are considered to realize an embedded and real-time implementation of the classifier.

A. Preprocessing

The PPG signals are first preprocessed with basic filtering. A high-pass filter is employed to remove offsets and low frequency wander. A low-pass filter is used to reduce noise and bandwidth. The filters are implemented as single-precision floating point, both are recursive filters. The HPF is a filter structure that provides a notch at 0 Hz, the width of the notch can be controlled by the α coefficient [11], $\alpha = 0.975$ was used. The LPF is an IIR Direct Form-II filter with a passband frequency of 48 Hz and a stopband frequency of 60 Hz. A Savitky-Golay [12] filter is used for additional PPG signal smoothing, with an M=13 (window length=37) and N=4 (polynomial order) configuration.

The preprocessing chain is illustrated in Fig. 1 and the result of the preprocessing is demonstrated in Fig. 2.

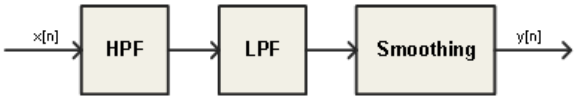

Figure 1. A block diagram of the preprocessing chain.

B. Features

The applications for this algorithm involves using PPG signals. The features are generated over each five-second window and averaged. After the PPG preprocessing, the PPG pulses are segmented on a beat-to-beat basis and the features extracted. The beat-to-beat segmentation can be achieved, for example, by using the approached outlined in [13].

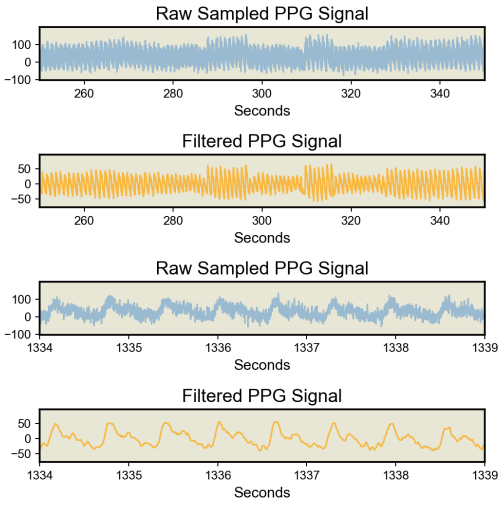

Figure 2. An example of the preprocessing on a photoplethysmography (PPG) signal.

The PPG pulse-rate segmentation used was previously proposed in [13]. After the pulse-rate segmentation based on [13] a derivative from the peak to the foot is taken; this derivative will provide location the of changes in the down slope. Features such as an estimate of the dicrotic notch location in time were identified including other waveform features on the downslope from the peak to the foot of the diastolic phase.

The amplitude and time index of the points in the PPG time-series signal were used as features to train the machine-learning algorithm.

The features from the signal are the primary inputs to the model used, however a subset of prior feature values are included to provide a history to each instance's set of features. The features are presented to the machine-learning algorithms at an asynchronous rate, when a peak of a signal is detected.

C. Classification of Hypovolemic Compromise

As presented in [6] a linear target of the LBNP can be used to train the model. The features previously described can be used to predict the current CRM state.

However, in a real-time device, reporting the predicted CRM along a linear objective is not needed. In place of modeling regression on the linear fit of hypolovemia derived from the LBNP studies, the physiological state may be described by percentile classifications of the CRM. Rather than predicting the CRM value, we segment the range of CRM values into three bins: good, fair and poor. The selected three classes were binned at respective thresholds of 100%-60%, 60%-30%, and 30%-0% These classifications are sufficient to provide the information required in a real-time device to

create a predictive indicator for progression toward and onset of hypovolemia.

A linear support-vector machine classifier (SVC) is used to predict the discretized CRM class bin. A linear and radial basis function (RBF) kernel SVC was trained for and evaluated. In addition to the SVC, a Random Forest and Ridge classifiers were trained and evaluated. The resources (target memory usage) and performance of each classifier is reviewed in the results section. Data from 59 LBNP study subjects were used, and the SVC was trained on 80% of the samples from these subjects and tested on the other 20%. To better appreciate the utility of the set of features being utilized, each feature's significance was measured by contribution to the macro-averaged area under the receiving-operator characteristic (AUROC) curve.

To further optimize the model, a causal, windowed statistical filter was implemented on the classifier predictions. A particle swarm optimization [14] technique was subsequently employed for the identification of an appropriate statistic and window size that would best smooth the outputs from the classifier. Specifically, the mean and median statistics were explored with smoothing windows between 0 and 25 samples wide with the objective being maximizing the macro-class-averaged F1-score. Consideration of the macro-averaged metrics (AUROC and F1-score), which weigh each prediction class equally, are significant to this application because there are fewer instances in the data set that belong to the "fair" class, yet robust identification of movement into this class is needed for effective intervention.

III. Results

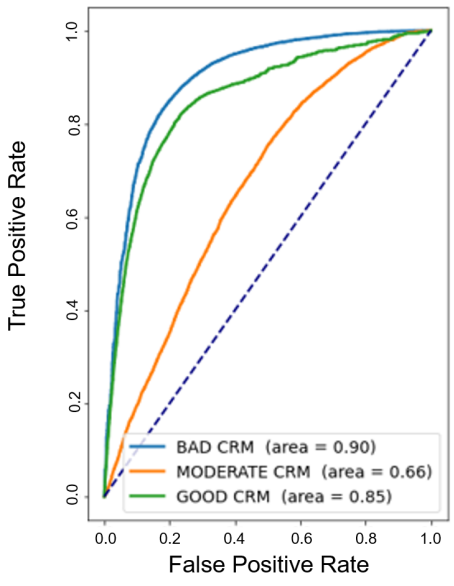

Figure 3. Linear SVC AUROC test data performance for predicting each class after training with all features.

Using the complete set of 34 features, the linear SVC achieves a macro-class averaged AUROC of 83% presented in Fig. 3 while generating coefficients which require 0.408 KBytes of storage. If only a good/bad binary classifier is used, the AUROC metric can be improved, however the *fair* state is critical in clinical applications for both prediction and intervention. A causal, moving-mean over 23-samples was identified as the optimal design for filtering the SVC's outputs, and the addition of post-prediction filtering increases the macro-averaged F1-score across these four test subjects from 0.5835 to 0.7223 illustrated in Fig. 4.

An optimization technique employed, was to rank the 34 features by most impactful to the macro-averaged AUROC metric. The ranking of features is presented in Fig. 5. A minimal improvement was observed in the AUROC, for each of the three classes, after 23/34 features are employed for training. TO further reduce the Linear SVC complexity, the input feature vector could be reduced to 23 features. Nonetheless, the full set of features is used to maximize the predictive ability of the fair state.

To provide a more comprehensive understanding of the trade-offs for selecting a linear SVC, a Random Forest, SVC Radial-Basis Function (SVC-RBF), and Ridge classifiers are examined and the performance is summarized in Table 1. For example, the Random Forest classifier explored achieves the best performance by all metrics of those in Table 1. However, the Random Forest classifier 2833.32 KBytes for storage, considerably more than the linear SVC and exceeds the budget for most low-energy microcontrollers. Optimal post-classification filtering utilized a 15-sample moving median and improved the Random Forest average F1-score, from the same four test subjects as with the SVC, from 0.7639 to 0.7994.

IV. Discussion

The work presented is focused on meeting the need to predict onset of hypovolemia for lightweight implementation on a wearable, and we demonstrate comparable performance to more computationally intensive models for prediction. The SVC used substantially less memory than other models. Although notable that the Random Forest demonstrated improved predictive ability in the fair class with 81% AUROC compared to the SVC's 66% fair AUROC, the Random Forest model used 1800% more memory than the SVC. However, all methods presented performed poorer than prior deep learning models [6].

Because of the time-series nature of the CRM, where the current state may be largely dependent on previous beat-to-beat states, networks like recurrent neural networks (RNN) and long short-term memory (LSTM) may be more appropriate. However, it was anticipated that both the RNN and the LSTM would result in much larger models, which would be difficult to implement on a low-energy microcontroller. The investigation of the recursive and other deep neural networks (DNN) would only be investigated if higher accuracy was needed. Other models may offer better fitting; however, the limitation to many low-energy platforms is that these generate large models beyond the capabilities of wearables.

Table 1. Model Performance Comparisons

Metric	SVC [15]	Random Forest [15]	SVM-RBF [15]	Ridge Classifier [15]
Training Time	16.92 s	1 min, 22 s	1 hr, 12 min, 33 s	2.11 s
Coefficient Space	0.408 KBytes	2.833 Mbytes	20.232 Mbytes	0.012 KBytes
F1-Score	0.5408	0.6730	0.3823	0.4802
Micro-AUROC	0.83	0.89	0.85	0.79
Macro-AUROC	0.80	0.88	0.87	0.77
Good AUROC	0.90	0.91	0.91	0.82
Fair AUROC	0.66	0.81	0.79	0.61
Poor AUROC	0.85	0.91	0.90	0.87

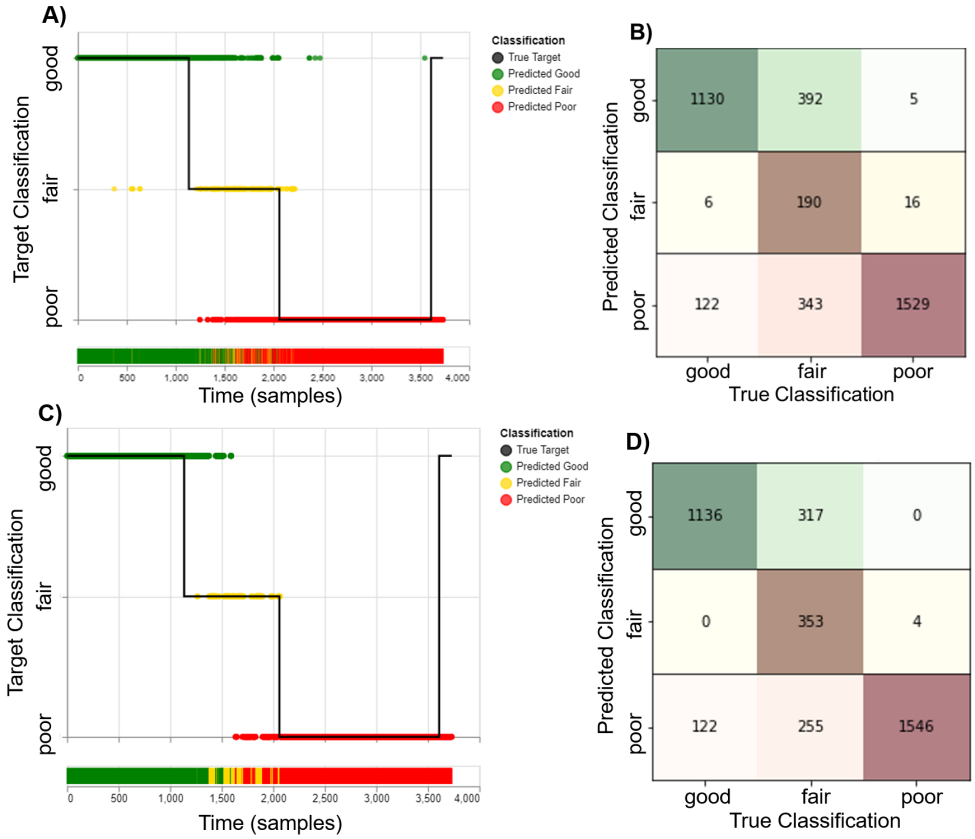

Figure 4. Example data for one subject from the test data is depicted for visualization of algorithm performance through time; the class predictions are color-coded with green representing "good", amber representing to "fair", and red representing "poor"
A & C) The time-series data for the subject with the true target in black and the predicted targets as a scatter plot (top), and a density plot for the predicted classes (bottom), where A) depicts the ouput from the classifier and **C)** depicts the result from applying a mean filter to the output depicted in A)
B & D) The corresponding confusion matrix for the three classes predicted for this subject; **B)** represents confusion matrix for the time-series data in A) and **D)** represents the confusion matrix for the time-series data in B).

The time-series known targets and predictions Fig. 4 demonstrate the relevance of the predictive capabilities in monitoring an individual's progression toward hypovolemia. Additionally, the ability to filter the output of the SVC model increases the performance to a range comparable with the more computationally demanding Random Forest model while maintaining a lightweight implementation. However, post filtering does come at the cost of up to a 25-beat delay in detecting a change between states (e.g. moving from good to fair). Nonetheless, the delay is not expected to interfere with triage and intervention, because the 25-beat delay, less than a minute, is significantly less than the treatment delay of undiagnosed conditions, which is many tens of minutes if not more [10]. As discussed, a linear SVM was used because of the low complexity so this algorithm may easily be realized in a small-resource, low-energy system.

Limitations to implementation exist due to the nature of the controlled conditions of the LBNP study data used. In ranking feature significance, heart rate was consistently ranked the highest contributing feature among classifiers. However, in these studies, the participants are lying down, so their heart rate would be primarily modulated due to the simulated hypovolemia. In reality, heart rate modulates for many reasons which are not specific to hypovolemia, so implementation may not be robust to conditions outside of hypovolemia and result in poor specificity. Additionally, this work utilized LBNP data in which participants were held at each level of hypovolemia for five minutes; this work's approach to optimizing time-series predictions may not generalize well to rapid hemorrhage conditions, although rapid

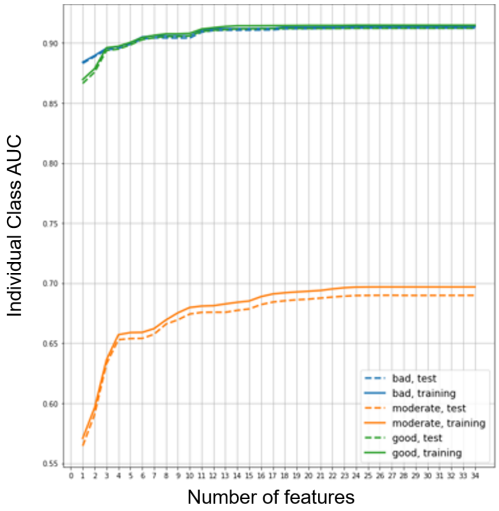

Figure 5. Linear SVC performance by measure of the individual class AUROC with increasing number of features for training.

transitions between states in subject were well characterized.

Additionally, a more comprehensive assessment of the computational cost is needed for determining constraints for implementing on a wearable device. The memory usage discussed herein results from a first-order approximation of the storage utilized by the model once generated. It does not address constraints that may be imposed through RAM usage nor is it a complete estimate of necessary read-only memory (ROM). Data from more subjects will assist in meeting this aim and further refinement of a lightweight algorithm.

Implementation may be further optimized through feature engineering for additional, heart-rate independent features that generate improved performance and by reducing the number of features necessary to segment and inform the classifier.

V. Conclusion

This work demonstrates the feasibility and utility of employing a lightweight machine-learning model for classification of hemorrhagic degree of compromise based on features extracted from a wearable PPG monitor. A linear SVC offers lightweight implementation with consistent performance and permits an analog scale for when an individual enters a compromised state that will be easy to interpret clinically. The model discussed herein achieved a macro-class averaged AUROC of 83% and macro-averaged F1-score of 0.5408. Data from more subjects will offer further refinement through feature engineering and filtering the SVC output, but the work thus far aids in and demonstrates progress toward development of a wearable device to provide real-time monitoring of pathological state.

References

[1] V. A. Convertino, G. Grudic, J. Mulligan, and S. Moulton, "Estimation of individual-specific progression to impending cardiovascular instability using arterial waveforms," *Journal of Applied Physiology*, vol. 115, no. 8, pp. 1196–1202, oct 2013. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/23928113/

[2] V. A. Convertino, J. T. Howard, C. Hinojosa-Laborde, S. Cardin, P. Batchelder, J. Mulligan, G. Z. Grudic, S. L. Moulton, and D. B. MacLeod, "Individual-Specific, Beat-to-beat Trending of Significant Human Blood Loss," *Shock*, vol. 44, no. Supplement 1, pp. 27–32, aug 2015. [Online]. Available: https://journals.lww.com/00024382-201508001-00005

[3] V. A. Convertino, R. W. Techentin, R. J. Poole, A. C. Dacy, A. N. Carlson, S. Cardin, C. R. Haider, D. R. H. III, C. C. Wiggins, M. J. Joyner, T. B. Curry, and O. T. Inan, "AI-enabled advanced development for assessing low circulating blood volume for emergency medical care: Comparison of compensatory reserve machine-learning algorithms," *Sensors*, vol. 22, no. 7, p. 2642, mar 2022.

[4] R. Nadler, V. A. Convertino, S. Gendler, G. Lending, A. M. Lipsky, S. Cardin, A. Lowenthal, and E. Glassberg, "The Value of Noninvasive Measurement of the Compensatory Reserve Index in Monitoring and Triage of Patients Experiencing Minimal Blood Loss," *Shock*, vol. 42, no. 2, pp. 93–98, aug 2014. [Online]. Available: https://journals.lww.com/00024382-201408000-00003

[5] T. E. Schlotman, M. R. Suresh, N. J. Koons, J. T. Howard, A. M. Schiller, S. Cardin, and V. A. Convertino, "Predictors of hemodynamic decompensation in progressive hypovolemia: Compensatory reserve versus heart rate variability," *The journal of trauma and acute care surgery*, vol. 89, no. 2S Suppl 2, pp. S161–S168, aug 2020. [Online]. Available: https://journals.lww.com/jtrauma/Fulltext/2020/08002/Predictors_of_hemodynamic_decompensation_in.24.aspx

[6] R. W. Techentin, C. L. Felton, T. E. Schlotman, B. K. Gilbert, M. J. Joyner, T. B. Curry, V. A. Convertino, D. R. Holmes, and C. R. Haider, "1d convolutional neural networks for estimation of compensatory reserve from blood pressure waveforms," in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2019, pp. 2169–2173.

[7] A. E. Schutte, H. W. Huisman, J. M. van Rooyen, N. T. Malan, and R. Schutte, "Validation of the finometer device for measurement of blood pressure in black women," *Journal of Human Hypertension*, vol. 18, no. 2, pp. 79–84, jan 2004.

[8] D. W. Eeftinck Schattenkerk, J. J. van Lieshout, A. H. van den Meiracker, K. R. Wesseling, S. Blanc, W. Wieling, G. A. van Montfrans, J. J. Settels, K. H. Wesseling, and B. E. Westerhof, "Nexfin Noninvasive Continuous Blood Pressure Validated Against Riva-Rocci/Korotkoff," *American Journal of Hypertension*, vol. 22, no. 4, pp. 378–383, 01 2009. [Online]. Available: https://doi.org/10.1038/ajh.2008.368

[9] D. A. Farcy and T. M. Osborn, *Chapter 46. Classification of Shock*. New York, NY: The McGraw-Hill Companies, 2012. [Online]. Available: accessemergencymedicine.mhmedical.com/content.aspx?aid=55815608

[10] J. E. Bloom, E. Andrew, L. P. Dawson, Z. Nehme, M. Stephenson, D. Anderson, H. Fernando, S. Noaman, S. Cox, C. Milne, W. Chan, D. M. Kaye, K. Smith, and D. Stub, "Incidence and outcomes of nontraumatic shock in adults using emergency medical services in victoria, australia," *JAMA Network Open*, vol. 5, no. 1, p. e2145179, jan 2022.

[11] R. Lyons, *Understanding Digital Signal Processing*. Prentice Hall, 2011. [Online]. Available: https://books.google.com/books?id=arVImAEACAAJ

[12] R. W. Schafer, "What is a savitzky-golay filter? [lecture notes]," *IEEE Signal Processing Magazine*, vol. 28, no. 4, pp. 111–117, 2011.

[13] V. S. Marks, C. L. Felton, R. W. Techentin, B. K. Gilbert, V. A. Convertino, M. J. Joyner, T. B. Curry, D. R. Holmes, and C. R. Haider, "Stockwell transform detector for photoplethysmography signal segmentation," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, 2018, pp. 1239–1243.

[14] L. J. V. Miranda, "PySwarms, a research-toolkit for Particle Swarm Optimization in Python," *Journal of Open Source Software*, vol. 3, 2018. [Online]. Available: https://doi.org/10.21105/joss.00433

[15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.