Short-Term Physiological Forecasting with Adaptive Covariance Matrix Estimation

1st Cem Okan Yaldiz
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

2nd Onur Selim Kilic
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

3rd Omer T. Inan
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA, USA

***Abstract*—Short-term physiological forecasting holds promise for applications requiring capture of rapid changes in physiological signals. Modeling these transient dynamics may enable the quantification of subtle physiological changes that can be critical in many real-time or sensitive contexts, such as assessing acute stress or closed-loop resuscitation. In this paper, we demonstrate the effectiveness of a time-varying autoregressive Kalman filter-based framework for short-term forecasting of physiological features. To eliminate the need for manual hyperparameter tuning, we integrate an adaptive mechanism that dynamically estimates the process and measurement noise parameters of the Kalman filter. Our results demonstrate that the proposed method outperforms baseline models by approximately 1 ms in predicting the next heartbeat's pre-ejection period and by 2–3 ms in predicting the next heartbeat's left ventricular ejection period. Moreover, we demonstrate that it can achieve this improved performance without hyperparameter tuning. This work provides a robust forecasting framework for tracking physiological features and generating new ones to capture short-term physiological variations.**

***Index Terms*—Autoregressive modeling, Kalman filter, physiological forecasting, seismocardiogram, electrocardiogram, adaptive estimation**

I. Introduction

Time series forecasting for physiological signals involves predicting future values of biosignals, or biosignal features, based on their historical temporal patterns. This enables modeling of short-term (i.e., millisecond-level) and long-term physiological changes, with applications in early anomaly detection, personalized health monitoring, and data-driven medical decision-making. Various forecasting methods, including autoregressive and deep learning models, have been explored in healthcare for diagnostic and prognostic purposes [1]. Most of the recent efforts focus on ECG waveform prediction [2]–[4] or certain features such as heart rate dynamics [5]. However, in the context of physiological forecasting, broader evaluations across diverse features derived from multiple sensing modalities are still needed to more effectively capture multimodal relationships. Additionally, although these features have been widely applied in various domains [6], [7], short-term forecasting at fine temporal resolutions (e.g., millisecond-level) remains underexplored.

This work was supported in part by the Office of Naval Research through Grant Number N000142512219.

To address this gap, [8] evaluated several algorithms for short-term cardiac feature forecasting, specifically targeting aortic opening and closing events, which have been demonstrated in the existing literature to provide value in assessing left ventricular function, ejection fraction, and blood volume decompensation status [9]–[11]. It has been proposed that such short-term forecasting not only improves temporal resolution but also enables the derivation of new features representing the model's belief based on historical data [8]. These features can be leveraged to quantify discrepancies between predictions and actual measurements, facilitating the development of new data-driven approaches for fine-grained physiological assessments. Among the methods examined, time-varying autoregressive Kalman filter-based models outperformed the others across noisy and clean settings, in both unimodal and multimodal configurations. Although the effects of modalities such as electrocardiogram (ECG) and seismocardiogram (SCG) have been examined [8], the work did not explore additional modalities that are commonly acquired alongside these signals. Additionally, another key limitation of [8] is the reliance on random search for tuning Kalman filter parameters—specifically, the process and measurement noise matrices, which is computationally expensive.

Building on this foundation, our study evaluates these forecasting models using a more diverse feature set obtained from three sensing modalities: ECG, SCG, and photoplethysmogram (PPG). To address the limitations associated with manual tuning of Kalman filter parameters, we introduce an adaptive update mechanism that dynamically adjusts the process and measurement noise parameters in real-time based on incoming data, eliminating the computationally expensive parameter tuning procedure while maintaining predictive performance.

In this paper, our key contributions are as follows:

- We validate the performance of time-varying autoregressive Kalman filter-based forecasters with multimodal inputs on a more diverse physiological feature set.
- We propose an adaptive mechanism for updating process and measurement noise parameters, removing the need for hyperparameter tuning and adding personalization.

With these contributions, we envision that this work paves the way towards future efforts which can leverage this short-

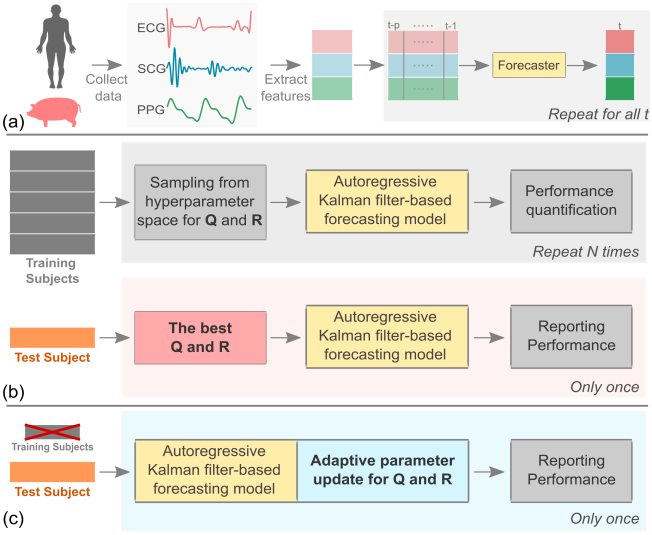

Fig. 1. Diagrams depicting the processes of feature forecasting and parameter optimization. (a) Pipeline from data collection to feature forecasting. (b) Diagram showing model evaluation via hyperparameter tuning with leave-one-subject-out cross-validation. (c) Diagram illustrating model evaluation using adaptive parameter estimation.

time forecasting capability to enable closed-loop interventions in response to rapid physiological shifts, including adaptive resuscitation protocols in response to changing left ventricular function and volume status.

II. Methods

A. Dataset

The dataset described in [8], [11]–[13] consists of ECG, PPG, and SCG recordings from six catheterized pigs undergoing controlled blood volume depletion to simulate hemorrhage. Following the preprocessing methods outlined in [8], [13], eleven features were extracted as per the approach in [11]. These features encompass both amplitude- and timing-based metrics, specifically: heart rate (HR); time-domain heart rate variability (HRV); HRV derived via the Poincaré method; frequency-domain HRV; PPG amplitude range; pulse arrival time (PAT) and its normalized form; pulse transit time (PTT); pleth variability index (PVI); pre-ejection period (PEP); and left ventricular ejection time (LVET). For a more detailed explanation of these features, readers are referred to [11].

B. Autoregressive Modeling of Feature Forecasting

A single-feature autoregressive model of order p, i.e., $AR(p)$, is expressed as

$$F_t^X = \phi_1^X F_{t-1}^X + \phi_2^X F_{t-2}^X + \cdots + \phi_p^X F_{t-p}^X + w_t, \quad (1)$$

where F refers to *feature* with X denoting the specific feature of interest to forecast, w_t is the process noise and ϕ_i, for $i = 1, 2, \ldots, p$ denote the autoregressive coefficients corresponding to the respective lagged terms. When dealing with multiple features, the model must be extended to incorporate cross-modal interactions.

C. Time-varying Autoregressive Kalman Filter Modeling

For next-feature forecasting, we adopt the time-varying autoregressive Kalman filter-based models introduced in [8]. Specifically, we use the unimodal "TV-SF-AR" (SF-KF) and the multimodal "TV-MF-AR" (MF-KF) models, which have demonstrated superiority in estimating next-beat aortic opening (AO) and aortic closing (AC) timings [8].

The state transition equation of a Kalman filter without external input is $\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \mathbf{w}_t$, where $\mathbf{x}_t$ is the state vector, $\mathbf{A}_t$ is the state transition matrix, and $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t)$ represents the process noise at time t. The measurement equation of a Kalman filter is $\mathbf{z}_t = \mathbf{C}_t \mathbf{x}_t + \mathbf{v}_t$, where $\mathbf{z}_t$ is the measurement vector, $\mathbf{C}_t$ is the measurement matrix and $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$ represents the measurement noise.

1) Single-Feature Kalman Filter (SF-KF): Following [8], this model follows the AR(p) process as in (1). We set $\mathbf{C}_t^X = \mathbf{F}_t^X = [F_{t-p}^X, \ldots, F_{t-1}^X]$ and $\mathbf{x}_t^X = [\phi_p^X, \ldots, \phi_1^X]^\top$. The state transition matrix $\mathbf{A}_t^X \in \mathbb{R}^{p \times p}$ and measurement $z_t^X \in \mathbb{R}$ are $\mathbf{I}_{p+q}$ and F_t^X for each target modality X. For a given feature X, the model relies solely on the historical values of X and updates the autoregressive weights ϕ_i^X, $i = 1, \ldots, p$, as defined in equation (1), enabling the model to adapt quickly to physiological changes by updating its parameters in real time.

2) Multi-Feature Kalman Filter (MF-KF): Unlike [8], which utilized only three features, our approach leverages a broader set of 11 features. Let $\mathcal{S} = \{X_1^{\mathcal{S}}, \ldots, X_{|\mathcal{S}|}^{\mathcal{S}}\}$ and $\mathcal{T} = \{X_1^{\mathcal{T}}, \ldots, X_{|\mathcal{T}|}^{\mathcal{T}}\}$ denote the sets of source and target features, respectively, with $|\mathcal{S}|$ and $|\mathcal{T}|$ representing their cardinalities. Target features are the variables we seek to estimate, whereas source features are included in the state to predict these targets. This distinction enhances the model's generalizability to different feature sets and constraints. Additionally, adjusting these features enables us to explore their impact on forecasting performance. The state vector $\mathbf{x}_t \in \mathbb{R}^{|\mathcal{S}|^2 p}$, encoding all auto- and cross-regressive weights between the source features can be written as $\mathbf{x}_t = \left[\mathbf{\Phi}^{X_1^{\mathcal{S}} \to X_1^{\mathcal{S}}}, \ldots, \mathbf{\Phi}^{X_{|\mathcal{S}|}^{\mathcal{S}} \to X_1^{\mathcal{S}}}, \ldots, \mathbf{\Phi}^{X_1^{\mathcal{S}} \to X_{|\mathcal{S}|}^{\mathcal{S}}}, \ldots, \mathbf{\Phi}^{X_{|\mathcal{S}|}^{\mathcal{S}} \to X_{|\mathcal{S}|}^{\mathcal{S}}}\right]^\top$ Each $\mathbf{\Phi}^{X_i^{\mathcal{S}} \to X_j^{\mathcal{S}}}$, $i, j = 1, 2, \ldots, |\mathcal{S}|$, contains p lag coefficients representing the influence of feature $X_i^{\mathcal{S}}$ on $X_j^{\mathcal{S}}$. The state transition matrix is taken as identity: $\mathbf{A}_t = \mathbf{I} \in \mathbb{R}^{|\mathcal{S}|^2 p \times |\mathcal{S}|^2 p}$. The measurement vector $\mathbf{z}_t = \mathbf{F}_t^{\mathcal{T}}$ consists of observations from the target features $\mathcal{T}$, and the corresponding measurement matrix $\mathbf{C}_t$ is given by:

$$\mathbf{C}_t = \begin{bmatrix} \mathbf{F}_t^{X_1^{\mathcal{T}}} & \mathbf{F}_t^{X_2^{\mathcal{T}}} & \ldots & \mathbf{F}_t^{X_{|\mathcal{T}|}^{\mathcal{T}}} & \mathbf{0}_W & \ldots & \mathbf{0}_W \\ \mathbf{0}_W & \mathbf{F}_t^{X_1^{\mathcal{T}}} & \mathbf{F}_t^{X_2^{\mathcal{T}}} & \ldots & \mathbf{F}_t^{X_{|\mathcal{T}|}^{\mathcal{T}}} & \mathbf{0}_W & \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0}_W & \ldots & \mathbf{0}_W & \mathbf{F}_t^{X_1^{\mathcal{T}}} & \mathbf{F}_t^{X_2^{\mathcal{T}}} & \ldots & \mathbf{F}_t^{X_{|\mathcal{T}|}^{\mathcal{T}}} \end{bmatrix}$$

where each $\mathbf{F}_t^{X_i}$ contains lagged observations from feature X_i, and $\mathbf{0}_W$ denotes a $1 \times W$ zero vector, with $W = |\mathcal{T}|p$, matching the dimension of the stacked $\mathbf{F}_p^{X_i}$ blocks. Note that this mathematical structure assumes, without loss of generality, that the first $|\mathcal{T}|$ elements of $\mathcal{S}$ correspond exactly to $\mathcal{T}$. In

other words, each target's own history is utilized to inform its prediction, which is then enhanced by information from other features. The inclusion of multimodality is the key distinction from *SF-KF*, enhancing both accuracy and robustness to noise by incorporating multiple sources of information.

D. Scenarios Examined

We examine both clean and noisy feature scenarios. In the noisy scenarios, following [8], we model the noise η_t^X as a mixture of Gaussian distributions $\mathcal{N}(0, \sigma^2)$, $\mathcal{N}(0, (3\sigma)^2)$, and $\mathcal{N}(0, (5\sigma)^2)$, each occurring with equal probability (i.e., one-third). This noise is added randomly to 25% of the data to simulate real-world conditions with occasional measurement errors. The parameter σ is determined by computing the moving standard deviation of the data. Our analysis focuses specifically on the target features PEP and LVET, as they are closely associated with AO and AC—the targets in [8]. The scenarios investigated is described below.

1) Naive: As a baseline for short-term forecasting, we include a naive approach that predicts the next value using the most recent observation. Specifically, it returns $\widehat{F}_t^X = F_{t-1}^X$ where X denotes any feature of interest. Its noisy counterpart is referred to as *N-Naive*.

2) Rolling Average (RA): As another baseline model, we include RA with parameter p, returning the average of the last p observations as the forecast, i.e., $\widehat{F}_t^X = (1/p)\sum_{i=1}^{p} F_{t-i}^X$. Its noisy counterpart is referred to as *N-RA*.

3) SF-KF: Two separate models are used—one for forecasting PEP and another for LVET—based on the previously described SF-KF, each relying solely on its own historical data for prediction. Its noisy counterpart is referred to as *N SF-KF*.

4) MF-KF: The previously introduced MF-KF model is applied to forecast both PEP and LVET, leveraging all 11 available features. The observation vector $\mathbf{z}$ includes only the observed values of PEP and LVET. We conduct two experiments involving noise inclusion: one where noise is partially added—applied only to SCG-derived features (e.g., PEP, LVET, and PTT)—and referred to as *PN MF-KF*, and another where noise is added to all features, called *N MF-KF*. The former simulates scenarios where only SCG data is affected by noise, while the latter represents conditions in which data from all devices are corrupted by noise.

E. Adaptive Estimation of Covariance Matrix Parameters to Remove Hyperparameter Tuning

In [8], the process noise and measurement noise covariance matrices, $\mathbf{Q}_t$ and $\mathbf{R}_t$, were tuned for each test subject j using cross-validation on the remaining subjects ($i = 1, \ldots, N;\ i \neq j$), requiring an extensive search over the hyperparameter space. Although random search was used to reduce computational burden, it remains prone to suboptimal solutions due to its inherent randomness. Moreover, tuning based on population-level data may not yield parameters optimal for the individual test subject. To address these limitations, we propose an adaptive strategy that estimates $\mathbf{Q}_t$ and $\mathbf{R}_t$ on the fly, based solely on incoming data, eliminating the need for manual or population-based tuning.

TABLE I
PEP AND LVET FORECAST ERRORS ACROSS ALL PIGS UNDER VARIOUS SCENARIOS

Scenario		PEP RMSE (ms)		LVET RMSE (ms)	
		Mean	Std	Mean	Std
Clean	Naive	3.693	2.126	10.649	3.876
	RA	3.564	1.447	9.584	2.646
	SF-KF	2.678	1.914	8.117	3.879
	MF-KF	**2.493**	1.754	**7.674**	3.697
Noisy	N-Naive	6.562	1.560	18.075	2.295
	N-RA	4.547	1.369	13.349	2.350
	N SF-KF	5.477	1.332	14.672	2.363
	PN MF-KF	4.815	1.283	13.782	2.581
	N MF-KF	**4.242**	1.431	**12.346**	2.933

At every time step we first compute the error $\widehat{\delta}_t = \widehat{\mathbf{x}}_{t|t} - \mathbf{A}_t\widehat{\mathbf{x}}_{t-1|t-1}$. This error term, $\widehat{\delta}_t$, is then used to update the process noise covariance matrix $\mathbf{Q}_t$ using

$$\mathbf{Q}_t = \frac{1}{M_Q}\sum_{i=1}^{M_Q}\widehat{\delta}_i\widehat{\delta}_i^\top, \tag{2}$$

where M_Q is a window size, empirically set to 10 to ensure stable estimates without excessive smoothing and short-term memory. While M_Q may be treated as a hyperparameter, tuning this single value is considerably simpler than adjusting all elements of $\mathbf{Q}_t$.

Similarly, for the measurement covariance matrix, we first compute the error $\widehat{\mathbf{v}}_t = \mathbf{z}_t - \mathbf{C}_t\widehat{\mathbf{x}}_{t|t-1}$. $\widehat{\mathbf{v}}_t$ is then used to update the measurement noise covariance matrix $\mathbf{R}_t$ using

$$\mathbf{R}_t = \frac{1}{M_R}\sum_{i=1}^{M_R}\widehat{\mathbf{v}}_i\widehat{\mathbf{v}}_i^\top, \tag{3}$$

where M_R denotes the window size, with a value of 10 chosen empirically, consistent with the reasoning outlined earlier.

III. RESULTS

Table I presents the mean and standard deviations of forecasting errors for PEP and LVET features. It is important to note that the covariance matrices in these scenarios are obtained through hyperparameter tuning, in line with the approach used in [8], to demonstrate the effectiveness of the methodology on a different feature set. In the noise-free scenario, *SF-KF* and *MF-KF* performed effectively on this new feature set. The similar performance of *SF-KF* and *MF-KF* indicates that the benefit of multimodality is not apparent under clean conditions, and a unimodal autoregressive model suffices. However, under noisy conditions, the value of multimodality becomes evident. Specifically, *N MF-KF* outperforms *PN MF-KF*, suggesting that incorporating additional, even non-target, modalities enhances forecasting performance. Overall, both multimodal models outperform the unimodal approach in the presence of noise. Furthermore,

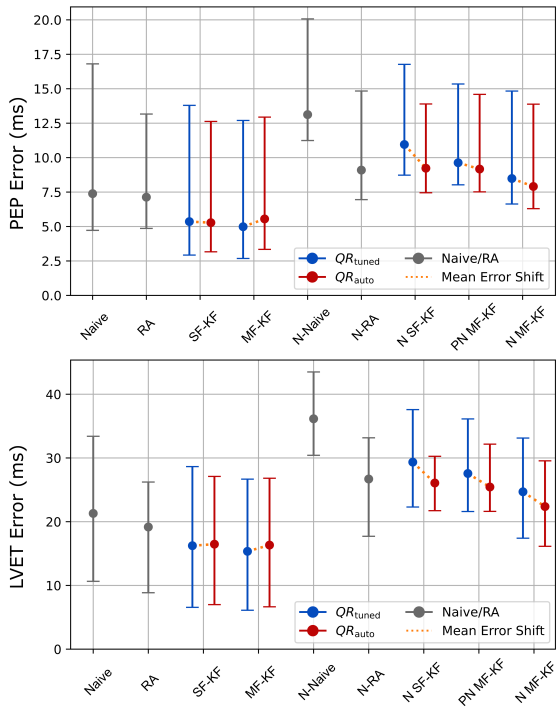

Fig. 2. Comparison of QR_{tuned} and QR_{auto} in Kalman filter-based scenarios for both target features along with performances of baseline models.

autoregressive Kalman filter-based models—particularly *MF-KF*—demonstrate superior performance compared to the baseline models. The RA model outperforms both unimodal and one of the multimodal autoregressive models, likely due to the smoothing effect of averaging, which helps mitigate the impact of occasional noise.

Figure 2 compares *SF-KF* and *MF-KF* with their counterparts that incorporate adaptive covariance matrix estimation. To facilitate comparison, we also show the ranges for the Naive and RA models, which do not incorporate covariance matrices. For both target features (PEP and LVET), the performance of QR_{auto} matches or even exceeds that of QR_{tuned}, indicating that the automated approach achieves comparable accuracy without requiring hyperparameter tuning. Notably, QR_{auto} eliminates the need for tuning entirely.

In multimodal models, the covariance matrix $\mathbf{Q} \in \mathbb{R}^{L \times L}$, where $L = |\mathcal{F}|^2 p$ and $|\mathcal{F}| = 11$ in this study, results in a large number of parameters to tune—L^2 in total. For unimodal models, $L = p$. To reduce complexity, [8] assumed no cross-lag or cross-modal relationships, limiting the estimation to the diagonal elements and reducing the number of tunable parameters to L. However, this simplification potentially omits useful information. Even with this reduction, hyperparameter tuning remains a combinatorial problem. [8] addressed this using a predefined search space and random search, terminating after a fixed number of trials—meaning not all configurations were explored. In contrast, QR_{auto} is fully data-driven, adaptively updating the matrix elements during training without such assumptions or manual effort, thus avoiding these challenges entirely. Furthermore, as the results show, this approach does not lead to any performance degradation. The same conclusion also applies to the measurement noise matrix $\mathbf{R}$.

IV. Conclusion

This work investigated the use of a time-varying autoregressive Kalman filter-based framework for unimodal and multimodal short-term forecasting of physiological features. It has been shown that the proposed model outperforms baseline forecasting methods. Additionally, the impact of incorporating multiple modalities under both clean and noisy conditions was examined. Importantly, our results demonstrated that the adaptive estimation of process noise and measurement noise parameters effectively removes the need for manual hyperparameter tuning. This adaptive approach maintains competitive performance, without relying on a computationally intensive tuning procedure.

References

[1] C. Bui, N. Pham, A. Vo, A. Tran, A. Nguyen, and T. Le, "Time series forecasting for healthcare diagnosis and prognostics with the focus on cardiovascular diseases," in *6th International Conference on the Development of Biomedical Engineering in Vietnam (BME6) 6*, pp. 809–818, Springer, 2018.

[2] H. Zacarias, J. A. L. Marques, V. Felizardo, M. Pourvahab, and N. M. Garcia, "Ecg forecasting system based on long short-term memory," *Bioengineering*, vol. 11, no. 1, p. 89, 2024.

[3] K. R. Prakarsha and G. Sharma, "Time series signal forecasting using artificial neural networks: An application on ecg signal," *Biomedical Signal Processing and Control*, vol. 76, p. 103705, 2022.

[4] M. Ali, P. W. Moore, T. Barkouki, and L. J. Brattain, "Zero-shot forecasting for ecg time series data using generative foundation models," in *2024 IEEE 20th International Conference on Body Sensor Networks (BSN)*, pp. 1–4, IEEE, 2024.

[5] M. Ludwig, K. Hoffmann, S. Endler, A. Asteroth, and J. Wiemeyer, "Measurement, prediction, and control of individual heart rate responses to exercise—basics and options for wearable devices," *Frontiers in physiology*, vol. 9, p. 778, 2018.

[6] B. Rim, N.-J. Sung, S. Min, and M. Hong, "Deep learning in physiological signal data: A survey," *Sensors*, vol. 20, no. 4, p. 969, 2020.

[7] P. J. Bota, C. Wang, A. L. Fred, and H. P. Da Silva, "A review, current challenges, and future possibilities on emotion recognition using machine learning and physiological signals," *IEEE access*, vol. 7, pp. 140990–141020, 2019.

[8] C. O. Yaldiz, D. J. Lin, A. H. Gazi, G. Cestero, C. Chen, B. K. Bracken, A. Winder, S. Lynn, R. Sameni, and O. T. Inan, "Real-time autoregressive forecast of cardiac features for psychophysiological applications," *IEEE Journal of Biomedical and Health Informatics*, 2025.

[9] R. P. Lewis, S. Rittogers, W. Froester, and H. Boudoulas, "A critical review of the systolic time intervals.," *Circulation*, vol. 56, no. 2, pp. 146–158, 1977.

[10] S. S. Ahmed, G. E. Levinson, C. J. Schwartz, and P. O. Ettinger, "Systolic time intervals as measures of the contractile state of the left ventricular myocardium in man," *Circulation*, vol. 46, no. 3, pp. 559–571, 1972.

[11] J. P. Kimball, J. S. Zia, S. An, C. Rolfes, J.-O. Hahn, M. N. Sawka, and O. T. Inan, "Unifying the estimation of blood volume decompensation status in a porcine model of relative and absolute hypovolemia via wearable sensing," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 9, pp. 3351–3360, 2021.

[12] J. Zia, J. Kimball, C. Rolfes, J.-O. Hahn, and O. T. Inan, "Enabling the assessment of trauma-induced hemorrhage via smart wearable systems," *Science advances*, vol. 6, no. 30, p. eabb1708, 2020.

[13] D. J. Lin, A. H. Gazi, J. Kimball, M. Nikbakht, and O. T. Inan, "Real-time seismocardiogram feature extraction using adaptive gaussian mixture models," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 8, pp. 3889–3899, 2023.