Complex Control Systems ISSN 2603-4697 (Online)

Hybrid Artificial Intelligence Model for R-Peak Forecasting in Electrocardiogram Signals

Dobromir Slavov[(1)], Ekaterina Popovska [(2)], and [(3)] Galya Georgieva-Tsaneva

[(1)(2)(3)] Institute of Robotics, Bulgarian Academy of Science, 1113 Sofia, Bulgaria

***Abstract*- Short-term cardio time series offer valuable insight into beat-level, moment-to-moment variations in heart rhythm and autonomic modulation. Beat-to-beat heart-rate sequences derived from electrocardiogram recordings demonstrate high non-stationarity and natural physiological variability, limiting the predictive stability of single-method data-driven approaches at short horizons. This study introduces the initial methodological milestone of a hybrid**

Artificial Intelligence framework for short-horizon heart-rate time-series analysis. The approach integrates rolling temporal representations learned by recurrent neural architectures with a compact set of physiologically informed descriptors extracted from R-peak annotations and derived R-R intervals. The forecasting task is formulated as a five-beat-ahead regression problem using only short preceding beat sequences, targeting embedded real-time monitoring scenarios rather than long-range signal prediction. The work benchmarks a baseline long short-term memory recurrent model against a hybrid feature-fusion extension, emphasizing trend adherence, temporal sensitivity and model stability under restricted context.

The findings confirm that lightweight hybrid fusion can enhance short-horizon responsiveness to transient fluctuations while preserving overall clinical plausibility and computational compactness. The methodology establishes a reproducible benchmark pipeline for short-context cardio time-series forecasting and supports future research extensions toward broader rhythm conditions and real-time AI monitoring pipelines.

***Keywords*—Artificial intelligence, cardio signals, deep learning, heart rate dynamics, hybrid models, nonlinear analysis, time series.**

I. Introduction

Short-term cardiovascular dynamics provide critical insight into the moment-to-moment regulation of heart function and autonomic nervous system activity. Time series derived from electrocardiographic recordings, beat-to-beat RR intervals, and instantaneous heart rate measurements are widely used for assessing physiological state, stress response, and early indicators of cardiac dysfunction [1], [2]. Unlike long-duration ECG monitoring, short cardio time series capture transient beat dynamics within limited observation windows, making their modeling particularly sensitive to signal non-stationarity, noise, and abrupt physiological modulation.

Traditional heart rate variability analysis methods rely on time-domain, frequency-domain, and nonlinear descriptors that remain clinically meaningful due to their interpretability and strong physiological grounding [3]. Metrics such as mean RR interval, variability dispersion, and entropy-based complexity indices have been applied for decades in ambulatory and clinical monitoring. However, these approaches do not learn temporal dependencies and may lose sensitivity to fast beat-level transients when the signal context is short.

Modern neural sequence models, including recurrent neural networks, one-dimensional convolutional networks, and attention-based architectures, have demonstrated the ability to learn representations directly from cardio signals [4], [5]. Their deterministic precision, however, may degrade when applied as standalone predictors on chronologically short windows, where data scarcity increases the likelihood of temporal oversmoothing and instability under transient fluctuations.

Hybrid artificial intelligence architectures attempt to mitigate these issues by integrating domain-informed variability markers extracted from RR dynamics into neural temporal embeddings [6], [7]. These strategies align with the broader scientific shift toward combining model-based signal meaning with data-driven learning, particularly for short-horizon regression formulations, where robustness and temporal sensitivity outweigh model scale or long-window classification.

This work investigates hybrid artificial intelligence models for short-horizon heart-rate time-series modeling. The study formulates a constrained beat-trend prediction task and evaluates a rolling adaptive RR baseline [7], a single-branch recurrent sequence model [4], and a hybrid recurrent model enhanced with standardized physiological variability descriptors [6], [7]. By emphasizing beat-level temporal precision, physiological plausibility, and noise robustness, the proposed framework supports the development of compact and reproducible cardiac AI pipelines suitable for real-time monitoring systems.

II. Data and Preprocessing

A. Dataset Description

The experimental evaluation in this study is conducted using real-world electrocardiographic recordings obtained from the MIT-BIH Normal Sinus Rhythm Database (NSRDB), which is part of the broader MIT-BIH PhysioNet collection [8], [9]. The database contains long-term ECG recordings acquired from subjects exhibiting normal sinus rhythm, making it suitable for the analysis of physiological heart rate dynamics without dominant pathological disturbances.

Each record consists of digitized ECG signals sampled at a frequency of 128 Hz, accompanied by expert-annotated R-peak locations. The availability of precise beat annotations enables reliable extraction of beat-to-beat intervals and subsequent construction of short-term heart rate time series. In this work, a subset of NSRDB records is used to focus on methodological validation, which is consistent

with the objectives of short-horizon prediction and hybrid model evaluation.

B. ECG Signal Processing and R-peak Annotation

Let $x(t)$ denote the discrete-time ECG signal sampled at frequency f_s. For each ECG record, R-peak locations are obtained directly from the provided annotation files, which have been validated and widely used in the literature [9]. This approach avoids potential errors introduced by automatic QRS detection algorithms and ensures reproducibility of the preprocessing pipeline.

Given a sequence of annotated R -peak sample indices

$$\{r_1, r_2, \dots, r_N\},$$

the corresponding R-peak times in seconds are computed as

$$t_i = \frac{r_i}{f_s}$$

This representation forms the basis for deriving RR interval and heart rate series used in subsequent modeling stages.

C. RR Interval and Heart Rate Time Series Construction

The beat-to-beat RR interval series is defined as:

$$RR_i = t_{i+1} - t_i, i = 1, \dots, N-1.$$

D. Windowing and Forecasting Setup

To formulate the short-horizon forecasting task, the cardio series is segmented into overlapping windows using a sliding window approach. For each sample, an input window of length L peaks is used to predict the subsequent H peaks:

$$\mathbf{x}_i = \{R_i, R_{i+1}, \dots, R_{i+L-1}\},$$
$$\mathbf{y}_i = \{R_{i+L}, R_{i+L+1}, \dots, R_{i+L+H-1}\}$$

In this study, the input length is fixed to $L = 30$ peaks, while the prediction horizon is set to $H = 10$ peaks, reflecting a short-horizon forecasting scenario. This formulation is motivated by practical monitoring applications, where near-future cardiac behavior is of greater relevance than long-term prediction [5].

The segmentation is performed chronologically without shuffling, and the resulting samples are divided into training, validation, and test sets using a time-based split. This strategy prevents information leakage across temporal boundaries and ensures that model evaluation reflects realistic deployment conditions.

E. Feature Normalization and Physiological Descriptors

For model training, heart rate input sequences are standardized using statistics computed from the training set only. Given the training mean μ and standard deviation σ, normalized inputs are obtained as

$$\tilde{R}R_i = \frac{RR_i - \mu}{\sigma}.$$

In addition to raw heart rate sequences, a small set of physiologically motivated descriptors is computed for each input window. These features include basic time-domain and short-term variability measures derived from the corresponding RR intervals, such as mean RR interval, standard deviation of RR intervals, and root mean square of successive differences. Such descriptors provide complementary information about shortterm autonomic regulation and have been shown to enhance robustness when combined with data-driven models [6], [11].

The preprocessing pipeline transforms raw ECG recordings into normalized short-term heart rate sequences and associated physiological descriptors suitable for hybrid modeling. By relying on expert-annotated R-peaks and beat-level segmentation, the proposed approach preserves physiologically meaningful temporal dynamics while enabling systematic evaluation of short-horizon prediction models.

III. Methodology

Let the beat-level heart rate time series derived from an ECG recording be denoted as:

$$\mathrm{T} = \{HR_1, HR_2, \dots, HR_T\},$$

where $\mathrm{T}_i \in \mathbb{R}$ represents the time in seconds of occurrence of the i-th R-peak t. The objective of this study is to address a short-horizon forecasting problem, in which the goal is to predict the future in which the moment of occurrence of future R-peaks is estimated based on a limited number of preceding events. Given an input window of length L, the prediction task is defined as:

$$\hat{\mathbf{y}}_t = f_\theta(\mathbf{x}_t),$$

Where:

$$\mathbf{x}_t = \{RR_t, HR_{t+1}, \dots, HR_{t+L-1}\}, \hat{\mathbf{y}}_t = \{\hat{R}R_{t+L}, \dots, \hat{H}R_{t+L+H-1}\},$$

and $f_\theta(\cdot)$ denotes a parameterized prediction model with parameters θ. In this work, the input window length is fixed to $L = 30$ beats, and the prediction horizon is set to $H = 10$ beats. This formulation reflects practical monitoring scenarios, where near-future cardiac behavior is of primary interest, and aligns with prior studies on short-term physiological forecasting [5], [12].

A. Baseline forecasting model

As a reference, a purely data-driven baseline model is implemented to establish a lower bound on prediction performance. The baseline model operates directly on normalized heart rate sequences and does not incorporate explicit physiological descriptors.

The baseline predictor is implemented using a recurrent neural network with long short-term memory (LSTM) units, which are well suited for modeling temporal dependencies in sequential data [13]. Given the input sequence $\tilde{\mathbf{x}}_t$, obtained through z-score normalization, the LSTM computes a hidden representation:

$$\mathbf{h}_t = \mathrm{LSTM}(\tilde{\mathbf{x}}_t),$$

which is subsequently mapped to the predicted output via a fully connected layer:

$$\hat{\mathbf{y}}_t = \mathbf{W}\mathbf{h}_t + \mathbf{b}.$$

The model is trained to minimize the mean squared error (MSE) between predicted and true future heart rate values:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{H}\sum_{i=1}^{H}\left(\hat{H}R_{t+L+i-1} - HR_{t+L+i-1}\right)^2$$

This baseline provides a strong yet interpretable benchmark for evaluating the added value of hybrid modeling strategies.

B. Physiological Feature Extraction

To incorporate domain knowledge into the prediction process, a set of physiologically meaningful descriptors is computed for each input window. These features are derived from the corresponding RR interval sequence and capture short-term variability characteristics of cardiac dynamics [3], [11]. Given an input window $\mathbf{x}_t$, the associated RR interval sequence is computed as:

$$\mathbf{x}_t = \{RR_t, RR_{t+1}, \dots, RR_{t+L-1}\}.$$

From this sequence, the following features are extracted:

- mean RR interval,
- standard deviation of RR intervals (SDNN),
- root mean square of successive differences (RMSSD),
- mean heart rate,
- standard deviation of heart rate.

These features form a low-dimensional physiological descriptor vector:

$$\mathbf{z}_t \in \mathbb{R}^5$$

which provides complementary information to the raw heart rate sequence and reflects autonomic regulation mechanisms.

C. Hybrid Model Architecture

The proposed hybrid model combines data-driven temporal modeling with physiologically motivated information through a mid-level feature fusion strategy [6], [14]. The architecture consists of two parallel branches: a sequence modeling branch that captures temporal dependencies in cardiac rhythm dynamics, and a physiological feature branch that encodes short-term variability characteristics derived from R–R intervals.

In the sequence branch, the normalized sequence of recent R–R intervals is processed by an LSTM network to produce a latent temporal embedding:

$$\mathbf{h}_t = \text{LSTM}\left(\tilde{RR}_t\right),$$

where $\tilde{RR}_t$ denotes the standardized input window of RR intervals.

In the physiological branch, the feature vector $\tilde{\mathbf{z}}_t$, obtained via standardization, is transformed using a fully connected layer:

$$\mathbf{g}_t = \phi(\mathbf{W}_z\tilde{\mathbf{z}}_t + \mathbf{b}_z)$$

where $\phi(\cdot)$ denotes a nonlinear activation function. The two representations are then concatenated:

$$\mathbf{u}_t = [\mathbf{h}_t \| \mathbf{g}_t]$$

and subsequently processed by dense layers to generate the final prediction of future RR intervals:

$$\hat{\mathbf{y}}_t = \psi(\mathbf{W}_u\mathbf{u}_t + \mathbf{b}_u).$$

This hybrid formulation allows the model to jointly exploit temporal patterns learned from data and structured physiological information, improving robustness in short and noisy time series settings.

D. Training Strategy

Both baseline and hybrid models are trained using the Adam optimization algorithm with early stopping based on validation loss [15]. To avoid information leakage, a chronological split is employed, dividing the dataset into training, validation, and test subsets. All normalization parameters are computed exclusively on the training set and subsequently applied to validation and test data.

Model performance is evaluated using mean absolute error (MAE) and root mean square error (RMSE), expressed in beats per minute, which preserve direct physiological interpretability.

The proposed methodology formulates short-horizon R-peak time prediction as a multistep regression problem and evaluates both purely data-driven and hybrid AI models. By integrating physiologically motivated descriptors with neural network-based temporal modeling, the hybrid approach aims to achieve improved robustness and interpretability in short-term RR time series analysis.

IV. Experimental Setup

A. Dataset Overview

All experiments use ECG records from the MIT-BIH Normal Sinus Rhythm Database as described in Section 2. The models are trained and evaluated on chronologically segmented R–R interval sequences for short-horizon regression tasks .

B. Signal Preprocessing and Segmentation

The ECG signals are first bandpass-filtered using a zero-phase Butterworth filter to suppress baseline wander and high-frequency noise. R-peak annotations provided by the database are used to derive RR intervals, defined as the temporal differences between consecutive R-peaks.

Based on the sequence of R-peaks, the inter-beat intervals (RR intervals) are computed as the difference between the occurrence times of two consecutive R-peaks. The resulting RR intervals form the input time series for the predictive models.

The resulting RR series are segmented into overlapping windows of fixed length $L = 30$ beats, followed by a prediction horizon of $H = 10$ beats. A sliding-window approach with a stride of one beat is applied, yielding a large number of training samples while preserving temporal continuity. Each input window is normalized using z-score normalization. Both metrics are computed over the entire prediction horizon and averaged across all test samples. Errors are reported in bpm, rather than percentages, to preserve clinical relevance and facilitate comparison with existing studies [18].

C. Experimental Protocol

Each experiment is repeated independently for each ECG record. Reported results correspond to the mean performance across all test segments for a given record. To

ensure reproducibility, random seeds are fixed for weight initialization and data shuffling. In addition to quantitative evaluation, qualitative analysis is performed by visually comparing predicted and ground-truth heart rate trajectories over selected short segments, highlighting the models' ability to capture rapid local fluctuations.

V. Results and Discussion

This section presents and analyzes the experimental results obtained from the proposed short-horizon heart rate (HR) prediction framework. The evaluation focuses on the ability of recurrent neural architectures to predict future RR intervals and the appearance times of upcoming R-peaks, formulated as a 10 -beat-ahead regression task on chronologically segmented, expert-annotated R-R interval sequences extracted from ECG recordings.

A. ECG-Derived R-Peaks Dynamics

Figure 1 illustrates a representative 6-second ECG segment from record 16 265 with annotated R-peaks. The signal exhibits stable morphology and clearly identifiable ventricular depolarization complexes, confirming the reliability of the provided annotations and pre-processing pipeline.

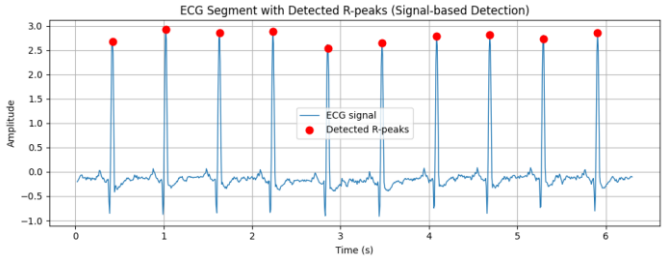

Figure 1. *ECG signal segment with detected R-peak locations*.

Based on the detected R-peaks, a sequence of R–R intervals was constructed by computing the temporal differences between consecutive peaks. The first 50 R–R intervals used as model input are shown in Figure 2. The sequence demonstrates moderate beat-to-beat variability, reflecting natural short-term fluctuations in cardiac rhythm even under normal sinus conditions. This non-stationary behavior motivates the use of adaptive and nonlinear prediction models rather than simple constant-rate assumptions.

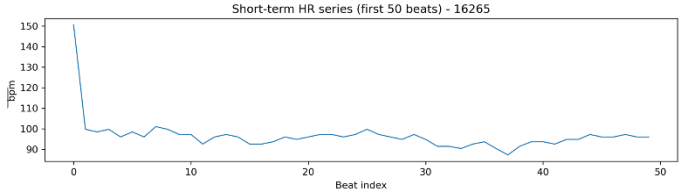

Figure 2. *Short-term beat-to-beat heart rate trajectory for the first 50 beats.*

B. Short-Horizon HR Prediction Performance

The prediction task was formulated as a multistep forecasting problem, where an input window of 30 consecutive R–R intervals was used to predict the occurrence times of the next 10 cardiac cycles. Two approaches were evaluated:

- an adaptive rolling R–R baseline model, and
- an LSTM-based recurrent neural network trained to predict future R–R intervals.

Figure 3 presents a representative example of short-horizon R-peak time prediction (adaptive rolling R–R baseline model). The true future R-peaks are shown together with the predicted peak locations obtained from both models. For visualization purposes, predicted times are projected onto the ECG waveform by mapping them to the nearest detected peak amplitudes.

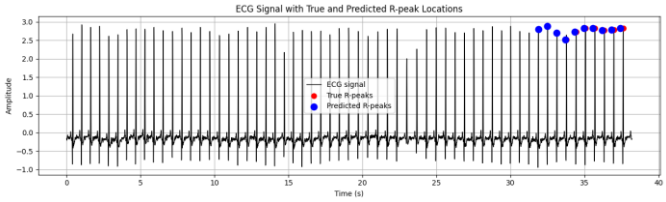

Figure 3. *True and predicted future R-peak locations for an ECG segment (baseline model).*

Figure 4 compares the true future R-peak locations with the predictions obtained from the adaptive rolling RR baseline and the LSTM-based model for a representative ECG segment. While both approaches successfully follow the general rhythm trend, the adaptive RR model exhibits increasing temporal deviation as the prediction horizon advances. In contrast, the LSTM model produces predictions that remain closely aligned with the true R-peak positions, particularly for the later beats. This visual comparison confirms the quantitative results, demonstrating the superior robustness of the LSTM model in short-horizon R-peak time forecasting.

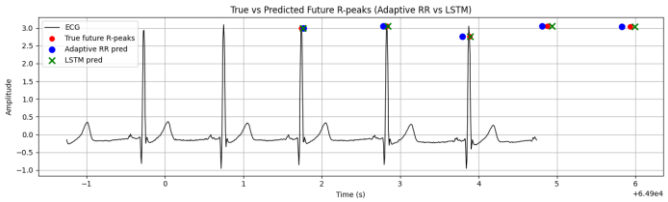

Figure 4. *Comparison of true and predicted future R-peak locations using Adaptive RR and Hybrid LSTM models.*

C. Model Comparison

Table I presents a quantitative comparison of three forecasting approaches for short-horizon R-peak time prediction over a 10-beat horizon: an adaptive rolling RR baseline, a data-driven LSTM model, and a hybrid LSTM model augmented with physiological variability features. Prediction accuracy is assessed using mean absolute error (MAE) and root mean square error (RMSE), expressed in

milliseconds, providing a direct measure of temporal prediction precision.

The adaptive rolling RR baseline achieved an MAE of 39.06 ms and an RMSE of 42.65 ms. These results indicate reasonable short-term extrapolation performance under normal sinus rhythm conditions, but also reflect the inherent limitations of relying primarily on the most recent RR interval. As expected, this simple approach is prone to gradual error accumulation in multistep prediction scenarios.

Introducing a purely data-driven LSTM model leads to a moderate improvement in prediction accuracy, reducing the MAE to 34.96 ms and the RMSE to 37.02 ms. This improvement demonstrates the benefit of recurrent neural modeling, which can exploit temporal dependencies across multiple preceding cardiac cycles rather than extrapolating from a single interval.

The hybrid LSTM model, which integrates temporal representations learned by the LSTM with physiologically motivated RR variability descriptors, achieves a substantial further reduction in prediction error. Specifically, the hybrid approach attains an MAE of 15.61 ms and an RMSE of 24.10 ms, representing more than a twofold improvement in MAE compared to the adaptive baseline. This pronounced performance gain highlights the complementary role of physiological features in constraining and stabilizing the learning process, particularly in short and noisy cardiac time series.

The comparison confirms that while recurrent neural networks improve short-horizon R-peak time prediction relative to simple adaptive baselines, the integration of physiological domain knowledge provides an additional and significant benefit. The hybrid architecture demonstrates enhanced robustness and precision, especially in scenarios characterized by limited temporal context, making it well suited for real-time cardiac monitoring applications. The hybrid model demonstrates enhanced temporal precision and improved short-horizon stability compared to baseline extrapolation and standalone LSTM trend predictors, while the overall beat-trend modeling remains clinically plausible and computationally compact.

In relative terms, the hybrid LSTM model reduces the mean absolute error by 60.0% compared to the adaptive rolling RR baseline and by 55.3% compared to the purely data-driven LSTM model. In addition, the hybrid approach achieves a 43.5% reduction in RMSE relative to the baseline, confirming its superior robustness in short-horizon R-peak time prediction.

Table I. *Model comparison for short-horizon R-peak time prediction (H = 10 beats)*

Model	MAE (ms)	RMSE (ms)
Baseline adaptive RR	39.06	42.65
LSTM	34.96	37.02
Hybrid LSTM + physiological features	15.61	24.10

Across all evaluated test segments, both the adaptive rolling RR baseline and the LSTM-based model achieve prediction errors within a physiologically acceptable range for short-horizon R-peak time forecasting. However, a clear performance difference is observed between the two approaches. While the adaptive RR model provides reasonable short-term extrapolation, the LSTM model consistently achieves substantially lower errors, particularly for the first few predicted beats.

Horizon-wise analysis reveals that the LSTM model exhibits significantly improved accuracy already at the earliest prediction steps (H = 1–3), indicating a higher sensitivity to immediate temporal context. This behavior reflects the ability of recurrent neural networks to exploit dependencies across multiple previous cardiac cycles, rather than relying solely on the most recent RR interval. In contrast, the adaptive RR baseline is more susceptible to gradual error accumulation as the prediction horizon increases.

D. Feature Fusion Implications for Short-Context Cardio Time Series

Previous studies have primarily focused either on classical HRV analysis using statistical descriptors [3], [18] or on deep learning approaches trained on long-duration recordings [8], [20]. In contrast, the present study explicitly targets short-horizon prediction, where limited temporal context poses significant modeling challenges.

The obtained results align with recent trends toward physiology-aware and physics-informed machine learning in biomedical applications. Unlike large transformer-based architectures, which typically require extensive training data, the proposed hybrid approach remains computationally lightweight and suitable for real-time cardiovascular monitoring systems.

E. Limitations and Future Work

Several limitations should be acknowledged. First, the experiments were conducted exclusively on normal sinus rhythm recordings; extending the analysis to pathological conditions may reveal additional challenges and opportunities for improvement. Second, the current study focuses solely on ECG-derived R-peak sequences.

Future work will explore adaptive and physiology-aware feature fusion strategies, as well as attention-based mechanisms to dynamically emphasize relevant temporal patterns. Additionally, multimodal extensions incorporating complementary signals, such as photoplethysmography or respiration, may further enhance robustness and prediction accuracy.

VI. Conclusion

This study investigated short-horizon prediction of future R-peak occurrence times using sequential models applied to ECG-derived R–R interval series. Focusing on limited observation windows and a 10-beat prediction horizon, the work addressed a practically relevant yet challenging problem characterized by strong non-stationarity and beat-to-beat variability.

A comparative evaluation between an adaptive rolling RR baseline and an LSTM-based recurrent neural model demonstrated that both approaches yield physiologically meaningful prediction errors, while the LSTM model achieves a substantial reduction in MAE and RMSE. These results highlight the advantage of recurrent temporal modeling in capturing short-term rhythm dynamics and mitigating error accumulation in multistep prediction tasks.

Beyond improved accuracy, the proposed framework emphasizes computational efficiency and interpretability, making it suitable for real-time cardiovascular monitoring applications. The presented methodology provides a reproducible benchmark for short-horizon R-peak time forecasting and supports the development of next-generation physiology-aware artificial intelligence systems for continuous cardiac monitoring.

ACKNOWLEDGMENT

This research was funded by the National Science Fund of Bulgaria (scientific project “Research on AI-based methods for modeling, analyzing and forecasting biomedical time series by applying a digital twin concept”), Grant Number КП-06-Н97/12, 10 December 2025.

References

[1] Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology, Heart rate variability: Standards of measurement, physiological interpretation, and clinical use, Circulation, vol. 93, no. 5, 1996, pp. 1043–1065.

[2] Clifford, G.D., Azuaje, F., McSharry, P., Advanced Methods and Tools for ECG Data Analysis, London, UK: Artech House, ch. 2, 2006, pp. 21–52.

[3] Shaffer, F., Ginsberg, J.P., An overview of heart rate variability metrics and norms, Frontiers in Public Health, vol. 5, 2017, pp. 1–17.

[4] Faust, O., et al., Deep learning for healthcare applications based on physiological signals: A review, Computer Methods and Programs in Biomedicine, vol. 161, 2018, pp. 1–13.

[5] Lim, B., Zohren, S., Roberts, S., Time-series forecasting with deep learning: A survey, Philosophical Transactions of the Royal Society A, vol. 379, no. 2194, 2021, pp. 1–25.

[6] Acharya, U.R., et al., Automated characterization of coronary artery disease using nonlinear features extracted from heart rate signals, Biomedical Signal Processing and Control, vol. 31, 2017, pp. 213–224.

[7] Natarajan, A., et al., A wide and deep learning approach for robust prediction of cardiovascular events, Proc. IEEE Conf. on Engineering in Medicine and Biology Society (EMBS), Honolulu, HI, USA, 2018, pp. 273–276.

[8] Goldberger, A.L., et al., *PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals*, Circulation, vol. 101, no. 23, 2000, pp. e215–e220.

[9] Moody, G.B., Mark, R.G., *The MIT-BIH Arrhythmia Database*, IEEE Engineering in Medicine and Biology Magazine, vol. 20, no. 3, 2001, pp. 45–50.

[10] Berntson, G.G., et al., *Heart rate variability: Origins, methods,* and interpretive caveats, Psychophysiology, vol. 34, no. 6, 1997, pp. 623–648.

[11] Acharya, U.R., et al., Heart rate variability: A review, Medical & Biological Engineering & Computing, vol. 44, no. 12, 2006, pp. 1031–1051.

[12] Box, G.E.P., Jenkins, G.M., Reinsel, G.C., Time Series Analysis: Forecasting and Control, Hoboken, NJ, USA: Wiley, ch. 3, 2015, pp. 89–132.

[13] Hochreiter, S., Schmidhuber, J., Long short-term memory, Neural Computation, vol. 9, no. 8, 1997, pp. 1735–1780.

[14] Bishop, C.M., Pattern Recognition and Machine Learning, New York, NY, USA: Springer, ch. 14, 2006, pp. 631–660.

[15] Kingma, D.P., Ba, J., Adam: A method for stochastic optimization, Proc. Int. Conf. on Learning Representations (ICLR), San Diego, CA, USA, 2015.

[16] Goldberger, A.L., Amaral, L.A.N., Glass, L., et al., *PhysioBank, PhysioToolkit, and PhysioNet*, Circulation, vol. 101, no. 23, 2000, pp. e215–e220.

[17] Hyndman, R.J., Athanasopoulos, G., *Forecasting: Principles and Practice*, Melbourne, Australia: OTexts, ch. 2, 2021, pp. 45–78.

[18] Shaffer, F., Ginsberg, J.P., *An overview of heart rate variability metrics and norms*, Frontiers in Public Health, vol. 5, 2017, pp. 1–17.