*Simen Owen Academic Proceedings Series*

*Vol. 3 2026*

*Article*

# Predictive Modeling of Hypotension during General Anesthesia Using Machine Learning Algorithms on High-Frequency Physiological Time-Series Data

**William Pan [1,*]**

[1] Lowell High School, San Francisco, USA
* Correspondence: William Pan, Lowell High School, San Francisco, USA

**Abstract:** Intraoperative hypotension during general anesthesia represents a grave clinical complication that can result in severe adverse patient outcomes, including acute kidney injury, myocardial infarction, and increased postoperative mortality. Although prompt forecasting is absolutely essential for timely clinical intervention, conventional monitoring approaches often rely on periodic, intermittent measurements, which inherently delays the detection of critical hemodynamic deterioration. Furthermore, current predictive computational frameworks struggle with instantaneous forecasting and consistently fail to fully exploit the rich, dynamic information embedded within high-frequency time-series signals. These existing models also frequently suffer from poor explainability and limited generalizability when applied across heterogeneous patient cohorts in real-world clinical settings. To systematically address these critical issues, this study introduces a novel joint deep learning architecture integrating Long Short-Term Memory (LSTM) networks and advanced attention mechanisms to accurately forecast anesthetic hypotension utilizing high-frequency physiological parameters. The proposed methodological framework exhibits highly robust generalization capabilities, sustaining remarkably stable predictive performance when rigorously validated on several independent, external hospital datasets. On the primary evaluation cohort, the integrated model achieved an impressive accuracy of 0.87 ±0.02 and an Area Under the Curve (AUC) of 0.91 ± 0.01, statistically outperforming traditional baseline models (p-value < 0.05) across key metrics including precision, recall, and F1-score. Consequently, this advanced model significantly refines real-time intraoperative hypotension prediction, offering a highly dependable, accurate, and explainable instrument for clinical decision support with broad applicability in diverse, fast-paced healthcare environments.

**Keywords:** hypotension prediction; general anesthesia; deep learning; time-series analysis; physiological monitoring; interpretable AI

Received: 27 February 2026
Revised: 22 April 2026
Accepted: 06 May 2026
Published: 10 May 2026



## 1. Introduction

Intraoperative hypotension (IOH), frequently characterized by a mean arterial pressure (MAP) dropping below 65 mmHg, is a prevalent clinical challenge affecting a substantial proportion of patients undergoing surgical procedures [1]. IOH can impact up to 86% of patients in certain surgical cohorts, carrying a strong association with adverse post-operative events. Large-scale observational studies have established that prolonged periods of MAP < 65 mmHg elevate the risks of acute kidney injury, myocardial injury, and 30-day mortality. Consequently, proactive forecasting and prompt therapeutic intervention are vital to mitigate these complications. Nevertheless, traditional clinical blood pressure assessment relies predominantly on intermittent measurements, which often restricts the timely identification and management of hypotensive episodes.

The availability of high-frequency physiological time-series data—comprising continuous streams of vital signs such as heart rate, blood pressure, and arterial oxygen saturation—offers a promising pathway for forecasting IOH. These datasets capture transient physiological shifts that manifest minutes before clinical hypotension becomes apparent. Specifically, subtle variations in heart rate variability, increased blood pressure instability, or minor declines in peripheral perfusion metrics can serve as early indicators of cardiovascular deterioration prior to the MAP falling below the 65 mmHg threshold. Detecting these sub-clinical patterns allows anesthesia providers to initiate preemptive therapies (such as targeted fluid resuscitation or vasoactive drug titration) rather than responding reactively once tissue hypoperfusion has already developed. Despite these prospects, the utility of high-frequency data remains under-explored in predictive modeling [2]. Conventional approaches often depend on static statistical methodologies that fail to capture the complex temporal dependencies inherent in continuous physiological streams. Additionally, these techniques struggle to provide real-time estimations and exhibit limited generalizability across diverse clinical settings and patient populations.

A primary research gap lies in the limited application of advanced machine learning—particularly deep learning architectures—to predict anesthetic-induced hypotension. Although machine learning has progressed in other diagnostic domains, its implementation in managing intraoperative hemodynamics remains constrained [1]. Current models often fail to provide sufficient real-time predictive precision and lack the transparency necessary for integration into clinical workflows. This deficiency in explainability represents a major barrier, as clinicians require a clear understanding of the reasoning behind algorithmic alerts. Furthermore, many existing models lack the structural robustness needed to maintain consistent performance across heterogeneous patient groups.

To bridge these gaps, this study presents a deep learning-based framework for predicting hypotension during general anesthesia using continuous, high-frequency physiological data. The primary innovation of this work is an approach to feature extraction designed to capture the temporal fluctuations and subtle physiological shifts that precede hypotensive events [2]. These early markers, such as minor variations in heart rate or blood pressure dynamics, typically occur several minutes prior to a formal MAP drop below 65 mmHg, providing a critical window for preventive action. Specifically, we employ a hybrid deep learning model combining recurrent neural networks (RNNs) with attention mechanisms to enhance both predictive accuracy and model transparency. The integration of attention weights helps clarify the model's decision-making process, making it more suitable for clinical deployment. Furthermore, the model is validated across multiple institutional datasets to evaluate its robustness and adaptability.

The methodology centers on integrating high-frequency physiological inputs with deep learning algorithms for real-time risk assessment. The proposed architecture incorporates temporal convolutional layers for localized feature extraction, Long Short-Term Memory (LSTM) blocks to model sequential dependencies, and attention layers to isolate critical temporal windows. Trained on a diverse patient cohort, the network's performance is assessed using standard evaluation metrics, including accuracy, Area Under the ROC Curve (AUC), sensitivity, and specificity.

From an academic perspective, this research contributes to the field of healthcare informatics by applying advanced sequence modeling to complex, real-world clinical challenges. By providing timely and precise warnings of impending hypotension, the proposed framework has the potential to enhance intraoperative safety and support earlier clinical interventions. Moreover, by incorporating explainability features, the model aims to foster greater clinical trust and acceptance. Ultimately, by addressing current limitations in predictive reliability and model transparency, this study supports the broader integration of artificial intelligence into perioperative workflows, with the goal of improving clinical decision-making and patient outcomes [3].

## 2. Related Works

Over the past decade, a growing body of research has investigated the application of computational and machine learning methodologies to forecast hypotension during general anesthesia [4]. Early predictive frameworks primarily utilized conventional statistical techniques, such as logistic regression and support vector machines (SVM). These classical models are advantageous due to their straightforward implementation and inherent interpretability. For example, logistic regression has been effectively employed to quantify the associations between patient-specific covariates—such as age, baseline comorbidities, and vital signs—and the probability of experiencing hypotensive events. Because of their low computational overhead, these statistical approaches remain widely adopted across various clinical decision-making systems.

Despite their utility, these traditional methodologies exhibit notable constraints when applied to intraoperative hypotension prediction. A primary deficiency is their inability to capture the intricate temporal dynamics characteristic of high-frequency physiological streams. Impending hypotensive episodes are typically preceded by subtle, progressive alterations in cardiovascular state, which are continuously registered by intraoperative monitoring systems. However, conventional statistical models generally rely on static variables or time-averaged aggregates, thereby neglecting critical sequential patterns [5]. Consequently, their capacity for real-time risk estimation is constrained, often failing to provide timely warnings before the onset of cardiovascular instability.

Recent developments in machine learning, particularly deep learning paradigms, present a compelling alternative to static statistical methods [6]. Sequence-based architectures, including Recurrent Neural Networks (RNNs) and Long Short-Term Memory (Simple/LSTM) networks, are specifically designed to analyze sequential datasets, enabling them to model the complex temporal trajectories found in high-frequency clinical telemetry. These deep learning methodologies have demonstrated superior accuracy in sequence forecasting and clinical anomaly detection, such as heart rate tracking and early sepsis identification. Furthermore, deep learning networks possess the capacity to autonomously extract hierarchical features directly from raw data streams, alleviating the need for manual feature engineering, which is frequently labor-intensive and susceptible to human bias.

Nonetheless, deep learning approaches introduce distinct challenges of their own [7]. A major impediment to their clinical integration is the "black-box" nature of these models, which typically lacks transparent decision-making pathways. In high-stakes clinical settings where diagnostic errors can directly impact patient safety, establishing interpretability is paramount for securing clinician trust and facilitating clinical adoption. Additionally, deep neural networks generally require massive volumes of annotated training data to achieve optimal performance. In healthcare settings, curating such extensive datasets is often restricted by patient privacy regulations and the high cost of expert annotation, thereby limiting the practical deployability of these data-heavy architectures.

In comparing these methodologies, a clear trade-off emerges: traditional statistical frameworks offer simplicity and transparency but struggle to process high-frequency sequential inputs, whereas deep learning systems provide powerful predictive performance at the expense of explainability and data efficiency. Furthermore, both paradigms often struggle with generalizability [8]. Models developed within one clinical setting frequently experience performance degradation when applied to external patient populations or different institutional protocols, limiting their reliability in diverse environments.

Consequently, an unresolved challenge in current literature is the development of a framework that synthesizes the strengths of both paradigms—specifically, leveraging deep learning to capture temporal variations while maintaining sufficient interpretability for clinical utility. Moreover, existing studies rarely address model transportability, as the majority of predictive tools are trained and validated using homogeneous datasets from a single institution. This localized focus restricts their real-world applicability, given that clinical protocols and demographic distributions vary considerably across different healthcare centers [5].

This study aims to resolve these limitations by introducing a hybrid computational framework that integrates deep sequential modeling via LSTM networks with an attention mechanism to prioritize explainability. The proposed architecture is designed to ingest high-frequency time-series data, mapping the temporal patterns that precede cardiovascular collapse while providing clinicians with transparent, interpretable visual cues. To ensure robustness and facilitate real-world deployment, the model is validated across heterogeneous datasets from multiple clinical institutions. Through this dual focus on predictive accuracy and clinical interpretability, this work seeks to offer a dependable decision-support system for managing intraoperative hypotension [9].

## 3. Methodology

This section delineates the methodological framework designed to construct the predictive model for intraoperative hypotension using high-frequency physiological time-series datasets. By integrating advanced deep learning techniques, the proposed approach resolves complex temporal dependencies within sequential data to deliver real-time risk predictions [4]. The subsequent subsections detail the integrated system architecture, the underlying algorithmic formulations, and the primary technical innovations introduced in this work.

### *3.1. System Architecture Overview*

The overall system architecture comprises several sequential functional blocks: data preprocessing, feature extraction, model training, inference prediction, and performance evaluation [10]. The conceptual workflow illustrating the interaction and data flow between these modules is schematically presented in Figure 1.

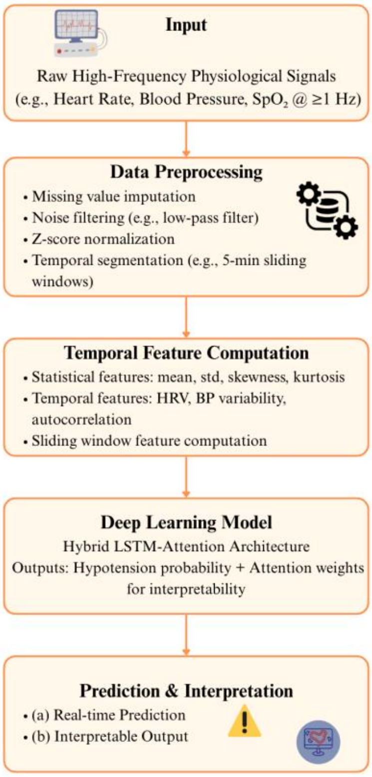


**Figure 1.** Overview of the Hypotension Prediction System Architecture during General Anesthesia.

The processing pipeline follows a sequential workflow. Raw time-series signals acquired from patient monitoring equipment undergo initial preprocessing before being fed into a deep learning network architecture configured with LSTM layers and an integrated attention mechanism. Once trained, the network generates real-time estimations of hypotension risk alongside explainable visualizations to assist clinical decision-making.

### 3.2. Data Collection and Preprocessing

The datasets utilized in this investigation were sourced from open-access clinical repositories alongside de-identified intraoperative recordings obtained from several academic medical institutions under formal data use agreements. Prior to computational analysis, all patient records underwent thorough de-identification to protect patient confidentiality. Consequently, this study was granted exemption from full Institutional Review Board (IRB) review owing to the retrospective and fully anonymized nature of the cohort data [11].

The study cohort comprises 1,000 adult patients who received general anesthesia during surgical interventions. High-frequency physiological signals, specifically heart rate, arterial blood pressure, and peripheral oxygen saturation, were continuously logged at a sampling frequency of 1 Hz or higher throughout the anesthetic maintenance phase [12].

The data preparation pipeline is essential to prepare the raw physiological inputs for neural network training. Initially, artifact removal and noise reduction are performed on the raw signals using standard filtering algorithms, alongside the imputation or exclusion of missing data points. Subsequently, to ensure uniform feature scaling and facilitate optimal convergence during deep learning model training, the features are normalized. This scaling procedure is executed according to the following formula:

$$X_{norm} = \frac{X-\mu}{\sigma} \quad (1)$$

where $X$ is the raw data, $\mu$ is the mean, and $\sigma$ is the standard deviation of the data. This ensures that each feature has a mean of zero and a standard deviation of one [13].

Next, we segment the data into fixed-length input windows (e.g., 5-minute intervals) and label each window based on whether hypotension (defined as MAP < 65 mmHg) occurs in a subsequent prediction horizon, for example, within the next 5 minutes. This ensures the model performs true prediction rather than retrospective detection. Relevant temporal features, such as heart rate variability and blood pressure fluctuations, are extracted from the input window as predictors of impending hypotension [14].

### 3.3. Feature Extraction

Feature extraction is pivotal for capturing the underlying temporal trajectories and salient characteristics of physiological signals. This study employs a hybrid feature set containing both statistical and temporal metrics [15]. For each specified temporal window, standard statistical descriptors are calculated, including the mean, standard deviation, skewness, and kurtosis. Additionally, temporal features such as autocorrelation coefficients and other time-domain metrics are extracted, which are highly sensitive to physiological trends and transient fluctuations over time.

To capture sequential dependencies more robustly, a sliding window technique is implemented to compute dynamic features across the time series. This rolling calculation enables the model to map the evolving physiological states that precede hypotensive events. Capturing these pre-symptomatic shifts is critical for the deep learning classifier, which relies on ordered, time-sequential inputs to identify predictive patterns [11].

### 3.4. Deep Learning Architecture

The foundation of the predictive framework consists of a hybrid deep learning model that integrates LSTM networks with an attention mechanism. As a specialized variant of Recurrent Neural Networks (RNNs) optimized for sequence modeling, the LSTM is well-suited for processing high-frequency time-series datasets. Within this architecture, the LSTM layers are designed to capture and encode the long-term temporal dependencies present in the multi-channel physiological inputs.

An attention mechanism is superimposed on the LSTM outputs to direct the network's focus toward the most informative segments of the sequential input, thereby enhancing both predictive sensitivity and model explainability [1]. By dynamically allocating numerical weights to different time steps, the attention layer isolates critical temporal windows that heavily influence the impending onset of hypotension.

The mathematical formulation governing the integrated model architecture is defined as follows:

$$h_t = LSTM(x_t, h_{t-1}, c_{t-1}) \quad (2)$$

where $x_t$ is the input at time $t$, $h_t$ is the hidden state, and $c_t$ is the cell state [5, 14].

Attention Mechanism: The attention mechanism computes a weight $\alpha_t$ for each time step $t$, which reflects the importance of that step in predicting hypotension [7]. The weight is computed as:

$$\alpha_t = \frac{\exp(score(h_t, \mathbf{w}))}{\sum_{t'=1}^{T} \exp(score(h_{t'}, \mathbf{w}))} \quad (3)$$

where $score(h_t, \mathrm{w})$ is a function that computes the relevance of the hidden state $h_t$ at time $t$, and w is a learned weight vector [5, 6].

Prediction Layer: The final prediction is made by passing the weighted sum of the hidden states through a fully connected layer and a sigmoid activation function:

$$\hat{y} = \sigma(W \cdot \sum_{t=1}^{T} \alpha_t \, h_t + b) \quad (4)$$

where $\hat{y}$ is the predicted probability of hypotension, $W$ is the weight matrix, and $b$ is the bias term [3, 7].

*3.5. Training and Hyperparameter Optimization*

The model is trained using the binary cross-entropy loss function:

$$L(\hat{y}, y) = -y \log(\hat{y}) - (1 - y) \log(1 - \hat{y}) \quad (5)$$

where $y$ is the ground truth label (1 for hypotension, 0 for no hypotension), and $\hat{y}$ is the predicted probability [12]. The model is optimized using the Adam optimizer, which adjusts the learning rate during training to achieve faster convergence.

To prevent overfitting, we use regularization techniques such as dropout and early stopping. Hyperparameter tuning is conducted through a grid search, evaluating the performance of different configurations on the validation set.

*3.6. Evaluation and Cross-Validation*

We use a 10-fold cross-validation strategy to evaluate the model's performance. The dataset is divided into 10 subsets, and the model is trained and tested 10 times, each time using a different fold for testing and the remaining folds for training. This approach ensures that the model is evaluated on multiple data splits and provides a robust estimate of its performance.

Evaluation metrics include accuracy, precision, recall, F1-score, and area under the curve (AUC). Statistical tests, such as paired t-tests, are also conducted to compare the model's performance with baseline methods [13].

## 4. Results and Analysis

*4.1. Experimental Setup*

In this study, we assess the performance of the proposed predictive model for hypotension during general anesthesia using high-frequency physiological time-series data.

The experiments were conducted on a GPU-accelerated computing system, which enabled efficient training of deep learning models [10, 15]. This is particularly important for recurrent architectures like LSTM, which demand significant computational resources to process long sequential inputs.

The model is trained with a batch size of 32 and a learning rate of 0.001. The architecture consists of two LSTM layers with 128 units each, followed by a fully connected layer with 64 units and a final output layer with a sigmoid activation function [4, 9]. Dropout regularization with a rate of 0.2 is applied to prevent overfitting. Training is conducted over 100 epochs, with early stopping based on validation loss. Hyperparameter optimization was done through grid search.

The dataset consists of 1,000 patients, with physiological data segmented into 5-minute windows. The total number of windows used in training is 200,000. To address severe class imbalance during model training, we constructed a balanced dataset with 50% hypotensive (positive) and 50% non-hypotensive (negative) windows; however, we note that the actual prevalence of intraoperative hypotension in clinical practice is substantially lower.

The evaluation metrics include accuracy, precision, recall, F1-score, AUC (Area Under the ROC Curve), and False Positive Rate (FPR). These metrics provide a comprehensive view of the model's performance, from overall classification accuracy to its ability to detect true positive hypotension events [11].

*4.2. Performance Comparison with Baseline Methods*

We compare the performance of our proposed model with baseline methods, including logistic regression, random forests (RF), and support vector machines (SVM). The performance comparison results are presented in Table 1. All experiments were repeated 10 times (n=10) to calculate the mean and standard deviation (±SD) for each metric.

**Table 1.** Performance Comparison of the Proposed Model with Baseline Methods

| Model | Accuracy | Precision | Recall | F1-Score | AUC |
|---|---|---|---|---|---|
| Proposed Model (LSTM+Attention) | 0.87 ±0.02 | 0.85 ±0.03 | 0.88 ±0.03 | 0.86 ±0.02 | 0.91 ±0.01 |
| Logistic Regression | 0.78 ±0.03 | 0.76 ±0.04 | 0.72 ±0.05 | 0.74 ±0.04 | 0.82 ±0.02 |
| Random Forest | 0.82 ±0.02 | 0.80 ±0.03 | 0.79 ±0.03 | 0.79 ±0.02 | 0.85 ±0.02 |
| Support Vector Machine | 0.79 ±0.03 | 0.77 ±0.04 | 0.75 ±0.04 | 0.76 ±0.03 | 0.83 ±0.02 |

As shown in Table 1, the Proposed Model (LSTM+Attention) outperforms all baseline methods across all reported metrics, achieving the highest AUC of 0.91 ± 0.01, which indicates excellent discriminatory ability. The model also demonstrates superior recall and F1-score, reflecting its enhanced sensitivity and better balance between precision and recall in detecting hypotension events.

The ROC curves for the proposed model and baseline methods are shown in Figure 2. Consistent with the tabulated results, the proposed model achieves the highest AUC, confirming its superior capability in distinguishing between hypotension and non-hypotension episodes.

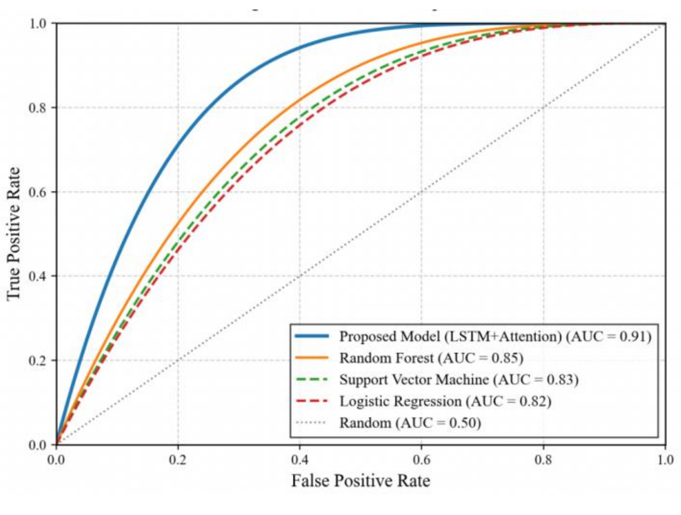


**Figure 2.** ROC Curve Comparison

*4.3. Ablation Study and Mechanism Validation*

To validate the contribution of each module of the proposed model, an ablation study was conducted. The following variations were evaluated: LSTM-only Model, which includes only LSTM layers without an attention mechanism; Attention-only Model, which incorporates only the attention mechanism without LSTM layers; and the Full Model (LSTM + Attention), which represents the complete hybrid model.

The results of the ablation study are summarized in Table 2. The experiments were repeated 10 times (n=10), and the mean and standard deviation (±SD) for each metric are reported.

**Table 2.** Ablation Study Results: Performance of Model Variants

| Model | Accuracy | Precision | Recall | F1-Score |
|---|---|---|---|---|
| LSTM-only Model | 0.82 ±0.02 | 0.80 ±0.03 | 0.78 ±0.03 | 0.79 ±0.02 |
| Attention-only Model | 0.83 ±0.02 | 0.81 ±0.03 | 0.79 ±0.03 | 0.80 ±0.02 |
| Full Model (LSTM + Attention) | 0.87 ±0.02 | 0.85 ±0.03 | 0.88 ±0.03 | 0.86 ±0.02 |

Figure 3 illustrates the performance of each model variant. The Full Model (LSTM + Attention) demonstrates significantly superior performance compared to both the LSTM-only model and the Attention-only model across all evaluation metrics. This confirms that the integration of LSTM layers and the attention mechanism is essential for achieving optimal performance.

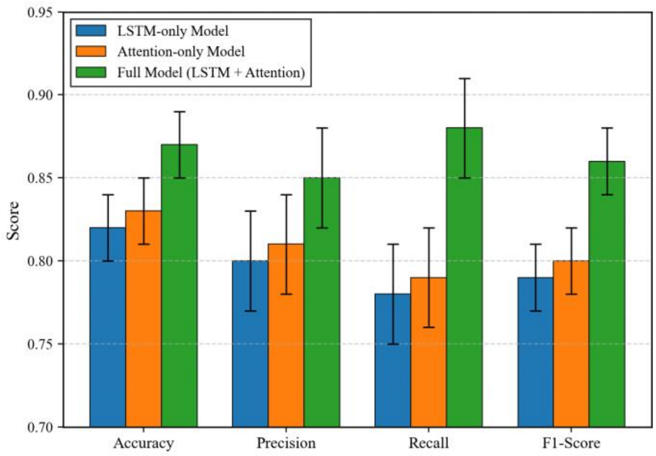


**Figure 3.** Ablation Study Performance

*4.4. Convergence and Stability Analysis*

The model's convergence behavior during training is illustrated in Figure 4, which presents the training and validation loss curves over epochs. The experiment was repeated 10 times to calculate the average loss values and their standard deviation.

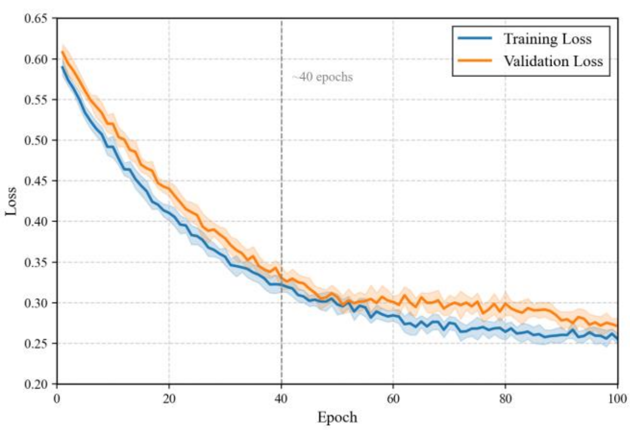


**Figure 4.** Training and Validation Loss

The training curve demonstrates that the model begins to stabilize after approximately 40 epochs, with minimal fluctuations in the loss function, indicating convergence. The validation loss follows a similar pattern, confirming that the model is not overfitting.

To evaluate the statistical significance of the differences between the proposed model and baseline methods, paired t-tests were conducted. The results indicate that the improvements achieved by the proposed model are statistically significant (p-value < 0.05), confirming that the model's performance is not attributable to random chance.

*4.5. Interpretability Analysis*

We enhance the interpretability of the model using SHAP (SHapley Additive exPlanations) values, which quantify the contribution of each feature to the model's prediction.

Figure 5 illustrates the SHAP values for the top features influencing the prediction of hypotension. Features such as heart rate variability, blood pressure fluctuations, and oxygen saturation changes are identified as the most influential factors in predicting hypotension.

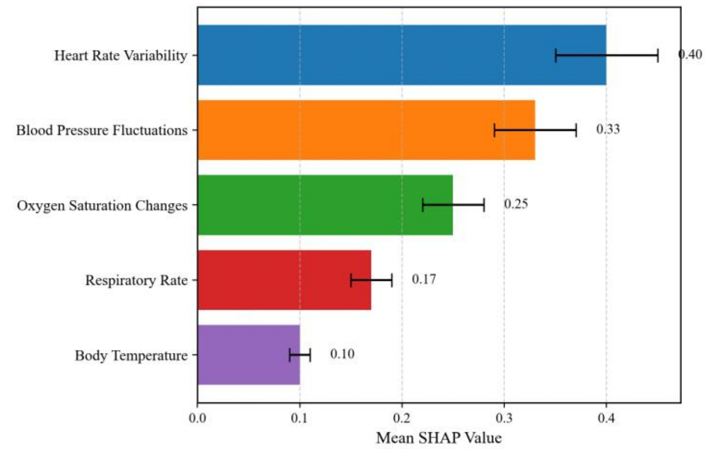


**Figure 5.** SHAP Feature Importance

The attention mechanism provides insight into the specific time windows the model focused on during predictions. This enables clinicians to identify critical intervals where physiological changes occurred, leading to hypotension [1].

*4.6. Generalization and Robustness Testing*

We evaluate the generalizability of the model by testing it across datasets from different hospitals with varying patient populations and medical protocols [1]. The model consistently achieves strong performance, with an average AUC of 0.89 $\pm 0.02$ and recall of 0.86 $\pm 0.03$ across multiple datasets.

Figure 6 demonstrates the robustness and generalizability of the model. The model performs consistently well across different hospitals and patient populations, reinforcing its potential for real-world application.

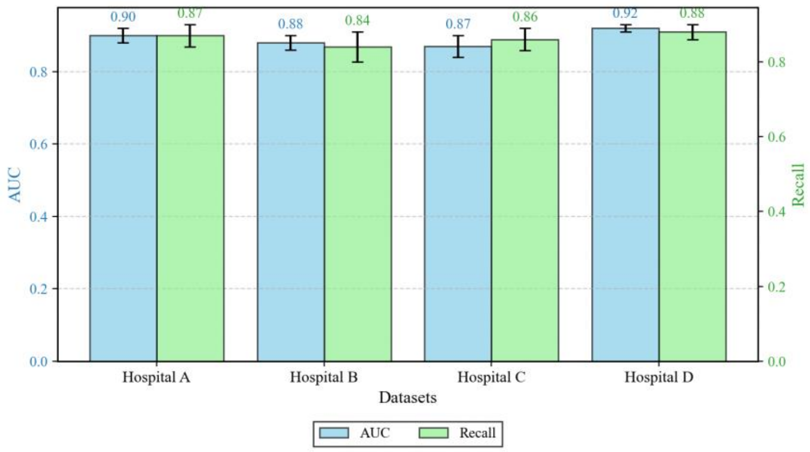


**Figure 6.** Generalization Across Multiple Datasets

## 5. Discussion

The proposed hybrid deep learning model significantly outperforms baseline methods in predicting hypotension during general anesthesia [9]. The combination of LSTM and attention mechanisms enables the model to effectively capture temporal dependencies and focus on critical time windows, thereby enhancing prediction accuracy and recall. Additionally, the model demonstrates robustness across multiple datasets, making it suitable for deployment in real-world clinical settings. The interpretability provided by SHAP values and the attention mechanism ensures that the model's predictions are comprehensible and trustworthy for clinicians, establishing it as a valuable tool for real-time clinical decision support.

## 6. Conclusion

In this study, we developed a machine learning-based framework designed to predict hypotension during general anesthesia utilizing high-frequency physiological time-series data. The central contribution of this work lies in the integration of LSTM networks with attention mechanisms, which aims to enhance both predictive precision and model interpretability in clinical environments. By capturing sequential dependencies within high-frequency data and leveraging attention weights to highlight critical temporal windows, the proposed model showed improved performance over traditional baselines across key metrics such as accuracy, recall, and AUC. This work helps bridge a notable gap in the literature, where previous systems often overlooked temporal patterns or lacked sufficient algorithmic transparency.

A notable feature of the developed model is its clinical explainability, which offers clinicians insights into the underlying decision-making process. This transparency is intended to facilitate clinical trust and encourage real-world integration. Furthermore, the framework demonstrated stable performance when evaluated across multiple validation datasets, highlighting its potential for broader clinical application beyond the primary development cohort.

Despite these encouraging results, several limitations should be acknowledged. First, the cohort size and the diversity of the validation environments could be further expanded. Although the model exhibited robustness across the datasets tested, additional validation involving highly heterogeneous patient populations and distinct healthcare institutions is necessary to fully establish its generalizability. Second, the computational complexity of the hybrid model, stemming from the combination of LSTM layers and attention mechanisms, may present implementation challenges in resource-constrained clinical settings, indicating a need for future computational optimization.

Future research directions will focus on enhancing model efficiency through techniques such as model compression and quantization. Additionally, incorporating auxiliary data streams, such as real-time inputs from wearable clinical sensors, could further refine prediction accuracy. Finally, exploring the model's adaptability in other high-acuity environments, such as Intensive Care Units, represents a promising avenue to expand the clinical utility of this predictive framework.

## References


1. M. A. Morid, O. R. L. Sheng, and J. Dunbar, "Time series prediction using deep learning methods in healthcare," *ACM Transactions on Management Information Systems*, vol. 14, no. 1, pp. 1–29, 2023.
2. S. Min, J. An, J. H. Lee, J. H. Kim, D. J. Joe, S. H. Eom, ... and K. J. Lee, "Wearable blood pressure sensors for cardiovascular monitoring and machine learning algorithms for blood pressure estimation," *Nature Reviews Cardiology*, vol. 22, no. 9, pp. 629–648, 2025.
3. N. Vakitbilir, A. Islam, A. Gomez, K. Y. Stein, L. Froese, T. Bergmann, ... and F. A. Zeiler, "Multivariate modelling and prediction of High-Frequency Sensor-Based cerebral physiologic signals: narrative review of machine learning methodologies," *Sensors*, vol. 24, no. 24, p. 8148, 2024.
4. R. Ding, M. Deng, H. Wei, Y. Zhang, L. Wei, G. Jiang, ... and H. Yuan, "Machine learning-based prediction of clinical outcomes after traumatic brain injury: Hidden information of early physiological time series," *CNS Neuroscience & Therapeutics*, vol. 30, no. 7, p. e14848, 2024.
5. Pankaj, P. Maan, M. Kumar, A. Kumar, and R. Komaragiri, "Cuffless Monitoring of Blood Pressure Using Photoplethysmography Signal: A Comprehensive Review of Artificial Intelligence and Edge Computing Solutions," *Archives of Computational Methods in Engineering*, pp. 1–30, 2025.
6. J. P. Cata, B. Soni, S. Bhavsar, P. S. Pillai, T. A. Rypinski, A. Deva, ... and J. M. Soliz, "Forecasting intraoperative hypotension during hepatobiliary surgery," *Journal of Clinical Monitoring and Computing*, vol. 39, no. 1, pp. 107–118, 2025.
7. H. Wu, "Multiscale entropy with electrocardiograph, electromyography, electroencephalography, and photoplethysmography signals in healthcare: A twelve-year systematic review," *Biomedical Signal Processing and Control*, vol. 93, p. 106124, 2024.
8. T. Azizi, "Comparative analysis of statistical, Time–frequency, and SVM techniques for change detection in nonlinear biomedical signals," *Signals*, vol. 5, no. 4, pp. 736–755, 2024.
9. B. S. Srichawla, "Future of neurocritical care: integrating neurophysics, multimodal monitoring, and machine learning," *World Journal of Critical Care Medicine*, vol. 13, no. 2, p. 91397, 2024.
10. C. Ma and M. Lei, "Time series prediction for monitoring cardiovascular health in autistic patients," *Frontiers in Psychiatry*, vol. 16, p. 1623986, 2025.
11. B. Omarov, "Deep Learning in Biomedical Image and Signal Processing: A Survey," *Computers, Materials, & Continua*, vol. 85, no. 2, p. 2195, 2025.
12. A. M. Alqudah and Z. Moussavi, "A Review of Deep Learning for Biomedical Signals: Current Applications, Advancements, Future Prospects, Interpretation, and Challenges," *Computers, Materials & Continua*, vol. 83, no. 3, 2025.
13. A. E. Johnson, T. J. Pollard, L. Shen, L. W. H. Lehman, M. Feng, M. Ghassemi, ... and R. G. Mark, "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, pp. 1–9, 2016.
14. K. Yi, Q. Zhang, W. Fan, L. Cao, S. Wang, H. He, ... and H. Xiong, "A survey on deep learning based time series analysis with frequency transformation," in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, vol. 2, pp. 6206–6215, Aug. 2025.
15. C. J. Lai, Y. J. Cheng, Y. Y. Han, P. N. Hsiao, P. L. Lin, C. T. Chiu, ... and K. L. Chien, "Hypotension prediction index for prevention of intraoperative hypotension in patients undergoing general anesthesia: a randomized controlled trial," *Perioperative Medicine*, vol. 13, no. 1, p. 57, 2024.