

# A Hybrid Multi-Factor Network with Dynamic Sequence Modeling for Early Warning of Intraoperative Hypotension

Mingyue Cheng<sup>1</sup>, Jintao Zhang<sup>1</sup>, Zhiding Liu<sup>1</sup>, Chunli Liu<sup>2\*</sup>

<sup>1</sup>State Key Laboratory of Cognitive Intelligence, University of Science and Technology of China

<sup>2</sup>Intelligent Interconnected Systems Laboratory of Anhui Province, Hefei University of Technology

mycheng@ustc.edu.cn, {zjt, zhiding}@mail.ustc.edu.cn, liuchunli@hfut.edu.cn

## Abstract

Intraoperative hypotension (IOH) prediction using past physiological signals is crucial, as IOH may lead to inadequate organ perfusion and significantly elevate the risk of severe complications and mortality. However, current methods often rely on static modeling, overlooking the complex temporal dependencies and the inherently non-stationary nature of physiological signals. We propose a Hybrid Multi-Factor (HMF) network that formulates IOH prediction as a dynamic sequence forecasting task, explicitly capturing both temporal dependencies and physiological non-stationarity. We represent signal dynamics as multivariate time series and decompose them into trend and seasonal components, enabling separate modeling of long-term and periodic variations. Each component is encoded with a patch-based Transformer to balance computational efficiency and feature representation. To address distributional drift from evolving signals, we introduce a symmetric normalization mechanism. Experiments on both public and real-world clinical datasets show that HMF significantly outperforms competitive baselines. We hope HMF offers new insights into IOH prediction and ultimately promotes safer surgical care. Our code is available at <https://github.com/Mingyue-Cheng/HMF>.

## 1 Introduction

Intraoperative mortality has decreased by a factor of 100 over the past century, making deaths during surgery a rare occurrence [Li *et al.*, 2009; Xue *et al.*, 2022]. However, mortality within the first month following surgery remains a significant concern, with approximately 2% of patients undergoing inpatient noncardiac surgery dying within 30 days postoperatively [Spence *et al.*, 2019]—amounting to more than 4 million deaths worldwide each year [Saugel and Sessler, 2021]. These postoperative deaths are most strongly linked to complications [Lee *et al.*, 2017], which are often triggered by intraoperative events. Among these, intraoperative hypotension (IOH) is a prevalent and clinically significant complica-




Figure 1: Illustration of MAP sequence: temporal dynamics and distribution shifts in intraoperative hypotension events.

tion [Kim *et al.*, 2023], characterized by a sustained reduction in mean arterial pressure<sup>1</sup> (MAP).

Previous studies [Hatib *et al.*, 2018a] have shown that IOH events can be predicted using machine learning applied to physiological signals. Clinical trials further indicate that early warnings and timely interventions help reduce hypotension severity and postoperative complications [Hwang *et al.*, 2023; Fernandes *et al.*, 2021]. Recent advances primarily rely on handcrafted features combined with classifiers such as logistic regression and random forests [Lee *et al.*, 2021; Davies *et al.*, 2020], but face two key limitations: (1) limited ability to capture temporal dependencies, and (2) lack of flexibility to handle complex real-world IOH scenarios. Additionally, fixed threshold-based IOH definitions may fail to reflect risk accurately in hypertensive or elderly patients, where more nuanced criteria are often needed.

To overcome these limitations, we reformulate IOH prediction as a multivariate time series forecasting problem. This formulation captures the temporal dynamics and evolving distributions inherent in physiological signals. Since hypotensive events are defined based on MAP series, accurate forecasting directly enhances risk identification. However, as illustrated in Figure 1, several challenges arise. First, MAP waveforms comprise trends and periodic components, which may require separate modeling for improved accuracy. Sec-

<sup>1</sup>MAP, derived from arterial blood pressure (ABP), is a key indicator of tissue perfusion. It is calculated as  $MAP = DBP + \frac{1}{3}(SBP - DBP)$  [Gauer, 1960], where systolic blood pressure (SBP) and diastolic blood pressure (DBP) represent the maximum and minimum arterial pressures within a cardiac cycle, respectively.

\*Corresponding author

ond, redundancy in long MAP sequences increases computational cost and complicates representation learning. Third, sudden fluctuations alter statistical properties such as mean and standard deviation, making it difficult to maintain consistent performance. Addressing these challenges is essential for a robust and adaptive IOH prediction framework.

We propose the Hybrid Multi-Factor (HMF) framework, which explicitly models temporal structure and mitigates distribution shifts. HMF employs sequence decomposition to separate trend and seasonal components, yielding a more structured representation of historical dynamics. By incorporating multi-factor physiological signals such as MAP and SBP, HMF captures both short-term fluctuations and long-range dependencies. A patch-based Transformer encoder is adopted for efficient sequence representation, while a symmetric normalization mechanism compensates for evolving input distributions. Extensive experiments on public and private intraoperative datasets demonstrate that HMF consistently improves predictive performance and effectively models multi-factor temporal interactions. The key contributions are summarized as follows:

- We reformulate IOH prediction as a multivariate time series forecasting problem and introduce HMF, which explicitly captures physiological signal dynamics.
- We highlight the challenges of IOH prediction based on historical physiological signal series, particularly the complexities in series structure and the non-stationary nature of physiological signals.
- We validate the effectiveness of the proposed approach through extensive experiments on two real-world datasets, demonstrating superior performance compared to existing methods.

## 2 Related Work

### 2.1 Intraoperative Hypotension Forecasting

Existing methods for intraoperative hypotension (IOH) prediction predominantly adopt a static modeling approach, extracting handcrafted features from physiological signals rather than directly modeling their temporal dynamics. Early studies focused on high-fidelity arterial pressure waveforms, leading to the Hypotension Prediction Index (HPI) [Hatib *et al.*, 2018b]. Machine learning approaches, including ensemble methods [Cherifa *et al.*, 2020] and gradient boosting [Kendale *et al.*, 2018], integrated preoperative and intraoperative factors but largely treated data as independent samples, overlooking temporal dependencies. Deep learning models, such as RNN-based predictors [Jeong *et al.*, 2019] and attention-based frameworks [Lu *et al.*, 2023], improved performance by capturing sequential patterns but primarily relied on fixed-length feature representations, failing to fully model non-stationary hemodynamic changes. Recent interpretable models [Ritter *et al.*, 2023] enhanced clinical applicability but remained feature-driven. Our work addresses these limitations by formulating IOH prediction as a multivariate time-series forecasting problem, explicitly capturing temporal dependencies and non-stationary dynamics in arterial blood pressure signals.




Figure 2: Feature extraction and temporal window setup for blood pressure trend prediction: context, skipped, and target windows.

### 2.2 Time Series and Sequence Forecasting

Time series forecasting is fundamental in domains such as finance, healthcare, and energy [Cheng *et al.*, 2025; Cheng *et al.*, 2024]. Traditional models like ARIMA [Hyndman and Khandakar, 2008] and exponential smoothing struggle with high-dimensional physiological signals, while deep learning methods, including LSTMs [Graves and Graves, 2012] and GRUs [Dey and Salem, 2017], improve long-term dependency modeling. Recent Transformer-based models enhance long-sequence forecasting effectiveness. Informer [Zhou *et al.*, 2021] optimizes self-attention mechanisms through sparse attention and Autoformer [Chen *et al.*, 2021] leverages series decomposition to explicitly model trend and seasonal components, which enhances interpretability and improves forecasting stability over long horizons. Lightweight alternatives, such as DLinear [Zeng *et al.*, 2023], demonstrate competitive performance by preserving temporal structures with simpler architectures. Due to the data sparsity issue [Cheng *et al.*, 2023], GPHT [Liu *et al.*, 2024] further introduces generative pretraining to improve forecasting transferability across datasets and horizons. While significant progress has been made in deterministic forecasting, probabilistic approaches such as diffusion-based models [Tashiro *et al.*, 2021] [Zhang *et al.*, 2025] are also gaining increasing attention for their ability to capture uncertainty. Complementarily, supervised architectures such as TimeDART [Wang *et al.*, 2025] highlight the potential of discriminative learning in capturing temporal patterns effectively. On the other hand, the introduction of LLM-based models [Jin *et al.*, 2023] [Zhou *et al.*, 2023] has enabled the integration of multimodal information. These advancements collectively provide a solid foundation for more accurate modeling of complex medical time series.

While existing models often struggle with real-time adaptability and the challenges posed by non-stationary physiological sequences in dynamic clinical settings, our work bridges this gap by combining efficient patch-based sequence encoding with structured modeling of physiological signals. This design enables rapid inference and robust adaptation to evolving patient states, thereby improving IOH prediction accuracy in dynamic clinical settings.

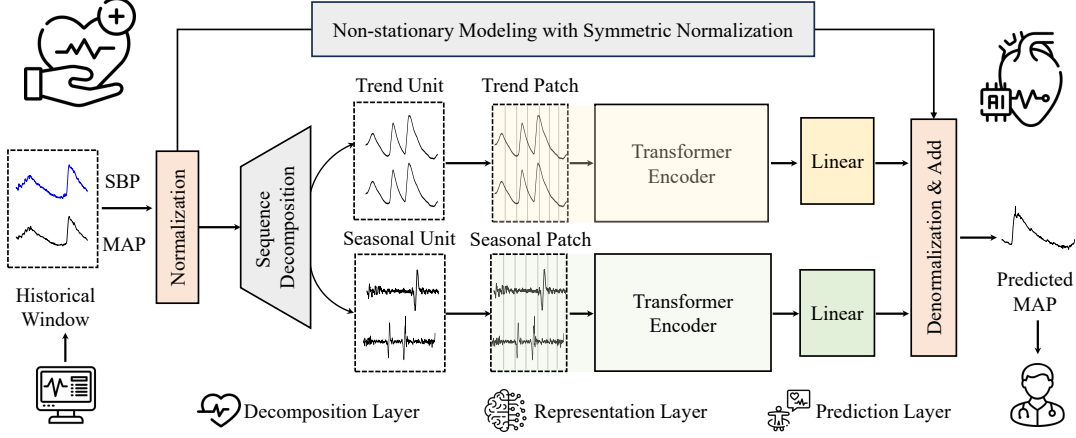


Figure 3: Proposed hybrid multi-factor model for intraoperative hypotension prediction.

### 3 Preliminaries

#### 3.1 Multi-factor Sequence Construction

To predict intraoperative hypotension (IOH), we analyze arterial blood pressure (ABP) dynamics to identify early indicators that precede hypotensive events, thereby enabling timely clinical intervention and risk mitigation. During surgery, ABP is continuously monitored at high frequency, capturing critical hemodynamic fluctuations. Prior studies have demonstrated a strong correlation between short-term variations in ABP and the onset of IOH.

Clinically, IOH is defined as a period when the mean arterial pressure (MAP) remains below 65 mmHg for at least one minute. Systolic blood pressure (SBP), the peak value in a cardiac cycle, is also considered due to its physiological relation to MAP. These two signals, MAP and SBP, serve as the primary predictive inputs for IOH forecasting. To construct reliable input sequences we extract clean ABP segments using a sliding window, remove artifacts and compute SBP and diastolic blood pressure (DBP), as shown in the Figure 2.

#### 3.2 Problem Definition

Building on the multi-factor feature construction framework, we formulate IOH risk prediction as a dynamic sequence forecasting task. Given real-time ABP monitoring, the goal is to forecast future mean arterial pressure (MAP) trends to enable early warnings of hypotensive events. Let  $\mathbf{X}_{\text{MAP}} = \{x_{\text{MAP}}^1, \dots, x_{\text{MAP}}^T\}$  and  $\mathbf{X}_{\text{SBP}} = \{x_{\text{SBP}}^1, \dots, x_{\text{SBP}}^T\}$  represent historical MAP and systolic blood pressure (SBP) sequences over a context window of length  $T$ . The objective is to predict MAP values  $\hat{\mathbf{Y}}_{\text{MAP}} = \{\hat{y}_{\text{MAP}}^{T+1}, \dots, \hat{y}_{\text{MAP}}^{T+\tau}\}$  over horizon  $\tau$ . A hypotensive event is flagged if predicted MAP drops below a threshold  $\theta_{\text{MAP}}$ . Further details on annotation and thresholding are provided in the supplementary material. By modeling multi-factor temporal dependencies, our approach improves IOH risk assessment and supports timely clinical decisions.

### 4 The Proposed HMF

#### 4.1 Framework Overview

The proposed Hybrid Multi-Factor (HMF) model predicts intraoperative hypotension (IOH) by leveraging MAP and SBP time series to capture both long-term trends and short-term fluctuations. As shown in Figure 3, the model consists of four key components: symmetric normalization, sequence decomposition, dynamic dependence modeling, and prediction. First, symmetric normalization reduces inter-patient variability. Sequence decomposition separates trend and seasonal components to enhance feature learning. Transformer encoders process each component independently to capture temporal dependencies. Finally, the model fuses learned representations and reconstructs the MAP trajectory for forecasting. By integrating multi-factor signals and sequence decomposition, the HMF model improves IOH prediction accuracy and adaptability to dynamic MAP variations.

#### 4.2 Symmetric Normalization

Non-stationarity in MAP and SBP series, driven by physiological variations during surgery, causes sudden shifts in statistical properties like mean and standard deviation [Kim *et al.*, 2021], challenging consistent IOH prediction. To mitigate these shifts and enhance reliability, we introduce a symmetric normalization module that ensures consistent feature scales across different surgical conditions, stabilizing model training and improving generalization in unseen cases.

Given a context window  $X \in \mathbb{R}^{L \times 2}$ , instance normalization is applied as:

$$X' = \frac{X - \mu_X}{\sigma_X}, \quad (1)$$

where  $\mu_X$  and  $\sigma_X$  are the mean and standard deviation of  $X$ . This transformation allows the model to focus on relative changes rather than absolute values, reducing the impact of abrupt fluctuations. Additionally, by normalizing across instances, the method mitigates inter-patient variability, ensuring stable feature distributions.

After prediction, de-normalization restores the original scale:

$$\hat{Y} = \hat{Y}' \cdot \sigma_X + \mu_X, \quad (2)$$

where  $\hat{Y}'$  is the predicted sequence and  $\hat{Y}$  is the final output. This ensures stability while preserving raw characteristics, facilitating reliable and interpretable IOH forecasting.

### 4.3 Decomposition Layer

The Arterial Blood Pressure (ABP) signal exhibits a complex waveform structure, consisting of multiple components like trend and periodic elements, which need to be individually modeled for accurate forecasting. After applying instance normalization, which results in the normalized sequence  $X'$ , it becomes essential to decouple the physiological series into its trend and seasonal components. This decomposition allows for more precise modeling of each component, improving the overall forecasting accuracy. We follow the methodology of previous work [Wu *et al.*, 2021] to achieve this decoupling, enabling the model to better capture the inherent patterns in the ABP signal, is computed using:

$$\text{Trend} = \text{AvgPool}(\text{Padding}(X')), \quad (3)$$

$$\text{Seasonal} = X' - \text{Trend}, \quad (4)$$

where AvgPool smooths the sequence by downsampling through averaging within a specified window, while Padding ensures full coverage across the entire sequence.

### 4.4 Representation Layer

**Patch Embedding.** To model the intricate dynamics of MAP and SBP series during surgery, we apply a patch embedding technique [Nie *et al.*, 2022] after the decomposition layer. This reduces sequence length, lowering computational complexity and improving efficiency. Patch-based modeling captures local patterns while mitigating noise and outliers, enhancing prediction robustness.

The decomposed series  $W \in \mathbb{R}^{L \times 2}$  is transformed into a compact representation  $W_{\text{patch}} \in \mathbb{R}^{\frac{L}{S} \times d_{\text{model}}}$  using three 1D convolutional layers:

$$W_{\text{patch}} = \text{Conv1D}(W), \quad (5)$$

where Conv1D applies three consecutive 1D convolutions to  $W$ . The embedding dimension  $d_{\text{model}}$  is chosen to extract key features, improving IOH prediction.

To encode temporal structures, we introduce learnable positional encodings:

$$W_{\text{pos}} = W_{\text{patch}} + \text{PositionalEncoding}(W_{\text{patch}}). \quad (6)$$

These encodings help preserve temporal dependencies, crucial for MAP forecasting and IOH detection, where event timing directly impacts prediction accuracy.

**Sequence Dependence Modeling.** Capturing cross channel dependencies is essential for accurate prediction of multivariate physiological time series [Han *et al.*, 2024]. Effective IOH prediction requires capturing both short-term and long-term dependencies in MAP and SBP series. The Transformer encoder, with its self-attention mechanism, is well-suited for this task, dynamically weighting the importance of different time segments to model complex temporal patterns.

After obtaining patch embeddings  $X_{\text{pos}}$  via patch embedding and positional encoding, the Transformer encoder processes them through self-attention and feedforward layers to learn sequence dependencies. Given positional embeddings  $X_{\text{pos}} \in \mathbb{R}^{\frac{L}{S} \times d_{\text{model}}}$ , the dependencies are computed as:

$$Z = \text{TransformerEncoder}(W_{\text{pos}}), \quad (7)$$

where  $Z \in \mathbb{R}^{\frac{L}{S} \times d_{\text{model}}}$  encodes cross-segment relationships in MAP and SBP. This formulation ensures that the model captures subtle fluctuations and interactions essential for dynamic sequence modeling. By leveraging self-attention mechanism, the Transformer enhances predictive capability, aligning with the multi-factor feature learning framework to improve IOH forecasting.

### 4.5 Prediction Layer

To ensure the temporal patterns captured by the Transformer encoder are effectively utilized in MAP predictions, we employ a linear layer to map the sequence representation  $Z$  to the predicted output. This transformation allows the model to preserve sequence dependencies while translating extracted features into precise forecasts. To maintain temporal coherence, we use a recurrent forecasting strategy where each predicted step conditions the next, ensuring consistency across the forecasted sequence. The mapping from  $Z \in \mathbb{R}^{\frac{L}{S} \times d_{\text{model}}}$  to  $Y \in \mathbb{R}^{\frac{L}{S} \times \frac{2TS}{L}}$  is formulated as:

$$Y = Z \cdot W + b, \quad (8)$$

where  $W$  and  $b$  denote the transformation matrix and bias vector, respectively. Linear transformations have been widely adopted in sequence forecasting [Zeng *et al.*, 2023], facilitating the direct conversion of encoded representations into accurate future MAP predictions.

### 4.6 Optimization Strategies

To improve prediction accuracy, we integrate a patch-based autoregressive method into our model. This approach begins with the final patch of decomposed component  $W$ , denoted as  $Y_0 = W_{\frac{L}{S}}$ , and sequentially generates subsequent patches using the autoregressive formula:

$$P(Y'_{i+1}) = \prod_{j=1}^i P(Y'_{j+1}|Y'_j), \quad (9)$$

in which  $Y'_{i+1}$  represents the next patch. The predicted sequence can be represented as:

$$\hat{Y}' = \{Y'_1, Y'_2, \dots, Y'_T\}. \quad (10)$$

Dataset	Patient Number	Sampling Rate (s)	Predicted Scope	Training Set	Validation Set	Testing Set
CH-OPBP	1,083	1	300	442,858	54,406	58,019
			600	489,273	58,810	66,071
			900	519,339	61,776	71,130
		3	100	148,274	18,243	19,401
			200	164,049	19,719	22,124
			300	174,311	20,764	23,848
VitalDB	1,522	3	100	466,402	61,158	58,042
			200	589,414	78,037	74,094
			300	689,650	91,902	87,685

Table 1: Statistics of training, validation, and testing sets across different datasets and prediction scopes.

By leveraging prior patches, this method improves temporal dependency modeling for MAP forecasting. The model is optimized by minimizing the mean squared error (MSE) between the predicted and actual sequences:

$$\text{MSE} = \frac{1}{T} \sum_{t=1}^T (Y_t - \hat{Y}_t)^2. \quad (11)$$

This optimization refines forecasting accuracy, enhancing IOH prediction effectiveness.




Figure 4: Visualization of data partitioning across different patients during surgery time.

## 5 Experiments

### 5.1 Experimental Setup

**Dataset Description.** We conduct experiments on two real-world clinical datasets: **CH-OPBP** and **VitalDB** [Lee, 2018]. The **CH-OPBP** dataset is collected from real-world intraoperative monitoring records of 3,422 patients, sampled continuously at 100Hz. The data acquisition process strictly adheres to ethical guidelines and protocols approved by the hospital ethics committee, ensuring full compliance with patient privacy regulations and relevant clinical research standards. After resampling to 1-second and 3-second intervals and filtering out records shorter than one hour, a total of 1,083 valid records remain. The data acquisition process also strictly follows patient de-identification protocols and has been formally approved by the Institutional Review Board of Anhui Provincial Hospital. Data collection spans the period from February 27, 2023, to August 4, 2023. The **VitalDB** dataset initially includes 6,388 intraoperative records, resampled uniformly at 3-second intervals. After excluding

records with more than 20% missing data, 1,522 records are retained for use. To ensure reliable model evaluation, both datasets adopt an 80%-10%-10% split for training, validation, and testing, while preserving temporal consistency. The **CH-OPBP** dataset supports multiple prediction horizons to evaluate forecasting flexibility, whereas the **VitalDB** dataset focuses specifically on 3-second resolution data. Table 1 summarizes data partitioning across different sampling rates and prediction scopes. Figure 4 visualizes the partitioning strategy, segmenting each patient’s surgical timeline for structured model training and evaluation.

#### Algorithm 1 IOH Events Detection

**Input:**  $O_{\text{pred}}$ : predicted sequence of MAP values,  $t$ : minimum duration of IOH,  $T_{\text{instance}}$ : length of the target window,  $L_{\text{actual}}$ : actual labels sequence.

**Output:**  $J_{\text{actual}}$ : boolean indicating the presence of an IOH event in the actual instance,  $J_{\text{prediction}}$ : boolean indicating the presence of an IOH event in the predicted instance.

```

1:  $J_{\text{actual}} = 0$ 
2:  $J_{\text{prediction}} = 0$ 
3: for  $i$  in range(0,  $T_{\text{instance}} - t + 1$ ) do
4:    $\text{sum}_{\text{actual}} = \text{sum}(L_{\text{actual}}[i : i + t])$ 
5:   if  $\text{sum}_{\text{actual}} > 0$  then
6:      $J_{\text{actual}} = 1$ 
7:   end if
8:    $\text{sum}_{\text{prediction}} = \sum_{j=0}^{t-1} (O_{\text{pred}}[i + j] \leq 65)$ 
9:   if  $\text{sum}_{\text{prediction}} > 0.8 * t$  then
10:     $J_{\text{prediction}} = 1$ 
11:   end if
12: end for
13: return  $J_{\text{actual}}, J_{\text{prediction}}$ 
    
```

**Compared Baselines.** To evaluate HMF, we compare it with traditional and deep learning-based forecasting models. For traditional methods, we use ARIMA [Hyndman and Khandakar, 2008] and Logistic Regression. ARIMA, a statistical time series model, is trained on 0.5% of the test data for efficiency. Logistic Regression, a standard classifier, is enhanced with 1,566 features extracted via tsfresh<sup>2</sup> and evaluated on 1% of the dataset. For deep learning baselines, we include LSTM [Graves and Graves, 2012], Transformer [Vaswani *et al.*, 2017], Informer [Zhou *et al.*, 2021],

<sup>2</sup><https://github.com/blue-yonder/tsfresh>



Datasets	Sample (s)	Model	MSE	MAE	AUC	Accuracy (%)	Recall (%)
CH-OPBP	1	Arima	130.2702	8.8526	0.5963	<b>77.32</b>	24.00
		LR	—	—	0.5054	76.08	35.10
		LSTM	118.1246	9.0985	0.5295	76.78	6.58
		Transformer	126.7972	9.3809	0.5919	75.63	22.43
		Informer	103.7028	8.0757	0.6452	72.26	35.38
		DLinear	125.6786	9.3232	0.5331	71.60	7.54
		HMF	<b>93.2677</b>	<b>7.5823</b>	<b>0.7352</b>	75.29	<b>67.98</b>
CH-OPBP	3	Arima	112.9281	8.1606	0.5928	75.70	25.65
		LR	—	—	0.6774	75.71	54.49
		LSTM	124.4213	9.8814	0.5000	75.53	0.00
		Transformer	104.8545	8.3441	0.5970	74.44	23.12
		Informer	111.0393	8.3883	0.6278	<b>79.81</b>	30.37
		DLinear	123.8899	9.2951	0.5413	74.10	11.98
		HMF	<b>86.4927</b>	<b>7.2828</b>	<b>0.7413</b>	61.38	<b>70.13</b>
VitalDB	3	Arima	257.3701	13.1127	0.5250	59.31	8.53
		LR	—	—	0.5595	62.60	33.47
		LSTM	188.7123	11.8613	0.5000	<b>75.62</b>	0.00
		Transformer	<b>158.7031</b>	10.6901	0.5040	73.51	0.93
		Informer	158.7873	10.8987	0.5003	75.01	0.05
		DLinear	175.1144	11.4968	0.5074	65.09	1.86
		HMF	165.7575	<b>9.3845</b>	<b>0.6468</b>	69.27	<b>45.87</b>

Table 2: Performance comparison of IOH prediction between our HMF and baseline models on two datasets .

and DLinear [Zeng *et al.*, 2023]. LSTM captures sequential dependencies via recurrence, while Transformer leverages self-attention for long-range interactions. Informer improves efficiency with sparse attention, making it suitable for long-sequence forecasting. DLinear, a lightweight alternative, models trend and seasonal components separately. Further implementation details and hyperparameter settings are in the supplementary material.

**Implementation and Evaluation Details.** The HMF model predicts future MAP values based on physiological forecasts of both MAP and SBP, utilizing a 15-minute context window as input. Predictions extend over 5, 10, and 15 minutes to capture short- and long-term trends. Model training minimizes Mean Squared Error (MSE), with evaluation using both MSE and Mean Absolute Error (MAE), focusing on hypotensive segments ( $L_{\text{actual}} = 1$ ). Although both MAP and SBP series are predicted, IOH assessment relies solely on the forecasted MAP trajectory. Following clinical guidelines [Wesselink *et al.*, 2018], a hypotensive event is identified when predicted MAP remains below  $\theta_{\text{MAP}} = 65$  mmHg for at least 1 minute. To enable effective early warning and facilitate clinical decision-making, a 2-minute skip window [Wijnberge *et al.*, 2020] is applied. Performance is assessed using accuracy, recall, and Area Under the Curve (AUC), with AUC serving as the primary metric for evaluating the model’s ability to differentiate hypotensive trends. Algorithm 1 details the IOH event detection process by comparing predicted and actual MAP sequences in a time-series forecasting framework. Further methodological details, including ground truth labeling and evaluation protocols, are provided in supplement material part.

## 5.2 Experimental Results

**Main Results Analysis.** Table 2 presents a comprehensive evaluation of predictive models for dynamic IOH prediction,

highlighting the advantages of the HMF model. The results indicate that HMF consistently outperforms baseline methods across datasets and sampling rates. Its ability to capture the intricate dynamics of blood pressure trends is reflected in improved predictive accuracy and classification performance. Notably, HMF demonstrates robustness in handling the non-stationary and complex nature of intraoperative blood pressure data, where traditional models often struggle. In the IOH prediction task, HMF shows strong performance in detecting hypotensive events. Its improvements across different sampling rates suggest adaptability to varying temporal granularities, a crucial factor in clinical applications. Compared to LSTM and Transformer, which perform well in some scenarios but face challenges with long sequence dependencies, HMF effectively models temporal dependencies and coupling effects. Informer and DLinear, despite their advantages in general forecasting tasks, show limitations in handling intraoperative data. Overall, HMF maintains stable performance across diverse conditions, demonstrating its potential to enhance IOH prediction in complex clinical settings.

**Ablation Study of Key Components.** To evaluate the impact of key components in the HMF framework, we conduct an ablation study on the CH-OPBP dataset. Table 3 presents the results, highlighting the roles of instance normalization and sequence decomposition in MAP forecasting and IOH detection. The full HMF model, integrating both components, achieves the best overall performance, suggesting their complementary effects in enhancing predictive capability. Removing instance normalization leads to a notable performance drop, indicating its importance in managing the non-stationary nature of intraoperative blood pressure data. Without normalization, the model struggles with variability, affecting its ability to detect IOH events consistently. Similarly, excluding sequence decomposition degrades performance, suggesting its role in capturing complex tempo-

Dataset	Model	MSE	MAE	AUC	Accuracy (%)	Recall (%)
CH-OPBP	HMF (Our model)	<b>86.4927</b>	<b>7.2828</b>	<b>0.7413</b>	75.53	<b>70.13</b>
	w/o instance normalization	106.9671	8.5005	0.5891	77.72	20.73
	w/o sequence decomposition	105.7231	8.4964	0.5750	<b>78.10</b>	17.50

Table 3: Ablation study on instance normalization and sequence decomposition in HMF on the CH-OPBP dataset.

ral dependencies. Without decomposition, the model may fail to identify underlying trends, limiting predictive accuracy. These findings underscore the importance of both components in improving model accuracy. Integrating instance normalization and sequence decomposition enhances generalization, leading to more effective IOH predictions.

Methods	Patient ID	MSE	MAE	AUC
Transfer	1	37.3299	5.6241	0.5375
	2	39.9698	6.2213	0.6100
Non-Transfer	1	18.6823	3.8187	0.2375
	2	47.6620	6.8317	0.5600

Table 4: Comparison of transfer and non-transfer learning on individual patients.



Methods	Feature	MSE	MAE	AUC
Transfer	Elderly	118.9462	9.1424	0.6089
	Young	38.0568	5.0047	0.8040
Non-Transfer	Elderly	66.5058	6.5874	0.7850
	Young	44.1721	5.2447	0.8540

Table 5: Comparison of transfer and non-transfer learning on elderly and young subgroups.

### Performance Analysis under Transfer Learning Settings.

Assessing model transferability across patients and demographic groups is essential for IOH prediction. Transfer learning enhances the robustness and generalizability of the HMF framework. To evaluate this, we apply the HMF model, trained on CH-OPBP, to new patient cohorts and age groups, comparing performance with and without transfer learning. Table 4 presents cross-patient transfer results, where a model trained on one patient is applied to another. AUC scores suggest that transfer learning preserves or enhances predictive performance, indicating the model’s ability to capture shared physiological patterns. This reduces reliance on patient-specific data while maintaining accuracy, though individual training may still offer advantages in distinct cases. Table 5 presents cross-group results between age cohorts. The non-transfer model performs better in elderly patients, suggesting that model better capture physiological complexities. In contrast, younger patients benefit from transfer learning, likely due to more homogeneous physiological responses. These findings underscore the importance of demographic considerations while demonstrating transfer learning’s potential to improve performance across diverse clinical settings.

**Parameter Sensitivity Analysis.** To optimize key parameters and enhance the Patch Encoder’s ability to capture temporal dependencies, we perform sensitivity analysis on the CH-OPBP dataset, focusing on model dimension  $d_{\text{model}}$  and


 Figure 5: Analysis of MSE and AUC performance for IOH segments across varying  $d_{\text{model}} \in \{64, 128, 256, 512, 1024\}$ .

 Figure 6: Analysis of MSE and AUC performance for IOH segments across varying  $S \in \{5, 10, 20, 25, 50\}$ .

patch length  $S$ . We first vary  $d_{\text{model}}$  with a fixed context window of 300 and prediction length of 100. As shown in Figure 5, performance is sensitive to  $d_{\text{model}}$ , with a clear optimum maximizing AUC. We then examine the effect of varying  $S$  while keeping  $d_{\text{model}}$  constant. Since  $S$  governs the model’s ability to capture local temporal patterns, its selection is important. Results in Figure 6 confirm that appropriate  $S$  values significantly affect forecasting accuracy. These findings guide the selection of optimal hyperparameters to ensure performance and highlight the importance of configuring the Patch Encoder to effectively model blood pressure dynamics.

## 6 Conclusion

In this paper, we proposed a novel approach to intraoperative hypotension (IOH) prediction by reformulating it as a dynamic sequence forecasting task, enabling proactive and continuous risk assessment. To address key challenges inherent in modeling arterial blood pressure dynamics, we introduced symmetric normalization to mitigate non-stationarity in mean arterial pressure (MAP) and systolic blood pressure (SBP) series, sequence decomposition to better capture complex temporal patterns, and a patch-based Transformer model to enhance representation learning and forecasting accuracy. Extensive experiments on real-world intraoperative datasets validated the effectiveness and robustness of our HMF framework, demonstrating its strong performance in improving early IOH forecasting under various clinical conditions.

## Acknowledgements

This research was supported by grants from the grants of Provincial Natural Science Foundation of Anhui Province (No.2408085QF193), USTC Research Funds of the Double First-Class Initiative (No. YD2150002501), the Fundamental Research Funds for the Central Universities of China (Grant No. PA2024GDSK0112, No. WK2150110032), Anhui Provincial Natural Science Foundation (No. 2308085MG226), the Key Technologies R&D Program of Anhui Province (No. 202423k09020039).

## References

- [Chen *et al.*, 2021] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12270–12280, 2021.
- [Cheng *et al.*, 2023] Mingyue Cheng, Qi Liu, Zhiding Liu, Hao Zhang, Rujiao Zhang, and Enhong Chen. Timemae: Self-supervised representations of time series with decoupled masked autoencoders. *arXiv preprint arXiv:2303.00320*, 2023.
- [Cheng *et al.*, 2024] Mingyue Cheng, Jiqian Yang, Tingyue Pan, Qi Liu, and Zhi Li. ConvtimeNet: A deep hierarchical fully convolutional model for multivariate time series analysis. *arXiv preprint arXiv:2403.01493*, 2024.
- [Cheng *et al.*, 2025] Mingyue Cheng, Zhiding Liu, Xiaoyu Tao, Qi Liu, Jintao Zhang, Tingyue Pan, Shilong Zhang, Panjing He, Xiaohan Zhang, Daoyu Wang, et al. A comprehensive survey of time series forecasting: Concepts, challenges, and future directions. *Authorea Preprints*, 2025.
- [Cherifa *et al.*, 2020] Ményssa Cherifa, Alice Blet, Antoine Chambaz, Etienne Gayat, Matthieu Resche-Rigon, and Romain Pirracchio. Prediction of an acute hypotensive episode during an ICU hospitalization with a super learner machine-learning algorithm. *Anesthesia & Analgesia*, 130(5):1157–1166, 2020.
- [Davies *et al.*, 2020] Simon James Davies, Simon Tilma Vistisen, Zhongping Jian, Feras Hatib, and Thomas WL Scheeren. Ability of an arterial waveform analysis-derived hypotension prediction index to predict future hypotensive events in surgical patients. *Anesthesia & Analgesia*, 130(2):352–359, 2020.
- [Dey and Salem, 2017] Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th MWSCAS*, pages 1597–1600. IEEE, 2017.
- [Fernandes *et al.*, 2021] Marta Priscila Bento Fernandes, Miguel Armengol de la Hoz, Valluvan Rangasamy, and Balachundhar Subramaniam. Machine learning models with preoperative risk factors and intraoperative hypotension parameters predict mortality after cardiac surgery. *Journal of Cardiothoracic and Vascular Anesthesia*, 35(3):857–865, 2021.
- [Gauer, 1960] OH Gauer. Kreislauf des blutes. *Lehrbuch der Physiologie des Menschen*, 28:65–186, 1960.
- [Graves and Graves, 2012] Alex Graves and Alex Graves. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, pages 37–45, 2012.
- [Han *et al.*, 2024] Lu Han, Han-Jia Ye, and De-Chuan Zhan. The capacity and robustness trade-off: Revisiting the channel independent strategy for multivariate time series forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [Hatib *et al.*, 2018a] Feras Hatib, Zhongping Jian, Sai Buddi, Christine Lee, Jos Settels, Karen Sibert, Joseph Rinehart, and Maxime Cannesson. Machine-learning algorithm to predict hypotension based on high-fidelity arterial pressure waveform analysis. *Anesthesiology*, 129(4):663–674, 2018.
- [Hatib *et al.*, 2018b] Feras Hatib, Zhongping Jian, Sai Buddi, Christine Lee, Jos Settels, Karen Sibert, Joseph Rinehart, and Maxime Cannesson. Machine-learning algorithm to predict hypotension based on high-fidelity arterial pressure waveform analysis. *Anesthesiology*, 129(4):663–674, 2018.
- [Hwang *et al.*, 2023] Eugene Hwang, Yong-Seok Park, Jin-Young Kim, Sung-Hyuk Park, Junetae Kim, and Sung-Hoon Kim. Intraoperative hypotension prediction based on features automatically generated within an interpretable deep learning model. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [Hyndman and Khandakar, 2008] Rob J Hyndman and Yeasmin Khandakar. Automatic time series forecasting: the forecast package for R. *Journal of statistical software*, 27:1–22, 2008.
- [Jeong *et al.*, 2019] Young-Seob Jeong, Ah Reum Kang, Woohyun Jung, So Jeong Lee, Seunghyeon Lee, Misoon Lee, Yang Hoon Chung, Bon Sung Koo, and Sang Hyun Kim. Prediction of blood pressure after induction of anesthesia using deep learning: A feasibility study. *Applied Sciences*, 9(23):5135, 2019.
- [Jin *et al.*, 2023] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, et al. Time-llm: Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728*, 2023.
- [Kendale *et al.*, 2018] Samir Kendale, Prathamesh Kulkarni, and Rosenberg et al. Supervised machine-learning predictive analytics for prediction of postinduction hypotension. *Anesthesiology*, 129(4):675–688, 2018.
- [Kim *et al.*, 2021] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2021.
- [Kim *et al.*, 2023] Junetae Kim, Sung-Hoon Kim, Eugene Hwang, Yong-Seok Park, Jin-Young Kim, and Sung-Hyuk Park. Intraoperative hypotension prediction based on features automatically generated within an interpretable deep learning model. *Authorea Preprints*, 2023.



- [Lee *et al.*, 2017] Hyung-Chul Lee, Ho-Geol Ryu, and Chul-Woo Jung. Performance measurement of intraoperative systolic arterial pressure to predict in-hospital mortality in adult liver transplantation. *Scientific reports*, 7(1):7030, 2017.
- [Lee *et al.*, 2021] Solam Lee, Hyung-Chul Lee, Yu Seong Chu, Seung Woo Song, Gyo Jin Ahn, Hunju Lee, Sejung Yang, and Sang Baek Koh. Deep learning models for the prediction of intraoperative hypotension. *British journal of anaesthesia*, 126(4):808–817, 2021.
- [Lee, 2018] Hyung-Chul et al Lee. Vital recorder—a free research tool for automatic recording of high-resolution time-synchronised physiological data from multiple anaesthesia devices. *Scientific reports*, 8(1):1527, 2018.
- [Li *et al.*, 2009] Guohua Li, Margaret Warner, Barbara H Lang, Lin Huang, and Lena S Sun. Epidemiology of anesthesia-related mortality in the united states, 1999–2005. In *The Journal of the American Society of Anesthesiologists*, volume 110, pages 759–765. The American Society of Anesthesiologists, 2009.
- [Liu *et al.*, 2024] Zhiding Liu, Jiqian Yang, Mingyue Cheng, Yucong Luo, and Zhi Li. Generative pretrained hierarchical transformer for time series forecasting. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2003–2013, 2024.
- [Lu *et al.*, 2023] Feng Lu, Wei Li, Zhiqiang Zhou, Cheng Song, Yifei Sun, Yuwei Zhang, Yufei Ren, Xiaofei Liao, Hai Jin, Ailin Luo, et al. A composite multi-attention framework for intraoperative hypotension early warning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 14374–14381, 2023.
- [Nie *et al.*, 2022] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.
- [Ritter *et al.*, 2023] Jodie Ritter, Xiaoyu Chen, Lihui Bai, and Jiapeng Huang. Predicting hypotension by learning from multivariate mixed responses. In *Proceedings of the International MultiConference of Engineers and Computer Scientists 2023*, 2023.
- [Saugel and Sessler, 2021] Bernd Saugel and Daniel I Sessler. Perioperative blood pressure management. *Anesthesiology*, 134(2):250–261, 2021.
- [Spence *et al.*, 2019] Jessica Spence, Yannick LeManach, Matthew TV Chan, CY Wang, Alben Sigamani, Denis Xavier, Rupert Pearse, Pablo Alonso-Coello, Ignacio Garutti, Sadeesh K Srinathan, et al. Association between complications and death within 30 days after noncardiac surgery. *Cmaj*, 191(30):E830–E837, 2019.
- [Tashiro *et al.*, 2021] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in neural information processing systems*, 34:24804–24816, 2021.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [Wang *et al.*, 2025] Daoyu Wang, Mingyue Cheng, Zhiding Liu, and Qi Liu. Timedart: A diffusion autoregressive transformer for self-supervised time series representation. In *Forty-second International Conference on Machine Learning*, 2025.
- [Wesselink *et al.*, 2018] EM Wesselink, TH Kappen, HM Torn, AJC Slooter, and WA Van Klei. Intraoperative hypotension and the risk of postoperative adverse outcomes: a systematic review. *British journal of anaesthesia*, 121(4):706–721, 2018.
- [Wijnberge *et al.*, 2020] Marije Wijnberge, Bart F Geerts, Liselotte Hol, Nikki Lemmers, Marijn P Mulder, Patrick Berge, Jimmy Schenk, Lotte E Terwindt, Markus W Hollmann, Alexander P Vlaar, et al. Effect of a machine learning-derived early warning system for intraoperative hypotension vs standard care on depth and duration of intraoperative hypotension during elective noncardiac surgery: the hype randomized clinical trial. *Jama*, 323(11):1052–1060, 2020.
- [Wu *et al.*, 2021] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 22419–22430. Curran Associates, Inc., 2021.
- [Xue *et al.*, 2022] Bing Xue, York Jiao, Thomas Kannampallil, Bradley Fritz, Christopher King, Joanna Abraham, Michael Avidan, and Chenyang Lu. Perioperative predictions with interpretable latent representation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4268–4278, 2022.
- [Zeng *et al.*, 2023] Aixin Zeng, Ming Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11121–11128, 2023.
- [Zhang *et al.*, 2025] Jintao Zhang, Mingyue Cheng, Xiaoyu Tao, Zhiding Liu, and Daoyu Wang. Conditional denoising meets polynomial modeling: A flexible decoupled framework for time series forecasting, 2025.
- [Zhou *et al.*, 2021] Haoyi Zhou, Shuai Zhang, Jiaxuan Peng, Sihang Zhang, Jianwen Li, Hongkai Xiong, and Weishi Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11106–11115, 2021.
- [Zhou *et al.*, 2023] Tian Zhou, Peisong Niu, Liang Sun, Rong Jin, et al. One fits all: Power general time series analysis by pretrained lm. *Advances in neural information processing systems*, 36:43322–43355, 2023.