

Hierarchical Information-Guided Spatio-Temporal Mamba for Stock Time Series Forecasting

Wenbo Yan^{1,2}, Shurui Wang^{1,2}, and Ying Tan^{1,3,4} (✉)

¹ School of Intelligence Science and Technology, Peking University, Beijing, China

² Computational Intelligence Laboratory

³ Institute for Artificial Intelligence

⁴ National Key Laboratory of General Artificial Intelligence

wenboyan@stu.pku.edu.cn, wsr@pku.edu.cn, ytan@pku.edu.cn

Abstract. Mamba has demonstrated excellent performance in various time series forecasting tasks due to its superior selection mechanism. Nevertheless, conventional Mamba-based models encounter significant challenges in accurately predicting stock time series, as they fail to adequately capture both the overarching market dynamics and the intricate interdependencies among individual stocks. To overcome these constraints, we introduce the Hierarchical Information-Guided Spatio-Temporal Mamba (HIGSTM) framework. HIGSTM introduces Index-Guided Frequency Filtering Decomposition to extract commonality and specificity from time series. The model architecture features a meticulously designed hierarchical framework that systematically captures both temporal dynamic patterns and global static relationships within the stock market. Furthermore, we propose an Information-Guided Mamba that integrates macro informations into the sequence selection process, thereby facilitating more market-conscious decision-making. Comprehensive experimental evaluations conducted on the CSI500, CSI800 and CSI1000 datasets demonstrate that HIGSTM achieves state-of-the-art performance.

Keywords: Stock Time Series Forecasting · Spatio-Temporal Mamba · Decomposition · Information-Guided.

1 Introduction

Stock time series forecasting, a pivotal component in investment decision-making, continues to be a primary research focus. The advancement of deep neural networks has driven the development of diverse models to tackle this challenge. Early approaches included RNN-based models such as TPA-LSTM[12] and CNN-based models like TCN[5]. Later, Transformer-based models, including Crossformer[23], iTransformer[9], and PatchTST[10], demonstrated robust performance in time series forecasting. With the introduction of Mamba[3], Mamba-based models such as Bi-Mamba[13] and TimeMachine[1] have emerged as highly effective solutions.

Mamba-based models demonstrate exceptional capabilities in efficient long-sequence modeling. However, given the distinctive characteristics of the stock

market, there exists a pronounced interdependence among stocks within the same market. Modeling each stock’s time series independently fails to capture the comprehensive market dynamics. Current Mamba models encounter two primary limitations in stock forecasting: 1. They are unable to model the influence of related stocks on each individual stock during node modeling, and 2. Their selection mechanisms depend exclusively on historical time series, lacking the capacity to integrate additional information to improve selection performance.

To address this issue, we propose the Hierarchical Information-Guided Spatio-Temporal Mamba (HIGSTM). HIGSTM incorporates a hierarchical structure that models the entire stock market from three perspectives: individual, temporal, and global. On one hand, it hierarchically aggregates neighborhood information of stock nodes, capturing both dynamic and static inter-stock relationships. On the other hand, we introduce an information-guided Mamba structure, which progressively extracts time step macro information and global macro information, integrating it into the sequence selection process.

We propose an Index-Guided Frequency Filtering Decomposition that transforms stock time series and indices into the frequency domain. Using stock indices, we derive filter parameters to decompose time series into commonality and specificity components. Each node is initially modeled independently via a Mamba block. Subsequently, the Temporal Information-Guided Spatio-Temporal Mamba (TIGSTM) integrates dynamic temporal information through: (1) leveraging sequence specificity to construct a sparse time-varying relationship graph, enabling dynamic neighborhood aggregation, and (2) aggregating time step macro information from commonality to guide Mamba’s sequence selection. The Global Information-Guided Spatio-Temporal Mamba (GIGSTM) incorporates static global information by: (1) aggregating specificity across all time steps to form a global static relationship graph for comprehensive neighborhood aggregation, and (2) consolidating commonality to generate global macro information for enhanced sequence selection guidance. In summary, the main contributions are as follows:

- We propose an Index-Guided Frequency Filtering Decomposition method to effectively extract commonality and specificity from time series.
- We introduce a Hierarchical Information-Guided Spatio-Temporal Mamba structure, which extracts node-related relationships and macro-level information across multiple perspectives, aggregates neighborhood information, and enhances sequence selection.
- We conduct comprehensive experiments on multiple real-world stock datasets, demonstrating the superior performance of our model.

2 Preliminary

2.1 Definition of Stock Spatio-Temporal Forecasting Problem

We formulate the **Stock Spatio-Temporal Forecasting Problem** as follows: Given historical data across T time steps for N stocks, where each stock at each

time step is described by F features, the dataset is represented as $\mathbf{X} \in \mathbb{R}^{N \times T \times F}$. The index time series for market are denoted as $\mathbf{I} \in \mathbb{R}^{T \times 1}$.

Moreover, stocks demonstrate complex interconnections that collectively constitute a graph structure. In this framework, each stock represents a node, while the correlation between stocks defines the edges. These relationships are captured by the adjacency matrix $\mathcal{G} \in \mathbb{R}^{N \times N}$, where \mathcal{G}_{ij} quantifies the correlation between stock i and stock j . When an edge connects stock i and stock j ($\mathcal{G}_{ij} \neq 0$), they are identified as neighboring stocks. The set of neighbors for stock i is formally defined as $\mathbf{U}_i = \{j \mid j \neq i \text{ and } \mathcal{G}_{ij} \neq 0\}$.

2.2 Definition of Broadcast

Considering the multi-scale information aggregation presented in the article, various data components may exhibit dimensional inconsistencies. Following the definition in Section 2.1, we normalize all variables to conform to the three-dimensional structure $[N, T, F]$. We introduce matrix broadcasting as the operation of replicating a matrix along absent dimensions to achieve the target $[N, T, F]$ format. This operation is formally expressed as:

$$\mathcal{B}^{(Dims)}(*) \quad (1)$$

For example, $\mathcal{B}^{(N)}(\mathbf{Z}_{t,f})$, where $\mathbf{Z}_{t,f} \in \mathbb{R}^{T \times F}$, denotes the operation of adding a dimension to $\mathbf{Z}_{t,f}$ and repeating it N times, resulting in a shape of $\mathbb{R}^{N \times T \times F}$.

2.3 Fast Fourier Transform and Amplitude Filter

The Fast Fourier Transform (FFT) is an efficient algorithm for computing the Discrete Fourier Transform (DFT), which maps a time-domain signal to the frequency domain. For ease of reference, we denote the FFT computation process as $FFT(*)$ and its inverse as $iFFT(*)$. Any frequency-domain matrix is denoted as \square^f , and its corresponding amplitude is denoted as $\square^{f,amp}$.

The amplitude filter selectively enhances or suppresses signals by directly manipulating the amplitude components in the frequency domain. Its inputs are a frequency-domain matrix and filtering parameters. For clarity, we represent the amplitude filtering operation as Θ .

2.4 Discretization and State Space Model

The computation process of Mamba involves three matrices: the state transition matrix \mathbf{A} , the input matrix \mathbf{B} , and the output matrix \mathbf{C} . Mamba employs the Zero-Order Hold (ZOH) method [3] to convert the continuous parameters \mathbf{A} and \mathbf{B} into discrete parameters $\overline{\mathbf{A}}$ and $\overline{\mathbf{B}}$. We denote this discretization process as:

$$\begin{aligned} \overline{\mathbf{A}}, \overline{\mathbf{B}} &= Discretization(\mathbf{A}, \mathbf{B}, \Delta) \\ \overline{\mathbf{A}} &= \exp(\Delta \mathbf{A}), \overline{\mathbf{B}} = (\exp(\Delta \mathbf{A}) - \mathbf{E})(\Delta \mathbf{A})^{-1}(\Delta \mathbf{B}) \end{aligned} \quad (2)$$

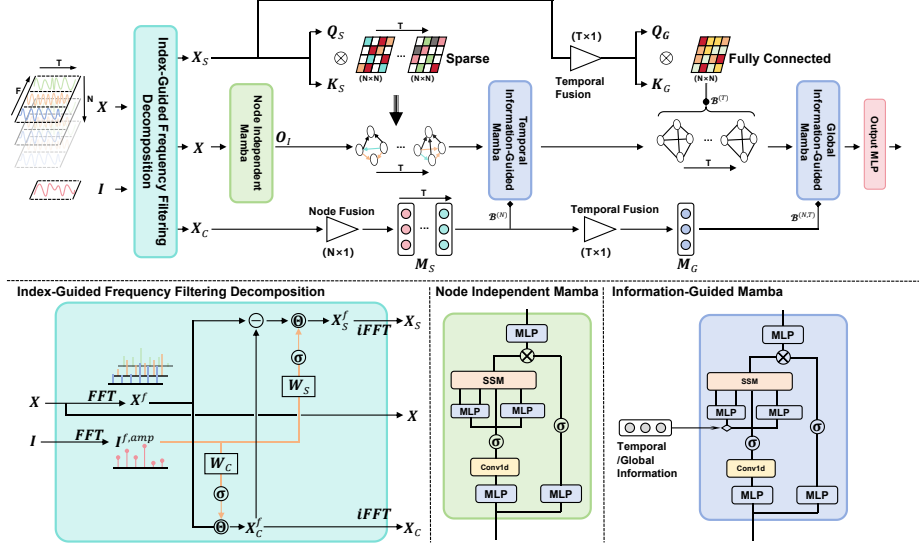


Fig. 1. The overview of the proposed Hierarchical Information-Guided Spatio-Temporal Mamba (HIGSTM)

where \mathbf{E} is the identity matrix. Additionally, we denote the computation process of the Structured State Space Model (SSM) as $\mathbf{O} = \text{SSM}(\bar{\mathbf{A}}, \bar{\mathbf{B}}, \mathbf{C})(\mathbf{X})$, with its core computation process being:

$$\begin{aligned} \mathbf{h}_t &= \bar{\mathbf{A}}\mathbf{h}_{t-1} + \bar{\mathbf{B}}\mathbf{x}_t \\ \mathbf{o}_t &= \mathbf{C}\mathbf{h}_t \end{aligned} \quad (3)$$

3 Method

In this section, we introduce our proposed **Hierarchical Information-Guided Spatio-Temporal Mamba (HIGSTM)**.

3.1 Index-Guided Frequency Filtering Decomposition

Stocks within the same market demonstrate both commonality, reflecting shared market trends and macroeconomic information that influences all stocks' future trajectories, and specificity, capturing individual characteristics that are more effective for identifying inter-stock relationships. Eliminating commonality results in sparser stock relationships, preserving only the strongest correlations.

However, commonality and specificity are intertwined and challenging to decompose directly without a reference. Our analysis reveals that stock indices, serving as indicators of macroeconomic, can effectively guide this decomposition. Leveraging this insight, we introduce the Index-Guided Frequency Filtering Decomposition. This approach begins by transforming both the stock time series \mathbf{X}

and the index time series \mathbf{I} into the frequency domain through the Fast Fourier Transform (FFT).

$$\begin{aligned}\mathbf{X}^f &= FFT(\mathbf{X}), \mathbf{I}^f = FFT(\mathbf{I}) \\ \mathbf{I}^{f,amp} &= |\mathbf{I}^f|\end{aligned}\tag{4}$$

where \mathbf{X}^f and \mathbf{I}^f are the frequency-domain representations of the stock series and the index series, respectively, $\mathbf{I}^{f,amp}$ is the amplitude of the index series.

From a frequency-domain perspective, decomposing commonality and specificity involves analyzing signal intensity across different frequencies. We utilize the index series' amplitude to derive the amplitude filter parameters, which are then applied to filter the stock series.

$$\mathbf{X}_c^f = \mathbf{X}^f \Theta \sigma(\mathbf{W}_c^T \mathbf{I}^{f,amp})\tag{5}$$

where \mathbf{W}_c is a linear transformation matrix, σ denotes the sigmoid activation function, Θ represents the filtering operation, and \mathbf{X}_c^f corresponds to the commonality. The filter parameters learned from index series' amplitude ensures that the commonality aligns more closely with macroeconomic market characteristics. Additionally, we focus on isolating the specificity from this decomposition.

$$\mathbf{X}_s^f = (\mathbf{X}^f - \mathbf{X}_c^f) \Theta (1 - \sigma(\mathbf{W}_s^T \mathbf{I}^{f,amp}))\tag{6}$$

where \mathbf{W}_s is a linear mapping matrix and \mathbf{X}_s^f denotes the specificity. Subtracting \mathbf{X}_c^f from \mathbf{X}^f removes the commonality, preserving each stock's specificity. Leveraging index information, we further eliminate common market trends by deriving an additional amplitude filter. Ultimately, we obtain the decomposed commonality series \mathbf{X}_c and specificity series \mathbf{X}_s via the inverse Fast Fourier Transform (iFFT).

$$\mathbf{X}_c = iFFT(\mathbf{X}_c^f), \mathbf{X}_s = iFFT(\mathbf{X}_s^f)\tag{7}$$

3.2 Node Independent Mamba

Following series decomposition, we initially model each stock's time series independently, deliberately excluding inter-stock correlations. This strategy prevents the premature introduction of neighboring node information, which could obscure intrinsic series features and compromise the stock's uniqueness. This methodology mirrors the subjective stock evaluation process: a stock's intrinsic characteristics fundamentally determine its future trajectory, and even in favorable market conditions, a fundamentally weak stock will not yield positive predictions.

We use a Mamba block [3] to independently model each time series, taking the raw stock time series \mathbf{X} as input. \mathbf{X} undergoes a one-dimensional convolution (Conv1d) to extract local features, followed by a Linear Projection that maps it

to the input matrix \mathbf{B}_{In} , the output matrix \mathbf{C}_{In} , and the discretized time step Δ_{In} .

$$\begin{aligned}\mathbf{X}_{In} &= \mu(\text{Conv1d}(\mathbf{X})) \\ \mathbf{B}_{In} &= \mathbf{W}_{B,I}^\top \mathbf{X}_{In}, \quad \mathbf{C}_{In} = \mathbf{W}_{C,I}^\top \mathbf{X}_{In} \\ \Delta_{In} &= \text{softplus}(\mathbf{W}_{\Delta,I}^\top \mathbf{X}_{In} + \mathbf{A}_{In})\end{aligned}\tag{8}$$

where $\mathbf{W}_{B,I}$, $\mathbf{W}_{C,I}$, and $\mathbf{W}_{\Delta,I}$ are linear transformation matrices, $\mu(*)$ is the SiLU activation function, $\text{softplus}(*)$ denotes the Softplus activation function, and \mathbf{A}_{In} is an optimizable matrix. Subsequently, the matrices \mathbf{A}_{In} and \mathbf{B}_{In} are discretized into $\bar{\mathbf{A}}_{In}$ and $\bar{\mathbf{B}}_{In}$.

$$\bar{\mathbf{A}}_{In}, \bar{\mathbf{B}}_{In} = \text{Discretization}(\mathbf{A}_{In}, \mathbf{B}_{In}, \Delta_{In})\tag{9}$$

Then, $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$, \mathbf{C} , and \mathbf{X} are fed into the State Space Model (SSM)[3].

$$\mathbf{O}_{In} = \text{SSM}(\bar{\mathbf{A}}_{In}, \bar{\mathbf{B}}_{In}, \mathbf{C}_{In})(\mathbf{X}_{In})\tag{10}$$

where the $\text{SSM}(*)$ process is detailed in Section 2.4. Within the Node Independent Mamba, we utilize the Mamba block to model each stock node independently, enabling exploration of intrinsic node characteristics while preventing neighboring information from obscuring the node's unique features.

3.3 Temporal Information-Guided Spatio-temporal Mamba

The stock market operates as an integrated system, necessitating the consideration of inter-stock interactions. We introduce the Temporal Information-Guided Spatio-temporal Mamba block (**TIGSTM**), capturing both inter-stock relationships within each time step and the influence of time step macro information on individual stocks. TIGSTM integrates two key components: Temporal Section Sparse Neighbor Aggregation and the Temporal Information-Guided Selective State Space Model.

Temporal Section Sparse Neighbor Aggregation We utilize the decomposed specificity series \mathbf{X}_s to capture inter-stock relationships. Given their dynamic evolution over time, we implement an attention mechanism to establish the correlation at each time step.

$$\begin{aligned}\mathbf{Q}_S &= \mathbf{W}_{Q,S}^\top \mathbf{X}_s, \quad \mathbf{K}_S = \mathbf{W}_{K,S}^\top \mathbf{X}_s \\ \mathcal{G}_S &= \text{softmax}(\text{topK}(\frac{\mathbf{Q}_S^\top \mathbf{K}_S}{\sqrt{d_{kS}}}, r)) + \mathbf{E}_S\end{aligned}\tag{11}$$

where $\mathbf{W}_{Q,S}$ and $\mathbf{W}_{K,S}$ denote linear mapping matrices, $\sqrt{d_{kS}}$ represents the dimension of the \mathbf{K}_S vector, \mathbf{E}_S is the identity matrix, r is a hyperparameter that controls the proportion of top neighbors, and $\text{topK}()$ ensures each node retains only the neighbors with the highest attention weights. The specificity series \mathbf{X}_s eliminates commonality, resulting in sparser yet more prominent node relationships. We enforce sparsity in the graph at each time step, restricting each

node to retain only r of its neighbors. We empirically set r to 30%. This enables each stock to aggregate information from its neighbors at each time step.

$$\mathbf{X}_{S,Agg} = \mathcal{G}_S \mathbf{O}_{In} \quad (12)$$

The time series after information aggregation is denoted as $\mathbf{X}_{S,Agg}$. **Temporal Information-Guided Selective State Space Model** Mamba’s Selection mechanism dynamically generates the parameters of the state space model (SSM) based on the time step, with the matrices being derived from the time series, significantly enhancing SSM performance. We introduce the Temporal Information-Guided Selective State Space Model, which integrates the time step macro information into the selection mechanism to guide sequence selection. Initially, we aggregate node information from the commonality series \mathbf{X}_c to form the time step macro information.

$$\mathbf{M}_S = \mathbf{W}_{F,S}^T (\mathbf{W}_{(N \times 1)}^T \mathbf{X}_c) + \mathbf{b}_{F,S} \quad (13)$$

where $\mathbf{W}_{(N \times 1)}$ is an $N \times 1$ mapping matrix that aggregates information from all stocks, and $\mathbf{W}_{F,S}$ is a linear mapping matrix that further blends the features to form the macro-level information \mathbf{M}_S for each time step. We incorporate \mathbf{M}_S into the input matrix and output matrix.

$$\begin{aligned} \mathbf{X}_S &= \mu(\text{Conv1d}(\mathbf{X}_{S,Agg})) \\ \mathbf{B}_S &= \langle \mathbf{W}_{B,S}^T \mathbf{X}_S, \mathcal{B}^{(N)}(\mathbf{M}_S) \rangle, \mathbf{C}_S = \langle \mathbf{W}_{C,S}^T \mathbf{X}_S, \mathcal{B}^{(N)}(\mathbf{M}_S) \rangle \\ \Delta_S &= \text{softplus}(\mathbf{W}_{\Delta,S}^T \mathbf{X}_S + \mathbf{A}_S) \end{aligned} \quad (14)$$

where $\mathbf{W}_{B,S}$ and $\mathbf{W}_{C,S}$ denote linear mapping matrices that project the input sequence $\mathbf{X}_{S,Agg}$ into the input and output matrices, \mathbf{A}_S represents a learnable state transition matrix, and $\langle \dots \rangle$ indicates matrix concatenation. We broadcast and concatenate the time step macro information \mathbf{M}_S into input matrix and output matrix. This ensures that the SSM process incorporates time step macro information for input and output decisions. Subsequently, we discretize the state transition matrix \mathbf{A}_S and input matrix \mathbf{B}_S and feed them into the SSM.

$$\begin{aligned} \overline{\mathbf{A}}_S, \overline{\mathbf{B}}_S &= \text{Discretization}(\mathbf{A}_S, \mathbf{B}_S, \Delta_S) \\ \mathbf{O}_S &= \text{SSM}(\overline{\mathbf{A}}_S, \overline{\mathbf{B}}_S, \mathbf{C}_S)(\mathbf{X}_S) \end{aligned} \quad (15)$$

where \mathbf{O}_S represents the final output of TIGSTM. The TIGSTM propagates information within localized neighborhoods, with each time step maintaining unique adjacency relationships to fully capture dynamic inter-stock correlation. Meanwhile, it extracts time step macro information to guide Mamba’s sequence selection mechanism.

3.4 Global Information-Guided Spatio-temporal Mamba

Beyond time-evolving dynamic information, static information also influences the change of the stock. To capture this, we introduce the Global Information-Guided Spatio-temporal Mamba Block (**GIGSTM**), which extracts both global stock correlation and global macro information.

Global Neighbor Aggregation Global Neighbor Aggregation models the global inter-stock relationships from the specificity series \mathbf{X}_s and enables each stock to aggregate information from its global neighbors. Initially, we aggregate the information across all time steps in the specificity series \mathbf{X}_s and construct the global stock correlations through an attention mechanism.

$$\begin{aligned}\mathbf{X}_{s,G} &= \mathbf{W}_{(T \times 1)}^\top \mathbf{X}_s \\ \mathbf{Q}_G &= \mathbf{W}_{Q,G}^\top \mathbf{X}_{s,G}, \mathbf{K}_G = \mathbf{W}_{K,G}^\top \mathbf{X}_{s,G} \\ \mathcal{G}_G &= \text{softmax}\left(\frac{\mathbf{Q}_G^\top \mathbf{K}_G}{\sqrt{d_{kG}}}\right) + \mathbf{E}_G\end{aligned}\quad (16)$$

where $\mathbf{W}_{(T \times 1)}$ denotes a $T \times 1$ mapping matrix that integrates information from all time steps to derive global specificity $\mathbf{X}_{s,G}$; $\mathbf{W}_{Q,G}$ and $\mathbf{W}_{K,G}$ represent linear mapping matrices; $\sqrt{d_{kG}}$ indicates the dimension of the \mathbf{K}_G vector; \mathbf{E}_G is the identity matrix; and \mathcal{G}_G represents the adjacency matrix of the learned global correlation graph. By aggregating multiple time step information, inter-node correlation becomes more comprehensive. Consequently, we relax the sparsity constraint on \mathcal{G}_G , allowing a fully connected structure. Finally, we broadcast global correlation across all time steps, enabling each stock to aggregate information from all others.

$$\mathbf{X}_{G,Agg} = \mathcal{B}^{(T)}(\mathcal{G}_G)\mathbf{O}_S \quad (17)$$

The time series after information aggregation is denoted as $\mathbf{X}_{G,Agg}$.

Global Information-Guided Selective State Space Model Global information plays a pivotal role in time step selection, as identical sequences may yield divergent decisions when incorporating versus excluding market context. We propose the Global Information-Guided Selective State Space Model, which enhances Mamba’s selection mechanism through global macro information. Initially, we extract global macro information from each time step’s macro data.

$$\mathbf{M}_G = \mathbf{W}_{F,G}^\top (\mathbf{W}_{(T \times 1)}^\top \mathbf{M}_S) + \mathbf{b}_{F,G} \quad (18)$$

where $\mathbf{W}_{(T \times 1)}$ denotes a $T \times 1$ mapping matrix that integrates time step macro information, and $\mathbf{W}_{F,G}$ represents a linear mapping matrix that fuses features to generate global macro information \mathbf{M}_G . We then incorporate \mathbf{M}_G into the input and output matrix to guide the sequence selection process.

$$\begin{aligned}\mathbf{X}_G &= \mu(\text{Conv1d}(\mathbf{X}_{G,Agg})) \\ \mathbf{B}_G &= \langle \mathbf{W}_{B,G}^\top \mathbf{X}_G, \mathcal{B}^{(N,T)}(\mathbf{M}_G) \rangle, \mathbf{C}_G = \langle \mathbf{W}_{C,G}^\top \mathbf{X}_G, \mathcal{B}^{(N,T)}(\mathbf{M}_G) \rangle \\ \Delta_G &= \text{softplus}(\mathbf{W}_{\Delta,G}^\top \mathbf{X}_G + \mathbf{A}_G)\end{aligned}\quad (19)$$

where $\mathbf{W}_{B,G}$ and $\mathbf{W}_{C,G}$ represent linear mapping matrices that project the input sequence $\mathbf{X}_{G,Agg}$ into the input and output matrices, \mathbf{A}_G denotes a learnable state transition matrix, and $\langle \dots \rangle$ indicates matrix concatenation. We propagate global macro information \mathbf{M}_G to each stock’s time step and concatenate it with both matrices. This integration ensures the SSM process incorporates

global macro information for sequence selection. Finally, we discretize the state transition matrix \mathbf{A}_G and input matrix \mathbf{B}_G for SSM processing.

$$\begin{aligned}\overline{\mathbf{A}}_G, \overline{\mathbf{B}}_G &= \text{Discretization}(\mathbf{A}_G, \mathbf{B}_G, \Delta_G) \\ \mathbf{O}_G &= \text{SSM}(\overline{\mathbf{A}}_G, \overline{\mathbf{B}}_G, \mathbf{C}_G)(\mathbf{X}_G)\end{aligned}\quad (20)$$

where \mathbf{O}_G represents the final output of GIGSTM. The GIGSTM integrates time step macro information to learn both global macro information and static global stock correlations. Through a fully connected graph, each stock accesses information from all others, enhancing its representation. Meanwhile, global macro information guides each stock’s time step selection, ensuring predictions incorporate historical patterns while aligning with market trends, thus improving stock modeling.

3.5 Prediction and Losses

Unlike conventional sequence prediction methods, we directly map all features output by GIGSTM to the prediction target through a linear layer. Considering the unique characteristics of stock prediction, we employ the mean-deviation prediction approach to achieve better forecasting performance [20].

$$\begin{aligned}\text{mean} &= \mathbf{W}_{\text{mean}}^\top \mathbf{O}_G + \mathbf{b}_{\text{mean}} \\ \text{dev} &= \tanh(\mathbf{W}_{\text{dev}}^\top \mathbf{O}_G + \mathbf{b}_{\text{dev}}) \\ \hat{\mathbf{Y}} &= \text{mean} + e^{\text{dev}}\end{aligned}\quad (21)$$

where \mathbf{W}_{mean} and \mathbf{W}_{dev} are linear mapping matrices, and \mathbf{b}_{mean} and \mathbf{b}_{dev} are the corresponding bias matrices. The prediction target is decomposed into mean and deviation predictions. We use the Pearson correlation coefficient loss $\mathcal{L}_{\text{pearson}}$ as the loss function to learn the ranking distribution at each time step.

$$\mathcal{L}_{\text{pearson}} = -\frac{(\mathbf{Y} - \bar{\mathbf{Y}})^\top (\hat{\mathbf{Y}} - \bar{\hat{\mathbf{Y}}})}{\sqrt{(\mathbf{Y} - \bar{\mathbf{Y}})^\top (\mathbf{Y} - \bar{\mathbf{Y}})} \cdot \sqrt{(\hat{\mathbf{Y}} - \bar{\hat{\mathbf{Y}}})^\top (\hat{\mathbf{Y}} - \bar{\hat{\mathbf{Y}}})}}\quad (22)$$

4 Experiments

4.1 Experimental Setup

Datasets We conducted experiments on three datasets: CSI500, CSI800, CSI1000. Brief statistical information is listed in Table 1. Detailed information about the datasets can be found in the appendix.

Baseline We compare our model with the following baselines: spatio-temporal models ASTGCN[4], MTGNN[18], DCRNN[7], STEMGNN[2], FC-STGNN[16] and MASTER[6], and time series models MICN[14], Filtarnet[21], iTransformer[9], TimeMixer[15], and TimesNet[17]. Where MASTER are specialized stock prediction models. Detailed descriptions of these models can be found in the appendix.

Table 1. The overall information for datasets

Dataset	Samples	Node	Partition
CSI500	3159	500	2677/239/243
CSI800	3159	800	2677/239/243
CSI1000	3159	1000	2677/239/243

Metrics We evaluate the performances of all baseline by nine metrics which are commonly used in stock prediction task including Information Coefficient (IC), Profit and Loss (PNL), Max Drawdown (MAXD), Sharpe Ratio (SHARPE), Win Rate (WINR), Profit/Loss Ratio (PL).

4.2 Main Results

We compare the proposed HIGSTM model with eight baselines across three datasets, and the experimental results are shown in Table 2. Overall, our method achieves the best performance in terms of IC, PNL, SHARPE, WINR, and PL on all datasets. For MAXD, HIGSTM just ranks second on the CSI500 dataset, slightly behind Filtarnet. On average, IC improves by 11%, SHARPE by 10%, and PNL and PL by over 6%. Additionally, the model’s trading win rate increases by at least 2%, which represents a substantial improvement for stock trading.

More specifically, HIGSTM significantly outperforms other models in IC, with improvements of up to 18%, indicating that our method provides superior predictions. Compared to temporal models, HIGSTM incorporates cross-sectional and global information to guide sequence modeling, allowing each sequence to consider both historical patterns and overall market trends. Compared to spatio-temporal models, our proposed decomposition method extracts sequence specificity, and the hierarchical structure progressively models both dynamic and static relationships among stocks. The improvements in WINR and PL demonstrate that our model has stronger predictive capabilities for top-performing stocks, enabling more accurate identification of high-potential stocks. Regarding MAXD, HIGSTM underperforms Filtarnet on the CSI500 dataset, suggesting that while HIGSTM delivers higher returns, it also introduces slightly higher risk. However, this is not necessarily a drawback, as evidenced by the SHARPE metric, which measures the ratio of return to risk. HIGSTM significantly outperforms other models in SHARPE, with improvements exceeding 15%, indicating that the increase in returns far outweighs the associated risk.

4.3 Ablation Study

To validate the effectiveness of our model, we conducted ablation study on the CSI500 and CSI1000 datasets, using IC (to measure prediction accuracy) and SHARPE (to measure the return-to-risk ratio) as evaluation metrics. Specifically, we compared the following variants:

- **w/o TIGSTM:** remove Temporal Information-Guided Spatio-temporal Mamba block

Table 2. Comparison results on CSI500, CSI800 and CSI1000 datasets. ↓ indicates that the smaller the metric is better. The best result is in bold. The suboptimal results are indicated with an underline.

		CSI500					
		IC	PNL	MAXD↓	SHARPE	WINR	PL
Spatio-Temporal	ASTGCN	0.0434	0.0811	0.1387	0.4676	0.5110	1.0899
	MTGNN	0.0503	0.0797	0.1582	0.6963	0.5232	1.1866
	DCRNN	0.0521	0.1401	0.1537	0.9265	0.5162	1.1688
	STEMGNN	0.0479	0.0708	0.1308	0.7475	0.4699	1.0375
	FC-STGNN	0.0420	0.0946	0.1441	0.5858	0.5124	1.1102
	MASTER	0.0567	<u>0.1766</u>	0.1086	<u>1.2904</u>	<u>0.5579</u>	<u>1.2493</u>
Temporal	MICN	<u>0.0668</u>	0.0834	0.1390	0.6408	0.5165	1.1197
	Filternet	0.0464	0.1345	0.0688	0.8863	0.5165	1.1688
	iTransformer	0.0400	0.1214	0.1412	0.8340	0.4835	1.1574
	TimeMixer	0.0472	0.1044	0.1436	0.7657	0.5041	1.1398
	TimesNet	0.0346	0.0516	0.1250	0.3374	0.5041	1.0621
HIGSTM		0.0791	0.2619	<u>0.1028</u>	1.4846	0.5596	1.3207
		CSI800					
		IC	PNL	MAXD↓	SHARPE	WINR	PL
Spatio-Temporal	ASTGCN	0.0431	0.0712	0.1380	0.5149	0.5027	1.1140
	MTGNN	0.0433	0.0939	0.1279	0.6682	0.5091	1.1409
	DCRNN	0.0373	0.0482	0.1374	0.3406	0.5098	1.0693
	STEMGNN	0.0364	0.0594	0.1289	0.3849	0.5032	1.0809
	FC-STGNN	0.0285	0.0312	0.1468	0.1803	0.5124	1.0322
	MASTER	<u>0.0628</u>	<u>0.1603</u>	<u>0.1116</u>	<u>1.2076</u>	0.4959	<u>1.2363</u>
Temporal	MICN	0.0380	0.0212	0.1493	0.1614	<u>0.5165</u>	1.0288
	Filternet	0.0302	0.0951	0.1255	0.6352	0.5083	1.1188
	iTransformer	0.0457	0.0314	0.1348	0.2192	0.4752	1.0397
	TimeMixer	0.0338	0.0487	0.1280	0.3028	0.5083	1.0548
	TimesNet	0.0253	0.0311	0.1427	0.2112	0.4793	1.0387
HIGSTM		0.0708	0.1918	0.1040	1.2602	0.5263	1.3617
		CSI1000					
		IC	PNL	MAXD↓	SHARPE	WINR	PL
Spatio-Temporal	ASTGCN	0.0694	0.1313	0.1295	0.9434	0.4932	1.0884
	MTGNN	0.0701	0.1108	0.1406	1.1415	0.5098	1.2393
	DCRNN	0.0726	0.1426	0.1614	1.1791	0.5531	1.2593
	STEMGNN	0.0688	0.1392	0.1021	1.1334	0.4867	1.2084
	FC-STGNN	0.0733	0.1514	0.1365	0.9339	<u>0.5537</u>	1.1709
	MASTER	0.0850	0.1843	<u>0.0995</u>	<u>1.3382</u>	0.5455	<u>1.2601</u>
Temporal	MICN	<u>0.0876</u>	0.1779	0.1467	1.2339	0.5455	1.2359
	Filternet	0.0723	<u>0.1883</u>	0.1312	1.2715	0.5413	1.2392
	iTransformer	0.0747	0.1642	0.1407	1.1638	0.5289	1.2226
	TimeMixer	0.0797	0.1580	0.1178	1.0681	0.5289	1.2007
	TimesNet	0.0766	0.1512	0.1092	1.0244	0.5372	1.1922
HIGSTM		0.0918	0.2178	0.0759	1.4173	0.5845	1.3566

- **w/o GIGSTM**: remove Global Information-Guided Spatio-temporal Mamba block
- **w/o T&GIGSTM**: remove Temporal and Global Information-Guided Spatio-temporal Mamba block
- **w/o Decomposition**: remove the Index-Guided Frequency Filtering Decomposition; both the relationships and macro-level information are extracted directly from the stock time series.

- **w/o Index**: when decomposing the stock time series, the decomposition is not guided by the stock indices.

The experimental results are shown in Table 3. It can be observed that removing either TIGSTM or GIGSTM leads to a significant decline in both IC and SHARPE, with similar magnitudes of degradation. This indicates that both temporal cross-sectional information and global information are crucial for prediction. When both TIGSTM and GIGSTM are removed, the model essentially loses its predictive capability, demonstrating the effectiveness of our proposed TIGSTM and GIGSTM. Additionally, removing the Decomposition process results in a 30% performance drop, highlighting the necessity of extracting commonality and specificity from the time series. Commonality enables better extraction of macro-level information, while specificity helps uncover more prominent relationships. Furthermore, when the Decomposition process is not guided by the Index, the model’s performance also declines by over 30%, underscoring the critical role of the Index in guiding the decomposition. Using the Index to learn filter parameters is an indispensable step.

Table 3. Ablation Study

	CSI500		CSI1000	
	IC	SHARPE	IC	SHARPE
w/o TIGSTM	0.0733	0.8262	0.0745	0.9026
w/o GIGSTM	0.0604	0.6124	0.0701	1.0075
w/o T&GIGSTM	0.0409	0.0333	0.0630	0.9172
w/o Decomposition	0.0539	0.6822	0.0594	0.4605
w/o Index	0.0451	0.5711	0.0592	0.7073
HIGSTM	0.0791	1.4846	0.0918	1.4173

4.4 Hyper-parameter Study

We conducted a detailed parameter study on the model on CSI500 dataset, and the experimental results are shown in Figure 2. We investigated the following parameters in HIGSTM: the hidden layer dimension d_{model} , the state dimension d_{state} , the convolution kernel dimension d_{conv} , the macro-level information dimension for each time step d_{te} , the global macro-level information dimension d_{ge} , and the correlation attention dimension d_{attn} .

The results show that the optimal value for d_{model} is 64. When d_{model} is smaller than 64, the dimension is insufficient to capture all features, while when it exceeds 64, the model tends to overfit, leading to a decline in performance. For d_{state} , although the performance differences are minor, a value of 16 is recommended considering computational efficiency. The optimal value for d_{conv} is 4, as larger convolution kernels can obscure local features, degrading performance.




Fig. 2. Hyper-parameter Study

For d_{te} , d_{ge} , and d_{attn} , the optimal dimension is 32. When these parameters are smaller than 32, the model cannot fully capture the required information, while values larger than 32 lead to overfitting, where the guiding information overshadows the unique characteristics of each node.

4.5 Analysis

Analysis of Characteristics of Temporal Section and Global Graphs We sampled 50 nodes and analyzed the characteristics of both the temporal section and the global graph through visualization. The visualized structure is shown in Figure 3(a). It can be observed that the graph learned in each temporal section is sparse and prominent. Each node has only a few neighboring nodes, and most nodes exhibit 1-2 highly correlated neighbors that stand out significantly. This aligns with the model’s design intention, which is to identify a small number of truly significant neighbors at each time point, preventing the information from being overshadowed by a large number of neighbors. On the other hand, the Global Graph identifies general static relationships, measuring correlations with all nodes, and does not exhibit significantly correlated neighbors.

Analysis of Dynamic Changes in Correlations We selected a stock (stock 0) and visualized its correlations with 100 stocks over time, as shown in Figure 3(b). It can be observed that, on one hand, the adjacency relationships of stock 0 change significantly over time steps, indicating that the temporal section graph can capture rapidly changing relationships based on different time steps. On the other hand, stock 0 shares the same neighbors across time steps, demonstrating that the temporal graph is not only capable of capturing fast-changing relationships at the time-step level, but also relationships that persist over a period of time.




Fig. 3. (a) Visualization of Temporal Section and Global Graphs. (b) Visualization of Sparse Correlation in Stock Time Series

5 Related Work

5.1 Temporal Models

Research in time series forecasting has focused on improving prediction accuracy through various models. Transformer-based models include Informer[24], and FEDformer[25], which use sparse attention and decomposition techniques. CNN-based models like TimesNet[17] extract patterns via frequency and time domain segmentation. MLP-based models such as TimeMixer[15] leverage basis approximation and multi-scale decomposition. Mamba-based models like TimeMachine[1] and Bi-Mamba[13] refine content selection. Other models include PatchTST[10], which segments time series into patches, and iTransformer[9], which redefines time embeddings.

5.2 Spatio-Temporal Models

Spatio-temporal forecasting integrates temporal and spatio information for better predictions. Convolution-based models such as MTGNN[18], StemGNN[2], use gating mechanisms for feature extraction. Complex architectures like ST-GDN[22] improve performance through advanced designs. RNN+GNN models like DCRNN[7] combine temporal and spatio modeling. Transformer-based models such as ASTGCN[4] and STTN[19] integrate attention mechanisms for capturing dependencies.

5.3 Stock Price Forecasting Models

Stock price forecasting has evolved from traditional methods like ANN[26] and regularization to deep learning. Models like MTDNN[8] have shown potential. Spatio-temporal frameworks like STHAN-SR[11] and MASTER[6] capture temporal and variable dependencies.

6 Conclusion

In this paper, we propose the Hierarchical Information-Guided Spatio-Temporal Mamba (HIGSTM) to address the limitations of Mamba in stock time series forecasting. HIGSTM decomposes the series into commonality and specificity through Index-Guided Frequency Filtering Decomposition, extracting macro-level information and inter-stock relationships separately. Then, it introduces a hierarchical structure: Node Independent, Temporal Information-Guided Spatio-temporal Mamba, and Global Information-Guided Spatio-temporal Mamba, progressively modeling node features, temporal dynamic correlations, and static global relationships. Our proposed Information-Guided Mamba incorporates time step macro-level information and global information to guide the sequence selection process, ensuring decisions align with the overall market. Experimental results on real-world stock market data fully demonstrate the effectiveness of our approach.

Acknowledgments. This work is supported by the National Natural Science Foundation of China (Grant No. 62276008, 62250037, and 62076010), and partially supported by the National Key R&D of China (Grant #2022YFF0800601).

References

1. Ahamed, M.A., Cheng, Q.: Timemachine: A time series is worth 4 mambas for long-term forecasting. In: ECAI 2024: 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain-Including 13th Conference on Prestigious Applications of Intelligent Systems. European Conference on Artificial Intelli. vol. 392, p. 1688 (2024)
2. Cao, D., Wang, Y., Duan, J., Zhang, C., Zhu, X., Huang, C., Tong, Y., Xu, B., Bai, J., Tong, J., et al.: Spectral temporal graph neural network for multivariate time-series forecasting. *Advances in neural information processing systems* **33**, 17766–17778 (2020)
3. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752* (2023)
4. Guo, S., Lin, Y., Feng, N., Song, C., Wan, H.: Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, pp. 922–929 (2019)
5. Lea, C., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks: A unified approach to action segmentation. In: *Computer vision–ECCV 2016 workshops: Amsterdam, the Netherlands, October 8-10 and 15-16, 2016, proceedings, part III 14*. pp. 47–54. Springer (2016)
6. Li, T., Liu, Z., Shen, Y., Wang, X., Chen, H., Huang, S.: Master: Market-guided stock transformer for stock price forecasting. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 38, pp. 162–170 (2024)
7. Li, Y., Yu, R., Shahabi, C., Liu, Y.: Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926* (2017)
8. Liu, G., Mao, Y., Sun, Q., Huang, H., Gao, W., Li, X., Shen, J., Li, R., Wang, X.: Multi-scale two-way deep neural network for stock trend prediction. In: *IJCAI*. pp. 4555–4561 (2020)

9. Liu, Y., Hu, T., Zhang, H., Wu, H., Wang, S., Ma, L., Long, M.: itransformer: Inverted transformers are effective for time series forecasting. arXiv preprint arXiv:2310.06625 (2023)
10. Nie, Y., Nguyen, N.H., Sinthong, P., Kalagnanam, J.: A time series is worth 64 words: Long-term forecasting with transformers. arXiv preprint arXiv:2211.14730 (2022)
11. Sawhney, R., Agarwal, S., Wadhwa, A., Derr, T., Shah, R.R.: Stock selection via spatiotemporal hypergraph attention network: A learning to rank approach. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 497–504 (2021)
12. Shih, S.Y., Sun, F.K., Lee, H.y.: Temporal pattern attention for multivariate time series forecasting. *Machine Learning* **108**, 1421–1441 (2019)
13. Tang, S., Ma, L., Li, H., Sun, M., Shen, Z.: Bi-mamba: Towards accurate 1-bit state space models. arXiv preprint arXiv:2411.11843 (2024)
14. Wang, H., Peng, J., Huang, F., Wang, J., Chen, J., Xiao, Y.: Micn: Multi-scale local and global context modeling for long-term series forecasting. In: The eleventh international conference on learning representations (2023)
15. Wang, S., Wu, H., Shi, X., Hu, T., Luo, H., Ma, L., Zhang, J.Y., Zhou, J.: Timemixer: Decomposable multiscale mixing for time series forecasting. arXiv preprint arXiv:2405.14616 (2024)
16. Wang, Y., Xu, Y., Yang, J., Wu, M., Li, X., Xie, L., Chen, Z.: Fully-connected spatial-temporal graph for multivariate time-series data. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 15715–15724 (2024)
17. Wu, H., Hu, T., Liu, Y., Zhou, H., Wang, J., Long, M.: Timesnet: Temporal 2d-variation modeling for general time series analysis. arXiv preprint arXiv:2210.02186 (2022)
18. Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., Zhang, C.: Connecting the dots: Multivariate time series forecasting with graph neural networks. In: Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 753–763 (2020)
19. Xu, K., Li, Y., Li, Y., Jiang, X., Duan, J., Wang, Y., Lin, Y.: Spatio-temporal transformer networks for traffic flow forecasting. *IEEE Transactions on Intelligent Transportation Systems* **22**(11), 7137–7147 (2021)
20. Yan, W., Tan, Y.: Double-path adaptive-correlation spatial-temporal inverted transformer for stock time series forecasting. arXiv preprint arXiv:2409.15662 (2024)
21. Yi, K., Fei, J., Zhang, Q., He, H., Hao, S., Lian, D., Fan, W.: Filtarnet: Harnessing frequency filters for time series forecasting. arXiv preprint arXiv:2411.01623 (2024)
22. Zhang, X., Huang, C., Xu, Y., Xia, L., Dai, P., Bo, L., Zhang, J., Zheng, Y.: Traffic flow forecasting with spatial-temporal graph diffusion network. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 15008–15015 (2021)
23. Zhang, Y., Yan, J.: Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In: The Eleventh International Conference on Learning Representations (2022)
24. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W.: Informer: Beyond efficient transformer for long sequence time-series forecasting. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 11106–11115 (2021)
25. Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., Jin, R.: Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In: International conference on machine learning. pp. 27268–27286. PMLR (2022)

26. Zou, J., Han, Y., So, S.S.: Overview of artificial neural networks. Artificial neural networks: methods and applications pp. 14–22 (2009)

Appendix

A Basic Informations

A.1 Datasets Feature

The stock features we utilize include: open, high, low, close, pre_close, change, pct_chg, vol, amount, turnover_rate, turnover_rate_f, and volume_ratio. pe, pe_ttm, pb, ps, ps_ttm, dv_ratio, dv_ttm, total_share, float_share, free_share, total_mv, and circ_mv. buy_sm_vol, buy_sm_amount, sell_sm_vol, sell_sm_amount, buy_md_vol, and buy_md_amount, sell_md_vol, sell_md_amount, buy_lg_vol, buy_lg_amount, sell_lg_vol, sell_lg_amount, buy_elg_vol, buy_elg_amount, sell_elg_vol, sell_elg_amount, net_mf_vol, and net_mf_amount, up_limit, down_limit, ma5, ma10, ma15, ma20, ma25 and industry.

A.2 Experimental Setup and Parameter Settings

Our experimental environment is as follows: Ubuntu 22.04, GPU NVIDIA GeForce RTX 3090 24G, CPU AMD EPYC7282, and 256GB of RAM. We conducted 5 repeated experiments. We set the parameters of HIGSTM as follows: the hidden layer dimension is $d_{model} = 64$, the state dimension is $d_{model} = 16$, the convolution kernel dimension is $d_{conv} = 4$, the macro-level information dimension for each time step is $d_{te} = 32$, the global macro-level information dimension is $d_{ge} = 32$, and the correlation attention dimension is $d_{attn} = 32$. The input time series length is $l_{in} = 16$.

A.3 Trading Strategies

Our simulated trading strategy involves selecting the top 10% of stocks each day, holding them for 10 days, and calculating the returns using equal-weighted simple interest for all stocks, while also taking into account transaction fees.

B Case Study

We present the stocks most correlated with 600519.SH in the temporal section and the global graph, as shown in Figure 4. The numerical codes in the figure correspond to individual stocks in the Chinese stock market. It can be observed that in the sparse temporal section correlation, the diversity of neighbors is low, but the correlations are prominent. The learned neighbors, such as 000596.SZ (Orange node), are stocks from the same liquor industry as 600519.SH, exhibiting




Fig. 4. Case: The neighbors of stock 600519.SH learned in the temporal section and the global graph.

significant correlations. In contrast, the Global Correlation shows higher diversity among neighbors, with no particularly dominant adjacency relationships. For instance, 000596.SZ (Orange node) has a high correlation due to belonging to the same industry as the benchmark stock, 603288.SH (Blue node) is highly correlated as it is also in the beverage sector, 601888.SH (Pink node) is correlated due to its involvement in downstream retail, and 601318.SH (Green node) is correlated as it is another industry leader.

C Daily Return Rate Analysis




Fig. 5. The Box Plot of the model's Daily Return Rate

The box plots of daily returns (Figure 5) show that HIGSTM achieves a significantly higher maximum return compared to other models, with its minimum return also exceeding all except MASTER. HIGSTM exhibits slightly higher return volatility, indicating greater risk, aligning with earlier findings. MASTER, in contrast, is a low-risk, low-return model, with its maximum return notably lower than HIGSTM's.