

Context-Aware Frequency-Embedding Networks for Spatio-Temporal Portfolio Selection

Ruirui Liu* Huichou Huang† Johannes Ruf‡ Qingyao Wu§

Abstract

Recent developments in the applications of deep reinforcement learning methods to portfolio selection have achieved superior performance to conventional methods. However, two major challenges remain unaddressed in these models and inevitably lead to the deterioration of model performance. First, asset characteristics often suffer from low and unstable signal-to-noise ratios, leading to poor learning robustness of the predictive feature representations. Second, the existing literature fails to consider the complexity and diversity in long-term and short-term spatio-temporal predictive relations between the feature sequences and portfolio objectives. To tackle these problems, we propose a novel **Context-Aware Frequency-Embedding Graph Convolution Network (Cafe-GCN)** for spatio-temporal portfolio selection. It contains three important modules: (1) frequency-embedding block that explicitly captures the short-term and long-term predictive information embedded in asset characteristics meanwhile filtering out the noise; (2) context-aware block that learns multiscale temporal dependencies in the feature space; and (3) multi-relation graph convolutional block that exploits both static and dynamic spatial relations among assets. Extensive experiments on two real-world datasets demonstrate that Cafe-GCN consistently outperforms proposed techniques in the literature.

1 Introduction

Portfolio selection [27, 29] is one of the major investment activities in the financial industry. In this task, investors observe the feature sequences of a basket of assets and make sequential portfolio decisions that aim to maximize long-term returns while minimizing portfolio risk. Early studies apply conventional machine learning methods, such as SVMs [14] and random forests [7]. With the rapid development of deep learning, various

deep neural networks are employed to capture the asset dynamics and the temporal dependency of feature sequences for financial problems, such as CNNs [11], RNNs [18], and transformers [12]. Moreover, financial markets exhibit complex and evolving comovement patterns across assets. As a result, spatio-temporal interdependence plays a pivotal role in portfolio selection (see Section A of the Online Appendix for motivation).

More recent studies aim to capture the dynamic spatial interdependence, e.g., using dilated convolutions [39] or attention mechanisms [28] to improve portfolio performance. However, these frameworks cannot be well adapted to the complex characteristics of the financial market. Specifically, the portfolio selection task still faces two main challenges that confine portfolio performance: (1) the asset feature sequences have a low and unstable signal-to-noise ratio, which renders it difficult for the model to learn robust predictive feature representations; (2) the feature sequences tend to exhibit complex long-term and short-term temporal relations as well as diverse spatial relations.

In this paper, we propose a novel model based on deep reinforcement learning (DRL), namely **Context-Aware Frequency-Embedding Graph Convolution Network (Cafe-GCN)**, to address the two mentioned challenges in portfolio selection task. The overall architecture of Cafe-GCN is shown in Figure 1. The Cafe-GCN primarily stacks multiple context-aware frequency-embedding (CAFE) spatio-temporal networks with different dilation rates to learn multiscale and diverse spatio-temporal feature representations that aim to capture the time-varying multi-view spatio-temporal predictive relations between the extracted feature sequences and portfolio objectives. It also considers historical portfolio weights in minimizing transaction costs across trading periods, and then directly outputs the current portfolio weights.

Specifically, the CAFE module designed to capture the spatio-temporal predictive relations of stock feature sequences consists of three key components: (1) The frequency embedding block transforms the stock feature sequences from the time domain to the frequency domain and learns stock features through a fully connected

*Brunel University of London and King’s College London, United Kingdom. Email: ruirui.liu@brunel/kcl.ac.uk.

†City University of Hong Kong and Bayescien Technologies, China. Email: huichou.huang@exeter.oxon.org.

‡London School of Economics and Political Science, United Kingdom. Email: j.ruf@lse.ac.uk.

§South China University of Technology, China. Email: qyw@scut.edu.cn. Corresponding author.

layer (FCL). It aims to remove high-frequency noise and to differentiate long-term predictive information from short-term predictive information, as fast-moving (slow-moving) signals correspond to short-term (long-term) predictability. This block alleviates the problem of low and unstable signal-to-noise ratios of stock feature sequences and increases the learnability of predictive features. (2) Standard attention mechanisms [22] perform matching between query and key vectors pointwise along the time dimension. However, this direct pointwise matching method overlooks the local contextual information and is susceptible to the influence of noisy data points, leading to abnormal capture of temporal features. To overcome this drawback, we propose the context-aware block that employs dilated convolutions to extract and further enhance local (short-term) temporal features and a multiscale attention mechanism to extract and further enhance global (long-term) temporal features. (3) The multi-relation graph convolutional block applies graph convolutional networks to extract multi-relation spatial features that apply industry, sector, and learnable time-varying adjacency matrices to capture both static and dynamic spatial relations among stocks. These three blocks lay a solid foundation for sequential portfolio selection.

Our main contributions are summarized as follows: First, we propose **Cafe-GCN**, a novel model based on deep reinforcement learning (DRL) method for spatio-temporal portfolio selection. Second, we design a CAFE module that establishes stable predictive relations between the feature sequences and the portfolio objective by improving the feature learnability and representational quality. This is achieved by denoising the frequency embedding that differentiates long-term predictive features from short-term predictive features. Third, we further enhance the spatio-temporal features by introducing multi-relation spatial features and capturing multiscale temporal features that consider both static and dynamic spatial relations meanwhile providing further fine-grained local (short-term) and coarse-grained global (long-term) predictive features. Last, extensive experiments on two real-world datasets demonstrate that our proposed model consistently outperforms other state-of-the-art (SOTA) solutions.

2 Related Work

Standard portfolio selection methods [13, 2] usually rely on expert knowledge in finance to make investment decisions using statistical analysis or traditional machine learning techniques. However, these methods heavily depend on specialized financial knowledge and struggle to adapt to the complex and changing financial market environment. With the advancement of deep learning,

numerous data-driven and neural network-based methods have been proposed to address financial problems [38, 10, 16, 8, 17]. These methods typically capture temporal and/or spatial asset characteristics from various multimodal and heterogeneous data sources, including fundamental price data [10, 24], news headlines [33], social media [5], relational graphs [30, 32, 3], etc.

2.1 Sequential Portfolio Selection Sequential portfolio selection methods focus on conducting a time-series analysis from the asset feature sequences that capture temporal features, without considering the interconnections between assets (i.e., spatial relations). Hence, such methods often rely on classical time-series networks, such as temporal convolutional network (TCN) [1], long short-term memory (LSTM) [35], and transformer [22]. Jiang *et al.* [6] apply a deterministic policy gradient-based reinforcement learning framework based on the ensemble of identical independent evaluators (EIIIE) to the portfolio selection task. It employs LSTM and CNN to independently extract temporal features from the stock market data, followed by a softmax layer to directly output the portfolio weights for each stock. Ye *et al.* [31] rely on a pre-trained neural network to extract temporal features from both market numeric and news text data for stock trend prediction and portfolio selection. Qin *et al.* [15] introduce the attention-based LSTM (ALSTM) model that incorporates attention mechanisms into LSTM by allocating attention at different time steps to capture temporal features. Zha *et al.* [37] propose a hierarchical policy-based reinforcement learning model, in which two independent agents respectively rely on macro and micro data to make sequential portfolio decisions.

2.2 Spatio-Temporal Portfolio Selection As spatial relations among assets play a pivotal role in portfolio selection, many studies explicitly consider cross-sectional relations in portfolio decisions. Wang *et al.* [25] propose an LSTM model with historical state attention (LSTM-HA) that introduces the cross-asset attention network (CAAN) to capture the spatial relations in addition to the attention-based temporal feature extraction. Zhang *et al.* [39] employ both LSTM and dilated convolutional networks to construct a dual-pathway model that separately extracts temporal and spatial features from stock sequences. Ye *et al.* [30] rely on graph convolutional networks (GCN) to capture the relations among stocks and gated recurrent units (GRU) to learn temporal features for stock movement prediction. Xu *et al.* [28] propose the relation-aware transformer (RAT), that captures local and global spatio-temporal features at each time point across stocks.

Wang *et al.* [26] rely on a TCN to extract individual stock temporal features and an attention-based GCN to capture spatial relations. Li *et al.* [9] leverage market information to automatically guide the spatio-temporal feature selection. Fan *et al.* [4] propose a simple MLP-based architecture that performs predictor mixing, followed by time mixing and stock mixing.

3 Problem Setting

At time t , given feature sequences \mathbf{X}_t and multi-relation graphs among stocks \mathbf{G} , the portfolio optimization process aims to train a policy network π that maps \mathbf{X}_t and \mathbf{G} to a vector of portfolio weights \mathbf{a}_t , i.e., $\mathbf{a}_t = \pi(\mathbf{X}_t, \mathbf{G})$ to maximize the risk-adjusted returns (see Section B of the Online Appendix).

In the framework of DRL-based portfolio selection, we consider the portfolio optimization process as a generalized Markov Decision Process by $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R})$. At each time t , the policy network (agent) obtains a current state $\mathbf{s}_t = \mathbf{X}_t \in \mathcal{S}$, and then, makes a portfolio decision according to the current state and previous portfolio decision $\mathbf{a}_t = \pi(\mathbf{s}_t, \mathbf{a}_{t-1}) \in \mathcal{A}$ with π denoting the policy network. Finally, the agent gets a reward from the financial market $r_t \in \mathcal{R}$ while the transition to the next state is stochastic $\mathbf{s}_{t+1} \sim \mathcal{T}(\mathbf{s}_t)$. The policy network serves as a representative agent that constantly optimizes its strategy based on the reward sequence.

Under the deep reinforcement learning-based framework of portfolio optimization, the following two general hypotheses are commonly proposed: (1) The investment of the representative agent is traded in a constantly highly liquid market, i.e., any trade can be executed immediately; (2) The investment of the representative agent is fairly small relative to the transactions of the whole market at each point of time and thereby does not influence market prices.

4 Methods

4.1 The Architecture of Cafe-GCN Stock feature sequences suffer from low and unstable signal-to-noise ratio issues and exhibit complex and changing spatio-temporal features. The proposed Cafe-GCN aims to alleviate these issues and capture the denoised multiscale and diverse spatio-temporal features, i.e., the long-term and short-term temporal relations as well as diverse spatial interconnections, for portfolio selection. Figure 1 illustrates the proposed Cafe-GCN architecture. Firstly, the model introduces a multi-relation GCN with static real-world relations in terms of stock industry and sector categorization along with the dynamically learnable GCN-based cross-stock relations. The Cafe-GCN learns multiscale and diverse spatio-temporal interdependence of stock feature sequences by overlaying

multiple CAFE modules with different dilation factors δ . It progressively expands the local receptive field of the CAFE to learn the temporal features from short-term to long-term and enhance the extraction of the global spatio-temporal interdependence. In addition, it incorporates the historical portfolio weights and transaction costs into the portfolio decision-making through a temporal attention layer. Then, the decision-making module takes the local and global spatio-temporal features as inputs to output the portfolio weights. Specifically, it flattens the temporal and feature dimensions of the spatio-temporal features using a flattened layer and then applies FCLs and a softmax layer to derive the portfolio weights. In the DRL-based portfolio selection framework, the Cafe-GCN acts as an investment agent. The portfolio decisions s/he makes are inputs into the stock market environment and generate a risk-adjusted reward function R , and s/he optimizes model parameters based on the reward function through the optimization process.

4.2 Context-Aware Frequency-Embedding Network

We introduce the CAFE module that aims to simultaneously learn the multiscale temporal features of individual stocks and diverse spatial features among stocks from stock feature sequences characterized by low and unstable signal-to-noise ratios. The overall architecture of this module is illustrated in Figure 2. CAFE is composed of three internal components, namely (1) a Frequency Embedding Block (FEB), (2) a Context-Aware Block (CAB), and (3) a Multi-Relation Graph Convolutional Block (MRGCB). Moreover, we also apply the residual technique and layer normalization to construct the CAFE module. As shown in Figure 1, the Cafe-GCN model includes three multiscale CAFE modules to capture different levels and scopes of spatio-temporal features. The first CAFE takes the initial stock feature sequences $\mathbf{X} \in \mathbb{R}^{N \times \Delta T \times P}$ as inputs, and the subsequent CAFE takes the outputs of the previous CAFE as inputs. For convenience, here we uniformly denote the inputs of different CAFE as the stock feature sequences \mathbf{H} . Then, the forward process of CAFE can be represented as

$$\begin{cases} \mathbf{H} := \mathcal{F}_{\text{FEL}}(\mathbf{H}) \\ \mathbf{H} := \mathcal{F}_{\text{LN}}(\mathcal{F}_{\text{CAL}}(\mathbf{H}) + \mathbf{H}) \\ \mathbf{H} := \mathcal{F}_{\text{LN}}(\mathcal{F}_{\text{MRGCL}}(\mathbf{H}) + \mathbf{H}), \end{cases}$$

where \mathcal{F}_{FEL} , \mathcal{F}_{CAL} , $\mathcal{F}_{\text{MRGCL}}$, and \mathcal{F}_{LN} denote the frequency embedding layer, context-aware layer, multi-relation graph convolutional layer, and layer normalization, respectively.




Figure 1: The architecture of the Context-Aware Frequency Embedding (CAFE) Graph Convolution Network (**Cafe-GCN**). Firstly, the model learns the multiscale and diverse spatio-temporal interdependence of stock feature sequences by overlaying multiple CAFE modules with different dilation factors δ . Meanwhile, it incorporates the historical portfolio weights and further extracts the global spatio-temporal interdependence that considers historical portfolio weights and transaction costs through a temporal attention layer. Finally, the decision-making module takes the local and global spatio-temporal features as inputs to output the portfolio weights.

4.2.1 Frequency-Embedding Block The FEB aims to filter high-frequency noise out of the stock feature sequences and differentiate long-term predictive features from short-term predictive features to alleviate the low and unstable signal-to-noise ratio issue, thereby improving the efficiency and effectiveness of the subsequent spatio-temporal feature extraction based on the ground that fast-moving (slow-moving) signals correspond to short-term (long-term) predictability. Specifically, given the stock feature sequences $\mathbf{H} \in \mathbb{R}^{N \times \Delta T \times D}$, the FEB first transforms \mathbf{H} from the time domain into the frequency domain using Discrete Fourier Transform (DFT) [21]. Here D denotes the hidden dimension of the model, and $\Delta T/2 + 1$ the size of the frequency dimension. Then the FEB employs a Low Pass Filter (LPF) to filter out high-frequency noise from the stock feature sequences. The above process can be represented as $\mathbf{H}_C = \text{FCL}(\text{LPF}(\text{DFT}(\mathbf{H})))$, where FCL denotes a fully-connected layer and $\mathbf{H}_C \in \mathbb{C}^{N \times (\Delta T/2 + 1 - M) \times D}$ are the stock feature sequences in the frequency dimension. Here M is the number of frequencies filtered out by the LPF.

After the LPF removes M high-frequency features from the stock features in the frequency dimension, we transform the frequency domain back to the temporal form with length ΔT through the Inverse Discrete Fourier Transform (IDFT). To this end, it is necessary to perform padding zero on the high-frequency part of the frequency domain features before apply-

ing IDFT. This process can be expressed as $\mathbf{H} = \text{FCL}(\text{IDFT}(\text{Padding}(\mathbf{H}_C)))$, where $\mathbf{H} \in \mathbb{R}^{N \times \Delta T \times D}$ is the final output of the FEB.

4.2.2 Context-Aware Block Standard attention mechanisms [22] involve three FCLs in the feature dimension to learn query, key, and value vectors separately. Subsequently, feature matching between query and key vectors is performed pointwise along the time dimension. However, this direct pointwise matching method overlooks the local contextual information and is susceptible to the influence of noisy data points, leading to the abnormal capture of temporal features. To overcome this drawback, the CAB block employs dilated convolutions [34] with a kernel size of 3 and a dilation rate of δ in the temporal dimension, replacing fully connected layers to learn the key vector $\mathbf{K} = \text{Dconv}_K(\mathbf{H}) \in \mathbb{R}^{N \times \Delta T \times D}$ and the value vector $\mathbf{V} = \text{Dconv}_V(\mathbf{H}) \in \mathbb{R}^{N \times \Delta T \times D}$, where Dconv denotes dilated convolution. The key and value vectors learned in this manner can capture the temporal information of the local context within the receptive field and mitigate the influence of noisy data points. Moreover, the dilated convolutions are capable of controlling the size of the local receptive field in the time dimension by adjusting the dilation rate δ , enabling the model to construct different key and value vectors for learning multiscale local temporal features.

This method first uses the dilated convolutions

to capture the key and value vectors containing local contextual information, aiming to mitigate the impact of noisy data points and introduce local contextual information. Then an FCL is applied to learn the query vectors $\mathbf{Q} = \mathbf{HW}_Q \in \mathbb{R}^{N \times \Delta T \times D}$. Subsequently, the CAB takes the pointwise matching of query vector \mathbf{Q} , context-aware key vectors, and value vectors as inputs to learn the global temporal features based on the attention mechanism

$$\mathbf{H} = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}} \right) \mathbf{V}.$$

4.2.3 Multi-Relation Graph Convolutional Block To learn the diverse spatial relations among stocks, the proposed MRGCB incorporates a dynamic learnable relation graph into the static realistic relation graph among stocks. It relies on graph convolutional networks to learn diverse relational features among stocks, as illustrated in component (3) of Figure 2.

We use static realistic relation graphs, including the industry relation graph $G_{\text{ind}}(V, E)$ and the sector relation graph $G_{\text{sec}}(V, E)$, where $v \in V$ represents stocks and $e \in E$ represents the relations between stocks. In the industry/sector relation graph, if stocks v_1 and v_2 belong to the same industry/sector, it establishes an edge connection between the two nodes. These static relation graphs categorize stocks based on their real-world connections and domain expert knowledge. However, these static relations cannot capture the fluctuations or temporary shifts in the relations among stocks. Therefore, we further introduces a learnable stock relation graph $G_{\text{dyn}}(V, E)$, which can be transformed into a dynamically updatable parameterized adjacency matrix $\mathbf{A}_{\text{dyn}} \in \mathbb{R}^{N \times N}$. This component relies on the dynamic adjacency matrix to learn the dynamic relations among stocks. Similarly, static relation graphs can also be transformed into static industry adjacency matrix $\mathbf{A}_{\text{ind}} \in \mathbb{R}^{N \times N}$ and sector adjacency matrix $\mathbf{A}_{\text{sec}} \in \mathbb{R}^{N \times N}$. Given these three types of relation graphs and the stock features $\mathbf{H} \in \mathbb{R}^{N \times \Delta T \times D}$, the first step is to reshape \mathbf{H} into $\mathbb{R}^{\Delta T \times N \times D}$, then using the graph convolutional networks to extract diverse spatial features. This forward computation process can be written as $\mathbf{H} = \mathbf{A}_{\text{dyn}} \hat{\mathbf{A}}_{\text{sec}} \hat{\mathbf{A}}_{\text{ind}} \mathbf{H}\mathbf{W} + \mathbf{b}$, where $\mathbf{W} \in \mathbb{R}^{D \times D}$ and $\mathbf{b} \in \mathbb{R}^D$ are the parameters of the linear transformation, whereas the normalization formulas for the industry adjacency matrix and the sector adjacency matrix are denoted by $\hat{\mathbf{A}}_{\text{ind}} = \mathbf{D}_{\text{ind}}^{-\frac{1}{2}} \mathbf{A}_{\text{ind}} \mathbf{D}_{\text{ind}}^{-\frac{1}{2}}$ and $\hat{\mathbf{A}}_{\text{sec}} = \mathbf{D}_{\text{sec}}^{-\frac{1}{2}} \mathbf{A}_{\text{sec}} \mathbf{D}_{\text{sec}}^{-\frac{1}{2}}$, respectively, and \mathbf{D} is the degree matrix of the adjacency matrix. Then \mathbf{H} is reshaped into $\mathbb{R}^{N \times \Delta T \times D}$.

4.3 The Decision-Making Process of Cafe-GCN

The context-aware frequency-embedding graph convolution network takes the stock feature sequences $\mathbf{X} \in \mathbb{R}^{N \times \Delta T \times P}$ as inputs, stacking three CAFE modules with different dilation factors δ to capture the multiscale and diverse spatio-temporal features $\mathbf{H} \in \mathbb{R}^{N \times \Delta T \times D}$. Subsequently, the Cafe-GCN model introduces historical portfolio weights $\mathbf{a}_h \in \mathbb{R}^{N \times \Delta T \times 1}$ that consider transaction costs to compute the current portfolio weights. Specifically, the spatio-temporal features \mathbf{H} are concatenated with the historical portfolio weights \mathbf{a}_h . Then, the model learns the query vector \mathbf{Q} , key vector \mathbf{K} , and value vector \mathbf{V} using three FCLs, and captures the spatio-temporal interdependence that consider historical portfolio weights and transaction costs through a standard temporal attention layer $\mathbf{H} = \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$.

In the decision-making module, Cafe-GCN flattens the temporal and feature dimensions of the spatio-temporal feature space and constructs portfolio weights $\mathbf{a} \in \mathbb{R}^N$ through an FCL and a softmax layer as $\mathbf{a} = \text{softmax}(\text{FCL}(\text{Flatten}(\mathbf{H})))$.

4.4 Reinforcement Learning for Portfolio Selection

The proposed Cafe-GCN model is based on a deep reinforcement learning framework, which interacts with the stock market environment to optimize model parameters. The reinforcement learning optimization method used in this paper is the Deterministic Policy Gradient (DPG) as in [20]. The DPG method optimizes model parameters directly based on a pre-designed reward (loss) function using gradient ascent (descent). Therefore, designing a reasonable reward function is crucial to optimizing the model effectively and efficiently. We introduce $u(t) = \mathbf{a}_t^\top \mathbf{y}_t - c_t - \eta \frac{\mathbf{a}_t^\top \boldsymbol{\Sigma}_t \mathbf{a}_t}{2}$, where the first term in the above equation considers the investment returns after transaction costs, and the second term is used for measuring investment risks. Here, \mathbf{a}_t represents the portfolio weight vector at time t , \mathbf{y}_t is the vector of stock returns in the next period, c_t denotes the transaction cost ratio, and $\boldsymbol{\Sigma}_t \in \mathbb{R}^{N \times N}$ is the covariance matrix of stock returns at time t . The parameter η is the risk aversion coefficient, where a higher η indicates a stronger aversion to investment risk, leading to a more conservative investment style. Conversely, a lower η implies a more aggressive investment style. When η equals 1 with the reward function assumed to be logarithmic, it transforms to the Sharpe Ratio (SR) [19] as a quadratic utility function of the first two moments of the portfolio return distribution. The average risk-adjusted reward over T periods is then given by $\mathcal{R} = \frac{1}{T} \sum_{t=1}^T u(t)$.

We take the negative of the reward function to calculate the loss function $\mathcal{L} = -\mathcal{R}$. Subsequently,

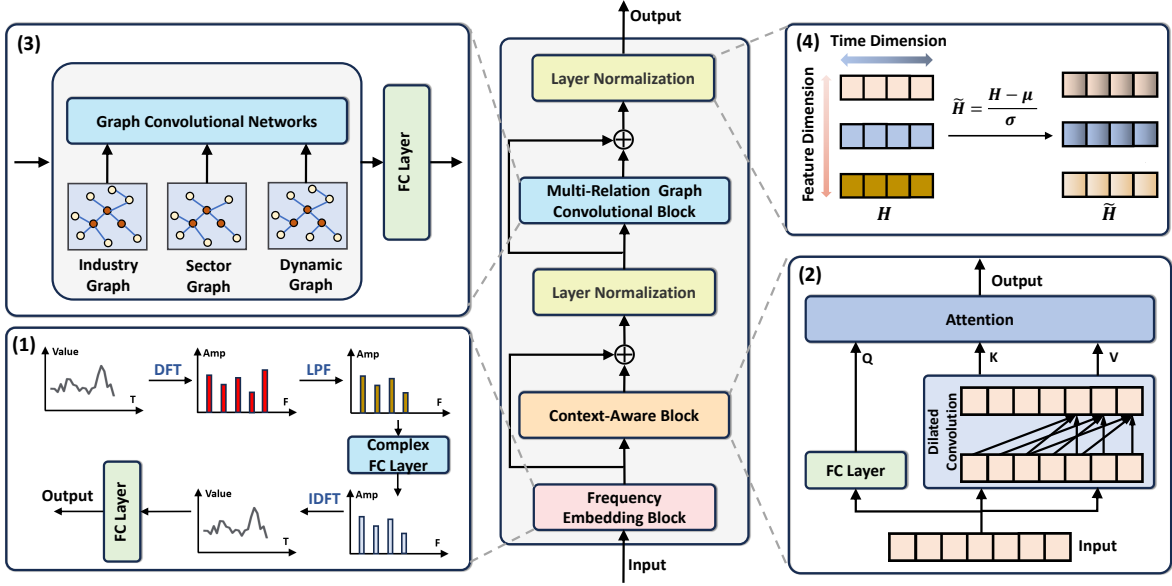


Figure 2: The architecture of the CAFE module. It has three components: FEB, CAB, and MRGCB. In addition, the residual technique and layer normalization are applied.

it uses gradient descent instead of gradient ascent to optimize the model parameters.

5 Experiment

5.1 Experimental Settings To evaluate the performance of our proposed Cafe-GCN model, we conduct comprehensive experiments on two real-world stock datasets, the CSI100 and NASDAQ100 datasets. The two stock datasets are collected from the China A-share market and the U.S. NASDAQ market, respectively. In this section, we first provide information about the datasets, the data preprocessing process, and the evaluation metrics used for the experiments. Then we briefly introduce the baseline models and the implementation details.

5.1.1 Datasets and Preprocessing Each of the CSI100 and NASDAQ100 datasets consists of sequential features of 100 stocks from 2011 to 2021. The raw features of each stock comprise four basic market prices: open, high, low, and close price. Industry and sector information for each stock is collected through the Wind and Bloomberg terminals. Both datasets are divided into training sets from March 2011 to February 2017, validation sets from March 2017 to February 2019, and test sets from March 2019 to February 2021. The CSI100 dataset spans 2,433 trading days and the NASDAQ100 dataset spans 2,517 trading days.

5.1.2 Baselines We compare the Cafe-GCN model with two major groups of SOTA baseline methods within the DRL-based framework for portfolio selection, including: (1) **Sequential Portfolio Selection:** EIEE [6], ALSTM [15], SFM [38], DLinear [36], SARL [31]; (2) **Spatio-Temporal Portfolio Selection:** GAT [23], AlphaStock [25], RAT [28], StockMixer [4], MASTER [9], PPN [39], MGCGRU [30], DeepTrader [26].

5.1.3 Evaluation Metrics DRL-based portfolio selection methods output portfolio weights for each trading period t based on stock returns \mathbf{y}_t . The portfolio return at time t can be calculated as $r_t = \mathbf{a}_t^\top \mathbf{y}_t - c_t$. Therefore, in a test set with T trading period, the model can continuously output a series of portfolio weights and obtain returns over multiple periods $\mathbf{R} = [r_1, r_2, \dots, r_T]$. We evaluate our proposed model and competing baselines by four performance metrics: Annualized Return (AR), Sharpe Ratio (SR), Maximum DrawDown (MDD), and Calmar Ratio (CR) (see Section D of the Online Appendix)

5.2 Performance Comparisons In this section, we compare the portfolio selection performance of our proposed Cafe-GCN with that of baseline models using four evaluation metrics from the experiments with two real-world datasets (see Section C of the Online Appendix for implementations). Table 1 presents the experimental results of each model, where the \uparrow symbol next to each metric indicates that the higher value of the metric,

Table 1: Experiment Results on CSI100 and NASDAQ100 Datasets.

	CSI100				NASDAQ100			
	AR \uparrow	SR \uparrow	MDD \downarrow	CR \uparrow	AR \uparrow	SR \uparrow	MDD \downarrow	CR \uparrow
ALSTM	0.134 \pm 0.005	0.641 \pm 0.034	0.451 \pm 0.001	0.303 \pm 0.026	1.577 \pm 0.062	0.831 \pm 0.004	1.161 \pm 0.079	1.361 \pm 0.038
EIIE	0.169 \pm 0.010	0.381 \pm 0.043	0.512 \pm 0.025	0.332 \pm 0.035	1.326 \pm 0.133	1.122 \pm 0.126	1.096 \pm 0.068	1.222 \pm 0.197
SFM	0.171 \pm 0.010	0.472 \pm 0.022	0.498 \pm 0.013	0.343 \pm 0.072	1.937 \pm 0.013	1.072 \pm 0.037	0.787 \pm 0.004	2.461 \pm 0.005
DLinear	0.164 \pm 0.001	0.460 \pm 0.003	0.389 \pm 0.013	0.467 \pm 0.071	2.128 \pm 0.057	0.863 \pm 0.132	0.964 \pm 0.073	2.437 \pm 0.105
SARL	0.104 \pm 0.028	0.448 \pm 0.055	0.434 \pm 0.031	0.253 \pm 0.056	2.131 \pm 0.029	1.006 \pm 0.145	0.593 \pm 0.059	3.593 \pm 0.068
GAT	0.205 \pm 0.005	0.698 \pm 0.026	0.428 \pm 0.006	0.463 \pm 0.064	3.505 \pm 0.238	0.991 \pm 0.082	1.078 \pm 0.006	3.250 \pm 0.241
AlphaStock	0.198 \pm 0.007	0.484 \pm 0.018	0.488 \pm 0.008	0.405 \pm 0.072	4.379 \pm 0.275	1.089 \pm 0.120	0.578 \pm 0.084	7.810 \pm 1.617
RAT	0.215 \pm 0.014	0.887 \pm 0.036	0.451 \pm 0.013	0.471 \pm 0.080	4.129 \pm 0.044	1.104 \pm 0.036	1.028 \pm 0.060	4.031 \pm 0.280
StockMixer	0.229 \pm 0.015	0.1512 \pm 0.012	0.368 \pm 0.014	0.627 \pm 0.102	4.576 \pm 0.497	1.087 \pm 0.018	1.284 \pm 0.178	3.687 \pm 0.901
MASTER	0.249 \pm 0.002	0.982 \pm 0.026	0.476 \pm 0.008	0.524 \pm 0.097	4.268 \pm 0.017	0.818 \pm 0.003	0.939 \pm 0.071	4.571 \pm 0.362
PPN	0.301 \pm 0.010	0.748 \pm 0.006	0.594 \pm 0.007	0.507 \pm 0.011	4.775 \pm 0.017	0.747 \pm 0.021	0.736 \pm 0.064	6.488 \pm 0.329
MGCGRU	0.321 \pm 0.019	0.556 \pm 0.027	0.411 \pm 0.063	0.795 \pm 0.073	5.191 \pm 0.185	0.811 \pm 0.125	0.582 \pm 0.623	8.916 \pm 0.394
DeepTrader	0.365 \pm 0.018	0.966 \pm 0.031	0.414 \pm 0.007	0.905 \pm 0.031	6.908 \pm 0.744	0.955 \pm 0.028	0.735 \pm 0.099	9.429 \pm 0.263
Cafe-GCN	0.404 \pm 0.003	1.060 \pm 0.060	0.330 \pm 0.018	1.228 \pm 0.078	6.251 \pm 0.045	1.310 \pm 0.006	0.460 \pm 0.003	13.590 \pm 0.178

the better performance of the model, while the \downarrow symbol means that the lower value of the metric, the better performance of the model. The best result for each metric is highlighted in bold. To run a credible and fair horse race, we train, validate, and test each model five times and take the average figures as the final results. The figures following the \pm symbol in the table represent the results' standard deviation of five experiments.

Table 1 presents the performance of Cafe-GCN and various competing baseline models from the experiments on the CSI100 and NASDAQ100 datasets. First, we find that the performance of the sequential portfolio selection methods (i.e., ALSTM, EIIE, SFM, DLinear, and SARL) is overall inferior to that of the spatio-temporal portfolio selection methods (i.e., GAT, AlphaStock, RAT, StockMixer, MASTER, PPN, MGCGRU, DeepTrader) across the four evaluation metrics in both datasets. This can be attributed to the fact that spatio-temporal models not only leverage the temporal dependencies in stock feature sequences but also capture spatial dependencies among stocks using graph convolution and/or attention mechanisms. These findings validate the significant role of capturing spatial relations among stocks for portfolio selection. As shown in Table 1, we can see from the performance of the models focusing on spatio-temporal relations that the models that aim to learn the dynamic spatio-temporal features of stock sequences, e.g., AlphaStock, PPN, and RAT, notably outperform those that rely on static relations, e.g., GAT and MGCGRU, in both datasets. It is evident that capturing the dynamic spatio-temporal interdependence is the key to satisfying portfolio performance.

As we can see from Table 1, our proposed Cafe-GCN significantly outperforms all competing baselines across all four evaluation metrics and two datasets except for the AR in NASDAQ100 (see Section E of the Online Appendix for the visualization of accumulative

wealth). Although the Cafe-GCN does not obtain the best result in the AR in NASDAQ100, it achieves the lowest MDD and the highest CR and SR, suggesting that it still has the highest mean-variance efficiency among all models, i.e., for a unit of risk in terms of volatility, it always achieves the highest return. Its SR is 16.8% (7.9%) higher than the runner-up model and its MDD is 20.4% (10.3%) lower than the runner-up model in the NASDAQ100 (CSI100) sample. Moreover, its CR figures imply that the annualized return is expected to be 13.6 (1.2) times the maximum drawdown of the portfolio strategy in the NASDAQ100 (CSI100) sample. The proposed Cafe-GCN shows strong risk mitigation capacity. These observations reveal the compelling robustness of the proposed Cafe-GCN in comparison with the competing baselines, in particular, in the downside risk protection of the portfolio strategy.

5.3 Ablation Studies This section presents extensive ablation experiments to validate the effectiveness of each component proposed in the Cafe-GCN. Table 2 shows the ablation experiment results of Cafe-GCN on both datasets. The symbol * indicates the model setups of Cafe-GCN.

Frequency-Embedding Block. The proposed FEB relies on the discrete Fourier transformation to convert stock sequences from the time domain to the frequency domain, aiming to remove high-frequency noise in stock sequences and differentiate long-term predictive features from short-term predictive features. It alleviates the negative impact of noise on subsequent spatio-temporal feature extraction procedures and helps to establish robust spatio-temporal relations between the stock features and portfolio objectives. In the table, the option FEB indicates the use of the component while w/o indicates that the component is not used. Compared to the setup that does not use the FEB,

Table 2: Ablation study of Context-Aware Frequency-Embedding Graph Convolution Network

		CSI100				NASDAQ100			
		AR \uparrow	SR \uparrow	MDD \downarrow	CR \uparrow	AR \uparrow	SR \uparrow	MDD \downarrow	CR \uparrow
FEB	w/o FEB*	0.388 \pm 0.006	0.997 \pm 0.031	0.380 \pm 0.024	1.019 \pm 0.093	5.886 \pm 0.133	1.106 \pm 0.001	0.478 \pm 0.010	12.305 \pm 0.126
		0.404 \pm 0.003	1.060 \pm 0.060	0.330 \pm 0.018	1.228 \pm 0.078	6.251 \pm 0.045	1.310 \pm 0.006	0.460 \pm 0.003	13.590 \pm 0.178
CAB	Atten CAB*	0.382 \pm 0.017	0.973 \pm 0.044	0.412 \pm 0.028	0.983 \pm 0.051	5.973 \pm 0.031	1.231 \pm 0.126	0.541 \pm 0.009	11.981 \pm 0.254
		0.404 \pm 0.003	1.060 \pm 0.060	0.330 \pm 0.018	1.228 \pm 0.078	6.251 \pm 0.045	1.310 \pm 0.006	0.460 \pm 0.003	13.590 \pm 0.178
MRGCB	w/o Ind	0.386 \pm 0.009	1.013 \pm 0.054	0.364 \pm 0.024	1.060 \pm 0.063	6.229 \pm 0.066	1.152 \pm 0.084	0.524 \pm 0.052	12.062 \pm 0.891
	w/o Sec	0.369 \pm 0.012	0.992 \pm 0.063	0.353 \pm 0.089	1.043 \pm 0.017	6.237 \pm 0.103	0.980 \pm 0.031	0.489 \pm 0.013	12.755 \pm 0.236
	w/o Dyn	0.305 \pm 0.004	0.752 \pm 0.105	0.335 \pm 0.021	0.913 \pm 0.067	5.731 \pm 0.021	0.973 \pm 0.038	0.538 \pm 0.006	11.753 \pm 0.153
	All*	0.404 \pm 0.003	1.060 \pm 0.060	0.330 \pm 0.018	1.228 \pm 0.078	6.251 \pm 0.045	1.310 \pm 0.006	0.460 \pm 0.003	13.590 \pm 0.178

its inclusion in the Cafe-GCN leads to considerable improvements across four performance metrics in both datasets. This observation confirms the usefulness of this component and justifies the point of view that fast-moving (slow-moving) signals correspond to short-term (long-term) predictability.

Context-Aware Block. The proposed CAB aims to enhance the ability to learn multiscale temporal features and mitigate the impact of noise in stock sequences on point-to-point matching in standard attention mechanism. To validate the importance of CAB in the Cafe-GCN, we compare experiments using CAB and those using the standard attention mechanism (‘Atten’ in Table 2) to learn the temporal features. The results demonstrate that using CAB to learn multiscale temporal information of stock sequences is essential for improving portfolio returns and reducing portfolio risks. In particular, performance metrics such as MDD and CR that measure the downside risk indicate that compared to using standard attention mechanism, CAB achieves significant improvements in risk mitigation in both datasets. These findings confirm the effectiveness of CAB in further reducing the negative impacts of noisy data points and generating high-quality multiscale temporal predictive features for the portfolio selection task.

Multi-Relation Graph Convolutional Block.

To capture diverse spatial relations among stocks, we propose the MRGCB. This component introduces industry and sector relation graphs and learns the time-varying relations among stocks through dynamic relation graphs. It integrates the features of correlated stocks through graph convolution networks to enhance the informativeness and responsiveness of spatial features for portfolio decision-making. To validate the effectiveness of the static industry and sector relation graphs as well as the dynamic relation graph in MRGCB, we compare four different model setups. Specifically, ‘w/o Ind’, ‘w/o Sec’, and ‘w/o Dyn’ indicate scenarios where the model does not use the industry relation graph, sector relation graph, and dynamic relation graph, respectively, to learn spatial relations. ‘All’

in the table represents the scenario where all relation graphs are used.

Compared to using all relation graphs, the portfolio performance of the Cafe-GCN without industry or sector relation graphs deteriorates in both datasets. It is worth noticing that when the dynamic relation graph is not used, the portfolio performance of the Cafe-GCN considerably declines by larger margins. Specifically, the average decrease in AR, SR, and CR metrics across both datasets is 16.4%, 27.3%, and 19.5%, respectively. As for the MDD metric, it performs similarly in the CSI100 dataset but drops by 16.9% in the NASDAQ100 dataset. Without dynamic relation graphs, the model can only build spatial relations through static relation graphs, which obviously cannot adapt to the evolving financial market states. These observations emphasize the paramount importance of the time-varying spatial relations. Overall, all three types of relation graphs in the MRGCB help the model learn diverse spatial relations among stocks to improve portfolio performance.

6 Conclusion

In this paper, we propose a novel context-aware frequency-embedding graph convolutional network for spatio-temporal portfolio selection. This method transforms the stock feature sequences using a frequency-embedding block that not only removes high-frequency noise but also distinguishes long-term predictive features from short-term predictive features. In addition, the context-aware block extracts comprehensive multiscale temporal features while the multi-relation graph convolutional block captures diverse static and dynamic spatial relations. Our proposed method significantly outperforms other baseline methods in two real-world datasets. Extensive ablation experiments demonstrate the effectiveness of different components in the proposed model.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) 62272172, Guangdong Basic and Applied Basic Research Foundation 2023A1515012920, and Zhuhai Science and Technology Plan Project (2320004002758).

References

- [1] S. BAI, J. Z. KOLTER, AND V. KOLTUN, *An empirical evaluation of generic convolutional and recurrent networks for sequence modeling*, arXiv preprint arXiv:1803.01271, (2018).
- [2] F. BLACK AND R. LITTELMAN, *Global portfolio optimization*, Financial Analysts Journal, 48 (1992), pp. 28–43.
- [3] C. CUI, X. LI, C. ZHANG, W. GUAN, AND M. WANG, *Temporal-relational hypergraph tri-attention networks for stock trend prediction*, Pattern Recognition, 143 (2023), p. 109759.
- [4] J. FAN AND Y. SHEN, *StockMixer: A simple yet strong MLP-based architecture for stock price forecasting*, in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, 2024, pp. 8389–8397.
- [5] F. FENG, X. HE, X. WANG, C. LUO, Y. LIU, AND T.-S. CHUA, *Temporal relational ranking for stock prediction*, ACM Transactions on Information Systems (TOIS), 37 (2019), pp. 1–30.
- [6] Z. JIANG, D. XU, AND J. LIANG, *A deep reinforcement learning framework for the financial portfolio management problem*, arXiv preprint arXiv:1706.10059, (2017).
- [7] L. KHAIDEM, S. SAHA, AND S. R. DEY, *Predicting the direction of stock market prices using random forest*, arXiv preprint arXiv:1605.00003, (2016).
- [8] J. LEE, R. KIM, S.-W. YI, AND J. KANG, *Maps: multi-agent reinforcement learning-based portfolio management system*, in Proceedings of the International Conference on International Joint Conferences on Artificial Intelligence, 2021, pp. 4520–4526.
- [9] T. LI, Z. LIU, Y. SHEN, X. WANG, H. CHEN, AND S. HUANG, *Master: Market-guided stock transformer for stock price forecasting*, in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, 2024, pp. 162–170.
- [10] Z. LIANG, H. CHEN, J. ZHU, K. JIANG, AND Y. LI, *Adversarial deep reinforcement learning in portfolio management*, arXiv preprint arXiv:1808.09940, (2018).
- [11] G. LIU, Y. MAO, Q. SUN, H. HUANG, W. GAO, X. LI, J. SHEN, R. LI, AND X. WANG, *Multi-scale two-way deep neural network for stock trend prediction*, in Proceedings of the International Conference on International Joint Conferences on Artificial Intelligence, 2021, pp. 4555–4561.
- [12] J. LIU, H. LIN, X. LIU, B. XU, Y. REN, Y. DIAO, AND L. YANG, *Transformer-based capsule network for stock movement prediction*, in Proceedings of the 1st Workshop on Financial Technology and Natural Language Processing, 2019, pp. 66–73.
- [13] H. MARKOWITZ, *Portfolio selection*, The Journal of Finance, 7 (1952), pp. 77–91.
- [14] R. K. NAYAK, D. MISHRA, AND A. K. RATH, *A naïve SVM-KNN based stock market trend reversal analysis for indian benchmark indices*, Applied Soft Computing, 35 (2015), pp. 670–680.
- [15] Y. QIN, D. SONG, H. CHENG, W. CHENG, G. JIANG, AND G. W. COTTRELL, *A dual-stage attention-based recurrent neural network for time series prediction*, in Proceedings of the International Joint Conference on Artificial Intelligence, 2017, pp. 2627–2633.
- [16] J. RUF AND W. WANG, *Neural networks for option pricing and hedging: A literature review*, Journal of Computational Finance, 24 (2020), pp. 1–46.
- [17] ———, *Hedging with linear regressions and neural networks*, Journal of Business & Economic Statistics, 40 (2022), pp. 1442–1454.
- [18] R. SAWHNEY, S. AGARWAL, A. WADHWA, AND R. SHAH, *Deep attentive learning for stock movement prediction from social media text and company correlations*, in Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2020, pp. 8415–8426.
- [19] W. F. SHARPE, *The Sharpe ratio*, The Journal of Portfolio Management, (1998), pp. 169–185.
- [20] D. SILVER, G. LEVER, N. HEES, T. DEGRIS, D. WIERSTRA, AND M. RIEDMILLER, *Deterministic policy gradient algorithms*, in Proceedings of International Conference on Machine Learning, PMLR, 2014, pp. 387–395.
- [21] D. SUNDARARAJAN, *The Discrete Fourier Transform: Theory, Algorithms and Applications*, World Scientific, 2001.
- [22] A. VASWANI, N. SHAZEER, N. PARMAR, J. USZKOREIT, L. JONES, A. N. GOMEZ, Ł. KAISER, AND I. POLOSUKHIN, *Attention is all you need*, Advances in Neural Information Processing Systems, 30 (2017).
- [23] P. VELIČKOVIĆ, G. CUCURULL, A. CASANOVA, A. ROMERO, P. LIO, AND Y. BENGIO, *Graph attention networks*, arXiv preprint arXiv:1710.10903, (2017).
- [24] H. WANG, S. LI, T. WANG, AND J. ZHENG, *Hierarchical adaptive temporal-relational modeling for stock trend prediction*, in Proceedings of the International Joint Conference on Artificial Intelligence, 2021, pp. 3691–3698.
- [25] J. WANG, Y. ZHANG, K. TANG, J. WU, AND Z. XIONG, *Alphastock: A buying-winners-and-selling-losers investment strategy using interpretable deep reinforcement attention networks*, in Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 1900–1908.
- [26] Z. WANG, B. HUANG, S. TU, K. ZHANG, AND L. XU, *DeepTrader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions embedding*, in Proceedings of the

- AAAI Conference on Artificial Intelligence, vol. 35, 2021, pp. 643–650.
- [27] W. XI, Z. LI, X. SONG, AND H. NING, *Online portfolio selection with predictive instantaneous risk assessment*, Pattern Recognition, 144 (2023), p. 109872.
- [28] K. XU, Y. ZHANG, D. YE, P. ZHAO, AND M. TAN, *Relation-aware transformer for portfolio policy learning*, in Proceedings of the International Joint Conferences on Artificial Intelligence, 2021, pp. 4647–4653.
- [29] P.-Y. YANG, Z.-R. LAI, X. WU, AND L. FANG, *Trend representation based log-density regularization system for portfolio optimization*, Pattern Recognition, 76 (2018), pp. 14–24.
- [30] J. YE, J. ZHAO, K. YE, AND C. XU, *Multi-graph convolutional network for relationship-driven stock movement prediction*, in Proceedings of the International Conference on Pattern Recognition, IEEE, 2021, pp. 6702–6709.
- [31] Y. YE, H. PEI, B. WANG, P.-Y. CHEN, Y. ZHU, J. XIAO, AND B. LI, *Reinforcement-learning based portfolio management with augmented asset movement prediction states*, in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, 2020, pp. 1112–1119.
- [32] T. YIN, C. LIU, F. DING, Z. FENG, B. YUAN, AND N. ZHANG, *Graph-based stock correlation and prediction for high-frequency trading systems*, Pattern Recognition, 122 (2022), p. 108209.
- [33] J. YOO, Y. SOUN, Y.-C. PARK, AND U. KANG, *Accurate multivariate stock movement prediction via data-axis transformer with multi-level contexts*, in Proceedings of the ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021, pp. 2037–2045.
- [34] F. YU AND V. KOLTUN, *Multi-scale context aggregation by dilated convolutions*, arXiv preprint arXiv:1511.07122, (2015).
- [35] Y. YU, X. SI, C. HU, AND J. ZHANG, *A review of recurrent neural networks: LSTM cells and network architectures*, Neural computation, 31 (2019), pp. 1235–1270.
- [36] A. ZENG, M. CHEN, L. ZHANG, AND Q. XU, *Are transformers effective for time series forecasting?*, in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, 2023, pp. 11121–11128.
- [37] L. ZHA, L. DAI, T. XU, AND D. WU, *A hierarchical reinforcement learning framework for stock selection and portfolio*, in Proceedings of the International Joint Conference on Neural Networks, IEEE, 2022, pp. 1–7.
- [38] L. ZHANG, C. AGGARWAL, AND G.-J. QI, *Stock price prediction via discovering multi-frequency trading patterns*, in Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2017, pp. 2141–2149.
- [39] Y. ZHANG, P. ZHAO, Q. WU, B. LI, J. HUANG, AND M. TAN, *Cost-sensitive portfolio selection via deep reinforcement learning*, IEEE Transactions on Knowledge & Data Engineering, 34 (2020), pp. 236–248.

Online Appendix for *Context-Aware Frequency-Embedding Networks for Spatio-Temporal Portfolio Selection*

Ruirui Liu* Huichou Huang† Johannes Ruf‡ Qingyao Wu§

A Motivation Demonstration

As shown in Figure 3, Stock B and Stock C are classified in the same industry and sector. Their feature sequences exhibit compelling similarities in trends and fluctuations, depending on the measuring domain and scale. Capturing these spatio-temporal features is critical to reveal the predictive relations between the feature sequences and portfolio objectives.

B Details for Problem Setting

We consider N ‘stocks’, one of them representing a cash asset.

Feature Sequences: At time t , the features of stock i are denoted by $\mathbf{x}_{t,i} \in \mathbb{R}^P$. The features of all stocks are denoted by $\mathbf{x}_t = [\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N}] \in \mathbb{R}^{N \times P}$. The feature sequences are denoted by $\mathbf{X}_t = \{\mathbf{x}_{t-\Delta T}, \dots, \mathbf{x}_{t-1}\} \in \mathbb{R}^{N \times \Delta T \times P}$, where ΔT is the look-back window size.

Stock Relation Graphs: The relationship between stocks is described by a graph $G(V, E)$. Each node $v \in V$ represents a stock, and each edge $e \in E$ between two nodes means that the two corresponding stocks are related. When the graph is a weighted graph, the edge weights represent the degree of relation between the two stocks.

Stock Returns: The percentage change in price at time t , i.e., 1-period stock returns, are represented as $\mathbf{y}_t = \frac{\mathbf{x}_{t+1}^c}{\mathbf{x}_t^c} - 1 \in \mathbb{R}^N$, where \mathbf{x}_t^c denotes the stock’s closing price.

Portfolio Weights: The portfolio decision at time t is represented by a vector of portfolio weights $\mathbf{a}_t = [a_{t,1}, \dots, a_{t,N}] \in \mathbb{R}_+^N$, where $a_{t,i} \geq 0$ and $\sum_{i=1}^N a_{t,i} = 1$. Here, $a_{t,i}$ denotes the proportion of wealth invested in the stock i at time t . We assume $a_{t,i} \geq 0$, i.e., do not allow short-selling.

Portfolio Returns: The portfolio returns after transaction costs at time t are given by $r_t = \mathbf{a}_t \mathbf{y}_t^\top - c_t$, where c_t is the product of the transaction costs in basis




Figure 3: Stock B and stock C are classified in the same industry and sector. Their feature sequences exhibit similar trends and fluctuations, and we aim for the model to capture as many similarities as possible in the dynamic spatio-temporal features of these two stocks.

points (bps) and the changes in portfolio weights.

C Implementation Details

The proposed Cafe-GCN and baseline models are implemented with PyTorch. For the experimental parameter settings, the lookback window size of feature sequence ΔT is set to 30, and the portfolio trading cost is set to 30 bps (recall that short-selling is not allowed). Regarding the design of hyperparameters for Cafe-GCN, the hidden dimension of the model D is set to 32, the number of frequencies filtered out by the LPF M is set to 3, the kernel size of the dilated convolution is 3, and the dilation rates of the three dilated convolutions in the Cafe-GCN are set to 1, 2, and 3, respectively. For the training process, we use the Adam optimizer to optimize the model with a learning rate of 0.02 in 200 training epochs, and the risk-aversion coefficient η is set to 1. As for the competing baselines, we follow the implementations of the original papers, and if necessary, adapt the setup to the datasets for the portfolio selection task.

D Details for Evaluation Metrics

- (1) **Annualized Return (AR):** annualized mean of the daily portfolio returns, i.e. $AR = \text{mean}(\mathbf{R}) \times 252$.
- (2) **Sharpe Ratio (SR):** risk-adjusted return. The SR is calculated as the annualized return divided by the annualized standard deviation of the daily portfolio returns, i.e., $SR = \frac{\text{mean}(\mathbf{R})}{\text{std}(\mathbf{R})} * \sqrt{252}$.
- (3) **Maximum DrawDown (MDD):** largest loss from peak to trough in the test period, i.e., $MDD = \max_{t_e > t_s} (\sum_{t=1}^{t_s} r_t - \sum_{t=1}^{t_e} r_t)$.
- (4) **Calmar Ratio (CR):** measurement of how much the MDD is covered by the portfolio strategy’s annualized return over the test period, i.e., $CR = \frac{AR}{MDD}$.

*Brunel University of London and King’s College London, United Kingdom. Email: ruirui.liu@brunel/kcl.ac.uk.


†City University of Hong Kong and Bayescien Technologies, China. Email: huichou.huang@exeter.oxon.org.

‡London School of Economics and Political Science, United Kingdom. Email: j.ruf@lse.ac.uk.


§South China University of Technology, China. Email: qyw@scut.edu.cn. Corresponding author.

E Visualization of Accumulative Wealth

Here we visualize the accumulative wealth of all compared portfolio strategies in two datasets in Figure 4. To better compare the proposed Cafe-GCN with the competing baselines, the returns to the baseline portfolio strategies are normalized to the same risk level (in terms of portfolio volatility) as Cafe-GCN. In the NASDAQ100 dataset, the Cafe-GCN model significantly outperforms other competing models. In the CSI100 dataset, most of the portfolio strategies have suffered dramatic losses in recent years while the portfolio performance of Cafe-GCN exhibits strong resilience in the crash and high-volatility market states.



(a) CSI100



(b) NASDAQ100

Figure 4: The accumulative wealth of the Cafe-GCN versus the competing baselines.