

STORM: A Spatio-Temporal Factor Model Based on Dual Vector Quantized Variational Autoencoders for Financial Trading

Yilei Zhao*
Nanyang Technological University
Singapore
yilei.zhao@ntu.edu.sg

Wentao Zhang*
Nanyang Technological University
Singapore
wt.zhang@ntu.edu.sg

Tingran Yang
Zhejiang University
Hangzhou, Zhejiang, China
3220102076@zju.edu.cn

Yong Jiang
Alibaba Group
Hangzhou, Zhejiang, China
jiangyong.ml@gmail.com

Fei Huang
Alibaba Group
Hangzhou, Zhejiang, China
f.huang@alibaba-inc.com

Wei Yang Bryan Lim[†]
Nanyang Technological University
Singapore
Bryan.limwy@ntu.edu.sg

Abstract

In financial trading, factor models are widely used to price assets and capture excess returns from mispricing. Recently, we have witnessed the rise of variational autoencoder-based latent factor models, which learn latent factors self-adaptively. While these models focus on modeling overall market conditions, they often fail to effectively capture the temporal patterns of individual stocks. Additionally, representing multiple factors as single values simplifies the model but limits its ability to capture complex relationships and dependencies. As a result, the learned factors are of low quality and lack diversity, reducing their effectiveness and robustness across different trading periods. To address these issues, we propose a Spatio-Temporal factor Model based on dual vector quantized variational autoencoders, named STORM, which extracts features of stocks from temporal and spatial perspectives, then fuses and aligns these features at the fine-grained and semantic level, and represents the factors as multi-dimensional embeddings. The discrete codebooks cluster similar factor embeddings, ensuring orthogonality and diversity, which helps distinguish between different factors and enables factor selection in financial trading. To show the performance of the proposed factor model, we apply it to two downstream experiments: portfolio management on two stock datasets and individual trading tasks on six specific stocks. The extensive experiments demonstrate STORM's flexibility in adapting to downstream tasks and superior performance over baseline models. **Code is available in PyTorch¹.**

CCS Concepts

• Information systems → Spatial-temporal systems; • Applied computing → Electronic commerce.

*These authors contributed equally to this work.

[†]Corresponding author.

¹<https://github.com/DVampire/Storm>



This work is licensed under a Creative Commons Attribution International 4.0 License.

WSDM '26, Boise, ID, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2292-9/2026/02

<https://doi.org/10.1145/3773966.3777972>

Keywords

Factor Model, VQVAE, Portfolio Management, Algorithmic Trading

ACM Reference Format:

Yilei Zhao, Wentao Zhang, Tingran Yang, Yong Jiang, Fei Huang, and Wei Yang Bryan Lim. 2026. STORM: A Spatio-Temporal Factor Model Based on Dual Vector Quantized Variational Autoencoders for Financial Trading. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3773966.3777972>

1 Introduction

In financial trading, factor models are fundamental tools for asset pricing and are widely used to predict asset returns. These models enhance pricing accuracy and risk management by identifying a set of key factors that explain excess returns. The well-known Fama-French three-factor model [7] extends the Capital Asset Pricing Model [26] with additional factors to capture market risk premiums. As financial theories evolve, researchers continue to discover new factors, broadening our understanding of market dynamics.

Recently, we have witnessed the rise of latent factor models [6, 11, 36, 37], connecting the factor model with the generative model, the variational autoencoder (VAE). These VAE-based models describe high-dimensional data (prices) to low-dimensional representations (factors), and learn factors self-adaptively. Although latent factor models have demonstrated substantial success in financial trading tasks, they still face several significant issues:

- **CH1: Limited Reflection of Market Complexity.** Latent factor models represent factors as single values and are inherently constrained by their insufficient capacity to capture the intricate complexity and nonlinearity of financial data, rendering them vulnerable to noise and non-stationarity, which compromises their predictive accuracy and stability.
- **CH2: Factor Inefficiency.** VAE-learned factors suffer from three inefficiencies: i) focusing mainly on cross-sectional factors while neglecting temporal information, ii) allowing noise in continuous latent spaces to overshadow meaningful signals, and iii) lacking independence among factors, which leads to multicollinearity and weak adaptability to varying market conditions.
- **CH3: Lack of Factor Selection.** Existing latent factor models primarily focus on generating factors without adequately differentiating between them. Furthermore, they neglect the crucial

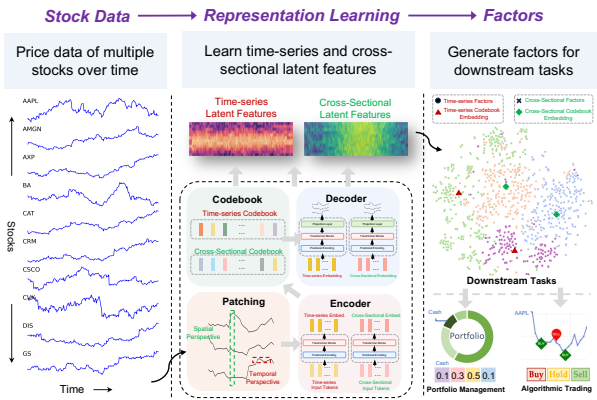


Figure 1: Overall architecture of STORM

process of factor selection, which is essential for identifying impactful factors, thereby limiting the model’s overall effectiveness and precision.

In order to address the challenges, we propose a Spatio-Temporal factor Model based on dual vector quantized variational autoencoders (VQ-VAE), named STORM, with the architecture shown in Figure 1. Unlike traditional scalar-valued financial factors with clear economic interpretations, STORM learns high-dimensional latent vectors that, while less interpretable capture complexity and nonlinearity inherent in financial market (CH1). Additionally, we develop a dual VQ-VAE architecture to capture cross-sectional and time-series features, considering both spatial and temporal perspectives². By integrating these features at both fine-grained and semantic levels, the model constructs more effective factors. Through diversity and orthogonality loss constraints, we ensure representation independence among learned representations (CH2). Furthermore, codebook embeddings act as cluster centers, serving as class tokens to categorize factor embeddings. This strategy provides clarity and transparency to the differentiation and selection process of factors (CH3). Specifically, our contributions are four-fold:

- We design a dual VQ-VAE architecture that systematically constructs cross-sectional and time-series factors by capturing both spatial and temporal dependencies.
- We leverage vector quantization techniques to improve factor embedding representation by reducing noise, enhancing diversity, and ensuring orthogonality. This approach significantly strengthens the factors’ predictive power and a more transparent factor selection process.
- We develop STORM as a versatile all-rounder capable of supporting three tasks simultaneously, whereas other methods tend to be specialized in specific tasks.
- Experiments on seven datasets from two U.S. stock markets, covering the stock future return prediction task (DJ30 and SP500), portfolio management (DJ30 and SP500), and algorithmic trading (five stocks), show that STORM outperforms all baselines across two prediction metrics and six standard financial metrics.

²We treat cross-sectional as spatial and time-series as temporal. These terms will be used interchangeably unless specified.

2 Related Work

2.1 Factor Model

Factor model has attracted extensive research focus from the finance community. Financial experts have carefully selected a wide range of factors from macroeconomic [26], market changes [7], individual stock fundamentals [9], and other perspectives to price assets in an attempt to capture excess returns from mispricing. Recently, latent factor models have emerged that improve investment returns by studying latent factors self-adaptively. Specifically, [11] learns latent factors depend on asset characteristics, [6] extracts cross-sectional factors by predicting future returns, and [37] learns market-augmented latent factors. However, these VAE-based latent factor models may suffer from a lack of robustness in representing low signal-to-noise ratio markets, leading to reduced reconstruction and generation capabilities. Furthermore, their effectiveness is constrained by an insufficient focus on the temporal perspective.

2.2 Financial Trading

Financial trading has attracted significant attention from both the finance and AI communities. Quantitative traders leverage mathematical models and algorithms to automatically identify trading opportunities [29]. Within the FinTech domain, key research areas encompass portfolio management (PM), algorithmic trading (AT), order execution [8], market making [28], etc. In this work, we focus on PM and AT, as they are fundamental to developing intelligent financial decision-making systems.

PM focuses on optimizing wealth allocation across assets [15]. Its development has evolved from simple rule-based methods [20] to prediction-based approaches using machine learning [13] and deep learning [21, 46] for predicting returns or price movements. Despite recent integration of reinforcement learning (RL) in PM [45], the latent factor models still predominantly rely on deep learning for market modeling and precise predictions.

AT, on the other hand, involves executing trades with algorithm-generated signals. Like PM, AT has seen a transition from rule-based and supervised learning techniques [39] to RL-driven approaches. In particular, RL-based methods [30] have gained prominence due to their ability to navigate complex sequential decision-making problems, further enhancing trade execution efficiency and adaptability in dynamic market environments.

2.3 Vector Quantized Variational Autoencoder

Vector Quantized Variational Autoencoder (VQ-VAE) has emerged as a fundamental framework in deep learning research, establishing methodological significance across core domains: representation learning, generative modeling, and time series analysis. Initially introduced by Van Den Oord et al.[32], VQ-VAE has showcased its capabilities in processing high-dimensional data across diverse modalities, including images [23], audio [43], and video [38], with a strong emphasis on learning discrete latent representations. Particularly noteworthy is the model’s demonstrated proficiency in capturing temporal dependencies, where recent studies like [31] have validated its superior performance in modeling sequential patterns compared to conventional continuous latent space approaches,

thereby advancing the development of robust generative architectures specifically optimized for temporal dynamics analysis.

3 Preliminaries

3.1 Factor Model

The factor model (e.g., Arbitrage Pricing Theory [24]) is defined to study the relationship between expected asset returns and factor exposures, focusing on the cross-sectional differences in expected asset returns.

$$\mathbb{E}[R_i^e] = \alpha_i + \beta_i \lambda. \quad (1)$$

Here, $\mathbb{E}[\cdot]$ denotes the expectation operator, and $\mathbb{E}[R_i^e]$ represents the expected excess return of asset i . β_i is factor exposure (or factor loading) of asset i , and λ is factor expected return (or factor risk premium). Lastly, α_i captures pricing error, reflecting any deviations from the model's predicted returns.

According to [6], the factor model is formulated from a spatial perspective, primarily examining cross-sectional variations in expected returns across different assets, rather than capturing the temporal evolution of individual asset returns. Therefore, we expand it from the temporal perspective and introduce the general functional form of the latent factor model:

$$y_{i,t} = \alpha_{i,t} + \sum_{k=1}^K \beta_{i,t}^k z_{i,t}^k + \epsilon_{i,t}, \quad (2)$$

where $y_{i,t}$ denotes the future return of stock i at timestep t , $z_{i,t}$ denotes K factor returns, and $\epsilon_{i,t}$ is the noise. From the spatial, Equation 2 can be expressed as $y_i = \alpha_i + \sum_{k=1}^K \beta_i^k z_i^k + \epsilon_i$, which we regard as the cross-sectional factor model. From the temporal perspective, Equation 2 can be expressed as $y_t = \alpha_t + \sum_{k=1}^K \beta_t^k z_t^k + \epsilon_t$, which we regard as the time-series factor model. The cross-sectional factor model focuses on the differences in returns of different assets at the same point in time, while the time-series factor model focuses on the changes in returns of a single asset at different time points.

3.2 Problem Formulation

To assess the effectiveness and generating ability of factors, we predict and evaluate the stock future returns. To show the adaptability of the latent factor model with multiple downstream tasks, we employ two financial trading tasks, i.e., *portfolio management* [35, 40, 45] and *algorithmic trading* [17, 30], to showcase the model's investment capabilities.

Observed Data. We utilize the stock's historical price data \mathbf{p} (i.e., open, high, low, and close) and technical indicators \mathbf{d} as observed variables $\mathbf{x} := [\mathbf{p}, \mathbf{d}] \in \mathbb{R}^{N \times W \times D}$, within a window size of W . The historical price data of stock i at time t is denoted as $\mathbf{p}_{i,t} := [p_{i,t}^o, p_{i,t}^h, p_{i,t}^l, p_{i,t}^c] \in \mathbb{R}^{D_1}$. The technical indicators which are calculated from price and volume values, are denoted as $\mathbf{d}_{i,t} \in \mathbb{R}^{D_2}$. Each stock is represented by $D = D_1 + D_2$ features per trading day, which we then use to predict the stock's future returns, defined as $y_{i,t+1} := \frac{p_{i,t+1} - p_{i,t}}{p_{i,t}}$.

Downstream tasks. These primarily encompass the two most widely studied areas in quantitative finance – *Portfolio management* and *Algorithmic trading*:

- *Portfolio management* task aims to construct an optimal portfolio based on the predicted returns to test the profitability of the

factor model. At each trading day t , the factor model outputs return predictions $\hat{y}_{i,t+1}$ for each stock i , which are then used by a portfolio allocation strategy, e.g., *TopK-Drop*, to generate the portfolio weight vector $\mathbf{w}_t = [w_t^0, w_t^1, \dots, w_t^{N-1}] \in \mathbb{R}^N$, where w_t^i represents the proportion of capital invested in stock i . To ensure full capital deployment, the allocation must satisfy the budget constraint $\sum_{i=0}^{N-1} w_t^i = 1$.

- *Algorithmic trading* task aims to execute buy, hold, and sell actions based on predicted asset states to balance returns and risks. We formulate it as a Markov Decision Process (MDP) under reinforcement learning scenario following [3, 44]. The MDP is constructed by a 5-tuple $(S, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where S and \mathcal{A} are sets of states and actions respectively, $\mathcal{T} : S \times \mathcal{A} \times S \rightarrow [0, 1]$ is the state transition function, $\mathcal{R} : S \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function where \mathbb{R} is a continuous set of possible rewards, and $\gamma \in [0, 1)$ is the discount factor. The latent factor embeddings \mathbf{Z} generated by the factor model are integrated into the observation set that constitutes the state space S , enabling the agent to incorporate predictive market signals. The goal is to learn a policy $\pi : S \rightarrow \mathcal{A}$ that maximizes the expected discounted cumulative reward $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r_t]$. At each trading day t for specific stock i , the agent takes action $a_{i,t} := \{a_{i,t}^b, a_{i,t}^h, a_{i,t}^s\} \in \mathcal{A}$ which means buy, hold and sell, according to the current environment $s_t \in S$.

3.3 Vector Quantized Variational Autoencoder

VQ-VAE[32] is different from VAE [14] as it quantizes the observed data into a discrete token sequence. The VQ-VAE contains three parts, encoder, decoder, and discrete codebook, denoted by ϕ_{enc} with parameter θ_{enc} , ϕ_{dec} with parameter θ_{dec} and $C = \{(k, e_k \in \mathbb{R}^H)\}_{k=1}^K$ with parameter $\{e_k \in \mathbb{R}^H\}_{k=1}^K$ respectively. Among them, K is the size of the codebook, H is the dimension of the vector, and e_k is the k -th vector of the latent embedding space. Given an input x , the encoder ϕ_{enc} firstly encodes it into a set of continuous feature vectors $Z = \phi_{\text{enc}}(x)$ and quantization operation $q(\cdot)$ passes each feature vector $z \in Z$ through the discretization bottleneck following by selecting the nearest neighbor embedding in the codebook C as its discrete code sequence Q_z .

$$Q_z = q(z, C) = c, \text{ where } c = \arg \min_{k \in [0, K-1]} \|z - e_k\|_2^2. \quad (3)$$

Then, the quantized feature for z denoted as Z_z^q , is obtained through $Z_z^q = e_{Q_z}$. The quantized vectors Z^q is fed into the decoder θ_{dec} to reconstruct the original data $x' = \theta_{\text{dec}}(Z^q)$.

4 Method

In this section, we represent our proposed latent factor model STORM, with the overall structure of the VQ-VAE depicted in Figure 2. In general, STORM models data into time-series and cross-sectional features, deriving factors by integrating both temporal and spatial perspectives. Firstly, we introduce the dual VQ-VAE architecture and explain the extraction of features in two modules. Then the factor module will fuse and align features and generate factor embeddings. Finally, we demonstrate the effectiveness of STORM through comprehensive evaluations on two critical downstream tasks: portfolio management and algorithmic trading.

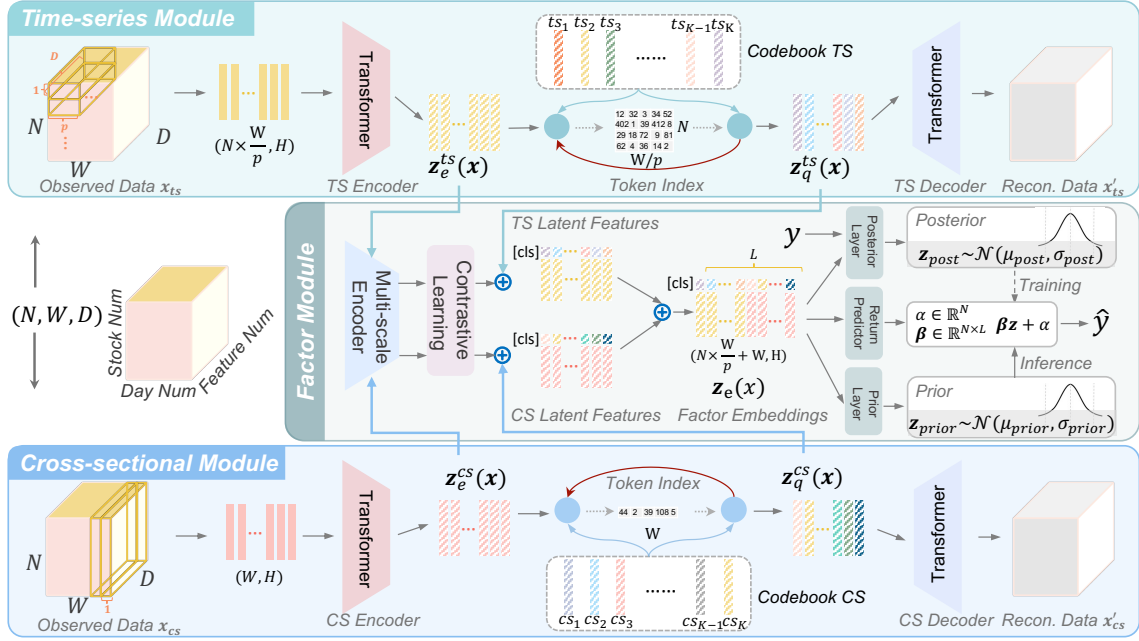


Figure 2: The architecture of our proposed STORM model

4.1 Dual VQ-VAE Structure

The dual VQ-VAE architecture processes stock data from both spatial and temporal perspectives by *patching* the data within each module. The *encoder* then extracts time-series (TS) and cross-sectional (CS) features, which are subsequently clustered using a learnable *codebook* for quantized representation. Finally, the *decoder* reconstructs the data, effectively capturing complex patterns and enhancing model’s forecasting performance.

4.1.1 Patching. In stock data analysis, it is essential to extract features from both temporal and spatial dimensions. Time-series features capture the dynamics of a stock over time, while cross-sectional features reveal the correlations among different stocks. To achieve this, we partition data into patches [19] within the TS and CS modules. In the TS module, the observed data \mathbf{x} is divided along the stock number dimension, with each patch containing data for one stock over p days (as illustrated in Figure 2). Similarly, in the CS module, the observed data \mathbf{x} is divided along the time axis, with each patch containing features of all stocks for a single trading day. Combining information from multi-scale features allows the model to capture both temporal and spatial dependencies, enhancing the depth and accuracy of market insights.

4.1.2 TS and CS Encoders. Considering the Transformer’s [33] remarkable ability to capture both local and global dependencies via self-attention and patching strategies, it is particularly well-suited for long-term time-series forecasting [19]. Therefore, we employ stacked transformer blocks as the encoders in both the TS and CS modules to capture complex time-series and cross-sectional patterns. This architecture enhances the accuracy and robustness of feature extraction, ultimately improving forecasting performance.

While transformers are widely applicable across domains such as natural language processing (NLP) and computer vision (CV),

STORM’s flexible encoder-decoder framework allows for the integration of alternative architectures, such as long short-term memory (LSTM) [34] and temporal convolutional networks (TCN) [41].

An encoder-decoder configuration comprising 4 and 2 blocks, respectively, each with 4 and 8 attention heads, remains relatively lightweight, avoiding excessive computational cost and a substantial increase in model complexity.

4.1.3 Codebook Construction and Optimization. The construction of the codebook involves a learnable embedding space designed to quantize the encoded features into discrete representations [32]. During training, the continuous latent features $\mathbf{z}_e^{ts}(\mathbf{x})$ and $\mathbf{z}_e^{cs}(\mathbf{x})$ are quantized into discrete space by mapping each vector to its nearest codebook entry, $\mathbf{z}_q^{ts}(\mathbf{x})$ and $\mathbf{z}_q^{cs}(\mathbf{x})$, based on minimizing the Euclidean distance. This process translates the encoded features into discrete tokens, enabling the model to leverage a finite set of vectors to capture complex patterns.

The VQ-VAE offers significant advantages in the latent factor model, primarily through its use of discrete codebook embeddings: i) *Discretization Benefits:* The discretization process inherent in VQ-VAE helps in clustering distinct and meaningful factors, providing orthogonality and diversity which improves upon the factors derived from traditional methods like FactorVAE. ii) *Explicit Factor Selection:* The discrete token indices facilitate an explicit factor selection process, identifying the most relevant factors that influence stock returns and thus enhancing prediction accuracy. iii) *Noise Reduction:* By filtering out irrelevant or redundant information during factor selection, VQ-VAE mitigates noise, contributing to improved model robustness and more reliable financial insights.

To encourage equal usage of the codebook vectors [1], we incorporate a diversity loss to ensure balanced representations across the codebook. It is achieved by maximizing the entropy of the averaged

softmax distribution over vectors for each codebook \bar{p}_g :

$$\mathcal{L}_{\text{div}} = \frac{1}{GK} \sum_{g=1}^G \sum_{k=1}^K \bar{p}_{g,k} \log \bar{p}_{g,k}, \quad (4)$$

where G represents the number of codebooks, set 2 in STORM. The diversity constraint helps enhance representational capacity, leading to better coverage of the input space.

However, the lack of independence among factors can introduce multicollinearity, making the model prone to overfitting and reducing its robustness across different market conditions. To address this issue, we impose an orthogonality constraint on the codebook, as recent studies [27] have shown that enforcing orthogonality allows discretized codes to maintain translation equivariance. The orthogonality loss is:

$$\mathcal{L}_{\text{ortho}} = \frac{1}{K^2} \left\| \ell_2(\mathbf{e})^\top \ell_2(\mathbf{e}) - I_K \right\|_F^2, \quad (5)$$

where $I_k \in \mathbb{R}^{K \times K}$ is the identity matrix, \mathbf{e} is a surrogate for two codebook embeddings \mathbf{ts} and \mathbf{cs} , $\ell_2(\mathbf{e})$ denotes L_2 -normalized embeddings, and $\|\cdot\|_F$ is the Frobenius norm. The orthogonal constraint ensures that factors are independent of each other, allowing the effect of each factor on returns to be individually explained, which is crucial for asset pricing analysis.

4.2 TS and CS Decoders

The decoder is responsible for reconstructing the original data \mathbf{x} from the quantized latent vectors $\mathbf{z}_q^{ts}(\mathbf{x})$ and $\mathbf{z}_q^{cs}(\mathbf{x})$. We utilize Transformer as decoders either, to generate constructed data \mathbf{x}'_{ts} and \mathbf{x}'_{cs} , which aim to closely approximate the original input data.

Finally, the encoders, decoders, and codebooks are jointly optimized by minimizing the following loss objectives:

$$\begin{aligned} \mathcal{L}_1 = & \lambda_{\text{ortho}} \mathcal{L}_{\text{ortho}} + \lambda_{\text{div}} \mathcal{L}_{\text{div}} + \|\mathbf{x} - \mathbf{x}'_{ts}\|_2^2 + \|\mathbf{x} - \mathbf{x}'_{cs}\|_2^2 \\ & + \left\| \text{sg}[\mathbf{z}_e^{ts}(\mathbf{x})] - \mathbf{z}_q^{ts}(\mathbf{x}) \right\|_2^2 + \left\| \text{sg}[\mathbf{z}_e^{cs}(\mathbf{x})] - \mathbf{z}_q^{cs}(\mathbf{x}) \right\|_2^2 \\ & + \left\| \text{sg}[\mathbf{z}_e^{ts}(\mathbf{x})] - \mathbf{z}_q^{ts}(\mathbf{x}) \right\|_2^2 + \left\| \text{sg}[\mathbf{z}_e^{cs}(\mathbf{x})] - \mathbf{z}_q^{cs}(\mathbf{x}) \right\|_2^2, \end{aligned} \quad (6)$$

where $\text{sg}[\cdot]$ is a stop-gradient operator.

4.3 Factor Module

The factor module aims to integrate the TS and CS latent features through *feature fusion and alignment*. Afterward, it utilizes *prior-posterior learning* to generate factor embeddings from latent features to predict stock future returns.

4.3.1 Feature Fusion and Alignment. As mentioned above we employ different patching methods to ensure that both time-series and cross-sectional features of the data were accurately captured. However, these methods result in two parallel processes that do not align at the feature level. To address the limitation, we employ a *cross-attention mechanism* to enhance the interaction and fusion of features at a fine-grained level, improving the model's dynamic understanding of the input data. Additionally, *contrastive learning* is used at the semantic level to enhance feature representation by emphasizing semantic similarities and differences.

Cross-Attention Mechanism. In order to fuse the features in the TS and CS modules at a fine-grained level, we leverage a multi-scale

encoder to combine information between two temporal and spatial patterns and produce stronger fusion features for factor generations, motivated by CrossViT [2]. The architecture incorporates cross-attention layers that explicitly learn correlations between two patterns, enabling the encoder to emphasize mutually informative regions while filtering out noise. This design balances accuracy and computational efficiency, ensuring the fused features yield richer and more discriminative signals for factor generation and that the resulting factors are more representative of market dynamics.

Contrastive Learning. To further enhance semantic consistency between TS and CS features, we integrate contrastive learning, drawing inspiration from its success in aligning multimodal representations [16, 22]. A contrastive loss is introduced within the factor module to encourage features from TS and CS modules that represent similar market conditions to cluster closely in the embedding space, while pushing apart features corresponding to different conditions. This dual-stream alignment improves the coherence of the factor space and enhances the model's ability to generalize across diverse market scenarios.

4.3.2 Prior-Posterior Learning. Through multi-scale encoder and contrastive learning layer, the CS and TS latent features are fused and aligned. In order to retain the codebook's categorical properties for factors, we incorporate the codebook embeddings $\mathbf{z}_q^{ts}(\mathbf{x})$ and $\mathbf{z}_q^{cs}(\mathbf{x})$ as extra [CLS]-tokens [5], which are appended to the latent features $\mathbf{z}_e^{ts}(\mathbf{x})$ and $\mathbf{z}_e^{cs}(\mathbf{x})$, respectively. Then, two latent features concat together to get the latent factor embeddings $\mathbf{z}_e(\mathbf{x})$.

Given the inherent volatility and complexity of the stock market, bridging the gap between noisy market data and an effective factor model remains a significant challenge [6]. Therefore, we introduce the prior-posterior structure to predict future returns and optimize latent factor embeddings. In the training stage, the posterior layer estimates posterior distribution of factor expected returns \mathbf{z}_{post} from the true future stock returns \mathbf{y} and the latent factor embeddings $\mathbf{z}_e(\mathbf{x})$:

$$[\mu_{post}, \sigma_{post}] = \phi_{FE}(\mathbf{y}, \mathbf{z}_e(\mathbf{x})), \quad (7)$$

and \mathbf{z}_{post} follows the independent Gaussian distribution. In the inference stage, the prior layer only use latent factor embeddings without any future information leakage to calculate prior distribution of factor expected returns \mathbf{z}_{prior} :

$$[\mu_{prior}, \sigma_{prior}] = \phi_{FP}(\mathbf{z}_e(\mathbf{x})). \quad (8)$$

The return predictor computes future returns based on the factor expected return \mathbf{z} :

$$\hat{\mathbf{y}} = \boldsymbol{\alpha} + \sum_{k=1}^K \boldsymbol{\beta}^k \mathbf{z}^k + \boldsymbol{\epsilon}. \quad (9)$$

The predicted returns $\hat{\mathbf{y}}$ can be directly integrated into portfolio optimization, algorithmic trading strategies, enabling dynamic rebalancing and improved risk management.

4.4 Downstream Tasks

Existing research on latent factor models [6, 37] primarily focuses on predicting future returns by analyzing cross-sectional characteristics of market stocks, which means the profitability of the factor model in financial investment is typically evaluated through

PM tasks. In contrast, STORM captures the cross-sectional features among stocks and the time-series features of individual stocks, enabling it to excel in both PM and single-asset AT tasks.

4.4.1 Portfolio Management. To evaluate the performance of the STORM and compare it with baseline methods, we follow the paradigm proposed by FactorVAE [6], focusing on return prediction and the PM task. Specifically, we utilize the factor decoder network to generate stock future returns \hat{y} , and then apply the *TopK-Drop* strategy³, which constructs a daily portfolio by selecting the top k stocks based on predicted returns, to backtest the factor model. Considering transaction costs in real markets, there is a turnover constraint that limits the number of changes in the portfolio composition between consecutive trading days. Specifically, when moving from \mathbf{w}_t to \mathbf{w}_{t+1} , at most d stocks can be replaced, where a replacement is defined as a stock whose weight changes from 0 (not held) to a positive value (newly bought) or from a positive value to 0 (fully sold). This constraint ensures that at least $k - d$ stocks remain in the portfolio, reducing excessive turnover.

4.4.2 Algorithmic trading. In STORM, $\mathbf{z}_e(\mathbf{x})$ is the stock latent factor embeddings encoded by the dual VQ-VAE architecture, including both time-series and cross-sectional factor embeddings in the data. The latent factor embeddings \mathbf{Z} are integrated into the observation set $\mathcal{O} = \{\mathbf{Z}, \mathcal{R}\}$, where \mathcal{R} is the reward function used to guide the agent’s learning and decision-making in the environment. By combining latent factor embeddings with reward, the agent can more accurately understand the market environment and adjust its strategy based on observations. In the AT task, since it involves trading a single stock, the features of other stocks are excluded. We use the Proximal Policy Optimization (PPO) algorithm [25] to optimize the policy. In each iteration, the agent interacts with the environment to collect trajectories, computes advantage estimates, and updates the policy parameters via PPO’s clipped surrogate objective. This optimization process ensures the stability and efficiency of policy updates, helping the agent gradually improve the decision-making ability and achieve better trading strategies.

5 Experiment

In this section, we evaluate the proposed STORM on real stock markets and conduct extensive experiments to address the following research questions.

RQ1 How does STORM perform on downstream tasks?

RQ2 How to evaluate the effectiveness of STORM’s learned factors?

RQ3 How do the key components contribute to the performance of STORM?

5.1 Experiment Settings

5.1.1 Datasets. We conduct experiments on two U.S. stock markets, SP500 and DJ30, using stock daily data that includes technical features based on *Alpha158* [39]. Both datasets span 16 years, from 2008-04-01 to 2024-03-31, encompassing global conditions, e.g., the 2007-2008 financial crisis and COVID-19. Datasets are chronologically divided into non-overlapping training (from 2008-04-01 to 2021-03-31) and test (from 2021-04-01 to 2024-03-31) sets.

³<https://qlib.readthedocs.io/en/latest/component/strategy.html>

5.1.2 Metrics. We compare STORM and baselines in terms of 6 financial metrics across the PM and AT tasks. The financial metrics include 2 profit criteria, annualized percentage yield (APY) and Cumulative Wealth (CW), 2 risk-adjusted profit criteria, Calmar ratio (CR) and annualized Sharpe ratio (ASR), and 2 risk criteria, maximum drawdown (MDD) and annualized volatility (AVO). The calculation formulas and meanings of these metrics are as follows:

- **CW** is the total returns yielded from a portfolio strategy: $CW_T = \prod_{i=1}^T (1 + r_i)$, where r_i is the net return.
- **APY** measures the average wealth increment that one portfolio strategy could achieve compounded in a year, which is defined as $APY_T = \sqrt[T]{CW_T} - 1$, where y is the number of years corresponding to T trading rounds.
- **MDD** measures the largest loss from a historical peak in the cumulative wealth to show the worst case, which is defined as: $MDD = \max_{i=0}^T \frac{P_i - R_i}{P_i}$, where $R_i = \prod_{j=1}^i \frac{V_j}{V_{j-1}}$ and $P_i = \max_{j=1}^i R_j$.
- **AVO** is the annualized standard deviation of daily returns and multiplied by \sqrt{AT} , where AT is the average trading rounds of annual trading days and $AT = 252$ for all the datasets.
- **ASR** is an annualized volatility risk-adjusted return defined as $ASR = \frac{APY - R_f}{AVO}$, where R_f is the risk-free return.
- **CR** measures the drawdown risk-adjusted return of a portfolio calculated as $CR = \frac{APY}{MDD}$.

Typically, to evaluate the effectiveness of the learned factors in the PM task, we adopt the Rank Information Coefficient (RankIC) and the Information Ratio of RankIC (RankICIR):

- **RankIC** is a ranking metric in finance, which measures the correlation between the predicted rankings and the actual returns. It is defined as:

$$\text{RankIC}_s = \frac{1}{N} \frac{(r_{\hat{y}_s} - \text{mean}(r_{\hat{y}_s}))^T (r_{y_s} - \text{mean}(r_{y_s}))}{\text{std}(r_{\hat{y}_s}) \cdot \text{std}(r_{y_s})}$$

$$\text{RankIC} = \frac{1}{T_{\text{test}}} \sum_{s=1}^{T_{\text{test}}} \text{RankIC}_s,$$

where T_{test} denotes the number of trading days in the test range, r_{y_s} and $r_{\hat{y}_s}$ represent the true and predicted ranks of stocks on the trading day s , respectively.

- **RankICIR** is the information ratio of RankIC, which measures the stability of prediction,

$$\text{RankICIR} = \frac{\text{mean}(\text{RankIC}_s)}{\text{std}(\text{RankIC}_s)}.$$

5.1.3 Baselines. As shown in Table 1, we compare STORM with nine methods across one prediction task and two downstream tasks. The baseline methods fall into three categories: ML & DL-based models, Factor models, and RL-based models. We evaluate STORM against ML&DL models and factor models in prediction and PM tasks, and against ML&DL models and RL methods in the AT task. Moreover, STORM is a versatile all-rounder that supports three tasks, unlike other methods specialized in specific tasks.

- **Market**
 - Buy-and-Hold (B&H) involves holding assets for an extended period, regardless of short-term market fluctuations, assuming that long-term returns will be more favorable.
- **ML&DL-based**

Table 1: Comparison of baselines across different tasks

Task	Market	ML & DL-based			RL-based			Factor model			Ours
	B&H	LGBM [39]	LSTM [42]	Transformer [39]	SAC [12]	PPO [25]	DQN [18]	CAFactor [11]	FactorVAE [6]	HireVAE [37]	STORM
Prediction	✓	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓
Portfolio Management	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Algorithmic Trading	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓

* ✓ indicates that the model is applicable, while ✗ signifies that it is NOT applicable to the corresponding task.

- LGBM [39] uses a series of tree models to predict price fluctuations and provide buy and sell signals.
- LSTM [42] utilizes long short-term memory to improve the accuracy of price predictions.
- Transformer [39] models leverage self-attention mechanisms to enhance the precision of price forecasts.

• RL-based

Table 2: Portfolio management task results of all models across six metrics (mean ± range, computed across 10 runs).

SP500 Dataset						
Strategies	Profit		Risk-Adj. Profit		Risk	
	APY↑	CW↑	CR↑	ASR↑	MDD↓	AVO↓
Market Index	0.058	1.184	0.228	0.142	0.254	0.410
LightGBM	0.059 ± 0.132	1.201 ± 0.456	0.304 ± 0.753	0.332 ± 0.755	0.238 ± 0.108	0.176 ± 0.008
LSTM	0.069 ± 0.020	1.221 ± 0.068	0.278 ± 0.065	0.371 ± 0.114	0.248 ± 0.042	0.186 ± 0.013
Transformer	0.076 ± 0.028	1.244 ± 0.098	0.389 ± 0.165	0.433 ± 0.161	0.198 ¹ ± 0.057	0.174 ± 0.004
CAFactor	0.075 ± 0.028	1.241 ± 0.100	0.342 ± 0.229	0.428 ± 0.165	0.223 ± 0.043	0.174 ± 0.006
FactorVAE	0.079 ± 0.025	1.256 ± 0.085	0.404 ± 0.177	0.460 ± 0.145	0.200 ± 0.040	0.173 ± 0.007
HireVAE	0.077 ± 0.029	1.249 ± 0.104	0.361 ± 0.180	0.448 ± 0.189	0.216 ± 0.048	0.172 ± 0.007
STORM	0.188² ± 0.055	1.683 ± 0.226	1.189 ± 0.661	1.052 ± 0.329	0.166 ± 0.050	0.171 ± 0.020
STORM-w/o-TS	0.090 ± 0.074	1.300 ± 0.277	0.503 ± 0.390	0.570 ± 0.339	0.181 ± 0.043	0.175 ± 0.025
STORM-w/o-CS	0.089 ± 0.045	1.294 ± 0.163	0.623 ± 0.385	0.592 ± 0.252	0.146 ± 0.073	0.167 ± 0.020
Improvement(%) ³	137.97	34.00	194.31	128.70	26.26	2.91

DJ30 Dataset						
Strategies	Profit		Risk-Adj. Profit		Risk	
	APY↑	CW↑	CR↑	ASR↑	MDD↓	AVO↓
Market Index	0.063	1.201	0.147	0.429	0.219	0.288
LightGBM	0.069 ± 0.040	1.221 ± 0.140	0.288 ± 0.262	0.430 ± 0.267	0.244 ± 0.047	0.160 ± 0.007
LSTM	0.060 ± 0.007	1.192 ± 0.023	0.243 ± 0.021	0.370 ± 0.043	0.248 ± 0.006	0.163 ± 0.005
Transformer	0.056 ± 0.013	1.179 ± 0.043	0.227 ± 0.067	0.367 ± 0.085	0.250 ± 0.033	0.154 ± 0.002
CAFactor	0.059 ± 0.015	1.186 ± 0.050	0.233 ± 0.058	0.382 ± 0.092	0.252 ± 0.021	0.153 ± 0.001
FactorVAE	0.076 ± 0.038	1.246 ± 0.135	0.352 ± 0.349	0.480 ± 0.235	0.225 ± 0.063	0.159 ± 0.005
HireVAE	0.072 ± 0.037	1.233 ± 0.132	0.298 ± 0.253	0.445 ± 0.256	0.247 ± 0.057	0.163 ± 0.012
STORM	0.148 ± 0.046	1.517 ± 0.188	1.396 ± 0.679	1.052 ± 0.297	0.108 ± 0.026	0.140 ± 0.014
STORM-w/o-TS	0.079 ± 0.057	1.259 ± 0.192	0.621 ± 0.425	0.603 ± 0.379	0.127 ± 0.038	0.142 ± 0.008
STORM-w/o-CS	0.073 ± 0.051	1.236 ± 0.169	0.533 ± 0.400	0.566 ± 0.343	0.138 ± 0.034	0.140 ± 0.017
Improvement(%)	94.74	21.75	296.59	119.17	52.00	8.50

¹ Underline indicates the best-performing baseline method result;

² **Bold** indicates the best performance among our proposed models (including variants);

³ Improvement of STORM over the best-performing baselines.

- SAC [12] is an off-policy actor-critic algorithm using entropy regularization and soft value functions.
- PPO [25] updates trading policies iteratively to balance exploration and exploitation, ensuring stability and efficiency.
- DQN [18] approximates the action-value function and make trading decisions from market data with a deep Q-network.

• Factor-model

- CAFEactor [11] learns latent factors from asset features.
- FactorVAE [6] is a state-of-the-art VAE-based latent factor model that learns optimal latent stock factors.
- HireVAE [37] is a state-of-the-art VAE-based latent factor model capturing market and inter-stock factors.

5.1.4 Implementation Details. All experiments are implemented using PyTorch and conducted on two NVIDIA RTX H100 GPUs with 80GB of memory each. The input data is organized into batches with dimensions (B, W, N, D) , where $B = 16$ is the batch size, N is the number of stocks (28 for DJ30, 408 for SP500), $W = 64$ denotes historical days, and $D = 152$ is the total number of features, including the OHLC data and technical indicators.

The TS module uses patches of size $(4, 1, D)$ (4-day sequences per stock), whereas the CS module uses $(1, N, D)$ (all stocks on a given day). Encoders (4 layers, 4 heads) and decoders (2 layers, 8 heads) are standard Transformer blocks with a shared embedding size of 256. Codebook sizes are searched in 128, 256, 512, 1024, 2048, with 512 performing best.

We adopt AdamW ($\text{lr}=1 \times 10^{-4}$, weight decay=0.05) with a linear warmup of 100 epochs and a decay to zero over 1000 epochs. The clip loss and TS/CS reconstruction losses are weighted by 1×10^{-3} ; commitment loss is 1.0; orthogonality, diversity, price prediction, and KL losses are weighted by 0.1.

For downstream tasks, we use Qlib’s⁴ TopK-Drop portfolio strategy ($k=5, d=3$) and PPO for trading. The pretrained TS/CS factors serve as states; all CS factors are retained while only TS factors relevant to traded stocks are used. The action space is buy, hold, sell. Rewards follow $r = \frac{v_{\text{post}} - v_{\text{pre}}}{v_{\text{pre}}}$ with $v = p \times m + c$, starting from 10^6 initial cash and transaction cost 10^{-4} .

The implementations of the ML & DL methods are based on Qlib [39]. As for other baselines, we adopt the default settings in their public implementations. Each experiment is repeated multiple times with different random seeds, and the mean and variance of the metrics are reported.

⁴<https://github.com/microsoft/qlib>

Table 3: Algorithmic trading task results on all models across six standard metrics (averaged over 5 runs).

Models	AAPL			JPM			IBM			INTC			MSFT		
	APY \uparrow	CW \uparrow	CR \uparrow	APY \uparrow	CW \uparrow	CR \uparrow	APY \uparrow	CW \uparrow	CR \uparrow	APY \uparrow	CW \uparrow	CR \uparrow	APY \uparrow	CW \uparrow	CR \uparrow
Buy&Hold	0.120	1.404	0.383	0.096	1.316	0.236	0.145	1.499	0.727	-0.117	0.690	-0.184	0.214	1.784	0.569
LightGBM	0.135	1.390	0.487	0.116	1.335	0.333	0.227	1.654	1.091	-0.042	0.880	0.038	0.267	2.068	0.637
LSTM	0.053	1.152	0.283	0.079	1.290	0.266	0.134	1.386	0.754	0.060	1.262	0.381	0.178	1.513	0.893
Transformer	± 0.015	± 0.042	± 0.101	± 0.019	± 0.053	± 0.023	± 0.141	± 0.406	± 0.69	± 0.119	± 0.429	± 0.62	± 0.122	± 0.352	± 0.822
DQN	0.083	1.240	0.512	0.133	1.384	0.614	0.131	1.377	0.782	0.079	1.290	0.458	0.138	1.397	0.726
SAC	± 0.099	± 0.284	± 0.864	± 0.07	± 0.201	± 0.509	± 0.072	± 0.207	± 0.541	± 0.174	± 0.68	± 0.824	± 0.118	± 0.338	± 0.643
PPO	0.135	1.374	0.510	0.105	1.305	0.607	0.139	1.400	0.802	0.061	1.185	0.442	0.166	1.475	0.534
STORM	± 0.075	± 0.226	± 0.239	± 0.132	± 0.383	± 0.673	± 0.055	± 0.158	± 0.270	± 0.106	± 0.305	± 0.436	± 0.035	± 0.100	± 0.107
STORM-w/o-TS	0.147	1.509	0.528	0.131	1.383	0.400	0.207	1.598	1.170	0.056	1.165	0.353	0.229	1.656	0.929
STORM-w/o-CS	± 0.021	± 0.060	± 0.052	± 0.038	± 0.111	± 0.198	± 0.057	± 0.163	± 0.335	± 0.210	± 0.6	± 0.994	± 0.114	± 0.331	± 0.585
Improvement(%)	55.782	23.062	45.076	30.827	17.124	2.117	3.965	14.4501	20.408	118.987	26.969	66.594	8.614	4.159	30.893

Models	AAPL			JPM			IBM			INTC			MSFT		
	ASR \uparrow	MDD \downarrow	AVO \downarrow	ASR \uparrow	MDD \downarrow	AVO \downarrow	ASR \uparrow	MDD \downarrow	AVO \downarrow	ASR \uparrow	MDD \downarrow	AVO \downarrow	ASR \uparrow	MDD \downarrow	AVO \downarrow
Buy&Hold	0.447	0.313	0.269	0.405	0.406	0.237	0.679	0.199	0.213	-0.324	0.635	0.359	0.777	0.376	0.275
LightGBM	0.950	0.309	0.017	0.921	0.386	0.015	1.822	0.175	0.011	0.098	0.553	0.023	1.456	0.370	0.017
LSTM	0.623	0.244	0.012	0.801	0.335	0.012	1.268	0.163	0.010	0.287	0.149	0.010	1.345	0.246	0.014
Transformer	± 0.163	± 0.072	± 0.003	± 0.185	± 0.061	± 0.002	± 0.924	± 0.017	± 0.001	± 1.02	± 0.101	± 0.007	± 0.842	± 0.124	± 0.003
DQN	0.972	0.239	0.010	1.332	0.226	0.010	1.336	0.182	0.010	0.604	0.257	0.014	1.206	0.182	0.011
SAC	± 0.148	± 0.087	± 0.006	± 0.607	± 0.100	± 0.003	± 0.688	± 0.100	± 0.003	± 0.789	± 0.194	± 0.006	± 0.734	± 0.134	± 0.006
PPO	0.974	0.288	0.016	1.085	0.212	0.010	1.415	0.161	0.009	0.609	0.243	0.014	1.194	0.309	0.015
STORM	± 0.472	± 0.159	± 0.007	± 1.270	± 0.113	± 0.003	± 0.403	± 0.049	± 0.004	± 0.584	± 0.310	± 0.009	± 0.148	± 0.079	± 0.004
STORM-w/o-TS	1.054	0.292	0.016	1.044	0.350	0.014	1.794	0.157	0.011	0.527	0.359	0.017	1.638	0.229	0.013
STORM-w/o-CS	± 0.223	± 0.022	± 0.002	± 0.276	± 0.085	± 0.001	± 0.458	± 0.068	± 0.005	± 1.061	± 0.191	± 0.004	± 0.736	± 0.215	± 0.005
Improvement(%)	27.704	4.184	-	19.144	8.491	-	15.642	16.561	11.111	95.238	-	-	2.076	-	-

5.2 Performance on Investment Tasks (for RQ1)

We use three types of metrics: profit, risk, and risk-adjusted profit to evaluate STORM’s investment performance on downstream tasks, i.e., portfolio management and algorithmic trading. The evaluation results are presented in Table 2 and Table 3.

For *profitability*, STORM outperforms all the baselines, with average improvements of 116.36% on APY over the best-performing baseline in the PM task and 43.64% over the best baseline methods on five stocks in the AT task. These results underscore the model’s capability to extract meaningful factors, and optimize decision-making for enhanced financial performance. For *risk resistance*, although STORM does not outperform the best baseline methods in terms of MDD and AVO on the INTC and MSFT datasets for the trading task, this can be attributed to the extreme price fluctuations during the COVID-19 event. Both stocks experienced a sharp decline of over 30%, followed by a subsequent rebound. Given STORM’s strong profitability capability, its trading strategy tends to aggressively capture trends, which naturally leads to larger drawdowns and increased volatility as a trade-off for higher potential returns.

Moreover, considering the *return and risk* simultaneously, STORM performs the best on ASR and CR with an average improvement of 123.94% and 245.45% respectively, over the best baseline result in the PM task, and surpasses the best baseline results across three stocks in the AT task. Such an outstanding ability to balance return and risk is particularly valuable in real-world investment, where achieving high returns while effectively managing risk is crucial for long-term financial success.

Table 4: Factor quality evaluation task results on RankIC and RankICIR (mean \pm range, computed across 10 runs).

Models	SP500		DJ30	
	RankIC \uparrow	RankICIR \uparrow	RankIC \uparrow	RankICIR \uparrow
<i>Future Return Prediction in Portfolio Management Task</i>				
LightGBM	0.027	0.274	0.031	0.272
LSTM	0.034	0.333	0.031	0.329
Transformer	0.035	0.340	0.033	0.343
CAFactor	0.037	0.356	0.040	0.380
FactorVAE	0.052	0.543	0.056	0.520
HireVAE	0.057	0.558	0.058	0.563
STORM	0.062	0.673	0.065	0.668
STORM-w/o-TS	0.053	0.513	0.055	0.563
STORM-w/o-CS	0.054	0.522	0.053	0.559
Improvement(%) ¹	8.772	20.609	12.069	18.650

¹ Improvement of STORM over the best-performing baselines.

5.3 Factor Quality (for RQ2)

To demonstrate the effectiveness and quality of the factors, we designed a stock future return prediction task. The goal is not to pursue time-series forecasting accuracy, but to evaluate the predictive quality and robustness of the factors in capturing dynamic market patterns. In this task, STORM shows substantial improvements over six methods across the SP500 and DJ30 datasets. We use

rank-based correlation metrics (i.e., RankIC, RankICIR) instead of error-based ones (e.g., MSE, MAE), as they better capture a factor’s ability to preserve the correct cross-sectional ordering of returns. As shown in Table 4, STORM achieves substantial average improvements of 14.690% on SP500 and 15.360% on DJ30, demonstrating its ability to capture more consistent and generalizable signals across diverse market conditions.

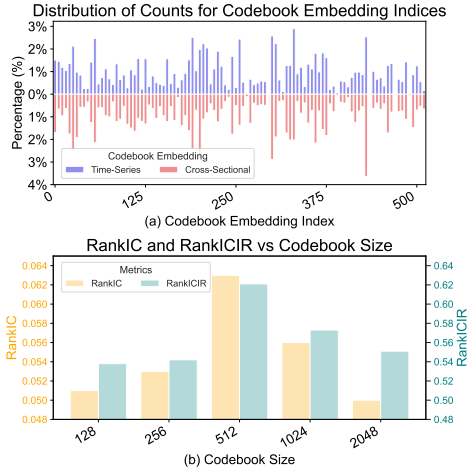


Figure 3: (a) Distribution of counts for codebook embedding indices. (b) Hyperparameter experiment results for different codebook sizes.

Factor Diversity. The diversity of factors is essential in financial investment, as it allows the model to capture the multi-dimensional characteristics of the market, avoiding redundancy or over-reliance on a single factor. In the codebook, each embedding represents a cluster center for a category of factor embeddings. To analyze the utilization rate of each factor category, we report the frequency of codebook embedding indices in the testing data. For clarity, we group every five indices into one interval, as illustrated in Figure 3(a). The results indicate a relatively even distribution of factor usage, highlighting the model’s effectiveness to maintain diversity during the factor selection process.

Codebook Design. Furthermore, we conduct a hyperparameter analysis on the DJ30 dataset to examine the effect of different codebook sizes, aiming to balance factor diversity and representational quality. An excessively large codebook may lead to sparse factor representations, limiting the effectiveness of factor utilization, while an overly small codebook could result in insufficient factor differentiation and information redundancy. Following standard practices, we evaluate five different codebook sizes, including 128, 256, 512, 1024, and 2048. As shown in Figure 3(b), a size of 512 achieves the best overall factor prediction performance, offering a favorable trade-off between expressiveness and efficiency.

5.4 Ablation Study (for RQ3)

We conducted ablation experiments with two simplified versions of STORM, namely STORM-w/o-TS and STORM-w/o-CS, which only extract cross-sectional and time-series factors, respectively. Specifically, the absence of either time-series or cross-sectional factors

significantly reduces the model’s effectiveness across various tasks. In the *stock return prediction* task, STORM outperforms both simplified models, with an average improvement of 16.50% on the RankIC metric and 23.79% on the RankICIR metric, highlighting the enhanced accuracy achieved by incorporating both time-series and cross-sectional factors. In the *PM* and *AT* tasks, STORM and two simplified models all performed well. Notably, STORM-w/o-TS exhibited the lowest AVO among the three models, although its profitability did not surpass that of STORM. These findings suggest that methods like FactorVAE [6], which focus solely on cross-sectional factors, may struggle to capture the full complexity of the financial market. Therefore, integrating both time-series and cross-sectional factors is essential for robust and accurate financial modeling.

5.5 Discussion on Performance

Efficiency. STORM runs efficiently in real-world daily trading without the ultra-low latency or specialized hardware needed for high-frequency systems, even on standard CPUs. This is enabled by a compact architecture with only a few transformer layers and 15.2M FP32 parameters on DJ30, enabling deployment on edge devices while supporting concurrent market analysis and trading.

Effectiveness. STORM delivers consistently strong performance across diverse tasks, including PM and AT, outperforming baselines that often specialize in one domain but fail to generalize. It balances return and risk, targeting comprehensive investment performance rather than over-optimizing for a single metric. Robustness and versatility across tasks justify the design of our proposed model.

Generalization. STORM generalizes across markets of varying scales and characteristics, excelling on both DJ30 (large-cap industrial blue chips) and SP500 (broad, diversified industries). These representative U.S. datasets cover concentrated and diversified indices, and their depth makes them a strong proxy for other markets like China or the U.K.. The flexible encoder-decoder architecture also allows using additional layers like Mamba [4, 10] and ensures applicability to diverse trading environments and asset classes.

6 Conclusions and Future Work

In this paper, we propose STORM, a spatio-temporal latent factor model leveraging a dual VQ-VAE framework to learn both time-series and cross-sectional factors. STORM ensures factor orthogonality and diversity, crucial for effective factor selection in financial trading. Unlike traditional factors, STORM learns high-dimensional latent vectors that, though less interpretable, capture richer market structures. Extensive experiments demonstrate STORM’s superiority over the state-of-the-art methods across a prediction task and two downstream tasks, showcasing robust performance and effectiveness. In the future, we plan to extend STORM by incorporating side information like news, to explore the effects of exogenous factors.

Acknowledgments

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL). It is also supported, in part, by Jinan-NTU Green Technology Research Institute (GreenTRI).

References

- [1] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems* 33 (2020), 12449–12460.
- [2] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. 2021. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 357–366.
- [3] Zhennan Chen, Zhicheng Zhang, Pengfei Li, Lingyue Wei, Shibo Feng, and Fan Lin. 2023. mTrader: A Multi-Scale Signal Optimization Deep Reinforcement Learning Framework for Financial Trading (S). In *SEKE*. 530–535.
- [4] Tri Dao and Albert Gu. 2024. Transformers are SSMs: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *International Conference on Machine Learning (ICML)*.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [6] Yitong Duan, Lei Wang, Qizhong Zhang, and Jian Li. 2022. Factorvae: A probabilistic dynamic factor model based on variational autoencoder for predicting cross-sectional stock returns. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4468–4476.
- [7] Eugene F Fama and Kenneth R French. 1992. The cross-section of expected stock returns. *The Journal of Finance* 47, 2 (1992), 427–465.
- [8] Yuchen Fang, Kan Ren, Weiqing Liu, Dong Zhou, Weinan Zhang, Jiang Bian, Yong Yu, and Tie-Yan Liu. 2021. Universal trading for order execution with oracle policy distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 107–115.
- [9] Joel Greenblatt. 2010. *The little book that still beats the market*. Vol. 29. John Wiley & Sons.
- [10] Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752* (2023).
- [11] Shihao Gu, Bryan Kelly, and Dacheng Xiu. 2021. Autoencoder asset pricing models. *Journal of Econometrics* 222, 1 (2021), 429–450.
- [12] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [13] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems* 30 (2017).
- [14] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [15] Bin Li and Steven CH Hoi. 2014. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)* 46, 3 (2014), 1–36.
- [16] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*. PMLR, 19730–19742.
- [17] Yang Liu, Qi Liu, Hongke Zhao, Zhen Pan, and Chuanren Liu. 2020. Adaptive quantitative trading: An imitative deep reinforcement learning approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2128–2135.
- [18] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [19] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730* (2022).
- [20] James M Poterba and Lawrence H Summers. 1988. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics* 22, 1 (1988), 27–59.
- [21] Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. 2017. A dual-stage attention-based recurrent neural network for time series prediction. *arXiv preprint arXiv:1704.02971* (2017).
- [22] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [23] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. 2019. Generating diverse high-fidelity images with vq-vae-2. *Advances in Neural Information Processing Systems* 32 (2019).
- [24] Stephen Ross. 1976. The arbitrage pricing theory. *Journal of Economic Theory* 13, 3 (1976), 341–360.
- [25] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [26] William F Sharpe. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19, 3 (1964), 425–442.
- [27] Woncheol Shin, Gyubok Lee, Jiyoung Lee, Eunyi Lyou, Joonseok Lee, and Edward Choi. 2023. Exploration into translation-equivariant image quantization. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [28] Thomas Spooner and Rahul Savani. 2020. Robust market making via adversarial reinforcement learning. *arXiv preprint arXiv:2003.01820* (2020).
- [29] Shuo Sun, Molei Qin, Wentao Zhang, Haochong Xia, Chuqiao Zong, Jie Ying, Yonggang Xie, Lingxuan Zhao, Xinrun Wang, and Bo An. 2024. TradeMaster: a holistic quantitative trading platform empowered by reinforcement learning. *Advances in Neural Information Processing Systems* 36 (2024).
- [30] Shuo Sun, Wanqi Xue, Rundong Wang, Xu He, Junlei Zhu, Jian Li, and Bo An. 2022. DeepScalper: A risk-aware reinforcement learning framework to capture fleeting intraday trading opportunities. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 1858–1867.
- [31] Sabera Talukder, Yisong Yue, and Georgia Gkioxari. 2024. TOTEM: Tokenized Time Series Embeddings for General Time Series Analysis. *arXiv preprint arXiv:2402.16412* (2024).
- [32] Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in Neural Information Processing Systems* 30 (2017).
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- [34] Jia Wang, Tong Sun, Benyuan Liu, Yu Cao, and Hongwei Zhu. 2021. CLVSA: A convolutional LSTM based variational sequence-to-sequence model with attention for predicting trends of financial markets. *arXiv preprint arXiv:2104.04041* (2021).
- [35] Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. 2021. 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. 643–650.
- [36] Zikai Wei, Bo Dai, and Dahua Lin. 2022. Factor investing with a deep multi-factor model. *arXiv preprint arXiv:2210.12462* (2022).
- [37] Zikai Wei, Anyi Rao, Bo Dai, and Dahua Lin. 2023. HireVAE: An Online and Adaptive Factor Model Based on Hierarchical and Regime-Switch VAE. *arXiv preprint arXiv:2306.02848* (2023).
- [38] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. 2021. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157* (2021).
- [39] Xiao Yang, Weiqing Liu, Dong Zhou, Jiang Bian, and Tie-Yan Liu. 2020. Qlib: An ai-oriented quantitative investment platform. *arXiv preprint arXiv:2009.11189* (2020).
- [40] Yunan Ye, Hengzhi Pei, Boxin Wang, Pin-Yu Chen, Yada Zhu, Ju Xiao, and Bo Li. 2020. Reinforcement-learning based portfolio management with augmented asset movement prediction states. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1112–1119.
- [41] Fisher Yu and Vladlen Koltun. 2015. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122* (2015), 1–10.
- [42] ShuiLing Yu and Zhe Li. 2018. Forecasting stock price index volatility with LSTM deep neural network. In *Recent Developments in Data Science and Business Analytics: Proceedings of the International Conference on Data Science and Business Analytics (ICDSBA-2017)*. Springer, 265–272.
- [43] Baoquan Zhang, Huaibin Wang, Chuyao Luo, Xutao Li, Guotao Liang, Yunming Ye, Xiaochen Qi, and Yao He. 2024. Codebook Transfer with Part-of-Speech for Vector-Quantized Image Modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7757–7766.
- [44] Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiase Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, et al. 2024. FinAgent: A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist. *arXiv preprint arXiv:2402.18485* (2024).
- [45] Wentao Zhang, Yilei Zhao, Shuo Sun, Jie Ying, Yonggang Xie, Zitao Song, Xinrun Wang, and Bo An. 2024. Reinforcement Learning with Maskable Stock Representation for Portfolio Management in Customizable Stock Pools. In *Proceedings of the ACM on Web Conference 2024*. 187–198.
- [46] Yu Zhao, Huaming Du, Ying Liu, Shaopeng Wei, Xingyan Chen, Fuzhen Zhuang, Qing Li, and Gang Kou. 2022. Stock movement prediction based on bi-typed hybrid-relational market knowledge graph via dual attention networks. *IEEE Transactions on Knowledge and Data Engineering* 35, 8 (2022), 8559–8571.