QuantVisExplorer: A Multi-Perspective Visual Analytics System for Quantitative Investment

Xi Huang[*]
School of Big Data and Software Engineering
Chongqing University

Xinchi Luo[†]
School of Big Data and Software Engineering
Chongqing University

Xuan He[‡]
School of Big Data and Software Engineering
Chongqing University

Baocheng Tang[§]
School of Big Data and Software Engineering
Chongqing University

Hongxing Qin[¶]
College of Computer Science
Chongqing University

Haibo Hu[‖]
School of Big Data and Software Engineering
Chongqing University

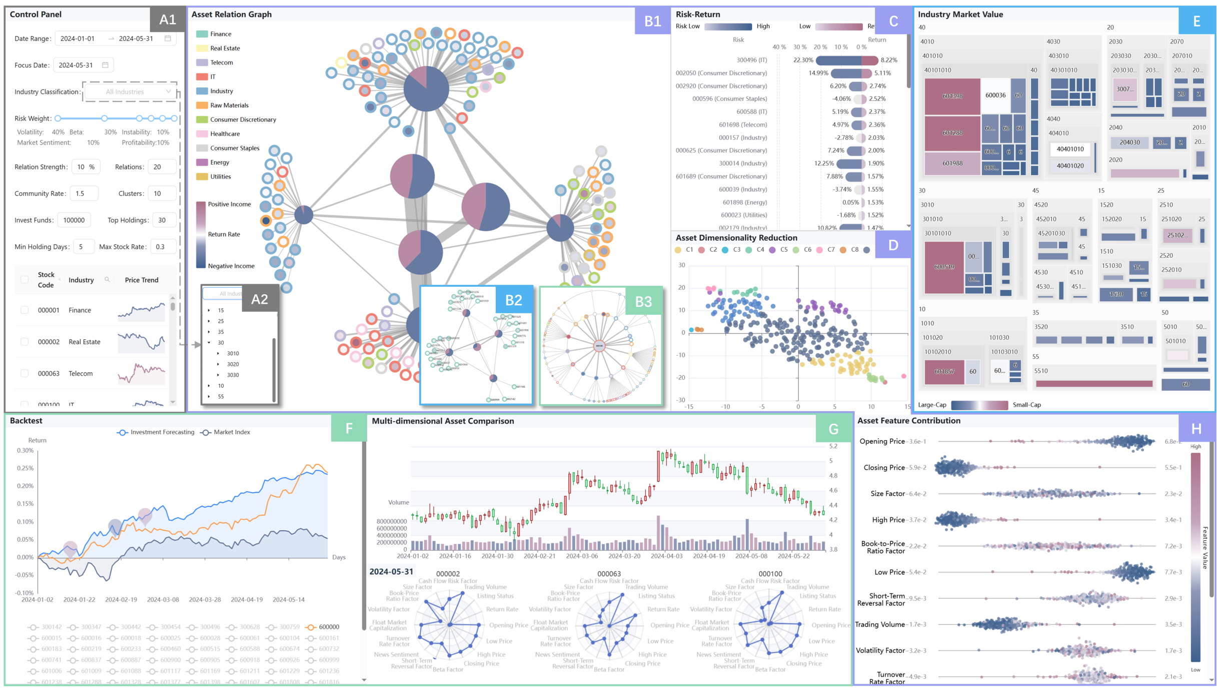

Figure 1: QuantVisExplorer system overview. **A1** Control Panel Supports global interaction functions at the market level, industry level and asset level **A2** Multiple choice tree selector helps users to view the industry structure and select a specific industry. **B** Asset Relation Graph View visualizes multi-faceted asset connections in different states: overview with clustered nodes (B1), industry-focused subgraph (B2), and asset-centric network (B3). **C** Risk-Return View ranks assets by performance and associated risk levels. **D** Asset Dimensionality Reduction View reveals feature-based clusters and outliers. **E** Industry Market Value View depicts hierarchical sector capitalization via treemap. **F** Backtest View simulates investment strategies and compares returns against benchmarks. **G** Multi-dimensional Asset Comparison View tracks selected assets using K-line and synchronized radar charts. **H** Asset Feature Contribution View illustrates feature influence on predictions using beeswarm plots.

Abstract

Quantitative investment increasingly relies on complex models like temporal graph networks for asset return prediction, yet these models often function as "black boxes", hindering user trust and effective decision-making. Concurrently, the multi-source, dynamic nature of financial data poses significant analytical challenges. To address these issues, we present QuantVisExplorer, an interactive visual analytics system designed to synergize a novel Multi-Relational Temporal Graph Fusion (MRTGF) model with a coordinated multi-view interface. The multi-perspective feature of our QuantVisExplorer enables users to conduct multi-level market overview analysis, specifically at the market, industry, and asset levels, explore complex, model-inferred inter-asset correlation networks, track multi-dimensional asset state evolution, and, crucially, understand the MRTGF model's prediction rationale through visualizations of feature contributions and an interactive backtesting module. This integration of a sophisticated predictive model with tailored visual analytics aims to demystify model behavior and enhance analytical capabilities. Case studies with domain experts and a user study demonstrate that QuantVisExplorer significantly improves insight generation, model understanding, and decision-

*e-mail: xihuang@stu.cqu.edu.cn
†e-mail: xc.luo@stu.cqu.edu.cn
‡e-mail: xuanhe@stu.cqu.edu.cn
§e-mail: tangbaocheng@stu.cqu.edu.cn
¶e-mail: qinhx@cqu.edu.cn
‖e-mail: haibo.hu@cqu.edu.cn

making efficiency in quantitative investment compared to traditional approaches. Our primary contribution is this validated visual analytics system that makes complex predictive models more transparent and actionable for financial analysts.

Index Terms: Financial visualization, quantitative investment, temporal graph network, visual analytics.

1 Introduction

Nowadays, increasing volatility and complexity in global financial markets, coupled with diverse data sources such as traditional prices to news and social media, challenge traditional investment paradigms [19]. Quantitative investment offers a data-driven alternative, using models and algorithms to analyze vast market data for scientific decision-making and return prediction [46]. Its core lies in the ability to uncover underlying market patterns and predict future asset returns through the analysis of multidimensional data.

Within quantitative investment, the effectiveness of return prediction serves as a crucial input for investment decisions and concurrently represents a long-standing technical challenge in the field. While various approaches, from traditional statistical models [10, 29, 9, 5] to machine learning [3, 6, 22, 8, 20] and graph-based techniques [41, 16, 32, 12], have been explored, they often struggle to dynamically model complex inter-asset relationships, effectively integrate diverse multi-source data (like news and sentiment), and consistently ensure prediction accuracy in evolving markets.

Although quantitative models can process vast amounts of information to provide predictions, their prevalent black-box mechanisms and lack of transparency hinder trust and understanding[13]. Moreover, while human analysts may outperform quantitative models in scenarios requiring accumulated professional experience[4], they also face limitations in processing high-dimensional dynamic information[18]. This underscores the need for explainable models and effective human-computer collaboration tools to prevent critical errors committed solely by either humans or models. Against this backdrop, explainable visual analytics targeting quantitative models offers a promising avenue to bridge this gap.

To make sense of complex financial data, visual analytics (VA) provides powerful tools for pattern discovery and decision support [37, 26]. Yet, current platforms, exemplified by the Bloomberg Terminal, primarily offer real-time data display with limited capabilities for unraveling intricate asset interrelationships or interpreting the outputs of advanced predictive models. This gap restricts deeper understanding and trust in quantitative strategies, motivating the development of VA systems that specifically address these challenges.

To address these challenges, we propose QuantVisExplorer, a visual analytics system integrating (1) a novel multi-relational temporal graph network (MRTGF) for asset return prediction, which fuses multi-source heterogeneous data and (2) a coordinated multi-view visual analytics framework for the interactive exploration and visual presentation of model prediction results, asset correlation structures, and internal model feature contributions. QuantVisExplorer is specifically designed to tackle the insufficient modeling of asset correlations in return prediction, to visually analyze and present key factors driving its model's decisions, and to mitigate cognitive friction during analytical decision-making. The use of this system facilitates a better understanding of market dynamics, exploration of complex inter-asset relationships, and interpretation of the model's predictive logic, thereby fostering more transparent, reliable, and efficient quantitative investment analysis and decision-making. The contributions of this paper are as follows:

- We propose a multi-source heterogeneous data processing method for quantitative investment, constructing a multi-dimensional asset data processing framework.
- A novel temporal graph network, MRTGF, that uniquely processes and fuses multi-source heterogeneous data to capture multi-relational asset dynamics, improving the accuracy and robustness of quantitative investment return predictions.
- A multi-view, interactive visual analytics system, QuantVisExplorer, designed to enhance the interpretability of the MRTGF model and support comprehensive exploration of quantitative investment data across market, asset correlation, temporal evolution, and model feature dimensions.

2 Related Work

2.1 Investment Return Prediction

Investment return prediction, a key objective of quantitative investment, requires deep mining and analysis of market data to forecast future asset returns and manage risk. Early approaches to return prediction relied on statistical models such as CAPM [10], APT [29], time series models such as ARIMA [31], GARCH [2], and multi-factor models [9, 5]. While interpretable, these methods often struggle with the non-linear dynamics and high dimensionality of modern financial data. To overcome the limitations of statistical methods in handling high-dimensional, non-linear data, and to capture data dynamics and temporal dependencies, past research employed machine learning and deep learning methods, including SVMs, RFs [3], and sequence models like RNNs/LSTMs [6], Transformers [36, 8, 22], and VAEs [20], have shown promise in capturing non-linearities. However, they can be prone to overfitting, require extensive feature engineering [38], or may not explicitly model inter-asset relationships. More recently, research has focused on exploring the complex relational structures between assets by employing Graph Neural Networks (GNNs), such as GCNs and GATs. This approach involves abstracting the financial market as a graph to uncover dynamic interaction effects among assets. These GNNs have been applied to model asset relationships and price co-movements [41, 32].

To address the pronounced dynamics inherent in financial data, GNNs have been further extended into Temporal Graph Networks (TGNs). TGNs are also widely employed in return prediction within quantitative investment. By combining GNNs with time series modeling, TGNs can simultaneously capture both the spatial dimension (asset correlations) and the temporal dimension (time series characteristics) features [11, 16, 12]. They analyze the temporal changes in asset prices while considering their co-movement effects with related assets. For example, Feng et al. [11] constructed stock correlation graphs based on Pearson correlation coefficients and combined them with decomposed time-series features, utilizing GAT and Temporal Convolutional Networks (TCN) for stock price prediction. Furthermore, some studies attempt to incorporate multi-source data, such as news and social media [17], to describe market information more comprehensively. For instance, news events can be represented as nodes connected to asset nodes to capture the potential impact of news on asset prices, thereby enhancing prediction accuracy.

However, most existing studies predominantly rely on predefined market structures and classifications to define static relationships, often overlooking potential latent or dynamic correlations among assets. This research aims to bridge this gap by further mining asset relationships through the analysis of financial news data. By identifying market hotspots and focus points, we capture timely, evolving individual relationships and reconstruct the asset network graph to better reflect market evolution patterns, ultimately seeking to enhance the accuracy and robustness of return predictions.

2.2 Financial Data Visualization

Visualization is crucial for understanding complex financial data, moving beyond simple line charts to reduce cognitive load [15]. Early efforts augmented line charts with additional encodings to display more data dimensions [30, 47, 33].

Beyond static displays, visualizing the dynamic evolution of financial data and enabling interactive exploratory analysis are essential for deeper insights [25, 7]. While general techniques for time-series visualization and exploration support continue to advance [35, 40], our focus shifts to interactive systems specifically designed for quantitative investment support.

Interactive visual analytics systems like Bitextract [44] for transaction patterns, and sPortfolio [42], iQUANT [43], and RankFIRST [14] for portfolio/factor analysis have demonstrated the value of VA in finance. However, few systems deeply integrate the interpretability of advanced prediction models like temporal graph networks or provide comprehensive support for understanding multi-relational asset dynamics alongside model explanations. QuantVisExplorer aims to fill this gap by offering seamless MRTGF model explanations alongside a multi-faceted visual exploration capability, supporting drill-down analysis of market performance, asset temporal evolution, inter-asset correlation networks, and prediction results.

2.3 Explainable Visual Analytics for Financial AI Models

The 'black-box' nature of complex quantitative models necessitates explainable AI (XAI) to build trust and improve decision-making.

Visual analytics for explainable investment models aims to reveal their internal states and complex transformations, thereby helping investors understand model outputs and assess reliability. Research in Visualization for AI (VIS4AI) plays a crucial role in this, assisting users in developing, understanding, and improving models, particularly in financial applications [34]. For instance, some VIS4AI approaches address the interpretability of specific neural network architectures relevant to finance. Lai et al. [21] proposed a workflow for rapidly prototyping visual interfaces for various neural networks, and TransforLearn [12] offers interactive exploration of Transformer models, which are increasingly used in financial forecasting. Beyond visual analytics tailored to model architectures, general XAI techniques like LIME [28] and SHAP [24] are common for model-agnostic explanations. Other research strives to make complex models more transparent. For example, Prasad et al. [27] proposed the 'Transform-and-Perform' framework to explain complex input-output relationships in high-dimensional problems. Closer to our system's goals, Zang et al. [45] designed the DMT-EV interactive interface, enabling investors to understand simplified model representations, analyze explanation outputs, and grasp the impact of hyperparameters. This aligns with QuantVisExplorer's focus on providing clear model insights through interactive visualization for financial decision-making. While these methods offer general model explanation capabilities, or visualize specific model types, there is a need for tailored visual analytics systems that can unpack the predictions of complex temporal graph networks like MRTGF in the context of financial data, linking feature importance, graph structure, and temporal dynamics.

3 REQUIREMENT AND PIPELINE

3.1 System Requirements Analysis

Building upon the challenges identified previously regarding return prediction accuracy, model interpretability, and cognitive load in quantitative investment, and through multiple rounds of in-depth interviews and discussions with domain experts about obstacles in existing analytical workflows, we derived the following key requirements for our visual analytics system:

R1: Efficient Market Situation Awareness. Provide comprehensive, multi-level overviews (market, industry) of returns and key statistics for rapid assessment of trends, risks, and opportunities.

R2: In-depth Exploration of Asset Correlations. Visualize explicit and data-driven dynamic asset correlations, supporting interactive exploration of network structures and the roles/influence of assets and groups.

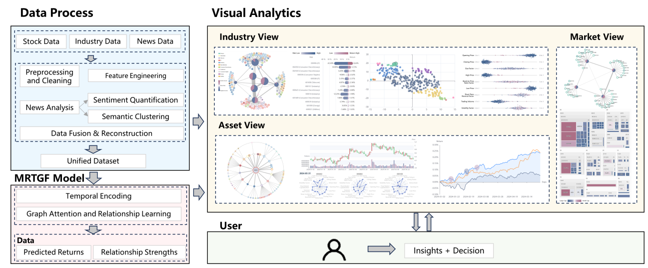

Figure 2: The QuantVisExplorer architecture and pipeline

R3: Tracking of Asset State Evolution. Offer interactive time-series visualizations to track and compare multi-dimensional asset state evolution at various granularities, aiding the identification of performance drivers.

R4: Transparent Interpretation of Model Prediction Process. Enhance model explainability by visualizing key feature contributions driving predictions. Integrate backtesting results to assess prediction reliability and rationale.

3.2 Pipeline

The QuantVisExplorer pipeline Fig. 2 proceeds as follows. Firstly, multi-source financial data, including stock data, industry data, and news data, undergo a data process stage. This involves preprocessing and cleaning, feature engineering, and news analysis (for sentiment quantification and semantic clustering), which culminates in a unified dataset. Secondly, the MRTGF model utilizes this dataset, employing temporal encoding and graph attention and relationship learning to generate predicted returns and relationship strengths. Finally, these outputs are channeled into the visual analytics interface. This interactive system features three coordinated views—the industry view, the market view, and the asset view—which empower the user to derive insights and make informed decisions.

4 METHODS

4.1 Data Processing

To comprehensively capture market dynamics and complex inter-asset correlations in support of subsequent quantitative investment return prediction and visual analytics, this study constructed a multi-source heterogeneous dataset. This dataset covers information from the Chinese A-share market for the period from September 30, 2022, to September 30, 2024, and primarily consists of three components: First, daily comprehensive stock data for the CSI 300 Index components, sourced from the RESSET Financial Research Database [1]. This includes 18 key daily trading indicators for constituent stocks (e.g., OHLC prices, volume), providing a foundation for time series analysis. Second, the CSI Industry Classification Standard (2021 version), also obtained from the RESSET database. This provides detailed hierarchical classification information used for analyzing industry correlations. Finally, related financial news data were obtained from stcn.com (Securities Times Network), a major financial news source.This data includes news headlines, body text, publication times, and keywords, and is used to extract market sentiment and semantic information.These raw datasets collectively form the data foundation for data processing, model training, and visual analytics conducted in this study.

To prepare this multi-source heterogeneous dataset for the subsequent return prediction model and visual analytics, we performed a series of data processing steps. First, raw data underwent rigorous preprocessing. This involved Cleaning stock trading records by employing linear interpolation alongside forward and backward

[1] www.resset.com

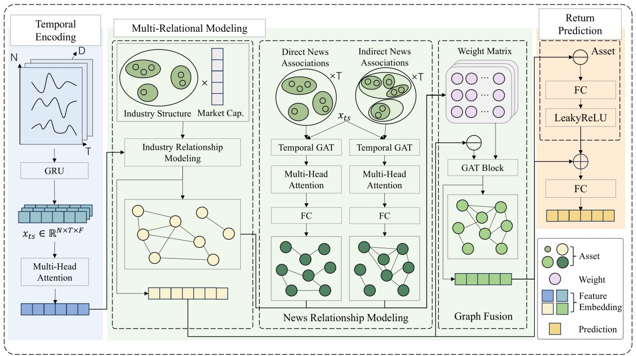

Figure 3: The overall architecture of MRTGF.

fill techniques to handle missing values, and removing assets that exhibited excessive missing data. Financial news texts were similarly cleansed by removing HTML tags and filtering articles that were overly short or excessively long, thereby ensuring information quality and processing efficiency. Second, to enrich the feature set for asset analysis, we incorporated seven well-established investment risk factors, including size (LOGCAP), Beta, volatility (measured by SMA120), turnover rate, short-term reversal, cash flow risk, and book-to-price ratio. Third, to leverage the rich information within financial news, we quantified news sentiment using the Baidu AI platform API and semantically clustered news articles. Specifically, news sentiment was numerically represented: 1 for positive, 0 for neutral, and -1 for negative. These representations were subsequently weighted by confidence scores. For semantic understanding, daily news texts were converted into embedding vectors via a finance-specific pre-trained model, FinBERT, and then clustered using the MiniBatchKMeans algorithm to identify latent topics and inter-asset relationships reflected in the news. Finally, we performed multi-source data reconstruction to create a unified data representation. This entailed temporally aligning the processed trading data, risk factors, and news-derived features like sentiment scores and semantic cluster IDs for each asset. A crucial step was establishing relevance between news articles and specific assets, or their respective industries if a direct asset match was not found, through a hierarchical keyword matching strategy. This allowed us to aggregate news sentiment at the asset level, considering both direct mentions and industry-wide sentiment, weighted by factors such as mention frequency and industry-level market capitalization. The resulting structured and enriched dataset for the CSI300 components, characterized by dimensions of (483, 282, 17) features, served as the input for our return prediction model and the QuantVisExplorer system.

4.2 The MRTGF Return Prediction Model

The primary goal of our underlying prediction task is to forecast the future return of financial assets. Given a set of N assets, let $X_t = [x_1^t, x_2^t, \ldots, x_N^t]$ represent the multi-dimensional feature matrix for these assets at time $t \in \{1, \ldots, T\}$, where $x_i^t \in \mathbb{R}^D$ is the D-dimensional feature vector for asset i. Let $A_t \in \mathbb{R}^{N \times N}$ denote the relationship matrix among assets at time t. The objective is to learn a mapping function $f(\cdot)$ such that $\hat{y}^{T+1} = f(X_1, \ldots, X_T; A_1, \ldots, A_T)$ where $\hat{y}^{T+1} \in \mathbb{R}^N$ is the predicted return score for each asset at the next time step $T+1$. These predictions, along with the modeled relationships, form a complex information space that is challenging for investors to interpret directly, motivating the need for our visual analytics system.

4.2.1 Architecture Overview

To generate reliable return predictions and insightful inter-asset relationships, we employ the Multi-Relational Temporal Graph Fusion (MRTGF) model, an architecture specifically designed for quantitative investment. As illustrated in Fig. 3, MRTGF integrates multi-source financial data through several key stages. It begins by processing historical asset features and constructing dynamic asset relationships from sources like industry classifications and financial news. These are then fed into a temporal graph network that captures both the temporal evolution of individual assets and the complex interplay between them. The model ultimately outputs predicted return scores for each asset and a quantified representation of the learned multi-faceted asset relationships. The richness and complexity of these outputs necessitate a visual analytics approach for effective exploration and comprehension by investment professionals.

4.2.2 Key Mechanisms

A core strength of MRTGF lies in its comprehensive modeling of the financial market. For Temporal Feature Extraction, the model processes time-series asset data, which includes historical prices, trading volumes, and derived factors, by employing Gated Recurrent Units (GRUs) to capture sequential dependencies and multi-head attention mechanisms to identify crucial temporal patterns influencing future returns. This allows MRTGF to discern evolving asset-specific characteristics over time. For Multi-Relational Modeling, which is pivotal for our visual analytics objectives, MRTGF constructs and reasons over diverse inter-asset relationships. These relationships are multifaceted: (1) Industry-based relationships are derived from official sector classifications to reflect inherent structural correlations and peer group effects. (2) News-driven direct relationships are established by linking assets that are co-mentioned or directly implicated within the same financial news events, thereby capturing event-specific impacts. (3) News-driven indirect (semantic) relationships connect assets through their shared exposure to similar semantic themes or topics. These are identified by performing textual analysis and semantic clustering on the corpus of financial news articles, allowing the model to capture broader market sentiment and thematic influences. Furthermore, MRTGF integrates these varied relationships using a GAT mechanism to effectively weigh and propagate information across the constructed asset network. These learned temporal dynamics and fused multi-faceted relationships provide a rich, yet complex, foundation for the subsequent visual exploration in QuantVisExplorer.

4.2.3 Outputs for Visual Analytics

The MRTGF model generates key outputs that are integral to the QuantVisExplorer system, facilitating in-depth visual analysis. Primarily, it produces predicted return scores for each asset, which indicate expected future performance and directly inform portfolio construction scenarios and performance evaluation within the visual interface. Secondly, it yields learned asset relationship strengths, quantifying the connections from each modeled relationship type, including industry-based, news-direct, and news-semantic, as well as a fused overall asset network graph. These quantified relationships form the backbone of interactive network visualizations, enabling users to explore market structure and the propagation of influence. The high dimensionality, temporal nature, and interconnectedness of these model outputs necessitate the visual analytics capabilities provided by QuantVisExplorer for effective user navigation and comprehension.

5 VISUALIZATION

5.1 Visual Task Analysis

Based on the analysis of the user requirements aforementioned (R1-R4), we have defined the following analytical tasks at the system level. These tasks delineate the core functionalities and support capabilities that the visual analytics system must provide to effectively address user challenges:

T1: Multi-level Market and Asset Feature Overview. Supporting layered exploration from macro (overall market structure, dominant groups) to micro (individual asset performance) perspectives, empowering users to rapidly identify patterns, trends, and anomalies. To achieve efficient awareness of the market situation (R1), the system must allow users to gain an overview of the market return distributions, volatility, and key indicators at both the industry and asset levels. Users should be able to explore asset characteristics (e.g., price, volume, risk factors) to understand similarities, identify clusters, outliers, and assess relative positioning.

T2: Support Visual Exploration and Analysis of Complex Inter-Asset Correlation Networks. To allow in-depth exploration of asset correlations (R2), the system must enable users to visually explore and analyze multiple types of dynamic relationships, including explicit domain-knowledge-based and implicit model-mined correlations. This involves providing interactive capabilities such as filtering, querying, and dynamic layout adjustments to help users identify network topology, key entities, and their influence.

T3: Enable Tracking and Comparison of Multi-dimensional Temporal Asset States. To address the need for tracking asset state evolution (R3), the system must enable users to monitor and compare the temporal changes of key multi-dimensional indicators (e.g., price, volume, risk factors) for selected assets or groups. This includes facilitating side-by-side comparisons to help identify evolutionary patterns, trends, and significant events over time.

T4: Facilitate Explainable Analysis and Validation of Model Prediction Results. To enhance model transparency and explainability (R4), the system must allow users to understand the rationale behind predictions and validate their effectiveness. This involves visually demonstrating the contribution of input features to prediction outcomes and providing an interactive backtesting environment. Users should be able to configure simulated trading strategies based on model predictions and visually assess their cumulative returns against benchmarks to gauge practical utility and reliability.

These system-level analytical tasks form the core basis for our subsequent visualization and interaction design, ensuring that the system's functionalities directly address and satisfy the analytical needs of the users.

5.2 System Overview

To address the analytical tasks T1-T4, we designed and implemented QuantVisExplorer, an interactive visual analytics system for explainable quantitative investment return prediction. QuantVisExplorer promotes a comprehensive understanding of market dynamics, inter-asset relationships, individual asset performance, and model prediction rationale through a coordinated multi-view interface, as illustrated in Fig. 1. The system is structured to support a multi-level analytical workflow, enabling users to seamlessly navigate from broad market assessments to granular asset-specific details.

At the market level, the system enables users to understand overall market structure, asset performance distributions, and feature-based asset clusters. For industry-level analysis, users can explore sector-specific dynamics and inter-industry linkages. At the asset level, the system allows for detailed comparison of the multi-dimensional performance of selected assets over time and assessment of model-driven trading strategies. The Asset Relation Graph View is used across these levels to visualize asset connectivity.

A central Control Panel underpins these explorations, offering tools for dynamic data filtering, temporal range selection, and adjustment of analytical parameters across all views, ensuring a cohesive and user-driven analytical experience.

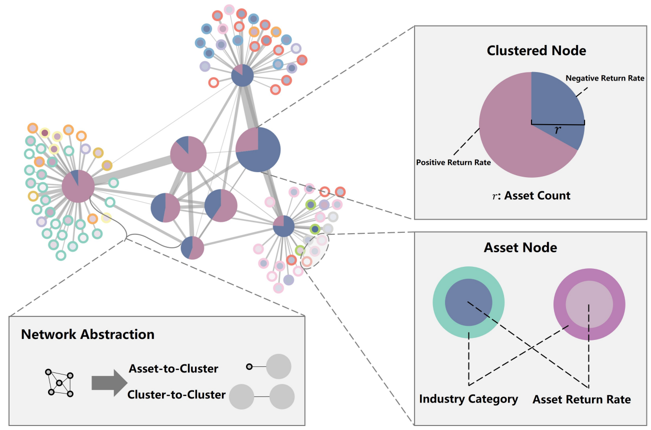

Figure 4: Element design of the asset graph view.

5.3 Control Panel

The Control Panel (Fig. 1A) serves as the central hub for user-driven data exploration and parameter adjustment, crucial for supporting iterative analysis across all tasks (T1, T2, T3, T4). It empowers users to define the analytical scope by selecting time ranges, specific industries through a multiple choice tree selector, and individual assets via an interactive table with preview trendlines and search functionality. Furthermore, it allows users to dynamically adjust parameters for various views, such as risk factor weights for the Risk-Return View, thresholds for relationship filtering in the Asset Relation Graph View, and strategy parameters for the Backtest View. This centralized control is key to enabling a flexible and personalized analysis workflow. The design ensures that users can easily tailor the visualizations to their specific interests and hypotheses, enhancing the system's utility and user engagement.

5.4 Multi-Level Analytical Views

QuantVisExplorer is designed to support a multi-level analytical workflow, guiding users from a broad market understanding down to granular asset-specific details and model explanations. The following sections describe the key visual components tailored for market-level, industry-level, and asset-level analysis, all referencing the main system interface shown in Fig. 1.

5.4.1 Market-Level Analysis Views

At the market level, the system provides a holistic perspective on overall market conditions, performance distributions, structural correlations, and broad factors influencing predictions.

The Risk-Return View (Fig. 1C) presents a ranked comparison of assets based on their returns and associated risks. Using a dual-axis chart design, asset returns (predicted or historical) are visualized as red horizontal bars extending to the right, while associated risk metrics, which represent a composite score from factors like volatility and Beta with weights adjustable in the Control Panel, are shown as blue bars extending to the left. Assets are sorted by return, enabling users to quickly identify top-performing or high-risk assets and understand risk-return trade-offs for initial screening (T1).

The Asset Dimensionality Reduction View (Fig. 1D) employs t-SNE to project high-dimensional asset features (including price, volume, and risk factors) onto a 2D scatterplot, where each point signifies an asset. Crucially, K-Means clustering is applied prior to projection to group similar assets, with these pre-defined clusters distinguished by color. This approach facilitates an understanding of feature-based market segmentation, allowing users to identify co-behaving asset groups, spot outliers, and understand market structure from a feature similarity perspective (T1, T2). Interactive highlighting and hover-for-details enhance its exploratory value.

The Asset Relation Graph View (Fig. 1B) is central to visualizing complex inter-asset relationships, innovatively displaying MRTGF model-learned relations, including industry affiliations and news-driven connections (T2). Its visual encoding is detailed in Fig. 4. In its Overview State(Fig. 1B1), designed for understanding macro-level market structure (T1), assets are initially grouped into community clusters via the Louvain algorithm[1]. As shown in Fig. 4, clusters are represented as pie charts where fill color (blue-red divergent for negative/positive returns) indicates internal asset return distribution, and pie size reflects asset count. Edges denote aggregated inter-cluster relationship strengths. Individual nodes (assets) have border colors encoding industry classification. The system defaults to highlighting key cluster types (e.g., highest positive/negative return proportions, largest/smallest asset counts) to simplify initial exploration. This abstraction is vital for identifying broad market segments and dependencies. General interactions like zooming, hover-for-details, and Control Panel-driven filtering (by strength/type) facilitate progressive insight discovery across all states of this view. The view's novelty lies in its multi-scale, adaptive representation of complex, model-derived relationships.

The Asset Feature Contribution View (Fig. 1H) visualizes the overall importance of input features to the MRTGF model's return predictions across all assets to enhance model transparency at a broad level (T4). This view utilizes a beeswarm plot design. Vertically, the 14 input features are listed, typically ordered by their total contribution (e.g., mean absolute SHAP value) to market-wide predictions. Horizontally, for each feature, data points representing assets are positioned by their SHAP value, with color encoding the actual feature value. This helps users identify which features are most influential market-wide and how their distributions impact predictions generally.

5.4.2 Industry-Level Analysis Views

For industry-level, users can in-depth study into sector-specific dynamics, examining composition, significance, and interconnections.

The Industry Market Value View (Fig. 1E), a treemap visualization, depicts the hierarchical structure of the market based on industry classifications and the relative market capitalization of industries and their constituent assets. Rectangle areas are proportional to market value, and their fill color is mapped to performance indicators like return rates. This effectively communicates both the scale and performance of different market sectors (T1). Users can interactively drill down for a more granular view.

When specific industries are selected via the Control Panel, the Asset Relation Graph View transitions to its Industry-Focused State (Fig. 1B2). Here, it displays relevant subgraphs where asset nodes within these industries become prominent (often with asset codes displayed), aiding the analysis of intra-industry and inter-industry connectivity (T2).

5.4.3 Asset-Level Analysis Views

At the asset level, the system facilitates in-depth examination of individual asset, their temporal states, model-driven strategy performance, and specific prediction explanations.

Detailed Asset Profiling and Comparison. To facilitate detailed tracking and comparison of multi-dimensional temporal asset states (T3), the Multi-dimensional Asset Comparison View (Fig. 1G) offers a synchronized visualization of selected assets. Its upper portion employs K-line charts for daily trading dynamics (OHLC prices, volume) over a user-defined period. Coordinated with this, the lower portion features parallel radar charts, one per asset, synchronized to a user-selected date on the K-line's time axis. Each radar chart visualizes multiple asset-specific features (risk factors, sentiment scores, etc.). This dual design enables identification of co-movements, performance divergences, and event impacts.

When an individual asset is selected, the Asset Relation Graph View activates its Asset-Centric State (Fig. 1B3). The target asset enlarged at the center and its direct and second-degree neighbors on concentric rings, clearly mapping its local influence network and relational pathways (T2, T3).

The Backtest View (Fig. 1F) addresses the need for validating model predictions and assessing practical utility (T4). It allows users to simulate investment strategies based on MRTGF predictions. The view displays cumulative profit and loss (P&L) curves (user strategy vs. market benchmark) with overlaid transaction markers (red for buy, blue for sell) on strategy or individual asset return lines. Strategy parameters are adjustable via the Control Panel (Fig. 1A), enabling evaluation of model effectiveness and strategy refinement.

For specific predictions, the Asset Feature Contribution View (Fig. 1H) explains why an individual asset received its prediction (T4). When an asset is selected, its data points are highlighted within the beeswarm plots, showing how its specific feature values contribute to its predicted return via their SHAP values. This helps users understand decision-making logic for individual cases and build trust.

6 EVALUATION

6.1 Experimental Study

To validate the effectiveness of our proposed MRTGF model for quantitative investment return prediction, we conducted comprehensive experiments on three real-world Chinese A-share market datasets: the SSE 50 Index (SZ50), CSI 300 Index (CSI300), and CSI 500 Index (ZZ500). The SZ50 (40 stocks, 6 features) and ZZ500 (206 stocks, 6 features) datasets spanned January 2021 - December 2024, sourced from BaoStock. The CSI300 dataset (146 stocks, 14 features after processing as per Section 4.1) covered September 2022 - September 2024.

We compared MRTGF against several representative baselines: Transformer, a hybrid model ALSTM+TRA [23], and graph-based model THGNN [39]. Model performance was evaluated using six standard metrics: Information Coefficient (IC), IC Information Ratio (ICIR), Rank Information Coefficient (Rank-IC), Rank-ICIR, Annualized Return (AR), and Information Ratio (IR). All models were implemented in PyTorch, trained using the Adam optimizer (learning rate 2×10^{-4})for 100 epochs, with results averaged over 10 runs. Key parameters included a hidden dimension of 64, 4 attention heads, and a 10-day input window, predicting the next day's return. Our backtesting strategy involved a 100,000 CNY initial capital, a 3-day trading cycle, top-K stock selection based on predictions ($K_{SZ50}=10, K_{CSI300}=30, K_{ZZ500}=50$), and a 30% maximum capital allocation per stock. The comparative results for prediction metrics are presented in Table 1, and backtesting results are in Table 2. These results indicate that MRTGF generally outperforms the baseline models across the three datasets, particularly in key metrics like IC, Rank-IC, and AR. For instance, on the CSI300 dataset, MRTGF achieved an IC of 0.070 and an AR of 0.310, markedly higher than the baselines. While some baselines showed strong performance on specific metrics or datasets, MRTGF demonstrated more consistent and superior overall performance, especially in backtesting. This robust predictive capability and its capacity to uncover multi-faceted asset relationships establish MRTGF as a solid foundation for the visual analytics facilitated by QuantVisExplorer.

6.2 Case Study

6.2.1 Market Trend Analysis

This case study, illustrated in Fig. 5, demonstrates how QuantVisExplorer can assist domain experts in identifying and understanding

Table 1: Comparative Experiment Results (Prediction Metrics). 'Trans.' for Transformer, 'A+T' for ALSTM+TRA.

Dataset	Model	IC	ICIR	R-IC	R-ICIR
SZ50	Trans.	0.067	0.33	0.071	0.41
	A+T	0.075	0.41	0.075	0.42
	THGNN	0.073	0.43	0.074	0.46
	MRTGF	**0.076**	**0.46**	**0.079**	**0.52**
CSI300	Trans.	0.040	0.28	0.041	0.30
	A+T	0.054	0.38	0.055	0.40
	THGNN	0.061	0.41	0.063	0.43
	MRTGF	**0.070**	**0.46**	**0.072**	**0.48**
ZZ500	Trans.	0.039	0.41	0.044	0.42
	A+T	0.044	0.38	0.047	0.46
	THGNN	0.050	0.36	0.059	0.42
	MRTGF	**0.061**	0.41	**0.067**	**0.44**

Table 2: Comparative Experiment Backtest Results. 'Trans.' for Transformer, 'A+T' for ALSTM+TRA.

	SZ50		CSI300		ZZ500	
Model	AR	IR	AR	IR	AR	IR
Trans.	0.261	1.40	0.157	1.16	0.101	1.21
A+T	0.271	1.70	0.206	1.13	0.106	1.18
THGNN	0.269	1.65	0.251	2.21	0.113	1.28
MRTGF	**0.277**	**1.79**	**0.310**	**2.40**	**0.122**	**1.39**

anomalous market trends and their underlying drivers by facilitating a multi-faceted visual exploration process. The study involved an expert (E1), a researcher specializing in macroeconomic market management, who used QuantVisExplorer to analyze the A-share market (CSI300 components) from March 30, 2024, to September 30, 2024. E1 began by examining broad market indicators. As

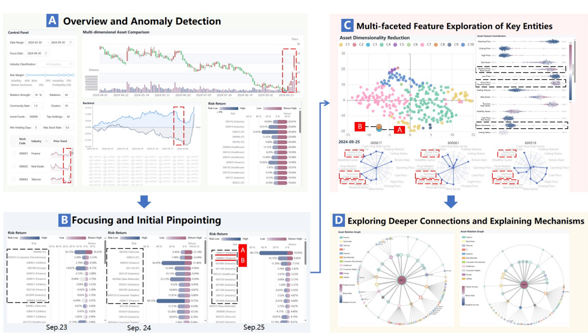

Figure 5: Visual analytics workflow for market trend analysis in QuantVisExplorer

shown in Fig. 5A , E1 primarily used the overview state of the Asset Relation Graph View and the Risk-Return View. An unusual, synchronized surge in asset returns became apparent in late September. Specifically, on September 30, the aggregated cluster nodes in the network view predominantly showed positive internal return distributions, with many individual assets exhibiting unusually high and similar returns (approaching 20%), a stark deviation from typical market behavior.

To investigate this anomaly, E1 narrowed the focus to the days leading up to this surge, detailed in Fig. 5B. The Risk-Return View for September 25 revealed an atypical concentration of high-performing assets from the industrial sector. Intriguingly, the top two performers, assets with stock codes '000617' and '600061' (henceforth Asset A and Asset B respectively), belonged to the financial sector. Selecting these two assets, E1 then consulted multiple views as depicted in Fig. 5C. The Asset Dimensionality Reduction View showed both Asset A and Asset B as distinct outliers. The Multi-dimensional Asset Comparison View, specifically its radar charts for September 25, highlighted their prominence in book-to-price ratio and news sentiment scores. This observation was corroborated by the Asset Feature Contribution View, which indicated that high positive news sentiment and favorable book-to-price ratios were significant positive contributors to their predicted returns on that day.

These visual cues led E1 to hypothesize a recent major positive news event impacting state-affiliated enterprises. A quick check of contemporaneous financial news confirmed a significant policy announcement by the People's Bank of China on September 24 regarding an interest rate cut, which historically boosts market sentiment and benefits state-owned enterprises.

To explore the connection between the high-performing financial assets Asset A, Asset B and industrial assets, E1 utilized the asset-centric state of the Asset Relation Graph View, as shown in Fig. 5D. Focusing on Asset A and Asset B individually, the view revealed strong learned relationships with several state-owned industrial enterprises, particularly in infrastructure and transportation sectors. E1 concluded that the co-movement was driven by the policy announcement, with the financial entities Asset A and Asset B providing crucial support such as financing, to these key industrial players, a connection visually explorable through their respective ego-centric networks.

This case study illustrates QuantVisExplorer's effectiveness in enabling an expert to move from detecting broad market anomalies to identifying specific affected assets, understanding the contributing factors through feature analysis, and finally, uncovering complex inter-asset relationships that explain the observed trends, all within an integrated visual analytics environment.

The study involved an expert (E2), a seasoned quantitative trading strategy developer, analyzing the model's backtest performance on the test set (January 1, 2024, to September 30, 2024).

E2 began by utilizing the Backtest View to assess the impact of different strategy parameters. By systematically adjusting the shortest holding period for assets, E2 observed through the visualized cumulative Profit & Loss (P&L) curves that a 5-day holding period struck an effective balance between capturing market trends and mitigating risks associated with overly frequent or infrequent trading, as shown in Fig. 6. Similarly, when varying the number of top-K predicted assets to include in the portfolio, the backtesting results indicated that while a smaller K could yield high returns, it also came with increased volatility. A larger K often led to more stable, excess returns, highlighting the trade-off between aggressive positioning and diversification.

6.2.2 Analysis of projected investment returns

Figure 6: Parameter-driven comparison of backtest strategies.

To further scrutinize the model's stock-picking capabilities, E2 focused on a strategy holding the top 10 predicted assets. The Backtest View consistently showed the strategy's cumulative returns outperforming the market benchmark while also showing that the model has good ability to obtain excess returns. An examination of the traded assets revealed a predominance of selections from the financial and industrial sectors. The visualization of transaction records highlighted distinct trading patterns; for instance, as shown in Fig. 7A financial stocks were traded frequently, but experienced a brief hiatus in mid-March, during which trading in industrial assets became more concentrated.

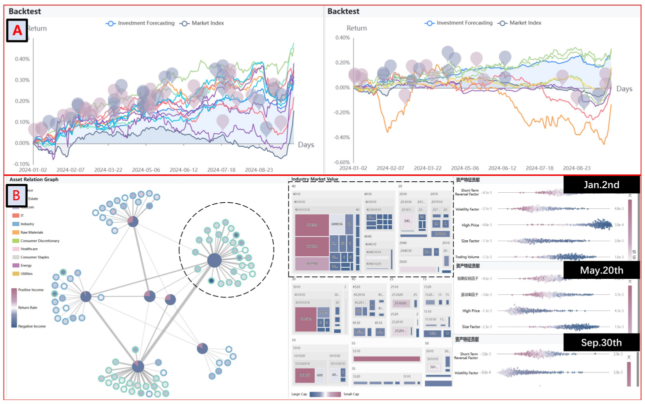

Figure 7: **A** Trading records of financial(left) and industrial(right) stocks. **B** Market performance of financial and industrial industry.

To understand this behavior, E2 consulted broader market views for mid-March. Fig. 7B indicated a generally negative sentiment and performance for both financial and industrial sectors during that specific period, even though the Industry Market Value View confirmed their continued market cap dominance. Turning to the Asset Feature Contribution View for some days, E2 noted that features such as short-term reversal, volatility and trading volume were consistently among the most influential for the model's predictions. E2 inferred that the model's tendency to favor these generally larger, more stable, and less volatile sectors was a result of these combined feature influences.

6.3 User Study

To evaluate QuantVisExplorer's effectiveness and usability, we conducted a user study with 14 participants from diverse backgrounds; 12 had prior investment experience. The participants were evenly divided into an experimental group using QuantVisExplorer and a control group using other common analysis tools, such as Excel and investment data terminals. Before the study, all participants received necessary prerequisite knowledge tailored to their professional backgrounds. During the hands-on session, each participant needed to complete multiple tasks detailed in Table 3, which are designed from the visual tasks (T1-T4) detailed in Section 5.1. All participants provided feedback, which included rating QuantVisExplorer on five 5-point Likert scale aspects (1=strongly disagree, 5=strongly agree). We designed five questions as follows:

Q1: Is the system's layout intuitive and its views practical?

Q2: Is system helpful for identifying investment opportunities?

Q3: Do the system's explanations enhance trust in the model?

Q4: Is the system effective for your quantitative analysis?

Q5: Does the system impose no additional interaction burden?

Table 3: User's analytical tasks in the experiment with a given specified period

No.	Task Description
T1-1	Find the industry with the best overall return.
T1-2	How many significant market fluctuations occurred?
T2-1	Identify the two most closely related industries.
T2-2	Find the association path between two specific assets.
T3-1	Find the five most suitable assets for investment.
T3-2	Identify potential alternative assets similar to a target asset.
T4-1	Find the feature that has the greatest impact on the model.
T4-2	Find the strategy parameter with the best return.

Subjective ratings were highly positive (average scores $>$ $4.0/5.0$), particularly for intuitive design, utility (Q1), meeting analytical needs (Q4), and enhancing model trust via explanations (Q3). QuantVisExplorer users also demonstrated significantly higher task correctness, especially for complex inter-asset relationship analysis (T2) and model prediction interpretation (T4), suggesting reduced cognitive load. Qualitative feedback reinforced these findings, with users praising the clarity of views like the Asset Relation Graph and the Interactive Backtesting Module. Constructive suggestions included performance optimization for very large datasets and real-time data support. Overall, the study indicates QuantVisExplorer effectively supports complex quantitative investment analyses, offering an intuitive, transparent, and efficient alternative to traditional tools, thus improving analytical efficiency and user confidence.

6.4 Discussion

Our case studies demonstrated the system's utility in enabling domain experts to uncover anomalous market trends, understand the driving factors behind model predictions, and interactively refine investment strategies through features like multi-faceted network exploration and integrated backtesting. The user study further corroborated these findings, indicating that QuantVisExplorer was perceived as intuitive and effective, significantly aiding users in comprehending complex financial data and model behaviors compared to traditional tools. Key strengths highlighted include the system's ability to integrate diverse information sources and provide clear visual pathways from market overviews to granular asset details and model explanations. While the results are promising, we acknowledge limitations such as the current prototype's performance with extremely large datasets and the absence of real-time data integration, which were noted in user feedback. Future work could focus on addressing these scalability challenges, enhancing the model's adaptability to unforeseen market shocks, and expanding the analytical capabilities to include more sophisticated financial instruments and automated insight recommendation features. Ultimately, QuantVisExplorer contributes to bridging the gap between complex quantitative models and the practical needs of investment professionals for transparent, explorable, and actionable insights.

7 Conclusion

In this work, we presented QuantVisExplorer, a novel visual analytics system that synergizes a multi-relational temporal graph network model with an interactive multi-view interface to support explainable quantitative investment. We demonstrated its capabilities in facilitating market trend analysis, asset relationship exploration, and model-driven strategy evaluation through comprehensive case studies and a positive user study. By providing intuitive visual access to complex model outputs and multi-faceted market data, QuantVisExplorer offers a significant step towards more transparent, trustworthy, and effective data-driven decision-making in the financial domain. We believe this approach not only empowers financial analysts but also contributes valuable design insights for future visual analytics systems targeting complex predictive modeling in other domains. Future directions include enhancing system scalability and exploring real-time data integration.

Acknowledgments

This work was supported by the National Science and Technology Major Project of China under Grant 2024ZD0526904, and National Natural Science Foundation of China under Grant 62272071.

References

[1] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008. 6

[2] T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327, 1986. 2

[3] L.-J. Cao and F. E. H. Tay. Support vector machine with adaptive parameters in financial time series forecasting. *IEEE Transactions on neural networks*, 14(6):1506–1518, 2003. 2

[4] S. Cao, W. Jiang, J. Wang, and B. Yang. From man vs. machine to man+ machine: The art and ai of stock analyses. *Journal of Financial Economics*, 160:103910, 2024. 2

[5] M. M. Carhart. On persistence in mutual fund performance. *The Journal of finance*, 52(1):57–82, 1997. 2

[6] R. Chiong, Z. Fan, Z. Hu, and S. Dhakal. A novel ensemble learning approach for stock market prediction based on sentiment analysis and the sliding window method. *IEEE Transactions on Computational Social Systems*, 10(5):2613–2623, 2022. 2

[7] Z. Deng, S. Chen, T. Schreck, D. Deng, T. Tang, M. Xu, D. Weng, and Y. Wu. Visualizing large-scale spatial time series with geochron. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):1194–1204, 2023. 3

[8] Q. Ding, S. Wu, H. Sun, J. Guo, and J. Guo. Hierarchical multi-scale gaussian transformer for stock movement prediction. In *IJCAI*, pp. 4640–4646, 2020. 2

[9] E. F. Fama and K. R. French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992. 2

[10] E. F. Fama and K. R. French. The capm is wanted, dead or alive. *The Journal of finance*, 51(5):1947–1958, 1996. 2

[11] R. Feng, S. Jiang, X. Liang, and M. Xia. Stgat: Spatial–temporal graph attention neural network for stock prediction. *Applied Sciences*, 15(8):4315, 2025. 2

[12] S. Gao, Y. Wang, and X. Yang. Stockformer: Learning hybrid trading machines with predictive coding. In *IJCAI*, pp. 4766–4774, 2023. 2, 3

[13] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42, 2018. 2

[14] H. Guo, M. Liu, B. Yang, Y. Sun, H. Qu, and L. Shi. Rankfirst: Visual analysis for factor investment by ranking stock timeseries. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 3

[15] H. A. He, J. Walny, S. Thoma, S. Carpendale, and W. Willett. Enthusiastic and grounded, avoidant and cautious: Understanding public receptivity to data and visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):1435–1445, 2023. 2

[16] Y.-L. Hsu, Y.-C. Tsai, and C.-T. Li. Fingat: Financial graph attention networks for recommending top-k k profitable stocks. *IEEE transactions on knowledge and data engineering*, 35(1):469–481, 2021. 2

[17] Y. Huang, N. Hu, K. Li, N. Wang, and Z. Lin. Extracting financial events from raw texts via matrix chunking. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 7035–7044, 2024. 2

[18] Y. Ji, C. Perin, and M. A. Nacenta. The effect of visual aids on reading numeric data tables. *IEEE Transactions on Visualization and Computer Graphics*, 2024. 2

[19] S. Ko, I. Cho, S. Afzal, C. Yau, J. Chae, A. Malik, K. Beck, Y. Jang, W. Ribarsky, and D. S. Ebert. A survey on visual analysis approaches for financial data. In *Computer Graphics Forum*, vol. 35, pp. 599–617. Wiley Online Library, 2016. 2

[20] K. J. Koa, Y. Ma, R. Ng, and T.-S. Chua. Diffusion variational autoencoder for tackling stochasticity in multi-step regression stock price prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 1087–1096, 2023. 2

[21] S. Lai, W. Luan, and J. Tao. Explore your network in minutes: a rapid prototyping toolkit for understanding neural networks with visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):683–693, 2023. 3

[22] 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, pp. 162–170, 2024. 2

[23] H. Lin, D. Zhou, W. Liu, and J. Bian. Learning multiple stock trading patterns with temporal routing adaptor and optimal transport. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 1017–1026, 2021. 6

[24] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017. 3

[25] M. P. Neto and F. V. Paulovich. Multivariate data explanation by jumping emerging patterns visualization. *IEEE Transactions on Visualization and Computer Graphics*, 30(2):1549–1563, 2022. 3

[26] E. Newburger and N. Elmqvist. Visualization according to statisticians: an interview study on the role of visualization for inferential statistics. *IEEE transactions on visualization and computer graphics*, 30(1):230–239, 2023. 2

[27] V. Prasad, R. J. van Sloun, S. van den Elzen, A. Vilanova, and N. Pezzotti. The transform-and-perform framework: Explainable deep learning beyond classification. *IEEE Transactions on Visualization and Computer Graphics*, 30(2):1502–1515, 2022. 3

[28] M. T. Ribeiro, S. Singh, and C. Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016. 3

[29] R. Roll and S. A. Ross. An empirical investigation of the arbitrage pricing theory. *The journal of finance*, 35(5):1073–1103, 1980. 2

[30] M. Schaefer, F. Wanner, R. Kahl, L. Zhang, T. Schreck, and D. A. Keim. A novel explorative visualization tool for financial time series data analysis. 2011. 2

[31] R. H. Shumway, D. S. Stoffer, R. H. Shumway, and D. S. Stoffer. Arima models. *Time series analysis and its applications: with R examples*, pp. 75–163, 2017. 2

[32] G. Song, T. Zhao, S. Wang, H. Wang, and X. Li. Stock ranking prediction using a graph aggregation network based on stock price and stock relationship information. *Information Sciences*, 643:119236, 2023. 2

[33] E. Sorenson and R. Brath. Financial visualization case study: Correlating financial timeseries and discrete events to support investment decisions. In *2013 17th International Conference on Information Visualisation*, pp. 232–238. IEEE, 2013. 2

[34] H. Subramonyam and J. Hullman. Are we closing the loop yet? gaps in the generalizability of vis4ml research. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):672–682, 2023. 3

[35] H. Tang, S. Wei, Z. Zhou, Z. C. Qian, and Y. V. Chen. Treeroses: outlier-centric monitoring and analysis of periodic time series data. *Journal of Visualization*, 22:1005–1019, 2019. 3

[36] 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. 2

[37] C. Walchshofer, V. Dhanoa, M. Streit, and M. Meyer. Transitioning to a commercial dashboarding system: Socio-technical observations and opportunities. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):381–391, 2023. 2

[38] H. Xia, H. Ao, L. Li, Y. Liu, S. Liu, G. Ye, and H. Chai. Ci-sthpan: pre-trained attention network for stock selection with channel-independent spatio-temporal hypergraph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, pp. 9187–9195, 2024. 2

[39] S. Xiang, D. Cheng, C. Shang, Y. Zhang, and Y. Liang. Temporal and heterogeneous graph neural network for financial time series prediction. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pp. 3584–3593, 2022. 6

[40] Y. Ye, R. Huang, and W. Zeng. Visatlas: An image-based exploration and query system for large visualization collections via neural im-

age embedding. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 3

[41] Z. You, P. Zhang, J. Zheng, and J. Cartlidge. Multi-relational graph diffusion neural network with parallel retention for stock trends classification. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6545–6549. IEEE, 2024. 2

[42] X. Yue, J. Bai, Q. Liu, Y. Tang, A. Puri, K. Li, and H. Qu. sportfolio: Stratified visual analysis of stock portfolios. *IEEE transactions on visualization and computer graphics*, 26(1):601–610, 2019. 3

[43] X. Yue, Q. Gu, D. Wang, H. Qu, and Y. Wang. iquant: interactive quantitative investment using sparse regression factors. In *Computer Graphics Forum*, vol. 40, pp. 189–200. Wiley Online Library, 2021. 3

[44] X. Yue, X. Shu, X. Zhu, X. Du, Z. Yu, D. Papadopoulos, and S. Liu. Bitextract: Interactive visualization for extracting bitcoin exchange intelligence. *IEEE transactions on visualization and computer graphics*, 25(1):162–171, 2018. 3

[45] Z. Zang, S. Cheng, H. Xia, L. Li, Y. Sun, Y. Xu, L. Shang, B. Sun, and S. Z. Li. Dmt-ev: an explainable deep network for dimension reduction. *IEEE transactions on visualization and computer graphics*, 30(3):1710–1727, 2022. 3

[46] H. Zhang, Z. Shi, Y. Hu, W. Ding, E. E. Kuruoğlu, and X.-P. Zhang. Optimizing trading strategies in quantitative markets using multi-agent reinforcement learning. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 136–140. IEEE, 2024. 2

[47] H. Ziegler, M. Jenny, T. Gruse, and D. A. Keim. Visual market sector analysis for financial time series data. In *2010 IEEE Symposium on Visual Analytics Science and Technology*, pp. 83–90. IEEE, 2010. 2