

Hybrid AI Models for Intelligent Investment Advisory System in the Saudi Stock Market

Student Name: Meshal Manfi Alkhaldi

Academic ID: 251001992

Midocean University

Program: Master of Science in Artificial Intelligence Updated

Thesis_(1)_AI900a

Research Proposal

Supervisor: Dr. Khaled Eskaf

Third Trimester 2025

Abstract

The Saudi Stock Market (Tadawul) has become one of the most dynamic and strategically important financial markets in the Middle East under the Kingdom's Vision 2030 program. This increased importance, along with growing participation from retail and institutional investors, has created a strong demand for intelligent, data-driven tools that can support investment decision-making. Traditional econometric models such as ARIMA and GARCH are limited in their ability to model the nonlinear, nonstationary, and regime-shifting nature of financial time series. In contrast, recent advances in Artificial Intelligence (AI) including deep learning, Transformer architectures, Graph Neural Networks (GNNs), and Reinforcement Learning (RL) have shown promising results in stock prediction and algorithmic trading. However, the majority of existing work focuses on global indices and on prediction or fully automated trading, rather than on human-centered investment advisory systems tailored to specific markets.

This research proposes a hybrid AI investment advisory system designed specifically for the Saudi Stock Market. The proposed framework integrates (a) deep learning and Transformer-based models to forecast short-term price movements and trend probabilities, (b) GNNs to model inter-stock and inter-sector relationships within Tadawul, and (c) RL agents to optimize trading decisions in a simulated environment. These components feed into an advisory engine that generates human-interpretable recommendations, including buy/sell/hold suggestions, confidence scores, and basic risk indicators. The system will be evaluated using both statistical metrics and financial performance indicators on historical Tadawul data. The expected contribution of this

study is a comprehensive, hybrid AI framework that advances academic research on AI in finance and offers practical value for investors and financial institutions in Saudi Arabia.

Hybrid AI Models for Intelligent Investment Advisory System in the Saudi Stock Market

1. Introduction

The financial markets of the twenty-first century are characterized by rapid information dissemination, high degrees of interconnectedness, and frequent structural changes. The Saudi Stock Market (Tadawul) has emerged as a central component of the Kingdom of Saudi Arabia's Vision 2030, which aims to diversify the national economy, attract foreign capital, and develop a sophisticated financial sector. Tadawul lists companies across multiple sectors, including energy, petrochemicals, banking, real estate, telecommunications, and consumer services. Each of these sectors is influenced by local macroeconomic policies, global commodity prices, geopolitical developments, and technological innovation. For investors, this environment presents substantial opportunities, but it also introduces significant uncertainty and complexity.

Many investors in Tadawul still rely heavily on traditional tools such as moving averages, oscillators, and discretionary chart analysis. While such techniques may be effective for experienced traders, they are limited in their ability to process large volumes of data or capture complex nonlinear patterns. Classical time-series models, including autoregressive and GARCH-type specifications, impose strong assumptions about linearity and stationarity that are often violated in real-world financial data. These limitations have motivated researchers to explore machine learning and AI-based techniques for forecasting stock prices, volatility, and risk. Deep learning models, particularly Long Short-Term Memory (LSTM) networks, have demonstrated superior performance compared with many traditional approaches when applied to financial time

series (Fischer & Krauss, 2018). LSTMs are able to capture temporal dependencies and long-range patterns in sequential data. However, LSTMs and other recurrent neural networks still face challenges when modeling very long horizons, multivariate series with many covariates, or highly irregular temporal dynamics. More recently, Transformer-based architectures have been proposed to address these challenges by using self-attention mechanisms to learn dependencies across time steps without relying on recurrent structures. Temporal Fusion Transformers (TFT), for example, integrate attention with variable selection and offer interpretable multi-horizon forecasts (Lim et al., 2021).

In addition to sequence models, Graph Neural Networks (GNNs) have emerged as powerful tools for learning from relational data. Financial markets naturally form graph structures, where nodes can represent stocks and edges can represent relationships such as sector membership, historical correlation, or supply chain linkages. By leveraging graph structures, GNN-based methods can model how shocks propagate through sectors and capture cross-sectional dependencies that are invisible to models that treat each stock independently (Chen et al., 2022; Hu et al., 2023). At the same time, Reinforcement Learning has gained traction in finance as a paradigm for learning trading strategies through trial-and-error interaction with a simulated environment (Zhang & Zhao, 2022; Liu et al., 2023). RL agents observe market states, take trading actions, and receive rewards based on realized profits and losses, enabling them to discover policies that balance risk and return.

Despite these advances, there remains a gap between state-of-the-art AI research and the tools available to investors in the Saudi market. Most academic work either focuses on

global markets or is framed as prediction or fully automated trading, with limited attention to human-centered advisory systems that output interpretable recommendations. This proposal seeks to bridge that gap by developing a hybrid AI investment advisory system tailored to the structure and characteristics of Tadawul.

2. Problem Statement and Research Gap

AI applications in finance can be broadly grouped into prediction-oriented systems and trading-oriented systems. Prediction-oriented systems use supervised learning to forecast prices, returns, volatility, or trends. Trading-oriented systems often use rule-based strategies or Reinforcement Learning agents to select buy, sell, or hold actions over time. Although both categories have demonstrated empirical success, they exhibit several limitations when considered from the viewpoint of investors in the Saudi Stock Market.

First, most models are developed and evaluated using data from mature markets such as the United States or Europe. These markets differ from Tadawul in terms of sector composition, liquidity, volatility profiles, trading hours, and regulatory environments. As a result, it is not clear whether methods that work well in these markets will transfer directly to the Saudi context without modification. There is relatively little academic literature focusing explicitly on AI models for Tadawul.

Second, a large proportion of existing models treat each stock as an independent time series, ignoring the web of relationships between firms and sectors. In practice, shocks to large-cap or sector-leading stocks can have cascading effects on related firms, and macroeconomic or policy events can influence entire segments of the market

simultaneously. GNNs offer a principled way to encode these relationships, but their use in modeling the Saudi market remains limited.

Third, many AI models output numerical forecasts or implicitly encode trading policies, but they do not present their outputs in a human-friendly format. Individual investors and financial advisors often require clear, interpretable recommendations—such as buy, sell, or hold—accompanied by confidence estimates and simple risk indicators. Systems that lack interpretability may be difficult to adopt in practice, especially when they function as black boxes.

In summary, there is a need for a hybrid AI framework that: (a) is tailored to the Saudi Stock Market, (b) incorporates relational information through graph-based modeling, (c) learns trading strategies using RL, and (d) exposes its outputs through an interpretable advisory layer. This proposal addresses that research gap.

3. Research Objectives and Research Questions

The overarching objective of this research is to design, implement, and evaluate a hybrid AI investment advisory system for the Saudi Stock Market. The proposed system will integrate deep learning, Transformer-based models, GNNs, and RL into a unified framework that supports data-driven, human-centered investment decisions.

The specific objectives of the study are as follows:

1. To develop deep learning and Transformer-based models for forecasting short-term price movements and trend probabilities for selected stocks listed on Tadawul.
2. To construct a graph representation of Tadawul and train GNN models that capture inter-stock and inter-sector relationships.
3. To implement a Reinforcement Learning trading agent that uses hybrid model outputs

as part of its state representation and learns policies to optimize risk-adjusted returns.

4. To design an advisory engine that converts internal model outputs into buy/sell/hold recommendations, confidence scores, and basic risk labels suitable for human investors.
5. To evaluate the overall system using statistical and financial performance metrics on historical data from the Saudi Stock Market.

The research will be guided by the following research questions:

1. Do hybrid AI models that combine deep learning, Transformers, and GNNs achieve better predictive performance for Saudi stocks than single-model baselines?
2. Can an RL agent, informed by hybrid predictive signals, improve the quality and robustness of trading decisions in Tadawul?
3. How can the outputs of multiple AI components be integrated into simple, interpretable investment recommendations for human users?
4. What levels of accuracy, profitability, and drawdown control can the proposed advisory system achieve in out-of-sample backtests?

4. Significance of the Study

This study is significant from both academic and practical perspectives.

Academically, it contributes to the growing body of literature on hybrid AI methods in financial markets by combining deep learning, Transformer architectures, GNNs, and RL in a single framework. While prior studies have investigated these technologies individually or in limited combinations—for example, LSTM for forecasting, GNN for relational learning, or RL for trading—few have attempted to integrate them into a unified advisory system. Furthermore, there is limited work targeting emerging markets

such as Saudi Arabia, where structural and behavioral characteristics differ from those of well-studied Western markets.

From a practical standpoint, the proposed system addresses the needs of investors and institutions in Saudi Arabia who seek AI-driven tools that do more than simply print numerical forecasts. By producing interpretable recommendations, confidence scores, and basic risk signals, the system can help investors understand and contextualize its outputs. In addition, the hybrid architecture may support financial institutions in developing more advanced robo-advisory services or decision support dashboards tailored to Tadawul.

5. Proposed Methodology

The research adopts a quantitative, experimental approach. The core methodology consists of four stages: data collection and preprocessing, development of hybrid predictive models, design of an RL trading agent, and implementation of the advisory engine followed by evaluation.

5.1 Data Collection and Preprocessing

Historical daily OHLCV data for a selected set of Saudi stocks and indices will be collected from Tadawul or licensed financial data providers. Sector classifications and, if feasible, simple sentiment indicators derived from financial news or social media will be included. Data preprocessing will involve handling missing values, filtering out extreme outliers, and computing technical indicators such as moving averages, Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and volatility measures. All variables will be normalized or standardized to facilitate training of neural networks.

5.2 Hybrid Predictive Models

The predictive component of the system will include several supervised learning models. Deep learning models (e.g., LSTM, GRU) will be trained on sequences of engineered features to predict short-term direction or trend probabilities. Transformer-based architectures, including Temporal Fusion Transformers, will be employed to capture long-range dependencies and complex covariate interactions. In parallel, a graph representation of the market will be created, where nodes represent individual stocks and edges encode sector membership or high historical correlation. GNN models will be trained on this graph to produce market-aware embeddings for each stock. These embeddings will then be fused with sequence-based model outputs to form rich feature representations.

5.3 Reinforcement Learning Trading Agent

An RL agent, such as PPO or A2C from the Stable-Baselines3 library, will be trained in a custom trading environment that simulates trading over historical data. At each time step, the agent will observe a state vector that includes recent price history, technical indicators, hybrid model predictions, and GNN embeddings. The agent's actions will be limited to buy, sell, or hold, perhaps with position sizing constraints to simplify the problem. The reward function will be based on realized returns, adjusted for transaction costs and penalizing excessive volatility or drawdowns. Training will proceed over multiple episodes, with hyperparameters tuned using validation sets.

5.4 Advisory Engine and Evaluation

The advisory engine will sit on top of the predictive and RL layers. Rather than automatically executing trades, the system will translate its internal signals into human-

readable outputs: recommended action (buy/sell/hold), confidence score derived from predictive probabilities or ensemble agreement, and a simple risk label based on historical volatility or expected drawdown. Evaluation will be conducted using out-of-sample test periods and will employ both statistical metrics (accuracy, F1-score, RMSE) and financial metrics (cumulative return, Sharpe ratio, maximum drawdown, win rate). Sensitivity analyses will examine how performance changes under different hyperparameter settings and market regimes.

6. Ethical Considerations and Limitations

The deployment of AI in financial decision-making raises important ethical considerations. First, it is essential that the proposed system not be presented as a guaranteed source of profit. Users must be clearly informed that all investments involve risk and that past performance does not guarantee future results. Second, transparency and interpretability should be prioritized. Although deep neural networks and RL agents are inherently complex, the advisory layer can incorporate explainability tools, such as feature importance scores or simple textual explanations, to help users understand why a given recommendation is made.

From a methodological standpoint, care must be taken to avoid common data biases, such as survivorship bias (excluding delisted stocks) or look-ahead bias (inadvertently using future information). The study will use strictly out-of-sample test sets and robust validation schemes to mitigate overfitting. Nevertheless, certain limitations remain. The models are trained on historical data and may not perform well during unprecedented events or structural breaks. Sentiment data for the Saudi market may be sparse or noisy. RL training can be computationally expensive and sensitive to hyperparameter choices.

Finally, the advisory system is intended as a decision support tool, not as a substitute for professional financial advice or the investor's own judgment.

8. References

- Chen, Y., Xu, Y., & Zhou, Y. (2022). Graph neural networks for stock movement prediction with relational modeling. *IEEE Transactions on Knowledge and Data Engineering*, 34(11), 5185–5198.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.
- Hu, S., Li, Z., & Wang, Y. (2023). Multi-cluster graph neural networks for stock price forecasting. *IEEE Access*, 11, 44567–44580.
- Lim, B., Zohren, S., & Roberts, S. (2021). Temporal fusion transformers for interpretable multi-horizon forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.
- Liu, Q., Wang, S., & Zhang, J. (2023). A deep reinforcement learning-based decision support system for automated stock trading. *Expert Systems with Applications*, 217, 119574.
- Nguyen, T., & Yoon, J. (2021). Financial news sentiment and deep neural networks for stock price forecasting. *Expert Systems with Applications*, 170, 114456.
- Wang, X., Zhou, L., & Li, P. (2022). A hybrid GARCH–LSTM model for stock market volatility forecasting. *Physica A: Statistical Mechanics and Its Applications*, 602, 127714.

Zeraatkar, Z., Pourbakhsh, S., & Wu, X. (2023). A multi-modal approach using a hybrid vision transformer and temporal fusion transformer model for stock price movement classification. *IEEE Access*, 11, 121040–121058.

Zhang, L., & Zhao, Q. (2022). Deep reinforcement learning for financial trading: A practical application. *Neural Computing and Applications*, 34(7), 5373–5391.