



Project Proposal

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

Prepared by:

Name: [Meshal Manfi Alkhalidi](#)

Student ID: [251001992](#)

Supervisor Name: [Dr. Khaled Eskaf](#)

Submitted in fulfillment of the requirements for the degree of Master of Science in Artificial
Intelligence

College/Department: [Collage of Informatics](#)

University: [Midocean University](#)

Year: 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 initiative. Increased participation from retail and institutional investors has created strong demand for intelligent, data-driven tools that support investment decision-making. Traditional econometric models such as ARIMA and GARCH are limited in modeling the nonlinear, nonstationary, and regime-shifting nature of financial time series. Recent advances in Artificial Intelligence (AI), including deep learning, Transformer architectures, Graph Neural Networks (GNNs), and Reinforcement Learning (RL), have demonstrated promising results in financial forecasting and trading. However, most existing work focuses on global markets or fully automated trading systems, with limited emphasis on human-centered investment advisory tools tailored to specific regional markets.

This project proposes a hybrid AI-based investment advisory system designed specifically for Tadawul. The framework integrates deep learning and Transformer-based models for short-term price movement forecasting, GNNs for modeling inter-stock and inter-sector relationships, and RL agents for optimizing trading decisions within a simulated environment. Rather than executing trades automatically, the system functions as a decision support tool that generates interpretable buy/sell/hold recommendations, confidence scores, and basic risk indicators. The system will be evaluated using statistical metrics and financial performance indicators on historical Tadawul data.

1. Project Overview

The Saudi Stock Market (Tadawul) has emerged as a strategically vital component of the Kingdom of Saudi Arabia's Vision 2030, reflecting national efforts to diversify the economy, attract foreign investment, and modernize the financial sector. Tadawul lists companies across multiple sectors—such as energy, petrochemicals, banking, real estate, telecommunications, and consumer services—each influenced by local macroeconomic policies, global commodity prices, and sector-wide shocks.

Despite this complexity, many investors continue to rely on traditional technical indicators or discretionary chart analysis. While such methods may provide useful signals, they are limited in their ability to process large volumes of data or capture nonlinear, nonstationary patterns and

regime shifts. Classical time-series models, including ARIMA and GARCH, impose assumptions about stationarity and linearity that are frequently violated in real-world market behavior, especially in emerging markets.

Recent AI approaches provide a path forward. Deep learning and Transformer-based architectures can capture complex temporal dependencies and multi-variable interactions, while GNNs can model relational information (e.g., sector membership or correlation structure) that is often ignored when each stock is treated independently. In addition, Reinforcement Learning offers a decision-making framework that can optimize actions under uncertainty by learning from simulated market interaction.

However, a key gap remains: many AI systems are designed either for pure prediction tasks or for fully automated trading, with limited focus on producing interpretable, human-centered investment recommendations in the Saudi market context. This project addresses that gap by proposing a hybrid AI advisory system for Tadawul that emphasizes actionable and explainable outputs (buy/sell/hold, confidence scores, and risk indicators) to support investors and institutions.

2. Problem Statement and Research Gap

AI applications in finance are typically categorized into prediction-oriented systems and trading-oriented systems. Prediction systems forecast prices, returns, or volatility using supervised learning, while trading systems learn policies that map market states to actions. Although both approaches demonstrate empirical success, they present limitations for investors in the Saudi Stock Market.

First, a substantial portion of the literature evaluates models using data from mature markets (e.g., the US and Europe). These markets differ from Tadawul in liquidity, volatility structure, sector composition, trading behavior, and regulatory context. As a result, findings from such settings may not transfer reliably to the Saudi market without adaptation.

Second, many forecasting models treat each stock as an independent time series, ignoring relational structures between firms and sectors. In reality, shocks propagate through sector networks and correlated groups, which motivates market-aware modeling using GNNs. Yet GNN applications targeting Tadawul remain limited.

Third, many systems output numerical forecasts or black-box policies, which are not easily interpretable by human investors. In practice, investors require clear recommendations accompanied by confidence estimates and basic risk indicators. Therefore, there is a need for a market-specific hybrid AI framework that combines temporal modeling, relational modeling, and sequential decision-making, while exposing interpretable advisory outputs.

3. Project Goal, Objectives, and Research Questions

Project Goal:

The primary goal of this project is to design, implement, and evaluate a hybrid AI-based investment advisory system tailored to the Saudi Stock Market (Tadawul), supporting human investors with interpretable, risk-aware recommendations (buy/sell/hold) rather than automated execution.

Specific Objectives:

- Develop deep learning and Transformer-based models to forecast short-term price movements and trend probabilities for selected Tadawul stocks.
- Construct a Tadawul market graph and train GNN models to capture inter-stock and inter-sector relationships.
- Implement a Reinforcement Learning trading agent that uses hybrid predictive signals and learns policies to optimize risk-adjusted returns in a simulated environment.
- Design an advisory engine that converts internal signals into buy/sell/hold recommendations, confidence scores, and basic risk indicators suitable for human investors.
- Evaluate the overall system using statistical metrics and financial performance indicators through strict out-of-sample backtesting.

Research Questions:

- RQ1: Do hybrid AI models (deep learning + Transformers + GNNs) improve predictive performance for Saudi stocks compared to single-model baselines?
- RQ2: Can an RL agent informed by hybrid predictive signals improve the robustness and quality of trading decisions in Tadawul?
- RQ3: How can outputs from multiple AI components be integrated into simple, interpretable investment recommendations for human users?

- RQ4: What levels of accuracy, profitability, and drawdown control can the proposed advisory system achieve in out-of-sample backtests?

4. Significance of the Study

Academically, this study contributes to hybrid AI research in finance by integrating deep learning, Transformer architectures, GNNs, and RL in a unified advisory framework. While prior work often examines these techniques individually, fewer studies combine them end-to-end for decision support, and even fewer target emerging markets such as Saudi Arabia.

Practically, the proposed system addresses the need for AI-driven tools that provide interpretable recommendations (rather than opaque forecasts or fully automated execution). By producing buy/sell/hold suggestions with confidence and basic risk indicators, the system can support investors and financial institutions in developing decision-support dashboards and robo-advisory services tailored to Tadawul.

5. Approach (Methodology)

5.1 Data Collection

Historical daily OHLCV data for a selected set of Tadawul-listed stocks and indices will be collected from Tadawul or licensed financial data providers. Sector classifications will be incorporated to support graph construction. Where feasible, auxiliary variables such as market indices and simplified sentiment indicators derived from financial news may be included. Data will be used strictly in historical form to prevent information leakage.

5.2 Data Processing

Preprocessing will include handling missing values, correcting data inconsistencies, filtering extreme outliers, and generating engineered features. Technical indicators such as moving averages, RSI, MACD, and volatility measures will be computed to enrich the feature space. Features will be normalized or standardized, and the dataset will be split chronologically into training, validation, and test periods to avoid look-ahead bias.

5.3 Hybrid Predictive Models

The predictive component will include deep learning models (LSTM/GRU) and Transformer-based architectures (e.g., Temporal Fusion Transformers) to capture long-range dependencies

and covariate interactions. In parallel, a market graph will be created where nodes represent individual stocks and edges encode sector membership and/or high historical correlation. GNN models will be trained to generate relational embeddings, which will be fused with sequence-based signals to create rich market-aware representations.

5.4 Reinforcement Learning Trading Agent

An RL agent (e.g., PPO or A2C) will be trained in a custom trading environment simulating historical trading. The state will include recent price history, technical indicators, hybrid model predictions, and GNN embeddings. The action space will be limited to buy/sell/hold (with optional position sizing constraints) to preserve interpretability. The reward will reflect returns net of transaction costs and include penalties for excessive volatility and drawdowns to encourage risk-aware policies.

5.5 Advisory Engine and Evaluation

The advisory layer will translate model and RL signals into human-readable outputs: recommended action (buy/sell/hold), confidence scores (from predictive probabilities/ensemble agreement), and a simple risk label (e.g., volatility-based). Evaluation will use statistical metrics (accuracy, F1-score, RMSE/MAE) and financial metrics (cumulative return, Sharpe ratio, maximum drawdown, win rate) under strict out-of-sample backtesting and rolling-window evaluation.

6. Project Steps and Timeline

The following timeline outlines the planned phases of the project over a 10-week period:

Week	Phase	Description
1	Data Collection & Setup	Collect Tadawul data, sector labels, and indices; set up development environment; initial data audit.
2	Data Cleaning & Preprocessing	Handle missing values/outliers; compute indicators; normalize features; create datasets with documentation.
3	Exploratory Data Analysis	Analyze patterns, correlations, sector behavior, and market regimes; finalize modeling plan.
4	Predictive Model Development	Implement LSTM/GRU and Transformer models; establish baselines and validation procedure.
5	Graph Construction & GNN Training	Build market graph (sector/correlation); train GNN for relational embeddings.
6	Hybrid Integration	Fuse temporal predictions with graph embeddings; evaluate hybrid vs single-model baselines.
7	RL Environment & Training	Implement trading environment; train PPO/A2C; tune reward design and hyperparameters.
8	Advisory Engine Development	Generate buy/sell/hold, confidence, and risk indicators; design user-friendly outputs.
9	Comprehensive Evaluation	Out-of-sample backtesting; compute statistical and financial metrics; robustness and sensitivity analysis.
10	Final Report & Presentation	Finalize write-up, figures, and tables; prepare presentation for defense/viva.

7. Expected Results

- Improved predictive performance from hybrid models relative to single-model baselines.
- Market-aware representations from GNN embeddings that capture inter-stock and inter-sector effects.
- An RL-based decision module that improves risk-adjusted performance compared to naive baselines.
- A human-centered advisory layer producing interpretable buy/sell/hold recommendations with confidence and risk indicators.
- Empirical insights on deploying hybrid AI decision-support systems in an emerging market context (Saudi Arabia).

8. Resources Needed

To successfully complete this project, the following resources will be essential across different stages of development, training, and evaluation:

Tools and Software

- Python Programming Language – for model development, data preprocessing, and evaluation.
- Jupyter Notebook – for interactive code execution and experimentation.
- Scikit-learn, Pandas, NumPy, Matplotlib – for data manipulation and visualization.
- PyTorch (and/or TensorFlow) – for deep learning and Transformer model implementation.
- PyTorch Geometric (or DGL) – for Graph Neural Network modeling.
- Stable-Baselines3 – for Reinforcement Learning (e.g., PPO, A2C).
- Google Colab / Local IDE – for development and prototyping.

Dataset

- Historical Tadawul OHLCV data – for selected stocks and market indices.
- Sector classification data – to build inter-stock and inter-sector relationships.
- Optional auxiliary datasets – indices (e.g., TASI) and simplified sentiment indicators (if feasible).

Computing Resources

- Local workstation or laptop – for initial experiments and testing.
- Optional access to cloud services (e.g., Google Colab Pro or equivalent) – for faster training of deep learning and RL models.

9. Literature Review

The following studies represent foundational literature that informs the development of this project and provides methodological guidance and benchmarking insights:

#	Reference	Title	Name of Researchers	Summary	Methodologies	Key Findings / Notes
1	Fischer & Krauss (2018)	Deep learning with LSTM networks for financial market predictions	T. Fischer; C. Krauss	Evaluated LSTM models for stock movement prediction and showed advantages in capturing nonlinear temporal dependencies.	LSTM; Deep Learning	LSTM outperformed traditional ML baselines in directional prediction (reported in study).
2	Lim et al. (2021)	Temporal Fusion Transformers for interpretable multi-horizon forecasting	B. Lim; S. Zohren; S. Roberts	Introduced TFT combining attention and variable selection to improve multi-horizon forecasting with interpretability.	Transformer; Attention; Variable Selection	Strong forecasting performance with interpretability mechanisms (feature importance/attention).
3	Chen et al. (2022)	Graph neural networks for stock movement prediction with relational modeling	Y. Chen; Y. Xu; Y. Zhou	Modeled stock relationships using graph structures to leverage cross-sectional dependencies.	Graph Neural Networks	Relational modeling improved prediction versus non-graph baselines in experiments.
4	Hu et al. (2023)	Multi-cluster graph neural networks for stock price forecasting	S. Hu; Z. Li; Y. Wang	Proposed clustering-aware GNN approach to better capture sector/group dynamics.	Multi-cluster GNN	Reported improved forecasting performance over standard GNN baselines.
5	Zhang & Zhao (2022)	Deep reinforcement learning for financial trading: a practical application	L. Zhang; Q. Zhao	Demonstrated practical deep RL trading approaches and considerations for reward design and risk control.	Deep RL (e.g., PPO/DQN)	RL can outperform rule-based strategies when trained and evaluated carefully.

Research Gap Summary:

- Most studies focus on mature markets, limiting direct transferability to Tadawul.
- Many approaches treat stocks independently, underutilizing relational information.
- Practical adoption is hindered by limited interpretability in prediction-only or automated trading systems.

This project addresses these gaps by integrating temporal, relational, and decision-making AI components into an interpretable advisory system tailored to the Saudi market.

10. Limitations and Risks

- Market regime shifts and structural breaks may reduce generalization to future conditions.
- Limited availability or noise in sentiment data for the Saudi market.
- Overfitting risk due to model complexity and multiple components.
- RL sensitivity to hyperparameters, reward design, and transaction cost modeling.

Mitigation strategies include strict chronological splitting, rolling-window backtesting, regularization and early stopping, sensitivity analysis across regimes, and conservative reward design with transaction costs and drawdown penalties.

11. Ethical Considerations

The proposed system is intended as a decision support tool and must not be presented as a guarantee of profit. Users should be clearly informed that investments involve risk and that historical performance does not ensure future outcomes. Transparency and interpretability will be prioritized through the advisory layer and, where applicable, explainability techniques. The study will mitigate common financial ML pitfalls such as survivorship bias and look-ahead bias by using strict out-of-sample protocols and reproducible evaluation.

12. References

1. 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.
2. 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.
3. Hu, S., Li, Z., & Wang, Y. (2023). Multi-cluster graph neural networks for stock price forecasting. *IEEE Access*, 11, 44567–44580.
4. Lim, B., Zohren, S., & Roberts, S. (2021). Temporal fusion transformers for interpretable multi-horizon forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.
5. 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.
6. Nguyen, T., & Yoon, J. (2021). Financial news sentiment and deep neural networks for stock price forecasting. *Expert Systems with Applications*, 170, 114456.
7. 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.
8. 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.
9. Zhang, L., & Zhao, Q. (2022). Deep reinforcement learning for financial trading: A practical application. *Neural Computing and Applications*, 34(7), 5373–5391.