

مشعل الخالدي

Meshal Manfi Alkhaldi.docx

 Midocean University

Document Details

Submission ID

trn:oid::2945:377785060

Submission Date

Apr 16, 2026, 10:55 PM GMT+2

Download Date

Apr 19, 2026, 6:55 PM GMT+2

File Name

Meshal Manfi Alkhaldi.docx

File Size

69.9 KB

28 Pages

4,858 Words

30,184 Characters

95% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Detection Groups



46 AI-generated only 94%

Likely AI-generated text from a large-language model.



1 AI-generated text that was AI-paraphrased 1%

Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.





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

Master Thesis

Prepared By

Name: Meshal Manfi Alkhaldi

Academic ID: 251001992

Submitted in partial fulfillment of the requirements for the degree of

Master of Science in Artificial Intelligence

College of Informatics

Midocean University

United Arab Emirates

Supervisor: Dr. Khaled Eskaf

Submission Date: April 2026

Acknowledgements

I would like to express my sincere gratitude to my mother for her prayers and unwavering support. I am also deeply thankful to my uncle for his continuous encouragement and motivation to pursue academic and professional excellence. My sincere appreciation goes to Aqeel for his constant pride in my achievements, which has been a strong source of motivation. I also extend my thanks to his brothers for their support and encouragement. Finally, I would like to express my heartfelt gratitude to my wife and children for their patience, understanding, and support throughout the completion of this thesis.

Abstract

This thesis presents a hybrid Artificial Intelligence framework for an investment advisory system tailored to the Saudi Stock Market (Tadawul). The study is motivated by the practical difficulty faced by investors in interpreting large volumes of market data, sector interactions, and changing price behaviour using conventional tools alone. Instead of depending on a single model family, the proposed framework combines sequence learning models, Transformer-based temporal forecasting, Graph Neural Networks (GNNs), and Reinforcement Learning (RL) in one decision-support pipeline.

The forecasting layer is designed to estimate short-term market direction from historical price and volume data enriched with technical indicators. A graph-based layer is then used to model structural links among listed stocks, such as sector membership and correlation patterns, so that cross-stock dependencies can be captured more explicitly. These predictive signals are integrated into a custom RL trading environment in which an agent learns buy, sell, and hold policies under market constraints and transaction costs. The final system does not aim to execute trades autonomously; instead, it translates internal model outputs into investor-facing recommendations accompanied by confidence and risk information.

The main contribution of the thesis lies in the integration of multiple AI paradigms into a single advisory architecture designed for the characteristics of Tadawul rather than for mature Western markets. The work also emphasizes interpretability by focusing on actionable recommendations instead of raw model scores alone. The framework is evaluated conceptually through predictive and financial performance criteria, including classification quality, cumulative return, Sharpe ratio, and maximum drawdown. The study concludes that a hybrid approach offers a stronger

basis for intelligent advisory support in the Saudi market than isolated forecasting or trading models used independently.

Keywords: Saudi Stock Market, Tadawul, Deep Learning, Temporal Fusion Transformer, Graph Neural Networks, Reinforcement Learning, Investment Advisory System

Table of Contents

Contents

Acknowledgements	2
Abstract.....	3
Table of Contents	5
List of Abbreviations	8
Chapter 1: Introduction.....	9
1.1 Background.....	9
1.2 Problem Definition	10
1.3 Research Objectives	11
1.4 Research Questions	11
1.5 Significance of the Study.....	12
1.6 Thesis Contribution	12
1.7 Thesis Organization.....	12
Chapter 2: Background and Related Work.....	13
2.1 Artificial Intelligence in Financial Markets	13
2.2 Limitations of Classical Financial Models	13
2.3 Deep Learning for Time-Series Forecasting	14
2.4 Transformer-Based Temporal Modeling.....	14

2.5 Graph Neural Networks in Financial Markets.....	15
2.6 Reinforcement Learning for Trading Decisions.....	15
2.7 Human-Centered Advisory Systems	16
2.8 Research Gap Summary	16
2.9 Comparative Summary.....	17
Chapter 3: Proposed Approach and Methodology	17
3.1 Research Design	17
3.2 Data Collection and Preprocessing.....	17
3.3 Predictive Modeling Layer	18
3.4 Fusion Strategy	19
3.5 Reinforcement Learning Environment	19
3.6 Advisory Engine	20
3.7 Evaluation Metrics.....	20
3.8 Ethical Considerations and Limitations.....	21
3.9 Proposed System Workflow	22
Chapter 4: Experimental Results and Discussion.....	22
4.1 Introduction	22
4.2 Experimental Setup	22
4.3 Baseline Models	23
4.4 Predictive Performance Analysis	23

4.5 Trading Performance Analysis	23
4.6 Interpretability and Advisory Quality	24
4.7 Discussion.....	24
4.8 Threats to Validity	24
4.9 Example Result Matrix.....	24
Chapter 5: Conclusion and Future Work.....	25
5.1 Conclusion.....	25
5.2 Main Contributions.....	25
5.3 Future Work.....	25
References	26
References	27

List of Abbreviations

- **AI:** Artificial Intelligence
- **Acc.:** Accuracy
- **DL:** Deep Learning
- **GNN:** Graph Neural Network
- **GRU:** Gated Recurrent Unit
- **LSTM:** Long Short-Term Memory
- **OHLCV:** Open, High, Low, Close, Volume
- **PPO:** Proximal Policy Optimization
- **RL:** Reinforcement Learning
- **RMSE:** Root Mean Square Error
- **TFT:** Temporal Fusion Transformer

Chapter 1: Introduction

1.1 Background

Financial markets operate in environments shaped by rapid information flow, dense interdependence, and persistent volatility. These characteristics make forecasting difficult because price series are nonlinear, noisy, and sensitive to structural change. Recent research has therefore moved away from relying on a single forecasting family and toward hybrid Artificial Intelligence (AI) pipelines that combine preprocessing, temporal learning, optimization, and decision logic to improve robustness (Jayanth et al., 2024; Jayanth & Manimaran, 2024; Sun et al., 2024).

Within this context, the Saudi Stock Market (Tadawul) provides an important setting for research. The market has grown in strategic relevance, includes sector-driven co-movement across banking, materials, energy, and telecommunications, and presents a realistic need for decision-support tools that can synthesize multiple signals into usable investment guidance. These characteristics make Tadawul an appropriate environment for studying both temporal forecasting and inter-stock dependence in one framework.

Recent advances in AI provide an opportunity to move beyond conventional forecasting tools. Deep learning models can learn temporal structure from historical market data, Transformer-based models can capture longer-range interactions across variables, Graph Neural Networks (GNNs) can represent stock-to-stock relationships, and Reinforcement Learning (RL) can transform predictive signals into sequential buy, sell, or hold decisions (Ansari et al., 2022; Ansari, 2024; Friday et al., 2025; Kabbani & Duman, 2022).

1.2 Problem Definition

A central problem in the literature is that many AI-based financial studies are developed using large international markets rather than the Saudi market. As a result, their assumptions about liquidity, sector behaviour, and market structure are not necessarily transferable to Tadawul. In addition, a substantial part of the literature focuses on either forecasting or automated trading, while fewer studies are framed explicitly as interpretable advisory systems for end users (Ansari et al., 2022; Kabbani & Duman, 2022).

A second limitation concerns representation of market structure. Many forecasting models treat each stock as an isolated time series even though stocks interact through sector membership, correlation structure, and broader market regimes. Graph-based research shows that relational modeling can improve forecasting by capturing interactions that independent models ignore (Ansari, 2024).

A third issue is practical usability. Existing systems often output numerical predictions or opaque policy signals that are difficult for investors to interpret. In real investment settings, users need recommendations that are concise, explainable, and supported by simple risk context. The challenge, therefore, is not only to improve predictive accuracy but also to convert model outputs into a human-centered advisory format that can support responsible decision-making (Ansari et al., 2022).

1.3 Research Objectives

- To build deep learning and Transformer-based forecasting models for selected Tadawul-listed stocks.
- To construct a graph-based representation of the market and generate relational embeddings using Graph Neural Networks.
- To design a Reinforcement Learning trading environment that uses hybrid predictive outputs as part of the state representation.
- To develop an advisory layer that converts internal model outputs into buy, sell, and hold recommendations with confidence and risk context.
- To evaluate the complete framework using both statistical and financial performance metrics.

1.4 Research Questions

- Do hybrid AI models outperform single-model baselines for forecasting Saudi stock movements?
- Can an RL agent informed by hybrid signals improve sequential trading decisions in Tadawul?
- How can outputs from multiple AI components be combined into simple and interpretable recommendations?
- What levels of profitability and drawdown control can the framework achieve in out-of-sample testing?

1.5 Significance of the Study

From an academic perspective, this study contributes to the literature on AI in finance by integrating sequence modeling, Transformer forecasting, relational learning, and sequential decision-making within one coherent framework. The novelty of the work does not lie in inventing each individual method from the ground up; rather, it lies in combining methods whose value has been demonstrated separately in prior studies and adapting them to the Saudi market context (Ansari et al., 2022; Ansari, 2024; Friday et al., 2025).

From a practical perspective, the thesis proposes a pathway toward an intelligent advisory system for investors and institutions operating in Tadawul. Instead of positioning AI as a fully autonomous replacement for human judgment, the study treats it as an analytical layer that enhances situational awareness, clarifies trade-offs, and communicates decisions in an accessible form.

1.6 Thesis Contribution

- A hybrid AI framework tailored to the Saudi stock market.
- A methodology for combining forecasting, relational learning, and decision policy learning.
- An interpretable advisory layer that communicates action, confidence, and risk.
- An evaluation structure that links machine learning performance to financial outcomes.

1.7 Thesis Organization

Chapter 1 introduces the research context, the problem being addressed, the study objectives, and the expected contribution. Chapter 2 reviews the theoretical background and the most relevant literature on AI-driven financial forecasting and trading. Chapter 3 presents the

proposed hybrid framework, including data preparation, predictive modeling, graph construction, RL design, and the advisory layer. Chapter 4 discusses the evaluation design and interprets the types of results that should be examined. Chapter 5 concludes the thesis and outlines directions for future development.

Chapter 2: Background and Related Work

2.1 Artificial Intelligence in Financial Markets

Artificial Intelligence has become a major direction in financial forecasting because stock data are nonlinear, noisy, and sensitive to changing market conditions. Recent studies increasingly combine statistical preprocessing, deep learning, and optimization instead of relying on a single model family. Hybrid forecasting models that merge smoothing, decomposition, or feature expansion with neural networks have reported stronger predictive performance than standalone baselines, which supports the broader shift toward hybrid AI in financial applications (Jayanth et al., 2024; Jayanth & Manimaran, 2024; Li et al., 2024; Sun et al., 2024; Swetha & Arya, 2025).

2.2 Limitations of Classical Financial Models

Classical models such as ARIMA and GARCH remain useful because they provide interpretable baselines and are still widely used in financial research. However, these models often struggle when data exhibit strong nonlinearity, regime changes, or heavy noise. Hybrid volatility work that combines GARCH-type modeling with LSTM shows that statistical structure remains valuable, but neural models are often needed to capture more complex behaviour in financial series (Koo & Kim, 2022). For this reason, classical methods are treated

in this thesis as reference points and complementary tools rather than as the final advisory engine.

2.3 Deep Learning for Time-Series Forecasting

Deep learning models, particularly LSTM, GRU, and BiLSTM variants, are widely used for stock forecasting because they can learn temporal patterns from rolling windows of price and volume data. Recent studies show that recurrent architectures become more effective when paired with stronger feature engineering or preprocessing. Examples include hybrid recurrent systems enhanced with exponential smoothing and Bayesian optimization, decomposition-based forecasting pipelines, and feature-expansion strategies that enrich model input before training (Jayanth et al., 2024; Jayanth & Manimaran, 2024; Li et al., 2024; Swetha & Arya, 2025). These findings are relevant to the present study because the first-term proposal already emphasized a forecasting layer based on engineered market signals rather than raw prices alone.

2.4 Transformer-Based Temporal Modeling

Transformer-based architectures have attracted attention because self-attention can model long-range temporal dependency and complex interactions among variables without relying purely on recurrence. The reviewed literature also shows that Temporal Fusion Transformer can be useful when the task involves multi-horizon movement classification and mixed feature sources. For example, Friday et al. (2025) combined a Vision Transformer and a Temporal Fusion Transformer to process historical and candlestick-derived information and reported strong classification performance across multiple indices. In this thesis, the Transformer component is retained as a forecasting stream, but it is used to strengthen the original proposal rather than to change the project into an image-centric study.

2.5 Graph Neural Networks in Financial Markets

Graph Neural Networks are especially relevant when stocks should not be treated as isolated time series. Ansari (2024) showed that clustering-based, multi-relation stock graphs can capture stock-to-stock interactions more effectively than simple independent forecasting, allowing a single graph model to forecast multiple stocks while exploiting shared structure. This idea fits the Saudi market context because sector-based co-movement is common in Tadawul, especially among banking, materials, energy, and telecommunications stocks. Accordingly, the GNN component in this thesis is justified as a relational layer that complements temporal forecasting instead of replacing it.

2.6 Reinforcement Learning for Trading Decisions

Reinforcement Learning reformulates trading as a sequential decision problem in which an agent observes market state, takes actions, and receives reward from the consequences of those actions. This is particularly relevant when the objective is not only to predict price movement but also to decide when to buy, sell, or hold under portfolio and transaction-cost constraints. Ansari et al. (2022) proposed a deep reinforcement learning-based decision-support system in which a forecasting network is concatenated with historical trends to improve the agent's observation state. Likewise, Kabbani and Duman (2022) formulated trading as a partially observed Markov decision process and showed that deep reinforcement learning can unify prediction and allocation into one decision framework. These studies support the inclusion of an RL decision layer in the present work while keeping the system as an advisory tool rather than a fully autonomous trader.

2.7 Human-Centered Advisory Systems

Although many published systems aim at higher forecasting accuracy or automated profit generation, fewer studies are designed around interpretable support for end users. The decision-support orientation of Ansari et al. (2022) is especially useful here because it shows that predictive signals can be used to support decisions without removing human oversight. This thesis follows that direction more closely than a fully automated execution model. The goal is therefore to generate advisory outputs such as buy, sell, or hold recommendations, together with confidence and basic risk context, in a form that can assist users in the Saudi stock market.

2.8 Research Gap Summary

The literature reviewed in this chapter points to four clear themes. First, hybrid deep learning models often outperform simpler baselines in noisy financial environments (Jayanth et al., 2024; Li et al., 2024; Sun et al., 2024). Second, Transformer-based models can add value when temporal dependency is long or multivariate (Friday et al., 2025). Third, graph-based methods improve representation when inter-stock relationships matter (Ansari, 2024). Fourth, reinforcement learning is useful when the problem is framed as sequential decision-making rather than pure prediction (Ansari et al., 2022; Kabbani & Duman, 2022). However, the surveyed studies are mostly developed on non-Saudi datasets and usually focus on only one or two of these dimensions at a time. The research gap addressed by this thesis is therefore the development of a Saudi-market-oriented hybrid advisory framework that remains close to the first-term proposal while integrating these ideas in a controlled and interpretable way.

2.9 Comparative Summary

Based on the reviewed studies, the most suitable direction for this thesis is not to replace the original proposal with a completely different architecture, but to refine it. Sequence models remain appropriate for short-term forecasting, Transformer models can strengthen temporal representation, GNNs capture market structure, and RL supports sequential buy, sell, and hold decisions (Ansari et al., 2022; Ansari, 2024; Friday et al., 2025). Their combination provides a coherent extension of the first-term proposal for Tadawul rather than an unnecessary departure from it.

Chapter 3: Proposed Approach and Methodology

3.1 Research Design

This study adopts a quantitative experimental design that directly extends the first-term proposal rather than redefining the project scope. The workflow begins with historical Tadawul data preparation, followed by feature engineering, development of a forecasting layer, construction of a simple relational market representation, training of a reinforcement learning decision module, and generation of final advisory recommendations. The design choice is consistent with prior literature showing value in hybrid forecasting pipelines and decision-support architectures, but it remains intentionally focused so that the thesis stays aligned with its original Saudi-market objective.

3.2 Data Collection and Preprocessing

The primary dataset consists of historical OHLCV data for selected Tadawul-listed stocks. To preserve consistency with the original proposal, the stock universe should include

representative companies from major Saudi sectors rather than attempting to model the entire market at once. Preprocessing includes handling missing observations, adjusting for data-alignment issues, and constructing rolling windows for sequential modeling. This approach is consistent with recent forecasting studies that rely on structured market histories as the base input for recurrent, Transformer, and hybrid models (Jayanth et al., 2024; Li et al., 2024; Swetha & Arya, 2025).

Feature engineering is deliberately kept practical and close to the first-term plan. The framework incorporates technical indicators such as moving averages, RSI, MACD, ATR, momentum, rolling volatility, and volume-derived measures. This choice is supported by recent literature showing that enriched technical feature sets improve model learning and that hybrid systems often benefit from preprocessing or expanded indicators before the forecasting stage (Moody et al., 2024; Swetha & Arya, 2025). Additional data such as sentiment may be considered in future work, but they are not made central to the present thesis in order to avoid drifting too far from the original project.

3.3 Predictive Modeling Layer

The predictive modeling layer contains two complementary streams. The first stream uses a recurrent model, such as LSTM or GRU, because these architectures remain strong baselines for financial sequence learning and are well supported in the reviewed literature. The second stream uses a Transformer-based temporal model to improve representation of longer-range dependencies and multivariate interactions. This design is supported by the general success of recurrent-hybrid forecasting models and by recent work using the Temporal Fusion Transformer for stock movement classification (Friday et al., 2025; Jayanth et al., 2024; Li et

al., 2024). In this thesis, however, the Transformer stream is used as a temporal enhancement to the original proposal, not as a shift toward a wholly different multimodal project.

To represent market structure without overcomplicating the implementation, a stock graph is constructed in which nodes represent selected Saudi stocks and edges reflect simple but meaningful relationships such as sector membership or sustained return correlation. A Graph Neural Network then learns relational embeddings from this graph. This choice is motivated by evidence that graph-based modeling can improve stock forecasting by capturing interactions that independent time-series models ignore (Ansari, 2024). The graph layer is therefore included as a supportive relational module that enriches the temporal forecast.

3.4 Fusion Strategy

The fusion stage combines outputs from the recurrent stream, the Transformer stream, and the graph-based representation into one hybrid feature space. A simple concatenation followed by a learned fusion layer is sufficient for the thesis because the goal is to show complementary value, not to introduce an unnecessarily complex fusion mechanism. This design remains faithful to the core logic of the first-term proposal: multiple AI components should contribute jointly to a stronger advisory signal for the Saudi market.

3.5 Reinforcement Learning Environment

A custom trading environment is used to simulate sequential decision-making over historical Tadawul data. At each time step, the agent observes recent market features together with the hybrid predictive signal and basic portfolio state information such as cash balance and current position. The action space is restricted to buy, sell, or hold in order to keep the advisory logic clear and aligned with the intended application. This setup is consistent with decision-support

and trading-automation studies that show how forecasting outputs can be integrated into the RL state representation (Ansari et al., 2022; Kabbani & Duman, 2022).

The reward function is designed to encourage profitable but realistic behaviour. It uses realized return while penalizing transaction costs, excessive turnover, and undesirable drawdowns. This is important because the literature on reinforcement learning for trading consistently shows that unrealistic reward design can produce unstable or misleading policies (Ansari et al., 2022; Kabbani & Duman, 2022). By including these penalties, the thesis keeps the RL component grounded as a decision-support mechanism rather than a purely theoretical trading exercise.

3.6 Advisory Engine

The advisory engine is the user-facing layer of the framework and is central to keeping the thesis practical. Its role is to transform internal model outputs into buy, sell, or hold suggestions together with a confidence score and a simple risk indication. When the recurrent model, Transformer stream, graph context, and RL policy are mutually supportive, recommendation confidence can increase; when signals conflict, the system can default to hold or lower-confidence advice. This design keeps the project close to its original goal of an intelligent investment advisory system for the Saudi stock market.

3.7 Evaluation Metrics

Evaluation is conducted using both predictive and financial criteria. On the predictive side, metrics such as accuracy, precision, recall, F1 score, RMSE, MAE, and MAPE may be used depending on whether the task is framed as movement classification or value forecasting. On the financial side, the analysis focuses on cumulative return, Sharpe ratio, maximum drawdown, win rate, and turnover. This two-level evaluation is consistent with the reviewed literature,

where forecasting studies report statistical error measures and trading studies report portfolio-oriented indicators (Ansari et al., 2022; Jayanth et al., 2024; Kabbani & Duman, 2022).

In addition to overall metrics, ablation analysis should be used to test whether each component adds meaningful value. Useful comparisons include recurrent-only forecasting, recurrent-plus-Transformer forecasting, recurrent-plus-graph forecasting, and the full hybrid framework with RL. This is important because the thesis is not merely claiming that more components are better; it is testing whether each addition contributes enough benefit to justify its inclusion in a Saudi-market advisory system.

3.8 Ethical Considerations and Limitations

The proposed framework is intended as a decision-support tool, not as a promise of guaranteed investment profit. Any practical deployment would therefore require clear communication about uncertainty, model limitations, and the continuing importance of human judgment. This is especially important in financial applications, where users may over-trust algorithmic outputs if they are presented without context (Ansari et al., 2022).

Methodologically, the study must guard against survivorship bias, look-ahead bias, and overfitting. It must also acknowledge that model performance can degrade during structural breaks or unusual market events. Additional limitations include the availability of high-quality Saudi market data, the sensitivity of RL training to hyperparameter choices, and the challenge of converting complex internal model behaviour into advisory outputs that remain both simple and informative.

3.9 Proposed System Workflow

- Collect and clean historical Tadawul market data.
- Generate technical indicators and market-structure features.
- Train sequence-based models and Transformer-based forecasting models.
- Construct a stock relationship graph and learn GNN embeddings.
- Fuse predictive signals into a hybrid market signal.
- Use the hybrid signal within a custom RL trading environment.
- Translate final outputs into buy/sell/hold recommendations with confidence and risk context.
- Evaluate predictive and financial performance on out-of-sample data.

Chapter 4: Experimental Results and Discussion

4.1 Introduction

This chapter explains how the proposed framework should be examined empirically and how the resulting evidence ought to be interpreted. Because the thesis combines forecasting, relational learning, and sequential decision-making, the results cannot be judged through a single metric. A useful discussion must therefore link model quality, trading behaviour, and advisory usefulness in one coherent analysis.

4.2 Experimental Setup

The experimental design is based on chronological train, validation, and test partitions in order to preserve the temporal nature of financial data. Models are first calibrated on historical

training windows, tuned on validation data, and then evaluated on unseen test periods. This sequence is important because random shuffling would create information leakage and lead to overly optimistic estimates of real-world performance.

4.3 Baseline Models

Appropriate baselines include a simple buy-and-hold strategy, a rule-based technical trading system, classical forecasting methods, and reduced versions of the proposed framework. These baselines establish whether the full hybrid architecture provides genuine value or merely reproduces what simpler methods can already achieve.

4.4 Predictive Performance Analysis

Predictive analysis should first compare the hybrid model against individual components. Improvement in directional classification, probability calibration, or forecast error would indicate that the combined architecture extracts richer information from the market state. However, the discussion should remain careful: a modest increase in predictive score is meaningful only if it is consistent across test periods and not concentrated in a narrow regime.

4.5 Trading Performance Analysis

The more practical question is whether stronger signals improve decision quality once they are embedded in the RL environment. The analysis should therefore examine whether the agent achieves more stable returns, lower drawdowns, and more disciplined position changes than baseline strategies. In a market-advisory setting, consistency and risk control may be as important as raw return (Ansari et al., 2022; Kabbani & Duman, 2022).

4.6 Interpretability and Advisory Quality

Interpretability remains a core evaluation dimension. The advisory output should be assessed in terms of clarity, internal consistency, and alignment with observed market conditions. Recommendations that are frequent but contradictory would reduce trust, whereas recommendations that communicate confidence and risk in a stable manner are more likely to be useful to investors.

4.7 Discussion

Several broader issues should be discussed when interpreting the findings. One is whether graph-based market structure adds value during sector-led movements or periods of broad co-movement. Another is whether the RL agent genuinely learns a robust decision policy or simply overfits to repeating historical patterns. A third is whether the advisory layer succeeds in simplifying the system without obscuring material uncertainty.

4.8 Threats to Validity

Threats to validity include assumptions about transaction costs, the selected stock universe, sample-period effects, and possible changes in market microstructure over time. The Saudi market may also exhibit regime shifts that affect sectors unevenly, which means conclusions drawn from one historical period should not be generalized without caution.

4.9 Example Result Matrix

The final version of the thesis should report completed result tables for both forecasting and financial evaluation. At a minimum, one table should compare predictive metrics across baseline and hybrid models, and another should summarize return, Sharpe ratio, drawdown, and

win rate for the advisory strategies. These tables should be accompanied by interpretation rather than presented as isolated numbers.

Chapter 5: Conclusion and Future Work

5.1 Conclusion

This thesis proposed a hybrid AI framework for an intelligent investment advisory system designed for the Saudi Stock Market. The framework brings together temporal forecasting, graph-based relational modeling, and Reinforcement Learning within one pipeline that produces human-readable recommendations rather than opaque signals alone. The study argues that such integration is well suited to Tadawul, where sector interaction, regime variation, and practical interpretability all matter for investment support.

5.2 Main Contributions

The main contribution of the work is methodological integration. Instead of treating forecasting, market-structure learning, and trading-policy design as separate tasks, the thesis connects them in a unified advisory architecture. A second contribution is contextual: the framework is tailored to the Saudi market rather than borrowed unchanged from better-studied foreign exchanges. A third contribution lies in the emphasis on explainable recommendations that retain human oversight.

5.3 Future Work

Future work can extend the framework in several directions. One natural extension is the inclusion of richer data sources such as intraday records, macroeconomic variables, corporate

disclosures, and Arabic financial-news sentiment. Additional work may also compare alternative graph-construction schemes, more advanced portfolio action spaces, and explainability techniques that make model reasoning easier to audit. Finally, a production-grade implementation could be tested as an interactive dashboard for investors or institutional analysts.

References

References

- Ansari, Y. (2024). Multi-Cluster Graph (MCG): A novel clustering-based multi-relation graph neural networks for stock price forecasting. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3476159>
- Ansari, Y., Yasmin, S., Naz, S., Zaffar, H., Ali, Z., Moon, J., & Rho, S. (2022). A deep reinforcement learning-based decision support system for automated stock market trading. *IEEE Access*, 10, 127469–127492. <https://doi.org/10.1109/ACCESS.2022.3226629>
- Friday, I. K., Pati, S. P., & Mishra, D. (2025). A multi-modal approach using a hybrid Vision Transformer and Temporal Fusion Transformer model for stock price movement classification. *IEEE Access*, 13. <https://doi.org/10.1109/ACCESS.2025.3589063>
- Jayanth, T., & Manimaran, A. (2024). Developing a novel hybrid model double exponential smoothing and dual attention encoder-decoder based bi-directional gated recurrent unit enhanced with Bayesian optimization to forecast stock price. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3435683>
- Jayanth, T., Manimaran, A., & Siva, G. (2024). Enhancing stock price forecasting with a hybrid SES-DA-BiLSTM-BO model: Superior accuracy in high-frequency financial data analysis. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3502175>
- Kabbani, T., & Duman, E. (2022). Deep reinforcement learning approach for trading automation in the stock market. *IEEE Access*, 10. <https://doi.org/10.1109/ACCESS.2022.3203697>

- Koo, E., & Kim, G. (2022). A hybrid prediction model integrating GARCH models with a distribution manipulation strategy based on LSTM networks for stock market volatility. *IEEE Access*, 10. <https://doi.org/10.1109/ACCESS.2022.3163723>
- Li, Y., Chen, L., Sun, C., Liu, G., Chen, C., & Zhang, Y. (2024). Accurate stock price forecasting based on deep learning and hierarchical frequency decomposition. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3384430>
- Moody, F., Jahangard Rafsanjani, A., Zarifzadeh, S., & Zare Chahooki, M. A. (2024). Fusion of technical indicators and sentiment analysis in a hybrid framework of deep learning models for stock price movement prediction. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3519556>
- Sun, Y., Mutalib, S., Omar, N., & Tian, L. (2024). A novel integrated approach for stock prediction based on modal decomposition technology and machine learning. *IEEE Access*, 12. <https://doi.org/10.1109/ACCESS.2024.3425727>
- Swetha, B., & Arya, K. (2025). Prediction of the stock market using a hybrid model based on feature expansion and LSTM-based algorithms. *IEEE Access*, 13. <https://doi.org/10.1109/ACCESS.2025.3631096>