An Intelligent System for Automated Portfolio Management using Graph Neural Networks and Semantic Bayesian Networks

Project Proposal

Insaaf Dhansay
dhnins001@myuct.ac.za
University of Cape Town
Cape Town, South Africa

Kialan Pillay
pllkia010@myuct.ac.za
University of Cape Town
Cape Town, South Africa

CCS CONCEPTS

• **Computing methodologies** → **Probabilistic reasoning**; **Neural networks**.

KEYWORDS

share evaluation, portfolio management, JSE, graph neural networks, ontologies, semantic bayesian networks

1 PROJECT DESCRIPTION

1.1 Introduction

Portfolio management is a complex decision-making process that aims to maximise return [27] through the proportional allocation of investable capital into identified financial securities. The initial share evaluation phase involves identifying securities with suitable characteristics for inclusion in a portfolio. Portfolio selection is the successive phase of portfolio management and determines the optimal allocation of the set of identified shares in a portfolio [6]. Financial markets are dynamic, volatile systems and portfolio management sub-tasks ordinarily require manual analysis by an investor with expert domain knowledge.

1.2 Current Limitations

The existing literature [15] [17] has developed solutions to automate portfolio management tasks using a range of Artificial Intelligence (AI) techniques; state-of-the-art has been achieved using deep neural network (DNN) models. However, these intelligent systems or architectures primarily incorporate black-box techniques that diminish the explainability of the system. By obscuring the rationale underpinning the transformation of inputs into outputs [3], these black-box systems do not provide adequate decision support to a decision-maker, since the results and machine decisions are not transparent to the user [8].

Semantic Bayesian networks, which are Bayesian networks (BN) enriched with knowledge represented in an ontology, are an explainable AI technique. In contrast to black-box approaches, explainable or glass-box approaches allow decision-makers to understand the rationale behind automated decisions through the explanation of outputs [18]. Semantic Bayesian networks have been incorporated into intelligent decision support systems (IDSS) to solve domain-specific problems. However, their application to share evaluation and portfolio management is not widely explored in the literature.

Graph neural networks (GNN) are a DNN model sub-class that capture non-temporal dependencies in multivariate time series data. Recent contributions [21] [14] using DNNs have illustrated that price forecasting can effectively function as a proxy for automated share evaluation. Similarly to semantic BNs, the literature is yet to investigate GNNs for price prediction and nor integration into a broader intelligent system to perform share evaluation decision support.

To overcome limitations on transparency and explainability, Drake [4] designed *INVEST*, a novel IDSS framework using ontologies and BNs for evaluating Johannesburg Stock Exchange (JSE)-listed shares. However, the framework was theoretically evaluated and no experimental results were obtained in the study. An IDSS implementation for share evaluation using semantic BNs is a notable absence in the state-of-the-art.

2 PROBLEM STATEMENT

Portfolio management is a decision-making process that ordinarily requires a significant level of expert knowledge on the part of the decision-maker. The literature [7] [14] [17] investigating automated share evaluation and portfolio selection using Artificial Intelligence has focused on black-box techniques that provide minimal decision support. A semantic Bayesian network-based IDSS has yet to be implemented for share evaluation, with the *INVEST* system the only notable theoretical contribution in this domain.

2.1 Motivation

This section discusses the motivation behind the aim and identified objectives of this project in the two research subfields.

2.1.1 Graph Neural Networks. Share price forecasting is an example of a non-structural GNN application. Non-structural applications involve domains in which the data has an implicit relational structure that cannot be graphically encoded. However, stock markets are dynamic systems that can exhibit intra-share correlations. A correlation matrix provides a method [12] for encoding potential non-temporal dependencies in multivariate financial time series data. However, the existing literature does not assess the applicability of GNNs for price prediction nor comprehensively evaluate the suitability of a correlation matrix for representing prior knowledge. Furthermore, the literature has yet to investigate the integration of GNNs into an IDSS for share evaluation. Completion of the outlined objectives will address these weaknesses in the literature and produce empirical results to baseline future research.

2.1.2 Semantic Bayesian Networks. Theoretical results alone are often insufficient to validate system performance and robustness; *INVEST* is a promising yet under-evaluated framework. The completion of the objectives will produce empirical results using the implemented system that can be used to validate the results presented by Drake. The literature has yet to propose a similar IDSS for share evaluation using semantic Bayesian networks, and hence no results exist to perform a comparative evaluation. The objectives will address this and support future research investigating automated share evaluation and decision support. Existing literature [19] [23] has identified the importance of risk in automating portfolio management tasks. Systematic risk has often been omitted in current solutions [19], often resulting in diminished returns in volatile markets. Completion of an additional objective will address this gap in the literature.

2.2 Aims & Objectives

The aim of this project is to implement the *INVEST* system and empirically evaluate its performance for automated JSE-listed share evaluation and decision support. Of equal importance, the implementation of the system will provide an experimental platform that will allow us to extend the system and obtain empirical results. We aim to extend the system by investigating: the application of GNNs for JSE-listed share price forecasting, the performance and explainability of different BN topologies, and the inclusion of GNNs in the intelligent system for share evaluation. The proposed experiments, discussed in Section 3, will evaluate GNNs and Semantic Bayesian networks to achieve a set of research objectives that are outlined in this proposal. The individual objectives addressed by this project are summarised and categorised below.

2.2.1 System.

(1) Design, implement and test the *INVEST* intelligent system framework in Python.

2.2.2 Graph Neural Networks.

(1) Evaluate the applicability of three state-of-the art GNN models for FTSE/JSE Top 40-share price forecasting.
(2) Assess the capability of a statistical correlation matrix to capture dependencies in the stock market and encode prior knowledge.
(3) Integrate a GNN predictive component into the *INVEST* IDSS and evaluate the share evaluation performance of the hybrid GNN+BN intelligent system.

2.2.3 Semantic Bayesian Networks.

(1) Evaluate the performance of the *INVEST* IDSS for share evaluation and decision support with JSE-listed shares using return metrics.
(2) Assess the explainability of the share evaluation decisions produced by the *INVEST* system implementation.
(3) Evaluate the share evaluation performance of an extended Bayesian network that includes systematic risk.
(4) Conduct an ablation study of the *INVEST* Bayesian network topologies both with and without simulated noisy data.

3 PROCEDURES AND METHODS

This section discusses the proposed methodology and experimental approaches to achieve the identified research objectives in each sub-field.

3.1 INVEST System

The *INVEST* system will be implemented using a modular design to support extensibility and rapid prototyping. System modules will have standard interfaces to minimise interdependencies and decouple development tasks. An agile software development methodology is selected to iteratively design, implement and evaluate the system. An iterative lifecycle is based on the successive refinement and enlargement of the system. This reduces risk through early feedback and system adaption over multiple iterations. In addition, we envisage certain architectural compromises to effectively map the proposed design to concrete software classes.

Python, a general-purpose programming language is the selected development platform for the system. Python is suitable for building modularised applications and has extensive support for third-party libraries. The `PyAgrum` library will be used to create and manage Bayesian networks, as it can perform efficient computations. The `PyTorch` machine learning library is used to implement the Graph neural network models. The system implementation will dually function as the experimental platform for this project and support empirical evaluation of the *INVEST* framework and the proposed extensions. The codebase will include auxiliary modules unrelated to the core share evaluation functionality of the system.

3.2 Graph Neural Networks

3.2.1 Models. We have selected Graph-WaveNet (GWN) [25], MTGNN [24] and Spectral-Temporal Graph Neural Network (StemGNN) [1] for testing in this project. These models have achieved state-of-the-art or near-parity performance on the evaluated datasets for both one-step-ahead and multi-step-ahead prediction tasks. The source code for these models is open-sourced and in the public domain. The source code will be cloned and each model re-implemented and adapted according to our requirements. Each model will be trained using the hyperparameters and procedures stated in the original contributions to facilitate an accurate assessment of applicability. Modifications to the experimental setup may be required that will diverge from the literature.

Although all models can sufficiently capture non-temporal dependencies in the data without the provision of prior information, we will use a statistical correlation matrix [12] to encode potential non-temporal dependencies in the data. The correlation matrix will be calculated using the share price data of FTSE/JSE Top 40 Share Index constituents in a preprocessing step. StemGNN is unable to accommodate any encoding of prior knowledge by design; it explicitly aims to learn correlations only from the data. The correlation matrix will be input into GWN and MTGNN models pre-training to assess its impact on learning and predictive performance.

3.2.2 Data. The proposed experiments will test the predictive performance of the GNN models using multivariate time series data of shares in the FTSE/JSE Top 40 Index. The Top 40 Index contains the 40 largest JSE-listed companies by market capitalisation. Market capitalisation is the current value of all outstanding shares.

3.2.3 Testing. To achieve our objectives, the following experiments are proposed to evaluate the performance and applicability of the selected GNN models. Each experiment will be run 10 times [24] and the average value of the metrics will be reported.

(1) This experiment will compare each model's performance on a one-step-ahead prediction task. This experiment will not incorporate the correlation matrix.
(2) This experiment will compare each model's performance on a multi-step-ahead prediction task with increasing horizons. This experiment will not incorporate the correlation matrix.
(3) This experiment will compare the performance of GWN and MTGNN with correlation matrix to StemGNN on a multi-step-ahead prediction task with increasing horizons.

3.3 Semantic Bayesian Networks

3.3.1 Ontology Design. To include systematic risk in the BN topology, the *INVEST* ontology `Factor` class requires an extension. The `Factor` class is constructed hierarchically with three sub-classes to represent sub-categories. The abstract Factor class top level contains four abstract classes relating to evaluation objectives. The second level represents Factor Type, while the lower level classes represent concrete observable factors. The `QualityFactor` abstract class (top-level) will be modified through the addition of systematic risk as a factor type (second-level) and the Beta coefficient as a concrete observable factor (lower-level). The Protégé-OWL tool will be used to modify the relevant ontology sub-classes. This tool produces an XML encoding of the ontology represented in OWL.

3.3.2 Networks. The BN construction is informed by the *INVEST* ontology, where factors of the ontology map to variables in the BN, and instances of these factors will map to states for the relevant factors. To effectively evaluate the *INVEST* system, we select the three networks proposed by Drake [4] for testing: Value Evaluation, Quality Evaluation and Investment Recommendation. The Investment Recommendation BN incorporates the output of the Quality and Value BNs to predict the final investment decision for a specific share. Using the modified ontology outlined in the previous section, the Quality BN topology will be extended by mapping the factor to a node. The aforementioned BNs will be evaluated in the experiments detailed in a successive section.

3.3.3 Data. The dataset used will consist of price and fundamental data of 36 shares listed on the JSE between 2012 - 2018: 17 shares from the General Industrials sector and 19 shares from the Consumer Services sector [4]. This dataset is selected to support accurate comparative evaluation and reproducibility of results.

To evaluate the robustness and stability of the system, we will introduce a component to simulate the inherent noise exhibited in complex systems such as the stock market. Random variation will be generated for each time series and used to test the *INVEST* system.

3.3.4 Testing. The following experiments are proposed to extend and evaluate the *INVEST* intelligent system framework:

(1) This experiment will evaluate the *INVEST* IDSS with no modifications to empirically reproduce Drake's [4] results. This experiment will replicate the experimental conditions in Drake's study.
(2) This experiment will conduct an ablation study by testing each Bayesian network (Value and Quality) in isolation to determine how it affects the performance of the system.
(3) This experiment will evaluate the predictive performance of each Bayesian network (Value and Quality) and the combined topology on simulated noisy share price data. This data determines the probability of the `FutureSharePerformance` BN node.
(4) This experiment will evaluate the *INVEST* IDSS over different holding periods [4]. The periods will include 12 weeks [5], two years and three years.
(5) This experiment will incorporate systematic risk as a quality factor into the ontology. The experiment will evaluate the computational efficiency and predictive performance of the Quality BN with the additional factor.

3.4 Unified Share Evaluation

This subsection outlines the methodology for testing the integration of Graph neural networks with the existing semantic BN-based intelligent system. The *INVEST* system extensions and experimental platform developed to run the experiments outlined in the previous sections are sufficient to conduct the proposed unified experiment.

3.4.1 Model Selection & Integration. Model selection will be performed manually; the best-performing GNN model identified in the price prediction experiments will be selected for inclusion in the hybrid intelligent system. To test the integration, the output of the GNN must be transformed into a format suitable for inclusion in the Bayesian network. The model output is continuous, whilst the BN requires each variable to consist of a set of discrete states. The continuous output of the model is converted into a discrete value using a piece-wise function, where $\hat{y_{i+1}}$ and y_i are the predicted future and current price respectively.

$$x = \begin{cases} 1 & \frac{\hat{y_{i+1}}}{y_i} \geq 1.01 \\ 0 & 0.99 < \frac{\hat{y_{i+1}}}{y_i} < 1.01 \\ -1 & \frac{\hat{y_{i+1}}}{y_i} \leq 0.99 \end{cases} \tag{1}$$

Note that whilst this procedure mimics the output of a movement prediction task [17], the GNN models only perform point prediction and the final output is computed using the above function. No modification of the model output layer is required. This discretised value will replace the existing input of the `FutureSharePerformance` BN node.

3.4.2 Data. The proposed experiment will use the same dataset tested in the semantics Bayesian network experiments. The dataset comprises 36 shares listed on the JSE in the General Industrials and Consumer Services sectors from 2012 - 2018. An identical dataset allows for an accurate comparison of share evaluation performance between the hybrid and standard system.

3.4.3 Testing. The following experiment is proposed to test the performance of the GNN+BN intelligent system for share evaluation and assess the inclusion of GNNs.

(1) Comparing the GNN+BN system's recommendations to those produced by the standard *INVEST* system, baselined against benchmark index performance.

3.5 Evaluation

This subsection presents the GNN model evaluation procedure and selected metrics for the GNN models and *INVEST* system evaluation. Testing and evaluation is a critical component of this project. Multiple evaluation metrics are selected to mitigate the effect of noisy metrics on the results and analysis. The evaluation metrics can be sub-divided into two classes: Error and Return metrics.

3.5.1 Graph Neural Network Model Evaluation. Following Li et al. [13], a 60:20:20 split of each dataset partition using a walk-forward evaluation procedure [10] will be used to train and validate the models. Walk-forward evaluation partitions the data into successive and overlapping train-validate-test sets. The selected procedure preserves the temporal dependencies in the data.

3.5.2 Error Metrics. Error metrics will be used to validate the out-of-sample performance of the GNN models for share price forecasting. The GWN, MTGNN and StemGNN models will be evaluated using Mean Absolute Percentage Error, Root Mean Square Error and Mean Absolute Error.

The expressions are given below, where $\hat{y}_i$ and y_i are the predicted and actual values.

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \left| \hat{y}_i - y_i \right| \tag{2}$$

$$MAPE = \frac{100}{n} \sum_{i=1}^{n} \left| \frac{\hat{y}_i - y_i}{y_i} \right| \tag{3}$$

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^{n} \left| \hat{y}_i - y_i \right|} \tag{4}$$

3.5.3 Return Metrics. Both Return and Risk-Adjusted Return metrics will be used to evaluate the performance of the semantic Bayesian networks for share evaluation. These metrics include Annual Return (AR), Compound Return (CR), Average Annual Return (AAR), Treynor Ratio (TR) and Sharpe Ratio (SR). This selection is informed by our requirements of reproducibility and to support robust evaluation of other BN topologies.

$$AR = \frac{Portfolio\ Value_{Year+1}}{Portfolio\ Value_{Year}} - 1 \tag{5}$$

$$CR = \left(\frac{Portfolio\ Value_{Year+N}}{Portfolio\ Value_{Year}}\right)^{\frac{1}{N}} - 1 \tag{6}$$

$$AAR = \frac{\frac{Portfolio\ Value_{Year+N}}{Portfolio\ Value_{Year}}}{N} \tag{7}$$

$$TR = \frac{r_p - r_f}{\beta_p} \tag{8}$$

$$SR = \frac{r_p - r_f}{\sigma_p} \tag{9}$$

The JSE All-Share Index, General Industrials (JGIND) and Consumer Services (JCSEV) are the selected benchmark indices to validate the share evaluation recommendations. Additionally, the explainability of the results produced by the implementation will be compared and evaluated with the explainability of the decisions produced by the theoretical framework presented by Drake.

4 ETHICAL, PROFESSIONAL AND LEGAL ISSUES

This project does not necessitate experiments involving human or animal subjects. Consequently, we do not anticipate any ethical or legal issues arising from this research. All software required to achieve the anticipated outcomes are publicly available for use. Relevant work referenced in this project and related outputs will be correctly cited to ensure copyright and intellectual property (IP) compliance. We do not envisage any legal or regulatory compliance issues. All IP arising from this project, including but not limited to: academic papers, software and supplemental code, is owned by the University of Cape Town (UCT). Copyright of the academic papers is vested in the authors, subject to certain rights of UCT.

5 RELATED WORK

This section introduces related literature that is useful in understanding the outlined methodology and motivation for the proposed experiments.

5.1 Share Evaluation & Portfolio Management

Portfolio management is a phased decision-making process that aims to achieve a set of investment objectives. These objectives principally revolve around risk minimisation [27] and value maximisation [20]. Share evaluation is the initial phase of this process. It involves identifying shares that have suitable risk-return characteristics [4] for inclusion in a portfolio of securities. The decision-maker has a reasonable expectation that the selected shares will generate excess returns. Excess returns (α) is a performance metric [9] defined as the ratio of portfolio return to a benchmark return.

5.2 Graph Neural Networks

5.2.1 Background. Data and the relationships between objects can be modelled graphically in almost every domain. Graph neural networks [16] (GNN) are a model class that extends supervised neural network methods to process graphically-encoded non-Euclidean data.

Mathematically, a graph is a pair $G = (V, E)$, where V the set of nodes (data objects) and E the set of edges (relationships between data). A feature vector X_v is associated with each node $v \in V$. The labels of a node v, edge $e = (v, u)$ or graph G are real-valued vectors. The adjacency matrix $A \in \mathbb{R}^{N \times N}$ is a mathematical representation of a graph G, with $A_{ij} > 0$ for $(v_i, v_j) \in E$ and $A_{ij} = 0$ for $(v_i, v_j) \notin E$.

GNN models are designed to perform two distinct predictive tasks: node classification and graph classification. Node classification refers to the model objective of learning a representation vector h_v of $v \in V$ and a function f, such that the predicted node label is $y_v = f(h_v)$ [26]. In node classification applications, f is conditional on v, and hence the classification or regression is dependant on the node properties [16]. Graph classification tasks instead learn a representation vector h_G and a function g to predict the graph label y_G [26]. The function g is independent of any node $v \in V$ and thus unconditional on the node-level properties of the graph.

5.2.2 Graph Structure Encoding. GNN applications can be categorised into structural and non-structural applications [28]. Structural applications involve data that exhibits an explicit relational structure. Non-structural applications involve data that has implicit or hidden relationships. Stock market data is an example of non-structured data in which modelling does not expose the relational schema. LSTM Relational Graph Convolutional Network [12] is an architecture that exploits intra-share statistical correlations to capture non-temporal dependencies and graphically encode this information in an adjacency matrix. Li et al. calculate this correlation matrix using historical price data under the assumption of share inter-dependency.

5.3 Bayesian Networks

BNs are graphical, decision modelling tools [2] that encodes a representation of probabilistic knowledge within a domain [11]. BNs provide a mechanism to deal with the inherent uncertainty in many domains, and therefore often integrated into ontology-driven models. BNs are suitable for building explainable systems [4] [22], addressing a key limitation identified in the literature: the use of black-box techniques that provide minimal decision support. BNs are modelled as a Directed Acyclic Graph (DAG), with causal links (arcs) between variables (nodes). Each non-root node has a conditional probability table (CPT) associated with it. This is quantified using Bayes's formula to obtain conditional probabilities.

5.4 INVEST System

5.4.1 Introduction. Traditional decision support systems integrated with AI techniques have resulted in intelligent decision support systems. Drake [4] proposed the *INVEST* IDSS to support medium-term realistic and flexible share evaluation decisions for investment professionals. The decisions implemented in the system consists of both factual and heuristic knowledge. In Drake's work, the author reasoned with expert knowledge with respect to share evaluation under the value investing approach. The *INVEST* system design follows a glass-box approach as opposed to black-box methods utilized in other systems within the financial domain, therefore incorporates explainability. The system is composed of two key AI components: an ontology and a BN.

5.4.2 Ontology. The *INVEST* ontology supplies investment professionals with a clear structure of useful information and articulates concepts and properties, which are required by the BN as evidence. Drake designed an ontology that consists of seven classes that are generic to share evaluation processes.

5.4.3 Bayesian Network. The BN supports investment professionals in their investment decisions and provides a mechanism to handle the uncertainty inherent to share evaluation decision making. Drake designed three decision networks to form an overall topology, which corresponds to the three sequential steps in the decision-making process: Value Evaluation, Quality Evaluation and Investment Recommendation. The Investment Recommendation Bayesian Network reflects the decision process performed by investors: evaluating whether the price is reasonable and the quality of the share, thus providing an investment recommendation for a specific share.

6 ANTICIPATED OUTCOMES

This section outlines the intended outcomes of the research project and presents the expected impact of our results. This section also identifies key indicators for evaluating project success.

6.1 Research

6.1.1 Graph Neural Networks. We expect the evaluated GNN models to produce similar levels of performance but that the StemGNN architecture will outperform for share price forecasting. We anticipate that the combination of the GNN predictive component with the semantic Bayesian network will improve the performance of the *INVEST* system for share evaluation.

6.1.2 Semantic Bayesian Networks. The results of the BN experiments are expected to be similar to those presented by Drake [4], thus providing evidence in support of the theoretical intelligent system framework for share evaluation and decision support. We expect that the network topology that incorporates systematic risk will produce increased share evaluation performance relative to the baseline results.

6.2 System

We expect the project to yield a robust, stable implementation of the proposed *INVEST* intelligent system using GNNs and BNs. We anticipate the software will be comprehensively tested and documented and include measurement functionality to support future automated share evaluation and portfolio management experiments.

6.3 Impact

6.3.1 Graph Neural Networks. The application of GNNs to stock market data is a notable gap in the literature. This research will contribute novel results that both evaluate the applicability of the model class to JSE-share price prediction and test the integration of GNNs into the *INVEST* system for automated share evaluation. The experiments will also evaluate the suitability of a correlation matrix for encoding non-temporal dependencies in the context of share evaluation. This research provides a baseline for further research into GNNs for portfolio management and other applications in the finance domain.

6.3.2 Semantic Bayesian Networks. The *INVEST* system is a theoretical IDSS framework for share evaluation. This research will address a significant weakness in the literature by empirically evaluating the performance of the proposed framework. The experimental results will form a baseline that supports the development of alternative network topologies and further investigation of semantic BNs for automated share evaluation and decision support. This research will also assess the contributions of individual components in the overall IDSS framework. Furthermore, the research will evaluate the performance of a novel Bayesian network that incorporates systematic risk.

6.3.3 System. The implementation of the *INVEST* system is a state-of-the-art contribution that will provide a basis for multiple directions of future research. The software will contain the functionality to support further GNN and BN experiments on varied datasets. In addition, the software can be adapted and extended for portfolio selection to create an end-to-end framework for automated portfolio management. The software will be open-sourced and documented to promote reusability and extension by other researchers.

6.4 Key Success Factors

The success of this research project is appraised by the delivery of the outputs specified in a successive section and identified factors itemised below:

- Implementation and testing of the *INVEST* system, including software documentation.
- Successful implementation of measurement functionality to support accurate data collection and evaluation.
- Successful implementation of the GNN and BN extensions to support the individual experiments.
- Production of sufficient experimental results that can be used to draw valid conclusions.

7 PROJECT PLAN

This section presents the project schedule, milestones and core deliverables. This section also outlines the identified risks, required resources and allocation of work.

7.1 Risks

The identified risks associated with this project, including strategies for risk monitoring, mitigation and management, are detailed in the risk matrix found in Appendix B. In summary, the estimated occurrence probabilities are low, although the risk realisation may significantly impact the project.

7.2 Timeline

The project is initiated on the 17th of May 2021 and concludes on the 18th of October 2021 with the submission of the project webpage. A complete project schedule, including work allocation and the dates of key milestones, is outlined in the Gantt chart presented in Appendix A.

7.3 Resources Required

Compute requirements for software development and experimentation will require access to personal computing devices with suitable development environments. Access to the UCT Honours computing laboratory workstations will function as a contingency for hardware device malfunction or loss. There are no human resources or software licenses required to achieve the anticipated outcomes of the project. All requisite software is publicly accessible for use without restriction. For stock market data that is not in the public domain, we require access to a Bloomberg Terminal.

7.4 Deliverables

The primary deliverables of this project are an academic paper, documented software, and supplemental experimental code. The academic paper presents the results of our research and is the core output. This project also consists of secondary intermediate deliverables delivered throughout the duration of the project. A summarised list of notable project deliverables is itemised below.

- Two literature reviews
- Project proposal
- Software feasibility demonstration
- Draft academic paper
- Academic paper
- System source code with documentation
- Project poster
- Project website

7.5 Milestones

Several milestones correspond to the submission of the listed deliverables in section 7.4. Other notable milestones include the implementation of the core system, completion of experiments, and final demonstration. The complete set of project milestones is included in the Gantt chart in Appendix B.

7.6 Work Allocation

Insaaf Dhansay and Kialan Pillay will collaborate equally on shared deliverables: poster and webpage development and demonstration preparation. Both team members will jointly design and implement the base system and measurement functionality. Kialan Pillay will implement the core and extended GNN functionality and conduct the GNN experiments. Insaaf Dhansay will implement the Bayesian networks and conduct the semantic BN experiments. Both team members will jointly run the unified share evaluation experiment. The experimental results will be recorded by the assignee and combined into a cohesive research output.

REFERENCES

[1] Defu Cao, Yujing Wang, Juanyong Duan, Ce Zhang, Xia Zhu, Conguri Huang, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, et al. 2021. Spectral Temporal Graph Neural Network for Multivariate Time-series Forecasting. (2021). arXiv:arXiv:2103.07719

[2] Riza Demirer, Ronald R Mau, and Catherine Shenoy. 2006. Bayesian Networks: A Decision Tool to Improve Portfolio Risk Analysis. *Journal of Applied Finance* 16, 2 (2006), 106.

[3] Derek Doran, Sarah Schulz, and Tarek R Besold. 2017. What Does Explainable AI Really Mean? A New Conceptualization of Perspectives. (2017). arXiv:arXiv:1710.00794

[4] Rachel Drake. 2020. *A semantic Bayesian network for automated share evaluation on the JSE.* Master's thesis. University of Cape Town.

[5] Stefan Drue. 2019. *A comparison of a factor-based investment strategy and machine learning for predicting excess returns on the JSE.* Master's thesis. University of Cape Town.

[6] Alberto Fernández and Sergio Gómez. 2007. Portfolio selection using neural networks. *Computers & Operations Research* 34, 4 (2007), 1177–1191.

[7] Xingyu Fu, Jin-Hong Du, Yifeng Guo, MingWen Liu, Tao Dong, and Xiuwen Duan. 2018. A machine learning framework for stock selection. (2018). arXiv:arXiv:1806.01743

[8] Randy Goebel, Ajay Chander, Katharina Holzinger, Freddy Lecue, Zeynep Akata, Simone Stumpf, Peter Kieseberg, and Andreas Holzinger. 2018. Explainable AI: the New 42?. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 295–303.

[9] Larry R Gorman, Steven G Sapra, and Robert A Weigand. 2010. The cross-sectional dispersion of stock returns, alpha, and the information ratio. *The Journal of Investing* 19, 3 (2010), 113–127.

[10] Kouame Hermann Kouassi and Deshendran Moodley. 2020. An Analysis of Deep Neural Networks for Predicting Trends in Time Series Data. In *Artificial Intelligence Research*. Springer International Publishing, 119–140.

[11] Sun-Mi Lee and Patricia A Abbott. 2003. Bayesian networks for knowledge discovery in large datasets: basics for nurse researchers. *Journal of Biomedical Informatics* 36, 4-5 (2003), 389–399.

[12] Wei Li, Ruihan Bao, Keiko Harimoto, Deli Chen, Jingjing Xu, and Qi Su. 2020. Modeling the Stock Relation with Graph Network for Overnight Stock Movement Prediction. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*. International Joint Conferences on Artificial Intelligence Organization, 4541–4547.

[13] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations (ICLR '18)*.

[14] Yilin Ma, Ruizhu Han, and Weizhong Wang. 2021. Portfolio optimization with return prediction using deep learning and machine learning. *Expert Systems with Applications* 165 (2021), 113973.

[15] Antonio Rafael Sabino Parmezan, Vinicius MA Souza, and Gustavo EAPA Batista. 2019. Evaluation of statistical and machine learning models for time series prediction: Identifying the state-of-the-art and the best conditions for the use of each model. *Information sciences* 484 (2019), 302–337.

[16] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE Transactions on Neural Networks* 20, 1 (2008), 61–80.

[17] Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu. 2020. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied Soft Computing* 90 (2020), 106181.

[18] Kacper Sokol and Peter A Flach. 2018. Glass-Box: Explaining AI Decisions With Counterfactual Statements Through Conversation With a Voice-enabled Virtual Assistant. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 5868–5870.

[19] Shivaar Sooklal, Terence L van Zyl, and Andrew Paskaramoorthy. 2021. DRICORN-K: A Dynamic RIsk CORrelation-driven Non-parametric Algorithm for Online Portfolio Selection. In *Southern African Conference for Artificial Intelligence Research*. Springer, 183–196.

[20] Jaap Spronk and Winfried Hallerbach. 1997. Financial modelling: Where to go? With an illustration for portfolio management. *European Journal of Operational Research* 99, 1 (1997), 113–125.

[21] Van-Dai Ta, Chuan-Ming Liu, and Direselign Addis Tadesse. 2020. Portfolio optimization-based stock prediction using long-short term memory network in quantitative trading. *Applied Sciences* 10, 2 (2020), 437.

[22] Chiu-Che Tseng and Piotr J Gmytrasiewicz. 2002. Real time decision support system for portfolio management. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*. IEEE, 1348–1356.

[23] Yang Wang, Dong Wang, and Thomas Fang Zheng. 2018. Racorn-k: risk-aversion pattern matching-based portfolio selection. In *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 1816–1820.

[24] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020. Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 753–763.

[25] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 1907–1913.

[26] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

[27] Yunan Ye, Hengzhi Pei, Boxin Wang, Pin-Yu Chen, Yada Zhu, Ju Xiao, and Bo Li. 2020. Reinforcement-learning Based Portfolio Management with Augmented Asset Movement Prediction States. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1112–1119.

[28] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81.

A GANTT CHART

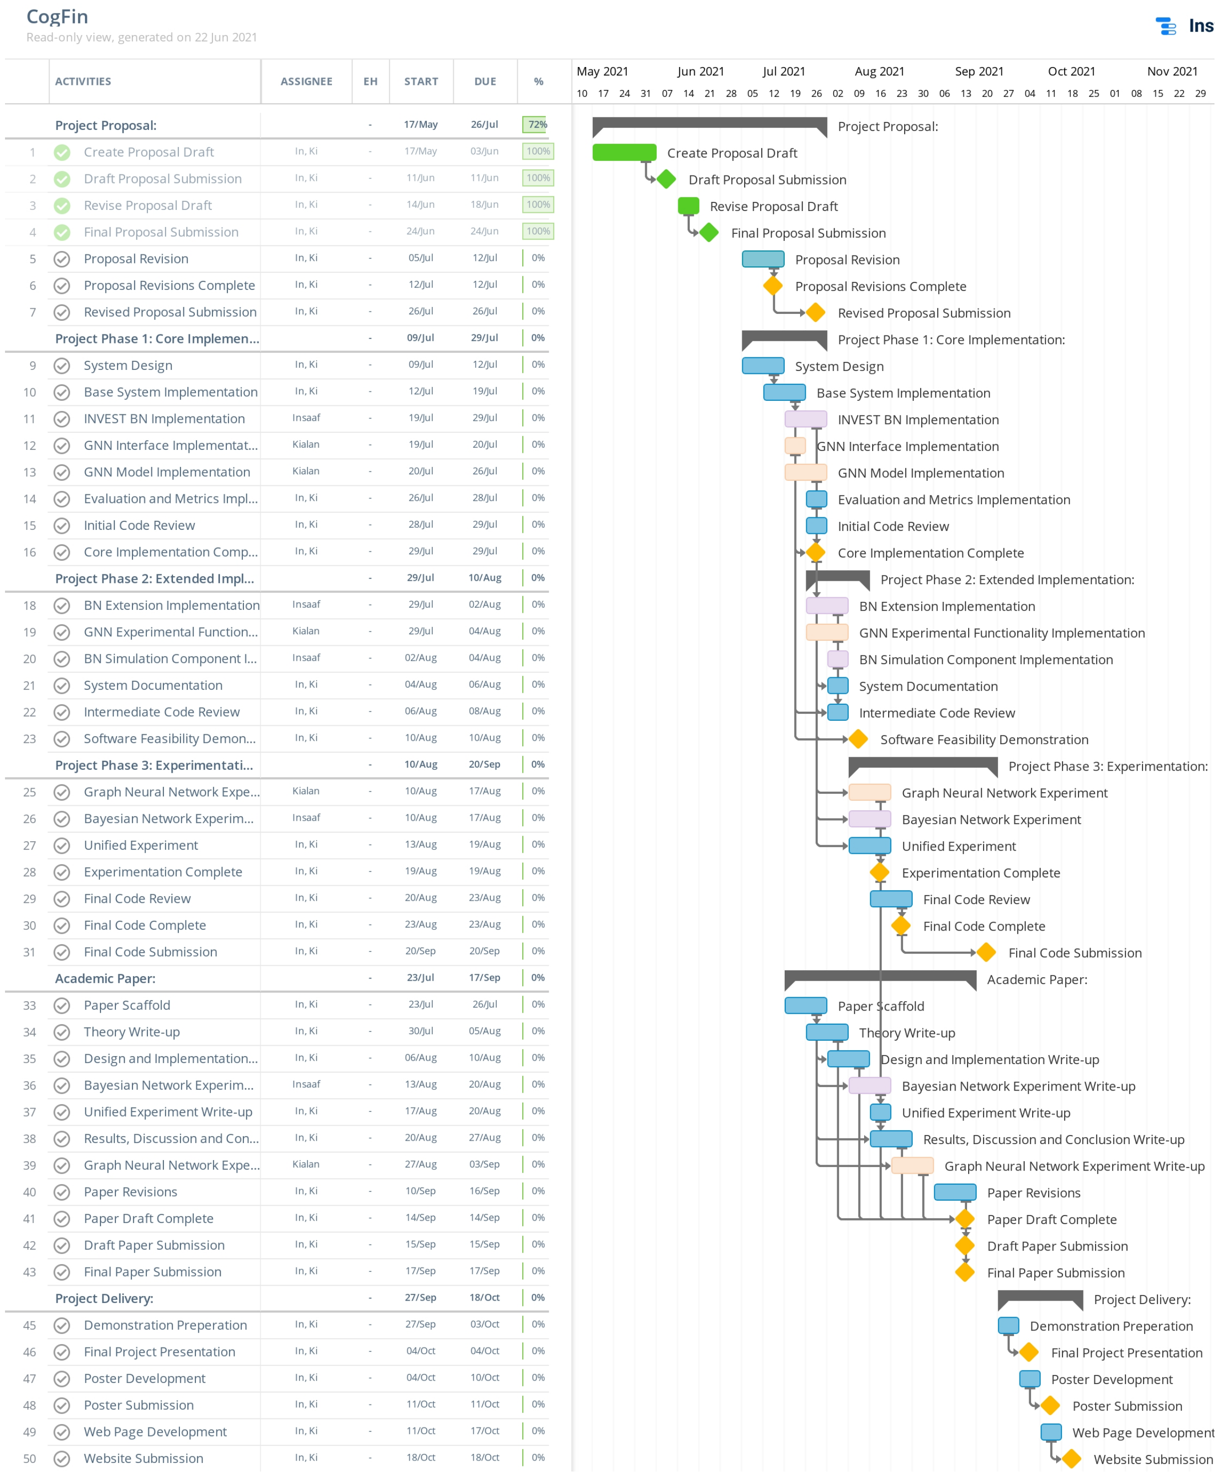

	ACTIVITIES	ASSIGNEE	EH	START	DUE	%
	Project Proposal:		-	**17/May**	**26/Jul**	**72%**
1	Create Proposal Draft	In, Ki	-	17/May	03/Jun	100%
2	Draft Proposal Submission	In, Ki	-	11/Jun	11/Jun	100%
3	Revise Proposal Draft	In, Ki	-	14/Jun	18/Jun	100%
4	Final Proposal Submission	In, Ki	-	24/Jun	24/Jun	100%
5	Proposal Revision	In, Ki	-	05/Jul	12/Jul	0%
6	Proposal Revisions Complete	In, Ki	-	12/Jul	12/Jul	0%
7	Revised Proposal Submission	In, Ki	-	26/Jul	26/Jul	0%
	Project Phase 1: Core Implemen...		-	**09/Jul**	**29/Jul**	**0%**
9	System Design	In, Ki	-	09/Jul	12/Jul	0%
10	Base System Implementation	In, Ki	-	12/Jul	19/Jul	0%
11	INVEST BN Implementation	Insaaf	-	19/Jul	29/Jul	0%
12	GNN Interface Implementat...	Kialan	-	19/Jul	20/Jul	0%
13	GNN Model Implementation	Kialan	-	20/Jul	26/Jul	0%
14	Evaluation and Metrics Impl...	In, Ki	-	26/Jul	28/Jul	0%
15	Initial Code Review	In, Ki	-	28/Jul	29/Jul	0%
16	Core Implementation Comp...	In, Ki	-	29/Jul	29/Jul	0%
	Project Phase 2: Extended Impl...		-	**29/Jul**	**10/Aug**	**0%**
18	BN Extension Implementation	Insaaf	-	29/Jul	02/Aug	0%
19	GNN Experimental Function...	Kialan	-	29/Jul	04/Aug	0%
20	BN Simulation Component I...	Insaaf	-	02/Aug	04/Aug	0%
21	System Documentation	In, Ki	-	04/Aug	06/Aug	0%
22	Intermediate Code Review	In, Ki	-	06/Aug	08/Aug	0%
23	Software Feasibility Demon...	In, Ki	-	10/Aug	10/Aug	0%
	Project Phase 3: Experimentati...		-	**10/Aug**	**20/Sep**	**0%**
25	Graph Neural Network Expe...	Kialan	-	10/Aug	17/Aug	0%
26	Bayesian Network Experim...	Insaaf	-	10/Aug	17/Aug	0%
27	Unified Experiment	In, Ki	-	13/Aug	19/Aug	0%
28	Experimentation Complete	In, Ki	-	19/Aug	19/Aug	0%
29	Final Code Review	In, Ki	-	20/Aug	23/Aug	0%
30	Final Code Complete	In, Ki	-	23/Aug	23/Aug	0%
31	Final Code Submission	In, Ki	-	20/Sep	20/Sep	0%
	Academic Paper:		-	**23/Jul**	**17/Sep**	**0%**
33	Paper Scaffold	In, Ki	-	23/Jul	26/Jul	0%
34	Theory Write-up	In, Ki	-	30/Jul	05/Aug	0%
35	Design and Implementation...	In, Ki	-	06/Aug	10/Aug	0%
36	Bayesian Network Experim...	Insaaf	-	13/Aug	20/Aug	0%
37	Unified Experiment Write-up	In, Ki	-	17/Aug	20/Aug	0%
38	Results, Discussion and Con...	In, Ki	-	20/Aug	27/Aug	0%
39	Graph Neural Network Expe...	Kialan	-	27/Aug	03/Sep	0%
40	Paper Revisions	In, Ki	-	10/Sep	16/Sep	0%
41	Paper Draft Complete	In, Ki	-	14/Sep	14/Sep	0%
42	Draft Paper Submission	In, Ki	-	15/Sep	15/Sep	0%
43	Final Paper Submission	In, Ki	-	17/Sep	17/Sep	0%
	Project Delivery:		-	**27/Sep**	**18/Oct**	**0%**
45	Demonstration Preperation	In, Ki	-	27/Sep	03/Oct	0%
46	Final Project Presentation	In, Ki	-	04/Oct	04/Oct	0%
47	Poster Development	In, Ki	-	04/Oct	10/Oct	0%
48	Poster Submission	In, Ki	-	11/Oct	11/Oct	0%
49	Web Page Development	In, Ki	-	11/Oct	17/Oct	0%
50	Website Submission	In, Ki	-	18/Oct	18/Oct	0%

Figure 1: Project Gantt Chart

B RISK MATRIX

Table 1: Risk Matrix

Risk	Probability	Impact	Consequence	Mitigation	Monitoring	Management
Loss or corruption of source code.	Low	Catastrophic	Considerable project delays and inability to meet all research objectives.	Regularly commit source code changes to VCS. Store backups of code on multiple devices.	Monitor the substance and frequency of source code commits by team members.	Retrieve the latest source code backups, commit code to the repository and reschedule tasks to offset lost/corrupt code.
Continuous change in research objectives or proposed experiments (Scope creep).	Medium	Medium	Project progress is dampened, wasted time on previously-performed tasks and potential project slippage. May impact the quality of the primary research outputs.	Document all research objectives, communicate scope with the project partner and verify scope with the supervisor.	Track any scope changes and verify if these are in accordance with the research objectives and methodology outlined in the proposal.	Prioritize core experiments and supporting features. Adjust the project schedule to determine what can be completed before the final deadline.
Project partner withdraws from the project.	Low	Medium	Certain research objectives may not be fulfilled due to incomplete experiments. Medium coupling exists between the research objectives and system implementation.	Assign tasks to achieve shared objectives such that dependencies are minimised; i.e. the other member has the knowledge and capacity to progress.	Regular check-ins with project partners, ensuring their interest in the project and commitment to the project is unchanged. Assess their social well-being and mental health.	Liaise with supervisor and re-evaluate research objectives and scope of the project. Exclude withdrawn member's contributions from project deliverables.
Lack of engagement with supervisor.	Low	Medium	Disapproval from the supervisor. May lead to an unsuitable and low-quality research project and scholarly outputs.	Regularly update supervisor on project progress through written communication. Schedule regular meetings for feedback.	Monitor the adherence and attendance of project team members and the supervisor to the agreed-upon meeting schedules.	Schedule a meeting with the supervisor to provide updates on project progress and gain feedback to progress with the project smoothly.
Inability to perform sufficient tests and validation.	Low	Critical	Incomplete project without experimental results. Research conclusions made will be significantly weakened.	Ensure adequate time has been allocated for experiment and system validation in the project plan and the outlined testing plan is complete.	Monitor project plan to ensure all milestones are met timeously and no project slippage occurs.	Prioritize completion of the core experiments and evaluation procedure and verify results with supervisor.