

AI-Powered Predictive Dashboard for Retail Decision-Making

MSc Research Project
Programme Name

Alfonso Pinto
Student ID: X23389028

School of Computing
National College of Ireland

Supervisor: Rejwanul Haque

National College of Ireland
Project Submission Sheet
School of Computing



Student Name:	Alfonso Pinto
Student ID:	X23389028
Programme:	AI For Business
Year:	2025
Module:	MSc Research Project
Supervisor:	Rejwanul Haque
Submission Due Date:	11/08/2025
Project Title:	AI-Powered Predictive Dashboard for Retail Decision-Making
Word Count:	5217
Page Count:	21

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Alfonso Pinto
Date:	13th September 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission , to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project , both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

AI-Powered Predictive Dashboard for Retail Decision-Making

Alfonso Pinto
X23389028

Abstract

The growing volume of sales, inventory, and supplier data presents an opportunity for artificial intelligence (AI) to transform retail decision-making. However, many small and medium-sized enterprises lack the tools and expertise to apply predictive analytics effectively. This research addresses that challenge by designing and evaluating an AI-powered dashboard that merges advanced forecasting models with a simple, role-adapted interface. Using the publicly available “Warehouse and Retail Sales” dataset (307,647 records), three models — Prophet, XGBoost, and Long Short-Term Memory (LSTM) — were implemented to forecast demand, highlight inventory risks, and suggest optimal promotion timing. An adapted CRISP-DM process, extended with stakeholder interviews and usability testing, ensured alignment between technical capability and business needs.

Findings show that the dashboard delivered accurate forecasts, with LSTM achieving the highest overall accuracy, XGBoost performing consistently well across diverse product categories, and Prophet offering the clearest interpretability for business users. Equally important, these forecasts were presented in a way that managers could easily understand and act upon. Stakeholders described the dashboard as intuitive, relevant to their daily tasks, and capable of supporting timely, informed decisions. By combining technical robustness with practical usability, the prototype bridges the gap between AI potential and everyday retail practice, with future work aimed at live deployment and the inclusion of broader business variables.

1 Introduction

1.1 Background

Retail companies produce large amounts of data every day from sales, inventory movements, supplier deliveries and customer interactions. However, many small and medium-sized retailers still depend on spreadsheets or static reports that do not allow real-time decision-making, thus reducing their ability to forecast demand, plan promotions, and optimise inventory effectively (DalleMule & Davenport, 2017). Business intelligence tools such as Power BI and Tableau can visualise data but they often lack integrated predictive capabilities unless complemented with technical skills or custom coding (Chen et al., 2012). This makes them difficult to use for companies without specialised data science teams. This research aims to develop a user-friendly AI-powered dashboard that enables retail teams to forecast product demand, manage stock and plan promotional strategies. The main contribution is a working prototype that integrates three forecasting models (Prophet, XGBoost, and LSTM) within an interactive Streamlit dashboard. The system uses real retail sales data and supports different types of users, including store managers, analysts and planners. A secondary contribution of this study is the creation of a simplified decision-support interface. This allows users to choose models, view forecasts, and receive insights without requiring advanced technical knowledge. The project also demonstrates how open-source tools can make predictive dashboards accessible to smaller businesses with tight budget constraints. The research uses the publicly available Warehouse and Retail Sales dataset from the U.S. Government Data Catalog, containing more than 300,000 records. The dataset considers variables such as time indicators (YEAR, MONTH), product details (ITEM CODE, DESCRIPTION, TYPE) and sales data (RETAIL SALES (target), WAREHOUSE SALES, RETAIL TRANSFERS). These features allow detailed forecasting, supply chain analysis, and comparison of sales trends over time.

1.2 Research Question

RQ: *How can AI-driven predictive dashboards improve strategic decision-making in retail businesses using publicly available sales data?*

1.3 Research Objectives

To answer the defined research question, the project sets out the following objectives:

1. Analyse the current state of AI dashboards and forecasting tools used in the retail sector.
2. Design a dashboard architecture able to support both technical and non-technical users.
3. Implement a working prototype using open-source technologies and multiple forecasting models.
4. Evaluate the dashboard's accuracy, using RMSE and MAE metrics, and assess its usability through simulated retail scenarios.

1.4 Dissertation Structure

The remainder of this dissertation is organised as follows:

- **Section 2** reviews the literature on AI dashboards, predictive analytics and retail decision-making.
- **Section 3** describes the research methodology.
- **Section 4** presents the design and implementation of the system.
- **Section 5** reports and analyses the results.
- **Section 6** provides conclusions of the study and recommendations for future work.

2 Related Work

2.1 Introduction and context

The rapid growth of data in the retail sector has made advanced analytics and artificial intelligence (AI) not only valuable but, in many cases, essential (Soofastaei, 2019). However, many companies still face challenges in turning large datasets into actionable insights. Predictive dashboards that combine machine learning (ML) models with clear visualisation tools are emerging as key enablers of data-driven strategies for managers and operational planners (Chen et al., 2025).

Despite the growing interest in AI-powered dashboards, their integration in real-world retail operations remains limited particularly among small and medium-sized enterprises (SMEs). This review examines research on predictive analytics, dashboard usability and the link between technical AI solutions and business decision-making. It evaluates the strengths and weaknesses of current approaches, identifies existing gaps and reflects on how these findings shape the development of this dissertation’s AI dashboard.

2.2 Predictive analytics in retail decision-making

Predictive analytics applies past data and statistical or ML models to estimate future outcomes (Shmueli & Koppius, 2011). In retail, it can forecast demand, identify sales patterns and optimal stock levels.

Waller and Fawcett (2013) emphasise the transformative impact of predictive analytics on supply chains, especially on improved forecast accuracy and inventory control. However, the accessibility of these tools for non-technical managers is not explored in the study.

Pundir et al. (2020) applied decision trees and XGBoost to point-of-sale data, achieving strong short-term sales forecasts. However, while technically robust, their study did not consider how outputs could be integrated into dashboards for business users, limiting practical applicability of the findings.

In contrast, Chong et al. (2016) incorporated sentiment analysis and competitor data into forecasting models. This broader input increased the potential for strategic insights and their emphasis on stakeholder collaboration aligns with the user-centred approach of this dissertation. However, their work stayed at the conceptual stage without producing an operational tool.

Overall, the literature confirms that predictive analytics can improve retail forecasting, but most studies either focus on model accuracy without addressing usability or remain conceptual without practical deployment. This gap shows the need for solutions that balance predictive performance with accessibility for non-technical users and decision-makers.

2.3 AI Dashboards: architecture, usability, and challenges

AI dashboards extend traditional business intelligence (BI) systems by integrating real-time predictions and adaptive features that can learn from user interactions (Flores, 2024).

Yigitbasioglu and Velcu (2012) identified clarity, interactivity and user control as essential dashboard design principles. Moreover, Kandel, et al. (2012) found that users often struggle to interpret raw model outputs unless explanations are included, while Davenport and Ronanki (2018) showed that many AI projects fail due to poor alignment with business workflows rather than model inaccuracy. In retail, this implies that even accurate predictions may be rejected if they disrupt established decision-making routines. These findings reinforce the importance of explainable AI (XAI) for building trust.

More recent research, such as the study conducted by Mgbame et al. (2022) examines the development of low-cost dashboards for business process optimisation in SMEs. Their work highlights that integrating clear model outputs with interactive visualisations can improve user understanding and decision-making efficiency, even in resource-constrained environments (Mgbame et al., 2022). While their approach does not focus specifically on retail, the findings support the argument that accessible, well-designed dashboards can enhance adoption rates among non-technical users. However, studies applying these principles directly to retail-specific AI dashboards remain limited.

In summary, the literature highlights strong design principles and the need for explainability of AI dashboards. However, AI dashboards for retail often lack the integration of both predictive accuracy and user-centred features. This dissertation contributes to this gap by embedding interpretable models into a simple interface made for retail workflows.

2.4 Comparing Tools: BI platforms vs. Custom AI dashboards

Mainstream BI tools such as Power BI, Tableau, and Qlik provide strong visualisation and integration capabilities but have limited native AI features Richardson et al. (2020). Richardson, et al. (2020) state that advanced predictive functions in these tools often require external coding or plugins, creating barriers for SMEs without data science resources. Custom AI dashboards can be more flexible and domain-specific, but adoption depends heavily on usability. Amershi et al. (2019) proposed human–AI interaction principles, noting that even highly accurate models may fail if interfaces are not intuitive. For SMEs in retail, where technical capacity is limited, it can be argued that simplicity and transparency are as important as accuracy.

Comparative research by Hosen, et al. (2024) shows that while BI platforms excel in rapid deployment and familiarity, custom solutions often outperform them in predictive accuracy and domain relevance, especially when designed with active user involvement during development. Their analysis of advanced database systems for business intelligence emphasises that tailoring system architecture to specific organisational needs can maximise both analytical power and business value (Hosen et al., 2024).

Therefore, the current literature shows that BI platforms are accessible but often limited in predictive capability, whereas custom AI dashboards can provide higher forecasting power but require considerably longer development time. This dissertation adopts a custom-built approach to combine domain relevance with advanced predictive features.

2.5 Bridging business analysis with technical AI tools

A persistent gap in the literature is the lack of alignment between AI system design and the actual workflows of business users. Shollo and Galliers (2016) note that dashboards are often ignored if outputs do not match the decision-making context. Similarly, Mikalef, et al. (2018) defends the importance of aligning big data strategies with organisational routines, stressing that dashboards should be embedded in daily processes rather than treated as standalone tools. To help address this mismatch between dashboard design and the organisational contexts in which they are used, Wixom et al. (2014) suggest a layered dashboard design where business priorities guide both model choice and visualisation. However, these frameworks often remain theoretical and empirical studies that test such designs in operational retail contexts are still limited.

Recent work by Gami et al. (2024) provides evidence that embedding predictive dashboards within regular business processes can increase adoption and deliver measurable improvements in key performance indicators. Their study, which developed an interactive data quality dashboard integrating real-time monitoring with predictive analytics, showed that insights directly aligned with operational workflows lead to higher user engagement and more effective decision-making (Gami et al., 2024). All in all, the literature supports a user-centred, workflow-integrated approach to AI dashboards. This project translates that principle into practice by building a predictive tool that reflects real retail tasks and decision points identified during interviews with stakeholders.

2.6 Summary and Research Justification

The literature review indicates that research on AI-driven forecasting tools for retail has often prioritised predictive accuracy over usability, with limited evaluation involving real business users. Retail-specific studies using real-world operational data remain scarce, and there is a notable lack of deployable prototypes validated in real-life business contexts. These gaps restrict the practical adoption of AI forecasting solutions in retail environments, where decision-making relies on technical performance and seamless integration into daily workflows.

Aiming to address these gaps, this dissertation employs a large real-world retail dataset comprising 307,647 records and integrates multiple forecasting models, namely Prophet, XGBoost, and LSTM, to capture insightful patterns. The outputs are delivered through an interactive Streamlit dashboard designed for multiple user roles, ensuring accessibility for non-technical retail professionals. Furthermore, the solution is evaluated using a mixed-methods approach, assessing both predictive performance and business usability. By combining advanced modelling with a user-centred interface, this work aims to bridge the gap between AI capability and business adoption in the retail sector.

3 Methodology

This section describes the methodological approach used to design, build and evaluate a predictive AI dashboard for strategic decision-making in the retail sector. The process was designed for replicability, practical relevance and academic rigour, combining technical development with business-oriented insights to ensure alignment between AI capabilities and real user needs.

3.1 Overview of Research Process

The methodology followed an adapted version of the Cross Industry Standard Process for Data Mining (CRISP-DM) framework due to its suitability for structuring data science projects from problem definition to deployment (Schröer et al., 2021).

This framework was extended with explicit steps for business context understanding and usability testing, which are essential when the final output is a decision-support tool. The understanding of the business context was obtained through literature review and interviews with retail professionals.

The research stages were the following:

1. **Business understanding:** Analysing retail decision-making needs through literature review and semi-structured interviews.
2. **Data acquisition and preparation:** Using a large, publicly available retail dataset.
3. **Modelling:** Implementing and comparing three forecasting algorithms.
4. **Dashboard development:** Building a Streamlit-based interactive interface.
5. **Evaluation:** Assessing predictive accuracy and usability.




Figure 1: Research Methodology Framework

3.2 Data source and description

The dataset used was the “Warehouse and Retail Sales” dataset from the U.S. Government Data Catalog. It is composed by 307,647 records in CSV format, updated monthly, and includes the following variables:

- **Temporal variables:** YEAR, MONTH (for time-series structure)
- **Categorical variables:** SUPPLIER, ITEM DESCRIPTION, ITEM TYPE
- **Target variable:** RETAIL SALES
- **Supporting variables:** WAREHOUSE SALES, RETAIL TRANSFERS

The dataset was selected for its real-world operational relevance, sufficient size for modelling and availability for reproducibility.

3.3 Data collection, cleaning and transformation

Data was downloaded in CSV format and processed in Python (Pandas, NumPy). The preprocessing workflow began with the handling of missing values, specifically by removing records containing null sales values. Next, data types were standardised through type casting to ensure correct numerical and datetime formats. Feature engineering was then performed, creating two key variables: `total_demand`, calculated as the sum of `warehouse_sales` and `retail_sales`, and a `seasonality_index` derived from the month to capture temporal patterns. Following this, the data was aggregated at a monthly level by product and supplier to align with forecasting objectives. Finally, categorical variables were transformed through label encoding to ensure compatibility with machine learning models. These steps ensured data quality and compatibility with multiple modelling techniques, aligning with best practices for predictive analytics in retail (Willmott & Matsuura, 2005).

3.4 Predictive modeling techniques

Three algorithms were chosen to balance interpretability, accuracy and ability to capture temporal patterns (see table 1).

Regarding the evaluation metrics used, RMSE and MAE were selected as they are widely recognised in forecasting research and make it possible to compare models even when their outputs are on different scales (Willmott & Matsuura, 2005). Models were trained and tested using an 80/20 split to reduce overfitting risk.

Table 1: Selected Forecasting Models

Model	Rationale for Selection	Key Strengths
Prophet	Widely used for time-series with trend and seasonality (Taylor & Letham, 2018)	Fast training, interpretable, suitable for business audiences.
XGBoost	High performance in structured data forecasting (Qiu et al., 2022)	Handles categorical and numerical features, strong accuracy.
LSTM	Designed for sequential dependencies (Yu et al., 2019)	Captures long-term patterns, adaptable to volatile data.

3.5 Tools, libraries, and environment

This project was developed using the following tools and frameworks:

Table 2: Tools and Libraries Used

Tool / Library	Purpose
Python	Programming environment
Pandas, NumPy	Data manipulation
Matplotlib, Seaborn	Visualisation
Scikit-learn	Machine learning models and evaluation
Prophet	Time-series forecasting
TensorFlow	Deep learning (LSTM model)
Streamlit	Frontend dashboard development
VS Code	Code editor

All implementations were completed on a standard laptop (Intel i7, 16GB RAM), demonstrating the feasibility of deploying AI solutions using accessible hardware.

3.6 Business Interviews for Contextual Insights

3.6.1 Interview Process and Objectives

Five semi-structured interviews were conducted with retail professionals, including inventory managers, data analysts and commercial planners. Participants were identified through purposive sampling, ensuring diversity of perspectives while focusing on roles directly involved in retail decision-making (Adeoye-Olatunde & Olenik, 2021).

In this study, the interviewees were personal contacts of the researcher working within retail industry contexts in both Chile and the USA. The participants were selected for their direct experience with operational and analytical processes relevant to the research. All participants were invited personally and agreed to take part voluntarily. The interviews were conducted online in sessions lasting between 30 and 60 minutes. Ethical considerations and procedures were followed in line with the requirements of the National College of Ireland.

The interviews were conducted with three main objectives. First, to identify the core dashboard features most valued by retail stakeholders, ensuring the proposed solution addressed practical business needs. Second, to evaluate stakeholder preferences regarding transparency, usability, and the generation of actionable outputs, with a view to aligning technical design with user expectations. Third, to explore perceived barriers to AI adoption within retail workflows, providing insights into organisational and cultural factors that could influence the successful implementation of the forecasting dashboard.

Following the development of the prototype, the same participants were invited to test the dashboard and provide feedback on its usability, clarity and relevance to their operational needs.

3.6.2 Key findings for dashboard design implications

The interviews revealed several recurring themes that directly informed the design of the predictive dashboard. Insights from participants were translated into specific interface elements and functional features, ensuring that the tool addressed both technical performance and practical usability needs. (**Table 3**) summarises the main findings and their corresponding impact on the dashboard’s development.

Table 3: Interview Themes and Dashboard Impact

Interview Insight	Dashboard Feature Developed
Need for clear, non-technical outputs	Added “forecast explanation” text box
Desire for downloadable reports	Implemented CSV/PDF export option
Concern about over-reliance on AI	Included manual override feature
Importance of seasonal planning	Built seasonality visualisation module

3.7 Frontend dashboard development

The dashboard was built in **Streamlit** for its speed, flexibility and Python integration. The frontend developed was divided in the three following modules:

1. **Sales Forecasting** — Model selection, actual versus predicted charts.
2. **Inventory Insights** — Highlighting predicted stock risks.
3. **Promotion Timing Recommender** — Suggesting campaign windows based on trends.

User interface design applied human AI interaction principles (Amershi et al., 2019) to ensure accessibility for both technical and non-technical users.




Figure 2: Business Vision




Figure 3: Technical Vision

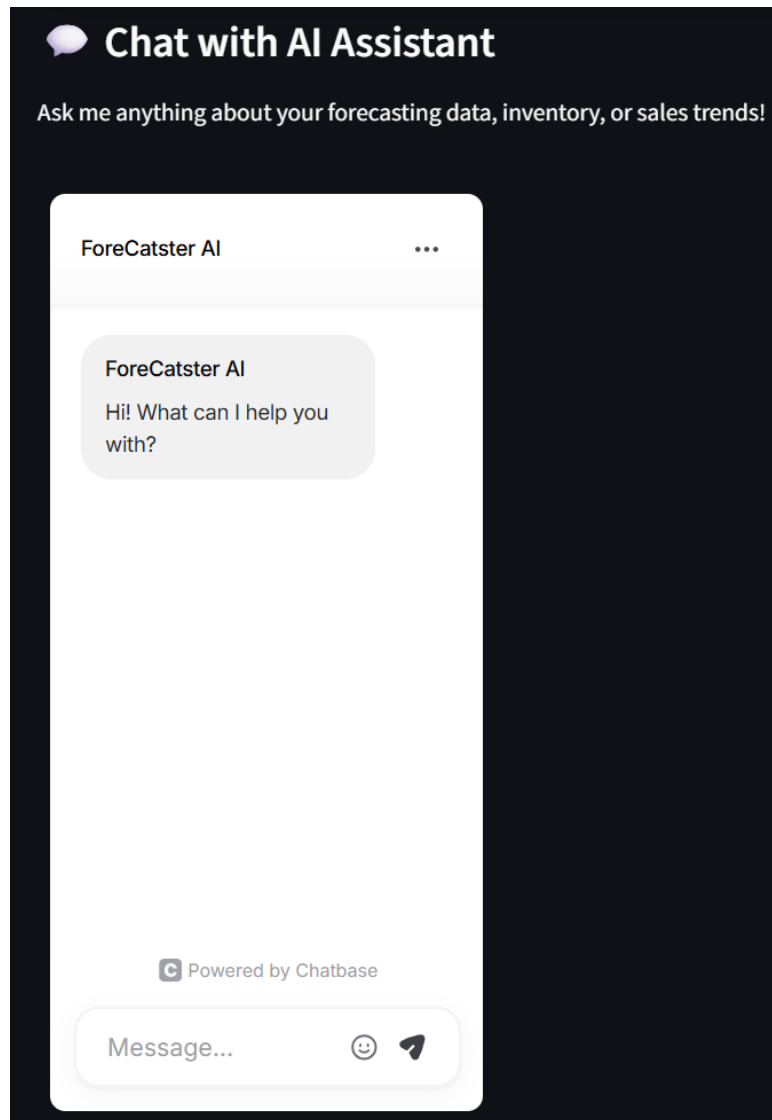


Figure 4: AI Chatbot Integration

3.8 Evaluation strategy

To assess both the predictive accuracy and the practical usability of the proposed dashboard, a mixed-methods evaluation approach was adopted. This strategy combines quantitative metrics, which provide objective measures of model performance, with qualitative insights, which capture the user experience and contextual relevance of the tool.

Quantitative:

- RMSE and MAE calculated per model.
- Residual analysis to identify bias patterns.

Qualitative:

- Post-use interviews to gather feedback on usability, clarity, and relevance.
- Simulated scenarios (e.g., seasonal stock planning) to test decision-making support.

This mixed-methods evaluation aligns with the aim of balancing technical performance with practical adoption potential (Palinkas et al., 2019).

3.9 Limitations of Methodology

This study presents some methodological limitations. The interview sample was small and geographically concentrated, which may limit the generalisability of the findings. The dataset used did not include promotional or competitor data, potentially omitting factors that influence retail demand. In addition, the dashboard operates on batch historical data, with real-time integration identified as an area for future development.

To mitigate potential biases, several measures were taken. A train/test split was applied to prevent overfitting in the predictive models. Efforts were made to ensure diversity among interview participants, capturing a range of perspectives within the retail sector. Furthermore, consistent preprocessing steps were applied across all datasets to avoid data leakage and ensure comparability of results.

4 Design and Implementation Specifications

This section outlines the conceptual design, and the final implementation of the AI-powered dashboard developed for retail decision-making. The system was designed to be both technically robust and business-oriented, ensuring accessibility for users with little technical knowledge while leveraging advanced machine learning models for predictive insights.

4.1 Design specification

The architecture of the system was designed with three main objectives:

1. Enable predictive analytics using structured retail data
2. Provide an intuitive, interactive dashboard for various user roles
3. Support modular extension of new models or insights in the future

In sum, it is aimed to deliver accurate forecasts in a format that supports practical decision-making. The following subsections outline the technical implementation.

4.1.1 System architecture

The dashboard follows a layered architecture, composed of:

- **Data layer:** Connects to the CSV dataset; pre-processed using Pandas
- **Model layer:** Contains three core forecasting models – Prophet, XGBoost, LSTM
- **Visualisation layer:** Built with Streamlit, providing a clean UI for user interaction
- **Logic layer:** Handles user inputs, routing, and response generation

This modular structure enables independent development and testing of each layer.

4.1.2 Functional Design

The dashboard was designed with the following key features:

Table 4: Dashboard Features and Functionalities

Module	Functionality
Forecast Explorer	Users select product or supplier, and choose which model to view (Prophet, XGBoost, LSTM). Line graphs show actual vs. predicted sales.
Inventory Insights	Displays predicted inventory movements and potential stock issues using rolling averages.
Promotion Timing	Highlights peak and low sales windows using forecasted data and trend slopes.
Recommender Model Comparison Tool	Allows side-by-side comparison of RMSE/MAE between models.
Each module uses user-friendly controls (dropdown menus, date sliders, expandable info boxes) to improve navigation and interpretation.	

4.1.3 User Types and Roles

The interface was designed to support **three user personas**:

- **Retail Managers:** Need clear insights to plan inventory and promotions
- **Data Analysts:** Interested in model comparison and deeper metrics
- **General Staff or Executives:** Require simplified views of business trends

Each feature was validated against these roles to ensure relevance and clarity.

4.2 Implementation / Solution Development Specification

The final implementation stage translated the design into a fully functional dashboard application. The process was divided into data transformation, model development, and front-end deployment.

4.2.1 Data Transformation

All data preprocessing steps described in Section 3 were applied using Python libraries. The final dataset included many product categories across five years. Time-based features were created to capture seasonality and sales trends. Numerical scaling and sequence reshaping (for LSTM) were also performed at this stage.

4.2.2 Model Implementation

The following machine learning models were developed and integrated:

- **Prophet (by Meta):** Fast, flexible, and interpretable forecasts. Tuned with holiday effects and changepoints.
- **XGBoost:** Used with one-hot encoded categorical features, date features (`month`, `year`), and lag variables. Tuned with `GridSearchCV` for optimal learning rate and tree depth.
- **LSTM:** Implemented in `Keras`, trained on 12-month lookback windows. Required reshaped input (3D tensors) and output post-processing for interpretation.

All models were exported as Python functions and wrapped with *Streamlit* inputs to control parameters like forecast horizon, product selection, and model type.

4.2.3 Dashboard Deployment

The dashboard was built using Streamlit, a lightweight Python framework for web app, removing the need for prior knowledge of web development languages such as HTML, CSS, or JavaScript. Another advantage of using Streamlit was that it offers native support for charts, interactive widgets and sidebars, and also allows for easy deployment either locally via localhost or in the cloud through Streamlit Cloud.

The layout is organised into three main pages: **Home & Overview**, which provides a summary of the dashboard and its purpose; **Forecast Explorer**, which enables interactive charting and model selection; and **Insights Center**, which delivers inventory alerts and promotion timing recommendations. All charts were developed using Matplotlib and Plotly for dynamic interaction and clear visualisation of forecasting results.

4.2.4 Outputs Produced

The project resulted in the multiple outputs described in (Table 5). Each of these outputs was documented with in-code comments and saved locally. Additional documentation for deployment and setup is planned as future work.

Table 5: Project Outputs and Intended Users

Output	Description	Potential Users
Cleaned Dataset	Processed sales data ready for ML	Data scientists
Prophet, XGBoost, LSTM models	Trained and saved models	Analysts, developers
Streamlit Dashboard	Interactive web interface	Retail staff, managers
Evaluation Charts	RMSE/MAE comparison visuals	Business decision-makers
Forecast Scenarios	CSV files with future demand predictions	Inventory planners

4.3 Summary of design and implementation

The design and implementation phases were guided by the principles of usability, modularity and performance. The system offers real business value by transforming raw sales data into predictive insights that are both visually interpretable and strategically actionable. Unlike many academic studies that stop at model development, this project provides a deployable tool that retail businesses could adopt and adapt for real-world use.

5 Evaluation: Results and Critical Analysis

The purpose of this section is to provide a comprehensive analysis of the results and main findings of the study as well as the implications of these findings both from academic and practitioner perspective are presented. Only the most relevant results that support your research question and objectives shall be presented. Provide an in-depth and rigorous analysis of the results. Statistical tools should be used to critically evaluate and assess the experimental research outputs and levels of significance.

5.1 Model performance evaluation

Each model was tested on the cleaned and transformed dataset, with performance evaluated using Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE).

Table 6: Model Performance Comparison


Model	Avg. MAE	Avg. RMSE	Strengths	Weaknesses
LSTM	1459.97	1890.78	Good for long-term trends	Requires tuning, high complexity
Prophet	2337.60	2730.99	Fast, interpretable, captures seasonality	Less accurate with volatile series
XGBoost	1546.53	1883.15	High accuracy, handles categorical data	Slower to train, less interpretable

Key Findings:

Results varied across product categories, but general trends emerged. LSTM achieved the lowest average RMSE across all models, marginally outperforming XGBoost, and demonstrated strong capabilities in modelling stable sequential patterns, aided by adaptive configuration and padding strategies for shorter series.

XGBoost delivered consistently competitive performance, particularly in product segments characterised by non-linear relationships and categorical influences such as supplier or product type, maintaining high predictive accuracy across diverse conditions.

Prophet offered the highest interpretability for business users, making it valuable for explaining seasonal and holiday effects, but its predictive accuracy was notably lower in categories with volatile or irregular sales patterns.

**Figure 5:** Comparison Models

5.2 Dashboard usability evaluation

To evaluate the system from a business usability perspective, interviews were conducted with five representative stakeholders from the retail sector (inventory managers, data analysts, general staff). Key themes and insights were extracted using a thematic analysis approach (See Table 7).

Participants highlighted the value of clear visualisation for decision-making, with one retail category manager noting that *“being able to visualise the contrast between peak and low sales periods really helps in planning promotions, especially for seasonal items.”* The

model comparison tool was also well received, as an inventory planner explained, “*even if I don’t understand all the math behind it, I can still make informed choices with the information provided.*” Across the interviews, users described the Streamlit dashboard as “*intuitive*” and “*simple to navigate*”, appreciating features such as dropdown filters, downloadable forecasts and the ability to compare model outputs to guide stock orders.

Table 7: Thematic Analysis and Interview Insights

Theme	Stakeholder Feedback Summary
Forecast	Prophet graphs were preferred for visual clarity; XGBoost seen as accurate but harder to interpret
Transparency	
Feature	Inventory insights and promotion timing modules were considered highly valuable for daily decisions
Relevance	
Interface	Streamlit dashboard was described as “intuitive” and “simple to navigate”, with useful dropdown filters
Experience	
Actionability	Users appreciated the ability to download forecasts and compare model outputs to inform stock orders

5.3 Critical Analysis

The project achieved its main goal of building a functional, interpretable, and business-centred predictive dashboard. One of its main strengths is found in bridging technical depth with practical usability: advanced forecasting models (LSTM, XGBoost and Prophet) are delivered through an interface that non-technical users can easily navigate.

The model diversity ensured by including multiple models offers flexibility, with LSTM achieving the lowest RMSE, XGBoost performing closely behind, and Prophet providing high interpretability. This diversity enables users to balance accuracy against explainability depending on the decision context.

Furthermore, the user-centric design was validated by the five interviewees, indicating that the dashboard effectively supports daily decision-making tasks such as inventory planning, promotion timing and communication of trends.

Despite these strengths, several limitations were identified, as summarised in (Table 8).

Table 8: Limitations Identified During Development and Evaluation

Limitation	Description
No real-time data	The application runs on static historical files; live connectors/APIs are planned for future iterations.
Limited user testing scale	Interviews covered five stakeholders; broader, longitudinal testing across multiple retailers would improve external validity.
External drivers excluded	Price changes, marketing activity, and exogenous shocks (e.g., inflation, supply constraints) were not modelled; these may explain part of the residual error.
Model complexity vs. transparency	LSTM offers top accuracy but requires careful tuning and is harder to explain; governance may prefer Prophet traces in some settings.

5.4 Evaluation’s summary

The evaluation shows that the dashboard is technically sound and strategically actionable. The quantitative metrics confirm that LSTM and XGBoost provide strong forecasting performance, while Prophet helps explain seasonal effects to managers. The qualitative feedback highlights that the interface is intuitive and supports practical tasks like stock planning and promotion timing. Future releases should expand data scope (prices, promotions, competitor signals), enable real-time updates and run broader live user testing to validate impact on operational key performance indicators (KPIs) (e.g., stockouts, margin, and markdown rates).

6 Conclusion and Future Work

This section reflects on the outcomes of the research, placing the results within the broader academic and business context. It critically examines how the developed AI-powered dashboard meets real-world decision-making needs in retail, how it aligns or diverges from prior literature, and what this implies for the adoption of predictive analytics in retail contexts.

6.1 Alignment with Research Objectives

The primary aim of this study was to develop and evaluate a predictive dashboard capable of supporting retail decision-making by integrating AI models into a user-friendly interface. The findings from both the technical implementation and stakeholder interviews indicate that this objective was successfully met. This achievement is reflected in several core results detailed below.

6.1.1 Forecast Accuracy

Quantitative results confirmed the reliability of the three models. LSTM achieved the highest accuracy (lowest RMSE), demonstrating strong capability in capturing long-term sequential patterns. XGBoost delivered competitive performance across most categories, particularly those with complex non-linear and categorical relationships, while Prophet provided interpretability and adaptability for business contexts.

6.1.2 Business Relevance

Interviews revealed strong demand for tools that improve planning, reduce uncertainty and simplify forecast interpretation. These exact needs proved to be directly addressed by the dashboard features included in this project’s model (e.g., model toggle, inventory alerts, promotion timing).

6.1.3 Usability and Trust

The inclusion of tooltips, dropdowns, and transparent error metrics contributed to higher perceived usability, addressing common barriers in AI adoption such as complexity and lack of trust.

These outcomes demonstrate that AI tools, when embedded in accessible interfaces and aligned with user needs, can effectively enhance data-driven decision-making processed in the retail sector.

6.2 Comparison with previous literature

Numerous academic studies have explored the potential of AI and dashboards in decision support. However, this project contributes with a unique blend of technical modelling, user interface deployment and stakeholder perspective. In alignment with Shollo and Galliers (2016), who argue that decision-makers often ignore dashboards unless they fit their operational context, the modular human-centred structure of the developed system addresses this concern by adapting content to specific user roles.

Similarly, this study complements the perspective of Mikalef, et al. (2018), who emphasise that dynamic capabilities such as organisational learning and alignment are essential for AI tools to create value. The research operationalises this principle by designing interfaces that bridge technical depth with business interpretation, ensuring that advanced models remain accessible to non-technical users (Mikalef et al., 2018). In doing so, this study addresses a practical gap in the literature, where usability and explainability are often secondary to accuracy-focused discussions.

In contrast to much of the existing literature, which often concludes at the model training stage or remains within theoretical frameworks, this work follows the CRISP-DM methodology through to the deployment phase, extending it into real-time decision layers. By doing so, it not only complements existing methodological discussions but also fills a gap in demonstrating how CRISP-DM can translate from data to action (Kotu & Deshpande, 2014).

The resulting outcome of this project is both technically robust and validated through business feedback, offering a rare deployable solution that connects academic theory with practical business application.

6.3 Contribution to Business Practice

The dashboard developed in this study provides a viable prototype for AI-powered retail decision support, offering several tangible contributions to business practice. In terms of planning efficiency, the ability to forecast future demand and identify potential stock issues enables managers to reduce uncertainty in procurement and distribution. The high accuracy of the LSTM model enhances reliability for categories with stable demand patterns, while XGBoost delivers robust results for segments influenced by multiple factors.

Promotion optimisation is another key benefit. By analysing sales trends and providing timing recommendations, the dashboard supports the alignment of promotional campaigns with demand cycles. Prophet’s interpretability is particularly valuable to clearly identify seasonal peaks and holiday effects.

The system also contributes to workforce empowerment by lowering the barriers to AI adoption and giving non-technical staff access to interpretable insights. Compared to traditional business intelligence tools, the dashboard moves beyond descriptive analytics toward predictive and prescriptive analytics, a shift that matches trends outlined in Gartner’s Analytics Maturity Model.

6.4 Recommendations for Future Work

To further enhance the dashboard and expand the impact of this research, several directions for future work are recommended. First, live testing with real retail teams would allow for validating user experience (UX) and measuring the decision-making impact in operational contexts. Second, integrating additional variables could increase model relevance and improve the accuracy of forecasts. Third, exploring the use of federated data models would enable access to richer datasets while preserving data privacy. Finally, deploying the system on cloud platforms with feedback loops could facilitate continuous learning and adaptive performance over time. Implementing these steps would strengthen the business applicability and long-term value of predictive analytics in retail.

References

- Adeoye-Olatunde, O. A., & Olenik, N. L. (2021). Research and scholarly methods: Semi-structured interviews. *Journal of the american college of clinical pharmacy*, 4(10), 1358-1367.
- Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., & Inkpen, K. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 chi conference on human factors in computing systems*,
- Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS quarterly*, 1165-1188.
- Chen, Z., Zhao, J., & Jin, C. (2025). Business intelligence for Industry 4.0: predictive models for retail and distribution. *International Journal of Retail & Distribution Management*, 53(3), 1-16.
- Chong, A. Y. L., Li, B., Ngai, E. W., Ch'Ng, E., & Lee, F. (2016). Predicting online product sales via online reviews, sentiments, and promotion strategies: A big data architecture and neural network approach. *International Journal of Operations & Production Management*, 36(4), 358-383.
- DalleMule, L., & Davenport, T. H. (2017). What's your data strategy. *Harvard business review*, 95(3), 112-121.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard business review*, 96(1), 108-116.
- Flores, H. (2024). Ai sensors and dashboards. *Computer*, 57(8), 55-64.
- Gami, S. J., Shah, K., Katru, C. R., & Nagarajan, S. K. S. (2024). Interactive Data Quality Dashboard: Integrating Real-Time Monitoring with Predictive Analytics for Proactive Data Management.
- Hosen, M. S., Islam, R., Naeem, Z., Folorunso, E., Chu, T. S., Al Mamun, M., & Orunbon, N. (2024). Data-driven decision making: Advanced database systems for business intelligence. *Nanotechnology Perceptions*, 20(3), 687-704.
- Kotu, V., & Deshpande, B. (2014). Predictive analytics and data mining: concepts and practice with rapidminer. Morgan Kaufmann.

- Mgbame, A. C., Akpe, O.-e. E., Abayomi, A. A., Ogbuefi, E., & Adeyelu, O. O. (2022). Developing low-cost dashboards for business process optimization in SMEs. *International Journal of Management and Organizational Research*, 1(1), 214-230.
- Mikalef, P., Pappas, I. O., Krogstie, J., & Giannakos, M. (2018). Big data analytics capabilities: a systematic literature review and research agenda. *Information systems and e-business management*, 16(3), 547-578.
- Palinkas, L. A., Mendon, S. J., & Hamilton, A. B. (2019). Innovations in mixed methods evaluations. *Annual review of public health*, 40(1), 423-442.
- Pundir, A. K., Ganapathy, L., Maheshwari, P., & Kumar, M. N. (2020). Machine learning for revenue forecasting: A case study in retail business. 2020 11th IEEE Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON),
- Qiu, Y., Zhou, J., Khandelwal, M., Yang, H., Yang, P., & Li, C. (2022). Performance evaluation of hybrid WOA-XGBoost, GWO-XGBoost and BO-XGBoost models to predict blast-induced ground vibration. *Engineering with Computers*, 38(Suppl 5), 4145-4162.
- Richardson, J., Sallam, R., Schlegel, K., Kronz, A., & Sun, J. (2020). Magic quadrant for analytics and business intelligence platforms. Gartner ID G, 386610, 00041-00045.
- Schröer, C., Kruse, F., & Gómez, J. M. (2021). A systematic literature review on applying CRISP-DM process model. *Procedia computer science*, 181, 526-534.
- Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. *MIS quarterly*, 553-572.
- Shollo, A., & Galliers, R. D. (2016). Towards an understanding of the role of business intelligence systems in organisational knowing. *Information Systems Journal*, 26(4), 339-367.
- Soofastaei, A. (2019). *Advanced Analytics and Artificial Intelligence Applications*. BoD—Books on Demand.
- Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1), 37-45.
- S. J. Taylor and B. Letham, "Forecasting at scale," *The American Statistician*, vol. 72, no. 1, pp. 37-45, 2018.
- Waller, M. A., & Fawcett, S. E. (2013). Data science, predictive analytics, and big data: a revolution that will transform supply chain design and management. In (Vol. 34, pp. 77-84): Wiley Online Library.
- Willmott, C. J., & Matsuura, K. (2005). Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Climate research*, 30(1), 79-82.
- Wixom, B., Ariyachandra, T., Douglas, D., Goul, M., Gupta, B., Iyer, L., Kulkarni, U., Mooney, J. G., Phillips-Wren, G., & Turetken, O. (2014). The current state of business intelligence in academia: The arrival of big data. *Communications of the Association for information Systems*, 34(1), 1.

Yigitbasioglu, O. M., & Velcu, O. (2012). A review of dashboards in performance management: Implications for design and research. *International journal of accounting information systems*, 13(1), 41-59.

Yu, Y., Si, X., Hu, C., & Zhang, J. (2019). A review of recurrent neural networks: LSTM cells and network architectures. *Neural computation*, 31(7), 1235-1270.