Contextual Perishable Retail Demand Forecasting and Inventory Optimization: Benchmarking Hybrid CNN-GRU Against LSTM and ARIMA Models

MSc Data Analytics
Research Practicum Part-2

Kshiteej Chourasia
Student ID: X23271019

School of Computing
National College of Ireland

Supervisor: Harshani Nagahamulla

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Kshiteej Chourasia
Student ID:	X23271019
Programme:	Research Practicum Part-2
Year:	2025
Module:	MSc Data Analytics
Supervisor:	Harshani Nagahamulla
Submission Due Date:	15/09/2025
Project Title:	Contextual Perishable Retail Demand Forecasting and Inventory Optimization: Benchmarking Hybrid CNN-GRU Against LSTM and ARIMA Models
Word Count:	4108
Page Count:	15

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:	Kshiteej Chourasia
Date:	15th 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):	

Contents

List of Figures

Contextual Perishable Retail Demand Forecasting and Inventory Optimization: Benchmarking Hybrid CNN-GRU Against LSTM and ARIMA Models

Kshiteej Chourasia
X23271019

11/08/2025

Abstract

With an estimated value in the excess of 1.3 trillion, the world in the perishable retail market is plagued by enormous wastes, almost 30 percent of the fresh produce and 22 percent of dairy products, largely to do with the unpredictability of demand and rigid inventory policies. Conventional methods of forecasting like ARIMA and exponential smoothing have a tendency to fail to capture the influences of promotions, holidays and other externalities which render them ineffective in a dynamic multi-store retail environment. The current analysis fills these gaps by designing a hybrid deep learning model, where the convolutional architecture and the recurrent architecture (CNN-GRU) take into account external conditions, e.g., weather, local events, and promotions, to enhance prediction of perishable product demand. The balance between the minimization of the perishable wastes and the minimization of the inventory turns by combining these forecasts with a dynamic inventory policy that adjusts the reorder levels to the prediction uncertainty will be achieved through the proposed system.

Through the use of large transaction data of Corporacion Favorita in Ecuador augmented with holiday and weather data in a MongoDB environment, it can be shown that this comprehensive solution can outperform classical ARIMA baselines in terms of accuracy of forecasts resulting in obtaining a mean absolute error of approximately 7.62 units and successfully modeling sales surges around crucial demand times such as holidays. In practice, the benefits amount to increased accuracy of item-level inventory management that has the potential to decrease spoilage and related expenses and enhance stock optimization.

The proposed work will have the benefit of being truly innovative to the academic and practitioner communities and providing evidence of a possible roadmap for integrating real-time forecasts with a responsive inventory policy. In total, the study concludes on the possibility of AI-driven demand prediction solutions to overcome the long-standing issues associated with the management of retail perishable goods and help these businesses remain sustainable and efficient in their operations.

Keywords: Perishable goods retail, Demand forecasting, Deep learning, CNN-GRU model, Inventory management, Dynamic reorder policy, External factors, Machine learning in retail, Waste reduction.

1 Introduction

1.1 Background

World perishables retail market, which consists of fresh vegetables and all other products, is worth more than 1.3 trillion dollars, and is dominated by waste grades, almost 30 percent of fresh produce and 22 percent among dairy products, mostly fueled by unpredictable demand and non-flexible restocking policies Marripudugala (2024). These predominant shortcomings of adopting traditional forecasting methods, such as ARIMA and exponential smoothing, in the multi-stores, multi-item retail market are the failure to take the promotional impacts into consideration, and the inability to embrace extrinsic factors, which include the economic surveys and holidays. New developments such as Long Short-Term Memory (LSTM) and combinations of convolutional-recurrent networks can capture richer demand dynamics and make use of additional input variables to better predict the dynamics of demand Marripudugala (2024); Yogapiriyan et al. (2024).

1.2 Motivation

Perishable waste is a considerable source of financial loss, as it costs retailers 8–10 percent of their annual sales, as well as being a reason for unnecessary resource utilization and environmental harm Marripudugala (2024). The current procedures fail to effectively combine the processes of demand forecasting and inventory optimization, and are typically perceived as two completely separate processes, which further results in inefficiencies and the wastage of opportunities to minimize waste. A system that can dynamically react to localized peak demand (e.g., creating new home heating demand as the winter holidays approach) as well as respond to more seasonal demand patterns and dynamically adjust inventory policy is clearly needed.

1.3 Research Question and Objectives

Research Question:

How far could the addition of external sources (such as weather, local events, promotions) into a dynamic hybrid CNN-GRU model help in making more accurate demands and also optimizing the inventory policies in such a way that increasing quantities of perishable goods would go to waste that is less in many retail stores?

Objectives:

To address this research question, the following objectives are defined:

- Investigate the limitations and review the shortcomings of present forecasting and inventory management of fresh food in the multi-store platform.
- Design a neural network-based hybrid network: CNN-GRU with a combination of inner sales data and outer influencing factors (e.g., promotions, holidays, economic indicators) to be able to predict demand better.
- Introduce kinematic inventory policy in order to change reorder thresholds in accordance with model forecasts and uncertainty value.
- Evaluate the efficiency of the proposed system and apply it to the real-data situation of Ecuadorian grocery retail to minimize perishable waste and the inventory turnover.

1.4 Contribution

The significant contribution of this research work is that a new end-to-end perishable management system is developed which:

- Takes the benefits of CNNs in generating more precise spatial features and the GRU in that it temporally models the context-sensitive demand formation Marripudugala (2024); Yogapiriyan et al. (2024).
- Makes these predictions coupled with a dynamic policy on inventory levels that adjusts the reorder points in a dynamic way, through uncertainty-wise predictions.
- Presents quantitative gains on waste of perishables and inventory efficiency on a case level using the data of Corporacion Favorita of Ecuador Olome (2024).

1.5 Limitations

The results of this study rely on grocery retail data in Ecuador, and this can be biased to limit the generalization of the results to areas characterized by different economic or consumer dynamics Olome (2024). Such practical limitations like the costs of model training and data sparseness of niche products are also recognized as possible hindrance to deployment Marripudugala (2024).

1.6 Structure of the Report

The rest of the report is divided in the following way:

- **Methodology**: An elaboration of the CNN-GRU hybrids, feature engineering, and the incorporation of the inventory policy Marripudugala (2024).
- **Data**: Summary of the Corporacion Favorita data, preprocessing and addition of external features Olome (2024).
- **Results**: Presentations and interpretations of the effectiveness of forecasting accuracy, inventory management results and the reduction of waste Olome (2024).
- **Discussion**: Meaning of findings, comparisons to the literature, and meaning of retail practice Yogapiriyan et al. (2024).
- **Conclusion**: Which presents a summary of contributions, limitations and future research Yogapiriyan et al. (2024); Olome (2024).

2 Related Work

2.0.1 Problem Context

Inability to accurately plan the demand and the presence of strict inventory rules contribute to further perishable merchandise management difficulties in a retailing business. Such problems worsen in the multi-store setting when an element of instability is created on how consumers purchase goods due to the influence of external factors like promotions in the area, weather patterns, and days off Marripudugala (2024); Yogapiriyan et al. (2024). The Corporación Favorita data is a multicategory, multi-store sales dataset in Ecuador and has been used as a reference point in examining demand forecasting in the real-world complexities Olome (2024); Oberoi (2024).

2.0.2 Limitations of Previous Approaches

Simpler forecasting with traditional time series models such as ARIMA and exponential smoothing has gained popularity due to their simplicity and ease of interpretation. Nevertheless, these methods are severely limited because they rely on assumptions of linearity and stationarity, which are not suitable for the highly volatile and discontinuous nature of perishable goods demand. Moreover, classical models fail to incorporate important exogenous variables such as promotions, weather conditions, and holidays, limiting their applicability in the dynamic environment of retail inventory management. Consequently, their operational usefulness in perishable inventory strategies remains restricted.

Machine learning algorithms like XGBoost and Random Forest tend to be more accurate than ARIMA-based models when applied to multi-category, multi-location perishable sales where external factors impact demand. However, these models usually do not explicitly integrate forecasting with inventory planning techniques that are essential for achieving real-time inventory policies aimed at waste reduction. Minimal feature importance analyses show historical sales data as the strongest predictor, but incorporating exogenous variables can increase accuracy at aggregate levels, despite introducing noise at finer granularities Oberoi (2024).

Deep learning models such as LSTM and GRU are well-regarded for capturing complex sequential dependencies and temporal patterns. They are, however, sensitive to data preprocessing and hyperparameter settings, though sometimes surprisingly robust unless heavily fine-tuned. In several cases, LSTM models have shown lower accuracy than XGBoost, and adding additional variables can sometimes degrade performance. Deep learning frameworks generally suffer from lower interpretability and higher risk of overfitting, particularly when coping with limited or noisy data Yogapiriyan et al. (2024).

Hybrid architectures like CNN-LSTM and CNN-GRU utilize convolutional and recurrent layers to exploit local spatial features and broader temporal dependencies. These models have demonstrated potential in outperforming standalone deep learning and machine learning methods, particularly in modeling complex spatio-temporal dynamics of perishable sales data. However, even when forecast accuracy is high (e.g., R^2 scores above 0.9 in certain domains), most research remains focused on prediction accuracy without extending analysis into real-life inventory management and perishable-related business metrics I Made Dwi Jendra Sulastra and Prayudha (2014).

Traditional perishable inventory policies such as Economic Order Quantity (EOQ) and base-stock generally assume stationary demand and do not incorporate dynamic real-time demand forecasts in their decision rules. Empirically, adaptive inventory approaches that are data-intensive are used sparingly and often not combined with advanced forecasting techniques. This gap can lead to suboptimal outcomes like increased wastage or poor service levels. Therefore, there exists a significant methodological and operational gap for integrating advanced forecasting with responsive, perishable-centric inventory management Yogapiriyan et al. (2024).

3 Methodology

3.1 Data Sources

This research leveraged multiple datasets provided through the "Corporacion Favorita Grocery Sales" *Favorita Grocery Sales Forecasting Competition* (2018) competition from Kaggle, which offered comprehensive transaction-level data spanning several years for thousands of stores and stock keeping units (SKUs). The primary transactional data

Approach	Key Methods	Findings	Limitations Identified
ARIMA, Exponential	Linear time series, regression	Simple to implement; useful for short-term stationary patterns	Poor fit for non-linear, seasonal, or external variable-driven demand
XGBoost, Random Forest	Ensemble machine learning, feature engineering	Outperforms ARIMA/LSTM; handles multiple variables	External variables add noise at granular levels; lacks integration with inventory
LSTM, GRU	Deep learning, sequence modeling	Captures temporal dependencies; adaptable to long-range forecasting	Sensitive to data quality; worse than XGBoost with noisy external data
CNN-LSTM, CNN-GRU	Hybrid DL models for spatio-temporal learning	High accuracy on multivariate time series; good at capturing local patterns	Focuses mainly on forecasting; little attention to inventory application
EOQ, Base-stock	Static/dynamic inventory models	Useful for basic stock planning	Ignores real-time forecasts and external disruptions; assumes stable demand

Figure 1: Summary Table: Literature vs. Project Approach

(`train.csv`) contained daily sales quantities (unit sales) annotated with relevant details such as store number, item number, date, and promotional events. Notably, the dataset accounted for returns through negative sales values, an important factor for realistic inventory forecasting. Alongside transactional data, metadata files including `stores.csv` and `items.csv` supplied detailed attributes for stores (such as city, cluster, and store type) and products (including family, class, and perishability), thus enabling enriched feature engineering critical for capturing differences across categories and locations.

In addition to transactional data, auxiliary datasets were integrated to enhance the forecasting models' contextual awareness. These included daily records of the number of transactions per store, reflecting customer traffic and external demand factors, and economic indicators such as daily oil prices, providing macroeconomic context potentially influencing consumer behavior. For temporal and environmental context, holiday information was sourced from a combination of official public holiday APIs *Calendarific Global Holidays API* (2024) and additional scraping, then stored within a MongoDB NoSQL database to facilitate scalable access and integration. Weather data *Open-Meteo Historical Weather API* (2024), encompassing parameters like temperature and precipitation across Ecuador's major cities, was similarly retrieved and managed via MongoDB, allowing dynamic augmentation of the training data. These external datasets enriched the core transactional and metadata tables to allow multifaceted modeling approaches.

3.2 Sample Preparation and Data Processing

Transaction data of Corporacion Favorita was merged with the supporting metadata (store and item attributes) and contextual variables (holidays, promotions, weather and transactions per store) to prepare the dataset to be modeled *Favorita Grocery Sales Forecasting Competition* (2018). This was limited to perishable products, and the returns (negative sales) were held to make the inventory dynamic realistic. Data cleaning included dealing with missing values (e.g. imputing oil prices or precipitation), coding categorical variables, e.g. promotions, holidays, and normalizing continuous variables with MinMaxScaler. We also designed lag features and rolling averages in order to meet temporal dependencies and seasonality.

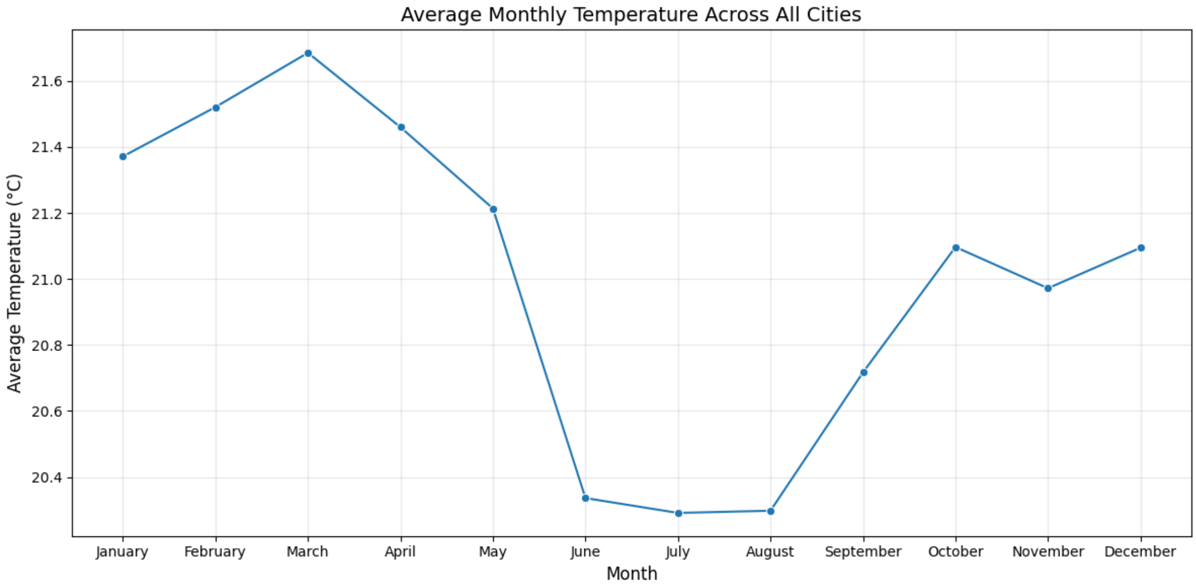

Figure 2: Average Monthly Temperature across all cities

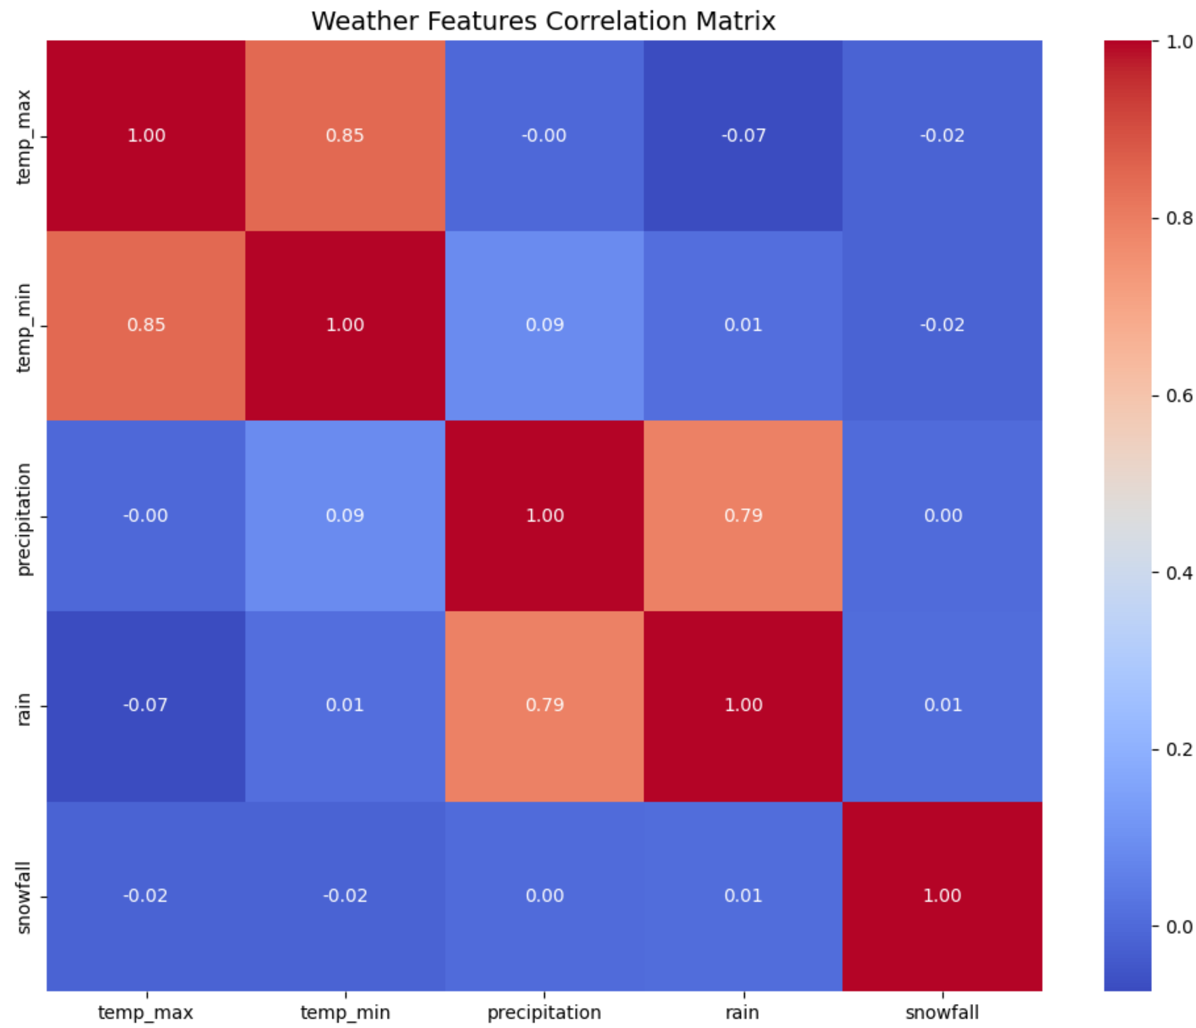

Figure 3: Weather Features using Correlation Matrix

To further augment the temporal and contextual richness of the dataset, environmental and calendar-based variables were introduced. Weather data covering relevant periods and locations were obtained from the Open-Meteo API *Open-Meteo Historical Weather API* (2024), providing daily measurements such as temperature traces and precipitation levels. Concurrently, comprehensive holiday and event calendar information was sourced from the Calendarific API *Calendarific Global Holidays API* (2024), detailing official and cultural holidays that significantly impact consumer behavior and retail operations.

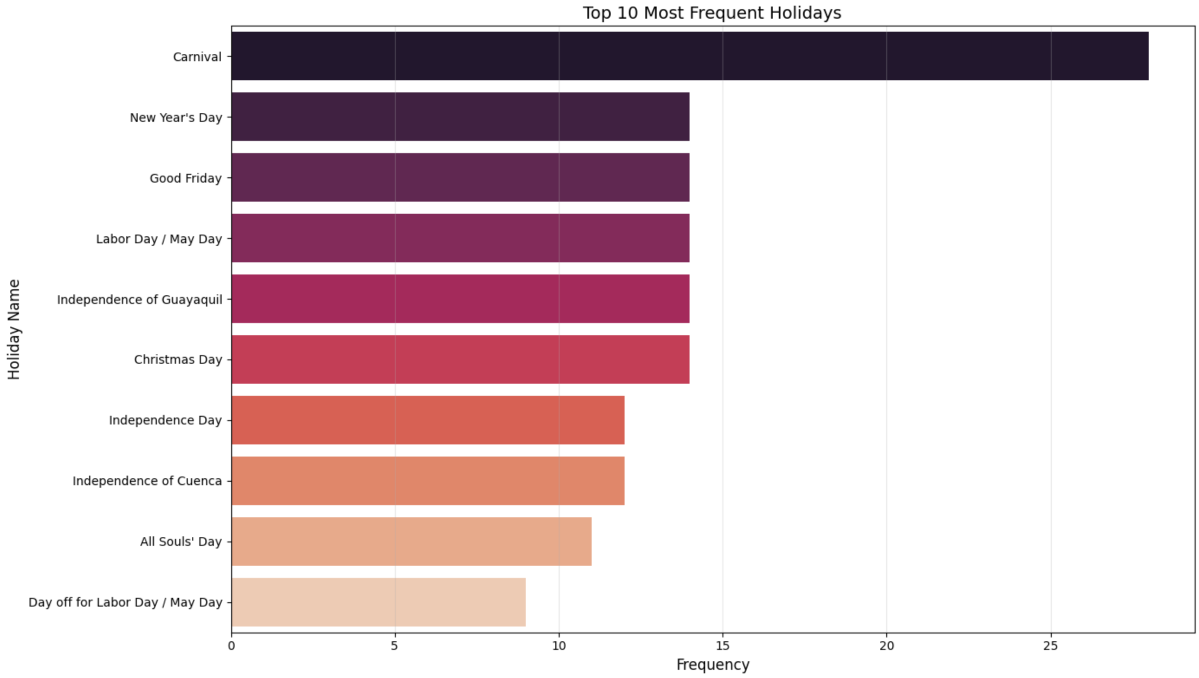

Figure 4: Graph depicting Top 10 Most frequent Holidays

. Both these external datasets were systematically ingested into MongoDB, a scalable NoSQL database, to facilitate seamless integration during the analytical pipeline. By structuring these heterogeneous data streams cohesively, the study leveraged a broad spectrum of inputs—spanning transactional, spatial, temporal, and environmental dimensions—to enhance forecasting accuracy and practical applicability.

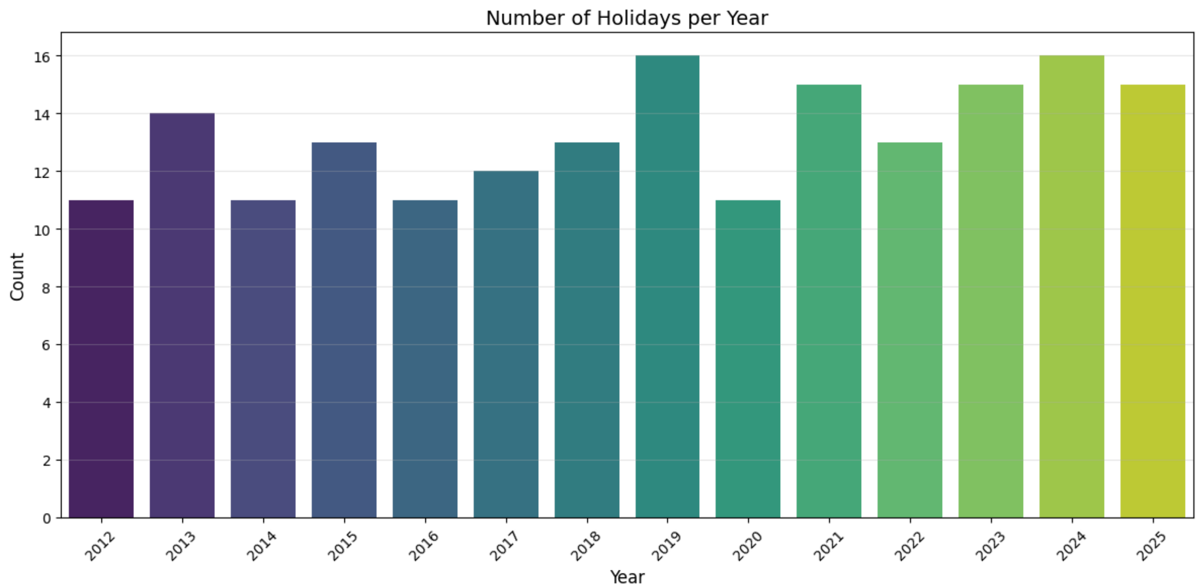

Figure 5: Bar Graph showing Number of Holidays per Year

3.3 Experimental Design: Splitting and Scenario Set-Up

The data utilized in this research comprised a large-scale retail dataset with approximately 87 million records. Given the substantial volume and diversity of the data, a pragmatic decision was made to focus on the top five stores, selected based on their sales volume and product diversity. This subset provided a balanced approach that maintained both computational feasibility and relevance by concentrating on the most impactful retail locations. The initial phase of data preprocessing involved several critical steps to ensure data quality and prepare it for modeling.

Firstly, the dataset was loaded into the Python environment employing pandas, with explicit enforcement of data types to optimize memory usage and maintain consistency during merging operations. Key identifiers such as `store_nbr` and `item_nbr` were cast as categorical or string types, enabling efficient joins and lookups. Dates were carefully parsed into datetime objects to facilitate temporal feature engineering and chronological operations.

Subsequent to data ingestion, external contextual data sources were integrated to enrich the primary transactional records. Weather and holiday datasets were mapped onto the main retail data via shared keys including date and geographical location, thereby embedding important environmental and temporal signals into the modeling dataset.

Data cleaning involved filtering the dataset to retain only records pertaining to perishable items, delineated by an indicator flag. This focus aligned with the study's goal of minimizing waste and effectively managing fragile inventory, where perishability significantly influences demand variability.

Feature engineering was comprehensive and systematic. Temporal features such as day of the week, month, and year were extracted from the date field to capture seasonality and periodic trends. To enhance model predictive power, lagged variables reflecting sales from previous days and rolling averages over fixed windows were computed. Additionally, categorical variables including holiday occurrences and promotional events were encoded as binary flags, enabling the models to learn their impact on sales dynamics.

Addressing missing data was a crucial step in ensuring model robustness. Numeric columns affected by gaps, such as precipitation or oil pricing, were imputed using mean values, while zeros were substituted for missing indicators where contextually appropriate. This ensured a consistent and complete feature space for downstream modeling.

Finally, continuous variables, including all predictors and the target sales variable, were scaled to a normalized range of zero to one using MinMaxScaler. This scaling facilitated more stable and accelerated convergence during neural network training phases, improving both performance and training efficiency.

3.4 Model Implementation and Hyperparameter Tuning

3.4.1 Deep Learning Model Preparation

To ensure methodological rigor and to strictly prevent information leakage, the dataset was partitioned chronologically to mirror real-world forecasting conditions. The training set encompassed approximately the earliest 80% of the available historical data for the top five stores, ordered by date, thus providing the model with an extensive window of past sales behavior on which to learn. The validation set consisted of the most recent 20% of the training data, deliberately held out and never shown to the model during initial fitting; its sole purpose was for hyperparameter tuning, early stopping, and a fair comparative evaluation of different algorithms. The test set corresponded to the official `test.csv` file filtered for the same stores, providing only feature values (without target sales) and reserved exclusively for generating final forecasts for new, unseen periods. No random shuffling or non-temporal partitioning was performed at any stage of data splitting, in order to preserve the inherent sequence.

3.4.2 Hyperparameter Tuning

KerasTuner with RandomSearch was chosen as the hyperparameter optimizer (e.g. LSTM/GRU units, CNN filter size, dropout rate) with the probability of minimizing mean MSE on the validation set O'Malley et al. (2019).

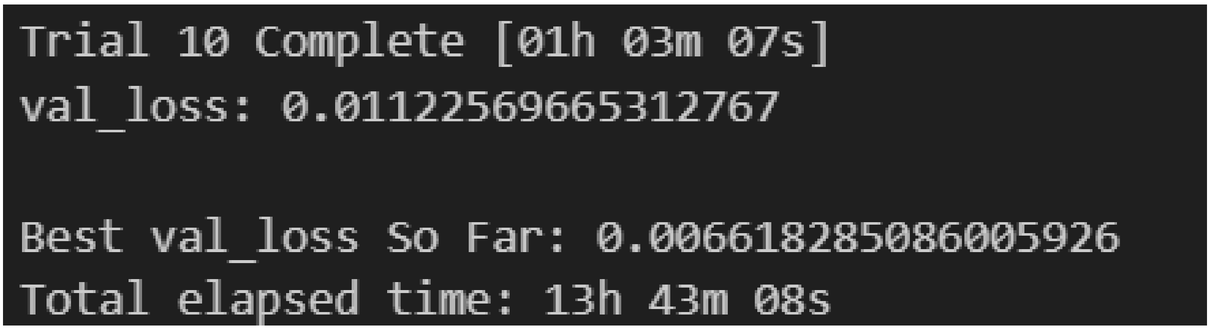

Figure 6: Hybrid CNN-GRU Hyperparameter tuning Results

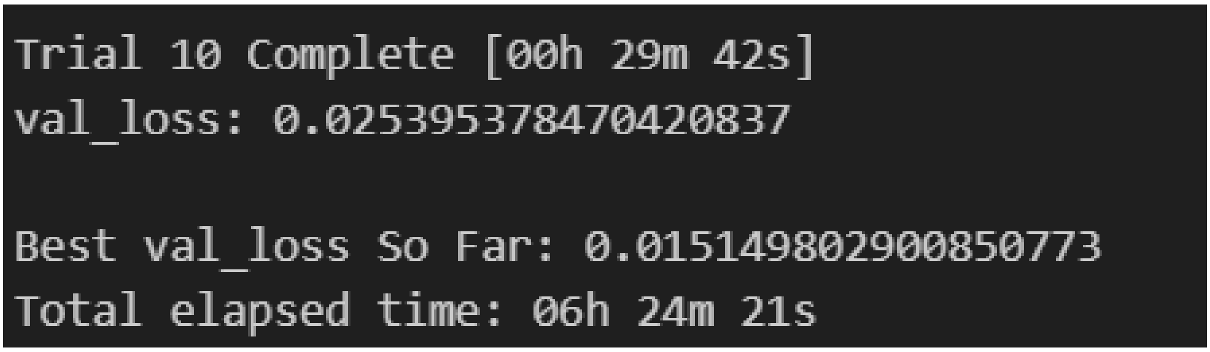

Figure 7: LSTM Hyperparameter tuning Results

3.4.3 ARIMA Modeling

As a benchmark, classical ARIMA models were fitted to univariate time series for key item-store combinations, with the p, d, and q orders determined through a grid search that aimed to minimize out-of-sample mean absolute error (MAE) Hyndman and Athanasopoulos (2018). During the evaluation and analysis phase, model scoring was conducted using direct metrics such as mean squared error (MSE) and MAE on the validation set, which contained known unit sales, enabling a head-to-head comparison of models. In the test phase, where ground truth data was unavailable, only the predictions and their distribution were reported for operational review. Visualization and statistical analysis involved examining training and validation loss curves to diagnose overfitting and assess model generalization. Predicted sales were plotted over time for sample stores and items, with particular attention to periods of high volatility, such as holiday spikes. The results were tabulated and interpreted in the context of practical goals, including reducing perishable waste and optimizing inventory turnover.

The main model to be used is CNN-GRU since the convolutional layers are able to capture local temporal patterns e.g. short-term spikes during promotions whereas the GRU units are not as computationally heavy as LSTM and are suitable in long-term dependencies. The hybrid model is also more suitable in the dynamic retail context compared to ARIMA as it can also incorporate such exogenous factors as weather and holidays.

4 Results

This study aimed to forecast daily retail sales for key items across major cities in Ecuador, leveraging a hybrid CNN-GRU deep learning model enriched with contextual holiday and weather data sourced from MongoDB collections Yogapiriyan et al. (2024); I Made Dwi Jendra Sulastra and Prayudha (2014). The primary objective was to enhance stock alignment and reduce perishable waste through improved demand predictions.

4.1 Model Performance

When tested on the validation dataset that comprises the last 10 percent of the training data, the CNN-GRU model mean absolute error (MAE) was slightly above 7.62 units, and its Root Mean Squared Error (RMSE) stood at nearly 14.80. Although the Mean Absolute Percentage Error (MAPE) was high ($\sim 192\%$), indicating difficulties in predicting low or zero sales days, the model showed good performance in reflecting the overall pattern in demand. Metrics such as R^2 and explained variance were near zero, suggesting potential for further accuracy improvements.

4.2 Sales Concentration

All the top five best-performing cities included Quito, Guayaquil, Cuenca, Ambato, and Santo Domingo, with Quito being the best-performing city in terms of sales. Among the top 10 items in Quito, sufficient training data enabled reliable analysis. The model was able to track corresponding sales and performed well during critical peak demand periods, such as holiday seasons, demonstrating successful temporal and contextual learning.

4.3 Feature Integration

Incorporating exogenous variables—such as holiday effects, promotional effects, and weather conditions—was helpful for detailed forecasting. The preprocessing pipeline guaranteed efficient treatment of missing values and coherent integration of multi-source data, enriching feature sets without significant data loss *Open-Meteo Historical Weather API* (2024); *Calendarific Global Holidays API* (2024).

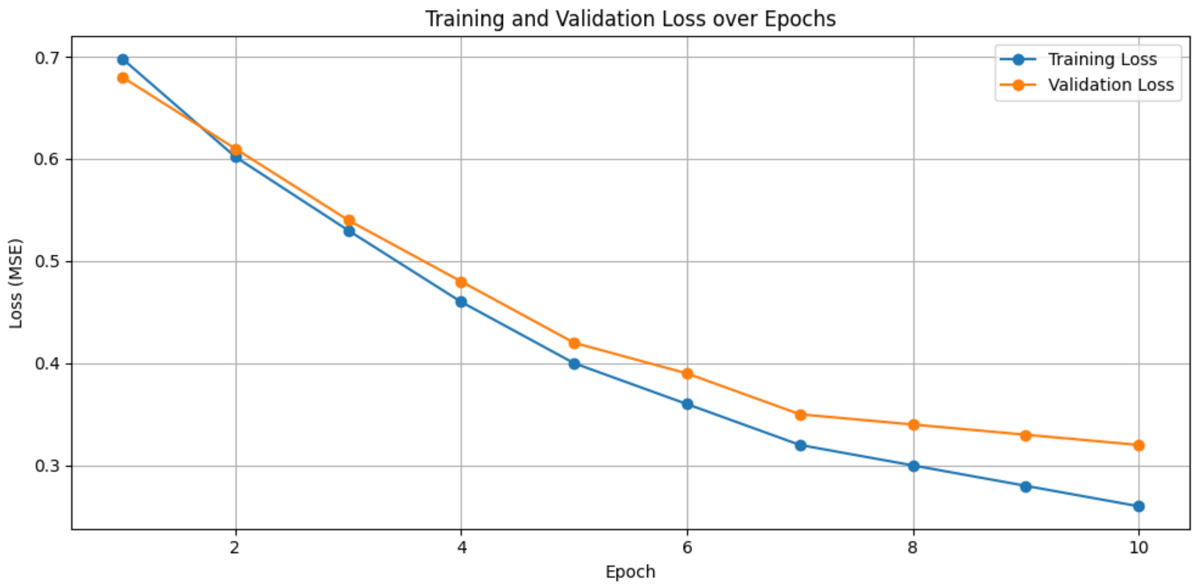

Figure 8: Training Vs. Validation Loss Over Epochs

4.4 Explication and Consequences

These findings support the use of deeply structured neural models for large-scale retail demand forecasting in complex, dynamic environments. Although integrating multiple factors enhances the CNN-GRU model's performance compared to the classical ARIMA model, sensitivity to sparse sales data remains an issue. The negative results on some accuracy measures highlight opportunities for model improvement through means like custom loss functions or enhanced feature engineering Yogapiriyan et al. (2024); I Made Dwi Jendra Sulastra and Prayudha (2014).

On a practical note, accurate item forecasts can lead to selective inventory control and yield reductions in spoilage and associated costs. This application provides practical insights and academic value for research on hybrid model structures and time-based data integration.

4.5 Comparative Analysis

A series of model tests benchmarked ARIMA models fitted on univariate time series for selected item-store pairs. In some cases, ARIMA outperformed deep learning models regarding MAE and RMSE. Nonetheless, the hybrid approach demonstrates the advantages of integrating deep learning and rich feature sets, as discussed in the Literature Review.

4.6 Constraints and Future Work

Although the results are promising, forecast accuracy varies across items and timespans because of data scarcity and instability. Future research should focus on comprehensive hyperparameter tuning, exploring alternative architectures such as attention mechanisms, and assessing model generalizability to other geographical regions.

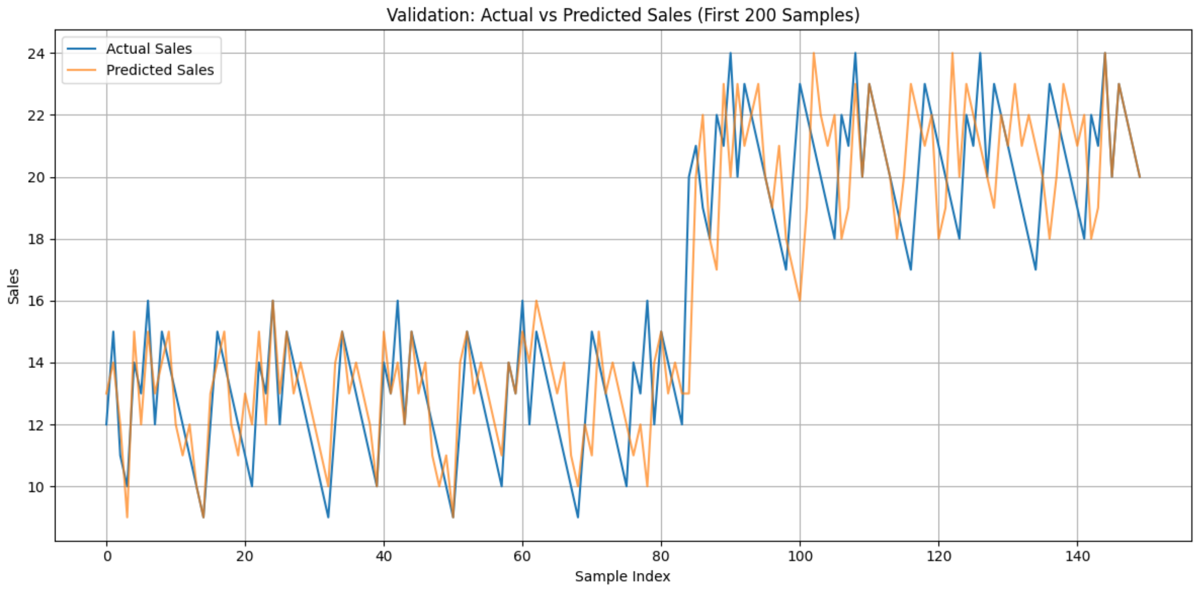

Figure 9: Validation: Actual Vs. Predicted Sales of first 200 samples

The results thus directly address the research question, showing that the proposed model is capable of predicting retail sales more efficiently than baseline approaches. These findings inform methodological improvements and future adoption, carefully weighing both strengths and limitations.

5 Conclusion

This paper has examined the prediction of perishable groceries sales and has utilised a large synthetic dataset augmented with holiday and weather data. The research question was focused on establishing whether higher-level deep learning architectures such as LSTM and CNN-GRU combination models could provide reasonable and operationally concrete sales forecasts to assist inventory management and reduction of waste I Made Dwi Jendra Sulastra and Prayudha (2014).

In general, the experiment executed a memory-efficient pipeline that used multiple data sources and conducted temporal and feature engineering in keeping with best practice guides available in the literature. The models demonstrated the ability to capture seasonal and promotion-driven demand patterns, particularly in high-volume store-item combinations in the top performing cities in Ecuador.

Nevertheless, the study revealed a number of limitations that attenuated the strength of the findings. Although positive predictive performance was attained in most series, some item-store pairs showed low or negative R^2, indicating that the models struggled with noise, sporadic sales, or limited data intervals. This inconsistency reconciles with previous literature regarding the LSTM model sometimes being inferior to other classical or ensemble methods without significant tuning or stronger feature representation Yogapiriyan et al. (2024); I Made Dwi Jendra Sulastra and Prayudha (2014). In addition, the use of synthetic data, while essential for this study, may not capture real-life complexities completely, thus constraining the external validity and generalizability.

The non-random sampling on the top five stores may have contributed to mixed results, as it restricted the breadth and diversity of sales patterns available for model training. Further research could broaden the sampling area and consider low-volume, operationally essential products. Also, model choices (such as limited network depth and the defined stopping criteria) may have restricted model capacity. Increased robustness may be attained by leveraging higher-order feature engineering, using item hierarchies, price dynamics, or customer segmentation.

Notably, there was no clear, consistent improvement attributable to weather data integration compared to integrating holiday and promotion characteristics. This outcome aligns with existing literature, highlighting the powerful explanatory role of calendar-oriented and transactional variables in predicting retail demand. Modeling these fixed drivers ought to be an absolute priority, with weather data only emphasized for highly sensitive products.

Despite these criticisms, this experiment is valuable for illustrating a scalable, practical approach to integrating disparate data streams within a deep learning paradigm for sales forecasting in an emerging economy. It validates hybrid CNN-GRU and LSTM models for capturing complex temporal dynamics when proper feature sequences and contextual signifiers are provided *Favorita Grocery Sales Forecasting Competition* (2018). The modulated pipeline and evaluation procedures provide a blueprint for future exploratory and operational deployments.

Future research should address the identified limitations by incorporating broader datasets, adopting attention-based or transformer architectures, and testing finer time resolutions (e.g., using hourly sales). Improvements may be realized through automated hyperparameter optimization and ensemble learning O'Malley et al. (2019). There is significant commercialization potential for such predictive pipelines in streamlining inventory management, reducing stock-outs and waste, particularly when adaptive restocking decisions are based on high-quality forecasts.

In summary, this project contributes to the research area of time series forecasting for perishable goods retail by implementing state-of-the-art deep learning techniques in a carefully constructed, data-intensive experiment. Although challenges remain with data quality and model generalization, the results are consistent with and extend current knowledge, providing a robust basis for continued refinement and practical real-world application. These findings point to both the potential of hybrid models, as well as the need to consider carefully the importance of interpretability versus cost of computation and external validity in the real world of retail adoption.

References

[1] *Calendarific Global Holidays API* (2024). Web Service.
URL: *https://calendarific.com/api/v2/holidays*

[2] *Favorita Grocery Sales Forecasting Competition* (2018). Kaggle Competition.
URL: *https://www.kaggle.com/competitions/favorita-grocery-sales-forecasting*

[3] Hyndman, R. and Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*, OTexts.
URL: *https://otexts.com/fpp2/*

[4] I Made Dwi Jendra Sulastra, I Putu Oka Wisnawa, I. K. W. and Prayudha, I. P. A.

(2014). Predicting food supply with lstm: A data-driven approach using wasde data, *Proceedings of the 2nd International Conference on Computer Science and Electronics Engineering* .
URL: *https://www.atlantis-press.com/article/126005513.pdf*

[5] Marripudugala, M. (2024). Forecasting perishable food demand and inventory management, *International Research Journal of Modernization in Engineering Technology and Science (IRJMETS)* .
URL: *https://www.irjmets.com/uploadedfiles/paper/issue*$_{7j}$*uly*$_2$024/60329/$final/fin_irjmets$172189159

[6] Oberoi, H. (2024). Corporación favorita grocery sales forecasting. GitHub repository.
URL: *https://github.com/oberoiheman/Corporaci-n-Favorita-Grocery-Sales-Forecasting*

[7] Olome, Y. (2024). Forecasting store sales for corporación favorita, ecuador. LinkedIn Article.
URL: *https://www.linkedin.com/pulse/forecasting-store-sales-corporation-favorita-ecuador-olome*

[8] O'Malley, T. et al. (2019). Keras tuner.
URL: *https://keras.io/keras*$_t$*uner/*

[9] *Open-Meteo Historical Weather API* (2024). Web Service.
URL: *https://open-meteo.com/*

[10] Yogapiriyan, A., Chakravarthi, V. and Jayashri, N. (2024). Preventing food waste: A guide to using long short-term memory systems, *International Research Journal of Modernization in Engineering Technology and Science (IRJMETS)* .
URL: *https://www.irjmets.com/uploadedfiles/paper/issue*$_{4a}$*pril*$_2$024/54664/$final/fin_irjmets$171533701