

Deep Learning Models for Sales Demand Forecasting In Retail Supply Chain Management System

¹Mr.K.Bala Krishna, ²Mr. K. Chiranjeevi,

Assistant Professor , Department of Computer Science & Engineering,
Geethanjali Institute of Science And Technology, Nellore, India

Abstract

Accurate demand forecasting is critical for effective supply chain management, especially in the retail sector where fluctuations in consumer behavior can significantly impact inventory and financial planning. This study aims to improve sales demand forecasting for 1,115 Rossmann stores across Europe by incorporating external variables often overlooked in traditional models. While prior research has primarily relied on historical sales data in a univariate context, our approach treats the problem as a multivariate forecasting task, integrating influential factors such as weather conditions, promotions, store location, and holidays. We propose a deep learning-based model, Sales Demand Forecasting using Weather Data (SDFW), which leverages a Gated Recurrent Unit (GRU) architecture optimized through Grid Search. Comparative analysis shows that SDFW outperforms the widely used Long Short-Term Memory (LSTM) model in terms of forecasting accuracy. The incorporation of weather-related features enables more precise prediction of demand trends, ultimately supporting better inventory management and operational efficiency for Rossmann stores.

Keywords: DL models, Sales demanding forecasting, supply chain management, GRU, LSTM

Introduction

This project reimagines retail demand forecasting by integrating deep learning with contextual weather intelligence to optimize supply chain performance for Rossmann Stores. Moving beyond conventional models that rely on historical sales alone, we introduce a Weather-Enhanced Demand Forecasting Framework powered by Gated Recurrent Units (GRUs) and Grid Search Optimization.

At its core, the project fuses multivariate data—sales, promotions, holidays, store attributes, and weather conditions—into a cohesive forecasting engine capable of learning complex, nonlinear relationships. By incorporating weather as a dynamic variable, the model anticipates demand fluctuations with greater accuracy, adapting to real-world events like snow, rain, or heatwaves that alter customer behavior.

The proposed model surpasses traditional Long Short-Term Memory (LSTM) networks in both performance and efficiency, particularly when applied to segmented data based on weather events. Experimental results show significantly improved forecasting accuracy, enabling more responsive inventory planning, reduced stockouts, and minimized overstocking.

Problem Statement

In the retail industry, demand forecasting plays a critical role in supply chain management, directly influencing financial performance, inventory control, and customer satisfaction. Accurate sales forecasting allows businesses to reduce overstocking and understocking, optimize resources, and enhance operational efficiency.

Rossmann, a leading drugstore chain with over 3,000 stores across 7 European countries, faces a significant challenge in predicting daily sales up to six weeks in advance. Store sales are influenced by multiple internal and external factors, including:

1. Promotional campaigns that drive short-term spikes in sales.

2. Competition from nearby stores affecting customer traffic.
3. School and state holidays influence shopping behavior.
4. Seasonal trends and local events create demand fluctuations

With thousands of individual store managers making sales predictions based on unique local conditions, the accuracy of results varies significantly, leading to inefficiencies in inventory planning and supply chain operations.

OBJECTIVES

Develop a Weather-Enhanced Forecasting Model Design and implement a deep learning model that integrates weather data (temperature, precipitation, visibility, etc.) with sales and operational variables to forecast demand more accurately.

Leverage Gated Recurrent Units (GRUs) for Time-Series Forecasting Utilize GRUs to capture long-term dependencies and nonlinear patterns in multivariate time-series data, offering a superior alternative to traditional and LSTM models.

Optimize Model Performance via Grid Search Apply systematic hyperparameter tuning using Grid Search to identify the best configuration for learning rate, batch size, activation functions, and architecture depth.

Segment Forecasting by Weather Conditions Create tailored forecasting models for different weather events (normal, rain, snow), enabling localized and event-responsive inventory planning.

Minimize Forecasting Error Achieve lower Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) compared to baseline models by enriching the model with environmental context and optimized architecture.

Enable Data-Driven Retail Decision Making Provide Rossmann with a robust forecasting tool to improve stock availability, reduce overstock, and optimize supply chain operations in real time.

Demonstrate Scalable AI Integration in Retail Showcase a replicable AI-driven forecasting framework that can scale across diverse locations and adapt to changing environmental and market conditions.

Literature Review

Demand forecasting has long been recognized as a critical function within supply chain and retail operations. Accurate forecasts enable businesses to optimize inventory, minimize costs, and enhance customer satisfaction. Over the decades, researchers have explored a range of techniques—from traditional statistical models to modern machine learning—to improve the precision of demand predictions.

However, much of the existing literature either treats forecasting as a univariate time series problem or overlooks the impact of external contextual variables, such as weather, promotions, and holidays. Moreover, while deep learning has gained traction, many implementations remain limited by lack of model optimization, exclusion of real-world events, or excessive complexity without scalable benefits.

This literature survey critically examines the evolution of demand forecasting methods, highlights their limitations, and identifies the innovation gap this project addresses. Specifically, it explores how the integration of Gated Recurrent Units (GRU), Grid Search optimization, and weather-based segmentation creates a more accurate, adaptive, and context-aware forecasting framework for the retail environment.

Application of Facebook's Prophet Algorithm for Successful Sales Forecasting Based on Real-World Data

Authors: Emir Žunić, Kemal Korjenić, Kerim Hodžić, Dženana Đonko

This paper introduces a sales forecasting framework based on Facebook's Prophet algorithm, specifically designed for the retail industry using real-world data from one of the largest retailers in Bosnia and Herzegovina. The model incorporates back testing and a classification module that evaluates the forecastability of each product in a portfolio. Prophet's strength lies in handling irregular sales patterns, seasonality, and outliers while delivering robust forecasts at scale.

The framework emphasizes practical deployment by targeting monthly and quarterly forecasts—periods highly relevant for production and inventory planning. Products are ranked based on sales volume, and forecastability is assessed using MAPE and PE metrics. The study demonstrates how Prophet, without complex tuning or exogenous variables, can serve as an effective forecasting tool for retail environments while maintaining scalability and simplicity.

A Sales Prediction Method Based on LSTM with Hyper-Parameter Search Authors: Yun Dai, Jinghao Huang

This paper proposes an LSTM-based demand forecasting model enhanced by a novel loss function and optimized through hyper-parameter search. The method targets challenges in real-world forecasting scenarios such as data sparsity, user-specific forecast preferences (e.g., overestimating to avoid stockouts), and the need for single robust models. The research is conducted using the Rossmann Kaggle dataset, with the model predicting both weekly and monthly sales.

Compared to baseline machine learning models and traditional LSTMs, the optimized LSTM model significantly improves performance, particularly on sparse data. Hyper-parameter tuning (via greedy strategy and sampling) is shown to boost forecasting accuracy while reducing overfitting and model complexity. The study concludes that their LSTM method, particularly with adjusted loss functions, better aligns with business needs in practical retail forecasting scenarios.

Gated Recurrent Unit with Genetic Algorithm for Product Demand Forecasting in Supply Chain

Management Authors: Jiseong Noh, Hyun-Ji Park, Jong Soo Kim, Seung-June Hwang

This paper introduces a hybrid deep learning model called GA-GRU, which combines a Gated Recurrent Unit (GRU) with a Genetic Algorithm (GA) for hyperparameter optimization. The focus is on enhancing forecasting accuracy for supply chain management by optimizing five critical parameters: window size, neuron count, batch size, epoch size, and learning rate. The model is validated through cross-validation, sensitivity analysis, and benchmarking against other methods like ARIMA, LSTM, and ANN.

GA-GRU achieves superior forecasting accuracy, demonstrating robustness and adaptability across real-world retail datasets. The GRU's structure enables fast training and effective handling of sequential patterns, while GA provides global search capabilities for optimal parameter tuning. This hybrid model is especially beneficial in highly dynamic retail settings, where demand patterns fluctuate, and standard tuning methods fall short.

This study explores the application of artificial intelligence (AI) and Machine Learning (ML) techniques, particularly Random Forest classification, to predict car insurance risks using publicly available datasets from Kaggle [7] [8]. By implementing feature extraction and classification methodologies, this research demonstrates the effectiveness of AI-driven predictive models in enhancing risk assessment accuracy and operational efficiency in the insurance sector.

Proposed Model

To overcome the limitations of traditional forecasting methods, the paper introduces a Weather-Enhanced Deep Learning Model for demand forecasting in Rossmann retail stores. The proposed system leverages a Gated Recurrent Unit (GRU) neural network architecture, optimized using Grid Search, to improve forecast accuracy by incorporating multivariate inputs, including weather data, promotions, holidays, and store-level information.

Key Features of the Proposed System:

Multivariate Time Series Forecasting: Instead of relying on univariate sales data, the system integrates various internal and external features, such as:

- 1 Weather conditions (temperature, snow, rain, etc.)
- 2 Holiday indicators (state and school holidays)
- 3 Promotions and promotional duration

4 Store location and competition distance

Deep Learning Architecture – Gated Recurrent Unit (GRU): GRU is a type of recurrent neural network (RNN) well-suited for time series prediction. It captures long-term dependencies and sequential patterns more efficiently than traditional models or even LSTM in some cases. GRU is chosen for its simplicity and computational efficiency compared to LSTM.

1. **Grid Search for Hyperparameter Optimization:** Grid Search is used to fine-tune the hyperparameters of the GRU model (e.g., number of units, activation functions, batch size, epochs). This ensures optimal performance by systematically testing combinations of parameters.
2. **Weather-Based Subset Analysis:** The dataset is further classified into subsets based on weather events (e.g., normal, rain, snow). Forecasting is performed separately on each subset using the GRU model, allowing for more accurate, condition-specific predictions.
3. **Preprocessing and Feature Engineering:** Includes handling missing values, encoding categorical variables, and generating new features such as mean/median sales per customer. Scaling is done using Min-Max normalization for model compatibility and performance.
4. **Model Evaluation:** Performance is measured using Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). The proposed GRU model demonstrates lower error rates and better generalization compared to the LSTM baseline.

Methodology

The methodology of the Weather-Enhanced Deep Learning Model focuses on accurately predicting daily sales for Rossmann stores by combining historical sales data with weather information. In the retail world, customer behavior is often influenced by external factors such as temperature, rainfall, or public holidays. By including weather data, the model becomes more intelligent and better at handling real-world conditions that affect sales.

This approach uses deep learning techniques like LSTM or GRU networks, which are well-suited for time-series forecasting. The model learns from past patterns, store-specific details, promotions, and local weather to make future predictions. The methodology also includes steps such as data collection, preprocessing, feature engineering, model training, evaluation, and deployment.

The goal is to build a system that not only predicts sales accurately but also helps Rossmann stores make better decisions around inventory, staffing, and marketing, especially during weather-sensitive periods.

System Architecture

The architecture of this forecasting system is structured as a multi-stage pipeline that begins with data acquisition and ends with predictive deployment. The process starts with gathering various data sources, including historical sales data from Rossmann stores, store metadata (like store type, location, and assortment), promotional schedules, calendar data (such as holidays and school vacations), and historical weather data for each store's geographic location. Weather data is collected via APIs or open datasets from sources like OpenWeatherMap or NOAA, and includes variables such as temperature, precipitation, humidity, and general weather conditions. Once the data is collected, an ETL (Extract, Transform, Load) process is employed to clean, normalize, and merge the datasets into a single, time-aligned table.

In the preprocessing and feature engineering phase, raw data is transformed into meaningful inputs. This includes generating time-based features (like day of the week, month, or whether a day is a holiday), calculating lagged sales and weather metrics (such as average sales over the past week), encoding categorical variables (such as store type or weather conditions), and normalizing numerical data. The result is a rich, structured dataset that reflects both temporal patterns and external influences like weather.

The processed data is then passed into the model input layer, where it's organized into tensors suitable for training. If the model is sequential (such as an LSTM or GRU), data is arranged into time-series windows to capture temporal dependencies. For more tabular-based models, features are organized per instance. Categorical features are typically embedded into dense vectors, while continuous features are fed directly into the model. These features are then processed by the core deep learning architecture.




Fig.1. System Architecture

DATA READING & DATA INTEGRATION

Data reading and integration form one of the most critical layers in the system architecture of the weather-enhanced deep learning model. This stage is responsible for bringing together different types of data from various sources and aligning them in a usable format for the model.

The system reads data from multiple key sources:

Sales data is fetched from Rossmann's internal databases, which include daily product-level sales for each store.

Weather data is collected using external APIs such as OpenWeatherMap or Meteostat, providing temperature, rainfall, wind speed, and other climate indicators.

Calendar data includes holidays, weekends, and school vacations, which often influence shopping behavior.

Store metadata contains information like store type, size, location, and whether any promotions are active.

DATA PREPROCESSING

Data preprocessing is a key step in the system architecture, as it prepares raw data for training and prediction. In the context of the weather-enhanced deep learning model for Rossmann stores, data preprocessing ensures that the combined sales, weather, and store-related data is clean, consistent, and useful for the model.

After data is read and integrated from multiple sources (like sales records, weather APIs, and calendar files), it often contains missing values, inconsistent formats, or noise that can negatively affect model performance. The

preprocessing step handles these issues by applying several techniques to transform raw data into structured and meaningful inputs.

Main tasks in preprocessing include:

Handling missing data: Filling or removing missing values in sales or weather columns using techniques like forward fill, backward fill, or mean imputation

Encoding categorical data: Converting store types, promo flags, and holiday markers into numerical format using label encoding or one-hot encoding.

Normalizing/Scaling: Adjusting values like temperature or sales so they fall within a similar range. This helps deep learning models train faster and more accurately.

Feature engineering: Creating new useful inputs such as: Lag features (e.g., sales from the past 7 days)

Rolling averages (e.g., average sales over the last month)

Weather categories (e.g., sunny, rainy, stormy)

Special day flags (e.g., weekend, holiday, school break)

Time alignment: Making sure that the weather, sales, and calendar data all match correctly by date and store ID.

Data filtering: Removing outliers or periods where stores were closed (no sales) to avoid misleading patterns.

DATASETS

The effectiveness of the Weather-Enhanced Deep Learning Model heavily depends on the quality and variety of datasets it uses. In the system architecture, multiple datasets are integrated to provide the model with a complete view of all the factors that can influence store sales. These datasets come from both internal sources (like sales records) and external sources (like weather APIs and public holidays).

1. Sales Dataset:

1. It contains historical daily sales for each Rossmann store.
2. Key columns include :Store ID, Date, Sales, Customers, Promo, Open/Closed.
3. This is the primary dataset used to train the model to learn sales patterns.

Store Dataset:

1. It contains static information about each store.
2. Columns include Store ID, Store Type, Assortment Type, Competition Distance, Promo2.
3. Used to give model context about each store's environment and behavior.

Weather Dataset:

Pulled from external APIs like OpenWeatherMap or Meteostat.

Includes daily weather conditions like Temperature, Humidity, Rainfall, Wind Speed, and Weather Description.

Mapped to each store based on its location to help the model learn how weather affects customer traffic and sales.

Calendar/Holiday Dataset:

It contains public holidays, school breaks, weekends, and special events. Helps the model understand unusual spikes or drops in sales that occur due to non-regular patterns.

Promotions Dataset:

1. Tracks when and where promotions were active.
2. Includes: Promo, Promo Interval, Promo2 Since Week/Year.
3. This allows the model to learn how marketing campaigns impact sales volumes.

Combined Dataset:

All the above datasets are merged together using common columns like Store ID and Date. The result is a **single, feature-rich dataset** where every row represents one store on one specific day with all related sales, weather, and promotion data.

SYSTEM IMPLEMENTATION

The implementation of a deep learning model enhanced by weather data for retail supply chains follows a pipeline-based architecture. It involves developing an automated, scalable system that can continuously ingest data, train the model, generate forecasts, and deliver predictions for decision-making in store operations and inventory planning.




Fig.2. System Modules

Performance Metrics

Mean Absolute Error (MAE)

It calculates the average difference between the calculated values and actual values. It is also known as scale-dependent accuracy as it calculates error in observations taken on the same scale. It is used as evaluation metrics for regression models in machine learning. It calculates errors between actual values and values predicted by the model. It is used to predict the accuracy of the machine learning mode

$$= \frac{1}{n} \sum_{i=1}^n |d_i - p_i| \quad (2)$$

RMSE is a square root of value gathered from the mean square error function. It helps us plot a difference between the estimate and actual value of a parameter of the model. Using RSME, we can easily measure the efficiency of the model.

RSME is a square root of the average squared difference between the predicted and actual value of the variable/feature. Let's see the following formula.

$$= \sqrt{\frac{1}{n} \sum_{i=1}^n (d_i - p_i)^2} \quad (3)$$

Where

Σ - It represents the "sum".

d_i - It represents the predicted value for the i^{th}

p_i - It represents the predicted value for the i^{th}

n - It represents the sample size.

R² Score (Coefficient of Determination)

The R² score, also known as the coefficient of determination, is a statistical measure used to evaluate the goodness-of-fit of a regression model. It quantifies how well the independent variable(s) explain the variance in the dependent variable.

Mathematically, R² is defined as:

$$R^2 = 1 - \frac{\text{SS}_{\text{res}}}{\text{SS}_{\text{tot}}}$$

Results & analysis




Fig.3. Homepage of Sales Demand Forecasting

	Store	DayOfWeek	Date	Sales	Customers	Open	Promo
1	S	5	2015-07-31	1045	925	1	1
2	S	5	2015-07-31	1004	825	1	1
3	S	5	2015-07-31	1014	927	1	1
4	S	5	2015-07-31	11995	1486	1	1
5	S	5	2015-07-31	1402	109	1	1
6	S	5	2015-07-31	1401	985	1	1
7	S	5	2015-07-31	11346	1414	1	1

Fig.4. Sample Training Data of Sales Demand Forecasting




Fig.5. Uploading the Test Data




Fig.6. GRU Model Sales Demand Predictions

Sample 10 Predictions		Model Metrics	
Date	Predicts	Metric	Value
17/09/2015 00:00	1520.6052246094	Model	GRU
17/09/2015 00:00	1673.2143554688	Train_MSE	746592.375
17/09/2015 00:00	1350.8774414062	Test_MSE	792950.5
17/09/2015 00:00	1421.7761230469	Train_RMSE	864.0557708
17/09/2015 00:00	1397.8916015625	Test_RMSE	890.4776808
17/09/2015 00:00	1416.6967773438	Train_MAE	646.5524292
17/09/2015 00:00	2148.4467773438	Test_MAE	647.2113037
17/09/2015 00:00	1824.4196777344	Train_R2	0.922469676
17/09/2015 00:00	2548.320214844	Test_R2	0.917825282
17/09/2015 00:00	2015.9808349609		

Fig.7. Sample Prediction & Model Metrics of GRU

Conclusion

The study concludes that accurate demand forecasting is crucial for the retail industry, particularly in optimizing inventory and meeting customer needs. It emphasizes the significance of using a multivariate approach—factoring in external and internal variables like weather, promotions, store locations, and holidays—rather than relying only on univariate sales data. The proposed model, a Gated Recurrent Unit (GRU) enhanced with Grid Search, outperforms the previously used Long Short-Term Memory (LSTM) models in accuracy. The GRU model demonstrated particularly better performance when weather-related subsets (like normal, rain, and snow) were separately considered, suggesting that environmental factors significantly influence sales. The key takeaway is that this weather-enhanced deep learning model can serve as a more reliable tool for sales demand forecasting in retail supply chain management, especially for Rossmann stores.

Future Scope

The weather-enhanced deep learning model for Rossmann demand forecasting shows strong predictive performance. Future research can extend its use to weather-sensitive industries like agriculture, fashion, and tourism to test generalizability. Enhancing the granularity of weather inputs—adding UV index, wind chill, and alerts—could boost accuracy. Real-time data integration would enable dynamic forecasting and operational agility. Replacing Grid Search with meta-heuristics like Genetic Algorithms, Bayesian Optimization, or Particle Swarm Optimization may yield better hyperparameter tuning. Deeper integration with supply chain systems (inventory/logistics) can transform the model into a prescriptive tool. Incorporating explainable AI methods would clarify feature impacts and improve stakeholder trust. Addressing missing store closure data remains critical for robustness.

References

1. D .Ni ,Z. Xiao, and M.K. Lim, “ A Systematic review of the research trends of machine learning in supply chain management,” *Int. J. Mach. Learn. Cybern.*, vol. 11, no. 7, pp. 1463–1482, Jul. 2020.
2. S. Punia, K. Nikolopoulos, S. P. Singh, J. K. Madaan, and K. Litsiou, “Deep learning with long short-term memory networks and random forests for demand forecasting in multi-channel retail,” *Int. J. Prod. Res.*, vol. 58, no. 16, pp. 4964–4979, Aug. 2020.
3. M. M. Taye, “Understanding of machine learning with deep learning: Architectures, workflow, applications and future directions,” *Computers*, vol. 12, no. 5, p. 91, Apr. 2023.
4. J. Noh, H.-J. Park, J. S. Kim, and S.-J. Hwang, “Gated recurrent unit with genetic algorithm for product demand forecasting in supply chain management,” *Mathematics*, vol. 8, no. 4, p. 565, Apr. 2020.
5. H.Hewamalage, C.Bergmeir, and K. Bandara, “Recurrent neural networks for time series forecasting: Current status and future directions,” *Int. J. Forecasting*, vol. 37, no. 1, pp. 388–427, Jan. 2021.
6. M. S. Amalnick, N. Habibifar, M. Hamid, and M. Bastan, “An intelligent algorithm for final product demand forecasting in pharmaceutical units,” *Int. J. Syst. Assurance Eng. Manage.*, vol. 11, no. 2, pp. 481–493, Apr. 2020.
7. Ramesh Chandra AdityaKomperla, “CAR INSURANCE RISK PREDICTION USING ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING, 2022 | *Global Journal of Advanced Engineering Technologies and Sciences*, 9(4), 1- <https://doi.org/10.29121/gjaets.2022.04.01>
8. Ramesh Chandra AdityaKomperla , “OPTIMIZING HEALTHCARE INSURANCE PREDICTIONS THROUGH DEEP LEARNING AND SCO-ENHANCED HYPERPARAMETER TUNING”. (2022). *Global Journal of Advanced Engineering Technologies and Sciences*, 9(11), 1- <https://doi.org/10.29121/gjaets.2022.11.01>