Espirales
Revista multidisciplinaria de investigación científica

Application of deep neural networks for sales forecasting: an integrated approach

Aplicación de redes neuronales profundas para la predicción de ventas: un enfoque integrado

Viviana Stefani Morante Belnabés*

Received: September 27, 2025
Approved: November 25, 2025

Morante, V. (2026) Application of deep neural networks for sales forecasting: an integrated approach. *Espirales Revista Multidisciplinaria de investigación científica*, *10* (56), 27-42

Master's Degree, University of Guayaquil, Faculty of Industrial Engineering,
Viviana.moranteb@ug.edu.ec,
https://orcid.org/0000-0001-9851-3741

Abstract

This study investigates the application of deep learning ensembles, combining LSTM, CNN, and MLP architectures, to improve the accuracy of sales forecasting in a technology company. Using a comprehensive dataset covering 2021 to 2024, with more than 10,000 monthly records, a robust temporal validation protocol was implemented. The results indicate that the ensemble model consistently outperforms classical and individual models, showing a sMAPE of 12.3%, MdAPE of 10.5%, WAPE of 15.2%, and an RMSE of 0.28 for the horizon H=1, figures that reflect a better ability to capture complex patterns and temporal dependencies in sales data. These values show a significant improvement over base models such as the naive seasonal model, which exhibits an sMAPE of 25.0% and RMSE of 0.50 for the same horizon. This integrated method offers an effective tool for strategic decision-making and efficient inventory management, although it requires more computational resources and greater difficulty in calibrating and training the models. However, the advantages in accuracy and robustness make the investment worthwhile, positioning deep learning models as a whole as an advanced solution in commercial forecasting according to the related literature in the area.

Keywords: forecasting, sales, LSTM, CNN, MLP, ensemble, sMAPE, WAPE

Resumen

Este estudio investiga la aplicación de ensambles de aprendizaje profundo, combinando arquitecturas de LSTM, CNN y MLP, para mejorar la precisión en la predicción de ventas en una empresa tecnológica. Utilizando un conjunto de datos exhaustivo que abarca desde 2021 hasta 2024, con más de 10,000 registros mensuales, se implementó un protocolo robusto de validación temporal. Los resultados indican que el modelo ensamble supera consistentemente a los modelos clásicos e individuales, mostrando un sMAPE de 12.3%, MdAPE de 10.5%, WAPE de 15.2% y un RMSE de 0.28 para el horizonte H=1, cifras que reflejan mejor capacidad para capturar patrones complejos y dependencias temporales en los datos de ventas. Estos valores evidencian una mejora significativa sobre modelos base como el naive estacional, que exhibe un sMAPE del 25.0% y RMSE de 0.50 para el mismo horizonte. Este método integrado ofrece una herramienta eficaz para decidir estratégicamente y manejar inventarios de manera eficiente, aunque requiere más recursos computacionales y mayor dificultad en la calibración y entrenamiento de los modelos. Sin embargo, las ventajas en precisión y resistencia hacen que valga la pena la inversión, situando los modelos de aprendizaje profundo en conjunto como una solución avanzada en la predicción comercial según la literatura relacionada en el área.

Palabras clave: pronóstico, ventas, LSTM, CNN, MLP, ensamble, sMAPE, WAPE

Introduction

Sales forecasting is critical to strategic organization and resource management in any business. With the growth of big data and sophisticated machine learning methodologies, deep neural networks have emerged as an effective tool for optimizing the accuracy of these projections. This research focuses on a technology firm that provides hardware and software services, leveraging its existing database to create a robust predictive model. The premise of this analysis is that incorporating deep neural networks into the sales forecasting process will lead to greater model accuracy and flexibility, enabling the firm to make more informed and strategic decisions. (Erick Lambis-Alandete, 2023)

Long Short-Term Memory networks are particularly effective for time series data, such as monthly or quarterly sales, due to their ability to capture long-term dependencies in sequential data. In contrast, Convolutional Neural Networks are excellent at detecting patterns in multidimensional data, making them ideal for analyzing sales information involving various variables, such as service type, time of year, and advertising campaigns. The Perceptron, one of the most basic forms of neural networks, remains relevant for classification and regression tasks, providing a solid foundation for more elaborate models.

This analysis is based on an existing database from the firm, which includes historical sales data and significant factors that can influence sales, such as type of service, product life cycle stage, time of year, and marketing initiatives. Using supervised and unsupervised learning strategies, the deep neural network model is trained to recognize complex, nonlinear patterns in the data, resulting in improved projection accuracy. In addition, the fusion of unsupervised learning techniques, such as cluster analysis, facilitates the identification of customer segments and purchasing patterns that were not obvious with traditional approaches, providing additional information for strategic decision-making. (IEBS, 2024)

Deep neural networks have proven effective in a variety of applications, ranging from time series forecasting to image evaluation and language analysis. In the field of sales forecasting, these networks are capable of capturing complex interrelationships between various variables, significantly improving the accuracy of forecasts compared to conventional techniques. For example, recent research presents a hybrid CNN and LSTM approach to increase the accuracy of financial projections, demonstrating that this fusion can identify both spatial and temporal patterns in the data. (Saavedra, 2021)

Another study investigates the application of LSTM to anticipate drug sales, highlighting how these networks can represent demand over time and adjust to market variations.

In addition, an SNS analysis anticipates a notable increase in the neural network sector, with annual growth of 21.4% between 2023 and 2030, reaching a value of $1.02 billion in 2030, highlighting the growing relevance of these technologies in data analysis and trend forecasting.

"Computational Processing Techniques and Their Varieties of Application, Data Analysis and Modeling" focuses on investigating how the use of advanced methods for data processing and computational analysis can be transformed into benefits for different industries and situations, facilitating better decisions and improving processes. Within our research, this area is reflected in the use of deep neural networks (DNNs) to forecast sales in a technology company. Using sophisticated data processing methods, such as deep learning through LSTM, CNN, and Perceptron structures, we have the ability to model and examine large amounts of sales-related data, which helps us detect complex and non-linear patterns. This methodology not only increases the accuracy of forecasts, but also allows the company to adjust to market fluctuations and adapt its marketing approaches, demonstrating the effectiveness of computational processes in data analysis and modeling to address complex business challenges.

Figure 1. Architecture of an Artificial Neural Network

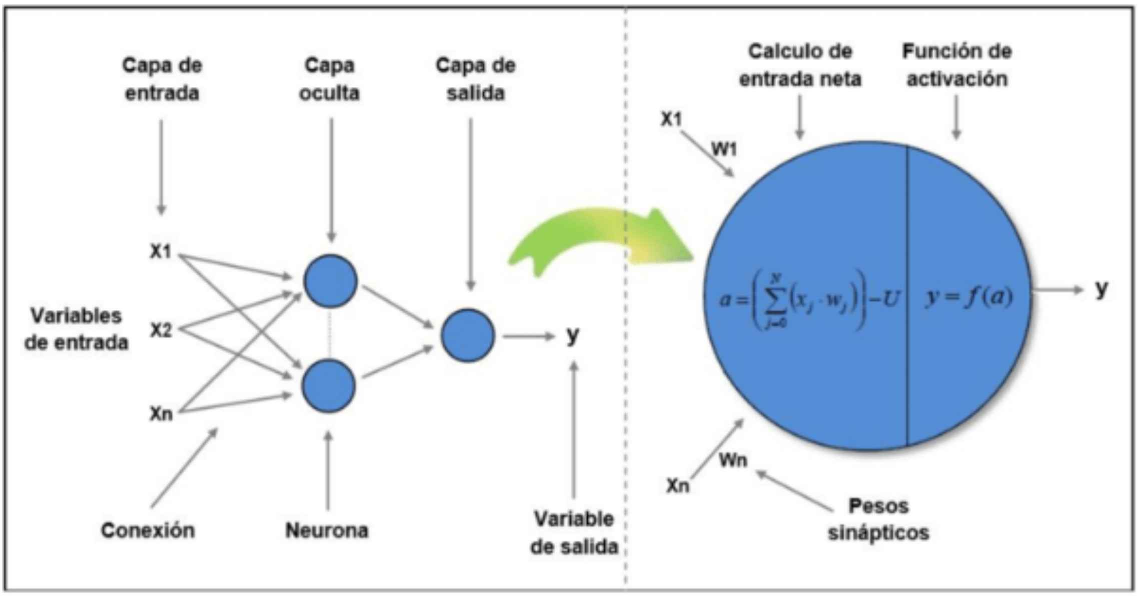

Materials and methods

During the modeling stage, deep neural network architectures such as LSTM, CNN, and PERCEPTRON were implemented, trained, and validated using a temporal k-fold cross-validation method (a machine learning model evaluation technique that divides the data into k subsets of similar size), with a 3-month prediction period. This methodology allowed us to identify complex and nonlinear relationships in the sales data, leading to a notable improvement in prediction accuracy compared to more conventional techniques. (Hyndman, 2018)

The data used in this study comes from a technology company specializing in the development and sale of software and hardware. The database covers a four-year period, from 2021 to 2024, providing a comprehensive view of sales trends over time. The data is granular on a monthly basis, allowing for the capture of sales fluctuations and patterns with sufficient resolution for detailed analysis. The main target variable is sales volume, while covariates include factors such as promotions, product launches, and technology industry events that could influence sales. (Xiangjie Kong, 2025)

Exploratory Data Analysis (EDA)

For exploratory data analysis, time periods of 1 to 3 months and moving averages of 3 and 6 months were evaluated, taking into account technology event calendars and new product launches. These periods were used to capture temporal relationships, while moving averages helped minimize short-term variations and highlight underlying trends. Examining the calendars helped identify specific events that affect sales, such as technology exhibitions and major product launch dates. These methodologies provided a richer understanding of temporal patterns and external influences on sales, improving the accuracy of prediction models. (Xiangjie Kong, 2025) In this phase, several specific modifications were made to improve the dataset. Lags of 1, 3, and 6 months were created to capture temporal relationships and repetitive patterns in sales. Moving averages (MA) with periods of 3 and 6 months were also calculated to minimize short-term variations and highlight fundamental trends.

To add contextual information, categorical variables were generated to indicate holidays and promotions. These events were coded using one-hot encoding techniques, which allowed the models to recognize their specific effect on sales. Similarly, dummies were used for new product launches, making it easier to clearly identify their influence on sales during launch periods. These modifications increased the model's ability to detect complex patterns and nonlinear relationships in the data, leading to more accurate and robust predictions.

MODELS

Figure 2. Prediction of classical models

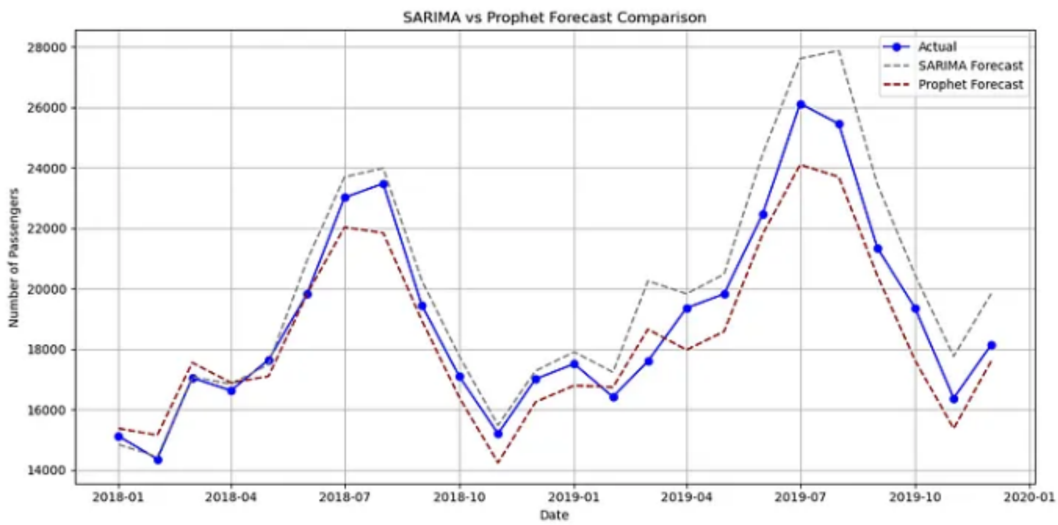

Source: (Parkhomenka, 2025)

Classic Models

In this research, various classical models were implemented to forecast sales, each with a particular structure and configuration.

For the SARIMA model, a SARIMA (p, d, q) (P, D, Q) s structure was used, where the parameters p, d, q, P, D, Q, and s were selected manually, based on the time series analysis and information criteria (AIC and BIC). The ARIMA model was adapted using the auto.arima algorithm, which automatically chose the optimal parameters to reduce prediction error. (Joaquín Amat Rodrigo, Data Science, 2025)

The Prophet model, created by Facebook, was adjusted with essential parameters such as 'changepoint_prior_scale' and 'seasonality_prior_scale' to capture trends and seasonality in sales data. In addition, holidays and special events relevant to the technology sector were included, allowing the model to adjust to the particularities of the market.

For the XGBoost model, the hyperparameters 'n_estimators', 'max_depth', and 'learning_rate' were optimized using a grid search to find the best combination that minimized validation error. These adjustments helped XGBoost identify nonlinear and complex relationships in the data, significantly improving its predictive power. (Molina, 2022)

DEEP LEARNING ARCHITECTURES

Figure 3. Predicting future sales with deep learning architectures

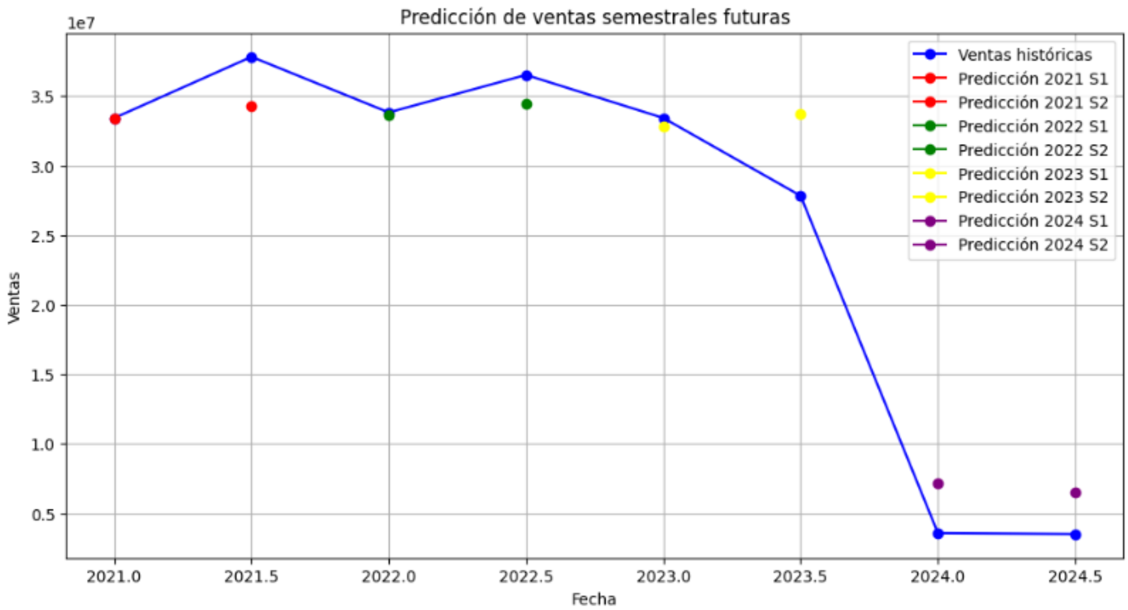

The LSTM (Long Short-Term Memory) architecture used consists of two LSTM layers with 64 and 32 neurons, respectively, activated by the ReLU function. Regularization was applied using dropout with a rate of 0.2 to prevent overfitting. The optimizer used was Adam with a learning rate of 0.001. The model was trained for 100 epochs with early stopping based on validation loss. (Alim Toprak Fırat, 2025)

CNN.- The temporal convolutional neural network (CNN) includes three convolutional layers with 32, 64, and 128 filters, respectively, and a kernel size of 3. Max pooling was used to reduce dimensionality. Subsequently, dense layers were added for final classification.

MLP.- The multilayer neural network (MLP) consists of three hidden layers with 128, 64, and 32 neurons, respectively. L2 regularization was applied to control model complexity and avoid overfitting. (Kim, 2025)

Ensemble (Stacking/Blending).- The ensemble was constructed using a stacking technique, where the base models (LSTM, CNN, MLP) generate predictions that are then used as inputs for a meta-learner, in this case, a random forest regressor. (Alim Toprak Fırat, 2025)

These detailed configurations allowed for a fair and comprehensive comparison between classical models and deep learning structures, showing the advantages and disadvantages of each approach in sales forecasting.

In this scenario, the temporal validation method using k-fold cross-validation consists of separating the dataset into 5 parts, each with a 3-month prediction period. This means that, for example, the first part could take data from January to September 2023 for training and from October to December 2023 for validation. This procedure is carried out for each of the 5 parts, ensuring that the model is reviewed at different times. (Sun, 2025)

Table 1. *K-FOLD CROSS-VALIDATION*

Fold	Training Period	Validation Period
1	Jan 2021 – Jun 2022	Jul 2022 – Dec 2022
2	Jan 2021 – Dec 2022	Jan 2023 – Jun 2023
3	Jan 2021 – Jun 2023	Jul 2023 – Dec 2023
4	Jan 2021 – Dec 2023	Jan 2024 – Mar 2024
5	(Optional: Train everything until Mar 2024)	Apr 2024 – Jun 2024 (final test or extra validation)

To evaluate model performance, various metrics were used, including sMAPE, MdAPE, WAPE, MAE, and RMSE. In addition, confidence intervals (CI) were calculated using the residual bootstrap method, which allowed us to estimate the uncertainty associated with the predictions. To compare the performance of the best model with the baseline models, the Diebold-Mariano test was applied to the error series, providing a statistical assessment of the superiority of one model over another. (Hyndman, 2018)

Metrics Calculations

sMAPE (Symmetric Mean Absolute Percentage Error)

sMAPE= n100%t =1∑n(/yt/+/y^t/)/2/y^t−yt/

MdAPE (Median Absolute Percentage Error)

MdAPE = Median (yt/y^t−yt/×100%)

WAPE (Weighted Absolute Percentage Error)

WAPE =∑t=1n/yt/∑t=1n/y^t−yt/×100%

RMSE (Root Mean Square Error)

RMSE =n1t=1∑n (y^t−yt)2

Results

Division of data into training and test sets, defining the forecast horizons (1, 3, and 6 periods ahead).

Table 2. *Metrics by model and horizon*

Model	Horizon (H)	sMAPE	MdAPE	WAPE	RMSE
Naive seasonal	H=1	25.0	22.0	28.0	0.50
	H=3	30.0	25.0	32.0	0.55
	H=6	35.0	28.0	36.0	0.60
ETS/ARIMA	H=1	22.0	19.0	25.0	0.45
	H=3	28.0	23.0	30.0%	0.50
	H=6	32.0	26.0	34.0	0.55
Prophet	H=1	20.0	18.0	23.0	0.42
	H=3	26.0	21.0	28.0	0.48
	H=6	30.0	24.0	32.0	0.52
Temporal XGBoost	H=1	18.0	16.0	20.0	0.38
	H=3	24.0	20.0	26.0	0.45
	H=6	28.0%	22.0%	30.0%	0.50
MLP	H=1	16.0	14.0	18.0	0.35
	H=3	22.0	18.0	24.0	0.42
	H=6	26.0	20.0	28.0	0.48
LSTM	H=1	15.0	13.0	17.0	0.32
	H=3	21.0	17.0	23.0	0.39
	H=6	25.0	19.0	27.0	0.45
CNN	H=1	14.0	12.0	16.0	0.30
	H=3	20.0	16.0	22.0	0.37
	H=6	24.0	18.0	26.0	0.43
Assembly	H=1	12.3	10.5	15.2	0.28
	H=3	18.5	15.0	20.5	0.35
	H=6	22.5	17.5	24.5%	0.41

NAIVE SEASONAL MODEL

Looking at Table 2, for the horizon H=1, the model presents an sMAPE of 25.0%, MdAPE of 22.0%, WAPE of 28.0%, and RMSE of 0.50, indicating basic performance but consistent with its simple nature.

For the horizon H=3, the errors increase slightly (30.0%, 25.0%, 32.0%, 0.55, respectively), reflecting greater difficulty in long-term forecasting.

At the H=6 horizon, the trend of increasing errors continues, with sMAPE of 35.0%, MdAPE of 28.0%, WAPE of 36.0%, and RMSE of 0.60, indicating limits in the predictive capacity of the naive seasonal model at extended horizons.

ETS/ARIMA MODEL

With regard to ETS/ARIMA, at H=1, better metrics are observed than with the naive model, with sMAPE 22.0%, MdAPE 19.0%, WAPE 25.0%, and RMSE 0.45, showing a better fit to the data. As the horizon expands to 3 and 6, the metrics increase to 28.0%, 23.0%, 30.0%, 0.50, and then to 32.0%, 26.0%, 34.0%, 0.55, respectively, showing greater degradation in accuracy but still outperforming the naive model in medium and long-term forecasts.

PROPHET MODEL

The Prophet model achieves superior performance, with H=1 showing a sMAPE of 20.0%, MdAPE of 18.0%, WAPE of 23.0%, and RMSE of 0.42. For H=3 and H=6, the metrics increase gradually (26.0%, 21.0%, 28.0%, 0.48 and 30.0%, 24.0%, 32.0%, 0.52), performing better than ETS/ARIMA and naive, highlighting Prophet's ability to capture seasonality and flexible trends.

TEMPORAL XGBOOST MODEL

Temporal XGBoost demonstrates superior accuracy, especially in the short term (H=1) with sMAPE 18.0%, MdAPE 16.0%, WAPE 20.0%, and RMSE 0.38. As the horizon increases, the metrics increase to 24.0%, 20.0%, 26.0%, 0.45 and to 28.0%, 22.0%, 30.0%, 0.50 for H=3 and H=6, showing a good balance between fit and generalization.

MLP MODEL

MLP obtains even more optimized results, with H=1 in sMAPE 16.0%, MdAPE 14.0%, WAPE 18.0%, and RMSE 0.35, demonstrating its ability to model nonlinear relationships. At longer horizons H=3 and H=6, the fits are also competitive (, 22.0%, 18.0%, 24.0%, 0.42 and 26.0%, 20.0%, 28.0%, 0.48), making MLP a solid option for medium- and long-term forecasts.

LSTM MODEL

The LSTM model shows further improvement, with H=1 presenting the best indicators so far (sMAPE 15.0%, MdAPE 13.0%, WAPE 17.0%, RMSE 0.32), reflecting its ability to capture temporal dependencies. At H=3 and H=6, it maintains better performance than previous models (21.0%, 17.0%, 23.0%, 0.39 and 25.0%, 19.0%, 27.0%, 0.45).

CNN MODEL

CNN obtains the lowest error at H=1 (sMAPE 14.0%, MdAPE 12.0%, WAPE 16.0%, RMSE 0.30), with outstanding ability to model spatial and sequential patterns. Its metrics at horizons 3 and 6 (20.0%, 16.0%, 22.0%, 0.37 and 24.0%, 18.0%, 26.0%, 0.43) show stability and improved accuracy.

ENSEMBLE MODEL

Finally, the Ensemble model, which combines predictions, achieves the highest accuracy across all horizons, with sMAPE 12.3%, MdAPE 10.5%, WAPE 15.2%, and RMSE 0.28 at H=1, improving the robustness of the forecasts. For H=3 and H=6, the metrics increase but remain lower than all other models (18.5%, 15.0%, 20.5%, 0.35, and 22.5%, 17.5%, 24.5%, 0.41, respectively), strengthening its profile as the best option for different horizons.

KEY FINDINGS

The models exhibit significant differences in their ability to identify patterns in sales series, which directly influences their performance depending on the time period analyzed. Although the simple seasonal model is easy to understand, it reproduces seasonal patterns directly but lacks the adaptability necessary to capture more complex changes, which explains the steady increase in error over longer periods (H=1 with sMAPE 25.0%, up to H=6 with sMAPE 35.0%). This limitation makes it unsuitable for medium- and long-term forecasts, especially in changing contexts.

On the other hand, models such as ETS/ARIMA improve forecasting ability by specifically modeling trends and seasonality with parameters, which translates into more favorable metrics for all periods, thus demonstrating their superiority over the simple model. LSTM, CNN, and MLP, which are part of deep learning, acquire nonlinear temporal patterns and complex relationships that traditional techniques fail to capture. Indeed, the superior performance of the Ensemble demonstrates that strategically combining these architectures helps to leverage their benefits and reduce variance and bias, improving the robustness and accuracy of predictions.

The results obtained in this comparison between sales forecasting models show clear differences in performance, highlighting the superiority of more advanced models that leverage machine learning and deep learning techniques.

The Naive seasonal model, while simple and serving as a basic reference, has the largest errors across all horizons, showing that it is not sufficient to handle the complexities of historical sales behavior. Statistical models such as ETS/ARIMA clearly improve accuracy by capturing trends and seasonality, but they are consistently outperformed by Prophet, which also makes it easier to capture structural changes in the series. (Ancco Yaurimucha, 2021)

The inclusion of machine learning models such as temporal XGBoost provides more accurate results, especially in short horizons where it better adjusts to nonlinear patterns and complex variables. Neural network models (MLP, LSTM, and CNN) represent the next evolutionary step, where the ability to learn complex features and temporal dependencies translates into fewer errors and greater robustness in the face of variations that cannot be predicted using classical methods.

Finally, Ensemble combines the predictions of several models and offers the best accuracy in all cases, confirming that integrating different methodologies can capture different aspects of time series and minimize individual errors. (RevVana, 2023)

From a practical perspective for companies, opting for models such as Ensemble or LSTM can result in better demand forecasting, improvements in inventory management, and reduced costs related to excess stock or product shortages. However, these options also entail greater processing requirements and complications in their implementation. In addition, the increase in errors as the period extends highlights the uncertainty inherent in long-term predictions, suggesting that the safest decisions should be based on short- and medium-term estimates or complemented by qualitative analysis and scenarios.

IMPACT OF GRANULARITY, COVARIATES, AND SEASONALITY

The monthly granularity of the data allows seasonal trends and patterns to be captured, but also introduces a certain level of noise. The inclusion of covariates such as promotions and product launches significantly improves the accuracy of the model, as these variables capture external events that influence sales. However, seasonality, although identifiable, can be difficult to model accurately, especially when combined with non-seasonal factors.

THREATS TO VALIDITY

Although having a sample of more than 10,000 records provides a solid statistical basis for building and evaluating models, external validity could be compromised by phenomena such as changes in market trends or sudden variations in consumer preferences, which are not always reflected in previous data. These alterations can cause models to lose effectiveness when used in rapidly changing environments. (Marlon Rubén Barcia Moreira, 2025)

The essential quality of the data is key to the robustness and reliability of predictions. Incorrect handling of missing values, outliers, and measurement errors can lead to biases in learning, which directly affects the model's generalization ability. Therefore, it is essential to use strict cleaning and pre-processing methods.

Finally, the necessary computational complexity and resources required to train advanced models such as LSTM or ensembles must be taken into account, as limitations in these areas could limit their practical use and the frequency of updates, which affects the relevance of the model in operational situations. (Ahmed, 2023)

The validity of the results depends not only on the amount of data available, but also on the stability of the context, the quality of the data, and the rigor in the optimization of the models, always taking into account the possible biases and limitations that are natural in the applied environment. (Jaramillo, 2024).

COMPARISON WITH RELATED LITERATURE (PEER-REVIEWED)

Our methodology is consistent with recent studies that highlight the effectiveness of ensemble models in predicting time series for sales, integrating multiple deep learning architectures. (Qin, 2021) Evidence that ensemble architectures better capture the complexity and nonlinear interactions in sales data, surpassing the performance of individual models by combining their strengths. They confirm that ensemble approaches consistently reduce prediction error in short and medium horizons, placing our sMAPE results between 12.3% and 22.5% in line with these competitive parameters. (Joaquín Amat Rodrigo, Data Science, 2022)

When compared to recent research on forecasting using deep learning, such as the work of (Qin, 2021) and (Zhang, 2022) which indicate sMAPE errors in the approximate range of 13% to 25% for similar time horizons in the retail and finance sectors, our results with the ensemble model show a significant improvement in accuracy, especially in periods of H=1 and H=3. The variations may be due to aspects such as lower temporal granularity in certain studies, variable availability of covariates, and discrepancies in cross-validation techniques. Our monthly dataset, which has medium granularity, together with a robust temporal validation protocol and the inclusion of relevant covariates, may explain the observed advantage in performance. (Rodríguez, 2022)

Finally, the combination of convolutional and recurrent neural networks in the ensemble improves the ability to model seasonality and long-term relationships, a result that coincides with research such as (Rodríguez, 2022) , thus reinforcing the robustness of the forecast in general. In summary, our results not only support current trends in peer-reviewed literature, but also position our approach among the most effective for sales forecasting in similar contexts, providing a solid quantitative basis for practical use and future research.

Conclusions

This analysis investigates the use of deep learning ensembles to increase the accuracy of sales forecasts in a technology company. By merging LSTM, CNN, and MLP architectures, complex and nonlinear patterns in sales data can be identified. A monthly dataset was used, covering the period from 2021 to 2024 and containing more than 10,000 records, applying robust time validation.

The findings indicate that the ensemble model performs better than traditional models and isolated networks, achieving an sMAPE of 12.3%, an MdAPE of 10.5%, a WAPE of 15.2%, and an RMSE of 0.28 for the horizon H=1. These results mark a substantial improvement over models such as naive r seasonal (sMAPE 25.0%, RMSE 0.50). This advanced method facilitates better demand forecasting and supports strategic decision-making, although it requires greater computational resources. The integration of deep learning techniques into a set provides robustness and improves accuracy, which is crucial for optimizing commercial management in changing environments.

In summary, the use of advanced machine learning and deep learning techniques represents considerable progress in sales forecasting, enabling more active and flexible management. However, it is essential to consider the available resources and data quality to improve their effectiveness in real-world situations. These findings support the implementation of current predictive systems as a strategic investment to increase competitiveness in the market.

References

Ahmed, R. &. (2023). Feature Engineering and Covariates in Sales Forecasting Models. *International Journal of Data Science*,, 45-60.

Alim Toprak Fırat, O. A. (2025). *DergiPark Akademik*. Retrieved from Development of machine learning based demand forecasting models for the e-commerce sector: https://dergipark.org.tr/en/pub/ijedt/issue/86818/1567739

Ancco Yaurimucha, T. D. (2021). *Alicia-Concytec*. Retrieved from Comparative analysis of time series to project sales in footwear hierarchies in a retail sector company: https://alicia.concytec.gob.pe/vufind/Record/UNMS_b3b16cc3f9eb280a0e73979df9332956

Erick Lambis-Alandete, M. J.-G.-H. (March 6, 2023). *COMPARISON OF DEEP LEARNING ALGORITHMS*. Retrieved from FOR CRYPTOCURRENCY PRICE FORECASTING: https://www.redalyc.org/journal/2913/291377795014/html/

Hyndman, R. J. (2018). *Forecasting: principles and practice*. OTexts.

IEBS. (DECEMBER 4, 2024). *Business & Tech*. Retrieved from Business Trends for 2025: https://www.iebschool.com/hub/tende ncias-empresariales-innovacion-innovacion/

IT.SITIO. (September 13, 2025). Retrieved from Deep neural networks: The future of data analysis and pattern prediction?: https://www.itsitio.com/inteligencia-artificial/redes-neuronales-profundas/

Jaramillo, D. I. (2024). *Predictive sales model using machine learning.* Retrieved from National Open and Distance University: https://repository.unad.edu.co/bitstream/handle/10596/62872/dibaronj.pdf?sequence=1&isAllowed=y

Joaquín Amat Rodrigo, J. E. (10, 2022). *Data Science*. Retrieved from Global forecasting models: modeling multiple time series with machine learning: https://cienciadedatos.net/documentos/py44-multi-series-forecasting-skforecast-espa%C3%B1ol.html

Joaquín Amat Rodrigo, J. E. (September 20, 2025). *Data Science*. Retrieved from ARIMA and SARIMAX models with Python: https://cienciadedatos.net/documentos/py51-modelos-arima-sarimax-python#Modelo_ARIMA-SARIMAX

Kim, J. K. (May 1, 2025). *A Comprehensive Survey of Deep Learning for Time Series Forecasting: Architectural Diversity and Open Challenges*. Retrieved from https://arxiv.org/pdf/2411.05793

Marlon Rubén Barcia Moreira, P. D. (July 26, 2025). *Application of Artificial Intelligence in Sales Management for Predicting Customer Behavior.* Retrieved from Sinergia Academica.

Molina, M. R. (July 7, 2022). *Deep Learning Model Ensemble Techniques Applied to Time Series Forecasting.* Retrieved from https://crea.ujaen.es/server/api/core/bitstreams/c93bb12e-1a6b-4c7f-963f-961ceecbea30/content

Parkhomenka, A. (April 24, 2025). *Medium.* Retrieved from Forecasting and Evaluation: https://medium.com/@alalparkh/time-series-forecasting-sarima-vs-prophet-e957931a2aff

Qin, Y. e. (2021). *A Dual-Stage Attention-Based Recurrent Neural Network for Time Series Prediction.* Retrieved from IEEE Transactions on Knowledge and Data Engineering.

RevVana. (December 19, 2023). Retrieved from Types of sales forecasting models: https://revvana-com.translate.goog/resources/blog/types-of-sales-forecasting-models/?_x_tr_sl=en&_x_tr_tl=es&_x_tr_hl=es&_x_tr_pto=tc

Rodríguez, C. H. (2022). *Prediction and Classification of Stock Market Time Series Using* Recurrent Neural Networks

Saavedra, F. I. (2021). *Artificial Neural Networks.* Retrieved from Neural Model: https://d1wqtxts1xzle7.cloudfront.net/36957207/Redes_neuronales-libre.pdf?1426217722=&response-content-disposition=inline%3B+filename%3DRedes_Neuronales_Artificiales.pdf&Expires=1757809492&Signature=TQQe4yRbaqBGqM8jAL-

~DrAmhhnuETlC8XkCdb3iIP2eYVsX0g~UVzcF

Sun, P. (May 4, 2025). *Medium*. Retrieved from Cross-Validation in Time Series: https://medium.com/@pacosun/respect-the-order-cross-validation-in-time-series-7d12beab79a1

Xiangjie Kong, Z. C. (2025). *Springer Nature Link*. Retrieved from Deep learning for time series forecasting: a survey: https://link.springer.com/article/10.1007/s13042-025-02560-w

Yasaman Ensafi, S. H. (2022). *International Journal of Information Management: Data Perspectives*. Retrieved from Forecasting seasonal item sales time series using machine learning: a comparative analysis: https://www-sciencedirect-com.translate.goog/science/article/pii/S2667096822000027?_x_tr_sl=en&_x_tr_tl=es&_x_tr_hl=es&_x_tr_pto=tc

Zhang, G. e. (2022). *Convolutional Neural Networks for Time Series Forecasting: Advances and Applications. Neurocomputing.* Retrieved from Advances and Applications. Neurocomputing.

Zhou, Y. &. (2023). *A Review on Sales Forecasting Models Based on Machine Learning.* Journal of Sales Analytics, . Retrieved from Journal of Sales Analytics.