Data and Metadata. 2025; 4:483
doi: 10.56294/dm2025483

ORIGINAL

A Hybrid HW-RFR Forecasting Model: Case of Moroccan Pharmaceutical sector

Un Modelo Híbrido de Previsión HW-RFR: El caso del sector farmacéutico marroquí

Mariam ATWANI[1,2], Mustapha HLYAL[2], Jamila EL ALAMI[1]

[1]LASTIMI Laboratory, Higher School of Technology in Sale, Mohammed V University, Rabat, Morocco.
[2]Logistics Center of Excellence (CELOG), Higher School of Textile and Clothing Industries, ESITH, Casablanca, Morocco

Cite as: ATWANI M, HLYAL M, EL ALAMI J. A Hybrid HW-RFR Forecasting Model: Case of Moroccan Pharmaceutical sector. Data and Metadata. 2025; 4:483. https://doi.org/10.56294/dm2025483

Submitted: 08-03-2024 **Revised:** 29-07-2024 **Accepted:** 30-11-2024 **Published:** 01-01-2025

Editor: Adrián Alejandro Vitón-Castillo

Corresponding author: Mariam ATWANI

ABSTRACT

Sales forecasting is an essential element of effective supply chain management, particularly in the pharmaceutical sector where continuous availability of drugs is crucial. This article examines sales forecasts for fluoxetine, an antidepressant available on the Moroccan market under six trade names and 14 different forms. The main objective of this study is to compare the effectiveness of four forecasting models, namely Prophet Facebook, ARIMA, GRU and Holt-Winters through their accuracy, and to propose a hybrid model that will contribute to improving the accuracy of demand forecasts.
Each model was applied individually to predict future sales, and evaluated using MAPE, MAE and RMSE metrics. Next, a hybrid model, integrating Holt-Winters and Random Forest Regressor methods, was developed to leverage the robustness of traditional models while improving predictive performance through machine learning techniques. The results of the study show that traditional models, such as ARIMA and Holt-Winters, offer a solid basis for sales forecasting. However, the hybrid HW-RFR (Holt-Winters Random Forest Regressor) model stands out for a significant improvement in forecast accuracy, demonstrating great robustness to fluctuations in fluoxetine demand. This article highlights the potential of hybrid models for forecasting pharmaceutical sales. The improved forecast accuracy achieved with the HW-RFR model provides stakeholders with more reliable information, enabling them to make informed decisions to optimize pharmaceutical supply chain management.

Keywords: Demand Forecasting; Pharmaceutical SC; Holt-Winters; Random Forest Regressor.

RESUMEN

La previsión de ventas es un elemento esencial de la gestión eficaz de la cadena de suministro, sobre todo en el sector farmacéutico, donde la disponibilidad continua de medicamentos es crucial. Este artículo examina las previsiones de ventas de fluoxetina, un antidepresivo disponible en el mercado marroquí bajo seis nombres comerciales y 14 formas distintas. El principal objetivo de este estudio es comparar la eficacia de cuatro modelos de previsión, a saber, Prophet Facebook, ARIMA, GRU y Holt-Winters por su precisión, y proponer un modelo híbrido que contribuya a mejorar la exactitud de las previsiones de demanda. Cada modelo se aplicó individualmente para predecir las ventas futuras, y se evaluó utilizando las métricas MAPE, MAE y RMSE. A continuación, se desarrolló un modelo híbrido, que integra los métodos Holt-Winters y Random Forest Regressor, para aprovechar la robustez de los modelos tradicionales y mejorar al mismo tiempo el rendimiento predictivo mediante técnicas de aprendizaje automático. Los resultados del estudio muestran que los modelos tradicionales, como ARIMA y Holt-Winters, ofrecen una base sólida para la previsión de ventas. Sin embargo, el modelo híbrido HW-RFR (Holt-Winters Random Forest Regressor) destaca por una mejora significativa de la precisión de las previsiones, demostrando una gran robustez frente a las fluctuaciones de la demanda de fluoxetina. Este artículo destaca el potencial de los modelos híbridos para predecir las ventas de productos farmacéuticos. La mejora de la precisión de las previsiones conseguida con el modelo HW-RFR proporciona a las partes interesadas una información más fiable, lo que les permite tomar decisiones con conocimiento de causa para optimizar la gestión de la cadena de suministro farmacéutica.

Palabras clave: Previsión de la demanda, SC farmacéutica, Holt-Winters, Regressor Random Forest.

INTRODUCTION

Demand forecasting is an essential discipline in supply chain management, enabling companies to anticipate future needs and optimize their resources. Accurate forecasting helps minimize the costs associated with overproduction or stock-outs, while improving customer satisfaction through better product availability.[1] In a context of increasing competition and fluctuating demand, companies need to adopt robust and resilient forecasting methods to maintain their competitive edge.[2]

Traditional forecasting methods, such as exponential smoothing, moving averages and Box-Jenkins approaches, assume that future demand will follow past trends.[3] Their effectiveness therefore depends on the reliability of historical data, which complicates forecasting for new products without historical data. However, disruptions in value chains and fickle customer preferences have raised concerns about the reliability of traditional forecasting models and their limitations in accurately predicting demand behavior. As a result, with the rapid evolution of data-driven environments, forecast accuracy is now essential. Researchers are now combining these traditional methods with artificial intelligence algorithms to improve accuracy.[4] The choice of modeling techniques depends on various factors.

In the pharmaceutical sector, demand forecasting is of particular importance. Medicines are critical products whose availability is vital to public health.[5] Accurate forecasting of drug sales ensures that sufficient quantities are available, thus avoiding shortages that could endanger patients' health. In addition, a good forecast helps to optimize stock levels, reduce storage costs and avoid product obsolescence, which is particularly critical for drugs with a limited shelf life. The global pharmaceutical market is one of the most dynamic and regulated, with sales reaching around $1,482 billion in 2022.[6] Innovation, an ageing population and an increase in chronic diseases fuel this growth. In Morocco, the pharmaceutical sector is in full expansion, with a strong local industry and significant imports. The country has several production units and a healthcare policy that promotes access to essential medicines. On the other hand, the antidepressant market represents a significant segment of the pharmaceutical sector. With the increase in mental health disorders worldwide, demand for these drugs has been growing steadily.[7] Forecasting sales of antidepressant drugs is therefore essential to ensure continued availability and to meet the growing needs of patients.

A preliminary study has focused on studies that have worked on forecasting demand in the pharmaceutical industry from 2018. Models used range from traditional methods such as ARIMA (AutoRegressive Integrated Moving Average), SMA (Simple Moving Average), Exponential Smoothing, Holt's Winter, to modern approaches such as deep neural networks (RNN, LSTM), random forests (RF), and machine learning techniques such as XGBoost and regression algorithms.

Data sources vary from internal pharmaceutical manufacturer data to global databases such as Kaggle, to electronic drug management systems (eLMIS). Data quality and richness strongly influence forecast accuracy. For example, studies using, in addition to sales history, exogenous data tends to provide more accurate results.[6,8,9,10,11,12,13,14]

Studies reveal that deep learning and hybrid models outperform regression models and statistical methods in terms of accuracy.[6,8,15,16] However, they require large amounts of data and high computing capacities. Approaches using combinations of methods, such as integrating symbolic learning with genetic programming, are also showing promising performance.[6]

In this study, the performance of four antidepressant drug sales forecasting models is compared: ARIMA (AutoRegressive Integrated Moving Average), Facebook Prophet, GRU (Gated Recurrent Unit), and Holt-Winters. These models are evaluated using three common metrics: MAPE (Mean Absolute Percentage Error), MAE (Mean Absolute Error) and RMSE (Root Mean Squared Error). In addition, a new hybrid model that combines the strengths of Holt-Winters and Random Forest Regression is proposed, to provide even more accurate forecasts. Thus, this article aims not only to evaluate existing models but also to explore innovative solutions for improving the forecasting of antidepressant drug sales, offering valuable insights for actors in the pharmaceutical sector.

METHOD

The present study concerns the pharmaceutical specialties of various antidepressants whose international non-proprietary name is fluoxitione. The study began with a statistical analysis of the time series, starting with data visualization and time series decomposition. Next, a descriptive and autocorrelation analysis was performed, to identify the mean, median, standard deviation, and quantiles, autocorrelation functions (ACF) and partial autocorrelation functions (PACF) etc. Finally, tested for stationarity was realized using the Augmented Dickey-Fuller (ADF) test.

The results of this statistical analysis enabled the identification of the most appropriate forecasting models for our data, namely AutoRegressive Integrated Moving Average (ARIMA), Holt-Winters, Prophet Facebook and Gated Recurrent Unit (GRU). Then compared the forecasts generated by these models was compared and the performance was assessed by calculating Mean Absolute Percentage Error (MAPE), Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). Finally, hybrid model combining Holt-Winters and Random Forest Regression was created.

Model description

AutoRegressive Integrated Moving Average ARIMA

The ARIMA model is a combination of autoregressive (AR) and moving-average (MA) methods, defined by three parameters: p, d and q, which correspond respectively to the order of autoregression(AR), the degree of differentiation to handle non-stationarities in the time series, and the order of the moving average (MA).[17] It is a commonly used method for dynamic time series forecasting.[16]

Holt-Winters (HW)

The Holt-Winters method is a time series analysis technique for modeling seasonal effects in data. It is classified as an exponential correction method. The formulation of this method is based on three distinct equations: level equation (6), trend equation (7), and seasonality equation (8). The HW model can be additive or multiplicative, its optimal parameters determined as α, β and γ respectively.[8,15]

Prophet Facebook

The Prophet model, developed by Facebook, is a flexible and robust method for forecasting time series. It is designed to handle time series with non-linear trends, seasonal components and holidays.[18,19] Prophet is characterized by three components: Tend, Seasonality and Holidays.

Gated Recurrent Unit (GRU)

Gated Recurrent Units (GRUs) are a type of recurrent neural network (RNN) designed to solve some of the problems encountered by traditional RNNs, notably the vanishing gradient problem when learning long sequences.[20,21] GRU is a simplified variant of Long Short-Term Memory (LSTM) and uses gates to control the flow of information, namely: Update gate, Reset gate, New memory state, Hidden state.

Hybrid model: Holt-Winters Random Forest Regressor (HW-RFR)

After comparing the accuracy of the forecasts generated by the above-mentioned models, it was clear that a more efficient model was needed to achieve a greater forecast accuracy. Consequently, a hybrid model was developed, which uses the Random Forest Regressor model to improve the sales forecasts generated by the HW model, by combining the forecasts of a Holt-Winters model with the residuals predicted by the Random Forest Regressor model. Figure 1 illustrates the steps involved in the HW-RFR model.

Holt-Winters forecast residuals are calculated as the difference between actual sales and the Holt-Winters forecast, equation 1:

$$Residues_t = Real\ sales_t - HW\ Forecasts_t \quad (1)$$

Where t represents time

And $HW\ Forecasts_t$ are represented by the equations (2) and (3).

$$\hat{Y}_{t+h} = l_t + hb_t + s_{t+h-m(k+1)} \quad (2)$$

$$\hat{Y}_{t+h} = (l_t + hb_t) \times s_{t+h-m(k+1)} \quad (3)$$

Where is the season length, h is the forecast horizon, and k is the number of full seasons between t and $t+h$.

https://doi.org/10.56294/dm2025483

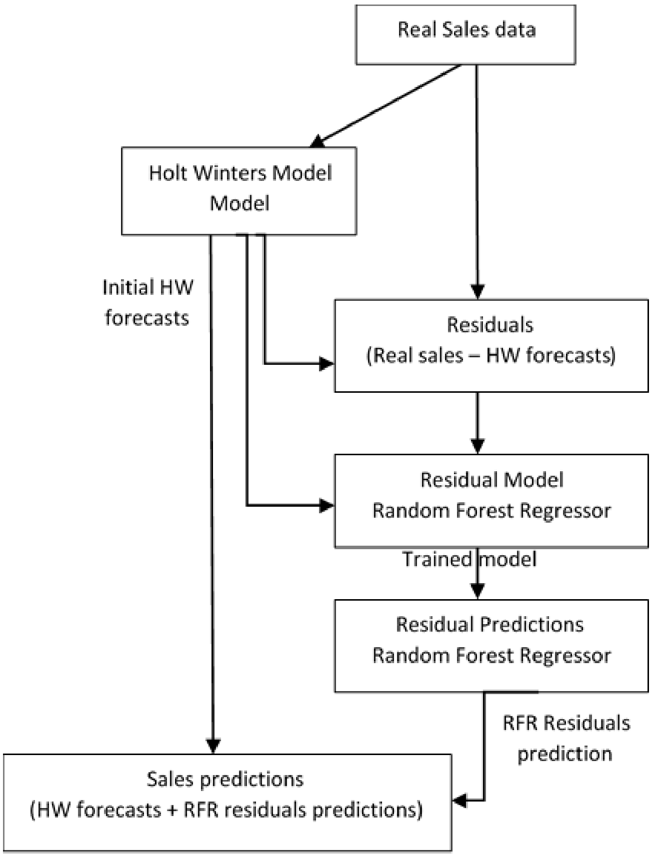

Figure 1. The HW-RFR model Flowchart

- Level equation is defined as follows:

$$l_t = \alpha(Y_1 - s_{t-m}) + (1-\alpha)(l_{t-1} + b_{t-1}) \quad (4)$$

- The trend equation is as follows:

$$b_t = \beta(l_l - l_{t-1}) + (1-\beta)b_{t-1} \quad (5)$$

- Seasonality is defined according to the equation:

$$s_t = \gamma(Y_t - l_t) + (1-\gamma)s_{t-m} \quad (6)$$

Where α, β and γ are optimal parameters.

Holt-Winters' additive method (equation (2)) includes additive seasonal variation as well as a linear trend over time. This approach is more effective at handling additive seasonality.[22] The multiplicative method (equation 3), on the other hand, incorporates multiplicative seasonal variation and a linear trend over time.[23,24,25] Next, the Random Forest Regressor model is trained on the residual training set using the time index as a feature according to equation 7:

$$\hat{R}_t = f(t) \quad (7)$$

Where f is the Random Forest Regressor model (equation 12) and $\hat{R}_t$ represents the predicted residuals. Finally, the Holt-Winters forecasts are combined with the predicted residuals to obtain the hybrid forecasts according to the following equation:

$$hybrid\ Forecasts_t = HW\ forecasts_t + \hat{R}_t \quad (8)$$

Figure 2 illustrates the architecture of the Random Forest Regressor model, which works by building a multitude of decision trees during training and averaging the predictions of these individual trees for the test observations. Each internal node represents a condition on a feature, each branch represents the result of the condition, and each leaf represents a predicted value.

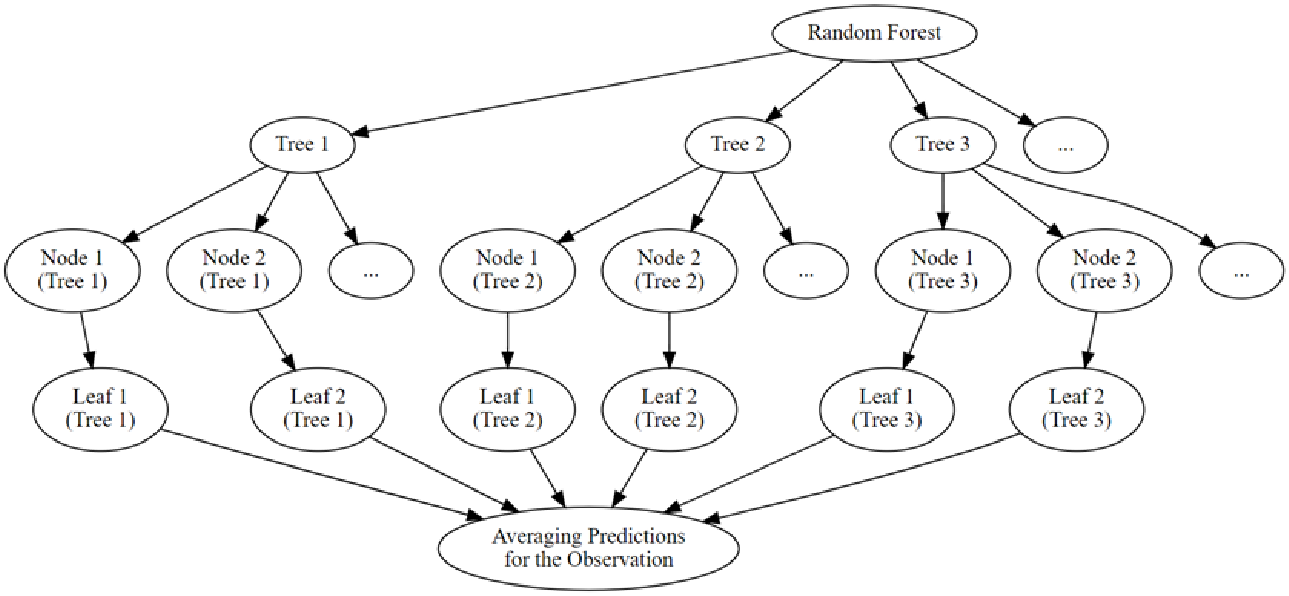

Figure 2. Random Forest Regressor model architecture

Consider a decision tree T. For a data point x with a feature t (time), the prediction of a decision tree is given by the equation 9.

$$\hat{R}_t^{\mathcal{T}} = \mathcal{T}(t) \quad (9)$$

Each decision tree in a random forest is constructed using a bootstrap sample of the training data. Consider $\{T_1, T_2, \ldots, T_B\}$ the B decision trees in the forest. The prediction of an individual tree T_b on a data point is given by equation 10.

$$\hat{R}_t^{\mathcal{T}_b} = \mathcal{T}_b(t) \quad (10)$$

The prediction of a tree is the average of the values of the leaves in which x is found, represented by the equation 11.

$$\hat{\mathcal{T}}_b(t) = \frac{1}{|L|} \sum_{i \in L} \mathcal{T}_i \quad (11)$$

Where *L* is the leaf containing x

The prediction of the random forest is then the average of the predictions of all the trees according to the equation 12.

$$\hat{R}_t = \frac{1}{B} \sum_{b=1}^{B} \hat{\mathcal{T}}_b(t) \quad (12)$$

Forecast accuracy assessment

To evaluate model performance and measure forecast accuracy, the following metrics were used:

MAPE

MAPE is used to assess the difference between predicted and observed values.[23] According to [10] The model is judged “very good” if MAPE<10 %, “good” if 10 % <MAPE <20 %, “acceptable” if 20 % < MAPE <50 % and finally “faulty and erroneous” if MAPE <50 %. The MAPE calculation formula is defined as follows:

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100 \quad (13)$$

Where *n* is the total number of observations, y_i the actual value for observation *i* et $\hat{y}_i$ the predicted value for observation *i*.

MAE

The MAE measures the average absolute error between actual and predicted values. It indicates the average magnitude of errors in predictions.[26]

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|y_t - \hat{y}_t| \quad (14)$$

Where y_t is actual demand for period t, $\hat{y}_t$ forecast demand for period t and n the total number of periods.

RMSE

RMSE is used to measure the accuracy of a forecasting model by calculating the mean of the squared errors between predicted and actual values.[11,12,13,14]

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2} \quad (15)$$

Where y_t is the actual demand for period t, $\hat{y}_t$ is the forecast demand for period t and n is the total number of periods. This measure is commonly used to assess the accuracy of forecasting models, as it gives an indication of the average size of forecast errors, weighting them equally, which is particularly useful for time series and regression models.[16]

RESULTS

Data collection

The data used for this study comes from the IQVIA™ Quintiles database, formerly known as IMS-health. This is a history of sales between September 2018 and April 2024, sold by pharmaceutical companies to wholesalers, pharmacies, hospitals and clinics. This article focuses on antidepressants whose International Nonproprietary Name is Fluoxitine, which is marketed in Morocco under six pharmaceutical specialties held by six laboratories, i.e. 14 drugs in all.

Data pre-processing

Data preparation is important as it gives models a better understanding of the data, and therefore leads to accurate forecasts.[27] Thus, well-prepared data makes the model more stable and less sensitive to minor variations in the data.

Missing values were processes using a forward fill method followed by a backward fill method. Each missing value x_t at time t is replaced by the last non-missing value before t, the forward fill, and if necessary, by the first non-missing value after t, the backward fill.

$$\begin{cases} x_{t-1} \text{ if } x_t \text{ is missing and } x_{t-1} \text{ is nonmising} \\ x_{t+1} \text{ if } x_t \text{ is missing and } x_{t-1} \text{ is missing} \end{cases}$$

To stabilize the variance, a logarithmic transformation was applied. For each value in the time series, the log transformation is as follows:

$$yt = \log(xt + 1) \quad (16)$$

Where y_t is the transformed series and we add 1 to compensate for the possibility of having log(0).

Then, for the 14 time series, the data were separated to two sets, using 80 % for training and 20 % for testing. Finally, to stabilize and accelerate training, and to ensure consistent performance of the Prophet FB, GRU and HW-RFR models, data normalization was applied using the Min-Max Scaling method (equation 17), which involves transforming the features to lie between 0 and 1.[28]

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (17)$$

Where X is the original value of the feature, x_{min} the minimum value of the feature in the dataset, x_{max} the maximum value of the feature in the dataset and X' is the normalized value of the feature, which lies between 0 and 1.

Hyperparameters tuning

AutoRegressive Integrated Moving Average ARIMA

The hyperparameters of the ARIMA model are p, which corresponds to the order of the autoregressive (AR) part, d which is equal to the number of differentiations needed to make the series stationary, and finally q, which corresponds to the order of the moving average (MA) part, these parameters can be deduced from the descriptive statistical analysis done earlier. However, to improve the accuracy of the ARIMA model, Python's "statsmodels" library was used to apply Grid Search and evaluate possible combinations using the Akaike Information Criterion (AIC) to identify the best-performing model. Table 1 describes the hyperparameters used for each time series.

Table 1. ARIMA model's hyperparameters

	Fluoxet 30 tablets	Fluoxet 60 tablets	Fluoxet 10 tablets	Fluoxet 20 tablets	Tuneluz 30 tablets	Tuneluz 20 tablets	Tuneluz 10 tablets	Serdep 28 tablets	Serdep 14 tablets	Vyoxet 28 tablets	Vyoxet 14 tablets	Fluzoft 28 tablets	Fluzoft 14 tablets	Fluctine 12 tablets
p	4	5	2	1	5	2	2	3	2	1	3	3	4	4
d	1	1	1	2	2	1	1	2	3	1	2	2	2	1
q	3	2	1	2	4	2	2	3	4	3	3	3	3	

Holt-Winters (HW)

Table 2 summarizes the hyperparameters selected for the HW model: Trend (additive or multiplicative), Seasonal (additive or multiplicative) and Smoothing Parameters: Alpha (level), Beta (trend), Gamma (seasonality).

Table 2. HW model's hyperparameters

	Fluoxet 30	Fluoxet 60	Fluoxet 10	Fluoxet 20	Tuneluz 30	Tuneluz 20	Tuneluz 10	Serdep 28	Serdep 14	Vyoxet 28	Vyoxet 14	Fluzoft 28	Fluzoft 14	Fluctine 12
Trend	Add	Add	Add	Add	Add	Mul	Mul	Add	Add	Add	Add	Add	Add	Add
Seasonal	Add	Add	Mul	Add	Add	Mul	Mul	Add	Add	Add	Add	Add	Add	Add
Seasonal Periods	12	12	12	12	12	12	12	12	12	6	12	12	12	12
Alpha	0	0	0	0	1	1	0	0	0	0	1	0	0	0
Beta	0,2	0,2	0,2	0,2	0,5	0,05	0,8	0,9	0,6	0,9	1	0,1	0,2	0,2
Gamma	0,8	0,8	0,5	0,7	0,2	0,1	0,3	0,85	0,6	1	0,1	0,7	0,4	0,3

Gated Recurrent Unit (GRU)

Optuna was used to tune the hyperparameters of the GRU model. This automated hyperparameter optimization library uses advanced optimization algorithms to find the best parameter combinations.

Table 3 shows the hyperparameters of the GRU model for each product:

Random Forest Regressor

To predict the residuals with the Random Forest Regressor model, Grid Search was used to exhaustively explore all the combinations of hyperparametric values grouped in table 4.

Table 3. GRU model's hyperparameters

	Fluoxet 30	Fluoxet 60	Fluoxet 10	Fluoxet 20	Tuneluz 30	Tuneluz 20	Tuneluz 10	Serdep 28	Serdep 14	Vyoxet 28	Vyoxet 14	Fluzoft 28	Fluzoft 14	Fluctine 12
Num layers	1	2	1	1	2	2	2	2	2	2	1	2	1	1
Units 0	64	32	96	96	128	128	96	96	64	32	64	32	32	32
Units 1		96			96	64	128	96	128	64		128		
Dropout 0	0,3	0,3	0,1	0,3	0,1	0,3	0,2	0,1	0,2	0,1	0,2	0,2	0,1	0,3
Dropout 1					0,3	0,2	0,3	0,1	0,1	0,3		0,1		
Units last	32	32	96	32	32	96	64	128	32	32	64	64	32	128
Dropout last	0,1	0,2	0,1	0,2	0,2	0,1	0,1	0,3	0,1	0,1	0,2	0,1	0,1	0,3
Dense units	32	64	48	48	16	48	32	16	48	32	32	64	64	32
Learning rate	0,0003	0,0006	0,0005	0,0023	0,0001	0,0007	0,0008	0,0022	0,0004	0,0070	0,0001	0,0061	0,0008	0,0001

Table 4. Random Forest Regressor model's hyperparameters

	Fluoxet 30	Fluoxet 60	Fluoxet 10	Fluoxet 20	Tuneluz 30	Tuneluz 20	Tuneluz 10	Serdep 28	Serdep 14	Vyoxet 28	Vyoxet 14	Fluzoft 28	Fluzoft 14	Fluctine 12
N estimators	100	300	300	300	100	300	300	100	100	300	100	200	300	300
Max depth	None	10	10	10	10	10	10	30	30	10	8	20	10	10
Min samples split	2	2	2	10	5	10	2	5	5	10	10	5	10	10
Min samples leaf	1	1	4	1	1	4	2	2	2	1	5	2	4	4
Max features	sqrt	sqrt	sqrt	sqrt	sqrt	sqrt	sqrt	log 2	log 2	sqrt	sqrt	sqrt	sqrt	sqrt

DISCUSSION

This section analyzes the results obtained from sales forecasts for the 14 drug presentations using four different models: Prophet FB, ARIMA, GRU and Holt-Winters (HW). To assess the performance of these models, three commonly used forecasting error metrics were calculated: MAPE, MAE and MSE. Table 5 summarizes these results.

Table 5. ARIMA, HW, Prophet FB and GRU model's assessment

	MAPE				MAE				RMSE			
	PROPHET FB	ARIMA	GRU	HW	PROPHET FB	ARIMA	GRU	HW	PROPHET FB	ARIMA	GRU	HW
FLUOXET 30	6,5 %	5,1 %	5,1 %	3,9 %	987	814	809	604	1 246	1 014	1 004	755
FLUOXET 60	7,9 %	8,5 %	10,6 %	4,9 %	207	225	280	125	236	265	318	153
FLUOXET 10	5,0 %	6,4 %	12,5 %	2,5 %	49	63	119	25	60	76	136	42
FLUOXET 20	80,0 %	12,0 %	15,4 %	21,6 %	313	35	50	25	346	47	57	30
TUNELUZ 30	6,1 %	5,7 %	6,8 %	4,9 %	448	413	505	354	644	603	616	490
TUNELUZ 20	23,2 %	23,3 %	17,5 %	18,6 %	91	89	70	62	113	98	84	77
TUNELUZ 10	228,3 %	229,0 %	149,3 %	86,8 %	183	197	112	60	198	205	131	77
SERDEP 28	39,5 %	7,1 %	7,4 %	8,6 %	2 422	420	441	496	2 732	536	548	604
SERDEP 14	48,9 %	7,6 %	10,4 %	9,1 %	360	57	80	77	379	72	94	88
VYOXET 28	87,6 %	74,1 %	54,8 %	45,3 %	994	799	1 016	717	1 317	1 117	826	912
VYOXET 14	67,2 %	25,9 %	25,8 %	28,9 %	128	46	52	55	21 870	64	69	76
FLUZOFT 28	8,6 %	10,0 %	18,3 %	8,3 %	73	87	147	68	92	101	169	92
FLUZOFT 14	22,5 %	29,5 %	37,6 %	14,7 %	17	21	26	12	22	23	27	17
FLUCTINE 12	95,7 %	25,7 %	8,4 %	34,1 %	486	131	39	160	530	141	50	216

The results of this study show that the Holt-Winters (HW) model proved to be the best performer among the models tested. In terms of MAPE, MAE, MSE and RMSE, it consistently demonstrated better accuracy for most drugs, namely, all Fluoxet presentations except Fluoxet 20, for which the best MAPE was provided by the ARIMA model, all Tuneluz presentations except Tuneluz 20, for which the best MAPE was provided by the GRU model, Fluzoft 28 and 14. The HW model also achieved the best MAPE for 10 of the 14 drugs studied, demonstrating its ability to minimize errors. It also demonstrated better accuracy with a MAPE below 10 % for 7 products, i.e. predictions judged “very good”, and “good” for 3 products according to the researchers’ scale.[10] The performance of the Holt-Winters model can be attributed to the data featuring trend and seasonality, since this model is particularly suited to time series with recurring demand cycles, thanks to its trend and seasonality components.

It is evident that the Pophet FB model underperformed with accurate forecasts (below 10) for only 5 products. This is due to the data not considering exogenous events, since the model includes vacation days in its composition. However, in this study, only historical sales data for the drugs was used. On the other hand, Prophet FB does not capture the effects of abrupt changes, while drug sales are influenced by sudden external factors such as stock-outs for imported drugs, price changes, mental health awareness campaigns, regulatory changes.

As for the GRU model, it has a complex structure that requires a large volume of data and longer training periods to capture temporal dynamics effectively, which hampered its performance in the current study. Finally, the ARIMA model showed average performance, with MAPE below 10 for 7 time series, and MAE and RMSE that are relatively acceptable.

Despite the good performance of the HW model, there is still room for improvement to further enhance forecast accuracy. Therefore, the development of a hybrid model was considered: the HW-RFR model, which combines Holt-Winters with Random Forest Regression. This hybrid model combines the advantages of Holt-Winters’ seasonal decomposition with the flexible modeling capabilities of Random Forests (RFR) to improve forecast accuracy.

Table 6. HW-RFR model’s assessment

	MAPE	MAE	RMSE
FLUOXET 30	2,3 %	639	826
FLUOXET 60	2,6 %	79	102
FLUOXET 10	3,4 %	63	85
FLUOXET 20	11,0 %	84	139
TUNELUZ 30	3,8 %	352	415
TUNELUZ 20	14,7 %	61	93
TUNELUZ 10	65,7 %	43	58
SERDEP 28	9,0 %	527	612
SERDEP 14	9,5 %	73	87
VYOXET 28	21,2 %	255	400
VYOXET 14	27,8 %	41	77
FLUZOFT 28	9,1 %	73	94
FLUZOFT 14	26,3 %	25	30
FLUCTINE 12	16,0 %	139	221

The results of the hybrid HW-RFR model (table 6) show a significant improvement in forecast accuracy compared to individual models for Fluoxet 30 tablets, 60 tablets, and 20 tablets, for all presentations of Tuneluz, and for Vyoxet. Additionally, the model achieves a low root mean square error for most of the data. This indicates that the HW-RFR model effectively reduces forecast errors by capturing both seasonal trends and complex variations in the time series. The development of the hybrid HW-RFR model represents a significant advancement in improving the accuracy of drug sales forecasts. By combining the strengths of Holt-Winters and Random Forests, this hybrid model offers a robust and accurate solution for forecasting complex time series, thus meeting the specific needs of the pharmaceutical sector. The promising preliminary results suggest that this approach deserves further exploration and validation to further optimize sales forecasts.

CONCLUSION

The aim of this study was to evaluate and improve sales forecasting models for fluoxetine, with a view to increasing accuracy to enable better pharmaceutical supply chain management. The results show that, among the standard models, the Holt-Winters model performed best in capturing seasonal trends. However, the development of the hybrid HW-RFR (Holt-Winters with Random Forest Regression) model addressed the limitations of the individual models, reducing prediction errors and capturing more complex variations in sales data.

This hybrid HW-RFR model, designed to adapt to specific fluctuations in fluoxetine sales, demonstrates a significant improvement in forecast accuracy. As a result, its adoption can not only strengthen inventory and supply management in the pharmaceutical sector, but also improve responsiveness to variations in demand.

For future research, several insights can be explored to further improve the accuracy of drug sales forecasts. First, an extensive validation of the hybrid HW-RFR model on more diverse and larger datasets is necessary to confirm its robustness and accuracy. Integrating exogenous factors, such as special events, mental health awareness campaigns, marketing campaigns, and socio-economic data, could also enhance forecast accuracy. Finally, comparative studies with other advanced machine learning and artificial intelligence techniques could help identify the most effective approaches for forecasting drug demand.

REFERENCES

1. KUMAR V, KUMAR R. The demand forecasting: A comparative review of conventional and non-conventional techniques. International Journal of Mechanical and Production Engineering Research. 2019.

2. Mediavilla MA, Dietrich F, Palm D. Review and analysis of artificial intelligence methods for demand forecasting in supply chain management. Procedia CIRP. 2022 Jan 1; 107:1126–31.

3. Atwani M, Hlyal M, Elalami J. A Review of Artificial Intelligence applications in Supply Chain. ITM Web of Conferences. 2022; 46:03001.

4. Ren S, Chan HL, Siqin T. Demand forecasting in retail operations for fashionable products: methods, practices, and real case study. Ann Oper Res. 2020 Aug 1;291(1–2):761–77.

5. Merkuryeva G, Valberga A, Smirnov A. Demand forecasting in pharmaceutical supply chains: A case study. In: Procedia Computer Science. Elsevier B.V.; 2019. p. 3–10.

6. Pharmaceutical Market Size, Share & Trends Report; 2030.

7. Hoefler R, Galvão TF, Ribeiro-Vaz I, Silva MT. Trends in Brazilian market of antidepressants: A five-year dataset analysis. Front Pharmacol. 2022 Oct 4;13.

8. İMECE S, BEYCA ÖF. Demand Forecasting with Integration of Time Series and Regression Models in Pharmaceutical Industry. International Journal of Advances in Engineering and Pure Sciences. 2022 Sep 30;34(3):415–25.

9. Zhu X, Ninh A, Zhao H, Liu Z. Cross-Series Demand Forecasting using Machine Learning: Evidence in the Pharmaceutical Industry

10. Poyraz İ, Gürhanlı A. Demand Forecasting with Time Series Analysis Using Drug Sales Data. International Journal of Engineering Research and Applications. 2020;10:51–4.

11. Mbonyinshuti F, Nkurunziza J, Niyobuhungiro J, Kayitare E. The prediction of essential medicines demand: A machine learning approach using consumption data in rwanda. Processes. 2022 Jan 1;10(1).

12. Kravets AG, Al-Gunaid MA, Loshmanov VI, Rasulov SS, Lempert LB. Model of medicines sales forecasting taking into account factors of influence. Journal of Physics: Conference Series. Institute of Physics Publishing. 2018.

13. Mbonyinshuti F, Nkurunziza J, Niyobuhungiro J, Kayitare E. Application of random forest model to predict the demand of essential medicines for non-communicable diseases management in public health facilities. Pan African Medical Journal. 2022 May 1;42.

14. Nowadly K, Jung S. Using Machine Learning Approaches to Improve Long-Range Demand Forecasting. 2020.

15. Siddiqui R, Azmat M, Ahmed S, Kummer S. A hybrid demand forecasting model for greater forecasting accuracy: the case of the pharmaceutical industry. Supply Chain Forum. 2022;23(2):124–34.

16. Cheng CY, Chiang KL, Chen MY. Intermittent Demand Forecasting in a Tertiary Pediatric Intensive Care Unit. J Med Syst. 2016 Oct 1;40(10).

17. Fourkiotis KP, Tsadiras A. Applying Machine Learning and Statistical Forecasting Methods for Enhancing Pharmaceutical Sales Predictions. Forecasting. 2024 Mar 1;6(1):170–86.

18. Rafferty G. Forecasting Time Series Data with Facebook Prophet: Build, improve, and optimize time series forecasting models using the advanced forecasting tool. Packt Publishing Ltd. 2021

19. Sivaramakrishnan S, Fernandez TF, Babukarthik RG, Premalatha S. Forecasting Time Series Data Using ARIMA and Facebook Prophet Models. Big Data Management in Sensing: Applications in AI and IoT. 2021 Jan 1;75–87.

20. Pirani M, Thakkar P, Jivrani P, 2022 undefined. A comparative analysis of ARIMA, GRU, LSTM and BiLSTM on financial time series forecasting. In: EEE International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE) [Internet]. 2022

21. Yamak PT, Yujian L, Gadosey PK. A comparison between ARIMA, LSTM, and GRU for time series forecasting. ACM International Conference Proceeding Series. 2019 Dec 20;49–55.

22. Tratar L, Mojškerc B, Toman A. Demand forecasting with four-parameter exponential smoothing. Int J Prod Econ. 2016; 181:162–73.

23. Burinskiene A. Forecasting Model: The Case of the Pharmaceutical Retail. Front Med (Lausanne). 2022 Aug 3;9.

24. Chen H, Qian C, Liang C, Kang W. An approach for predicting the compressive strength of cement-based materials exposed to sulfate attack. PLoS One. 2018 Jan 1;13(1).

25. Emrul Kays HM, Karim ANM, Hasan M, Sarker RA. Impact of Initial Level and Growth Rate in Multiplicative HW Model on Bullwhip Effect in a Supply Chain. Data and Decision Sciences in Action. 2018;357–68.

26. Hodson TO. Root-mean-square error (RMSE) or mean absolute error (MAE): when to use them or not. Geoscientific Model Development. Copernicus GmbH.2022; 15: 5481–7.

27. El Filali A, Lahmer EH Ben, El Filali S, Kasbouya M, Ajouary MA, Akantous S. Machine Learning Applications in Supply Chain Management: A Deep Learning Model Using an Optimized LSTM Network for Demand Forecasting. International Journal of Intelligent Engineering and Systems. 2022 Apr 1;15(2):464–78.

28. Shantal M, Othman Z, Bakar AA. A Novel Approach for Data Feature Weighting Using Correlation Coefficients and Min-Max Normalization. Symmetry (Basel). 2023 Dec 1;15(12).

FINANCING
The authors did not receive financing for the development of this research.

CONFLICT OF INTEREST
The authors declare that there is no conflict of interest.

AUTHORSHIP CONTRIBUTION
Conceptualization: Mariam ATWANI, Mustapha HLYAL, Jamila EL ALAMI.
Data curation: Mariam ATWANI.
Formal analysis: Mariam ATWANI.
Research: Mariam ATWANI.
Methodology: Mariam ATWANI.
Project management: Mustapha HLYAL.
Resources: Mariam ATWANI.
Software: Mariam ATWANI.
Supervision: Mustapha HLYAL, Jamila EL ALAMI.
Validation: Mustapha HLYAL, Jamila EL ALAMI.
Display: Mariam ATWANI.
Drafting - original draft: Mariam ATWANI.
Writing - proofreading and editing: Mariam ATWANI, Mustapha HLYAL.