ISBN: 978-981-94-1156-6
2024 the 14th International Workshop on Computer Science and Engineering (WCSE 2024)
doi: 10.18178/wcse.2024.06.021

On the ES-LSTM Forecasting Model for Optimizing Drug Inventory Management: A Preliminary Attempt

Kittikorn Sriwichai [1], Janejira Laomala [1] and Sayan Kaennakham [2 +]

[1] Department of Interdisciplinary Science and Internationalization, Institute of Science, Suranaree University of Technology, Nakhon Ratchasima, 30000, Thailand;

[2] School of Mathematics and Geoinformatics, Institute of Science, Suranaree University of Technology, Nakhon Ratchasima, 30000, Thailand;

Abstract. This study introduces and evaluates the ES-LSTM, a forecasting model that integrates exponential smoothing with LSTM networks to optimize drug inventory management. The ES-LSTM model was tested against traditional LSTM models using pharmaceutical sales data, demonstrating superior accuracy in capturing both peak and regular demand patterns. The findings indicate that ES-LSTM could revolutionize inventory management in healthcare, making it a promising area for further investigation.

Keywords: ES-LSTM, Drug Inventory Management, Exponential Smoothing

1. Introduction

Accurate drug demand prediction is not merely beneficial but essential for ensuring effective drug inventory management, which has a direct impact on the efficiency of healthcare operations and the quality of patient care [1]. Traditional forecasting methods, including time series analysis and regression models [2], have frequently fallen short in addressing the complexity of pharmaceutical demand, particularly when confronted with abrupt shifts in demand or dealing with sparse datasets. These challenges can lead to either significant overstock or dangerous shortages, both of which are detrimental to healthcare outcomes and economic efficiency. The severity of these issues has underscored the necessity for more sophisticated predictive tools capable of handling the intricate and often unpredictable nature of drug demand. In this context, the advent of machine learning algorithms has marked a revolutionary advancement in forecasting methodologies [3]. Among these, neural networks, including Long Short-Term Memory (LSTM) networks, and decision trees stand out for their ability to assimilate and analyze vast and varied datasets. These range from real-time health statistics to insights gleaned from social media, providing a more nuanced and dynamic approach to forecasting.

Long Short-Term Memory (LSTM) networks are a distinguished class of machine learning models, excelling in their capacity to model complex time series data and capture long-term dependencies [4]. Their prowess makes them highly suitable for a variety of applications across different domains such as finance, healthcare, and natural language processing. Despite their capabilities, the deployment of LSTMs requires rigorous data preprocessing to effectively manage challenges like missing values, noise, and the intrinsic variability of time series data [5]. Normalizing data to align with the LSTM's sensitivity to input scale is also crucial. This preprocessing stage is not merely a preliminary step but a foundational component that

+ Corresponding author. Tel.: +0-4422-4383; fax: +0-4422-4185.
E-mail address: sayan_kk@g.sut.ac.th

significantly influences the LSTM's ability to learn from the data accurately and deliver reliable predictions [6].

In our research, we are exploring the integration of exponential smoothing (ES) with Long Short-Term Memory (LSTM) networks, a method we call 'ES-LSTM'. This approach aims to combine ES's ability to smooth short-term fluctuations with LSTM's capacity for capturing long-term dependencies, enhancing time-series prediction accuracy. Our goal is to improve predictions in applications such as drug demand and sales forecasting, potentially revolutionizing forecasting methods in healthcare and beyond.

2. Methodology

2.1 The Long Short-Term Memory (LSTM)

Long Short-Term Memory (LSTM) network [7], a type of recurrent neural network (RNN) used for sequential data processing. Each component of the LSTM cell is highlighted in Figure 1, showing how data flows through the network at time t, with the description of each of the essential compoents is provided in Table 1.

Table 1: Overview of LSTM network components: Mathematical formulations and functional descriptions

Component	Math formula	Function
Forget Gate (f_t)	$f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right)$	Decides what information to discard from the cell state. It uses the sigmoid activation function (σ).
Input Gate (i_t)	$i_t = \sigma\left(W_i \cdot [h_{t-1}, x_t] + b_i\right)$	Decides what new information to store in the cell state. It also involves a sigmoid layer.
Candidate Layer ($\tilde{C}_t$)	$\tilde{C}_t = \tanh\left(W_C \cdot [h_{t-1}, x_t] + b_C\right)$	Creates a candidate vector of new information to add to the cell state, using the tanh activation function.
Cell State Update	$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$	Combines the forget gate's output and the input gate's proposal to update the cell state (C_t).
Output Gate (O_t)	$O_t = \sigma\left(W_o \cdot [h_{t-1}, x_t] + b_0\right)$ $h_t = O_t * \tanh(C_t)$	Decides the next hidden state (h_t), using the sigmoid function and the updated cell state passed through tanh.

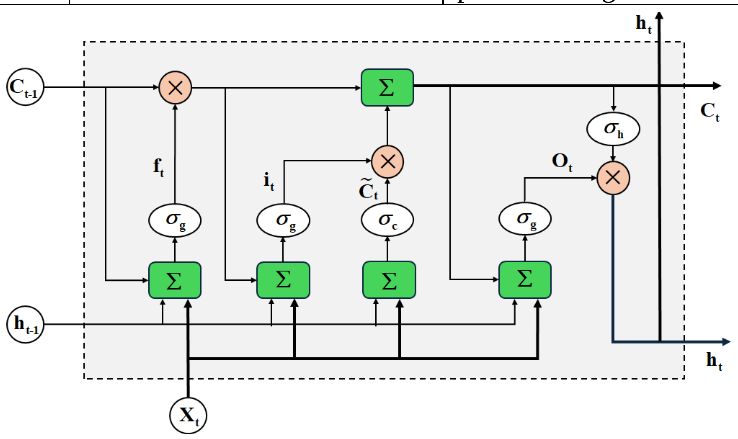

Figure 1: An architecture of a Long Short-Term Memory (LSTM) network.

2.2 The Exponential Smoothing (ES)

Exponential smoothing (ES) [8] is a rule of thumb technique for smoothing time series data using the exponential window function. Its popularity stems from the simplicity of calculation, intuitive appeal, and good performance on a wide range of time series. The simplest form of exponential smoothing is applied when the time series does not exhibit trend or seasonal patterns, which is often called "Single Exponential Smoothing." The formulas for the Single Exponential Smoothing are given by

$$S_t = \alpha x_t + (1-\alpha) S_{t-1}, \tag{1}$$

where

- S_t is the smoothed statistic, the forecast for the next period
- x_t is the actual value at time t
- S_{t-1} is the value of the smoothed statistic for the previous period
- α is the smoothing factor of the series, $0 \leq \alpha \leq 1$.

3. Experimental Setupus

3.1 The Algorithm

Figure 2 depicts a data processing sequence starting with preprocessing, including exponential smoothing, followed by a series of LSTM layers interspersed with dropout layers for regularization and ReLU activations for non-linearity. This structured pipeline culminates in an output, demonstrating a methodical approach to enhance data handling and prediction.

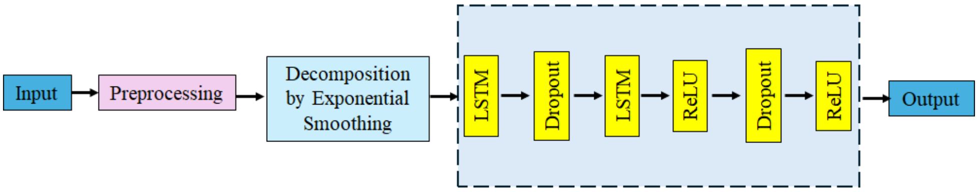

Figure 2: Data processing pipeline with the sequential flow.

3.2 The Dataset

Since the actual dataset representing the real demand for medicine in a healthcare unit is not available for this initial research, we used a dataset of drug sales instead, assuming that the statistical behaviors of both are broadly similar. The dataset comprises 600,000 transaction records from 2014 to 2019, detailing sales date, time, drug brand, and quantity from a pharmacy Point-of-Sale system. It includes 57 drugs categorized under the Anatomical Therapeutic Chemical (ATC) system, with this study focusing specifically on the M01AB category, non-steroidal anti-inflammatory and antirheumatic drugs, particularly acetic acid derivatives. The dataset is available at:
https://www.kaggle.com/datasets/milanzdravkovic/pharma-sales-data.

3.3 The Error Measurement Norms

To measure the performance of the proposed strategy, three well-known error measurement norms are utilized and they are; Mean Absolute Error (MAE): $\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$, Mean Squared Error (MSE): $\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$, and Root Mean Squared Error (RMSE): $\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$.

4. Main Results and General Dissucions

Table 2 and Table 3 show that the ES-LSTM consistently outperforms the traditional LSTM model. In the 2-unit setup, the MAE for monthly forecasts decreased dramatically from 112.38029 in LSTM to just 0.56420 in ES-LSTM, highlighting the effectiveness of integrating exponential smoothing into LSTM. Similar improvements are seen in other periods and error metrics, such as the reduction in RMSE for weekly predictions from 15.26822 in LSTM to 0.97782 in ES-LSTM, which underscores the added value of exponential smoothing in reducing prediction errors across all tested intervals.

Table 2: Performance comparison of 2-unit LSTM vs. 2-unit ES-LSTM across different periods.

Error Norm	Hourly		Daily		Weekly		Monthly	
	LSTM	ES-LSTM	LSTM	ES-LSTM	LSTM	ES-LSTM	LSTM	ES-LSTM
MAE	0.54293	0.24174	0.81048	1.73186	11.77241	0.75947	112.38029	0.56420
MSE	0.81509	0.15614	2.21820	5.05532	233.11858	0.95614	13405.92555	0.98753
RMSE	0.90282	0.39515	1.48936	2.24840	15.26822	0.97782	115.78396	0.99375

Table 3: Performance comparison of 4-unit LSTM vs. 4-unit ES-LSTM across different periods.

Error Norm	Hourly		Daily		Weekly		Monthly	
	LSTM	ES-LSTM	LSTM	ES-LSTM	LSTM	ES-LSTM	LSTM	ES-LSTM
MAE	0.3076	0.2276	1.5362	1.5273	5.0883	4.9724	0.3076	0.2276
MSE	0.1539	0.1326	3.7413	3.7331	39.4985	37.176	0.1539	0.1326
RMSE	0.3923	0.3642	1.9342	1.9321	6.2847	6.0972	0.3923	0.3642

When comparing the 2-unit and 4-unit LSTM architectures, the results clearly favor the more complex 4-unit configuration, which generally exhibits lower error rates. For instance, the RMSE for monthly forecasts in the 4-unit LSTM is significantly lower, at 0.3923, compared to 115.78396 in the 2-unit LSTM. This trend is consistent across all periods and metrics, illustrating that an increase in the number of LSTM units enhances the model's ability to capture and analyze data trends more effectively, thereby improving overall forecast accuracy. This suggests that optimizing the architecture complexity based on the forecasting horizon is crucial for achieving the best model performance.

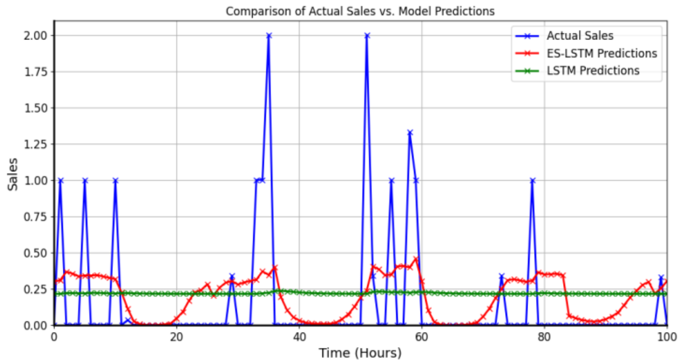

Figure 3: Comparison of actual drug sales with predictions between standard LSTM and enhanced ES-LSTM models over a 100-hour period.

Figure *3* provides a clear comparison between two predictive models, the standard LSTM (LES) and the enhanced ES-LSTM (ES-LES), against actual drug sales data across a 100-hour period. While the actual sales, marked by blue stars, show notable spikes indicating high demand, the standard LSTM model, represented by green circles, consistently underestimates these peaks, suggesting a limitation in its ability to adapt to sudden demand changes. In contrast, the ES-LSTM predictions, shown with red crosses, more accurately mirror the actual sales, especially in capturing the peaks' timing and magnitude. This improved alignment is likely due to the integration of exponential smoothing in the ES-LSTM, which enhances its capability to manage short-term fluctuations effectively, demonstrating its potential for more reliable forecasts in critical areas like healthcare drug inventory management.

5. Conclusion

The ES-LSTM forecasting model has proven effective in optimizing drug inventory management, offering a substantial improvement over standard LSTM models by accurately predicting demand spikes and maintaining stability in typical sales fluctuations. This research suggests that ES-LSTM could significantly enhance healthcare inventory practices, reducing both shortages and overstock. However, further research is required to refine and adapt the model for broader healthcare applications.

Based on this preliminary attempt, we recognize the need to investigate further, including incorporating Holt-Winters and ARIMA baselines, detailing dataset splitting methods, using alternative smoothing for seasonality, comparing ES-LSTM with other models, and optimizing LSTM units.

6. Acknowledgements

This work was supported by (i) Suranaree University of Technology (SUT, http://www.sut.ac.th), (ii) Thailand Science Research Innovation (TSRI, https://www.tsri.or.th), and (iii) National Science, Research and Innovation Fund (NSRF) (NRIIS number 179269). The grant recipient is S. Kaennakham who would like to express his sincere gratitude for their support.

7. References

[1] X. Zhu, A. Ninh, H. Zhao, and Z. Liu. Demand forecasting with supply-chain information and machine learning: Evidence in the pharmaceutical industry. *Production and Operations Management.* 2021, **30**(9): 3231-3252.

[2] S. R. Bhatta, P. Adhikari, and R. Byanjankar. Choice of regression models in time series data. *Economic Journal of Development.* 2020, **29**:1-2.

[3] T. N. Angula and A. Dongo. Assessing the impact of artificial intelligence and machine learning on forecasting medication demand and supply in public pharmaceutical systems: a systematic review. *GSC Biological and Pharmaceutical Sciences.* 2024, **26**(2): 140-150.

[4] B. Lindemann, T. Müller, H. Vietz, N. Jazdi, and M. Weyrich. A survey on long short-term memory networks for time series prediction. *Procedia Cirp.* 2021, **99**: 650-655.

[5] D. Ageng, C.-Y. Huang, and R.-G. Cheng. A short-term household load forecasting framework using LSTM and data preparation. *Ieee Access.* 2021, **9**: 167911-167919.

[6] P. S. Muhuri, P. Chatterjee, X. Yuan, K. Roy, and A. Esterline. Using a long short-term memory recurrent neural network (LSTM-RNN) to classify network attacks. *Informatio.* 2020, **11**(5): 243.

[7] A. Graves and A. Graves. Long short-term memory. *Supervised sequence labelling with recurrent neural networks.* 2012: 37-45.

[8] E. S. Gardner Jr. Exponential smoothing: The state of the art. *Journal of forecasting.* 1985, **4**(1): 1-28.