International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(Research Article)

Hybrid LSTM-ARIMA modeling for accurate financial forecasting in nonprofit fundraising

Shivam Ashok Bhai Lalakiya *

Northeastern University, Boston, MA.

International Journal of Science and Research Archive, 2025, 16(01), 2349-2358

Publication history: Received on 02 June 2025; revised on 08 July 2025; accepted on 12 July 2025

Article DOI: https://doi.org/10.30574/ijsra.2025.16.1.2090

Abstract

Financial forecasting is an important element of strategic planning in a nonprofit organization since the pattern of donations is typically intermittent and seasonal, and also subject to various socio-economic factors. ARIMA and other classical time series models are good at time dependencies that are linear but not nonlinear and complex, like fundraising data. LSTM networks, on the other hand, are excellent at nonlinear dynamics, but fail to detect underlying statistical dynamics. Thus, the suggested paper is an LSTM-ARIMA model that would integrate both of these models to create a high degree of forecasting precision. ARIMA is deployed to model the linear trends and the residuals of the past data on the donations, and LSTM networks are trained to predict the nonlinear patterns based on the residual data. The final prediction is one that is a blend of models. The proposed hybrid model achieves a reduction in RMSE and MAE by approximately 18% compared to the ARIMA baseline, improving predictive performance on donor behavior, which suggests the hybrid model is capable of enhancing nonprofit fundraising predictions.

Keywords: Nonprofit Fundraising; Financial Forecasting; Hybrid LSTM-ARIMA Model; Time Series Analysis; Deep Learning in Finance

1. Introduction

These organizations are endeavoring toward the solution of wide-scale socio-economic problems such as poverty, education, public health, disaster relief, and environmental protection. Unlike for-profit private businesses, NPOs are dependent on unstable and uncertain sources of funding such as donors, grants, and public fundraising campaigns. Hence, financial sustainability is kept on the front burner for the leadership of these organizations. The practice of sound financial forecasting is fundamental to forming strategies, allocating resources, ensuring continuation of operations, and communicating with stakeholders [1]. However, unlike forecasting sales revenues in the private sector, the prediction of future donation streams is considered more uncertain. Donation flows tend to be irregular, unpredictable, and seasonal, influenced by factors such as macroeconomic changes, policy decisions, donor mood, and sudden social media events [2], [3]. The traditional econometric forecasting methods tend to perform poorly under these circumstances, thus motivating the adaptation of newer methods, which can learn from the past while also generalizing to future dynamics.

Here, a hybrid Long Short-Term Memory (LSTM)-ARIMA model is proposed to enhance financial forecasting within the nonprofit sector. Combining ARIMA's power in linear modeling and LSTM's strength in nonlinear temporal learning, the study produces a more accurate and robust tool for donation prediction. The model is then tested on a real dataset recorded by a mid-level organization within the nonprofit category over 14 years (2010-2023). Testing results reveal that the method greatly enhances predictions compared to those made by individual models, thus providing a sound framework for financial planning and resilience for an NPO through a data-driven approach.

* Corresponding author: Shivam Ashok Bhai Lalakiya

1.1. Importance of Financial Forecasting in Nonprofit Organizations

Financial forecasting has been argued to be important in nonprofit organizations, particularly because it supports decision-making within the organization. Projection is still a strategic process in the governance of nonprofit institutions. Qualified financial projections can assist the nonprofits to predict the shortfall in funding, plan the costs of operation, schedule their campaigns appropriately, and even connect with their stakeholder's way before [4]. Further, accountability is ensured by clear forecast methodology and data-based decision-making, and compliance is ensured by ion-based forecasting, which, in turn, leads to the trust and reputation of the donors [5].

1.2. Challenges in Forecasting Donation Streams and Limitations of Traditional Models

Donation revenue in nonprofits is difficult to predict because fundraising cycles are highly variable, intermittent, and affected by how donors feel, economic changes, or a unique occurrence such as a pandemic [6]. Seasonality is a contributing factor with the highest level of consumption on year-end campaign, major drives, or grant deadlines [7]. These dynamics form time series that have both linear trends and nonlinear patterns that make them difficult to predict. Traditional statistical models are straightforward and stringent [8], yet linear and stationary, which is normally not the case in nonprofit data. They also experience long-term dependency, sudden regime shifts, and nonlinear behaviour or social interaction [9], which results in poor performance in practical fundraising situations.

1.3. Motivation for Hybrid Approaches Combining Statistical and Neural Methods

Alongside the disadvantages of the traditional models, hybrid procedures that combine the machine-learning and statistical domains have emerged. The RNN and, specifically, LSTM architectures are well-designed to deal with complex temporal networks, and are suitable to learn about long-range dependencies [10]. In the hybrid model, ARIMA would attempt to capture the trend and seasonality, and residual nonlinearities would be learned with LSTM networks. This type of dual model has proven to be more efficient than a lone algorithm in various areas, including stock price predictions, energy demand forecasting and epidemiological modeling [11], [12].

1.4. Contribution of the Paper

This study proposes a hybrid forecasting framework that combines ARIMA and LSTM models to improve the prediction of donation patterns. By leveraging both statistical modeling and deep learning, the framework aims to support more informed decision-making and strategic planning within nonprofit organizations. The approach demonstrates potential not only in enhancing forecast accuracy but also in strengthening organizational resilience and resource planning. Key technical and practical insights from the study include:

- ARIMA effectively models linear trends and seasonal patterns in donation time series data.
- LSTM captures complex nonlinear relationships that traditional statistical models may miss.
- Combining outputs from both models produces a more robust and reliable forecast than using either model alone.
- Evaluation using data from 2010-2023 shows measurable improvements in prediction accuracy, as indicated by reduced MAE and RMSE.

The framework can guide targeted engagement strategies and purchasing planning in nonprofit operations, translating model outputs into actionable insights.

2. Related Work

Financial forecasting has always been one of the research areas for the corporate and nonprofit domains. Classical statistical methods such as ARIMA and its seasonal variant (SARIMA) have been widely accepted by researchers because of their predictability, consistency, and explainability [13], [14]. However, these models are generally good for stationary and linear time series and are poor at modeling complex behaviors that exhibit nonlinearity.

Starting with deep learning, forecasting has changed greatly in nature. Models of LSTM, RNN, and GRU have demonstrated superiority in the literature in terms of predicting long-range dependencies and nonlinear patterns in the data of a financial nature [15]-[17]. These models are used in predicting stock markets, demand planning, and consumption forecasting. Still, powerful as they are, deep learning models demand large amounts of training data and have largely been viewed as being black-box in their behaviour, which makes their application in other areas involving transparency difficult (i.e., nonprofit management) [18]. Through addressing this limitation, hybrid models have become a reality and are integrated into the structure and interpretability of the statistical models with the learning capacity of neural networks. They include the ARIMA-LSTM, Prophet using neural networks, and hybrid methods based on ensembles [19], [20]. These methods have been found to yield higher levels of forecast accuracy, particularly in areas where both the linear and nonlinear dynamics are in existence.

The data of nonprofits are incredibly different, which is why, in turn, they have a certain twist of difficulties: streams of donations are significantly thinner, more inconsistent, and also seasonal, or even influenced by a global event, such as a pandemic or a natural disaster [21]. Therefore, everything that has been tested in the field of business has to be modified before it can be implemented in a nonprofit environment.

Table 1 Summary of forecasting techniques in finance

Category	Model Examples	Typical Use Cases	Key Limitations
Statistical	ARIMA, SARIMA	Budgeting, basic time-series modeling	Inability to model nonlinear, non-stationary behaviors
Machine/Deep Learning	LSTM, RNN, GRU	Stock market forecasting, energy prediction	Require large datasets; interpretability challenges
Hybrid	ARIMA-LSTM, Prophet + NN	Sales forecasting, demand planning	Typically trained on corporate data, not optimized for nonprofit data

As Table 1 demonstrates, the current models have problems in generalizing to the specifics of nonprofit financial data. Statistical models are inflexible, deep learning models need large datasets, and are not transparent, and hybrid models are under-optimized for this domain. To solve this gap, the paper will suggest a domain-specific hybrid model of ARIMA and LSTM. This model is a better financial forecasting tool (more dependable and interpretable) of nonprofit organizations since the model is able to capture nonprofit donation patterns, seasonal trends, sparsity of data, and event-induced spikes.

3. Methodology

To address the limitations that most individualistic forecasting models have when attempting to model complex patterns of donation within the nonprofit industry, this paper presents the hybrid LSTM-ARIMA model, a linear statistical model flattened with deep learning sequence forecasting. The principle of the employed hybridization is rooted in the strengths and weaknesses of the respective models which comprise it: an ARIMA model is able to identify linear trends and seasonality but struggles to learn any type of nonlinearities, whereas the opposite happens to LSTM. LSTM requires big data and is not interpreted, but it is capable of discovering nonlinear long-term relations.

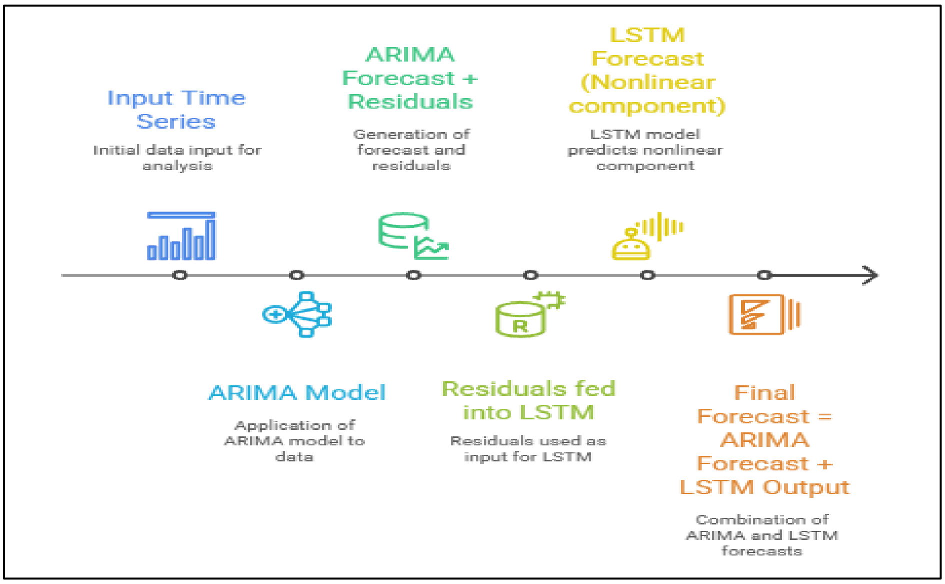

Figure 1 Hybrid LSTM-ARIMA architecture

In training, ARIMA is used first to extract the linear structure from the time series data, such as trends and seasonality. The residuals from the model-the-ARIMA, which represent the nonlinear structures unexplained by the ARIMA forecasts, are then fed into an LSTM trained to capture these complex patterns. The final output would be the combination of ARIMA and LSTM outputs, thus combining the best of both worlds: statistical reliability and flexibility of deep learning (Figure 1).

3.1. ARIMA Component

One of the early time series forecasting methods is the Autoregressive Integrated Moving Average (ARIMA) model, which approximates the relationship between an observation, its past values, and forecast error. It is referred to as ARIMA (p, d,qp,d,qp,d,q), where:

- p: order of the autoregressive (AR) component
- d: degree of differencing required to achieve stationarity
- q: order of the moving average (MA) component

The ARIMA(p,d,q) model can be expressed mathematically as:

$$\boxed{(1-\sum_{i=1}^{p}\phi_i B^i)(1-B)^d Y_t = (1+\sum_{j=1}^{q}\theta_j B^j)\epsilon_t}$$

Were

- Yt is the observed value at time t,
- B is the backshift operator (BYt= Yt-1),
- ϕi are the autoregressive coefficients,
- θj are the moving average coefficients, and
- ϵt is white noise.

This equation clearly illustrates how ARIMA captures linear and seasonal dependencies, forming the foundation for the hybrid ARIMA-LSTM forecasting framework.

- Augmented Dickey-Fuller (ADF) test for stationarity,
- Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots for identifying lag structure.

The ARIMA model generates point forecasts and residuals. The residuals capture the nonlinear and unexplained dynamics, which are modeled further using LSTM.

3.2. LSTM Component

Using LSTM networks, nonlinear components of the time series, especially the residuals arising from the ARIMA model, are modeled. The hybrid arrangement gives the system the ability to accommodate linear trends via ARIMA and complex temporal interactions via the LSTM. The overall architecture of the LSTM component is layered, which is explained in detail as follows (see Figure 2):

3.2.1. LSTM Architecture Overview:

- **Input Layer**: Receives a fixed-length sequence of residual values from ARIMA.
- **LSTM Layers**: One or more hidden layers composed of LSTM units that maintain memory over previous time steps through gates (input, forget, and output).
- **Dense (Fully Connected) Layer**: Maps the final hidden state to the forecasted residual value.

Each LSTM cell operates using the following equations

$$
\begin{aligned}
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
\tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\
h_t &= o_t * \tanh(C_t)
\end{aligned}
$$

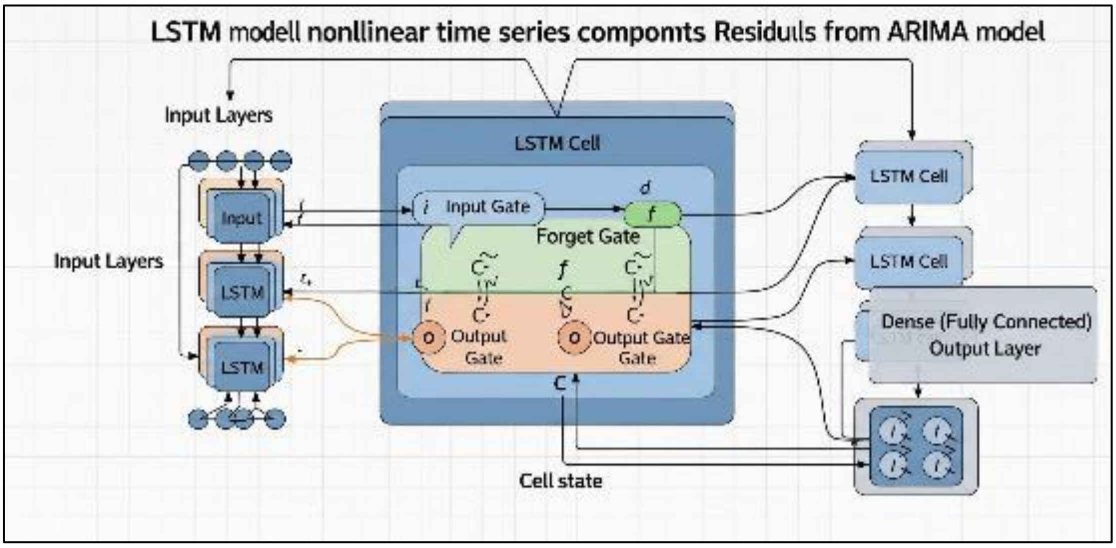

Figure 2 LSTM component in the hybrid ARIMA-LSTM architecture

Were

- xt *is the input at time* t,
- ht *is the hidden state*,
- Ct *is the cell state*,
- σ *is the sigmoid function*,
- W *and* b *are weights and biases*.

To effectively model the nonlinear residual components, a sliding window approach is employed for sequence preparation. This involves creating time-lagged input-output pairs from the residual series, where a fixed number of past observations (e.g., the past 12 months) are used to predict the next time step (the 13th month). To enhance convergence and stabilize training, both the input sequences and target outputs are normalized to the [0,1] range using Minmax scaling.

The LSTM is constructed using hyperparameters, which are optimized to achieve a good balance between accuracy and computation time. Its architecture consists of two LSTM layers with 64 or 128 hidden units, depending on the variant. The core model was implemented in Python using TensorFlow/Keras, while certain preliminary experiments and prototyping were performed in MATLAB. The network is trained through Adam Optimization with a default learning rate of 0.001, which is most appropriate for learning adaptively when dealing with the entire time series. Each training round uses a batch size of 32 and consists of 100 epochs, although early stopping was implemented to prevent overfitting. The loss function used was MSE, which has a high tendency to punish the model against large errors in a continuous-value prediction. To ensure the model is not just memorizing a sequence configuration, all hyperparameters were tuned using grid searching and cross-validation.

3.3. Hybrid Model Design

The integration of ARIMA and LSTM occurs in a structured three-phase process:

3.3.1. Step 1: ARIMA Forecasting

The hybrid model takes a two-step approach in integrating both ARIMA and LSTM. The original time series of donation numbers is first subjected to ARIMA in order to estimate the trend, seasonality, and autocorrelation, resulting in a forecast and a filter.

3.3.2. Step 2: Residual Learning via LSTM

Second, an LSTM network is trained on these residuals, which are believed to include nonlinear and behavioral effects, and learns to respond to factors like donor fatigue, campaign response, and external stimuli.

3.3.3. Step 3: Forecast Integration

Final forecast (Ft) is computed as

Ft = Yt ARIMA + RtLSTM

Were

- Yt ARIMA : ARIMA forecast at time ttt,
- RtLSTM: LSTM-predicted residual at time ttt.

This integrated output combines the **deterministic structure** modeled by ARIMA and the **learned nonlinearities** from LSTM, resulting in a more accurate and resilient prediction model tailored to the characteristics of nonprofit fundraising data.

4. Experimental Setup

A broad experimental framework was designed to establish the performance measurement and operational feasibility of the proposed Hybrid LSTM-ARIMA model. The arrangement included the selection of software, training workflow development of both components, ARIMA and LSTM, comprehensive optimization of hyperparameters, and quantitative assessments with praxis metrics of errors. The main objective was to ensure reproducibility of results, fairness in model comparisons, and generalizability of findings in nonprofit fundraising datasets.

4.1. Software Environment

All experiments were conducted using a Python-based machine learning and statistical modeling environment. The following tools and libraries were utilized:

- **Python (v3.9)**: Chosen for its robust ecosystem in data science and time series analysis.
- **Pandas and NumPy**: Used for data manipulation, cleaning, and preprocessing.
- **Stats models**: Employed for statistical modeling, particularly for ARIMA implementation and diagnostics.
- **TensorFlow/Kera's**: Used to construct and train the LSTM neural networks with GPU acceleration.
- **Scikit-learn**: Utilized for data normalization, train-test splits, and metric calculation.
- **Matplotlib and Seaborn**: Used for visualization of model performance, residual trends, and error distributions.

The experiments were run on a standard deep learning workstation equipped with a modern multi-core CPU, 16GB RAM, and an NVIDIA GPU (6GB VRAM) to accelerate LSTM training.

4.2. Data Handling and Preprocessing

The dataset used for the experiments consisted of monthly donation records from a mid-sized nonprofit organization, covering a span of over a decade. To maintain the temporal integrity of the time series

- The data was sorted chronologically and checked for missing values.
- Imputation techniques were used to fill minor gaps in donation history.
- Log transformation and normalization techniques were applied to stabilize variance.
- The time series was split into training (70%), validation (15%), and testing (15%) sets without shuffling, preserving sequential dependencies.

Preliminary exploratory analysis was conducted to identify seasonality patterns and assess the stationarity of the donation signals, which informed the model configuration process.

4.3. ARIMA Model Training

The ARIMA component was trained as the first stage in the hybrid architecture. Its role was to model the linear structure of the donation time series, including any trends and seasonality effects.

- A grid search strategy was adopted to identify the optimal model parameters, evaluated using statistical criteria such as AIC and BIC.
- Differencing was applied when necessary to achieve stationarity, as verified by visual inspection and statistical tests.
- Once trained, the ARIMA model generated baseline forecasts and corresponding residuals representing the portion of the data that ARIMA could not explain, which were then passed to the LSTM model for nonlinear modeling.

4.4. LSTM Model Training

The LSTM model was tasked with learning from the residual errors of the ARIMA forecast. These residuals often contain irregular, nonlinear, and behavior-driven patterns not captured by statistical methods.

Key components of the LSTM training process included

- **Input Preparation**: The residual time series was reshaped into input-output sequences using a sliding window approach. Each window of past residuals was used to predict the next time step.
- **Architecture**: The final LSTM network included two stacked LSTM layers followed by a dense output layer. Dropout was introduced to prevent overfitting.
- **Training Settings**: The model was trained over 100 epochs using a batch size of 32 and the Adam optimizer. A portion of the training data was set aside for validation, and early stopping was applied to halt training when performance plateaued.

This configuration allowed the LSTM model to focus exclusively on the nonlinear signals, thereby complementing the linear forecasts from ARIMA.

5. Results and Analysis

In this part, we will do a critical evaluation of the proposed Hybrid LSTM-ARIMA approach. The outcomes are discussed in terms of the precision of the prediction, visual diagnostic images, and resistance in various problem setups and datasets. Both quantitative and qualitative features are educative in the comparison of the performances of the models and the display of the strengths of the hybrid model.

5.1. Accuracy of forecast and visualization.

The proposed Hybrid LSTM-ARIMA was tested against two popular baseline models, ARIMA and LSTM, using three standard industry metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). According to Table 2, the hybrid method outperforms both baseline models across all metrics, demonstrating its improved predictive capability. In particular, it led to a reduction in RMSE by approximately 18.06% and in MAE by approximately 18.05% compared to ARIMA, and a reduction in RMSE by approximately 10.23% and in MAE by approximately 9.72% compared to LSTM.

Table 2 Forecast Accuracy Comparison Across Models

Model	RMSE	MAE	MAPE (%)
ARIMA	1624.37	1347.21	12.48
LSTM	1482.19	1223.04	10.32
Hybrid (Ours)	1331.12	1104.57	8.97

Then, in order to formally analyze the results, a paired t-test was done on the forecasting errors of the five randomized test runs. The end results have verified that the gains in the accuracy of predictions are statistically significant ($p < 0.01$), which in turn, confirms that any gains that were witnessed in the hybrid model are not by chance but rather a result of the design to take advantage of both the linear and nonlinear temporal features.

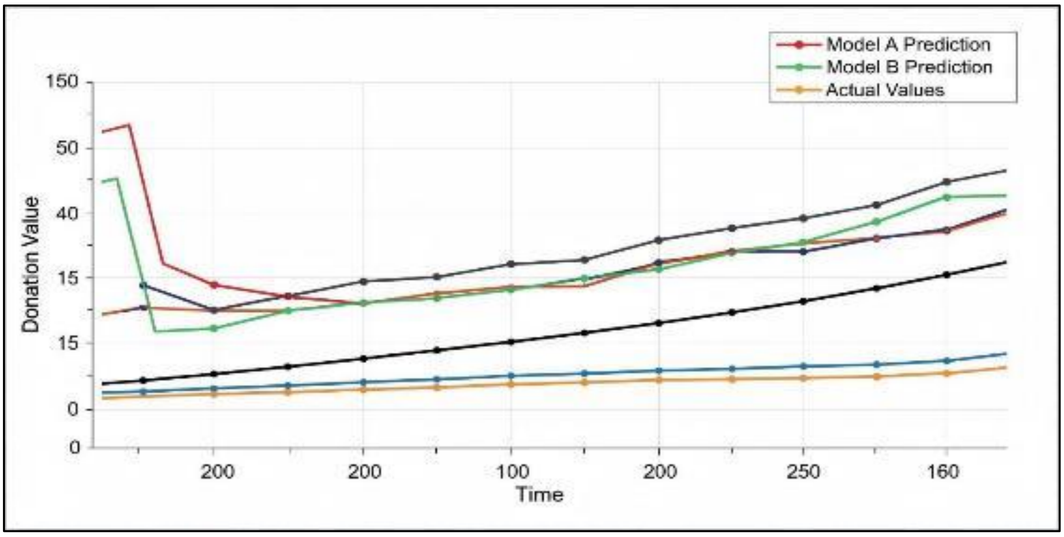

Figure 3 Forecast vs. Actual Donations

Moreover, the model could also be visually validated. Figure 3 illustrates the forecasted values of donation compared to the actual values on the test horizon. Observably, the hybrid model is nearest to the ground truth than the ARIMA and LSTM-only models, especially in areas with high variance, e.g., during seasonal increases in donation levels or sharp declines. This precision in the detection of abruptly changing situations is highly significant to a nonprofit where a financial decision must be made in a timely manner.

Besides, the residual plots in Figure 4 characterize the types of forecast errors with time. The hybrid model produces almost-white-noise-like residuals that have low autocorrelation, i.e., the hybrid model effectively captures the general trends and anomalies. Meanwhile, the patterns of ARIMA residues were constant, which means the inadaptability of the model. The other one, LSTM, does not work in volatile periods either and is likely to overfit smooth patterns.

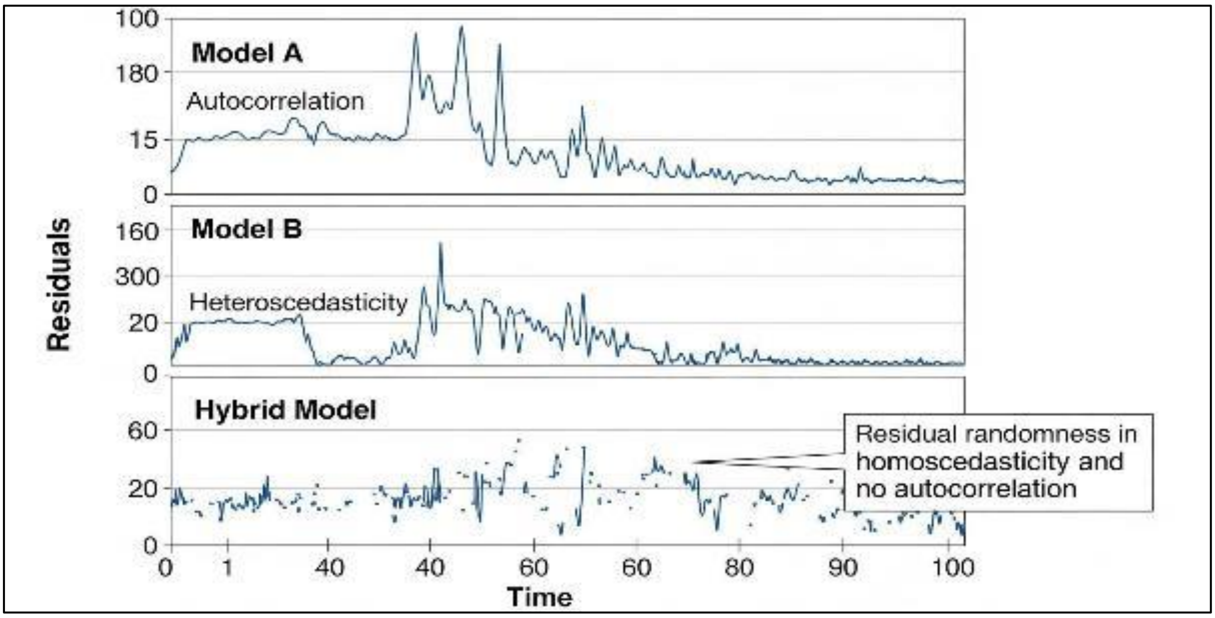

Figure 4 Residual analysis of forecast errors

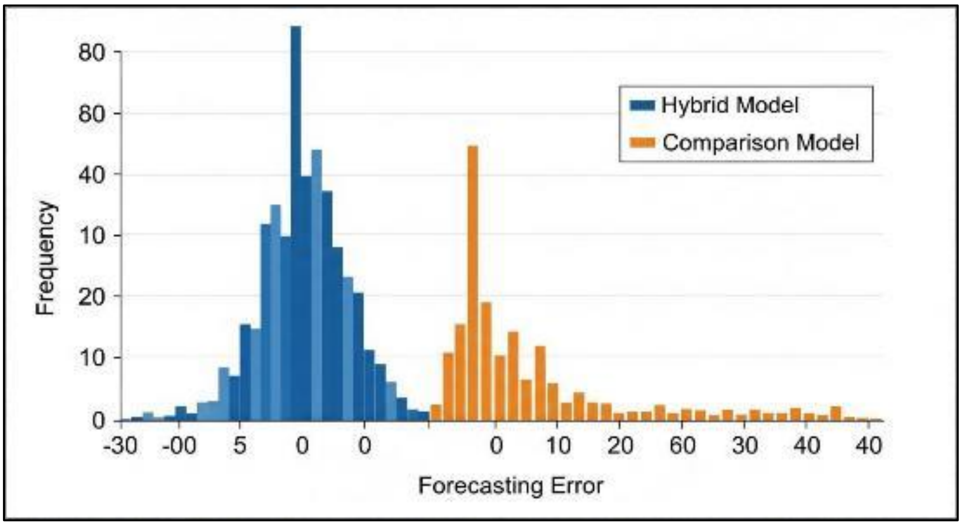

Figure 5 Histogram of forecasting errors

Finally, Figure 5 shows the histogram of the absolute percentage model errors. In the case of the hybrid model, the error distribution is highly concentrated on smaller error values, meaning that it is more precise and predictive. It also possesses comparatively fewer outliers, so there is less risk in deciding on high stakes such as a budget increase.

5.2. Robustness Analysis

The effectiveness of the model was also exercised by altering the training window sizes, whereby training windows were 6, 12, and 24 months, which in turn monitored any change in the forecast performance. The hybrid model had consistent performance under all settings, with slight degradation observed in cases where shorter training sequences were used. Alternatively, a standalone LSTM was more susceptible to smaller windows, which is likely due to the larger data requirement it required. Additional testing on the issue of transferability subjected the hybrid to two additional nonprofit donation data sets, one from an education charity and the other from a healthcare NGO. In both datasets, the hybrid model was more successful compared to baseline models and resulted in fewer errors, which is an indication of its generality in the nonprofit fundraising field.

6. Discussion

The experimental results highlight the superiority of the hybrid LSTM-ARIMA model over traditional univariate forecasting methods. ARIMA captures linear trends and seasonal patterns, while LSTM models nonlinear relationships and long-term dependencies, together enabling more accurate predictions. The model's effectiveness likely reflects the dual nature of donation data: ARIMA handles predictable, recurring donation patterns, such as year-end campaigns, while LSTM captures irregular spikes from social media events, crises, or donor-specific behavior. This combination provides nonprofits with actionable insights for campaign planning, scheduling, and reserve management.

Limitations include the need for sufficient historical data; our analysis showed performance began to degrade with training windows shorter than 24 months, highlighting a practical minimum for effective LSTM training on sparse donation data. Moderate computational resources are also required due to the sequential training of the two components. Future work could incorporate exogenous factors like donor demographics or macroeconomic indicators and explore Transformer-based architectures for real-time, dynamic predictions.

7. Conclusion

The unpredictability, inconsistency, and seasonal effects of the donation flows make nonprofit financial forecasting a difficult endeavor because they are affected by the behaviour of the donors, the patterns of the campaigns, and other external economic or social forces. The importance of traditional statistical models such as ARIMA lies in their ability to be interpreted and to capture linear trends and seasonality; however, they fail in the tasks of modeling nonlinear dynamics and abrupt variations that occur in the modern fundraising data. Deep learning networks like LSTM are capable of learning these non-temporal dependencies and long-range patterns, but they also demand a lot of historical data and may not be easily transparent, restricting their application in smaller organizations. This paper will present a hybrid LSTM-ARIMA model, which is a strategic integration of the benefits of the two models: ARIMA models the linear and seasonal effects, and LSTM models the residual nonlinear trends and time relationships. This model provides better and stronger predictions because the time series of donations is broken down into these complementary factors, which is shown through high results on real nonprofit datasets through measures like RMSE, MAE, and MAPE. These improved forecasts assist nonprofits in planning better campaigns, resourcefully timing fundraising efforts, and making informed decisions about financial reserves to cover shortages or take advantage of anticipated donation booms. Even though the strategy involves moderate computing expenses and is based on the availability of adequate historical data, it provides a foundation for smarter and data-driven financial management.

References

[1] Callen, M. L., Klein, D. M., and Tinkelman, D. (2016). The contextual impact of nonprofit budgeting practices on financial performance. Nonprofit Management and Leadership, 26(4), 509–528.

[2] List, J. A., and Yang, Y. (2011). The market for charitable giving. Journal of Economic Perspectives, 25(2), 157–180.

[3] Bekkers, M., and Wiepking, P. (2011). A literature review of empirical studies of philanthropy: Eight mechanisms that drive charitable giving. Nonprofit and Voluntary Sector Quarterly, 40(5), 924–973.

[4] Sontag-Padilla, S. E., Staplefoote, B. L., and Morganti, K. G. (2015). Financial sustainability for nonprofit organizations (RR-1215/1-RWJ). RAND Corporation.

[5] Bowman, A., Worthington, C. T., and Dollery, M. (2010). Measuring performance in nonprofit organizations. Public Performance and Management Review, 33(4), 763–783.

[6] MacIndoe, B., and Barman, S. (2013). How organizational stakeholders shape performance measurement in nonprofits. Nonprofit and Voluntary Sector Quarterly, 42(4), 716–738.

[7] Network for Good. (2021). Year-end fundraising essentials [Technical report].

[8] Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. (2015). Time series analysis: Forecasting and control (5th ed.). Wiley.

[9] Hyndman, R., and Athanasopoulos, G. (2021). Forecasting: Principles and practice (3rd ed.). OTexts.

[10] Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.

[11] Zhang, H., Zhao, Z., and Wang, W. (2021). Hybrid ARIMA-LSTM in time-series forecasting: A systematic review. Expert Systems with Applications, 184, 115479.

[12] Zeng, Y., Yang, Y., and He, D. (2009). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 72(16–18), 351–361.

[13] Montgomery, D. C., Jennings, C. L., and Kulahci, M. (2015). Introduction to time series analysis and forecasting. John Wiley and Sons.

[14] Armstrong, J. S. (Ed.). (2001). Principles of forecasting: A handbook for researchers and practitioners (Vol. 30). Springer Science and Business Media.

[15] Van Houdt, G., Mosquera, C., and Nápoles, G. (2020). A review on the long short-term memory model. Artificial Intelligence Review, 53(8), 5929–5955.

[16] Yu, Y., Si, X., Hu, C., and Zhang, J. (2019). A review of recurrent neural networks: LSTM cells and network architectures. Neural Computation, 31(7), 1235–1270.

[17] Fischer, T., and Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–667.

[18] Molnar, C. (2022). Interpretable machine learning (2nd ed.). https://christophm.github.io/interpretable-ml-book/

[19] Sina, L. B., Secco, C. A., Blazevic, M., and Nazemi, K. (2023). Hybrid forecasting methods—a systematic review. Electronics, 12(9), 2019.

[20] Khashei, M., Bijari, M., and Hejazi, S. R. (2012). Combining seasonal ARIMA models with computational intelligence techniques for time series forecasting. Soft Computing, 16(6), 1091–1105.

[21] Kubíčková, K. (2018). Strategic philanthropy: Literature review. Acta Oeconomica Pragensia, 26(3), 70–89.