

Exploring the Use of Recurrent Neural Networks for Time Series Forecasting

Jasmin Praful Bharadiya,
Doctor of Philosophy Information Technology,
University of the Cumberland, USA

Abstract:- Recurrent Neural Networks (RNNs) have become competitive forecasting methods, as most notably shown in the winning method of the recent M4 competition. However, established statistical models such as exponential smoothing (ETS) and the autoregressive integrated moving average (ARIMA) gain their popularity not only from their high accuracy, but also because they are suitable for non-expert users in that they are robust, efficient, and automatic. In these areas, RNNs have still a long way to go. We present an extensive empirical study and an open-source software framework of existing RNN architectures for forecasting, and we develop guidelines and best practices for their use. Recurrent neural networks have been effectively used to predict outcomes from irregular time series data in a variety of industries, including medicine, traffic monitoring, environmental monitoring, and human activity detection. The paper focuses on two widely used methods for dealing with irregular time series data: missing value imputation during the data preprocessing stage and algorithm modification to deal with missing values directly during the learning process. Models that can handle problems with irregular time series data are the only ones that are reviewed; a wider variety of models that deal more widely with sequences and regular time series are not included.

Keywords: *Time Series Forecasting, Recurrent Neural Networks, Deep-Latent Variable Models, Sensitivity Analysis and Time Series Data Prediction.*

I. INTRODUCTION

Recurrent neural networks are well suited to supervise learning problems where the dataset has a sequential nature. Time series forecasting should not be an exception. RNNs are essentially neural networks with memory. They can remember things from the past, which is obviously useful for predicting time-dependent targets. Yet, applying them to time series forecasting is not a trivial task. The aim of this article is to review some practices that can help RNNs deliver useful forecasts.

The objective of supervised learning is to predict something from data. A training set of an output (target) and some input variables is fed to an algorithm that learns to predict the target values. The output may be categorical (classification) or continuous (regression). The task of the algorithm is to

deliver high quality predictions, all by itself, extracting the required knowledge solely from the available data. Neural networks are a popular framework for supervised learning a network-like system of weighted summations and differentiable functions that can learn astoundingly complex things. Usually, variants of gradient descent together with backpropagation (chain rule) are used to find the optimal values of the network weights. This is all simple and intuitive, yet the resulting networks are usually difficult for people to understand. There are so many weights and connections that we just wonder how the system produced the results. The reason for their ever high popularity is simple: they are good. They can learn arbitrarily complex functions, and they often provide excellent predictions for pretty difficult machine learning problems.

RNNs are neural networks for sequential data hereby we apply them to time series. The main idea behind recurrent neural networks is using not only the in-put data, but also the previous outputs for making the current prediction. This idea makes a lot sense — we could build neural networks passing values forward in time. However, such simple solutions usually do not work as expected. They are hard to train and they are forgetful. Rather, we need to have a system with some kind of memory. There are two popular and efficient RNN models that work really well: long short-term memory and gated recurrent unit.

Long short-term memory is a gated memory unit for neural networks. It has 3 gates that manage the contents of the memory. These gates are simple logistic functions of weighted sums, where the weights might be learnt by backpropagation. It means that, even though it seems a bit complicated, the LSTM perfectly fits into the neural network and its training process. Time series modeling is a key area of scientific research. Forecasting the future state of a given sequence is a critical part of real-world problems in a wide range of areas.

There are a number of possible reasons for the underperformance of NNs in the past, one being that individual time series themselves usually are too short to be modelled using complex approaches. Another possibility may be that the characteristics of time series have changed over time, so that even long time series may not contain enough relevant data to fit a complex model. Thus, to model sequences by complex approaches, it is essential that they have adequate length and that they are generated from a comparatively stable system. Also, NNs are further criticised for their black-box. Thus,

forecasting practitioners have traditionally opted for more straightforward statistical techniques.

However, we are now living in the era of big data. Companies have gathered a plethora of data over the years, which contain important information about their business patterns. Big data in the context of time series does not necessarily mean that the individual time series contain lots of data. Rather, it typically means that there are many related time series from the same domain. In such a context, univariate forecasting techniques that consider individual time series in isolation may fail to produce reliable forecasts. They become inappropriate in the context of big data, where a single model can learn from many similar time series simultaneously. By contrast, more complex models such as NNs benefit most from the availability of massive amounts of data.

➤ *RNNs in Time Series Forecasting:*

- **Stock Market Prediction:** RNNs have been used to forecast stock prices based on historical price and volume data. By learning from the patterns and trends in the past, RNNs can make predictions that assist investors in making informed decisions
- **Energy Demand Forecasting:** RNNs have demonstrated success in predicting energy demand, aiding in load balancing and resource allocation. Accurate forecasts enable better planning, cost optimization, and efficient energy distribution.
- **Weather Forecasting:** RNNs have shown promise in weather prediction by analyzing historical weather data. They can capture complex atmospheric patterns and provide forecasts for temperature, precipitation, and other meteorological variables.
- **Sales and Demand Forecasting:** RNNs have been applied to predict future sales and demand patterns in various industries. By considering historical sales data, seasonality, and promotional events, RNNs help businesses optimize inventory management and plan production accordingly.

II. MATERIALS AND METHODS

The systematic assessment of RNN, specifically LSTM, involves modeling scenarios established to represent a typical development progression. The model response to changing geospatial input is assessed in each scenario. This involves changing the temporal resolution parameter to the next smallest increment possible while maintaining the comparison of the same LC data layer across all scenarios. For each change of temporal resolution, the experiments are performed on one of the different datasets containing a different number of classes. The aim is to observe if there are trends in model response on changed temporal resolution despite model optimizations.




Fig 1 Systematic Assessment of RNN

III. RNN MODEL DEVELOPMENT

A baseline RNN model undergoes optimizations to create three modeling scenarios. In this work, a stacked LSTM model comprised of three layers and 32 neurons per layer forms the baseline RNN model. This is because stacked DL models have proven their ability to learn increasingly complex relationships. The input layer is compatible with the one-hot encoded sequences. One-hot encoding involves converting LC class labels to vectors containing zeroes and a single non-zero value (1) at the index corresponding to the class ID. Each input sequence is a matrix of $N \times M$ dimensions corresponding to respective cells across the study area, where N denotes the number of timesteps and M denotes the number of possible categories represented by the one-hot encoded vector. The output layer of the respective models produces an $M \times 1$ vector containing the probabilities of each LC class forecasted to occur at the next timestep for a single cell. The position in the output vector featuring the highest probability is selected, with the position in the vector corresponding to the forecasted class label.

The use of recurrent neural networks (RNNs) in time series forecasting has gained significant attention in recent years. This article delves into the potential of RNNs as a powerful tool for predicting future values in time series data. By leveraging their ability to capture temporal dependencies, RNNs offer promising solutions to address the challenges of accurate forecasting in various domains.

➤ *RNN Architecture:*

The article begins by explaining the architecture of recurrent neural networks. Unlike feed-forward neural networks, RNNs have feedback connections that enable the network to retain information about previous inputs. This characteristic allows RNNs to process sequential data and learn from its temporal patterns. The key component of RNNs is the hidden state, which serves as memory to store and propagate information throughout the sequence.

➤ *Long Short-Term Memory (LSTM):*

The article further explores a specific type of RNN called Long Short-Term Memory (LSTM). LSTMs are designed to address the vanishing gradient problem by incorporating specialized memory cells. These cells possess input, forget, and output gates, enabling LSTMs to capture long-term dependencies in the time series data effectively. The ability to

retain relevant information over long time horizons makes LSTMs particularly valuable for forecasting tasks.

➤ *Training and Evaluation:*

The article outlines the training and evaluation process for RNNs in time series forecasting. It emphasizes the significance of proper data preparation, including normalization and splitting the data into training and testing sets. The training process involves optimizing model parameters using techniques like backpropagation through time (BPTT) and gradient descent. The evaluation of RNN models typically involves metrics such as mean absolute error (MAE) or root mean square error (RMSE) to assess the accuracy of predictions.

IV. CASE STUDIES AND RESULTS

To display the effectiveness of RNNs in time series forecasting, the article presents several case studies. These studies span diverse domains such as financial markets, energy demand, and weather prediction. The results demonstrate that RNNs, particularly LSTMs, consistently outperform traditional forecasting methods. The ability of RNNs to capture complex patterns and adapt to changing dynamics in time series data contributes to their superior performance.

➤ *Advantages of RNNs in Time Series Forecasting:*

- **Handling Variable-Length Sequences:** Unlike traditional forecasting methods, RNNs can handle time series data with variable-length sequences. This flexibility allows them to effectively model and predict future values even in cases where the length of the sequence varies.
- **Capturing Long-Term Dependencies:** RNNs, particularly LSTMs, excel at capturing long-term dependencies in time series data. By maintaining a memory of past inputs through their hidden state, they can capture and leverage patterns that span across multiple time steps, enabling accurate forecasting even for complex and non-linear time series.
- **Non-Linear Modeling:** Time series data often exhibits non-linear patterns and relationships. RNNs, with their ability to model non-linear mappings, offer a powerful solution for capturing and representing these complex dynamics, thereby enhancing the accuracy of predictions.
- **Adaptability to Changing Patterns:** Time series data frequently undergoes changes in patterns and underlying dynamics. RNNs can adapt to such changes by continually updating their model parameters based on new input data. This adaptability allows them to provide reliable forecasts even in dynamic and evolving time series scenarios.

V. LIMITATIONS AND CHALLENGES

➤ *Computational Complexity*

RNNs, especially LSTMs, can be computationally expensive to train, particularly when dealing with large-scale

time series datasets. Training RNNs may require significant computational resources and time, making it essential to consider the computational constraints when applying them to real-world forecasting problems.

➤ *Overfitting*

RNNs, like other machine learning models, are susceptible to overfitting, where the model becomes too closely tuned to the training data and performs poorly on unseen data. Regularization techniques such as dropout and early stopping can help mitigate overfitting in RNNs.

➤ *Need for Sufficient Training Data*

RNNs typically require a sufficient amount of training data to capture meaningful patterns and generalize well to unseen data. In some cases, limited or insufficient training data can lead to suboptimal performance of RNNs in time series forecasting tasks.

➤ *Recurrent Neural Networks (RNNs)*

have significantly advanced time series forecasting, but their potential for future advancements remains promising. Here, we explore some potential directions for the future development and enhancement of RNNs in time series forecasting:

➤ *Improved Architectures*

Researchers are actively working on developing novel RNN architectures that can enhance the modeling capabilities of RNNs for time series forecasting. Variants such as Gated Recurrent Units (GRUs) and Attention-based RNNs are gaining attention for their ability to capture complex temporal patterns and focus on relevant information.

➤ *Hybrid Approaches:*

Combining RNNs with other advanced techniques such as Convolutional Neural Networks (CNNs) or Transformers can leverage the strengths of multiple architectures for time series forecasting. Hybrid models can capture both local and global patterns, allowing for more accurate predictions and improved generalization.

➤ *Handling Uncertainty:*

Accounting for uncertainty is essential in time series forecasting. Bayesian RNNs and Variational Autoencoders (VAEs) are being explored to capture uncertainty estimates in predictions. These approaches can provide probabilistic forecasts and confidence intervals, aiding decision-making under uncertainty.

➤ *Attention Mechanisms:*

Further advancements in attention mechanisms within RNNs can enable more focused and adaptive modeling. Attention mechanisms allow the model to weigh different parts of the time series differently, giving more emphasis to relevant patterns and reducing the impact of noise or irrelevant information.

➤ *Transfer Learning and Pretrained Models:*

Transfer learning, where knowledge from one domain is transferred to another, can be leveraged in time series forecasting. Pretrained models on large-scale time series datasets can capture general patterns, reducing the need for extensive training on specific datasets and improving efficiency.

➤ *Handling Big Data and Parallel Computing:*

Efficient processing of large-scale time series data is crucial. Future advancements in RNNs will likely focus on leveraging parallel computing techniques and distributed systems to handle big data more effectively. This includes exploring strategies for distributed training and inference to accelerate computations.

➤ *Interpretability and Explainability:*

Enhancing the interpretability of RNNs for time series forecasting is an ongoing research direction. Techniques such as attention visualization, layer-wise relevance propagation, and saliency maps can provide insights into the model's decision-making process, making it more transparent and explainable.

➤ *Online Learning and Adaptive Forecasting:*

RNNs can be further developed to accommodate online learning scenarios where new data becomes available continuously. Adaptive forecasting techniques that dynamically update the model's parameters based on new observations can be explored to handle evolving time series data.

The future of RNNs in time series forecasting holds tremendous potential for advancements. Improved architectures, hybrid approaches, uncertainty handling, attention mechanisms, transfer learning, parallel computing, interpretability, and adaptive forecasting are all areas that researchers are actively exploring. As these developments continue, RNNs are expected to provide even more accurate and reliable predictions, empowering decision-makers in various domains to make informed choices based on the analysis of time series data.

VI. RECOMMENDATION

Based on the information provided, here are some recommendations related to recurrent neural networks (RNNs) and time series forecasting:

Familiarize yourself with RNN architectures: Gain a solid understanding of the different RNN architectures, particularly Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs). Explore how these architectures handle temporal dependencies and capture patterns in time series data.

Experiment with open-source RNN frameworks: Utilize open-source software frameworks for RNNs, such as TensorFlow or PyTorch. These frameworks provide pre-built

RNN modules and libraries that can facilitate the development and training of RNN models for time series forecasting.

Preprocess your time series data: Preprocessing plays a crucial role in preparing time series data for RNNs. Consider techniques such as normalization, missing value imputation, and feature engineering to enhance the quality and relevance of the input data.

Split your data into training and testing sets: Ensure that you have separate datasets for training and testing your RNN models. This division allows you to evaluate the performance and generalization capabilities of your models on unseen data, providing a more accurate assessment of their forecasting accuracy.

Optimize model parameters: Experiment with different hyperparameters and model configurations to optimize the performance of your RNN models. Techniques like grid search or Bayesian optimization can help you systematically explore the parameter space and find the best settings for your specific forecasting task.

Evaluate model performance: Use appropriate evaluation metrics such as mean absolute error (MAE) or root mean square error (RMSE) to assess the accuracy of your RNN models. Compare the performance of your RNN models against traditional forecasting methods to validate their effectiveness.

Consider ensemble approaches: Explore ensemble methods that combine multiple RNN models or different forecasting techniques to improve the accuracy and robustness of your predictions. Ensemble approaches can help mitigate individual model biases and provide more reliable forecasts.

Stay updated on advancements: Keep up-to-date with the latest research and developments in RNNs and time series forecasting. Follow conferences, journals, and online communities dedicated to machine learning and forecasting to stay informed about new techniques, architectures, and best practices.

Be mindful of computational resources: Take into account the computational requirements of training and deploying RNN models, especially when dealing with large-scale time series datasets. Consider using distributed computing or cloud resources to handle the computational complexity efficiently.

Document and share your findings: As you experiment and develop RNN models for time series forecasting, document your methodologies, findings, and insights. Consider sharing your experiences through blog posts, research papers, or online communities to contribute to the collective knowledge in the field. Remember that these recommendations serve as general guidelines. It's important to adapt and tailor them to your specific domain, dataset, and forecasting objectives.

VII. CONCLUSION

In conclusion, recurrent neural networks (RNNs) have emerged as powerful tools for time series forecasting. Their ability to capture temporal dependencies and model complex patterns in sequential data makes them well-suited for predicting future values in time series datasets. By leveraging RNN architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs), researchers and practitioners have made significant advancements in the field of time series forecasting. To effectively utilize RNNs for time series forecasting, it is crucial to understand the underlying concepts and techniques associated with these models. Familiarizing yourself with RNN architectures, preprocessing techniques, and evaluation metrics is essential for building accurate and reliable forecasting models. Open-source frameworks like Tensor Flow and PyTorch provide valuable resources for implementing RNNs and experimenting with different configurations.

When working with time series data, it is important to preprocess the data appropriately, splitting it into training and testing sets, and optimizing the model parameters. Evaluating the performance of RNN models using appropriate metrics and comparing them against traditional forecasting methods helps assess their effectiveness. Additionally, exploring ensemble approaches and staying updated with the latest advancements in the field can further enhance the accuracy and robustness of time series forecasts. It's worth noting that the computational resources required for training and deploying RNN models can be significant, particularly when dealing with large-scale datasets. Considering the availability of distributed computing or cloud resources can help address these computational challenges efficiently.

Ultimately, by following best practices, documenting findings, and sharing knowledge within the machine learning and forecasting communities, we can continue to push the boundaries of time series forecasting with recurrent neural networks. RNNs offer immense potential for accurately predicting future values in various domains, enabling us to make more informed decisions and gain valuable insights from time-dependent data.

REFERENCES

- [1]. Bharadiya , J. P., Tzenios, N. T., & Reddy , M. (2023). Forecasting of Crop Yield using Remote Sensing Data, Agrarian Factors and Machine Learning Approaches. *Journal of Engineering Research and Reports*, 24(12), 29–44. <https://doi.org/10.9734/jerr/2023/v24i12858>
- [2]. Bharadiya, J. (2023). Artificial Intelligence in Transportation Systems A Critical Review. *American Journal of Computing and Engineering*, 6(1), 34 - 45. <https://doi.org/10.47672/ajce.1487>
- [3]. Bharadiya, J. . (2023). A Comprehensive Survey of Deep Learning Techniques Natural Language Processing. *European Journal of Technology*, 7(1), 58 - 66. <https://doi.org/10.47672/ejt.1473>
- [4]. Bharadiya, J. . (2023). Convolutional Neural Networks for Image Classification. *International Journal of Innovative Science and Research Technology*, 8(5), 673 - 677. <https://doi.org/10.5281/zenodo.7952031>
- [5]. Bharadiya, J. . (2023). Machine Learning in Cybersecurity: Techniques and Challenges. *European Journal of Technology*, 7(2), 1 - 14.
- [6]. Bharadiya, J. . (2023). The Impact of Artificial Intelligence on Business Processes. *European Journal of Technology*, 7(2), 15 - 25. <https://doi.org/10.47672/ejt.1488>
- [7]. Bharadiya, J. P. (2023, May). A Review of Bayesian Machine Learning Principles, Methods, and Applications. *International Journal of Innovative Science and Research Technology*, 8(5), 2033-2038. DOI: <https://doi.org/10.5281/zenodo.8002438>
- [8]. Bharadiya, J. P. (2023, May). A Tutorial on Principal Component Analysis for Dimensionality Reduction in Machine Learning. *International Journal of Innovative Science and Research Technology*, 8(5), 2028-2032. DOI: <https://doi.org/10.5281/zenodo.8002436>
- [9]. Desell, T.: Large scale evolution of convolutional neural networks using volunteer computing. *CoRR abs/1703.05422* (2017). <http://arxiv.org/abs/1703.05422>
- [10]. Eck, D., Schmidhuber, J.: A first look at music composition using lstm recurrent neural networks. *Istituto Dalle Molle Di Studi Sull Intelligenza Artificiale* 103, (2002)
- [11]. ElSaid, A., El Jamiy, F., Higgins, J., Wild, B., Desell, T.: Optimizing long short-term memory recurrent neural networks using ant colony optimization to predict turbine engine vibration. *Appl. Soft Comput.* 73, 969–991 (2018)
- [12]. Findell, K.L.; Berg, A.; Gentine, P.; Krasting, J.P.; Lintner, B.R.; Malyshev, S.; Shevliakova, E. The impact of anthropogenic land use and land cover change on regional climate extremes. *Nat. Commun.* 2017, 8, 1–10
- [13]. Gharbia, S.S.; Alfatah, S.A.; Gill, L.; Johnston, P.; Pilla, F. Land use scenarios and projections simulation using an integrated GIS cellular automata algorithms. *Model. Earth Syst. Environ.* 2016, 2, 1–20.
- [14]. Nallamothu, P. T., & Bharadiya, J. P. (2023). Artificial Intelligence in Orthopedics: A Concise Review. *Asian Journal of Orthopaedic Research*, 6(1), 17–27. Retrieved from <https://journalajorr.com/index.php/AJORR/article/view/164>
- [15]. Schreinemachers, P.; Berger, T. Land use decisions in developing countries and their representation in multiagent systems. *J. Land Use Sci.* 2006, 1, 29–44.
- [16]. Werbos, P.J.: Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78(10), 1550–1560 (1990)