International Journal of Advance Interdisciplinary Research
Vol. 2, Special Issue 1, (Jan-March) 2026, pp. 257-263, e-ISSN: 3107-913X
DOI: https://doi.org/10.66095/ijair.2026.v2.S1.28

RNN-Based Forecasting for Dynamic Portfolio Rebalancing: Deep Learning Approach to Tactical Asset Allocation

Yashika Mishra and Richa Mehrotra
Department of Mathematics, Amity School of Applied Sciences, Amity University Lucknow Campus, Lucknow (U.P.) – 226028, India
***E-mail: yashikamishra262@gmail.com, rmehrotra@lko.amity.edu**

Abstract: Deep learning is becoming increasingly popular in finance and time series prediction, including attempts to predict asset prices and the best weights for a trading portfolio. Portfolio rebalancing is an important concept in finance to ensure an investor's portfolio reflects the risk level and time horizon they are comfortable with. Classical rebalancing methods, e.g., the equal weight (1/N) approach, are predominantly static in nature and are based on fixed timings or thresholds. These approaches, however, fail to incorporate the time dependent and dynamic nature of financial markets in which asset values are affected by macroeconomic events as well as other stochastic events that change over time.

Recurrent Neural Networks (RNNs) for tactical asset allocation and dynamic portfolio rebalancing are investigated in this work. This work uses a data driven approach that uses Recurrent Neural Networks (RNNs) to forecast short term asset returns and make rebalancing decisions based on these forecasts. The RNN model generates the dynamic rebalancing signal, which adapts to portfolio weights and also times the rebalancing decisions. This work uses exchange traded funds (ETFs) to construct a portfolio whose weights are periodically adjusted according to predicted future returns. Our results demonstrate that the RNN based strategy outperforms traditional equal weight benchmarks in terms of cumulative returns and Sharpe ratio, while maintaining acceptable risk levels.

Keywords: Recurrent Neural Networks, RNNs, Portfolio Optimization

1. Introduction

Portfolio rebalancing plays an essential role in the management of portfolios, by continuously adjusting portfolio weights in response to changes in market conditions, asset performance, and investor objectives. Market performance can change asset values and an investor's goal over a period of time, prompting frequent adjustments, selling overvalued assets and acquiring undervalued assets. Projected returns, asset correlations and risk levels evolve and change over time, which makes the financial market exceptionally volatile and uncertain. Classical portfolio rebalancing measures, such as equal weight (1/N) relies on the pre-determined rules. These static strategies help control transaction costs and maintain diversification, but they do not fully account for the sudden market shifts and change in the investor objectives, failing to account for non-linear dependencies in asset returns.

To address these limitations of static techniques, dynamic approach is introduced, which integrates forecasting models and deep learning algorithms to adjust the portfolio allocations proactively rather than reactively. The recent advancement in the field of machine learning and deep learning has strengthened these dynamic rebalancing approaches, enabling the investors to aim to manage risks more effectively, seize short-term opportunities, and prevent portfolios from deviating too far from their intended risk-return profiles. By applying dynamic techniques, it enables tactical approach to asset allocation in portfolio rebalancing.

Among these dynamic techniques, Recurrent Neural Networks (RNNs) have emerged as especially promising due to their capacity to represent sequential dependencies and capture non-linear dynamics in time series data. This work focuses on a deep learning framework that employs Recurrent Neural Networks (RNNs) to simulate sequential dependencies in financial time series and forecast short-term asset returns. Unlike prior work, our model jointly learns portfolio weights and rebalancing timing.

2. Past Works

Rule-based and optimization approaches, mean-variance optimization (MVO), threshold-based rules, and fixed-interval rebalancing are traditionally the popular choices for addressing the problem of portfolio rebalancing. These methods are interpretable and computationally tractable, but typically depend on static assumptions about the structure of covariance and return distributions that cannot be easily adapted to rapidly changing financial markets.

Recurrent neural networks (RNNs) and long short-term memory (LSTM) network have been prevailing in financial time-series modelling. There has been more recent work on generalizing the usage of Fischer-style sequence models to portfolio-level applications [3]. Wu introduced a dynamic portfolio optimization method using LSTM that boosts both cumulative returns and Sharpe ratios compared to traditional benchmarks [1]. He used expected asset returns to tweak portfolio weights. His research shows just how important it is for recurrent architectures to pick up on patterns over time, it makes a real difference when making portfolio decisions.

Recent research has started looking into end-to-end learning for portfolio management. Yang et al. [2], along with follow up studies in *IEEE Access and Expert Systems with Applications*, show that deep reinforcement learning can actually adjust portfolio weights spontaneously as market conditions change. But most of these deep learning and reinforcement learning methods still stick to a fixed rebalancing schedule, even though they're supposed to be flexible. This work steps in to change that. It introduces a recurrent neural network that handles both short-term return predictions and rebalancing decisions at the same time. So, it pushes deep learning based portfolio strategies forward by letting them adapt their rebalancing in real time, not just at preset intervals.

3. Methodology

3.1 Data Description and Preprocessing

The U.S. sector-based exchange-traded funds (ETFs) XLK (Technology), XLF (Financials), XLY (Consumer Discretionary), XLE (Energy), and XLV (Health Care) are used in the empirical analysis. These ETFs are chosen to preserve adequate liquidity and data availability while representing a variety of eco-nomic sectors. Yahoo Finance provides daily closing price information from January 2020 to December 2025. Daily prices are resampled to weekly frequency using Friday closing prices to eliminate high-frequency noise. Next, each ETF's weekly log returns are calculated. To provide numerical stability and enable neural network optimization, the return series are normalized using Min–Max scaling prior to model training. The scaled return matrix serves as the primary input for further sequence modelling.

3.2 Sequence Construction and Rebalancing Signal

The model is trained using rolling sequences of historical data in order to capture temporal relationships in asset returns. A three-month lookback window that is frequently utilized in tactical asset allocation is reflected in each input sequence, which is made up of the prior weeks' normalized returns across all ETFs. Two prediction targets are established for every sequence: a binary rebalancing signal and the return vector for the following week. There is no look-ahead bias because the rebalancing signal is exclusively built using historical market data. Specifically, a rolling volatility measure is derived using previous returns, and rebalancing is triggered when current average volatility exceeds a dynamically defined threshold based on historical dispersion.

3.3 RNN Architecture and Learning Framework

Return forecasting and rebalancing behaviour are concurrently learned using a multi-output Recurrent Neural Network (RNN). The network is composed of a dropout layer to prevent overfitting after a *SimpleRNN* layer with 64 hidden units and a hyperbolic tangent activation function. Two output heads receive the shared hidden representation: a sigmoid-activated layer for calculating the likelihood of rebalancing and a linear layer for forecasting next-week ETF returns. The Adam optimizer is used to train the model, employing bi-nary cross-entropy loss for the rebalancing signal and mean squared error loss for return prediction.

3.4 Portfolio Evaluation

A rolling out-of-sample simulation that is started with a fixed capital base is used to assess portfolio performance. Returns are compounded weekly based on realized asset returns and the model-generated weights. The same asset universe and evaluation period are used to create a static equal-weight portfolio for comparison. For each fold of the cross-validation procedure, the dynamically rebalanced portfolio was evaluated in terms of Compound Annual Growth Rate (CAGR), Sharpe Ratio, and Maximum Drawdown (MDD) [6]. A thorough comparison between the suggested dynamic method and the benchmark is made possible by the final findings, which are presented as averages across cross-validation folds.

4. Result and Analysis

The proposed RNN-based dynamic portfolio rebalancing framework was evaluated using a 5-fold cross-validation setup on weekly returns of sectoral ETFs *(XLK, XLF, XLY, XLE, XLV)*. The model simultaneously learned to forecast short-term asset returns and generate a binary rebalancing signal.

The average classification accuracy with all folds was **0.9119**, thus indicating that the model has more than a 91% true positive rate of detecting rebalance periods. Such high forecast accuracy indicates the good ability of RNN to capture temporal dependencies in sector-level ETF returns and to differentiate between stable and volatile market environments. This is crucial for portfolio management where too frequent rebalancing leads to higher transaction costs and can also lead to delayed rebalancing, which can subject the portfolio to more risk.

International Journal of Advance Interdisciplinary Research
Vol. 2, Special Issue 1, (Jan-March) 2026, pp. 257-263, e-ISSN: 3107-913X
DOI: https://doi.org/10.66095/ijair.2026.v2.S1.28

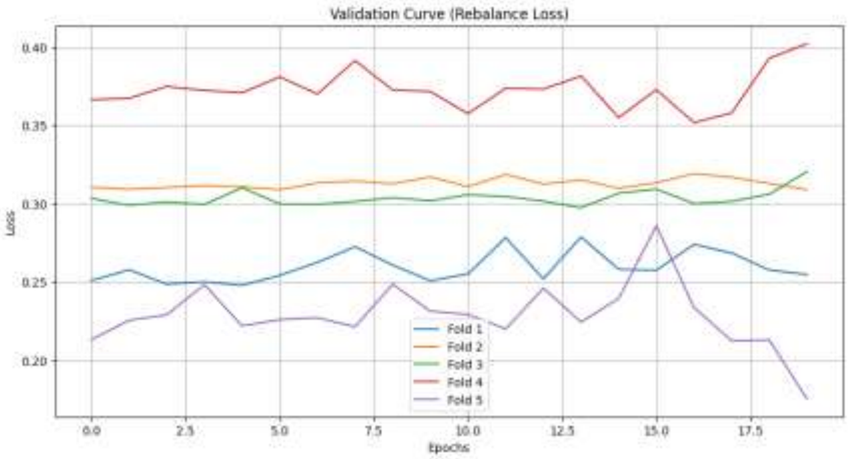

Figure 1: Validation Curve (Rebalancing Loss across Folds)

The validation loss curves of the rebalancing output are presented in Figure 1 for the five folds. Some variance appears in the cross-validation folds, notably in fold four, which has a slightly higher loss value. But overall validation losses are fairly constant with no indication of divergence or runaway overfitting. The validation curves behave fairly consistently suggesting that the selected network architecture is enough for this dataset's size and complexity.

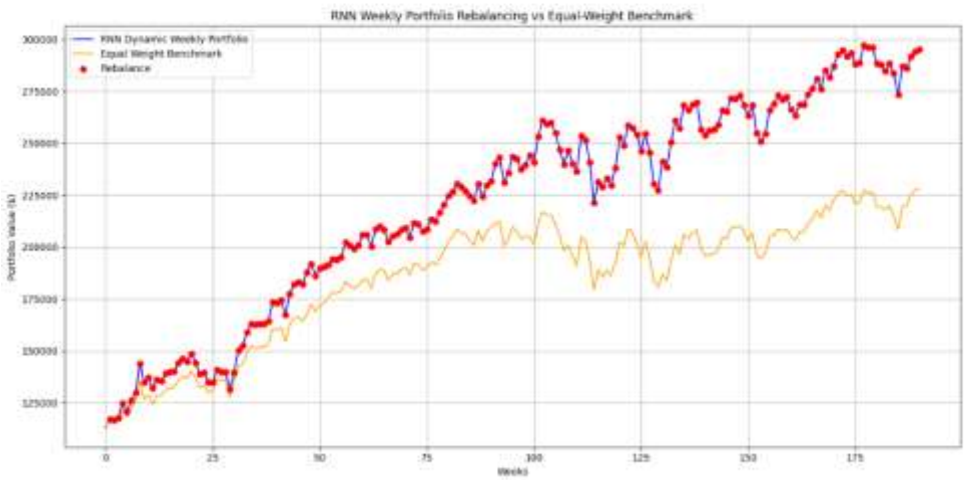

Figure 2: RNN Weekly Portfolio vs Equal-Weight Benchmark

In Figure 2, the RNN-driven dynamic portfolio surpasses the equal-weight benchmark throughout the entire investment period. The two portfolios show upward trends but the dynamically rebalanced portfolio achieves higher capital growth and responds better to market changes. The RNN-based strategy reaches a much higher portfolio value by simulation end because it shows the economic impact of this method.

International Journal of Advance Interdisciplinary Research
Vol. 2, Special Issue 1, (Jan-March) 2026, pp. 257-263, e-ISSN: 3107-913X
DOI: https://doi.org/10.66095/ijair.2026.v2.S1.28

In figure 2, the red markers denote the rebalancing of the events triggered by the RNN model. The events show specific time patterns because they occur during periods when market volatility increases which include market drawdowns. This behavior matches financial knowledge because rebalancing works best when market conditions switch from one state to another instead of during times of market stability. The model achieves better outcomes because its selective rebalancing feature allows it to choose when to rebalance instead of following set rebalancing times.

The equal-weight benchmark offers investors a simple and reliable investment option, but it is not flexible enough. The system does not profit from brief increases in sector performance, and it continues to be vulnerable to underperforming industries for long stretches of time. Because active portfolio rebalancing generates higher value when investors hold their investments for longer periods of time, the two asset paths begin to diverge. While the system does not profit from short-term improvements in sector performance, it remains vulnerable to underperforming industries for long periods of time.

Quantitative performance metrics averaged across all folds are summarized in Table 1.

Table 1 Performance metrics of the model averaged across all folds.

Average Metric	RNN Portfolio
Sharpe Ratio	1.176
Max Drawdown	-11.37%
CAGR	23.71%

CAGR is a measure of the rate of investment return that must be made on the beginning balance of an investment account to end up with the final ending balance in its respective future period, assuming the profits were reinvested at the end of each period. The Sharpe Ratio was utilized in assessing performance rather than risk by examining weekly portfolio returns and estimating the excessive returns for each unit of risk.

Maximum Drawdown was calculated as the greatest reduction in portfolio value between peaks over any of the assessment periods. This is particularly important for portfolio

International Journal of Advance Interdisciplinary Research
Vol. 2, Special Issue 1, (Jan-March) 2026, pp. 257-263, e-ISSN: 3107-913X
DOI: https://doi.org/10.66095/ijair.2026.v2.S1.28

management as it shows how successfully the strategy can fare in turbulent market conditions whilst limiting significant financial losses. In addition, the study team assessed each period of consolidation by computing all three input performance metrics *(CAGR, Sharpe Ratio and Maximum Drawdown),* and then averaging these into a true measure of portfolio performance. The analysis of financial performance over a series of evaluation periods eliminates the ability of the metric to be manipulated by one positive period and also provides continued measures regardless of market conditions.

5. Conclusion

In this work, we considered a tractable model using recurrent neural networks for dynamic portfolio rebalancing in a multi-asset setting. In summary, the results indicate that introducing temporal dependencies and adaptive rebalancing decisions through recurrent models is a promising approach for tactical asset allocation. Although the analysis is confined to a small number of sector ETFs and single recurrent architecture, results provide groundwork for further works where broader set of assets, different types of recurrent networks and explicit transaction costs may be included. The framework proposed in this paper is a move towards the development of more flexible, data-based portfolio management methodologies that capture more appropriately the time-changing character of financial markets.

6. References

[1] Songwei Wu (2023) Dynamic Optimization f Investment Portfolio Based on Self-Attention Mechanism – LSTM Model Prediciton. Journal of Data Science, p 45-59

[2] Yang, H., Liu, X., Wang, J. (2020) Deep reinforcement learning for automated stock trading and portfolio management. IEEE Access, 8, 226237–226248

[3] Fischer, T., Krauss, C. (2018) Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669.

[4] Lim, B., Zohren, S. (2021) Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A, 379(2194), 20200209.

[5] Zhou, Y., Liang, J., & Dong, G. (2023) Learning dynamic portfolio rebalancing strategies using recurrent neural networks. Applied Soft Computing, 132, 109859.

[6] Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2022) Deep hedging. Quantitative Finance, 22(3), 1–23.