Journal of Advanced Robotics, Autonomous Systems and Human-Machine Interaction
ISSN: 3070-6491

PRIME OPEN ACCESS

Volume 2, Issue 3

Research Article

Date of Submission: 24 June, 2026
Date of Acceptance: 24 July, 2026
Date of Publication: 10 Aug, 2026

Hedge Fund Replication with Deep Neural Networks and Generative Adversarial Networks

Devin M. Chatterji*

Illinois Institute of Technology, Stuart School of Business, USA

***Corresponding Author:** Devin M. Chatterji, Illinois Institute of Technology, Stuart School of Business, USA.

Citation: Chatterji, D. M. (2026). Hedge Fund Replication with Deep Neural Networks and Generative Adversarial Networks. *J Adv Robot Auton Syst Hum Mach Interact, 2*(3), 01-15.

Abstract

This study investigates whether modern deep learning techniques can improve hedge fund replication by modeling the nonlinear, dynamic, and regime-dependent relationships that characterize hedge fund returns more effectively than traditional linear factor models. Monthly returns from the Hennessee Hedge Fund Index (HHFI) and twenty-one hedge fund strategies were analyzed using fixed-weight and rolling-window regression models, deep neural networks (DNN), recurrent neural networks (RNN), long short-term memory networks (LSTM), and generative adversarial networks (GAN). An expanded universe of sixty liquid financial factors was incorporated into the generative adversarial network framework to enhance factor selection and portfolio construction. Model performance was evaluated using out-of-sample tracking error together with paired statistical tests to determine whether observed differences in performance were statistically significant. The deep learning models consistently outperformed traditional regression approaches across the majority of hedge fund strategies. Among the conventional neural network architectures, recurrent neural networks generally achieved the lowest tracking error. The generative adversarial network models produced the strongest overall performance, with the six-factor and fifteen-factor configurations delivering the most accurate replication results. In selected hedge fund strategies, tracking error was reduced by as much as ninety-six percent relative to traditional regression models, and the improvements over regression, recurrent neural networks, and long short-term memory networks were statistically significant. These findings demonstrate that adversarial learning provides a more effective framework for estimating the nonlinear and time-varying factor exposures that drive hedge fund returns than conventional linear methods. The proposed approach offers a scalable, transparent, and cost-effective methodology for constructing hedge fund replication portfolios while demonstrating the broader value of deep learning techniques for financial modeling, portfolio management, and asset allocation.

Keywords: Hedge Fund Replication, Deep Neural Networks, Generative Adversarial Networks, Artificial Intelligence, Portfolio Optimization

Introduction

Hedge fund replication aims to create a portfolio that mimics hedge fund strategy returns. However, transparency is a major challenge because hedge funds often operate like black boxes. Despite their diversification benefits, hedge funds have leverage, investment restrictions, low transparency, reduced liquidity, and high fees compared to traditional asset management.

Given high management and performance fees, investors increasingly question what they are paying for. Institutional investors are particularly wary of paying alpha fees for beta returns, typically 2% management and 20% incentive fees. Fund managers argue that their fees are justified by their skills, leading to higher returns, but assessing this skill is challenging. Even with an impressive track record, there is no consensus on whether it persists and contributes to high fees in this competitive industry.

Hedge fund replication products introduced by investment banks and companies offer hedge fund-like returns at lower costs. These products aim to provide similar returns using rule-based strategies and liquid assets, as seen with the "Merrill Lynch Factor Index" and "Goldman Sachs Absolute Return Tracker Index." Hedge fund replication appeals to traditional fund managers seeking to provide hedge fund-like performance with greater transparency, liquidity, and accessibility. Replication portfolios must consist of liquid, transparent, scalable, and cost-efficient instruments across various asset classes.

The extant literature on factor-based hedge fund replication remains constrained by its reliance on predominantly linear specifications, which are theoretically ill-equipped to approximate the nonlinear, state-dependent, and dynamically evolving factor exposures characteristic of hedge fund strategies. Consequently, persistent uncertainty remains regarding both the identification of relevant factors and stability of their functional relationships over time. This study extends the theoretical framework of replication by introducing deep neural networks and generative adversarial networks as nonlinear universal approximators applied to hedge fund returns, empirically demonstrating that these architectures yield materially improved replication performance relative to canonical linear models.

The remainder of this paper is organized as follows. Section II reviews the literature on hedge fund replication. Section III describes the research methodology, including the dataset, regression models, DNNs, RNNs, LSTMs, GANs, and evaluation metrics. Section IV presents the empirical results, including performance comparisons across all models and statistical significance testing. Section V discusses the findings, interprets the performance of various deep learning architectures, and provides the economic intuition underlying the superior performance of the GAN framework. Finally, Section VI summarizes the principal findings, contributions, and directions for future research.

Motivation

The interest in hedge fund replication began in the late 1990s. Several authors have explored replication approaches. Karavas et al. (2003) extend this research to five hedge fund strategies and examine the replication of equally-weighted portfolios using a rolling 24-month window [1-6].

This study is motivated by Hasanhodzic and Lo's (2006) regression of monthly returns for 1610 hedge funds against six factors. They developed two types of clones: fixed-weight and rolling-window clones. They concluded that these methods worked well for certain hedge-fund categories but not all. They also suggested that more advanced nonlinear methods could improve performance and goodness-of-fit.

Subsequent studies on factor-based hedge fund replication build on Hasanhodzic and Lo's (2006) methodology and findings. The consensus is that, to improve replication products, one must

• Develop a dynamic model to capture the nonlinear nature of hedge fund returns.
• Avoid using a set of predefined assets.
• Identify the right set of factors for changing market conditions.
• Choose the strategies to be cloned.

Researchers have begun to leverage deep neural networks to address these challenges. DNNs can detect and exploit interactions in data that traditional financial theories cannot (Heaton et al., 2016). They model complex nonlinearities and interpret nonparametric shocks in financial econometrics. Consequently, machine learning (ML) and deep learning models are increasingly studied for various financial applications, including asset pricing, credit, equity forecasting, fixed income and trading [7-41].

Gu, Kelly, and Xiu (2020) performed a comparative analysis of ML methods, finding that these techniques offer better predictions of expected returns than traditional forecasting methods. Notable outcomes from neural network (NN) research include the following.

• NN-based strategies outperform others.
• DNNs are strong classifiers when trained across multiple markets.
• NNs predict financial time series movements even trained only on plain time series data.
• LSTM networks outperform memory-free methods such as random forest and logistic regression.
• Key factors differ from those in linear models when accounting for nonlinearity.
• NNs using relevant features achieve the best results.
• RNNs surpass classical feed-forward models.
• Complex models like LSTMs could boost forecasting power.

This study focuses on applying GANs to hedge fund replication. Since the development of GANs in 2014, few studies have applied them to finance research. Zhou et al. (2018) used a combination of LSTM and convolutional neural networks (CNNs) for adversarial training to predict stock prices, achieve better accuracy, and reduce forecast errors. Koshiyama et al. (2019) employed Conditional GANs (cGANs) for trading strategy calibration and found that cGANs outperformed traditional methods when they failed to generate alpha.

Motivated by the persistent limitations of linear factor models in capturing the nonlinear, regime-dependent structure of hedge-fund payoff functions, this study evaluates whether adversarial learning can offer a more expressive replication framework. Specifically, it asks, Can GAN-based factor selection provide a superior approximation to hedge-fund payoff functions compared with the linear factor models traditionally employed in portfolio theory?

Methods

Data

This study utilizes monthly return data from the Hennessee Hedge Fund Index (HHFI), together with returns from 21 individual hedge fund strategies spanning January 2003 through December 2013. The HHFI represents approximately 3,500 hedge funds and serves as the benchmark for this study. Six liquid financial risk factors were initially selected following Hasanhodzic and Lo (2006): the MSCI EAFE Index, Barclays Aggregate Bond Index, S&P GSCI Commodity Index, S&P 500 Index, Russell 2000 Index, and NASDAQ Index. These factors were chosen because they provide broad exposure to equity, fixed income, commodities, size, and international markets, while remaining liquid and readily investable. Table 1 lists the 21 hedge fund strategies.

21 Hedge Fund Strategies		
Asia Pacific	Fixed Income	Market Neutral
Convertible Arbitrage	Growth	Merger Arbitrage
Distressed	Healthcare and Biotech	Multiple Arbitrage
Emerging Markets	High Yield	Opportunistic
Europe	International	Short Biased
Event Driven	Latin America	Technology
Financial Equities	Macro	Value

Table 1: Hedge Fund Strategies

For the GAN analysis, the factor universe was expanded to 60 liquid exchange-traded funds representing multiple asset classes, sectors, geographic regions, and investment styles. Candidate factors were required to have inception dates prior to December 31, 2002, complete return histories over the study period, and an average daily trading volume exceeding $50 million.

Baseline Regression Models

I run fixed, 24-month rolling, and 36-month rolling regressions on the dataset to establish the baseline, as has been performed in previous studies. In-sample forecasts were obtained for 12 months, from January 2013 through December 2013. Only the previous 120 months were used for each monthly forecast for 2013. I then run an out-of-sample forecast of returns at time t+1, January 2014, using these regression models.

The regression model, which I use for the fixed, 24-month, and 36-month rolling regressions, has the same format as Hasanhodzic and Lo (2006).

$$R_{it} = \alpha_i + \beta_{i1}\text{RiskFactor}_{1t} + \cdots + \beta_{iK}\,\text{RiskFactor}_{Kt} + \varepsilon_{it} \quad (1)$$

Here, K = 6 because these are the six risk factors. MSCI EAFE (USD) Price Index, Barclays Aggregate Bond Index (BOND), S&P GSCI Commodity Total Return Index (CMDTY), S&P 500 Index, Russell 2000 Index, and NASDAQ Index. These six factors were selected for two reasons:

• They provide a reasonably broad cross-section of risk exposures for typical hedge funds (stocks, bonds, currencies, commodities, credit, and size) and

• Each of the factor returns can be realized through relatively liquid instruments.

This model decomposes a hedge fund strategy's return R_{it} into several components, as opposed to the returns of an index, as in previous studies. Forecast accuracy was evaluated using tracking error, defined as the standard deviation of active returns relative to the target strategy, relative to both the HHFI benchmark and each individual hedge fund strategy.

Deep Neural Networks

For this study, I create five models (NN1-NN5), each with an increasing number of hidden layers. I estimate a sequence of feed-forward DNNs that use six monthly factor returns as inputs and produce one-step-ahead forecasts of hedge fund strategy returns. The model depth was gradually increased to evaluate the marginal benefit of deeper architectures. Each model was trained using rolling windows that mirror the regression framework to ensure comparability across the methods. Forecast quality is evaluated using the out-of-sample tracking error.

Recurrent Neural Networks and Long Short-Term Memory Networks

RNNs and LSTMs were evaluated to incorporate temporal dependence in financial returns. Both models used the same six factor inputs as the feed-forward networks while maintaining sequential information through recurrent hidden states. The models were trained using the mean squared error loss with the Adam optimizer, and the prediction performance was evaluated using tracking error against both the HHFI benchmark and individual hedge fund strategies.

Generative Adversarial Networks

GANs were developed using a generator-discriminator architecture trained through adversarial minimax optimization. The generator selects factor combinations from the expanded factor universe to produce synthetic hedge fund returns, whereas the discriminator attempts to distinguish generated returns from actual hedge fund returns. Both networks employed Adam optimization, binary cross-entropy loss, dropout regularization, batch normalization, and early stopping to improve convergence and reduce overfitting.

The GAN models were evaluated using four different factor set sizes (GAN-6, GAN-15, GAN-30, and GAN-60). Additional experiments examined training periods of eleven, five, three, and one years to determine the effect of historical window length on replication performance.

Evaluation Metrics and Statistical Analysis

Model performance is evaluated using tracking error. Performance was compared against both the HHFI benchmark and individual hedge fund strategies.

To determine whether the observed improvements reflected genuine performance gains rather than sampling variation, paired t-tests and Wilcoxon signed-rank tests were conducted across the strategies. The statistical significance was evaluated for comparisons between the GAN models and the best-performing regression, RNN, and LSTM models.

Results

Baseline Regression Performance

Table 2 shows a comparison of the fixed-weight, 24-month rolling and 36-month rolling returns for each strategy. In addition, all values in Table 2 are intended to be evaluated relative to the HHFI return of 12.22%.

Typically, individual factor weights are estimated by minimizing the tracking error or the sum of squared deviations (OLS) between the target return and the replication portfolio return. Table 3 presents the tracking error results. These baseline regression results are comparable to those of previous research, particularly Hasanhodzic and Lo (2006), Amenc et al. (2003), Amenc et al. (2007), and Tancar et al. (2012).

Strategy Name	Fixed Regression Replicated Portfolio	24-month Rolling Regression Replicated Portfolio	36-month Rolling Regression Replicated Portfolio	Strategy Return
Asia Pacific	5.99%	6.02%	5.93%	9.84%
Convertible Arbitrage	4.70%	3.24%	3.92%	6.50%
Distressed	9.10%	7.21%	7.76%	15.17%
Emerging Markets	7.58%	7.12%	6.92%	7.13%
Europe	6.35%	7.88%	7.26%	17.64%
Event Driven	9.93%	10.71%	10.64%	14.08%
Financial Equities	8.27%	8.75%	7.97%	21.21%
Fixed Income	3.01%	2.56%	2.77%	4.19%
Growth	10.48%	11.66%	11.23%	17.01%
Healthcare and Biotech	11.12%	11.08%	12.24%	30.25%
High Yield	5.19%	4.26%	4.83%	7.63%
International	7.81%	8.99%	8.63%	8.54%
Latin America	10.20%	9.83%	9.34%	-2.39%
Macro	-0.18%	1.94%	1.54%	-0.69%
Market Neutral	2.75%	2.42%	2.53%	7.21%
Merger Arbitrage	2.99%	4.33%	4.14%	7.49%
Multiple Arbitrage	3.73%	3.48%	3.58%	10.72%
Opportunistic	6.52%	8.56%	7.64%	14.99%
Short Biased	-18.27%	-18.84%	-19.42%	-32.68%
Technology	11.26%	12.95%	12.45%	16.41%
Value	11.30%	11.58%	11.44%	19.42%

Table 2: Replicated Portfolio Returns

Strategy Name	Tracking Error - Index			Tracking Error - Strategy		
	Fixed	24 Month	36 Month	Fixed	24 Month	36 Month
Asia Pacific	1.28%	1.36%	1.41%	1.85%	1.90%	1.90%
Convertible Arbitrage	4.44%	1.15%	1.15%	6.57%	1.04%	1.12%
Distressed	6.99%	0.71%	0.69%	6.99%	1.06%	1.04%
Emerging Markets	5.75%	1.29%	1.37%	6.35%	1.38%	1.45%
Europe	4.34%	0.94%	0.97%	5.89%	1.55%	1.60%
Event Driven	4.23%	0.58%	0.61%	5.74%	0.88%	0.90%
Financial Equities	4.40%	0.62%	0.73%	7.59%	1.67%	1.79%
Fixed Income	6.50%	1.20%	1.18%	4.78%	0.70%	0.73%
Growth	4.71%	0.53%	0.56%	4.77%	0.79%	0.88%
Healthcare and Biotech	4.50%	0.67%	0.60%	7.45%	2.02%	2.02%
High Yield	4.39%	0.98%	0.95%	5.67%	0.61%	0.61%
International	4.65%	1.07%	1.14%	5.39%	1.44%	1.49%
Latin America	7.23%	1.60%	1.61%	12.39%	1.78%	1.77%
Macro	6.57%	1.20%	1.25%	5.88%	0.72%	0.75%
Market Neutral	6.77%	1.21%	1.24%	4.36%	0.71%	0.73%
Merger Arbitrage	6.15%	0.92%	0.96%	3.66%	0.44%	0.46%
Multiple Arbitrage	5.07%	1.06%	1.07%	5.04%	0.89%	0.91%
Opportunistic	4.15%	0.75%	0.85%	5.29%	1.42%	1.50%
Short Biased	21.40%	3.61%	3.64%	6.20%	1.48%	1.46%
Technology	5.21%	0.63%	0.70%	5.13%	1.09%	1.14%
Value	4.77%	0.58%	0.60%	4.20%	1.12%	1.17%

Table 3: Regression Tracking Error

Deep Neural Network Performance

Five feed-forward neural network architectures were evaluated against regression benchmarks. Tracking error comparisons against the HHFI benchmark (Table 4) showed that 13 of the twenty-one strategies (62%) achieved lower tracking errors than the regression models. When evaluated against individual hedge fund strategies (Table 5), eight strategies (38%) outperformed the regression benchmarks, indicating that additional temporal modeling may improve replication performance.

Strategy Name	Fixed	24 Month	36 Month	DNN1	DNN2	DNN3	DNN4	DNN5
Asia Pacific	1.28%	1.36%	1.41%	0.99%	1.10%	1.34%	1.27%	1.27%
Convertible Arbitrage	4.44%	1.15%	1.15%	1.04%	1.09%	1.18%	1.19%	1.18%
Distressed	6.99%	0.71%	0.69%	1.00%	1.00%	1.06%	1.07%	1.09%
Emerging Markets	5.75%	1.29%	1.37%	1.24%	1.04%	1.14%	1.14%	1.12%
Europe	4.34%	0.94%	0.97%	0.96%	0.98%	1.16%	1.19%	1.18%
Event Driven	4.23%	0.58%	0.61%	0.83%	0.91%	1.09%	1.08%	1.07%
Financial Equities	4.40%	0.62%	0.73%	0.71%	1.10%	1.13%	1.11%	1.10%
Fixed Income	6.50%	1.20%	1.18%	1.16%	1.10%	1.17%	1.16%	1.19%
Growth	4.71%	0.53%	0.56%	0.66%	1.12%	1.13%	1.17%	1.13%
Healthcare and Biotech	4.50%	0.67%	0.60%	0.84%	0.91%	1.08%	1.08%	1.08%
High Yield	4.39%	0.98%	0.95%	0.91%	0.97%	1.10%	1.12%	1.09%
International	4.65%	1.07%	1.14%	1.04%	0.96%	1.10%	1.10%	1.09%
Latin America	7.23%	1.60%	1.61%	1.71%	1.78%	1.37%	1.06%	1.06%
Macro	6.57%	1.20%	1.25%	1.24%	1.20%	1.39%	1.17%	1.16%
Market Neutral	6.77%	1.21%	1.24%	0.99%	1.10%	1.34%	1.27%	1.27%

Merger Arbitrage	6.15%	0.92%	0.96%	0.83%	1.05%	1.09%	1.15%	1.15%
Multiple Arbitrage	5.07%	1.06%	1.07%	0.78%	1.04%	1.09%	1.14%	1.14%
Opportunistic	4.15%	0.75%	0.85%	0.63%	1.05%	1.06%	1.13%	1.12%
Short Biased	21.40%	3.61%	3.64%	3.28%	1.95%	1.71%	1.71%	1.68%
Technology	5.21%	0.63%	0.70%	0.85%	1.10%	1.12%	1.13%	1.12%
Value	4.77%	0.58%	0.60%	0.81%	1.09%	1.11%	1.12%	1.12%

Table 4: DNN Tracking Error Against Index

Strategy Name	**Fixed**	**24 Month**	**36 Month**	**DNN1**	**DNN2**	**DNN3**	**DNN4**	**DNN5**
Asia Pacific	1.85%	1.90%	1.90%	2.10%	2.10%	2.10%	2.10%	2.10%
Convertible Arbitrage	6.57%	1.04%	1.12%	0.81%	0.61%	0.62%	0.60%	0.59%
Distressed	6.99%	1.06%	1.04%	1.14%	1.18%	1.31%	1.31%	1.33%
Emerging Markets	6.35%	1.38%	1.45%	1.56%	1.51%	1.69%	1.70%	1.68%
Europe	5.89%	1.55%	1.60%	1.43%	1.63%	1.75%	1.77%	1.76%
Event Driven	5.74%	0.88%	0.90%	0.93%	0.97%	1.18%	1.15%	1.15%
Financial Equities	7.59%	1.67%	1.79%	1.37%	1.85%	1.99%	2.02%	2.01%
Fixed Income	4.78%	0.70%	0.73%	1.00%	0.75%	0.72%	0.71%	0.74%
Growth	4.77%	0.79%	0.88%	0.80%	1.50%	1.68%	1.72%	1.70%
Healthcare and Biotech	7.45%	2.02%	2.02%	2.05%	2.12%	2.58%	2.57%	2.59%
High Yield	5.67%	0.61%	0.61%	0.54%	0.70%	0.65%	0.67%	0.65%
International	5.39%	1.44%	1.49%	1.63%	1.64%	1.43%	1.42%	1.43%
Latin America	12.39%	1.78%	1.77%	1.92%	2.03%	2.38%	2.54%	2.54%
Macro	5.88%	0.72%	0.75%	0.91%	0.80%	0.85%	1.15%	1.14%
Market Neutral	4.36%	0.71%	0.73%	0.57%	0.61%	0.77%	0.69%	0.69%
Merger Arbitrage	3.66%	0.44%	0.46%	0.49%	0.41%	0.47%	0.52%	0.52%
Multiple Arbitrage	5.04%	0.89%	0.91%	0.53%	0.73%	0.75%	0.80%	0.80%
Opportunistic	5.29%	1.42%	1.50%	1.61%	1.69%	1.89%	1.97%	1.97%
Short Biased	6.20%	1.48%	1.46%	2.22%	2.91%	3.01%	3.01%	3.04%
Technology	5.13%	1.09%	1.14%	1.51%	1.63%	1.66%	1.65%	1.62%
Value	4.20%	1.12%	1.17%	1.18%	1.65%	1.72%	1.72%	1.73%

Table 5: DNN Tracking Error Against Strategies

Recurrent Neural Network and LSTM Performance

The RNN and LSTM architectures were evaluated to determine whether sequential learning improved replication performance. Against the HHFI benchmark (Table 6), the regression models and LSTM each produced the best results for seven strategies. However, a comparison against individual hedge fund strategies (Table 7) demonstrated that the RNN consistently achieved the lowest tracking error, outperforming all competing models in 11 of the twenty-one strategies. These findings suggest that incorporating temporal dependencies substantially improves the replication accuracy.

Strategy Name	**Fixed**	**24 Month**	**36 Month**	**DNN1**	**DNN2**	**DNN3**	**DNN4**	**DNN5**	**RNN**	**LSTM**
Asia Pacific	1.28%	1.36%	1.41%	0.99%	1.10%	1.34%	1.27%	1.27%	1.76%	0.90%
Convertible Arbitrage	4.44%	1.15%	1.15%	1.04%	1.09%	1.18%	1.19%	1.18%	1.00%	0.63%
Distressed	6.99%	0.71%	0.69%	1.00%	1.00%	1.06%	1.07%	1.09%	0.75%	1.04%
Emerging Markets	5.75%	1.29%	1.37%	1.24%	1.04%	1.14%	1.14%	1.12%	0.85%	0.97%
Europe	4.34%	0.94%	0.97%	0.96%	0.98%	1.16%	1.19%	1.18%	0.80%	0.80%
Event Driven	4.23%	0.58%	0.61%	0.83%	0.91%	1.09%	1.08%	1.07%	1.14%	0.76%
Financial Equities	4.40%	0.62%	0.73%	0.71%	1.10%	1.13%	1.11%	1.10%	0.91%	0.85%
Fixed Income	6.50%	1.20%	1.18%	1.16%	1.10%	1.17%	1.16%	1.19%	0.51%	0.57%

Growth	4.71%	0.53%	0.56%	0.66%	1.12%	1.13%	1.17%	1.13%	0.92%	0.62%
Healthcare and Biotech	4.50%	0.67%	0.60%	0.84%	0.91%	1.08%	1.08%	1.08%	1.05%	1.00%
High Yield	4.39%	0.98%	0.95%	0.91%	0.97%	1.10%	1.12%	1.09%	0.78%	0.75%
International	4.65%	1.07%	1.14%	1.04%	0.96%	1.10%	1.10%	1.09%	0.97%	0.97%
Latin America	7.23%	1.60%	1.61%	1.71%	1.78%	1.37%	1.06%	1.06%	1.41%	1.35%
Macro	6.57%	1.20%	1.25%	1.24%	1.20%	1.39%	1.17%	1.16%	0.78%	0.73%
Market Neutral	6.77%	1.21%	1.24%	0.99%	1.10%	1.34%	1.27%	1.27%	1.00%	0.78%
Merger Arbitrage	6.15%	0.92%	0.96%	0.83%	1.05%	1.09%	1.15%	1.15%	1.06%	0.81%
Multiple Arbitrage	5.07%	1.06%	1.07%	0.78%	1.04%	1.09%	1.14%	1.14%	0.62%	0.51%
Opportunistic	4.15%	0.75%	0.85%	0.63%	1.05%	1.06%	1.13%	1.12%	1.02%	0.82%
Short Biased	21.40%	3.61%	3.64%	3.28%	1.95%	1.71%	1.71%	1.68%	3.39%	3.13%
Technology	5.21%	0.63%	0.70%	0.85%	1.10%	1.12%	1.13%	1.12%	1.01%	0.81%
Value	4.77%	0.58%	0.60%	0.81%	1.09%	1.11%	1.12%	1.12%	0.97%	0.79%

Table 6: RNN/LSTM Tracking Error Against Index

Strategy Name	**Fixed**	**24 Month**	**36 Month**	**DNN1**	**DNN2**	**DNN3**	**DNN4**	**DNN5**	**RNN**	**LSTM**
Asia Pacific	1.85%	1.90%	1.90%	2.10%	2.10%	2.10%	2.10%	2.10%	1.36%	1.90%
Convertible Arbitrage	6.57%	1.04%	1.12%	0.81%	0.61%	0.62%	0.60%	0.59%	1.00%	1.02%
Distressed	6.99%	1.06%	1.04%	1.14%	1.18%	1.31%	1.31%	1.33%	0.64%	0.92%
Emerging Markets	6.35%	1.38%	1.45%	1.56%	1.51%	1.69%	1.70%	1.68%	0.96%	1.40%
Europe	5.89%	1.55%	1.60%	1.43%	1.63%	1.75%	1.77%	1.76%	0.52%	0.78%
Event Driven	5.74%	0.88%	0.90%	0.93%	0.97%	1.18%	1.15%	1.15%	1.18%	0.78%
Financial Equities	7.59%	1.67%	1.79%	1.37%	1.85%	1.99%	2.02%	2.01%	0.78%	1.11%
Fixed Income	4.78%	0.70%	0.73%	1.00%	0.75%	0.72%	0.71%	0.74%	0.69%	0.84%
Growth	4.77%	0.79%	0.88%	0.80%	1.50%	1.68%	1.72%	1.70%	1.14%	0.87%
Healthcare and Biotech	7.45%	2.02%	2.02%	2.05%	2.12%	2.58%	2.57%	2.59%	1.25%	1.26%
High Yield	5.67%	0.61%	0.61%	0.54%	0.70%	0.65%	0.67%	0.65%	0.37%	0.63%
International	5.39%	1.44%	1.49%	1.63%	1.64%	1.43%	1.42%	1.43%	1.43%	1.56%
Latin America	12.39%	1.78%	1.77%	1.92%	2.03%	2.38%	2.54%	2.54%	2.04%	2.31%
Macro	5.88%	0.72%	0.75%	0.91%	0.80%	0.85%	1.15%	1.14%	0.99%	1.32%
Market Neutral	4.36%	0.71%	0.73%	0.57%	0.61%	0.77%	0.69%	0.69%	0.44%	0.47%
Merger Arbitrage	3.66%	0.44%	0.46%	0.49%	0.41%	0.47%	0.52%	0.52%	0.42%	0.28%
Multiple Arbitrage	5.04%	0.89%	0.91%	0.53%	0.73%	0.75%	0.80%	0.80%	0.51%	0.57%
Opportunistic	5.29%	1.42%	1.50%	1.61%	1.69%	1.89%	1.97%	1.97%	1.74%	1.34%
Short Biased	6.20%	1.48%	1.46%	2.22%	2.91%	3.01%	3.01%	3.04%	1.41%	2.08%
Technology	5.13%	1.09%	1.14%	1.51%	1.63%	1.66%	1.65%	1.62%	1.73%	1.18%
Value	4.20%	1.12%	1.17%	1.18%	1.65%	1.72%	1.72%	1.73%	0.99%	0.92%

Table 7: RNN/LSTM Tracking Error Against Strategies

GAN Performance

GAN models incorporating expanded factor sets were evaluated using 6-, 15-, 30-, and 60-factor configurations. The results are summarized in Table 8.

GAN-6 and GAN-15 consistently produce the lowest tracking errors across most hedge fund strategies. Combined, these two models outperformed approximately 86% of all the competing regression and neural network models. In contrast, GAN-60 generally produced inferior performance, suggesting that larger-factor universes introduce unnecessary complexity rather than improving predictive accuracy.

Strategy Name	**RNN**	**LSTM**	**GAN-6**	**GAN-15**	**GAN-30**	**GAN-60**
Asia Pacific	1.36%	1.90%	1.48%	**1.17%**	1.48%	2.56%
Convertible Arbitrage	1.00%	1.02%	0.25%	**0.10%**	0.21%	1.72%
Distressed	0.64%	0.92%	**0.24%**	0.35%	0.25%	1.64%

Emerging Markets	0.96%	1.40%	0.52%	**0.38%**	0.48%	1.46%
Europe	0.52%	0.78%	**0.27%**	0.71%	0.69%	1.59%
Event Driven	1.18%	0.78%	0.21%	0.40%	**0.17%**	2.72%
Financial Equities	0.78%	1.11%	0.76%	**0.34%**	0.34%	1.24%
Fixed Income	0.69%	0.84%	**0.19%**	0.40%	0.23%	1.21%
Growth	1.14%	0.87%	**0.45%**	0.61%	0.84%	2.39%
Healthcare and Biotech	1.25%	1.26%	1.41%	**1.17%**	1.18%	2.78%
High Yield	0.37%	0.63%	**0.08%**	0.88%	0.98%	1.67%
International	1.43%	1.56%	**0.41%**	0.46%	0.70%	2.65%
Latin America	2.04%	2.31%	0.89%	**0.70%**	0.99%	3.29%
Macro	0.99%	1.32%	**0.69%**	0.86%	0.72%	1.28%
Market Neutral	0.44%	0.47%	**0.12%**	0.73%	0.58%	1.98%
Merger Arbitrage	0.42%	0.28%	0.14%	0.16%	**0.07%**	1.24%
Multiple Arbitrage	0.51%	0.57%	**0.06%**	0.49%	0.24%	1.68%
Opportunistic	1.74%	1.34%	0.92%	0.94%	**0.42%**	2.69%
Short Biased	1.41%	2.08%	**1.37%**	2.23%	2.31%	3.14%
Technology	1.73%	1.18%	0.83%	**0.72%**	0.78%	2.74%
Value	0.99%	0.92%	**0.36%**	0.99%	1.08%	2.47%

Table 8: GAN Tracking Errors

Significant improvements were observed in the tracking error; for example, the GAN-15 model reduced the tracking error by 90% in Convertible Arbitrage compared to the RNN model, while the GAN-6 model yielded up to 88% lower errors in other strategies. The performance improved even further when compared with regression models, indicating the GAN model's robust efficacy, especially when leveraging both the minimax game and an expanded factor set.

Further evaluation focused on four strategies: Asia Pacific, Europe, Fixed Income, and Technology. The GAN model identifies specific factors for each, with diverse assets that enhance the predictive capability. For example, Asia Pacific's portfolio includes factors such as EWT (Taiwan), EWS (Singapore), and EWJ (Japan), while Fixed Income favors bonds and mid-cap equities. Table 9 shows the factors selected by the GAN model that resulted in the best performing forecasted return results.

Asia Pacific		**Europe**	**Fixed Income**	**Technology**	
EWC	IEV	EWA	Bonds	EWJ	IYW
EWI	IXN	EWG	IEF	EWT	IYZ
EWJ	IYR	EWI	IJH	EWW	LQD
EWL	SHY	EWU	IXJ	EWY	NASDAQ
EWS	USD	IXC	NASDAQ	IJH	VTI
EWT	VTI	XLF	SHY	IJR	XLK
EWW	XLB			IXN	XLV
EWZ				IYE	

Table 9: Final Factor Set of Four Strategies

Reduced Historical Window Analysis

To examine the effect of the training window length, GAN models were re-estimated using five-, three-, and one-year historical datasets for the four representative hedge fund strategies. The results presented in Table 10 demonstrate that shorter historical windows further reduced the tracking error, with one-year training windows producing improvements of up to 96% relative to the best-performing regression models. This finding aligns with those of previous studies, Andres (2018), showing that models using shorter data periods can improve prediction accuracy in dynamic markets [21].

Strategy Name	**11 Years**		**5 Years**		**3 Years**		**1 Year**	
	GAN-6	**GAN-15**	**GAN-6**	**GAN-15**	**GAN-6**	**GAN-15**	**GAN-6**	**GAN-15**
Asia Pacific	1.48%	**1.17%**		**0.084%**		0.644%		0.384%
Europe	**0.27%**	0.71%	0.073%		0.090%		**0.061%**	
Fixed Income	**0.19%**	0.40%	0.182%		0.207%		**0.044%**	
Technology	0.83%	**0.72%**		0.381%		0.423%		**0.054%**

Table 10: 5-3-1 Year Tracking Error Results

Statistical Performance Comparisons

To formally evaluate whether the reductions in tracking error observed for the GAN models reflect genuine performance improvements rather than sampling variation, paired statistical comparison tests were conducted using strategy-level tracking errors from each model. Because each strategy is replicated by multiple models over the same horizon, comparisons are naturally paired. I first computed paired t-tests on the differences in tracking errors between each benchmark model and GAN variants. To relax distributional assumptions and account for the skewness and heavy-tailed behavior typical of hedge-fund return processes, I also implement Wilcoxon signed-rank tests, which assess whether GAN models systematically produce lower tracking errors than their counterparts do. The paired tests across strategies indicate that GAN-based models, particularly GAN-6 and GAN-15, generate statistically significant reductions in tracking error relative to both the RNN and LSTM benchmarks. Wilcoxon signed-rank $p < 0.01$ in all cases, $p < 0.001$ for most comparisons. The magnitude of the improvement is economically meaningful, averaging 30–60 basis points of tracking-error reduction across strategies. Detailed test statistics and p-values for all model comparisons are presented in Table 11.

Comparison	Mean TE difference (positive = GAN better)	Paired-t p-value	Wilcoxon p-value	Interpretation
RNN vs GAN-6	**0.47%**	0.000011	0.0000091	GAN-6 significantly reduces TE
LSTM vs GAN-6	**0.57%**	0.00000014	0.0000014	Very strong evidence GAN-6 is better
RNN vs GAN-15	**0.32%**	0.012	0.00079	GAN-15 significantly better
LSTM vs GAN-15	**0.41%**	0.002	0.00030	GAN-15 significantly better
RNN vs GAN-30	**0.33%**	0.0135	0.0063	Statistically better, but weaker
LSTM vs GAN-30	**0.42%**	0.00036	0.00043	Clearly better

Table 11: Statistical Significance of Tracking Errors

Across specifications, the results show statistically significant improvements, particularly for GAN-6 and GAN-15, indicating that performance gains are both statistically and economically meaningful.

Discussion

Deep Neural Networks

Interest in machine learning, particularly Artificial Neural Networks (ANN), for finance has grown significantly in both academia and industry. However, their complexity makes them difficult to interpret and highly parameterized, often referred to as "black boxes."

Neural networks provide a flexible data-driven framework for modeling nonlinear relationships between hedge fund returns and risk factors. In contrast to linear specifications, neural networks allow mapping from factors to payoffs to vary across states of the world and to capture interaction effects that would otherwise be omitted.

The universal approximation theorem (Gelenbe, et al., 1999) demonstrates that multi-layer perceptrons with hidden layers can approximate any continuous function. In this study, neural networks were used to assess whether more expressive architectures can improve hedge fund replication accuracy relative to traditional regression models.

To maintain clarity in the main text, the technical and theoretical background of the network architectures, activation functions, optimization, and learning mechanics are summarized in Appendix A.

Recurrent Neural Network and LSTM

Recurrent Neural Networks are well suited to financial time series because they incorporate information from previous observations when generating forecasts. Unlike feed-forward networks, RNNs maintain a hidden state that evolves over time, allowing the model to condition each prediction on recent market history. This characteristic is particularly relevant for hedge fund strategies that often exhibit persistence and regime dynamics. However, RNNs can be challenging to train because of their depth, which can lead to the vanishing gradient problem, where the gradient of the loss function becomes too small to update the network effectively, limiting their ability to look far back in sequences [43].

To address the vanishing gradient problem, I evaluated LSTM architectures, which introduce gating mechanisms that help retain relevant information across longer horizons. Although these models are theoretically capable of capturing more complex temporal dependencies, they require greater data depth and parameter regularization.

Each recurrent model uses the same factor inputs, and a rolling evaluation design is applied to the feed-forward networks. Performance is again summarized using the tracking error against individual hedge fund strategies and benchmark index. The implementation details for both architectures, including notation, conceptual intuition, and training mechanics, are provided in Appendix A.

Generative Adversarial Networks

GANs, introduced by Goodfellow et al. (2014), initially focused on image generation but are now applied across various fields previously dominated by neural networks. GANs are generative models that learn to approximate a data distribution through unsupervised learning, thus allowing the creation of new plausible samples from the original dataset.

In GANs, two models are trained together: a generator (G) that attempts to mimic the real data distribution and a discriminator (D) that classifies inputs as real or fake. These models compete in a zero-sum game, where the generator seeks to produce realistic data to "fool" the discriminator, which, in turn, refines its ability to distinguish generated data from real samples.

The architecture of the generator can vary (e.g., multi-layer perceptrons, CNNs) as long as its input-output dimensions match the latent space and real data. Through a minimax game with the value function V(D,G), the generator maximizes the likelihood of its outputs being classified as real, whereas the discriminator aims to minimize it. This setup allows both models to improve via gradient-based optimization, with the generator producing data increasingly similar to the real data over time. Fig. 1 shows a representation of the created GAN model.

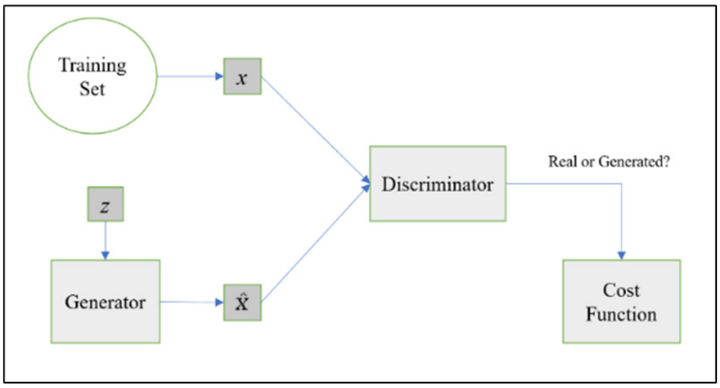

Figure 1: Generative Adversarial Model

A specific case, known as adversarial nets, involves both models using multilayer perceptrons and can be trained using only backpropagation and dropout algorithms. No approximate inference or Markov chains are necessary. Here, the training set will consist of actual hedge fund strategy returns (x). These will be known to the discriminator as 'real'. The generator selects a set of six factors from an expanded factor data set (z) and creates a hedge fund strategy return ($\hat{X}$) which gets presented to the discriminator. If the returns ($\hat{X}$) are not similar enough to the actual returns (x) then the discriminator will label them as "generated" or "fake". However, the generator takes this rejection from the discriminator and makes another attempt at creating a return ($\hat{X}$) with a new set of six factors from z and presents this to the discriminator. The generator iteratively adjusts to produce returns closer to real data until the discriminator identifies them as authentic.

GAN Training Procedure and Hyperparameters

To train the GAN, I adopt a minimax optimization framework consistent with Goodfellow et al. (2014). The discriminator D and generator G are optimized via alternating gradient descent applied to the standard adversarial loss:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D\big(G(z)\big)\right)\right] \quad (2)$$

where x represents realized hedge-fund strategy returns and z represents candidate factor return sets sampled from the expanded factor universe.

Formally, the generator loss is:

$$L_G = -\mathbb{E}_{z\sim p_z(z)}\left[\log\left(D\big(G(z)\big)\right)\right] \quad (3)$$

and the discriminator loss is:

$$L_D = -\mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] - \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D\big(G(z)\big)\right)\right] \quad (4)$$

These modifications help stabilize training and align the GAN with the heavy-tail characteristics of asset return distributions.

Both G and D use the Adam optimizer, which is standard for adversarial architectures and empirically effective for financial time-series learning. The GAN was trained using a binary cross-entropy loss, Adam optimization (learning rate 0.0002, β_1=0.5, β_2=0.999), and dropout regularization to ensure convergence under heavy-tailed financial return distributions. The optimization parameters used, their values, and rationale are listed in Table A1 in the Appendix.

To ensure convergence and avoid mode collapse, the following techniques are implemented: dropout layers (0.2 - 0.3) in both G and D, batch normalization inside G layers, and early stopping after 250 epochs of no discriminator loss improvement. These techniques are motivated by empirical evidence in financial GAN literature, especially for noisy, autocorrelated time series.

These architectural and optimization decisions reflect the economic nature of hedge-fund payoffs: dynamic, regime-sensitive factor exposures, volatility clustering and nonlinear return surfaces, and Non-Gaussian tail behavior in alternative investment strategies. The GAN framework allows the model to discover sparse and nonlinear factor portfolios that resemble hedge-fund risk premia more closely than linear or feed-forward models.

Factor-Set Expansion and Empirical GAN Performance

Expanding the factor set to 60 (from an initial 6) allowed for a diverse asset representation and enhanced GAN modeling for multiple strategies. This expanded dataset consists of liquid, transparent, and scalable financial instruments from sources such as iShares ETFs and SPDR funds. Regarding liquidity, the factors were selected based on an inception date no later than 12/31/2002, returns were available for the entire 11-year period, and the average amount of transactions per day was greater than $50 million.

Economic Intuition Behind GAN Performance

GANs outperform others because they implicitly learn state-dependent nonlinear payoff mappings between factors and hedge-fund returns. The linear factor models assume constant betas and additive shocks. However, hedge funds change exposures across regimes (risk-on vs. risk-off), embed optionality (stop-loss rules, leverage constraints), dynamically rebalance in response to volatility and liquidity conditions, and concentrate risk in tails rather than in means.

Economically, these behaviors generate asymmetric and path-dependent payoff functions that resemble synthetic options, not linear portfolios. A GAN learns the distribution of these realized payoffs, not just the conditional mean, by forcing the generator to produce factor combinations whose payoff characteristics cannot be statistically distinguished from true hedge-fund outcomes. This adversarial setup allows the GAN to discover nonlinear factor loadings, latent interactions among factors, regime-specific exposures, and risk-tail structure embedded in manager behavior.

Thus, the GAN does not just fit returns better; it approximates the economic decision rules managers use to dynamically trade risk. This is why the tracking error and replication are improved.

Conclusion

Previous research has utilized fixed, 24-month, and 36-month rolling regressions. Replicating these regressions across each hedge fund strategy, benchmark HHFI, and 21 strategies confirms Hasanhodzic and Lo (2006): regression-based hedge fund replication returns are generally inferior to the actual returns.

This study evaluates the application of deep learning techniques to hedge fund replication using regression models, DNNs, RNNs, LSTMs, and GANs. Consistent with previous literature, regression models established effective baseline replication, but were limited in capturing the nonlinear and dynamic behavior characteristics of hedge fund returns.

Among the conventional neural network architectures, RNNs produced the strongest overall performance when replicating individual hedge fund strategies. Incorporating recurrent learning within the GAN framework further improved the replication accuracy, with GAN-6 and GAN-15 consistently producing the lowest tracking errors across multiple hedge fund strategies. Additional analysis demonstrates that shorter historical training windows further enhance forecasting performance, suggesting that more recent market information provides greater predictive value for hedge fund replication.

This study shows that if using this approach to hedge fund replication, an investor could generate hedge fund-like returns that are no more than a few basis points off the actual returns, and in some cases, may obtain even higher returns.

Overall, the findings demonstrate that adversarial learning provides a more effective framework for modeling nonlinear regime-dependent hedge fund behavior than traditional linear factor models. These results suggest that strategy-specific replication, instead of an index, using recurrent GAN architectures offers a scalable and transparent alternative for constructing hedge fund replication portfolios, while substantially reducing tracking errors. Future research may extend this framework through additional factor universes, alternative deep learning architectures, and Explainable AI techniques to further improve the interpretability and replication performance.

Appendix

Appendix A. Neural Network Architecture and Details

Appendix A provides a background on neural network architectures and the terminology referenced in Section 4 to support reproducibility and clarity.

A.1 Overview of Feed Forward Neural Networks

Feed-forward neural networks (multi-layer perceptrons) consist of an input layer, one or more hidden layers, and an output layer. Each hidden layer applies a weighted linear transformation, followed by a nonlinear activation. An example of a three-layer neural network is shown in Fig. A1, where six inputs (x_i) represent monthly factor returns. Each input is fully connected to the nodes in the hidden layer, which in turn connects to the output layer that forecasts returns at time t+1. Fig. A2 shows a model with three hidden layers, each of which can vary in the number of nodes. The networks used in this study employ Rectified Linear Unit (ReLU) activations, which help stabilize the training and mitigate vanishing gradient effects in deeper models.

Activation functions, such as ReLU, are nonlinear transformations, which previous research has found to rapidly reduce dimensionality, improving the model's performance. The output layer linearly transforms the output of the hidden layer, and without this nonlinearity, the network is reduced to a generalized linear model. The ReLU function is defined as follows:

$\mathrm{ReLU}(x_j) = \max(x_j, 0)$.

Deep learning predictors overcome the "curse of dimensionality" using univariate activation functions, such as ReLU, and the architecture depth is determined by the number of layers, denoted as l. These models use six inputs, representing monthly factor returns, and employ ReLU activation, the Adam optimizer, and metrics such as the Mean Squared Error (MSE) and Mean Absolute Error (MAE) for evaluation. The output layer generates the average of 50 runs to forecast returns at time t+1.

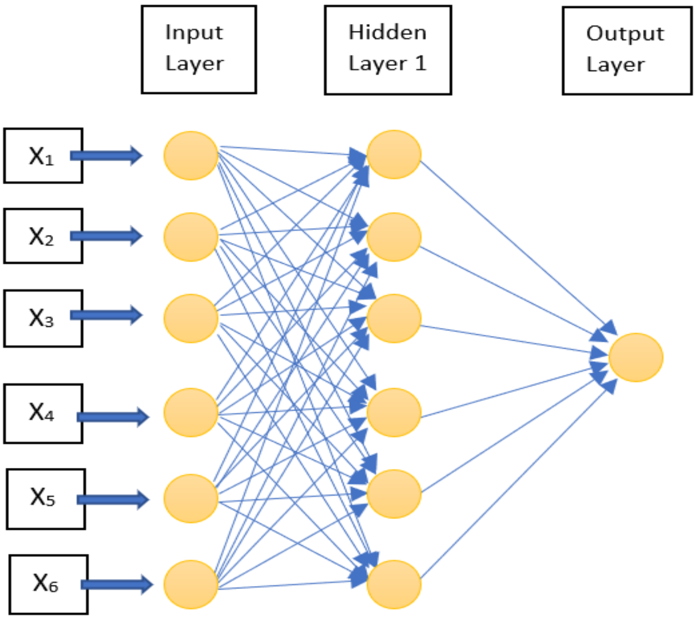

Figure 1: One-Layer Neural Network Model

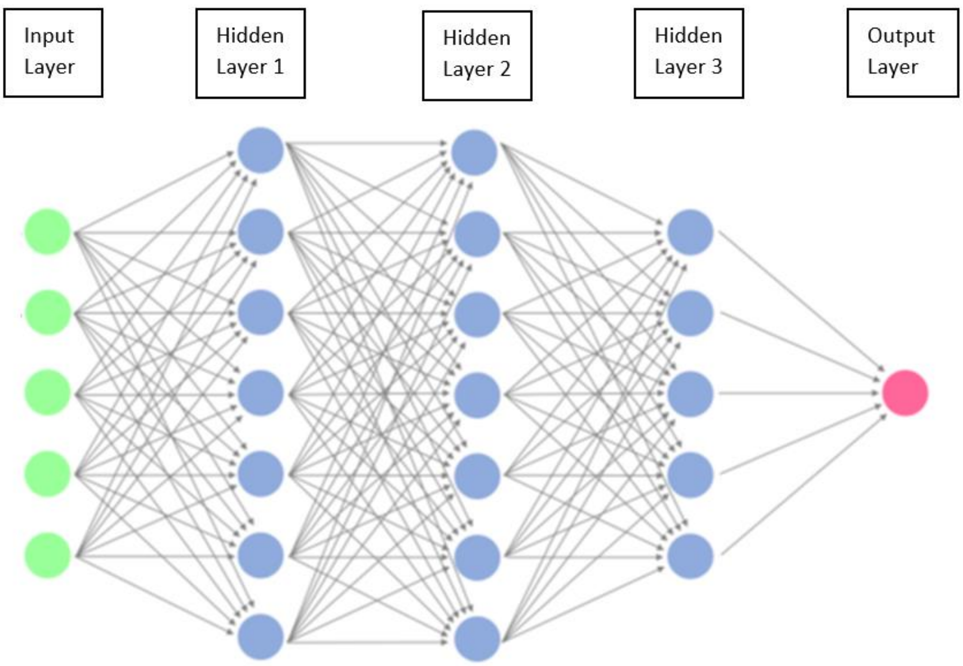

Figure 1: Multi-Layer Neural Network Model

Recurrent Neural Networks

Recurrent Neural Networks extend feed-forward architectures by introducing a hidden state that is updated at each time step. This hidden state allows the model to incorporate information from previous returns and factors when producing new forecasts. This makes them ideal for tasks in which inputs and outputs depend on past data, which is a common scenario in financial data. In a vanilla RNN, the current state is a nonlinear transformation of the current input (x_t) and the previous hidden state (h_{t-1}). Formally, the hidden state evolves through a nonlinear transformation of the current input and the prior state. This formulation is analogous to a generalized autoregressive process but with greater modeling flexibility, with the ReLU activation function used for nonlinearity. RNNs can capture short-run temporal dependence and regime transitions, which linear models cannot.

Long Short Term Memory Networks

LSTM networks augment the RNN framework with input, output, and forget gates that regulate the information flow. To create an LSTM model, the current input ($x_t = I_t$) and hidden state (h_t) are used to generate a memory cell (c_t) controlled by the input and forget gates. The output gate controls the amount of information stored in the hidden state utilizing the sigmoid function. These gates help to address the vanishing gradient problem and allow relevant features to persist across longer horizons. LSTMs are powerful when long memory dynamics drive returns, but they may be more data-intensive and sensitive to hyperparameter selection than simpler RNNs.

Training, Hyperparameters, and Evaluation

All neural network models were trained using the Adam optimizer with mean squared error loss. Hyperparameters, such as learning rate, batch size, and number of epochs, were tuned through cross-validated experiments to balance bias and variance. To ensure comparability, the networks were trained on rolling estimation windows that mirror the regression benchmarks. Out-of-sample performance is evaluated using tracking error and mean absolute error, with model selection based on minimizing.

Interpretation and Robustness Notes

Neural networks can capture nonlinear interactions and state-dependent exposures that resemble option like payoff structures commonly observed in hedge fund strategies. Robustness checks included sensitivity to window length, alternative evaluation metrics, and repeated runs to smooth stochastic training variations. The additional empirical diagnostics linked to these tests are discussed in Section 5.

Parameter	Value	Rationale
Learning rate (α)	0.0002	Conventional for GAN stability; prevents mode collapse
β_1	0.5	Improves stability by reducing momentum on gradients
β_2	0.999	Maintains adaptive learning rate smoothing
Batch size	64	Balance between gradient fidelity & computational efficiency
Epochs	5,000	Ensures convergence across factor sets

Table 1: Optimization Parameters

References

1. Fung, W., & Hsieh, D. A. (1997). Empirical characteristics of dynamic trading strategies: The case of hedge funds. The review of financial studies, 10(2), 275-302.
2. Fung, W., & Hsieh, D. A. (2002). Asset-based style factors for hedge funds. Financial Analysts Journal, 58(5), 16-27.
3. Fung, W., & Hsieh, D. A. (2004). Hedge fund benchmarks: A risk-based approach. Financial Analysts Journal, 60(5), 65-80.
4. Amenc, N., El Bied, S., & Martellini, L. (2003). Predictability in hedge fund returns (corrected). Financial Analysts Journal, 59(5), 32-46.
5. Amenc, Noel, Walter Gehin, Lionel Martellini, and Jean-Christophe Meyfredi. "The myths and limits of passive hedge fund replication." EDHEC Risk and Asset Management Research Centre, Working Paper (2007).
6. Agarwal, V., & Naik, N. Y. (2004). Risks and portfolio decisions involving hedge funds. The Review of Financial Studies, 17(1), 63-98.
7. Amenc, N., Gehin, W., Martellini, L., & Meyfredi, J. C. (2008). Passive hedge fund replication: A critical assessment of existing techniques. The Journal of Alternative Investments, 11(2), 69.
8. Amenc, N., Martellini, L., Meyfredi, J. C., & Ziemann, V. (2010). Passive hedge fund replication–Beyond the linear case. European Financial Management, 16(2), 191-210.
9. Bacmann, J. F., Held, R., Jeanneret, P., & Scholz, S. (2008). Beyond factor decomposition: Practical hurdles to hedge fund replication. The Journal of Alternative Investments, 11(2), 84.
10. Gupta, B., Szado, E., & Spurgin, W. (2008). Performance characteristics of hedge fund replication programs. The Journal of Alternative Investments, 11(2), 61.
11. Tancar, R., Poddig, T., & Ballis-Papanastasiou, P. (2012). Hedge Fund Replication: The Asymmetric Way. The Journal of Alternative Investments, 15(1), 68.
12. Heaton, J. B., Polson, N. G., & Witte, J. H. (2016). Deep learning in finance. arXiv preprint arXiv:1602.06561.

13. Hutchinson, J. M., Lo, A. W., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks. The journal of Finance, 49(3), 851-889.
14. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223-2273.
15. Li, Y., & Forsyth, P. A. (2019). A data-driven neural network approach to optimal asset allocation for target based defined contribution pension plans. Insurance: Mathematics and Economics, 86, 189-204.
16. Galindo, J., & Tamayo, P. (2000). Credit risk assessment using statistical and machine learning: basic methodology and risk modeling applications. Computational economics, 15(1), 107-143.
17. Tsai, C. F., & Wu, J. W. (2008). Using neural network ensembles for bankruptcy prediction and credit scoring. Expert systems with applications, 34(4), 2639-2649.
18. Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. Journal of Banking & Finance, 34(11), 2767-2787.
19. Aggarwal, S., & Aggarwal, S. (2017). Deep investment in financial markets using deep learning models. International Journal of Computer Applications, 162(2), 40-43.
20. Sadhwani, A., Giesecke, K., & Sirignano, J. (2021). Deep learning for mortgage risk. Journal of Financial Econometrics, 19(2), 313-368.
21. Lam, M. (2004). Neural network techniques for financial performance prediction: integrating fundamental and technical analysis. Decision support systems, 37(4), 567-581.
22. Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015, July). Deep learning for event-driven stock prediction. In Ijcai (Vol. 15, pp. 2327-2333).
23. Di Persio, L., & Honchar, O. (2016). Artificial neural networks approach to the forecast of stock market price movements. International Journal of Economics and Management Systems, 1(Anno 2016), 158-162.
24. Andres, J. R. (2018). Beating Momentum: Using Neural Networks to Uncover Additional Information in Past Returns. Available at SSRN 3113070.
25. 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.
26. Feng, G., He, J., & Polson, N. G. (2018). Deep learning for predicting asset returns. arXiv preprint arXiv:1804.09314.
27. Torres, D. G., & Qiu, H. (2018). Applying recurrent neural networks for multivariate time series forecasting of volatile financial data. KTH Royal Institute of Technology: Stockholm, Sweden.
28. Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. The Review of Financial Studies, 34(2), 1046-1089.
29. Nunes, M., Gerding, E., McGroarty, F., & Niranjan, M. (2018). Artificial neural networks in fixed income markets for yield curve forecasting. Available at SSRN 3144622.
30. https://cs229.stanford.edu/proj2018/report/76.pdf
31. Chen, A. S., Leung, M. T., & Daouk, H. (2003). Application of neural networks to an emerging financial market: forecasting and trading the Taiwan Stock Index. Computers & Operations Research, 30(6), 901-923.
32. Quek, C., Pasquier, M., & Kumar, N. (2008). A novel recurrent neural network-based prediction system for option trading and hedging. Applied Intelligence, 29(2), 138-151.
33. von Spreckelsen, C., von Mettenheim, H. J., & Breitner, M. H. (2014). Real-time pricing and hedging of options on currency futures with artificial neural networks. Journal of Forecasting, 33(6), 419-432.
34. Gerlein, E. A., McGinnity, M., Belatreche, A., & Coleman, S. (2016). Evaluating machine learning classification for financial trading: An empirical approach. Expert Systems with Applications, 54, 193-207.
35. Lu, D. W. (2017). Agent inspired trading using recurrent reinforcement learning and lstm neural networks. arXiv preprint arXiv:1707.07338.
36. Navon, A., & Keller, Y. (2017). Financial time series prediction using deep learning. arXiv preprint arXiv:1711.04174.
37. Ghoshal, S., & Roberts, S. (2017). Reading the tea leaves: a neural network perspective on technical trading. In KDD 2017 mining and learning from time series workshop.
38. Dixon, M. (2018). Sequence classification of the limit order book using recurrent neural networks. Journal of computational science, 24, 277-286.
39. Tsantekidis, A., Passalis, N., Tefas, A., Kanniainen, J., Gabbouj, M., & Iosifidis, A. (2017, August). Using deep learning to detect price change indications in financial markets. In 2017 25th European signal processing conference (EUSIPCO) (pp. 2511-2515). IEEE.
40. Lim, B., Zohren, S., & Roberts, S. (2019). Enhancing time series momentum strategies using deep neural networks. arXiv preprint arXiv:1904.04912.
41. Chalvatzis, C., & Hristu-Varsakelis, D. (2019). High-performance stock index trading: making effective use of a deep LSTM neural network. arXiv preprint arXiv:1902.03125.
42. Gelenbe, E., Mao, Z. H., & Li, Y. D. (1999). Function approximation with spiked random networks. IEEE Transactions on Neural Networks, 10(1), 3-9.
43. Hochreiter, S. (1998). The vanishing gradient problem during learning recurrent neural nets and problem solutions. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 6(02), 107-116.
44. Aggarwal, S., & Aggarwal, S. (2017). Deep investment in financial markets using deep learning models. International Journal of Computer Applications, 162(2), 40-43.
45. Goodfellow, I. (2016). Nips 2016 tutorial: Generative adversarial networks. arXiv preprint arXiv:1701.00160.
46. Hasanhodzic, J., & Lo, A. W. (2006). Can hedge-fund returns be replicated?: The linear case. The Linear Case

(August 16, 2006).
47. Koshiyama, A., Firoozye, N., & Treleaven, P. (2021). Generative adversarial networks for financial trading strategies fine-tuning and combination. Quantitative Finance, 21(5), 797-813.
48. Liew, J. K. S., & Mayster, B. (2018). Forecasting etfs with machine learning algorithms. SSRN.
49. Schneeweis, T., Kazemi, H., & Karavas, V. (2003). Eurex Derivative Products in Alternative Investments: The Case for Hedge Funds. unpublished working paper, CISDM Working Paper.
50. Zhou, X., Pan, Z., Hu, G., Tang, S., & Zhao, C. (2018). Stock market prediction on high-frequency data using generative adversarial nets. Mathematical Problems in Engineering, 2018(1), 4907423.