

DeepUnifiedMom: Unified Time-series Momentum Portfolio Construction via Multi-Task Learning with Multi-Gate Mixture of Experts

Joel Ong^a, Dorien Herremans^a

^a*Singapore University of Technology and Design, 8 Somapah Rd, Singapore 487372*

Abstract

This paper introduces **DeepUnifiedMom**, a deep learning framework that enhances portfolio management through a multi-task learning approach and a multi-gate mixture of experts. The essence of **DeepUnifiedMom** lies in its ability to create unified momentum portfolios that incorporate the dynamics of time series momentum across a spectrum of time frames—a feature often missing in traditional momentum strategies. Our comprehensive backtesting, encompassing diverse asset classes such as equity indexes, fixed income, foreign exchange, and commodities, demonstrates that **DeepUnifiedMom** consistently outperforms benchmark models, even after factoring in transaction costs. This superior performance underscores **DeepUnifiedMom**'s capability to capture the full spectrum of momentum opportunities within financial markets. The findings highlight **DeepUnifiedMom** as an effective tool for practitioners looking to exploit the entire range of momentum opportunities. It offers a compelling solution for improving risk-adjusted returns and is a valuable strategy for navigating the complexities of portfolio management.

Keywords: Deep Learning, Forecasting, Multi-Gate, Mixture-of-Experts, Portfolio Construction, Momentum

*Joel Ong

Email addresses: joel_ong@mymail.sutd.edu.sg (Joel Ong),
dorien_herremans@sutd.edu.sg (Dorien Herremans)

1. Introduction

Time-series momentum (TSMOM) strategies are a systematic approach in finance that leverages the persistence of asset returns over time. These strategies aim to exploit the continuation of underlying trends by establishing long positions during uptrends and short positions during downtrends (Jegadeesh & Titman, 1993, 2001). The concept of momentum has garnered extensive attention in financial literature, underscoring its importance. Research by Moskowitz et al. (2012) highlights the effectiveness of TSMOM, showcasing impressive risk-adjusted returns by simply buying assets with positive past 12-month returns. Further studies across various asset classes corroborate these findings, emphasizing the robustness of TSMOM strategies (Georgopoulou & Wang (2016); Levine & Pedersen (2016); Hurst et al. (2017); Malitskaia (2020)). At the core of TSMOM strategies is volatility scaling, a crucial method for managing exposure to market volatility (Baltas & Kosowski, 2012). Through adjustments in exposure levels during periods of low and high volatility, TSMOM strategies effectively mitigate the risk of significant losses during market turbulence (Harvey et al., 2020). This approach has proven invaluable in enhancing Sharpe ratios, curbing extreme tail returns, and limiting maximum drawdowns in portfolios of risky assets (Barroso & Santa-Clara, 2015; Daniel & Moskowitz, 2016; Ong & Herremans, 2023).

However, TSMOM strategies often fail to account for the interactions between different assets within a portfolio (Nilsson, 2015; Pu et al., 2023). This oversight can lead to excessive risk exposure, as these strategies treat each asset in isolation without considering how they correlate and interact. For example, the simultaneous momentum trends in equities, commodities, and currencies can amplify overall portfolio risk if not managed cohesively. Overlooking these interactions can diminish diversification benefits and increase the likelihood of significant drawdowns during market turbulence.

Moreover, various asset classes, and even individual assets within those classes, exhibit distinct momentum dynamics with differing trend speeds. This variation makes risk allocation challenging, as applying a one-size-fits-all approach to trend speed can result in less-than-ideal investment outcomes (Levine & Pedersen, 2016; Duan, 2023). To address this, some practitioners implement multiple TSMOM portfolios, each tailored to a specific trend speed, and distribute capital among them (Tzotchev, 2018; Zambrano & Rizzolo, 2022; Goulding et al., 2023). Nevertheless, this approach can

still result in inefficient capital distribution, indicating a need for more refined methods that can adeptly navigate the varying momentum speeds of different assets and asset classes for more effective capital allocation.

We propose an innovative approach to bridge the research gaps identified earlier inspired by recent advancements in deep learning for portfolio construction. The proposed method utilizes a deep Multi-task Learning framework combined with a Multi-gate Mixture-of-Experts architecture to develop a momentum portfolio (Jacobs et al., 1991; Ma et al., 2018). Since its development more than three decades ago, the Mixture-of-Experts (MoE) approach has become foundational in numerous research areas and has recently been pivotal in advancing the field of natural language processing in large language models (LLMs) (Fedus et al., 2022; Zoph et al., 2022; He et al., 2022; Gale et al., 2022; Shen et al., 2023). Our approach, which we call **DeepUnifiedMom**, aims to seamlessly integrate momentum opportunities across various speeds, enhancing the efficacy of traditional momentum strategies. Our main contributions are: i) introducing a novel Multi-task Learning framework with a Multi-gate Mixture-of-Experts architecture, which facilitates end-to-end learning for multi period portfolio construction to enhance momentum portfolio performance. ii) This study represents the first implementation and examination of Multi-task Learning combined with a Multi-gate Mixture-of-Experts approach, specifically within portfolio construction. iii) We provide a comprehensive experimental analysis to evaluate and understand the performance outcomes of this innovative methodology against existing momentum strategies and various portfolio construction techniques.

The paper is organized as follows. In Section 2, we discuss existing work on constructing both classical and deep-learning momentum portfolios. Section 3 presents the proposed **DeepUnifiedMom** model. Next, in Section 4 we present the setup of the different experiments, including a description of the dataset, benchmark models, and proposed backtesting strategy. In Section 5, the results of our experiments are presented. Finally, in Section 6 we summarize our findings and suggest directions for future research.

2. Related Work

Research in deep learning in finance is well-established, with numerous studies leveraging these techniques to enhance prediction accuracy, portfolio optimization, and risk assessment. Zhao & Yang (2023) introduce a hybrid model, SA-DLSTM, combining emotion-enhanced convolutional neural net-

works (ECNN), denoising autoencoders (DAE), and long short-term memory (LSTM) models to predict stock price movements by analyzing sentiment from user-generated comments. Wang et al. (2022) propose a portfolio construction model integrating the KMV model with a multiobjective water cycle algorithm, enhancing portfolio evaluation and stability by incorporating financial data from listed companies. Min et al. (2021) address conservatism in worst-case robust portfolio optimization by suggesting hybrid models that use LSTM and XGBoost to forecast market movements and generate hyperparameters for modeling. Lin et al. (2022) develop a multiagent-based deep reinforcement learning framework for portfolio management, featuring a two-level nested agent structure and a custom reward function to optimize trading decisions and risk transfer behaviors. Lastly, Ozbayoglu et al. (2020) provide a comprehensive survey of deep learning applications in finance, categorizing models and identifying future research opportunities. These studies collectively demonstrate advancements in applying deep learning to financial applications, highlighting the potential for innovative techniques to improve financial model robustness and performance. However, none of these studies specifically focus on momentum portfolio construction.

Moskowitz et al. (2012) introduced the concept of time-series momentum (TSMOM), demonstrating that the excess returns of an asset over the past 12 months strongly predict its future performance. Since introducing this concept, it has become a conventional practice for practitioners to implement momentum-based portfolios (Asness et al., 2014; Hurst et al., 2017; Baltas & Kosowski, 2021). In recent years, the use of deep learning approaches in portfolio construction has gained increasing popularity (Lim et al., 2019; Zhang et al., 2020; Kang et al., 2022; Yu et al., 2022; Pu et al., 2023; Yuan et al., 2023). More notably, in the space of deep learning for portfolio construction, Wood et al. (2022) proposed deep-learning architecture that improves upon traditional time-series momentum and mean-reversion strategies. With multiple attention heads (Vaswani et al., 2017), it tracks diverse market regimes across timescales and offers interpretability by highlighting influential factors and key time steps, refining trading strategies. The attention mechanism helps to enhance learning of long-term dependencies and adaptability to new market conditions like the SARS-CoV-2 crisis. Ong & Herremans (2023) first proposed the application of deep multi-task learning (MTL) in momentum portfolio construction, which incorporates auxiliary tasks related explicitly to volatility forecasting (Parkinson, 1980; Garman & Klass, 1980; Rogers & Satchell, 1991; Yang & Zhang, 2000). The findings highlighted that a com-

prehensive MTL framework, encompassing all proposed auxiliary tasks, not only enhances the risk-adjusted performance of portfolios but also notably reduces maximum drawdowns compared to deep learning models without auxiliary tasks or those with a single auxiliary task. This approach underscores the critical role of effectively selecting auxiliary tasks in MTL settings to optimize portfolio outcomes.

Despite extensive research in the field, a significant gap persists in the development of a unified momentum portfolio capable of adapting to trends of varying speeds. Previous studies have typically focused independently on fast (less than one month), medium (three to six months), or slow (six months to one year) momentum strategies, rather than integrating these approaches into a single, dynamic framework. A unified momentum approach aims to capture momentum across this spectrum, seamlessly adjusting to the changing pace of trends within assets and asset classes in the portfolio. Addressing this gap in the literature represents a crucial advancement that could significantly enhance our understanding of momentum-based trading and lead to more robust portfolio management strategies. By developing such a portfolio, we can better exploit the full range of momentum opportunities in financial markets, resulting in superior risk-adjusted returns compared to existing momentum strategies in the literature.

This work aims to bridge this research gap by presenting a deep learning approach to constructing a unified momentum portfolio. We begin by applying the principles of multi-task learning to train three task-specific networks, each specializing in predicting the forward momentum signal score for one month, three months, and six months ahead. These networks generate portfolios representing fast, medium, and slow momentum strategies. The outputs from these task-specific networks are then fed into a final network, the Capital Allocation Module, which determines the weight allocation for each portfolio. By distributing risk according to the allocations provided by the Capital Allocation Module, we create a unified momentum portfolio that effectively exploits and leverages both short-term and long-term trends across assets and asset classes in the portfolio.

3. Methodology

3.1. Overview

This work proposes to include Multi-Gate Mixture of Experts with a Multi-Task Learning Architecture (Ma et al., 2018). This novel approach

combines Long-Short Term Memory (LSTM) (Hochreiter & Schmidhuber, 1997) modules with the Multi-Gate Mixture of Experts (MoMME) framework (Jacobs et al., 1991) to construct a unified momentum portfolio.

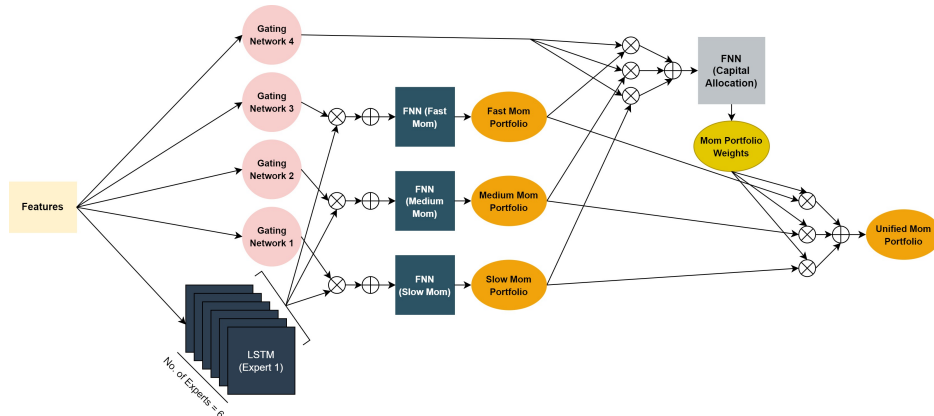


Figure 1: The figure illustrates the proposed architecture, highlighting the flow from shared LSTM experts through task-specific gating and FNN layers and culminating in a final FNN that determines portfolio weights. These weights allocate risk across the portfolios generated by the task-specific FNN layers, which include three momentum portfolios (Fast, Mid, and Slow), each tailored to different trend speeds. The overarching goal of the final unified momentum portfolio is to capitalize on diverse market trends strategically. The symbol \otimes followed by \oplus denotes the weighted sum of the outputs by the gating network with either the LSTM experts or the FNN task-specific network output.

In our proposed architecture, the LSTM experts serve as shared layers forming the backbone of our multi-task learning framework. This setup enables effective parameter sharing across various task-specific pathways yielding two key benefits (Caruana, 1997; Sebastian Thrun, 1998; Ruder, 2017). Firstly, the shared LSTM experts facilitate a more efficient learning process by leveraging commonalities among tasks and consolidating learning efforts. This approach accelerates the training process and enhances overall model performance by drawing on a broader base of data insights. Secondly, the use of shared experts enhances the model’s generalization capabilities. By exposing the model to a variety of tasks within the same learning process, it becomes less prone to overfitting on any single task (Ghosn & Bengio, 1996; Baxter, 2000; Ruder, 2017; Liebel & Körner, 2018). Overall, the integration of shared LSTM experts within our architecture underscores a strategic approach to harnessing the complexities of financial data.

Each task-specific network has a dedicated gating network, a cornerstone

of the Multi-Gate Mixture of Experts (MoMME) framework (Ma et al., 2018). In our work, these gating networks are specialized one-layer FNNs with a softmax activation function. They receive the same feature inputs as the LSTM experts and output a set of weights that sum to one. These weights determine the reliance on the corresponding LSTM experts. By selectively activating relevant LSTM experts, the gating networks enhance the performance of task-specific networks. This selective activation enables more effective learning, as each task-specific network can focus on constructing momentum portfolios tailored to specific speeds or timeframes, ultimately improving the performance of the constructed momentum portfolios. Here, three task-specific networks are trained to minimize the root mean square error (RMSE) between their outputs and the forward-looking TSMOM signals with one-month, three-month, and six-month lookback timeframes. The outputs of these task-specific networks yield the fast, medium, and slow momentum portfolios. Finally, we have a task-specific network called the Capital Allocation Network, supported by a gating network. The gating network receives feature inputs and assigns appropriate weights to the outputs of each task-specific network. The weighted outputs are then fed into the Capital Allocation Network, which is trained to allocate weights to the fast, medium, and slow momentum portfolios generated by the three preceding task-specific networks. This results in a final unified momentum portfolio that strategically capitalizes on opportunities across various market trends.

The output produced by the Capital Allocation Network serves as a set of weights assigned to the fast, medium, and slow momentum portfolios generated by the preceding task-specific networks. These weights determine the allocation of capital across the different portfolios, reflecting the model’s strategic decisions on how to distribute resources among various market trends. By optimizing these weights, the Capital Allocation Network aims to construct a final unified momentum portfolio that effectively captures opportunities across diverse market conditions. Essentially, the portfolio weights determined by the Capital Allocation Network represent the model’s assessment of the relative importance and potential profitability of each momentum portfolio. In summary, the output of the Capital Allocation Network, in conjunction with the portfolio weights for the fast, medium, and slow momentum portfolios, collectively yield the final unified momentum portfolio. This integrated approach enables the model to adaptively allocate resources and strategically capitalize on market trends, ultimately enhancing portfolio performance.

To sum up, our proposed deep mixture of experts’ multi-gate and multi-task learning architecture generates three distinct momentum portfolios, each designed to capture momentum at different timeframes. The Capital Allocation Network also allocates weights to these three momentum portfolios, constructing the final unified momentum portfolio. This approach enables the creation of a unified momentum portfolio in a single step, with the model optimized in an end-to-end fashion. A single-step, end-to-end optimized model is superior because it ensures that all components are trained simultaneously, allowing for seamless integration and interaction between different parts of the model. This approach enhances overall performance by effectively capturing dependencies and relationships within the data, resulting in a more cohesive and robust final portfolio. Our experimental results, detailed in Section 5, substantiate this claim.

3.2. Multi-Task Learning Network

By categorizing momentum into fast (one-month), medium (three-month), and slow (six-month) categories, each task-specific network within our Multi-Task Learning framework is trained to construct portfolios that capture momentum at their respective time frame. This segmentation enhances the model’s ability to detect and leverage the subtle variations in momentum within each asset class and individual asset, thereby improving the precision and relevance of its predictions for future momentum returns across diverse assets. Each task-specific network is represented by a Feedforward Neural Network (FNN), trained to predict the forward-looking time-series momentum (TSMOM) signal. The TSMOM signal is essentially the forward return of an asset, adjusted for its volatility. This adjustment takes into account the risk associated with the asset, providing a normalized measure of momentum that is more comparable across different assets. During the training process, the objective is to minimize the difference between the predicted TSMOM signal and the actual forward-looking TSMOM signal (the ground truth, denoted as \hat{y}_t^i). Specifically, we minimize the Root Mean Squared Error (RMSE) between the predicted output y_t^i and the ground truth \hat{y}_t^i . The RMSE is a commonly used metric for regression tasks, providing a measure of the average magnitude of the errors between predicted and actual values.

$$L_{\text{RMSE}} = \frac{1}{B} \sum_{t \in B} l_{\text{RMSE}_t} \quad (1)$$

where,

$$l_{\text{RMSE}_t} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_t^i - y_t^i)^2}$$

where L_{RMSE} represents the overall RMSE loss over a batch of size B . The term l_{RMSE_t} denotes the RMSE at specific time t in the batch. In this context, n is the total number of assets in the portfolio and y_t^i is the prediction output by the task-specific network for the forward-looking TSMOM signal for asset i at time t . Moreover, the forward-looking TSMOM signal for each asset i at time t , denoted as \hat{y}_t^i represents the ground truth, or the actual observed value of the forward-looking TSMOM signal for each asset i at time t . This ground truth is used during the training phase of your model to compare against the predicted output y_t^i . It is defined as:

$$\hat{y}_t^i = \text{TSMOM}_t^i = \frac{r_{t+1,t+s}^i}{\sigma_{t+s}^i} \quad (2)$$

In this formulation, $r_{t+1,t+s}^i$ represents the returns from $t + 1$ to $t + s$, and σ_{t+s}^i is the standard deviation of the returns, both calculated at a future window s beyond time t . Here, the window s varies according to the speed category of the momentum being analyzed: 20 trading days for **DeepUnifiedMom(Fast)**, 60 trading days for **DeepUnifiedMom(Medium)**, and 120 trading days for **DeepUnifiedMom(Slow)**. This variation allows each task-specific network to fine-tune its learning process to the particular momentum time frame it addresses, thus enhancing the model’s ability to adapt to the distinct market dynamics associated with each speed category. Finally, the return of the fast, medium and slow portfolio can be calculated as follows:

$$r_{t,t+1}^\rho = \frac{1}{n} \sum_{i=1}^n y_{t-1,t}^{\rho,i} \times r_{t,t+1}^i \quad (3)$$

where ρ represents the **DeepUnifiedMom-Fast**, **Medium** and **Slow** portfolio, n is the number of assets in the portfolio, $y_{t-1,t}^{\rho,i}$ is the output of the task-specific network (indicating the weight allocated to asset i for the given momentum timeframe) at time t , $r_{t,t+1}^i$ is the one day return of the asset i and $r_{t,t+1}^\rho$ is the ρ portfolio’s return at time t .

3.3. Capital Allocation Network

The objective of the Capital Allocation Network (CAN) is to generate weights for allocating capital across the various momentum portfolios produced by the task-specific networks. It is implemented as a specialized feed-forward neural network (FNN) with a tanh activation function at each intermediate layer and a softmax activation function at the final layer to ensure the output sums to 1. By applying the weights generated by the CAN to DeepUnifiedMom-Fast, Medium and Slow, we obtain the final unified momentum portfolio, which we term DeepUnifiedMom(CAN). The unified momentum portfolio’s return can be calculated as follows:

$$r_{t,t+1}^{\cup} = \sum_{\rho \in P} w_{t-1,t}^{\rho} \times r_{t,t+1}^{\rho} \quad (4)$$

where P is the set of fast, medium and slow momentum portfolios, $w_{t-1,t}^{\rho}$ is the weight predicted by the CAN for portfolio ρ at time t , $r_{t,t+1}^i$ is the one day return of the asset i and $r_{t,t+1}^{\cup}$ is the return of the unified momentum portfolio.

Since the objective of the CAN differs from those of the task-specific networks, we use the Sharpe Ratio (Sharpe, 1994) as the objective function for training. By utilizing the Sharpe Ratio, as outlined in Equation 5, we direct the model to learn how to generate portfolios optimized for risk-adjusted returns from the input features. The Sharpe Ratio measures an investment’s performance relative to a risk-free asset, adjusting for risk, and offers a comprehensive metric for evaluating the trade-off between risk and return. Incorporating this ratio into the model’s learning process ensures that the constructed portfolios aim to maximize returns while minimizing risk, resulting in superior risk-adjusted performance (Lim et al., 2019; Zhang et al., 2020; Ong & Herremans, 2023).

$$L_{\text{Sharpe Ratio}^{\cup}} = -\frac{\mathbb{E}[r^{\cup}]}{\sigma_{r^{\cup}}} \quad (5)$$

where $\mathbb{E}[r^{\cup}]$ and $\sigma_{r^{\cup}}$ are the mean and standard deviation of the unified momentum portfolio’s realised returns, respectively. The high noise-to-signal ratio in financial data significantly increases the risk of overfitting in deep learning models (Vitells & Gross, 2011; Harvey & Liu, 2014; Bailey et al., 2015; López de Prado, 2018; Israel et al., 2020). During training, we may




Figure 2: Relationship between the Sharpe Ratio with Soft Capping Mechanism and the Standard Sharpe Ratio with Threshold = 0.01

encounter instances where certain batches have an extremely high noise-to-signal ratio. Fitting to this noise can result in a high Sharpe Ratio. Consequently, the model, aiming to maximize the Sharpe ratio, may overfit to these noisy patterns, allocating more weight to these instances and performing well on the training data but failing to generalize to new data. We introduce a modified objective function called the Sharpe Ratio with a Soft Capping Mechanism to mitigate this risk. It first caps the Sharpe ratio at a specified threshold value, ensuring that any value above this threshold is limited. For values that exceed this threshold, the function computes the excess and applies a logarithmic transformation, which reduces the impact of these extreme values by making them grow more slowly. Similarly, it ensures that the Sharpe ratio does not fall below the negative of this threshold by capping the lower end and applying a logarithmic transformation to values below the threshold. This combination of capping and logarithmic adjustments smooths out the extremes, as shown in the equation below:

$$L_{\text{SR}_{\text{soft}}} = -(L + \log(1 + U_e) - \log(1 - L_e)) \quad (6)$$

where,

$$\begin{aligned} U &= \min(\text{SR}, \tau) \\ U_e &= \max(\text{SR} - \tau, 0) \\ L &= \max(U, -\tau) \\ L_e &= \min(\text{SR} - \tau, 0) \end{aligned}$$

Here, SR_{soft} represents the modified Sharpe ratio with soft capping mechanism, SR represents the original Sharpe ratio, and τ is the threshold, which is set to 0.01 during the training process. The resulting Sharpe Ratio with Soft Capping mechanism can be seen in Figure 2. The logarithmic transformation applied to values exceeding the threshold moderates their growth, reducing the impact of extreme values. This smoothing effect stabilizes the training process by preventing abrupt changes in the model’s behavior due to outliers. Consequently, the model is encouraged to focus on more consistent and reliable patterns in the data, leading to better generalization. As a result, the model is less likely to overfit to noise and more likely to capture true underlying signals that are relevant to the objective at hand. Our experimental results, detailed in Section 4 will substantiate the effectiveness of training our proposed architecture with the Sharpe Ratio with a Soft Capping Mechanism. This comparison with models trained without the mechanism will highlight the improvement of performance and generalization capabilities when using the Sharpe Ratio with the Soft Capping Mechanism.

3.4. Loss function

$$L_{\text{total}} = L_{\text{SR}_{\text{soft}}} + \sum_{\rho \in P} L_{\text{RMSE}}^{\rho} \quad (7)$$

Putting it all together, the final loss function of our model, which we minimize during training, can be written as Equation (7). In this equation, L_{total} combines two components. The first component, $L_{\text{SR}_{\text{soft}}}$, represents the loss for the CAN, calculated as the negative Sharpe ratio with a soft capping mechanism. The second component is the sum of the RMSE losses L_{RMSE}^{ρ} for each task-specific network. The overall loss function integrates these elements to guide the training process and optimize the model’s performance.

4. Experimental Setup

4.1. Dataset

Individual futures contracts are subject to expiration dates and varying levels of liquidity, which can hinder the practical analysis of long-term trends. To overcome this challenge, we rely on the Pinnacle Data Corp CLC database as our primary data source for evaluating the proposed model. The data we used in our experimentation spans from January 1990 to December 2023, providing daily data and encompassing over three decades of historical information. The Pinnacle Data Corp CLC database offers a comprehensive continuous price history for 49 futures contracts across diverse asset classes, such as commodities, currencies, fixed income, and equity index futures. We leverage the continuous contract history of each asset, constructed through end-to-end concatenation and price adjustment using the backward-ratio method, to ensure robust analysis of long-term trends.

4.2. Feature Set

We derive a set of time-series momentum features from the daily settled price of the continuous futures by taking the log returns ($r_{t-d,t}^i$) over the past 3 trading days, 5 trading days, 10 trading days, 21 trading days, 63 trading days, 126 trading days, and finally 252 trading days:

$$r_{t-d,t}^i = \ln \frac{P_t^i}{P_{t-d}^i} \quad (8)$$

where $r_{t-d,t}^i$ is the natural logarithm of the d -day return of the asset i at day t , P_t^i is the settled price of asset i at time t and P_{t-d}^i is the settled price of asset i , d trading days ago at time t . To normalize the returns and account for the variability in market conditions, we scale the calculated log returns by the asset's volatility. This approach ensures that the returns are standardized, allowing for a more equitable comparison across different assets and time periods. The scaled return, $\hat{r}_{t-d,t}^i$, is computed as follows:

$$\hat{r}_{t-d,t}^i = \frac{r_{t-d,t}^i}{\sigma_{t-d,t}^i} \quad (9)$$

where $\hat{r}_{t-d,t}^i$ represents the volatility-normalized return over the d -day period for asset i at day t , $r_{t-d,t}^i$ is the log return as previously defined, and $\sigma_{t-d,t}^i$ denotes the volatility of the asset over the d -day period.

Following the approach of Ong & Herremans (2023), our methodology for feature creation is guided by two principal considerations. Firstly, we aim to preserve the essence of time-series momentum by utilizing features that align closely with those employed in the construction of time-series momentum portfolios, as outlined by Moskowitz et al. (2012). Secondly, and of greater significance, we intentionally limit the complexity of our feature engineering to ensure that the observed performance of the portfolios is primarily attributed to the efficacy of our architectural design, rather than to the ingenuity or specificity of the features used. This approach underscores our commitment to validating the inherent strength and adaptability of the architecture in capturing momentum trends, rather than leveraging elaborate feature engineering to enhance portfolio performance artificially.

4.3. Benchmark Models

The concept of Time-Series Momentum (TSMOM) portfolios, as introduced by Moskowitz et al. (2012), forms the cornerstone of our benchmarking process. These portfolios operate on the principle of buying or selling assets based on their performance over the past 12 months. To comprehensively assess our model’s efficacy, we have meticulously crafted a suite of TSMOM portfolios, each tailored to capture distinct momentum horizons:

- TSMOM(1): based on the past one month’s returns.
- TSMOM(3): based on the past three month’s returns.
- TSMOM(6): based on the past six month’s returns.
- TSMOM(12): based on the past twelve month’s returns.
- TSMOM(1,4): An equal-weighted combination of the 1, 2, 3, and 4-month TSMOMs.
- TSMOM(5,8): An equal-weighted combination of the 5, 6, 7, and 8-month TSMOMs.
- TSMOM(9,12): An equal-weighted combination of the 9, 10, 11, and 12-month TSMOMs.
- TSMOM(1,12): An equal-weighted combination of the 1 to 12-month TSMOMs.

This wide array of TSMOM portfolios serves as a benchmark, enabling us to evaluate our proposed model against a spectrum of momentum-based investment strategies across different timeframes. To support our claim that a unified momentum portfolio, which can capitalize on momentum opportunities across various timeframes, `DeepUnifiedMom(CAN)` should outperform all TSMOM benchmark strategies. In addition to that, to rigorously evaluate the final unified momentum portfolio’s performance constructed by the `DeepUnifiedMom(CAN)`, we established two benchmark portfolios: `DeepUnifiedMom(EQWT)` and `DeepUnifiedMom(MVO)`. These carefully chosen benchmarks will help us assess the effectiveness of the unified momentum portfolio constructed by the CAN compared to existing standard portfolio construction techniques.

- `DeepUnifiedMom(EQWT)`: This portfolio equally distributes weights from portfolios constructed by `DeepUnifiedMom-Fast`, `Medium` and `Slow`.
- `DeepUnifiedMom(MVO)` utilizes Mean-Variance Optimization (MVO) (Markowitz, 1952) to construct a portfolio that maximizes the Sharpe ratio.

Constructing a portfolio with equal weighting is a straightforward heuristic that does not involve any optimization process. In contrast, constructing a final portfolio using MVO involves a second optimization process by using the historical returns of `DeepUnifiedMom-Fast`, `Medium` and `Slow` portfolios to calculate the expected returns and covariance matrix, which are then optimized by maximizing the Sharpe ratio to construct the final portfolio. The key drawback of such approaches is that they do not consider the interactions between the components holistically. This fragmented optimization can lead to suboptimal overall performance because each step is optimized in isolation without accounting for all components’ interdependencies and joint effects. Finally, to substantiate our earlier claim that our proposed model’s single-step, end-to-end optimization of a final unified momentum portfolio is more optimal, the final portfolio constructed by the `DeepUnifiedMom(CAN)` should outperform the portfolios constructed using both the `DeepUnifiedMom(EQWT)` and `DeepUnifiedMom(MVO)` methods.

4.4. *Backtest Specifications*

In the following section, we present the backtest results obtained from our proposed model, trained using an expanding window cross-validation

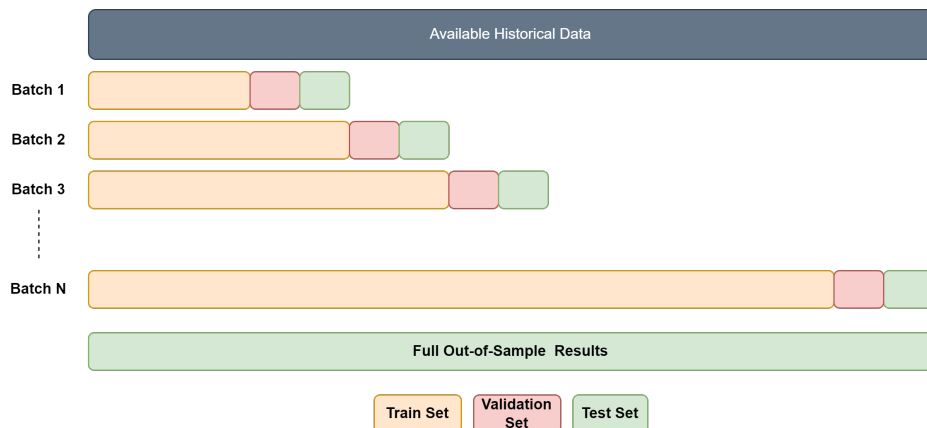


Figure 3: The diagram illustrates the expanding window cross-validation approach used in our model’s training.

approach with the first batch of training data spanning a period of 10 years, as illustrated in Figure 3. Throughout the training phase, 20% of the data was allocated for validation purposes. The trained models were then utilized to construct portfolios for the test set, with each portfolio corresponding to a year’s worth of out-of-sample data. This process was repeated 24 times, resulting in an out-of-sample backtest period from January 2000 to December 2023.

Parameters	Values
Number of LSTM Layers	1, 2, 3
LSTM Hidden Units	64, 126, 252, 512
Number of LSTM Experts	3, 6, 9, 12
Task-Specific Network Layers	2, 3, 4
Task-Specific Network Hidden Units	64, 126, 252, 512

Table 1: Hyperparameter Search Space.

During the backtesting phase, our model underwent training on the designated training dataset. We employed Stochastic Gradient Descent (SGD) with the Adam optimizer to minimize the loss functions. We conducted a grid search on the validation set, guided by the parameter search space outlined in Table 1. The training process was designed to conclude after 20 epochs; however, we incorporated an early stopping mechanism that halts training if

there’s no improvement in validation loss over 5 consecutive epochs. Utilizing an expanding window approach for out-of-sample training and validation, the final training iteration—covering data from January 1990 to December 2023—required approximately 1 hour on a system equipped with an NVIDIA GeForce RTX 2090.

5. Performance Evaluation

As shown in Table 2, our final unified momentum portfolio generated by `DeepUnifiedMom(CAN)`, outperforms all TSMOM benchmark models, including `DeepUnifiedMom-Fast`, `Medium`, and `Slow`, as well as the final portfolios generated by `DeepUnifiedMom - EQWT` and `MVO`. `DeepUnifiedMom(CAN)` achieves a Sharpe ratio of 2.33 and a Sortino ratio of 3.88 while incurring a maximum drawdown of -1.02%. In comparison, the best TSMOM benchmark strategy, `TSMOM(1,12)`, achieves a Sharpe ratio of 1.07 and a Sortino ratio of 1.58, with a drawdown of -2.01%. This demonstrates the superior performance of `DeepUnifiedMom(CAN)` in terms of both risk-adjusted returns and drawdown management.

Additionally, when comparing the portfolios generated by the task-specific networks `DeepUnifiedMom-Fast`, `Medium`, and `Slow`, to the final unified momentum portfolio, the latter consistently outperforms the former on a risk-adjusted basis, notably achieving a much lower maximum drawdown. The best-performing task-specific network, `DeepUnifiedMom(Slow)`, achieves a Sharpe ratio of 1.54 and a Sortino ratio of 2.41 while incurring a maximum drawdown of -3.62%, which is almost 3.5 times larger than the maximum drawdown incurred by the final unified momentum portfolio. When we compare the performance of `DeepUnifiedMom-Fast`, `DeepUnifiedMom-Medium`, and `DeepUnifiedMom-Slow` against TSMOM benchmark strategies, the results are less appealing. Across the board, `DeepUnifiedMom-Fast`, `Medium`, and `Slow` incur much higher maximum drawdowns compared to TSMOM strategies. Running a portfolio that incurs significantly larger drawdown risk without corresponding performance compensation may not be an appealing strategy for many portfolio managers.


Overall, the performance of the final unified momentum portfolio generated by `DeepUnifiedMom(CAN)` compared to the benchmark TSMOM strategies and `DeepUnifiedMom-Fast`, `Medium`, and `Slow` supports the claim that a portfolio capable of capitalizing on a spectrum of momentum opportunities results in a more robust portfolio, contributing to better risk-adjusted

performance and minimizing maximum drawdown.


Portfolios	Ann. Return (%)	Ann. vol (%)	Sharpe	Sortino	Max DD (%)
TSMOM(1)	1.06	1.44	0.73	1.09	-4.60
TSMOM(3)	1.16	1.46	0.80	1.16	-3.37
TSMOM(6)	1.18	1.45	0.82	1.20	-2.91
TSMOM(12)	1.48	1.47	1.01	1.46	-3.28
TSMOM(1,4)	1.27	1.27	1.00	1.49	-2.05
TSMOM(5,8)	1.19	1.37	0.87	1.27	-2.80
TSMOM(9,12)	1.40	1.40	1.00	1.45	-2.60
TSMOM(1,12)	1.30	1.22	1.07	1.58	-2.01
DeepUnifiedMom(Fast)	1.85	1.39	1.34	2.06	-6.32
DeepUnifiedMom(Medium)	1.32	1.34	0.99	1.47	-3.52
DeepUnifiedMom(Slow)	2.14	1.40	1.54	2.41	-3.62
DeepUnifiedMom(CAN)	1.92	0.82	2.33	3.81	-1.02
DeepUnifiedMom(EQWT)	1.79	0.77	2.31	3.71	-0.99
DeepUnifiedMom(MVO)	1.91	1.11	1.72	2.69	-3.13

Table 2: Backtest results (net) for the period from January 2000 to December 2023, with transaction costs set at 3 basis points. The DeepUnifiedMom model results presented here were obtained using the Sharpe Ratio with a Soft Capping mechanism as the loss function during training. Max DD stands for maximum drawdown.


Figure 5 shows that the weights assigned by the DeepUnifiedMom(CAN) to the DeepUnifiedMom portfolios —Fast, Medium, and Slow— maintain remarkable consistency throughout the year. Moreover, the weight allocation strategy implemented by the DeepUnifiedMom(CAN) markedly diverges from traditional equal weighting approaches, demonstrating that DeepUnifiedMom(CAN) does not rely on a basic equal weighting scheme. However, DeepUnifiedMom(CAN) outperforms DeepUnifiedMom(EQWT) by a narrow margin, achieving a Sharpe ratio of 2.33 compared to 2.31, and a Sortino ratio of 3.81 compared to 3.71. While DeepUnifiedMom(EQWT) incurred a slightly smaller maximum drawdown of -0.99%. Comparatively, both the DeepUnifiedMom(CAN) and DeepUnifiedMom(EQWT) portfolios considerably outperform the DeepUnifiedMom(MVO) approach in terms of risk-adjusted returns and maximum drawdown. The Sharpe ratio of DeepUnifiedMom(MVO) portfolio is only 1.72, significantly lower than the 2.33 of the DeepUnifiedMom(CAN) portfolio. The maximum drawdown for the DeepUnifiedMom(MVO) portfolio stands at -3.13%, which is less favorable than the drawdowns experienced by the equal weighted TSMOM portfolios such as TSMOM(1,4), TSMOM(5,8), TSMOM(9,12), and TSMOM(1,12). These findings suggest that the DeepUnifiedMom(MVO) approach may be suboptimal given that




(a) DeepUnifiedMom(CAN) versus TSMOMs.



(b) DeepUnifiedMom(CAN) versus Equal Weight TSMOMs.



(c) DeepUnifiedMom(CAN) versus DeepUnifiedMom - Fast, Medium and Slow.



(d) DeepUnifiedMom(CAN) versus Baseline Portfolio Allocations: Equal Weight and MVO.

Figure 4: Portfolios' Cumulative Returns (%) from January 2000 to December 2023.




Figure 5: The figure illustrates the portfolio weight allocations determined by the Capital Allocation Task Specific Network within the DeepUnifiedMom framework across the Fast, Medium, and Slow portfolios from January 2000 to December 2023.

both DeepUnifiedMom(CAN) and DeepUnifiedMom(EQWT) perform much better. While criticism of MVO is well-established, our results lend further support to these critiques, emphasizing the need for ongoing refinement and exploration of alternative strategies (Michaud & Michaud, 2007; Bailey & Lopez de Prado, 2012; Jurczenko & Teiletche, 2015; Lopez de Prado, 2016). In conclusion, the performance analysis of DeepUnifiedMom(CAN) in comparison to benchmark strategies and DeepUnifiedMom portfolios —Fast, Medium, and Slow, underscores its effectiveness in portfolio management. By outperforming TSMOM benchmarks and task-specific networks in terms of risk-adjusted returns and drawdown management, DeepUnifiedMom(CAN) demonstrates its ability to capitalize on momentum opportunities across a spectrum of market conditions. This suggests that a unified approach to portfolio construction, such as DeepUnifiedMom(CAN), leads to a more robust portfolio, contributing to improved risk-adjusted performance and minimized maximum drawdown. These findings highlight the potential of DeepUnifiedMom(CAN) as a practical and promising solution for investors seeking enhanced portfolio performance in dynamic market environments.

Table 3 reveals that when DeepUnifiedMom(CAN) is trained using the Sharpe Ratio with a Soft Capping Mechanism, the resulting portfolios consistently perform better than those trained using only the Sharpe Ratio. Specifically, DeepUnifiedMom(CAN) trained with the modified Sharpe Ratio achieved a Sharpe Ratio of 2.33 and a Sortino Ratio of 3.81, compared to 2.14 and 3.43, respectively, for those trained with the standard Sharpe Ratio. Additionally, the task-specific networks constructing the Fast, Medium, and Slow portfolios also outperform their counterparts. Overall, the results

Portfolios	Ann. Return (%)	Ann. vol (%)	Sharpe	Sortino	Max DD (%)
Sharpe Ratio with Soft Capping Mechanism (Threshold = 0.01)					
DeepUnifiedMom(Fast)	1.85	1.39	1.34	2.06	-6.32
DeepUnifiedMom(Medium)	1.32	1.34	0.99	1.47	-3.52
DeepUnifiedMom(Slow)	2.14	1.40	1.54	2.41	-3.62
DeepUnifiedMom(CAN)	1.92	0.82	2.33	3.81	-1.02
DeepUnifiedMom(EQWT)	1.79	0.77	2.31	3.71	-0.99
DeepUnifiedMom(MVO)	1.91	1.11	1.72	2.69	-3.13
Sharpe Ratio					
DeepUnifiedMom(Fast)	1.73	1.42	1.22	1.87	-6.33
DeepUnifiedMom(Medium)	1.20	1.39	0.87	1.28	-3.61
DeepUnifiedMom(Slow)	2.03	1.43	1.42	2.20	-4.00
DeepUnifiedMom(CAN)	1.80	0.84	2.14	3.43	-1.10
DeepUnifiedMom(EQWT)	1.67	0.79	2.11	3.33	-1.21
DeepUnifiedMom(MVO)	1.72	1.17	1.47	2.26	-3.64

Table 3: Here are the backtest metrics (net) for the period from January 2000 to December 2023, with transaction costs set at 3 basis points. The DeepUnifiedMom model results presented here were obtained using the Sharpe Ratio with a Soft Capping mechanism as the loss function during training.

are promising, indicating that further research into improved Sharpe Ratio objective functions for training deep learning model is worthwhile.

6. Conclusion

The proposed DeepUnifiedMom framework represents a significant advancement in applying deep learning in portfolio management, adeptly addressing the limitations of traditional momentum strategies. At its core, DeepUnifiedMom leverages advanced deep learning, employing a multi-task learning approach and a multi-gate mixture of experts to construct unified momentum portfolios. Our extensive backtesting, spanning various asset classes such as equity indexes, bonds, currencies, and commodities, has consistently shown that DeepUnifiedMom surpasses benchmarks, maintaining its superior performance even after accounting for transaction costs. This highlights its ability to construct a final portfolio that accounts for a wide spectrum of momentum opportunities in the financial market in an end-to-end fashion. This new approach to using deep learning in investment strategies showcases the benefits of advanced computational techniques in financial decision-making and outcomes. The model was developed using Python and Pytorch; the framework is accessible for review and utilization, with its source code publicly available online¹. Future research will focus on enhanc-

¹https://github.com/joelowj/unified_mom_mmoe

ing the DeepUnifiedMom framework by incorporating sparsity into the gating mechanism and utilizing more sophisticated deep learning architectures, such as the Transformer model, for time-series analysis. Additionally, efforts will be made to integrate explainable AI techniques into the portfolio construction process, aiming to increase the transparency and interpretability of the DeepUnifiedMom framework. This progression will not only refine the model's performance but also bolster user trust and understanding of how AI-driven decisions are made within the portfolio management context.

References

- Asness, C., Frazzini, A., Israel, R., & Moskowitz, T. (2014). Fact, fiction, and momentum investing. *The Journal of Portfolio Management*, *40*, 75–92.
- Bailey, D., & Lopez de Prado, M. (2012). The sharpe ratio efficient frontier. *The Journal of Risk*, *15*, 3–44. doi:10.21314/JOR.2012.255.
- Bailey, D. H., Borwein, J. M., Lopez de Prado, M., & Zhu, Q. J. (2015). The probability of backtest overfitting. *Journal of Computational Finance*, . Forthcoming.
- Baltas, A., & Kosowski, R. (2012). Improving time-series momentum strategies: The role of trading signals and volatility estimators. *SSRN Electronic Journal*, .
- Baltas, N., & Kosowski, R. (2021). Demystifying time-series momentum strategies: Volatility estimators, trading rules and pairwise correlations. In *Market momentum theory and practice* (p. 30–63). Wiley.
- Barroso, P., & Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics*, *116*, 111–120.
- Baxter, J. (2000). A model of inductive bias learning. *J. Artif. Int. Res.*, *12*, 149–198.
- Caruana, R. (1997). Multitask learning. *Machine Learning*, *28*, 41–75.
- Daniel, K., & Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics*, *122*, 221–247.

- Duan, S. (2023). Performance of time-series momentum strategy: Us evidence. *Advances in Economics, Management and Political Sciences*, 35, 45–54. doi:10.54254/2754-1169/35/20231722.
- Fedus, W., Zoph, B., & Shazeer, N. (2022). Switch transformers: scaling to trillion parameter models with simple and efficient sparsity. *J. Mach. Learn. Res.*, 23.
- Gale, T., Narayanan, D., Young, C., & Zaharia, M. (2022). Megablocks: Efficient sparse training with mixture-of-experts. [arXiv:2211.15841](https://arxiv.org/abs/2211.15841).
- Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *The Journal of Business*, 53, 67–78.
- Georgopoulou, A. A., & Wang, G. J. (2016). The trend is your friend: Time-series momentum strategies across equity and commodity markets. *SSRN Electronic Journal*, .
- Ghosn, J., & Bengio, Y. (1996). Multi-task learning for stock selection. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems*. MIT Press volume 9. URL: https://proceedings.neurips.cc/paper_files/paper/1996/file/1d72310edc006dadf2190caad5802983-Paper.pdf.
- Goulding, C. L., Harvey, C. R., & Mazzoleni, M. (2023). Momentum turning points. *SSRN Electronic Journal*, .
- Harvey, C. R., Hoyle, E., Rattray, S., & van Hemert, O. (2020). Strategic risk management: Out-of-sample evidence from the covid-19 equity selloff. *SSRN Electronic Journal*, .
- Harvey, C. R., & Liu, Y. (2014). Evaluating trading strategies. *The Journal of Portfolio Management*, 40, 108–118.
- He, J., Zhai, J., Antunes, T., Wang, H., Luo, F., Shi, S., & Li, Q. (2022). Fastermoe: modeling and optimizing training of large-scale dynamic pre-trained models. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming PPOPP '22* (p. 120–134). New York, NY, USA: Association for Computing Machinery. URL: <https://doi.org/10.1145/3503221.3508418>. doi:10.1145/3503221.3508418.

- Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*, 1735–1780. doi:10.1162/neco.1997.9.8.1735.
- Hurst, B., Ooi, Y. H., & Pedersen, L. H. (2017). A century of evidence on trend-following investing. *The Journal of Portfolio Management*, *44*, 15–29.
- Israel, R., Kelly, B., & Moskowitz, T. (2020). Can machines “learn” finance? *Journal of Investment Management*, *18*. URL: <https://joim.com/article/can-machines-learn-finance/>.
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation*, *3*, 79–87. doi:10.1162/neco.1991.3.1.79.
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*, 65–91.
- Jegadeesh, N., & Titman, S. (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of Finance*, *56*, 699–720.
- Jurczenko, E., & Teiletche, J. (2015). Active risk-based investing. *SSRN Electronic Journal*, . doi:10.2139/ssrn.2592904.
- Kang, Y., Chen, L., Jia, N., Wei, W., Deng, J., & Qian, H. (2022). A cwgan-gp-based multi-task learning model for consumer credit scoring. *Expert Systems with Applications*, *206*, 117650. URL: <https://www.sciencedirect.com/science/article/pii/S0957417422009538>. doi:<https://doi.org/10.1016/j.eswa.2022.117650>.
- Levine, A., & Pedersen, L. H. (2016). Which trend is your friend? *Financial Analysts Journal*, *72*, 51–66.
- Liebel, L., & Körner, M. (2018). Auxiliary tasks in multi-task learning. *CoRR*, *abs/1805.06334*. arXiv:1805.06334.
- Lim, B., Zohren, S., & Roberts, S. (2019). Enhancing time series momentum strategies using deep neural networks. *arXiv preprint arXiv:1904.04912*, . doi:10.48550/ARXIV.1904.04912.

- Lin, Y.-C., Chen, C.-T., Sang, C.-Y., & Huang, S.-H. (2022). Multiagent-based deep reinforcement learning for risk-shifting portfolio management. *Applied Soft Computing*, *123*, 108894. URL: <https://www.sciencedirect.com/science/article/pii/S1568494622002678>. doi:<https://doi.org/10.1016/j.asoc.2022.108894>.
- Ma, J., Zhao, Z., Yi, X., Chen, J., Hong, L., & Chi, E. H. (2018). Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining KDD '18* (p. 1930–1939). New York, NY, USA: Association for Computing Machinery. URL: <https://doi.org/10.1145/3219819.3220007>. doi:10.1145/3219819.3220007.
- Malitskaia, Y. (2020). Uncovering momentum. *SSRN Electronic Journal*, .
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, *7*, 77–91.
- Michaud, R., & Michaud, R. (2007). Estimation error and portfolio optimization: A resampling solution. *Journal of Investment Management*, *Vol. 6*, pp. 8 – 28.
- Min, L., Dong, J., Liu, J., & Gong, X. (2021). Robust mean-risk portfolio optimization using machine learning-based trade-off parameter. *Applied Soft Computing*, *113*, 107948. URL: <https://www.sciencedirect.com/science/article/pii/S156849462100870X>. doi:<https://doi.org/10.1016/j.asoc.2021.107948>.
- Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics*, *104*, 228–250.
- Nilsson, L. (2015). Trend following - expected returns. *SSRN Electronic Journal*, .
- Ong, J., & Herremans, D. (2023). Constructing time-series momentum portfolios with deep multi-task learning. *Expert Systems with Applications*, *230*, 120587.
- Ozbayoglu, A. M., Gudelek, M. U., & Sezer, O. B. (2020). Deep learning for financial applications : A survey. *Applied Soft Computing*, *93*, 106384. URL: <https://www.sciencedirect.com/science/>

article/pii/S1568494620303240. doi:<https://doi.org/10.1016/j.asoc.2020.106384>.

- Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business*, *53*, 61–65.
- Lopez de Prado, M. (2016). Building diversified portfolios that outperform out of sample:. *The Journal of Portfolio Management*, *42*, 59–69. doi:10.3905/jpm.2016.42.4.059.
- López de Prado, M. (2018). *Advances in Financial Machine Learning*. Wiley.
- Pu, X., Roberts, S., Dong, X., & Zohren, S. (2023). Network momentum across asset classes. [arXiv:2308.11294](https://arxiv.org/abs/2308.11294).
- Rogers, L. C. G., & Satchell, S. E. (1991). Estimating variance from high, low and closing prices. *The Annals of Applied Probability*, *1*, 504–512.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *CoRR*, [abs/1706.05098](https://arxiv.org/abs/1706.05098). [arXiv:1706.05098](https://arxiv.org/abs/1706.05098).
- Sebastian Thrun, L. P. (1998). Learning to learn: Introduction and overview. *Learning to Learn*, (pp. 3–17).
- Sharpe, W. F. (1994). The sharpe ratio. *The Journal of Portfolio Management*, *21*, 49–58.
- Shen, S., Hou, L., Zhou, Y., Du, N., Longpre, S., Wei, J., Chung, H. W., Zoph, B., Fedus, W., Chen, X., Vu, T., Wu, Y., Chen, W., Webson, A., Li, Y., Zhao, V., Yu, H., Keutzer, K., Darrell, T., & Zhou, D. (2023). Mixture-of-experts meets instruction tuning:a winning combination for large language models. [arXiv:2305.14705](https://arxiv.org/abs/2305.14705).
- Tzotchev, D. (2018). Designing robust trend-following system: Behind the scenes of trend-following. *SSRN Electronic Journal*, .
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc. volume 30.

- Vitells, O., & Gross, E. (2011). Estimating the significance of a signal in a multi-dimensional search. *Astroparticle Physics*, *35*, 230–234. doi:10.1016/j.astropartphys.2011.08.005.
- Wang, J., Zhang, H., & Luo, H. (2022). Research on the construction of stock portfolios based on multiobjective water cycle algorithm and kmv algorithm. *Applied Soft Computing*, *115*, 108186. URL: <https://www.sciencedirect.com/science/article/pii/S1568494621010346>. doi:<https://doi.org/10.1016/j.asoc.2021.108186>.
- Wood, K., Giegerich, S., Roberts, S., & Zohren, S. (2022). Trading with the momentum transformer: An intelligent and interpretable architecture. URL: <https://arxiv.org/abs/2112.08534>.
- Yang, D., & Zhang, Q. (2000). Drift-independent volatility estimation based on high, low, open, and close prices. *The Journal of Business*, *73*, 477–492.
- Yu, P., Wynter, L., & Lim, S. H. (2022). Federated reinforcement learning for portfolio management. In H. Ludwig, & N. Baracaldo (Eds.), *Federated Learning: A Comprehensive Overview of Methods and Applications* (pp. 467–482). Cham: Springer International Publishing. URL: https://doi.org/10.1007/978-3-030-96896-0_21. doi:10.1007/978-3-030-96896-0_21.
- Yuan, C., Ma, X., Wang, H., Zhang, C., & Li, X. (2023). Covid19-mlsf: A multi-task learning-based stock market forecasting framework during the covid-19 pandemic. *Expert Systems with Applications*, *217*, 119549. URL: <https://www.sciencedirect.com/science/article/pii/S0957417423000507>. doi:<https://doi.org/10.1016/j.eswa.2023.119549>.
- Zambrano, E. A., & Rizzolo, C. (2022). Long-only multi-asset momentum: Searching for absolute returns. *SSRN Electronic Journal*, .
- Zhang, Z., Zohren, S., & Roberts, S. (2020). Deep learning for portfolio optimisation. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3613600.

- Zhao, Y., & Yang, G. (2023). Deep learning-based integrated framework for stock price movement prediction. *Applied Soft Computing*, 133, 109921. URL: <https://www.sciencedirect.com/science/article/pii/S156849462200970X>. doi:<https://doi.org/10.1016/j.asoc.2022.109921>.
- Zoph, B., Bello, I., Kumar, S., Du, N., Huang, Y., Dean, J., Shazeer, N., & Fedus, W. (2022). St-moe: Designing stable and transferable sparse expert models. [arXiv:2202.08906](https://arxiv.org/abs/2202.08906).