Advances in Artificial Intelligence and Machine Learning; Research 3 (4) 1671-1684 Received 15-9-2023; Accepted 24-11-2023; Published 1-12-2023

Noxtrader: Lstm-Based Stock Return Momentum Prediction For Quantitative Trading

Hsiang-Hui Liu morrisliuting@gmail.com
Department of Computer Science
National TsingHua University
300044 Hsinchu, Taiwan.

Han-Jay Shu jerryshu98@gmail.com
Department of Electrical Engineering and Computer Science
National TsingHua University
300044 Hsinchu, Taiwan.

Wei-Ning Chiu william08162010@gmail.com
Department of Computer Science
National TsingHua University
300044 Hsinchu, Taiwan.

Corresponding Author: Hsiang-Hui Liu

Abstract

We introduce NoxTrader, a sophisticated system designed for portfolio construction and trading execution with the primary objective of achieving profitable outcomes in the stock market, specifically aiming to generate moderate to long-term profits. The underlying learning process of NoxTrader is rooted in the assimilation of valuable insights derived from historical trading data, particularly focusing on time-series analysis due to the nature of the dataset employed. In our approach, we utilize price and volume data of US stock market for feature engineering to generate effective features, including Return Momentum, Week Price Momentum, and Month Price Momentum. We choose the Long Short-Term Memory (LSTM) model to capture continuous price trends and implement dynamic model updates during the trading execution process, enabling the model to continuously adapt to the current market trends. Notably, we have developed a comprehensive trading backtesting system — NoxTrader, which allows us to manage portfolios based on predictive scores and utilize custom evaluation metrics to conduct a thorough assessment of our trading performance. Our rigorous feature engineering and careful selection of prediction targets enable us to generate prediction data with an impressive correlation range between 0.65 and 0.75. Finally, we monitor the dispersion of our prediction data and perform a comparative analysis against actual market data. Through the use of filtering techniques, we improved the initial -60% investment return to 325%.

Keywords: Artificial intelligence, Machine learning, Portfolio management, Quantitative finance, Time-Series analysis.

Citation: Hsiang-Hui Liu, et al. Noxtrader: Lstm-Based Stock Return Momentum Prediction For Quantitative Trading. Advances in Artificial Intelligence and Machine Learning. 2023;3(4):95.

1. INTRODUCTION

In an era characterized by rapid advancements in algorithmic and machine learning-based trading, financial markets are undergoing a profound transformation. While many papers discuss how to enhance market prediction accuracy using artificial intelligence models, there is a scarcity of literature addressing their practical application in real markets or their failures when applied in such contexts. Therefore, this paper is dedicated not only to creating highly accurate models but also to achieving successful real-world applications in financial markets as the ultimate objective.

Certain existing papers provide us with a strong knowledge foundation and offer examples of integrating artificial intelligence models with financial market data. “Machine Learning Approaches in Stock Price Prediction: A Systematic Review” [1], provides a profound overview of the current use of artificial intelligence models in market prediction, including machine learning models such as SVM and random forest, as well as deep learning models like LSTM and RNN, giving us an initial comprehensive understanding. Additionally, “101 Formulaic Alphas” [2], demonstrates methods of utilizing market data and emphasizes the significance of feature engineering in the field of financial market prediction.

Unlike prevailing methodologies, NoxTrader leverages the predictive capabilities of LSTM networks and supervised learning techniques to discern intricate patterns within historical data, effectively capturing the nuanced fluctuations in market prices. NoxTrader introduces a unique perspective on label generation by adopting the concept of “return momentum”, which is the difference of return between two consecutive days, as a predictive label with appropriate filter, as opposed to the more customary use of raw returns. This novel target of prediction distinguishes us significantly from other research papers and entities employing LSTM models for market forecasting, resulting in a substantial enhancement in our market performance.

The subsequent sections of this paper are carefully structured to provide a comprehensive exposition of NoxTrader’s inner workings. We delve into the intricacies of each constituent unit, meticulously detailing the process of feature generation, predictive modeling, and the creation of a robust backtest environment. Our narrative is further enriched by a diverse array of empirical experiments, designed to showcase the results garnered from NoxTrader’s operational deployment and substantiate its potential profitability. Finally, we make a complete discussion and conclude it with some possible future improvements we’re going to make.

2. METHODS

NoxTrader’s implementation consists of three primary components. The initial component, termed “Feature Generation,” assumes the role of crafting distinct facets of features and subsequently conveying them to the model. The subsequent component, denoted as the “Prediction” module, bears the responsibility of harnessing the previously generated features. It assigns scores to individual stocks within our screener based on these features. Following the evaluation of all stocks, a mechanism for portfolio construction and strategy backtesting becomes imperative. This function, referred to as “Backtest,” facilitates the process of portfolio assembly and strategy evaluation. In the subsequent sections, a comprehensive exposition of each of these three components is presented.

2.1 Feature Generation

The historical data supplied by the yfinance platform embodies a set of fundamental raw metrics, encompassing the opening, closing, highest, and lowest values of each individual stocks. Notably, these initial data points, while serving as a foundational basis, lack the requisite depth to facilitate effective learning and precise predictive outcomes within the model. As such, a compelling imperative exists to undertake a transformative process on this raw dataset, thereby bringing in some enlightening features that are primed to empower the model's learning capacity.

In the spirit of holistic pattern recognition, NoxTrader undertakes a multifaceted approach to feature engineering. We collect market data for US companies with a market capitalization of more than 200 billion dollars and draw upon a rich spectrum of insights that collectively enhance the model's ability. Our data collection encompasses various aspects, focusing primarily on three key areas: fundamental analysis, chips, and technical analysis. Given that our main predictive target is 'return momentum'—a concept rooted in technical analysis—we've decided to utilize features exclusively from this domain. We believe that aligning the features closely with the output enhances the model's learning efficiency. Consequently, our feature selection is primarily centered around momentum-based and trend-based attributes. The subsequent enumeration explain the diverse array of features meticulously integrated into the NoxTrader framework:

> Returns—Showing how much the closing price has grown compared to its previous value, usually from the day before. This measure helps the model understand the changes in stock values over time more effectively.
>
> ReturnMomentum—Building upon the Returns paradigm, ReturnMomentum augments the model's insight by quantifying the differential between the current day's Returns and the Returns observed on the preceding day. This parameter encapsulates the intraday dynamics that could potentially exert influence on future stock performance.
>
> ReturnAcceleration—Delving deeper into the temporal dynamics, ReturnAcceleration discerns the variance in ReturnMomentum from one day to the next, thereby encapsulating the intricate curvature of the stock valuation trajectory. This higher-order derivative augments the model's ability to capture evolving trends.
>
> WeekPriceMomentum—Harnessing a broader temporal horizon, WeekPriceMomentum appraises the growth rate of the closing price compared to its state a week prior. This temporal frame of reference imparts a long-range perspective on valuation trends, arming the model with insights into sustained momentum [2, 3].
>
> MonthPriceMomentum—Extending the purview even further, MonthPriceMomentum represents the growth rate of the closing price in relation to its value a month preceding. This elongated temporal context serves as a harbinger of extended trends, enriching the model's predictive prowess [2, 3].
>
> VolumeVelocity—Recognizing the pivotal role of trading volume, VolumeVelocity indexes the growth rate of trading volume in comparison to the preceding day. This feature offers insights into market sentiment and potential shifts in supply-demand dynamics [4].

These intricate features collaboratively contribute to the enhancement of NoxTrader's learning efficacy. By seamlessly integrating these multifaceted metrics, the model is able to glean a more profound comprehension of market nuances, thereby enhancing its predictive power and heightening the precision of its projection.

2.2 Prediction Method

In this section, we detail the methodology employed for predicting stock price changes using a Long Short-Term Memory (LSTM) network. The dataset, feature extraction, model architecture, loss function, performance evaluation, and the rationale behind the chosen approach are discussed.

The dataset comprises individual data instances, each consisting of two main components: features and labels. The features encapsulate information from the stock market for the past 10 days, including features mentioned in part A. Notably, the features incorporate data from the current day as well as the preceding 9 days. The labels represent the difference between return of two consecutive days, namely return momentum. A training set of 240 such instances is constructed, chosen based on utilizing historical data from the past year. Since stock market only opens 5 days a week, approximately 20 trading days correspond to a month.

Our approach employs the Long Short-Term Memory (LSTM) network as the primary model architecture. Specifically, our model incorporates a Long Short-Term Memory (LSTM) layer with 128 units. We've opted for a timestep of 25 in this layer, as this medium range aligns well with our analytical goals. Following the LSTM, the model includes a fully connected layer that outputs a single unit, serving as the foundation for our predictions. We explored deeper network architectures but did not observe significant improvements in outcomes compared to this simpler structure. Therefore, prioritizing generalizability, we selected this more straightforward model as our final configuration. LSTMs excel at capturing sequential dependencies, a crucial aspect for modeling stock prices with inherent temporal patterns. This capability allows the model to effectively consider historical contexts when predicting future stock returns. Additionally, LSTMs exhibit proficiency in handling irregularities commonly found in financial markets, such as sudden changes in volatility or unexpected events, providing adaptability to dynamic market conditions. The model's automatic feature extraction eliminates the need for manual engineering, enabling it to discern and utilize relevant features from historical stock price data. Moreover, LSTMs showcase robustness to the evolving temporal dynamics of financial markets, making them a preferred choice for capturing changing market conditions compared to traditional machine learning models.

The Mean Squared Error (MSE) is chosen as the loss function for the LSTM model. This selection aligns with the objective of minimizing the discrepancy between predicted and actual stock return momentum. The MSE quantifies the average squared difference between predicted and actual values, enabling the model to learn optimal parameters that minimize this error.

To assess the performance of the model, we not only employ the MSE loss function but also calculate the correlation between our predictions and the true labels. This correlation metric ranges predominantly between 0.65 and 0.75, signifying a meaningful correspondence between predicted and actual trends. Importantly, considering the temporal nature of stock data, retraining the model is necessary every 10 days to ensure its adaptability to evolving market patterns. The average predicted correlation comparisons between LSTM and other traditional models are showed in FIGURE 1.

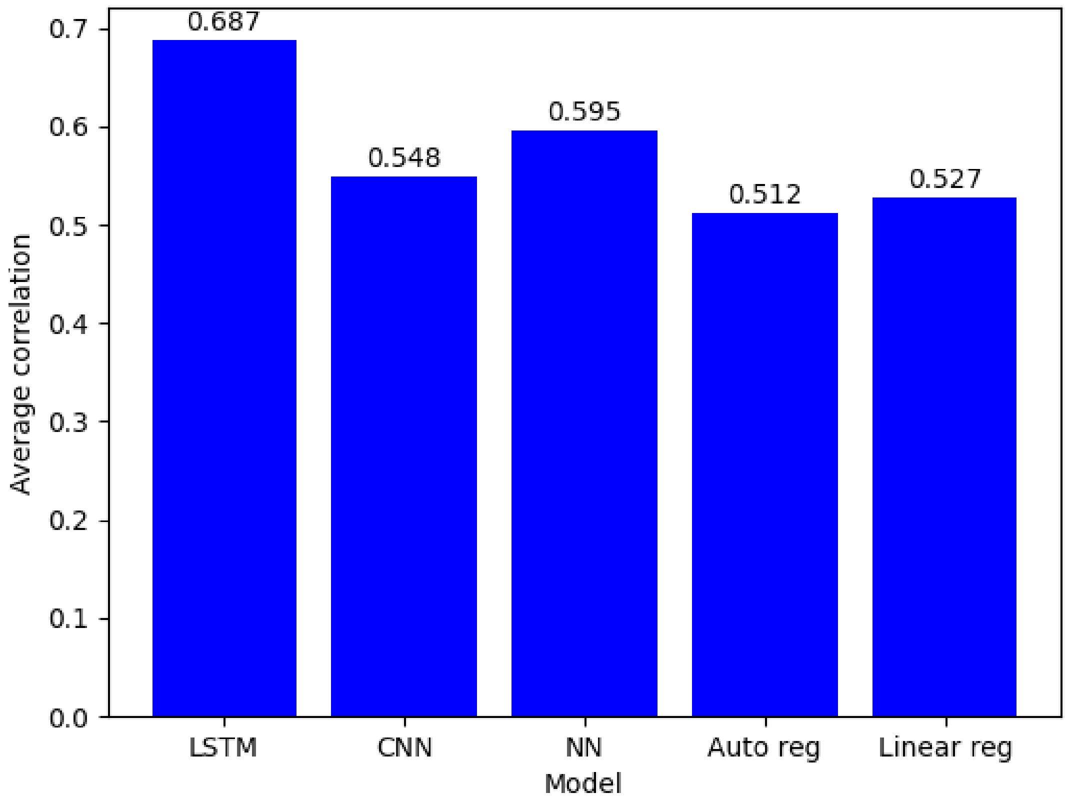

Figure 1: average predicted correlation comparison

While a predictive horizon beyond 10 days might be appealing, we observed a decline in correlation beyond this point. Specifically, if we use the same model to make predictions for 40 days, the correlation for the initial 20 days significantly surpasses the latter 20 days, implying reduced accuracy for longer prediction horizons. Consequently, forecasting stock prices over an entire year would necessitate the training of 24 separate models, each specialized for a specific 10-day prediction window. As for predicting stock momentum across various volatility patterns, we strategically employed sectors as a benchmark for discerning distinct patterns, given their propensity to exhibit diverse market behaviors. Our dataset encompasses a comprehensive set of 9 different sectors, as illustrated in TABLE 1. Delving into the performance outcomes depicted in FIGURE 2, our LSTM model demonstrates a remarkable capability to grasp and adapt to volatility patterns across these diverse sectors. This analysis underscores that our model not only excels in adapting to stocks manifesting specific patterns but also showcases a robust overall ability to uncover underlying clues within a myriad of sectors. The versatility exhibited by our model across sectors substantiates its effectiveness in comprehending and leveraging various market dynamics. This adaptability is a pivotal aspect, illustrating the model's capacity to discern nuanced patterns within different industry landscapes, thereby fortifying its credibility for widespread applicability.

Table 1: Tested sectors and abbreviation

Sector name	Abbreviation
Technology	TEC
Telecommunication	TEL
Healthcare	HC
Financials	FIN
Consumer Discretionary	CD
Consumer Staples	CS
Industrials	IND
Energy	ENE
Utilities	UTIL

2.3 BackTest Environment

The presented backtesting framework is designed with the intention of harnessing the outcomes generated by our model. This is accomplished through the conversion of model-generated labels into corresponding stock positions, followed by the simulation of trading activities under conditions resembling those of the real market. The outlined system comprises two sections: label-to-position conversion and performance evaluation.

The section of label-to-position conversion consists of a two-fold procedure, involving filtration and capital allocation. In the filtration procedure, the output y_i for each score x_i generated by the model and a threshold t is assigned as:

$$y_i = \begin{cases} x_i + t & x_i < -t \\ x_i - t & x_i > t \\ 0 & otherwise \end{cases}$$

x_i that falls within $(-t,\ t)$ are classified as "no-trade" and would be assigned to 0. The remaining candidates will undergo an adjustment process, involving a subtraction by t if the value is positive, and an addition by t if the value is negative.

In the capital allocation procedure, the allocated c_i for each y_i generated from the filtration process is assigned as:

$$c_i = \frac{y_i}{\sum_{j=1}^{n} y_j} \times Equity$$

On a daily basis, each candidate will be allocated a position value. The position value is determined by multiplying the total equity by their respective weighted average.

For performance evaluation, the computation of portfolio gains and losses entails the multiplication of positions with actual market returns. The subsequent benchmarks are introduced to provide enhanced comprehension of the portfolio's performance:

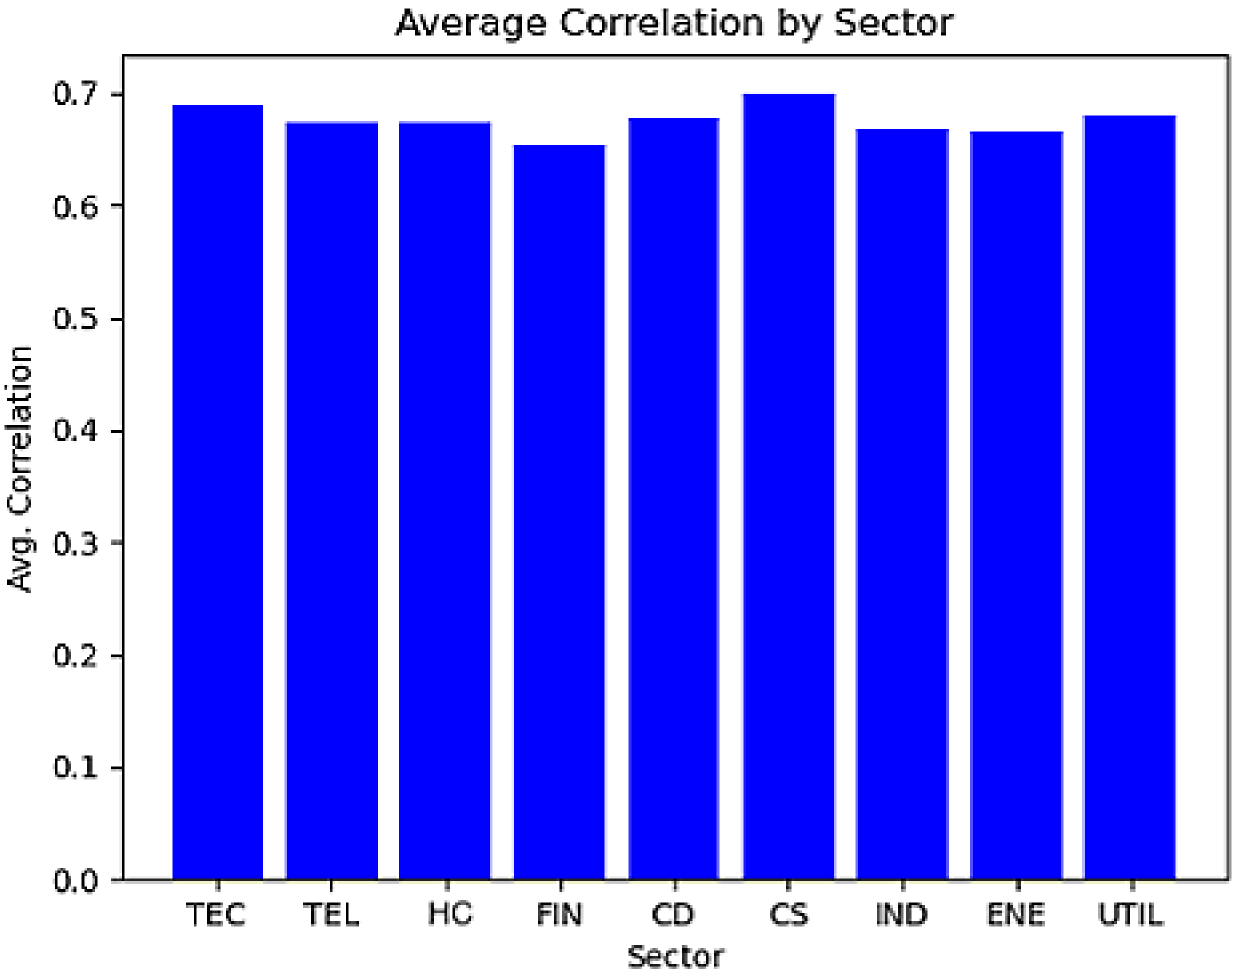

Figure 2: Average predicted correlation among different sectors

In the market days—the ratio of the number of days in the market to the total number of days.

Position Qualified—the ratio of the number of position in the market to the total number of position.

Annual Returns—the total returns earned by an investment in a year, considering compounding effects.

Win Rate—the percentage of successful trades among all position.

Max Drawdown—the largest percentage decline in an investment's value from its peak to the lowest point.

3. EXPERIMENT RESULT

The following experiment will demonstrate our approach to selecting model labels, establishing evaluation criteria, and ultimately formulating a comprehensive strategy.

3.1 Refining Label Representation

During the first stage of our research, we take "returns" as the predicted label. Although the overall correlation stood at around 0, it was intriguing to observe that the predicted trends of upward and downward movements exhibited some semblance to the true labels. Thus, a pivotal change was introduced in the way we represented the true labels. Shifting from the initial approach of using stock returns, we adopted return momentum as the new label representation. This alteration yielded remarkable results as the correlation coefficient surged to an impressive 0.6. Please refer to FIGURE 3. The solid line represents the return momentum of true market while the broken line represents the predicted return momentum. This change reaffirmed the importance of label representation in the predictive accuracy of the model. However, a concerning pattern emerged when the predicted results were inversely transformed to the original returns. The correlation plummeted back to negligible levels, puzzling us further.

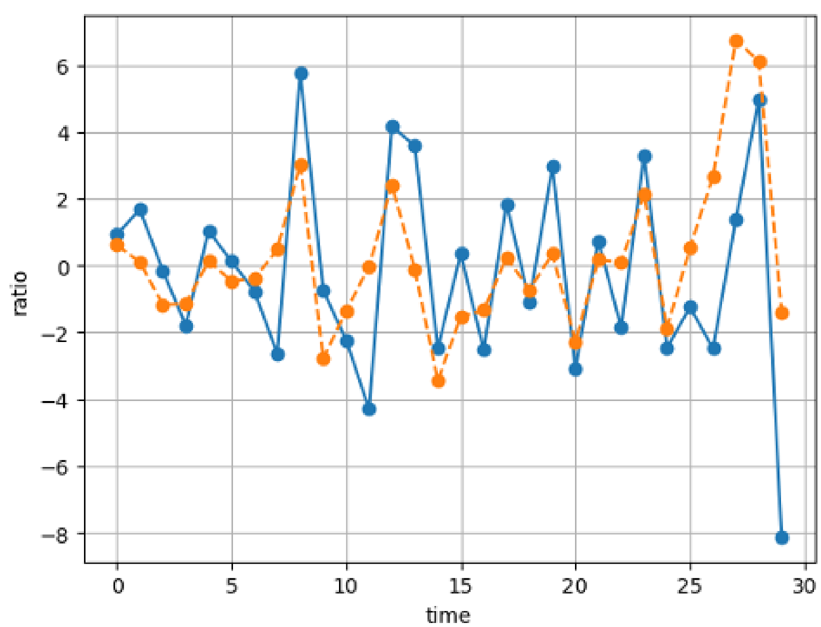

Figure 3: Correlation between Predicted Return Momentum and True Labels

3.2 Return Momentum Backtesting

Our focus shifted towards understanding the stack discrepancy observed during the transformation from differentials to actual price changes. Intriguingly, employing the true return momentum as inputs for backtesting proved highly effective. Please refer to FIGURE 4. This unexpected success hinted at an inherent capability of return momentum to profit effectively without transforming back to return. Despite this, using the predicted return momentum for backtesting yielded disappointing results, indicating a major discrepancy between the model's predictive power and its application to the actual data. Please refer to FIGURE 5.

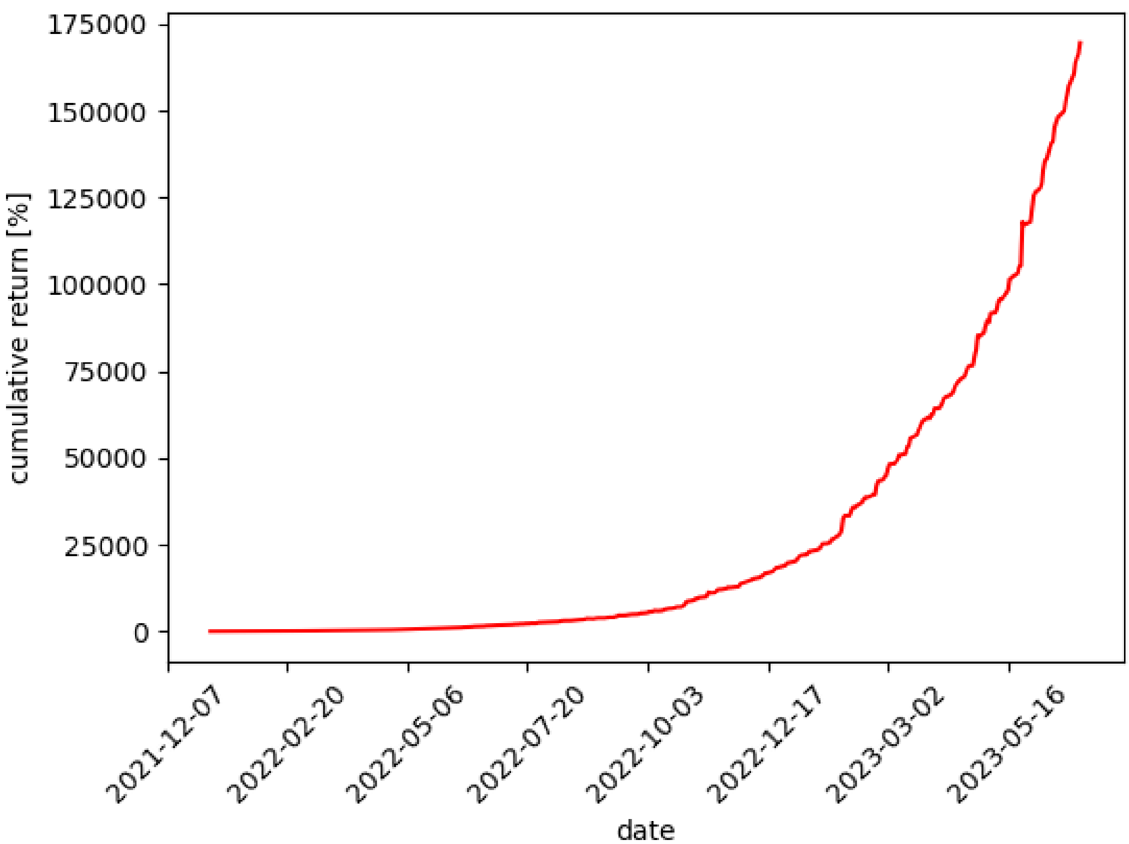

Figure 4: True Label of Return Momentum as Inputs for Backtesting

3.3 Correlation-Label Relationship

To evaluate the viability of the correlation between predicted and true labels, stocks exhibiting a correlation exceeding 0.7 within each group are selected for a backtesting exercise. Each group represents a four-month testing period. The findings reveal that the observed correlation is indeed viable, substantiated by an annual return of 122.95% and a maximum drawdown of 11.14%.

Our investigation delves into establishing the connection between labels and correlation, building on the understanding that high correlation signifies superior performance. During this analysis, we incorporate the standard deviation of labels within each distinct group. It reveal a discernible

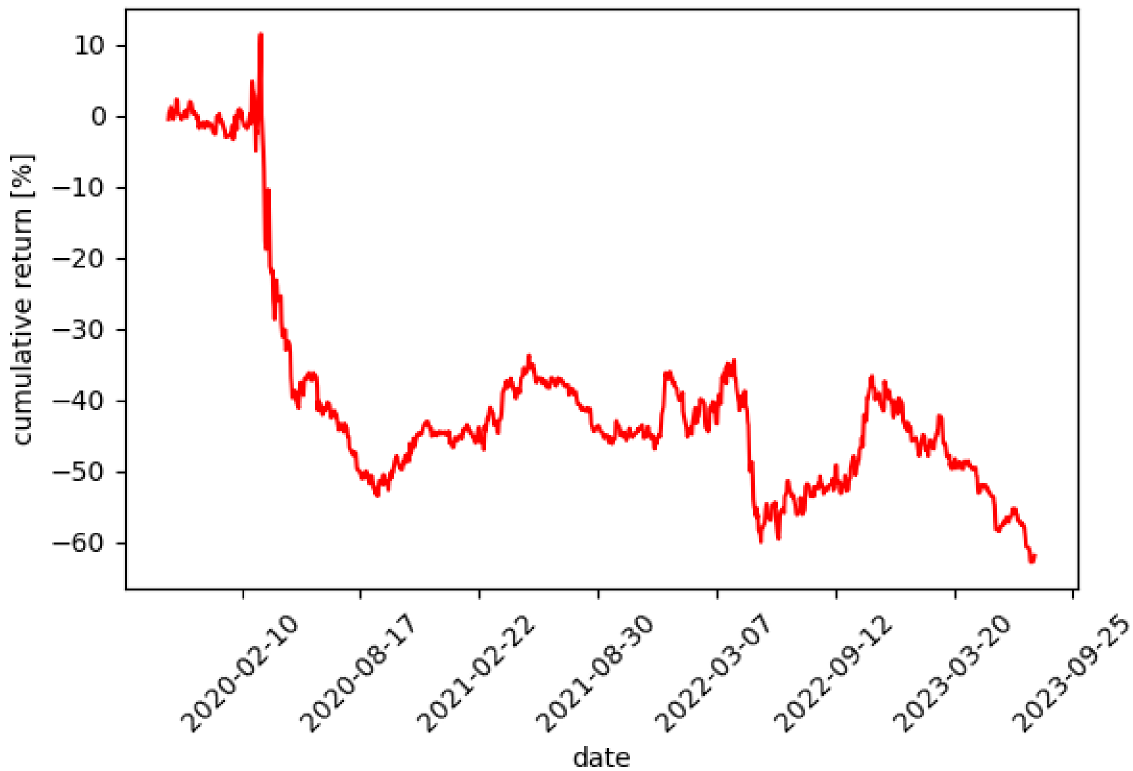

Figure 5: Predicted Return Momentum as Inputs for Backtesting

pattern: instances of heightened standard deviation correlate with relatively high correlation. For an in-depth visual representation, please refer to FIGURE 6.

3.4 Final Strategy

We observed that data with a higher standard deviation exhibits a higher correlation. In other words, more dispersed data tends to yield better performance. Consequently, we hypothesize that labels with smaller numerical values may contribute to reducing the data dispersion. Through experiments, we identified a threshold for segregating data into categories of large and small values. In our experiments, we noticed a significant improvement in strategy performance when excluding data with absolute label values below 7. Additionally, for the remaining labels with absolute values greater than 7, we subtracted 7 to create new labels. This adjustment increased sensitivity to differences within the range of 7 to the label in the allocation of funds. The empirical outcomes are striking: a cumulative return of 325.38% over a span of six years, an annual return of 37.72%, and a maximum drawdown of 23.84%. The visual representation of a gracefully undulating curve, akin to an exponential curve, serves as compelling evidence of the pronounced positive impact of

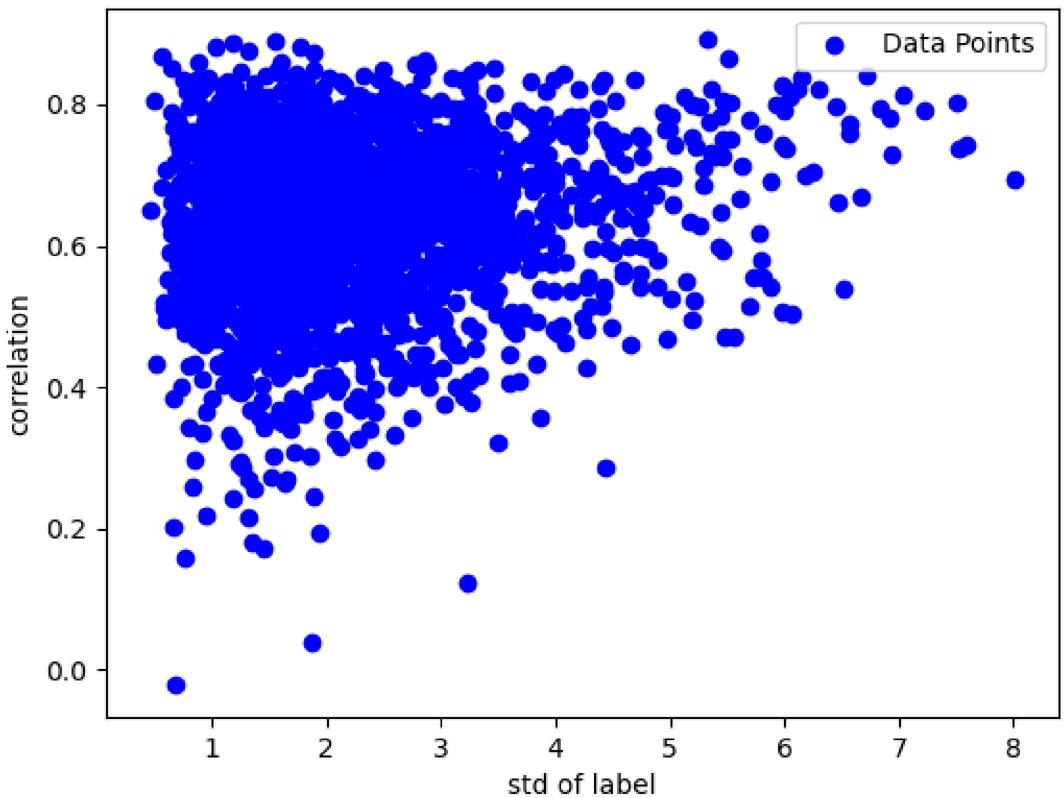

Figure 6: Relation between Standard Deviation of Labels and Correlation

our filtration approach. A comprehensive tabulation can be found in TABLE 2, while FIGURE 7, visually reinforces these outcomes.

The backtesting results reveal the following phenomena. Out of 1140 trading days, we engaged in trading on 41.5% of the days, averaging one trade every 2.4 days. Simultaneously, our trading success rate is 61%. Maintaining a high success rate for a strategy with a trading frequency of 41.5% is not an easy task. The total trading duration over six years further validates that our experimental results are not merely coincidental. Regarding risk and return, the total return is 352.38%, with a maximum drawdown of 23.84%. This implies that for every 1% of return assumed, the strategy yields a 14.78% return. Converting the total return to an annualized rate, the strategy demonstrates an annual return of 37.72%. This categorizes the strategy as high-yield, offering not only a hedge against inflation risks but also providing additional income for investors.

We analyze the factors contributing to the success of the final strategy. Firstly, the primary determinant is the selection of the target for prediction. Return momentum provides us with robust predictive performance. Serving as a capital allocation factor, return momentum can reflect market fluctuations. In cases where low numerical return momentum is discarded, it aligns more closely

Table 2: Result data of final Strategy

Property	Value
Start Date	2018-10-14
End Date	2023-04-17
Market Days	1140
In the Market Days	41.491228%
Position Qualified	2.353445%
Commision	0.01%
Equity Initial	1000000
Equity Final	4253800
Return	352.380000%
Ann. Return	37.718672%
Win Rate	61.090909%
Max Drawdown	23.837605%

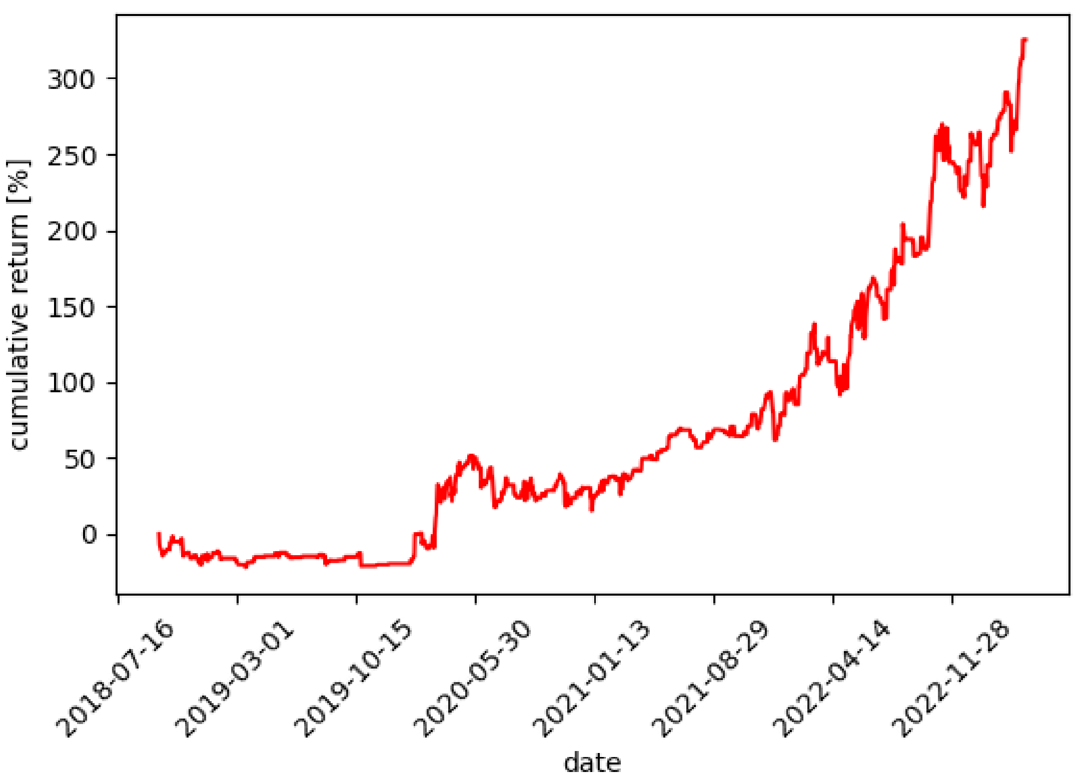

Figure 7: Cumulative Return Chart of Final Strategy

with the daily market trends. Secondly, we incorporate appropriate price and trading volume information as features, opting for Long Short-Term Memory (LSTM) among numerous machine learning models to incorporate relevant market information from the past several days. The formulation of the loss function enables us to approach accurate predictions effectively. Finally, the statistical methods we employ involve multiple experiments on predicted labels and model performance. These experiments provide us with an in-depth understanding of the data characteristics and model performance, leading to the development of the final filtering and capital allocation methods.

4. CONCLUSION

In contrast to conventional methodologies, NoxTrader diverges by introducing an innovative approach to label generation. This entails the incorporation of the "return momentum" concept as a predictive label combined with carefully defined filters. This distinctive strategy yields notable outcomes that underscore its significance and efficacy. The feature engineering process and model selection of the LSTM model effectively leverage the favorable properties of return momentum to create a stable model performance. The filtration method and the capital allocation method is also crucial in this strategy as it pertains to risk control. Through experiments, we identify a threshold that ensures the quality of model performance and allocate weights after the filtration process, reducing the risk associated with uncertain numerical values participating in capital allocation. Consequently, larger return momentum values receive larger capital allocations, allowing more capital to be allocated to stock positions with a higher likelihood of profitability. Through the synergy of these various methods, this distinctive strategy produces noteworthy outcomes, highlighting its significance and efficacy.

The current experimental approach has certain limitations. The threshold in the filtration process is a specific value derived from our experiments, which may not perfectly align with financial instruments other than U.S. stocks. This is due to the varying magnitudes of price fluctuations across different financial products, and the characteristics of each product can impact the model's predictive performance. Therefore, our future research will focus on two main directions. Firstly, we aim to extend the experiments to various financial instruments in different countries. The current plan is to expand into cryptocurrency and the Chinese stock market, with the intention of further broadening our scope to other markets. Secondly, we will work on the generalization of the filtration process, seeking a formula for defining threshold that is applicable to various products. Finding a threshold suitable for different commodities will help us more efficiently apply the strategy to a wide range of financial markets. Most importantly, we will continue to monitor the performance of this strategy in the U.S. stock market and consistently explore optimization opportunities.

In closing, NoxTrader stands as a testament to the potential within algorithmic trading models. Through this comprehensive exposition, it is evident that NoxTrader's journey is just beginning, with the horizon teeming with opportunities for innovation, refinement, and ever-greater achievements in the realm of model trading.

References

[1] Soni P, Tewari Y, Krishnan D. Machine Learning Approaches in Stock Price Prediction: A Systematic Review. J Phys Conf Ser. 2022;2161:012065.

[2] Kakushadze Z. 101 Formulaic Alphas. Wilmott. 2016;2016:72-81.

[3] Asness CS, Moskowitz TJ, Pedersen LH. Value and Momentum Everywhere. J Finance. 2013;68:929-985.

[4] Lee CM, Swaminathan B. Price Momentum and Trading Volume. J Finance. 2000;55:2017-2069.