

Algorithmic Trading and Short-term Forecast for Financial Time Series with Machine Learning Models; State of the Art and Perspectives

Dakota Joiner
Computer Science
Okanagan College
Kelowna, Canada
0000-0002-3094-0015

Amy Vezeau
Computer Science
Okanagan College
Kelowna, Canada
0000-0002-9390-0053

Albert Wong
Mathematics and Statistics
Langara College
Vancouver, Canada
0000-0002-0669-4352

Gaétan Hains
LACL
Université Paris-Est
Créteil, France
0000-0002-1687-8091

Youry Khmelevsky
Computer Science
Okanagan College
Kelowna, Canada
0000-0002-6837-3490

Abstract—Stock price prediction with machine learning is an oft-studied area where numerous unsolved problems still abound owing to the high complexity and volatility that technical-factors and sentiment-analysis models are trying to capture. Nearly all areas of machine learning (ML) have been tested as solutions to generate a truly accurate predictive model. The accuracy of most models hovers around 50%, highlighting the need for further increases in precision, data handling, forecasting, and ultimately prediction. This literature review aggregates and concludes the current state of the art (from 2018 onward) with specifically selected criteria to guide further research into algorithmic trading. The review targets academic papers on ML or deep learning (DL) with algorithmic trading or data sets used for algorithmic trading with minute to daily time scales. Systems that integrate and test sentiment and technical analysis are considered the best candidates for an eventual generalized trading algorithm that can be applied to any stock, future, or traded commodity. However, much work remains to be done in applying natural language processing and the choice of text sources to find the most effective mixture of sentiment and technical analysis. The best models being useless on themselves, we also search for publications about data warehousing systems aggregating financial factors impacting stock prices. A brief review in this area is included in this regard.

Index Terms—Machine learning, deep learning, algorithmic trading, data warehousing.

I. INTRODUCTION

A profitable algorithmic stock trading algorithm will benefit from a forecasting system that can produce accurate short-term forecasts. Based on this premise, we propose this research project to leverage our previous experience

We would like to acknowledge and thank Brenda Naka, Co-op Education, and Grants-in-Aid (GIA) Fund, Okanagan College, for supporting our research.

in building short-term forecasting models using machine learning (ML) algorithms [1]–[5].

As the first step in the research project, an in-depth review of existing literature on the development of forecasting models for price of financial assets using machine learning models was conducted and reported here. As a disclaimer, true comparability across different models is quite difficult to achieve when many projects use different measures to report their results. Given that there is no standard for comparison on different analytical methods without some degree of normalization, we focused our review on models that has a high degree of accuracy in predicting price as opposed to research that only mentions trading results and are therefore too dependent on the trading context. Drilling down to model-level specifics and the improvements made to best-of-class models are also goals for this review.

II. SCOPE OF THE REVIEW

This review includes recent (mostly from 2018 onward) academic papers in the use of ML or DL for financial asset price forecast. Two categories of research are considered: They must pertain to either price forecast algorithms or data sets used for developing price forecasts using ML. Algorithmic trading strategies aimed at real-time trading or high-frequency trading are not considered in this review.

As well, research that utilize either technical or sentiment analysis are included in the review. Technical analysis relates to the considerations of quantitative factors (trends, volume, etc.) included in the development of a price forecast. Sentiment analysis, on the other hand, is the quantification of the "sentiment", either positive or

negative, of the market participants. Also, special attention is given to research that has achieved significant results in producing accurate short term (for example daily) price forecasts. In the course of the review, a strong discriminating criterion has evolved to favour algorithms that were rated for its predictive power or its trading results.

Any research that focus on macro economic analysis of the financial market and without the use of ML applications are excluded in the review.

III. REVIEW OF THE EXISTING ALGORITHMIC TRADING MODELS FOR FINANCIAL TIME SERIES

There exists an incredibly large body of research papers on algorithmic trading including testing and results of various ML models using neural networks (NN), random forest (RF), support vector regression (SVR), XGBoost, and long short-term memory (LSTM) algorithms. This review is roughly organized into sections that specifically review LSTM and SVR models, as those seem to be the most heavily represented, and the other remaining models together.

Pawar [6] outlines the basic steps in setting up a machine learning model for testing stock price predictions. Other research groups have also carried out their own literature and historical reviews in recent years highlighting the current state at the time [7]–[10].

Attempts at machine learning models in financial time series forecasts are hardly a new topic area, with literature dating as far back as 2011 on how these models might be utilized to trade profitably [11]. Specifically, feature extraction and trading strategy design are discussed in detail as a basis for autonomous trading. Pricope [12] notes the preponderance of tested ML models that show little real-world applicability or have conditions too unrealistic to be useful. They remark that many of the reviewed models seem good in theory but not in practice. Many models, such as that seen in [13], attempt to match predictions to real data, but a considerable price and time discrepancy exist between the predicted values and actual values. The discrepancy could also be off by several hundred dollars (or any given unit of currency). A forecast model serves little purpose if the predicted value on a certain day only match the previous day's closing price for a given stock or if the model does not produce accurate prediction consistently. Research such as this shows the considerable difficulty in developing a ML model that produces accurate predictive results and could be used in the real world for the purpose of increasing trading profit. In the same vein, Cartea, Jaimungal, and Sánchez-Betancourt [14] employed reinforcement learning (RL) techniques to create statistical arbitrage strategies for electronic markets. The learning paradigm used was double deep Q learning network (DDQN), a new variant of reinforced deep Markov models (RDMM). Their simulated results shows the financial performance of the strategy but

leave room for future work on the use of the model with real market data. This is yet another example of a ML modelling and strategy that seems promising but is not proven in the real world.

Still another research provides enhancements to computational execution of trading algorithms without posting trading results [15]. This is a noteworthy area of pursuit, particularly if short-term models cannot produce results within that time frame, and offer no real-world testing. In work by Li and Peng [16], reinforcement learning concepts are applied to design simple trading strategies with quadruples on 11 currency pairs from these countries (EU, USA, JP, CND, AUS). The results are measured with profitability which is not informative in terms of stock-value prediction power.

Other research groups, like that in [17], tested many neural networks (NN) and ML models to see which is the best for a given data set. They used back propagation NN, radial basis function NN, general regression NN, SVR, and Least Squares-SVR. The testing data was weekly adjusted close price of three individual stocks: Bank of China, Vanke A, and Kweichou Moutai with mean square error (MSE) and mean absolute percentage error (MAPE) as criteria. Though the authors gave largely inconclusive results, it was noted that back propagation had the best results, at least among the tested models, with an MAPE under 5%.

In 2018, Chen and He [18] investigated the reliability of deep learning methods based on a 6-layer convolutional neural network (CNN) at predicting prices on the Chinese stock market. They set the time scale to be a year and input to the opening price, high price, low price, closing price, and the volume for historical stock data sets from the Chinese stock market. The results obtained showed an accuracy of about 73%. Results at this level begin to approach a usable state, but still require fine-tuning to be considered truly reliable and predictive.

Many works have experimented, with success, on stock price predictions with financial news. In [19], the authors examined the relationship between measured sentiment from some media messages about a company and its stock price. The authors used natural language processing (NLP) to perform sentiment analysis on messages from Twitter, called tweets. When a positive message appears, the algorithmic trading system opens a long position and negative messages, a short position. Once a position was opened, there were three options. If a new tweet appeared with the same positive sentiment, the duration of the position was extended by some 'X' amount of time. If an opposing sentiment tweet appeared, the position was closed, and if there were no further tweets, the position was closed after 'X' minutes. The authors found that positive results were gained from this method. The best time to leave the position open was for 60 minutes, anything longer and the performance deteriorated until it

hit losses at an average of 108 minutes. This research was performed with a shorter duration than what was originally planned and has a noticeably small sample size, but it still highlights the effect of sentiment on market performance as a determining factor to be leveraged for profit.

Similarly, Von Beschwitz, Keim, and, Massa [20] investigated how news analytics (as a subarea of sentiment analysis), not the information contained in the news, affected stock prices. It also investigated how news analytics were related to high frequency trading and low frequency trading. They used the RavenPack software to analyze how the Dow Jones Newswire categorizes article topics (such as acquisitions) or a listed company was downgraded by a rating agency. RavenPack also provided the composite sentiment score, the event sentiment score, and relevance. Key takeaways from the research reveal the ability of RavenPack to “see through” positive stock results in a news event to provide a truer analysis of stock performance.

In [21], the authors created a new approach that incorporated technical timeline data along with financial news. The resulting blended ensemble model is a combination of two recurrent neural nets followed by a fully connected neural net. When applied to price forecasting of stocks in the S&P 500, performance improvements between 30% and 50% on evaluation metrics like MSE, MDA, precision, recall and f1-score show that more complex data sources can help predict next-day stock prices without adding noise.

Weng, Lu, Lang, Megahed, and Martinez [22] sought to predict short-term (1 day or 1 to 10 days) stock value using a rich hybrid data set from many available online sources, including Google search engine results, published news, technical stock indicators, historical stock values and Wikipedia article information. The data sets were scraped using public APIs and the TTR (Technical Trading Rules) R package. Neural networks, SVR, boosted regression tree (BRT), and random forest models were tested individually and as sets of ensembles to determine the most effective combination. The results for single day trading on the Citigroup test case saw an $MAPE \leq 1.5\%$, while the 10-day test showed the best results with an $MAPE \leq 1.89\%$ utilizing BRT. These results are comparable to traditional non-ML methods and show considerable promise when compared to other models discussed so far. Among the reviewed models, the performance here is the best we have observed.

Nan, Perumalm, and Zaiane [23] formulated a reinforcement learning model that incorporates the sentiment analysis of news headlines pertaining to stocks and knowledge graphs. They found that sentiment analysis gave better results from the Sharpe ratio of 2.4 for MSFT, 2.2 for AMZN, and close to 2 for TSLA. It is important to note that the stock trading bot described was limited to buying or selling one stock in a day which likely limits the amount

of profit that could be made.

Ren, Wu, and Liu [24] developed a stock-exchange prediction technique that integrates sentiment analysis into an Support Vector Machine (SVM) algorithm, while taking into account the day of the week in the pre-processing of the data. As a proof of concept, they reported the accuracy of prediction reaches 90% with a rise of 18.6% after introducing sentiment data. These results are plausible, but pertain specifically to predicting the trend direction in stock prices, rather than the stock price themselves. Another approach in the prediction of stock values with public sentiment was by Batra and Daudpota [25]. They created sentiment scores through a SVM model using customer comments about Apple products on a social network to predict next day’s stock movement. The model was found to have an accuracy of 76%.

Khan et al. [26] ran an analysis on the need and applicability of social media data to complement a stock time series in the prediction of stock values. Specific stocks such as Red Hat and IBM were used to evaluate the sensitivity of their stock prices to real-time social media messages. The highest level of predictive accuracy using a RF model and with social media is found to be around 81% and 75% with financial news.

In a slight change of pace, an analysis and proof of feasibility of price prediction for cryptocurrencies [27] compared the utilization of RF, SVM, and NN with input from either purely sentiment data (from Twitter) or hybrid technical-sentiment data. Many studies concerning cryptocurrency focus solely on Bitcoin. Valencia et al. [27] have chosen to expand their models to include three other cryptocurrencies: Ethereum, Ripple, and Litecoin. The conclusion is that sentiment data could only predict with a level of accuracy the price of some of the cryptocurrency while hybrid data gave better predictive performance for all three at a maximum of 72% accuracy. A multilayer perceptron (MLP) neural net performed best overall when compared to SVR and RF, both of which performed worse than a random prediction. While no results showed significant performance improvement over random prediction, the result demonstrated that the inclusion of sentiment data from Twitter does increase accuracy in predicting market direction (though not price or volume.)

Mallqui and Fernandes [28] set out to demonstrate the efficacy of predicting Bitcoin prices using artificial NN, SVM, and Ensemble algorithms. The accuracy of the different models tested varied between 42% to 63% at the maximum. These results are approximately in line with many of the other models reviewed. The authors note that the incredible volatility of Bitcoin prices presents specific challenges in the prediction of price and the direction of price movement.

In [29], the authors explored a novel approach that combined both DL and RL to create a neural net for making decisions in algorithmic trading. The deep learning

perception model was trained on real data. Enhancements to RL models with DL improves performance for all baseline tests. Moreover, the novel model provided a 6% higher rate of return on investment over the best performing trading algorithm at the time. In another attempt at improving existing models through a combination of analytical techniques, Lei, Peng, and Shen add a residual network to the classical moving average convergence divergence (MACD) analysis to minimize the unpredictability inherent in MACD trading strategies [30]. Residual networks were used to further validate trading signals from MACD analysis of the Chinese stock market between 2017 and 2019. When compared to traditional models, MACD-KURT demonstrates a predictive trading accuracy of 70%.

A different type of neural net model is applied to Chinese stock-exchange CSI300 data with better prediction results than “non-hybrid” methods such as statistical models, SVM, RNN, or LSTM-NN. The approach is called multi-filters NN and combines convolution with recurrent layers and multiple input features (multi-variable inputs). The architecture uses a CNN and a RNN in parallel then batch normalization. These steps are then repeated twice before a fully-connected CNN and a “SoftMax” layer are applied. This model achieve a maximum 42% rate of return on investment in the simulated financial markets [31].

To bridge the gap between research and the real world, Mukerji et al. [32] created their own simulated asset market to analyze algorithmic trading (AT) on specific outcomes in the market. In their simulations, they gradually increase the market presence of AT from 0% to 100%. They found that 10% of AT trades is most beneficial with respect to market liquidity. Further increases in the AT percentage contributed little to liquidity. When the authors simulated a rise in trading pairs to see the relationship between price and the fundamental value of the stock, they found that price is correlated with trading pairs but has little correlation with fundamental value.

An alternative approach, coined Trading Deep Q-Network algorithm (TDQN) by Théate and Ernst [33], is the combination of DeepQ + Deep Reinforcement Learning to further the success of deep reinforcement learning (DRL) on determining the optimal trading positions. The training of the reinforcement learning (RL) agent is entirely based on the generation of artificial trajectories from a limited set of stock market historical data. The assessed performance of the TDQN exceeded the average benchmark trading strategies like buy and hold and mean reversion with moving averages. The average Sharpe ratio for TDQN was 0.404 compared to a range of -0.331 to 0.369 for the benchmark methods.

Orimoloye, Sung, Ma, et al. [34] created a comprehensive dataset from 32 countries with 34 financial indices over 6 years to compare deep feedforward neural nets (DNN) with shallow architectures like one-layer NN on their predictive capacity. At daily time scale, the deep

models are not better but they have significantly better predictive accuracy when using minute-frequency data. They also noted that the rectified linear unit (ReLU) activation function was consistently better than the tanh hidden layer activation function (TANH) for stock price prediction.

Park and Lee [35] demonstrated the effectiveness of offline/online state representation learning and imitative reinforcement learning (SIRL-Trader) in long term prediction. The SIRL-Trader had the best average return rate (57.8%) and Sharpe ratio (1.06) compared to the other state of the art baseline methods. In the ablation studies for the components of SIRL-Trader it was found that all components contribute to performance, but the offline SRL and gate structure were crucial to the performance obtained.

Ponomarev, Oseledets, and Cichocki utilize an asynchronous advantage actor-critic (A3C) model to test viability of the trading efficiency of reinforced learning models on stocks in the Russian market. They do not mention results in terms of accuracy or predictive power, but instead post annual rate of return of the model at 66% per year if commission is subtracted from any trading transactions. Interestingly, these authors seem to have the intention of selling the model if it is successful (as per the inclusion of commission on buy and sell trades), and as such have built cost considerations into the model that includes only predicted profitable trades above the commission price for the end user and with penalties for trading inaction. The trading environment is modelled as a Markov chain and viewed as a game with costs/rewards. Neural nets are used to produce its cost function and optimal traversal strategy [36].

Sawhney et al. [37] created a novel model called hyperbolic stock graph attention network (HyperStock-GAT) to demonstrate the effectiveness of using temporal hyperbolic graph learning on Riemannian manifolds. HyperStock-GAT was evaluated against a variety of baseline models on returns, Sharpe ratio, ranking, and normalized discounted cumulative gain. HyperStock-GAT performed 12% better than more standard methods, and 3% better than similar Euclidean graph-based methods.

The authors in this paper [38] introduces NARX models (non-linear autoregressive network with exogenous inputs) to produce trading orders (buy/sell/hold) from financial time series data. The research shows that NARX improves on SVM models. An ICA (independent component analysis) is used for pre-processing and denoising the data to avoid overfitting; It decomposes the process into independent components and extracts the “noise-like” ones. ICA is found to improve trading performance (comparing to no pre-processing) and to improve on principal component analysis (PCA) by about 20% in trading performance. In the experiments, mean squared error is replaced by trading performance as error measure. The paper introduces some

clear ideas, but it is difficult to compare its results because of its non-standard approach. It could be interesting to reproduce some of its results for further verification of performance against other models.

Kumar, Dogra, Utreja, and Yadav [39] carry out tests comparing predictability using SVM, RF, K-nearest neighbours, Naïve Bayes and softmax models. Random forests were found to perform best on large datasets (accuracy 72%), while Naïve Bayes was found to perform best on small datasets (accuracy 64%). The accuracy of all models improved with the number of technical indicators. The small number of stocks tested and the absence of a neural net model leaves open the question of how comparison of results was carried out.

In another application, prediction of stock prices tested on NYSE and S&P 500 data (2015-2017) and synthetic data using a novel feedforward neural net (FFNN) with forward conditional probability distributions was carried out by Levendovsky and team [40]. The restrictions in the training data and scope of stocks selected for portfolios (only 5 stocks with small movements in the mean) limits the potential effectiveness of this trading strategy. Very specific hypotheses are made about the trading strategies, and the results, though positive in profitability, are not easily compared and are only marginally greater than that of the buy and hold strategy.

So-called pairs-trading is extended to more than two stocks by Liang, Thavanasweran, Yu, Hoque, and Thulasiram utilizing a completely different strategy in comparison to those discussed so far [41]. An improved strategy is designed using dynamic “non-Gaussian” filters and is shown to improve on standard trading strategies that use Kalman filters. The advanced signal-processing techniques could be of general use, but the profit quality metric is not easily compared with prediction quality metrics used for evaluating ML algorithms. In the same year, virtually the same research team [42] compared classical trading strategies from the angle of their sensitivity to volatility. Most common algorithmic trading strategies follow technical indicators based on desirable trends. Therefore, the authors goal was to create a data driven way to estimate algorithmic volatility and the appropriate distribution of returns. To accomplish this task, they introduced fuzzy Bollinger bands, fuzzy SR estimates, and fuzzy SR forecasts. Identifying an appropriate distribution of strategy returns is found useful in trading. Both of these test cases utilize the Sharpe ratio for means of reporting profitability, which in recent times has been questioned in terms of its usefulness in favour of other ratios [43], though the results are still promising.

Nabipour et al. [44] analyzed ten years of historical data from the Tehran Stock Exchange to predict stock market trends employing ten machine learning models: decision tree, RF, Adaptive Boosting (Adaboost), eXtreme gradient boosting (XGBoost), support vector classifier

(SVC), Naïve Bayes, K-Nearest neighbors (KNN), logistic regression, artificial neural network (ANN), recurrent neural network (RNN), and LSTM. This data was employed in two ways, as continuous data from calculating the indicators by stock trading values and converting indicators into binary data. The research found that for both continuous and binary data, RNN and LSTM outperformed other algorithms. However, when binary data is evaluated, the difference between models is lessened. This team is able to post maximum accuracy in the range of 86%, a number higher than most have been able to reach at the time of this writing. In support of these results, Vijh, Chandola, Tikkiwal, and Kumar [45] note that ANN performs better than RF in their tests.

A. Review of LSTM Models

In this section, all papers reviewed utilize variations of a LSTM algorithm as the modelling technique and, therefore, descriptions of the models will be omitted for the sake of brevity unless there are particularly interesting features of note that distinguish the model or its performance.

Many of these models focus on predicting a single value of a given stock, but Ding and Qin generate a model capable of predicting the opening, lowest, and highest price of a stock for the next day with a reported accuracy > 95% on three separate Chinese stock indices [46]. These results seem incredible compared to the performance of many other models reviewed and warrant a closer look.

Chalvatzis and Hristu-Varsakelis develop an LSTM model trained to accomplish more than minimizing the MAPE with sub-1% MAPE. They also implement the model on a simulated market of the S&P 500, DJIA, NASDAQ, and Russell 2000 stock indices and realize short-term hypothetical profits in excess of 300% [47]. The Eapen, Verma, and Bein research group was able to show the superiority of bidirectional LSTM over SVR to the tune of a 9% increase in predictive accuracy [48]. Utilizing deep NN as an extension to LSTM modelling, Yu and Yan [49] were able to achieve an average predictive accuracy of 58% with maximum accuracy of 63% on the S&P 500 using a method focusing on breakdown, resolution, and recombination of noisy, time-dependent data to a workable state.

Attempting to tackle the problem of developing a predictive model with integrated hybrid technical and sentiment analysis, Li, Wu, and Wang [50] use NLP for sentiment analysis and an LSTM model for technical analysis in the prediction of Hong Kong stock prices. The hybrid technique is found to outperform baseline models (SVR and multiple kernel learning) as well as those that use either technical or sentiment data as input. However, the overall accuracy remains low with most results coming at sub 50%. They also demonstrate the separate effects that the individual sentiment data sources have on predictive power: domain-specific news dictionaries produce the

greatest improvements in accuracy at over 80%. As well, multiple information sources with respect to price data produce better results over any single information source.

Chan et al. [51] develop a new module, called multicollinearity reduction module (MRM), to improve the performance of LSTM models in predictive power. This module mitigates the effects of multicollinearity between the factors included in an LSTM model. When compared to the standalone LSTM model, the addition of the MRM produced 47% higher returns without sacrificing training time, though the actual prediction accuracy remains below 50% for both the LSTM and the LSTM+MRM. The development of this module shows promise as an addition to any LSTM model. Li, Zheng, and Zheng [52] also propose an addition to the LSTM model with the combination of asynchronous advantage actor-critic and stacked denoising autoencoders. With these additions, the trading algorithms employed see a marked increase in return over plain LSTM, though as with many other models, the price predictive accuracy hovers around 50%. On a similar note, Chen, Chen, and Fushimi [53] were able to demonstrate a simple, baseline improvement using LSTM over both locally weighted regression and buy and hold methods with positive returns and high directional accuracy in periods of high volatility in market price.

Over several years and in a number of research papers, Mehtab and Sen [54]–[56] were able to use several variations on LSTM models to predict opening and closing prices for NIFTY 50 on the National Stock Exchange of India (NSE) using the previous one week of opening and closing prices as input data. Additionally, they were able to successfully employ NLP for sentiment analysis from Twitter alongside a self-organizing fuzzy neural network (SOFNN) for prediction of closing prices using the same one-week-at-a-time input approach. Each of these three models excelled in predicting trending directions to a high degree of accuracy.

While not bringing in a particularly novel application of LSTM and RNN, the 2019 work of Mohan, Mullapudi, Sammeta, Vijayvergia, and Anastasiu [57] should be noted for the sheer volume of data processed. Over 265,000 news articles were collected and analyzed for sentiment analysis alongside 5 years worth of pricing data on the S&P 500. The introduction of a metric dubbed ‘polarity’ increased the accuracy of price prediction. High volatility was shown to negatively affect the predictive power when price and sentiment data were taken together. No clear ‘best’ model can be determined from their results, only that each tested model outperforms older benchmarks.

Qiu and team [58] introduce a wavelet transform, not dissimilar to Fourier transforms, to their LSTM model to denoise the volatile data set and separate useful data signal from the noise with the goal of overcoming long-term dependence issues. The performance of the WLSTM+Attention model was excellent in terms of error

(MSE, MAE, and RMSE) compared to LSTM along with a gated recurrent unit (GRU). With the addition of sentiment analysis, these techniques are likely to perform very well in stock price prediction. Similar to Qiu, the authors Li, Shen, and Zu [59] focus on reducing the noise of volatile stock prices in their LSTM model. Rather than applying a wavelet transform, they incorporate extra input gates in a method dubbed multi-input LSTM (MI-LSTM). This model marginally outperforms LSTM alone in terms of MAE as well as in profit comparison. From these two similar research areas, it is probable that attention mechanisms and noise reduction in signal input are likely to result in increases in prediction accuracy and increased profit.

B. Review of SVR Models

In this section, brief comments on the literature selected that use support vector regression (SVR) are provided.

In SVR, research parameter optimization is a critical challenge as the selection of the wrong kernel results in poor forecasting performance. Zhang, Teng, and Chen [60] produced improved performance by combining SVR with a modified firefly algorithm (MFA). The MFA is developed and employed to optimize the SVR parameters. The resulting performance is the best with an average MAPE of 0.0023%. The forecasted next day prices are very close to actual prices, but that is to be expected for such a short prediction window. It is reasonable to assume next day prices can be predicted with higher accuracy than longer time scales. Henrique, Sobreiro, Kimura [61] applied SVR to the prediction of stock-market values at mid- to low-frequency time scales and found similar results. The predictive accuracy improves during low volatility periods when a dynamic model-updating technique was applied: the predicted values track close to the real values resulting in smaller RMSE and an improved prediction. Additionally, it was observed that increasing the prediction time scale to the minute actually reduced predictive accuracy, to the extent of being inferior to that of a random prediction.

Song, Zhou, and Han [62] test the predictive power of five neural networks including SVR and least squares support vector machine regression (LS-SVMR) to see which provides the best price predictions for three individual stocks. Back propagation was found to have the best results whereas SVR was found to be the second best model. Back propagation displayed MAPE values between 0.019-0.049. Results posted by Song show similar day-shift issues as mentioned before: if all values are predicted well but shifted to be off by a day, the “weight” of the results displayed is diminished greatly.

Nikou et al. [63] seek to evaluate the price predictive power of ML models for iShares MSCI United Kingdom exchange-traded fund over a three and a half year period. Testing five different models (ANN, SVR, RF, DL, and

LSTM), the results of the study show that LSTM was the most accurate in forecasting the close price, with SVR following closely second. It was noted that subjecting the different models to deep learning lessened the error in predicting price in all methods, with SVR performing the best after the DL treatment.

IV. REVIEW OF DATA SETS, DATA PREPROCESSING, AND STORAGE

We briefly review a paper related to machine learning that instead focuses on storage and processing of stock trading data from a data warehouse or warehouse-like structures such as data lakes or other cloud-hosted data storage services. Relatively little research has been published in this topic and as such, we have identified a gap in the literature where our previous work with integrating data warehousing and ML may lead to fruitful discovery concerning the prediction of stock prices [4].

Maji, Mondal, Goto, Debnath, and Sen [64] identify a warehousing model and analytical framework for analyzing the parameters that impact the stock market. Several factors intrinsic to the stock market are stored alongside external factors such as the price of oil or gold. They implement an OLAP data warehouse with a snowflake schema to perform data summary, consolidation, and aggregation. No conclusions have been drawn about the results of integrating internal and external factors within a single data warehouse, but rather leave an opening for further research and discussion on statistical measures that can be applied to the stored data.

V. FUTURE WORK

We have three future goals pertaining to this research: build a data warehouse with targeted financial time series, build a novel short-term forecasting model with machine learning, and develop an automated evaluation system to assess effectiveness of short-term trading strategies. Compute Canada, an advanced research computing (ARC) organization with high-powered computing capabilities, has graciously provided a one year grant for access to vast computing resources necessary for complex machine learning projects. When paired with the existing Computer Science Department's Data Centre at Okanagan College [65]–[72], we have access to ample computing resources to achieve all three of our outlined research goals. With supporting computing power, we can tackle the wide breadth of raw natural language data and raw financial data for the non trivial task of NLP and the demands of careful data extraction, cleaning, integration, and model selection.

VI. CONCLUSIONS

In this paper, we presented a review on the current state of research into how ML and DL are being leveraged in algorithmic trading for their predictive powers. Long short-term memory models with neural networks seemed to be

the best performers among the different models reviewed. Interestingly, the models that integrated sentiment and technical analysis were more capable to accurately predict stock price than those that used only sentiment or technical analysis alone. To a human, it is easy to predict that analysis of say, Elon Musk's tweets is likely to have an impact on stock prices for companies with which he is involved. However, a model does not have the same intuitive power and instead relies of training and actual data accrued over time to be able to make a similar prediction. Nevertheless, previous works do not point to a unique or maximal set of variables that models should include, so we will closely investigate this selection process and its effect on stock value forecasting, the most portable of quality measures.

REFERENCES

- [1] G. Hains, C. Mazur, J. Ayers, J. Humphrey, Y. Khmelevsky, and T. Sutherland, "The WFast's Gamers Private Network (GPN®) Performance Evaluation Results," in *2020 IEEE International Systems Conference (SysCon)*. IEEE, 2020, pp. 1–6.
- [2] C. Mazur, J. Ayers, J. Humphrey, G. Hains, and Y. Khmelevsky, "Machine Learning Prediction of Gamer's Private Networks (GPN®S)," in *Proceedings of the Future Technologies Conference*. Springer, 2020, pp. 107–123.
- [3] A. Wong, C. Chiu, G. Hains, J. Humphrey, Y. Khmelevsky, C. Mazur, and H. Fuhrmann, "Gamers Private Network Performance Forecasting - From Raw Data to the Data Warehouse with Machine Learning and Neural Nets," 2021. [Online]. Available: <https://arxiv.org/abs/2107.00998>
- [4] A. Wong, D. Joiner, C. Chiu, M. Elsayed, K. Pereira, Y. Khmelevsky, and J. Mahony, "A Survey of Natural Language Processing Implementation for Data Query Systems," in *2021 IEEE International Conference on Recent Advances in Systems Science and Engineering (RASSE)*, 2021, pp. 1–8.
- [5] A. Wong, C. Chiu, A. Abdulgapul, M. N. Beg, Y. Khmelevsky, and J. Mahony, "Estimation of Hourly Utility Usage Using Machine Learning," in *SysCon 2022 (accepted for the publication)*, 2022.
- [6] K. Pahwa and N. Agarwal, "Stock Market Analysis using Supervised Machine Learning," in *2019 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COMITCon)*, 2 2019, pp. 197–200.
- [7] M. A. Shetty, "A Study On Algorithmic Trading And Automated Trading System — Uses And Challenges." *Jamshedpur Research Review*, vol. 1, no. 44, pp. 76–81, 2021.
- [8] P. Gomber and K. Zimmermann, "Algorithmic trading in practice," in *The Oxford Handbook of Computational Economics and Finance*, 2 2018, pp. 311–332.
- [9] C. PÄUNA, "Automated Trading Software - Design and Integration in Business Intelligence Systems," *Database Systems Journal*, vol. IX, pp. 22–28, 2018.
- [10] B. Huang, Y. Huan, L. D. Xu, L. Zheng, and Z. Zou, "Automated trading systems statistical and machine learning methods and hardware implementation: a survey," pp. 132–144, 2019.
- [11] G. Nuti, M. Mirghaemi, P. Treleaven, and C. Yingsaeree, "Algorithmic trading," *Computer*, vol. 44, no. 11, pp. 61–69, 2011. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8786132>
- [12] T.-V. Pricope, "Deep Reinforcement Learning in Quantitative Algorithmic Trading: A Review," *arXiv preprint*, vol. abs/2106.0, 2021. [Online]. Available: <https://arxiv.org/abs/2106.00123>
- [13] I. Parmar, N. Agarwal, S. Saxena, R. Arora, S. Gupta, H. Dhi-man, and L. Chouhan, "Stock Market Prediction Using Machine Learning," in *2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC)*, 12 2018, pp. 574–576.
- [14] A. Cartea, S. Jaimungal, and L. Sanchez-Betancourt, "Deep Reinforcement Learning for Algorithmic Trading," *SSRN Electronic Journal*, 3 2021. [Online]. Available: <https://papers.ssrn.com/abstract=3812473>

- [15] P. Rys, R. Ślepaczuk, and others, "Machine learning in algorithmic trading strategy optimization-implementation and efficiency," 2018.
- [16] X. Li and Z. Peng, "A novel algorithmic trading approach based on reinforcement learning," in *Proceedings - 2019 11th International Conference on Measuring Technology and Mechatronics Automation, ICMTMA 2019*. Institute of Electrical and Electronics Engineers Inc., 4 2019, pp. 394–398.
- [17] Y.-G. Song, Y.-L. Zhou, and R.-J. Han, "Neural networks for stock price prediction," 2018. [Online]. Available: <https://arxiv.org/abs/1805.11317>
- [18] S. Chen and H. He, "Stock Prediction Using Convolutional Neural Network," *IOP Conference series: materials science and engineering*, vol. 435, p. 012026, 11 2018. [Online]. Available: <https://doi.org/10.1088/1757-899x/435/1/012026>
- [19] R. Gómez-Martínez, C. Prado-Román, M. de la Cruz, and others, "ALGORITHMIC TRADING SYSTEM ON NASDAQ TWEETS," *Review of General Management*, vol. 29, no. 1, 2019.
- [20] B. Von Beschwitz, D. B. Keim, and M. Massa, "First to "Read" the News: News Analytics and Algorithmic Trading," *The Review of Asset Pricing Studies*, vol. 10, no. 1, pp. 122–178, 2020. [Online]. Available: <https://academic.oup.com/raps/article/10/1/122/5555424>
- [21] Y. Li and Y. Pan, "A novel ensemble deep learning model for stock prediction based on stock prices and news," *International Journal of Data Science and Analytics*, vol. 13, pp. 139–149, 2022. [Online]. Available: <https://doi.org/10.1007/s41060-021-00279-9>
- [22] B. Weng, L. Lu, X. Wang, F. M. Megahed, and W. Martinez, "Predicting short-term stock prices using ensemble methods and online data sources," *Expert Systems with Applications*, vol. 112, pp. 258–273, 12 2018.
- [23] A. Nan, A. Perumal, and O. R. Zaiane, "Sentiment and Knowledge Based Algorithmic Trading with Deep Reinforcement Learning," *arXiv preprint*, vol. abs/2001.0, 1 2020. [Online]. Available: <https://arxiv.org/abs/2001.09403v1>
- [24] R. Ren, D. D. Wu, and T. Liu, "Forecasting Stock Market Movement Direction Using Sentiment Analysis and Support Vector Machine," *IEEE Systems Journal*, vol. 13, no. 1, pp. 760–770, 3 2019.
- [25] R. Batra and S. M. Daudpota, "Integrating StockTwits with sentiment analysis for better prediction of stock price movement," in *2018 International Conference on Computing, Mathematics and Engineering Technologies (iCoMET)*, 3 2018, pp. 1–5.
- [26] W. Khan, M. A. Ghazanfar, M. A. Azam, A. Karami, K. H. Alyoubi, and A. S. Alfakeeh, "Stock market prediction using machine learning classifiers and social media, news," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–24, 3 2020. [Online]. Available: <https://link.springer.com/article/10.1007/s12652-020-01839-w>
- [27] F. Valencia, A. Gómez-Espinosa, and B. Valdés-Aguirre, "Price Movement Prediction of Cryptocurrencies Using Sentiment Analysis and Machine Learning," *Entropy*, vol. 21, no. 6, 2019. [Online]. Available: <https://www.mdpi.com/1099-4300/21/6/589>
- [28] D. C. Mallqui and R. A. Fernandes, "Predicting the direction, maximum, minimum and closing prices of daily Bitcoin exchange rate using machine learning techniques," *Applied Soft Computing*, vol. 75, pp. 596–606, 2 2019.
- [29] K. Lei, B. Zhang, Y. Li, M. Yang, and Y. Shen, "Time-driven feature-aware jointly deep reinforcement learning for financial signal representation and algorithmic trading," *Expert Systems with Applications*, vol. 140, p. 112872, 2 2020.
- [30] Y. Lei, Q. Peng, and Y. Shen, "Deep Learning for Algorithmic Trading: Enhancing MACD Strategy," in *Proceedings of the 2020 6th International Conference on Computing and Artificial Intelligence*, ser. ICCAI '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 51–57. [Online]. Available: <https://doi.org/10.1145/3404555.3404604>
- [31] W. Long, Z. Lu, and L. Cui, "Deep learning-based feature engineering for stock price movement prediction," *Knowledge-Based Systems*, vol. 164, pp. 163–173, 1 2019.
- [32] P. Mukerji, C. Chung, T. Walsh, and B. Xiong, "The Impact of Algorithmic Trading in a Simulated Asset Market," *Journal of Risk and Financial Management*, vol. 12, no. 2, p. 68, 4 2019.
- [33] T. Théate and D. Ernst, "An application of deep reinforcement learning to algorithmic trading," *Expert Systems with Applications*, vol. 173, p. 114632, 7 2021.
- [34] L. O. Orimoloye, M. C. Sung, T. Ma, and J. E. Johnson, "Comparing the effectiveness of deep feedforward neural networks and shallow architectures for predicting stock price indices," *Expert Systems with Applications*, vol. 139, p. 112828, 1 2020.
- [35] D. Y. Park and K. H. Lee, "Practical Algorithmic Trading Using State Representation Learning and Imitative Reinforcement Learning," *IEEE Access*, vol. 9, pp. 152 310–152 321, 2021.
- [36] E. S. Ponomarev, I. V. Oseledets, and A. S. Cichocki, "Using Reinforcement Learning in the Algorithmic Trading Problem," *Journal of Communications Technology and Electronics*, vol. 64, no. 12, pp. 1450–1457, 2 2019. [Online]. Available: <https://link.springer.com/article/10.1134/S1064226919120131>
- [37] R. Sawhney, S. Agarwal, A. Wadhwa, and R. Shah, "Exploring the scale-free nature of stock markets: Hyperbolic graph learning for algorithmic trading," in *The Web Conference 2021 - Proceedings of the World Wide Web Conference, WWW 2021*. Association for Computing Machinery, Inc, 4 2021, pp. 11–22.
- [38] A. Ceffer, J. Levendovszky, and N. Fogarasi, "Applying Independent Component Analysis and Predictive Systems for Algorithmic Trading," *Computational Economics*, vol. 54, pp. 281–303, 7 2019. [Online]. Available: <https://link.springer.com/article/10.1007/s10614-017-9719-z>
- [39] I. Kumar, K. Dogra, C. Utreja, and P. Yadav, "A Comparative Study of Supervised Machine Learning Algorithms for Stock Market Trend Prediction," in *2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*, 4 2018, pp. 1003–1007.
- [40] J. Levendovszky, I. Reguly, A. Olah, and A. Ceffer, "Low Complexity Algorithmic Trading by Feedforward Neural Networks," *Computational Economics*, vol. 54, pp. 267–279, 7 2019. [Online]. Available: <https://link.springer.com/article/10.1007/s10614-017-9720-6>
- [41] Y. Liang, A. Thavaneswaran, N. Yu, M. E. Hoque, and R. K. Thulasiram, "Dynamic Data Science Applications in Optimal Profit Algorithmic Trading," in *Proceedings - 2020 IEEE 44th Annual Computers, Software, and Applications Conference, COMPSAC 2020*. Institute of Electrical and Electronics Engineers Inc., 7 2020, pp. 1314–1319.
- [42] A. Thavaneswaran, Y. Liang, Z. Zhu, and R. K. Thulasiram, "Novel data-driven fuzzy algorithmic volatility forecasting models with applications to algorithmic trading," in *IEEE International Conference on Fuzzy Systems*, vol. 2020-July. Institute of Electrical and Electronics Engineers Inc., 7 2020, pp. 1–8.
- [43] S. Farinelli, M. Ferreira, D. Rossello, M. Thoeny, and L. Tibiletti, "Beyond Sharpe ratio: Optimal asset allocation using different performance ratios," *Journal of Banking & Finance*, vol. 32, no. 10, pp. 2057–2063, 2008.
- [44] M. Nabipour, P. Nayyeri, H. Jabani, S. Member, and A. Mosavi, "Predicting Stock Market Trends Using Machine Learning and Deep Learning Algorithms Via Continuous and Binary Data; a Comparative Analysis," *IEEE Access*, vol. 8, pp. 150 199–150 212, 2020. [Online]. Available: www.tsetmc.com
- [45] M. Vijh, D. Chandola, V. A. Tikkiwal, and A. Kumar, "Stock Closing Price Prediction using Machine Learning Techniques," *Procedia Computer Science*, vol. 167, pp. 599–606, 1 2020.
- [46] G. Ding and L. Qin, "Study on the prediction of stock price based on the associated network model of LSTM," *International Journal of Machine Learning and Cybernetics*, vol. 11, no. 6, pp. 1307–1317, 2020. [Online]. Available: <https://doi.org/10.1007/s13042-019-01041-1>
- [47] C. Chalvatzis and D. Hristu-Varsakelis, "High-performance stock index trading via neural networks and trees," *Applied Soft Computing*, vol. 96, p. 106567, 11 2020.
- [48] J. Eapen, D. Bein, and A. Verma, "Novel Deep Learning Model with CNN and Bi-Directional LSTM for Improved Stock Market Index Prediction," in *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*, 1 2019, pp. 264–270.
- [49] P. Yu and X. Yan, "Stock price prediction based on deep neural networks," *Neural Computing and Applications 2019* 32:6, vol. 32, pp. 1609–1628, 4 2020. [Online]. Available: <https://link.springer.com/article/10.1007/s00521-019-04212-x>

- [50] X. Li, P. Wu, and W. Wang, "Incorporating stock prices and news sentiments for stock market prediction: A case of Hong Kong," *Information Processing & Management*, vol. 57, no. 5, p. 102212, 9 2020.
- [51] J. Yi-Le Chan, S. Mun Hong Leow, K. Thye Bea, W. Khuen Cheng, S. Wai Phoong, Z.-W. Hong, J.-M. Lin, and Y.-L. Chen, "A Correlation-Embedded Attention Module to Mitigate Multicollinearity: An Algorithmic Trading Application," *Mathematics*, vol. 10, no. 8, 2022. [Online]. Available: <https://doi.org/10.3390/math10081231>
- [52] Y. Li, W. Zheng, and Z. Zheng, "Deep Robust Reinforcement Learning for Practical Algorithmic Trading," *IEEE Access*, vol. 7, pp. 108 014–108 022, 8 2019.
- [53] G. Chen, Y. Chen, and T. Fushimi, "Application of deep learning to algorithmic trading," Stanford University, Tech. Rep., Tech. Rep., 2017. [Online]. Available: <http://...>, Tech. Rep., 2021.
- [54] S. Mehtab and J. Sen, "A Robust Predictive Model for Stock Price Prediction Using Deep Learning and Natural Language Processing," *SSRN Electronic Journal*, 12 2019. [Online]. Available: <https://papers.ssrn.com/abstract=3502624>
- [55] S. Mehtab, J. Sen, and A. Dutta, "Stock Price Prediction Using Machine Learning and LSTM-Based Deep Learning Models," in *Machine Learning and Metaheuristics Algorithms, and Applications*. Singapore: Springer Singapore, 2021, pp. 88–106.
- [56] S. Mehtab and J. Sen, "Stock Price Prediction Using Convolutional Neural Networks on a Multivariate Timeseries," *arXiv preprint*, vol. abs/2001.0, 1 2020. [Online]. Available: <http://arxiv.org/abs/2001.09769> <http://dx.doi.org/10.36227/techrxiv.15088734.v1>
- [57] S. Mohan, S. Mullanpudi, S. Sammeta, P. Vijayvergia, and D. C. Anastasiu, "Stock Price Prediction Using News Sentiment Analysis," in *2019 IEEE Fifth International Conference on Big Data Computing Service and Applications (BigDataService)*, 4 2019, pp. 205–208.
- [58] J. Qiu, B. Wang, and C. Zhou, "Forecasting stock prices with long-short term memory neural network based on attention mechanism," *PLoS ONE*, vol. 15, no. 1, 2020. [Online]. Available: <https://doi.org/10.1371/journal.pone.0227222.g001>
- [59] H. Li, Y. Shen, and Y. Zhu, "Stock Price Prediction Using Attention-based Multi-Input LSTM," in *Proceedings of The 10th Asian Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Zhu and I. Takeuchi, Eds., vol. 95. PMLR, 4 2018, pp. 454–469. [Online]. Available: <https://proceedings.mlr.press/v95/li18c.html>
- [60] J. Zhang, Y.-F. Teng, and W. Chen, "Support vector regression with modified firefly algorithm for stock price forecasting," *Applied Intelligence*, vol. 49, no. 5, pp. 1658–1674, 2019. [Online]. Available: <https://doi.org/10.1007/s10489-018-1351-7>
- [61] B. M. Henrique, V. A. Sobreiro, and H. Kimura, "Stock price prediction using support vector regression on daily and up to the minute prices," *Journal of Finance and Data Science*, vol. 4, no. 3, pp. 183–201, 9 2018.
- [62] G. Duan, M. Lin, H. Wang, and Z. Xu, "Deep Neural Networks for Stock Price Prediction," in *arXiv preprint*. arXiv, 2022, pp. 65–68. [Online]. Available: <https://arxiv.org/abs/1805.11317>
- [63] M. Nikou, G. Mansourfar, and J. Bagherzadeh, "Stock price prediction using DEEP learning algorithm and its comparison with machine learning algorithms," *Intelligent Systems in Accounting, Finance and Management*, vol. 26, no. 4, pp. 164–174, 2019.
- [64] G. Maji, D. Mondal, T. Goto, N. C. Debnath, and S. Sen, "A Data Warehouse Based Modelling Technique for Stock Market Analysis," *International Journal of Engineering & Technology*, vol. 7, no. 3, pp. 165–170, 2018. [Online]. Available: www.sciencepubco.com/index.php/IJET
- [65] Y. Khmelevsky, "SW Development Projects in Academia," *WCCCE 2009 - Proceedings of the 14th Western Canadian Conference on Computing Education*, vol. 1, no. 250, pp. 60–64, 2009. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1536292>
- [66] Y. Khmelevsky and V. Voytenko, "Cloud computing infrastructure prototype for university education and research," in *Computing*. ACM Press, 2010, pp. 1–5. [Online]. Available: <https://doi.org/10.1145/1806512.1806524>
- [67] Y. Khmelevsky, V. Ustimenko, G. Hains, C. Kluka, E. Ozan, and D. Syrotovsky, "International collaboration in SW engineering research projects," in *Proceedings of the 16th Western Canadian Conference on Computing Education - WCCCE '11*, 2011.
- [68] Y. Khmelevsky, "Research and Teaching Strategies Integration at Post-secondary Programs," in *Proceedings of the 16th Western Canadian Conference on Computing Education*, ser. WCCCE '11. New York, NY, USA: ACM, 2011, pp. 57–60. [Online]. Available: <http://doi.acm.org/10.1145/1989622.1989638>
- [69] Y. Khmelevsky, G. G. Hains, and C. Li, "Automatic Code Generation Within Student's Software Engineering Projects," in *WCCCE '12*, 2012, pp. 29–33.
- [70] Y. Khmelevsky and V. Voytenko, "Hybrid Cloud Computing Infrastructure in Academia," in *WCCCE 2015 - the 20th Western Canadian Conference on Computing Education, At May 8-9, 2015*. Vancouver Island University (VIU), Nanaimo, British Columbia, Canada., 2015.
- [71] Z. Ouimet, H. Caswell, Y. Khmelevsky, R. Bartlett, and A. Needham, "Game servers deployment automation case study," in *2016 Annual IEEE Systems Conference (SysCon)*, 2016, pp. 1–7.
- [72] Y. Khmelevsky, "Ten Years of Capstone Projects at Okanagan College: A Retrospective Analysis," in *Proceedings of the 21st Western Canadian Conference on Computing Education*. New York, NY, USA: ACM, 2016, pp. 7:1–7:6. [Online]. Available: <http://doi.acm.org/10.1145/2910925.2910949>