To appear in *Quantitative Finance*, Vol. 00, No. 00, Month 20XX, 1–6

Machine Learning Perspectives in Finance

Germán G. Creamer*†‡, Gary Kazantsev±, and Tomaso Aste∓

†School of Business, Stevens Institute of Technology, NJ, USA
‡Visiting Scholar at Stern School of Business, New York University, NY, USA
±Bloomberg, NY, USA
∓University College London, London, UK

(*June 2019*)

1. Introduction

During the last decade we have witnessed a rapid expansion of artificial intelligence (AI) applications and machine learning (ML) algorithms in an increasingly broad range of problems in finance. This development is fueled by a unique confluence of factors: an exponentially growing computational capacity that is available for enterprises, and similarly exponential growth in the amount of machine-readable data, along with improvements in the state of the art which allow ML and AI applications that were impractical ten or twenty years ago. There is an ambitious feeling emerging across industry and academia that some cognitive processes can be automated via ML and AI, radically expanding automation of the services industry and finance in particular.

However, finance is very different from other domains, such as image recognition, where ML an AI have been developed and successfully deployed. ML and AI need training sets and a lot of reliable and consistent data to make machines learn their models. But all financial data are not alike. We can characterize financial information as big data because of its large volume (financial time series data easily scales into petabytes), velocity (much of financial data is high-frequency), and variety (numerical, categorical, text, images, etc.). This data, exhibits complex behavior: nonstationarity, nonlinear interactions, heteroscedasticity, and biases. The research goal in this domain is to find in this data relevant patterns that could be used for investment, risk management or trading decisions. Time series analysis and traditional statistics can facilitate the process of understanding, modeling, and forecasting the behavior of financial assets. Present day developments of AI and ML algorithms provide novel approaches and perspectives such as feature selection in high dimensional data that mixes large structured and unstructured datasets, and incorporates a large number of linear and nonlinear features. Some of them are reported in this special issue.

Financial systems can provide an excellent opportunity for AI and ML researchers to develop and test new algorithms due to the availability of datasets and the persistent demand for novel solutions to complex new questions. Many financial problems do not require a forecast of a continuous variable, as they can be formulated as classification problems. For Instance, a trader may only need to anticipate the market direction within a certain probable range rather than with a direct forecast of returns. This perspective influenced many researchers during the 1990s and early 2000s, leading them to apply ML algorithms to problems of algorithmic trading and market microstructure. Among the most common algorithms used were neural networks (Trippi and DeSieno 1992, Choi *et al.* 1995), traditional convex models such as support vector machines (Tay and

*Corresponding author. Email: gcreamer@stevens.edu

Cao 2001), reinforcement learning (Nevmyvaka *et al.* 2006, Moody and Saffell 2001, Dempster and Leemans 2006, Bates *et al.* 2003), boosting (Creamer and Freund 2010) and genetic algorithms (Dempster *et al.* 2001, Allen and Karjalainen 1999). The main areas of application for many of these algorithms had been the equity markets, futures, and foreign exchange (Chaboud *et al.* 2014). It is difficult to estimate the true extent of the application of these methods in the industry, as publication of results has been sparse and limited in scope. We hope that this special issue will be a reflection of a new trend that is, an expanding and diversifying collaboration between academic and industry researchers. This kind of cooperation cannot help but move the industry forward by openly communicating current applications and enabling dialogue.

Recently, the range of AI and ML applications in finance has been extended to other markets. Some of the papers in derivative pricing using ML algorithms are oriented to the replication of the Black-Scholes (BS) model for pricing derivative instruments (Hutchinson *et al.* 1994, Amilon 2003). They fill a critical gap in the literature as they help to price derivatives using a data-driven approach without satisfying all the assumptions of the BS model. There are still many unsolved questions, such as finding an optimal model for the volatility smile. Other new developments by AI researchers include developing artificial agents for automated investment recommendation and portfolio optimization systems (Decker *et al.* 1996, Seo *et al.* 2004, Creamer 2015). These models have enabled the current generation of robo-advisors (Baker and Dellaert 2018) that recommend portfolio allocations based on the clients characteristics and market trends such as "Warren" from Kensho, an AI startup acquired by Standard and Poor's. In parallel, many large organizations have also explored the use of AI and ML methods for customer management, primarily to target new customers, anticipate their consumption patterns, define customer segmentation, develop recommendation systems for products, and reduce the attrition rate, thereby offering new incentives and products to clients with a high propensity to leave. This line of research has led to the highly competitive and profitable area of digital marketing (Domingos and Richardson 2001, Richardson and Domingos 2006, Leskovec *et al.* 2006, Hill *et al.* 2006).

Risk management is another major area of application for ML methods in finance. The early models were mostly oriented to develop credit score systems using decision trees or other interpretable models that segmented the customers according to their risk level and demographic characteristics. These models were developed in-house and complemented the well-known Fair Isaac Corporation (FICO) score or other risk indicators provided by the large risk information providers. The application to other areas of risk in financial institutions was limited because of the restrictions imposed by the regulators and the methods that were acceptable by the early Basel accords. After the credit crisis of 2008 and thanks to the flexibility of Basel III, many new methods combining ML and social network analysis have been explored to evaluate systemic risk (Angelini *et al.* 2008, Koyuncugil and Ozgulbas 2012, Bao and Datta 2014, Birch and Aste 2014, O'Halloran *et al.* 2015, Fraiberger 2016, Manela and Moreira 2017, Van Liebergen 2017, Hanley and Hoberg 2019, Mamaysky and Glasserman 2019) as well as alternative calculations of Value-at-Risk and its components: credit, market and operational risk. In the last few years, mainstream econometricians and finance researchers have also incorporated the use of ML methods to solve problems of asset pricing and financial forecasting. Some of the most important topics of research are the development of new factor analysis models that extend the original Fama-French approach and incorporate a diverse and extensive set of features for asset pricing models.

The combination of all these different applications in risk management, algorithmic trading, financial forecasting, asset pricing, portfolio optimization, customer management, and digital marketing, among other areas, has led to the current wave of FinTech/RegTech innovation. A plethora of startups are currently developing new solutions using AI and ML algorithms that have more flexibility and efficiency than many traditional large institutions (Financial Conduct Authority 2015, Institute of International Finance 2015, Douglas W. Arner and Buckley 2017, Rohner and Uhl 2017, Philippon 2017, Basel Committee on Banking Supervision 2017, Financial Stability Board 2017). Because of these changes, the financial corporations have created their FinTech companies or units, and they are embracing the use of AI and ML algorithms to find new solutions for problems of finance.

2. New Perspectives

This special issue brings together papers that review some of the above finance topics that address unresolved questions or propose new methods using innovative AI and ML perspectives. Additionally, all the papers include an empirical dimension that can help practitioners to explore new approaches to optimize their financial decisions. The current issue contains papers that address the following areas of the problem domain:

(i) Microstructure:
  a) Liquidity: The capacity to anticipate the liquidity demand and supply of the limit order book is vital for the exchanges as they want to provide enough liquidity to the market, and for the traders to find new liquidity-driven strategies. Chua and Chen deal with this problem approximating the liquidity demand and supply in the limit order book through a Vector Functional AutoRegressive (VFAR) framework.
  b) Structure: A significant challenge for microstructure literature is the simulation of market microstructure features such as order imbalance or mean reversion as they can undergo significant change among the different financial products. Ju and Kim develop a neural network model using the limit order book and path dependent features to examine the lead-lag relationship between the spot and futures markets for the Korea Exchange. The model is capable of learning market microstructure features such as spread-volatility correlation, order imbalance and mean reversion. Likewise, Matsushima uses a convolutional neural network with order-based inputs to forecast the price trend of the Tokyo Stock Exchange. A distinctive aspect of this model is its capacity to learn features of the order book.
  c) Universal model: A practical question for forecasting and the development of trading strategies is to build either a model for every stock or a universal model constructed with data from a large number of stocks. Sirignano and Cont use a large-scale Deep Learning model applied to a high-frequency database for U.S. equities to uncover non-parametric evidence for the existence of a universal and stationary relationship between price formation and order flow. This universal model outperforms asset-specific models.

(ii) Forecasting:
  a) Prediction markets: The wisdom of the crowd has received increasing interest in the last few years as it has shown predictive value in the case of elections, recommendation systems, etc. Bottazzi and Giachini predict the use of the crowd in financial markets finding that prices do not converge to the true probabilities in a repeated prediction market model based on fractional Kelly traders with heterogeneous beliefs. However, the average price of the crowd approximates the true probability with lower information loss than any individual belief.
  b) Stock price forecast: Following the tradition of the early papers that forecast the price direction with technical indicators, and ML algorithms, Chen and Ge predict the Hong Kong stock market price direction using the Attention mechanism in a Long-Short Term Memory (LSTM) network. The Attention mechanism allows the neural network to focus on the most critical inputs leading to an improvement of the tuning process and the final prediction in comparison to a simple LSTM model.
  c) Market cycle: A significant challenge for macro and financial economists is to anticipate substantial market changes. Procacci and Aste propose an unsupervised algorithm, Inverse Covariance Clustering (ICC), to cluster multivariate time series into different market states in a consistent temporal way. In their approach each market state is

defined by a correlation network, characterizing the interdependencies between assets in each cluster. Their experiments accurately predict bull and bear market periods of the constituents of the Russell 1000 index.

(iii) Causality and Testing:

a) Causality: The default method of evaluating causality in the econometrics and finance literature is the Granger causality test. Creamer and Lee propose and assess a multivariate distance nonlinear causality test (MDNC) using the partial distance correlation in a time series framework. The test can detect nonlinear lagged relationships between time series, and when integrated with ML methods, can improve their forecasting power.

b) Investment strategy: A major practical problem for the investment community is the generation of overly optimistic, although unrealistic, investment strategies due to multiple testing. Lopez de Prado and Lewis evaluate this problem and propose an unsupervised learning algorithm that determines the number of effectively uncorrelated trials required for filtering out false investment strategies

(iv) Option Pricing:

a) Volatility smile: A key research topic in option pricing literature is the volatility smile. In this respect, Halperin proposes an Inverse Reinforcement Learning algorithm based only on prices and traders actions, but not rewards. This algorithm is an extension of prior work, which describes a QLBS model, combining Q-learning with the Black-Scholes model to reduce the option pricing and hedging problem to a rebalancing problem of a replicating portfolio of an option made of cash and stock. Halperin also indicates how the QLBS model can price a portfolio of options and provide a data-driven solution to the BS volatility smile problem.

In closing, as financial econometrics developed to find new solutions for financial problems, AL and ML also grew to solve problems in specific domains. Considering the significant effort that researchers in the financial industry and academia invested in AI and ML, and the dynamic and "big data" nature of the financial markets, we expect that future innovative development in this area may also come from finance. We hope that this special issue contributes to this new AI and financial ML literature and provides an opportunity to test novel algorithms on new finance problems.

We thank the authors who have contributed to this special issue and the referees who have helped to select and improve the quality of the papers, Patrick Jardine for proof-reading the article, and Drs. Michael Dempster and Jim Gatheral, the Editors-in-Chief of Quantitative Finance, for their support and advise in the development of this special issue.

References

Allen, F. and Karjalainen, R., Using genetic algorithms to find technical trading rules. *Journal of Financial Economics*, 1999, **51**, 245–271.

Amilon, H., A neural network versus Black Scholes: a comparison of pricing and hedging performances. *Journal of Forecasting*, 2003, **22**, 317–335.

Angelini, E., di Tollo, G. and Roli, A., A neural network approach for credit risk evaluation. *The quarterly review of economics and finance*, 2008, **4**, 733–755.

Baker, T. and Dellaert, B.G.C., Regulating Robo Advice Across the Financial Services Industry. *Iowa Law Review*, 2018, **103**, 713.

Bao, Y. and Datta, A., Simultaneously Discovering and Quantifying Risk Types from Textual Risk Disclosures. *Management Science*, 2014, **60**, 1371–1391.

Bates, R., Dempster, M. and Romahi, Y., Evolutionary Reinforcement Learning in FX Order Book and Order Flow Analysis. In *Proceedings of the IEEE International Conference on Computational Intelligence for Financial Engineering, Hong Kong, March 20-23, 2003*, pp. 355–362, 2003 (IEEE: Hong Kong).

Birch, A. and Aste, T., Systemic Losses Due to Counterparty Risk in a Stylized Banking System. *Journal of Statistical Physics*, 2014, **156**, 998–1024.

Chaboud, A.P., Chiquoine, B., Hjalmarsson, E. and Vega, C., Rise of the machines: Algorithmic trading in the foreign exchange market. *Journal of Finance*, 2014, **69**, 2045–2084.

Choi, J.H., Lee, M.K. and Rhee, M.W., Trading S&P500 stock index futures using a neural network. In *Proceedings of the 3rd Annual International Conference on Artificial Intelligence Applications on Wall Street*, pp. 63–72, 1995 (New York).

Creamer, G. and Freund, Y., Automated trading with boosting and expert weighting. *Quantitative Finance*, 2010, **10**, 401–420.

Creamer, G.G., Can a Corporate Network and News Sentiment Improve Portfolio Optimization using the Black Litterman Model?. *Quantitative Finance*, 2015, **15**, 1405–1416.

Decker, K., Sycara, K. and Zeng, D., Designing a multi-agent portfolio management system. In *Proceedings of the Proceedings of the AAAI Workshop on Internet Information Systems*, 1996 (AAI Press: Menlo Park,CA).

Dempster, M. and Leemans, V., An automated FX trading system using adaptive reinforcement learning. *Expert Systems with Applications: Special issue on financial engineering*, 2006, **30**, 534–552.

Dempster, M., Payne, T.W., Romahi, Y. and Thompson, G., Computational learning techniques for intraday FX trading using popular technical indicators. *IEEE Transactions on neural networks*, 2001, **12**, 744–754.

Domingos, P. and Richardson, M., Mining the network value of customers. In *Proceedings of the KDD '01: Proceedings seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, San Francisco, California, pp. 57–66, 2001 (ACM: New York, NY, USA).

Douglas W. Arner, J.B. and Buckley, R.P., FinTech, RegTech and the Reconceptualization of Financial Regulation. *Northwestern Journal of International Law and Business*, 2017, **37**, 371–413.

Fraiberger, S., News Sentiment and Cross-Country Fluctuations. In *Proceedings of the Proceedings of 2016 EMNLP Workshop on Natural Language Processing and Computational Social Science*, November, 2016 (Austin, TX).

Hanley, K.W. and Hoberg, G., Dynamic Interpretation of Emerging Systemic Risks. *Review of Financial Studies*, 2019 Forthcoming.

Hill, S., Provost, F. and Volinsky, C., Network-based Marketing: Identifying likely adopters via consumer networks. *Statistical Science*, 2006, **21**, 256–276.

Hutchinson, J.M., Lo, A.W. and Poggio, T., A nonparametric approachto pricing and hedging derivative securities via learning networks. *Journal of Finance*, 1994, **49**, 851–889.

Koyuncugil, A.S. and Ozgulbas, N., Financial early warning system model and data mining application for risk detection. *Expert Systems with Applications*, 2012, **39**, 6238–6253.

Leskovec, J., Adamic, L.A. and Huberman, B.A., The dynamics of viral marketing. In *Proceedings of the EC '06: Proceedings 7th ACM conference on Electronic commerce*, Ann Arbor, Michigan, USA, pp. 228–237, 2006 (ACM: New York, NY, USA).

Mamaysky, H. and Glasserman, P., Does Unusual News Forecast Market Stress?. *Journal of Financial and Quantitative Analysis*, 2019 Forthcoming.

Manela, A. and Moreira, A., News implied volatility and disaster concerns. *Journal of Financial Economics*, 2017, **123**, 137–162.

Moody, J. and Saffell, M., Learning to trade via direct reinforcement. *IEEE Transactions on Neural Networks*, 2001, **12**, 875–889.

Nevmyvaka, Y., Feng, Y. and Kearns, M., Reinforcement Learning for Optimized Trade Execution. In *Proceedings of the Proceedings of the 23rd International Conference on Machine Learning*, ICML 2006, Pittsburgh, Pennsylvania, USA, pp. 673–680, 2006 (ACM: New York, NY, USA).

O'Halloran, S., Maskey, S., McAllister, G., Park, D.K. and Chen, K., Big Data and the Regulation of Financial Markets. In *Proceedings of the IEEE/ACM International conference on Advances in Social Networks Analysis and Mining*, August, 2015 (Paris, France).

Philippon, T., The Fintech Opportunity. Technical report, Bank for International Settlements, 2017 Working Paper No. 655.

Richardson, M. and Domingos, P., Markov logic networks. *Machine Learning*, 2006, **62**, 107–136.

Rohner, P. and Uhl, M., Robo-Advisors vs. Traditional Investment Advisors - An Unequal Game. *Journal of Wealth Management*, 2017 Forthcoming.

Seo, Y.W., Giampapa, J. and Sycara, K., Financial News Analysis for Intelligent Portfolio Management. Technical report CMU-RI-TR-04-04, Robotics Institute,Carnegie Mellon University, 2004.

Tay, F. and Cao, L., Application of support vector machines in financial time series forecasting. *Omega*, 2001, **29**, 309–317.

Basel Committee on Banking Supervision, Sound Practices: Implications of Fintech developments for banks and bank supervisors. Technical report, Bank for International Settlements, 2017 Working Paper.

Financial Conduct Authority, Call for Input: Supporting the development and adoption of RegTech. Technical report, Financial Conduct Authority, 2015 Working Paper.

Financial Stability Board, Artificial intelligence and machine learning financial services. Technical report, Financial Stability Board, 2017 Working Paper.

Institute of International Finance, Regtech: exploring solutions for regulatory challenges. Technical report, Institute of International Finance, 2015 Working Paper.

Van Liebergen, B., Machine Learning: A Revolution in Risk Management and Compliance?. *The CAPCO Institute Journal of Financial Transformation*, 2017, **45**, 60–67.

Trippi, R.R. and DeSieno, D., Trading equity index futures with a neural network. *The Journal of Portfolio Management*, 1992, **19**, 27–33.