

DISCUSSION PAPER SERIES

DP14525
(v. 3)

ARTIFICIAL INTELLIGENCE IN ASSET MANAGEMENT

Söhnke Bartram, Jürgen Branke and Mehrshad
Motahari

**FINANCIAL ECONOMICS
INTERNATIONAL MACROECONOMICS AND FINANCE**



ARTIFICIAL INTELLIGENCE IN ASSET MANAGEMENT

Söhnke Bartram, Jürgen Branke and Mehrshad Motahari

Discussion Paper DP14525
First Published 24 March 2020
This Revision 01 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics
- International Macroeconomics and Finance

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Söhnke Bartram, Jürgen Branke and Mehrshad Motahari

ARTIFICIAL INTELLIGENCE IN ASSET MANAGEMENT

Abstract

Artificial intelligence (AI) has a growing presence in asset management and has revolutionized the sector in many ways. It has improved portfolio management, trading, and risk management practices by increasing efficiency, accuracy, and compliance. In particular, AI techniques help construct portfolios based on more accurate risk and returns forecasts and under more complex constraints. Trading algorithms utilize AI to devise novel trading signals and execute trades with lower transaction costs, and AI improves risk modelling and forecasting by generating insights from new sources of data. Finally, robo-advisors owe a large part of their success to AI techniques. At the same time, the use of AI can create new risks and challenges, for instance as a result of model opacity, complexity, and reliance on data integrity.

JEL Classification: G11, G17

Keywords: Algorithmic trading, Machine Learning, Lasso, neural networks, deep learning, decision trees, random forests, SVM, evolutionary algorithms, NLP

Söhnke Bartram - s.m.bartram@wbs.ac.uk
University of Warwick and CEPR

Jürgen Branke - juergen.branke@wbs.ac.uk
University of Warwick

Mehrshad Motahari - m.motahari@jbs.cam.ac.uk
University of Cambridge

Acknowledgements

Helpful comments and suggestions by Florian Bardong (SysAMI Advisors), Gurvinder Brar (Macquarie), Marie Briere (Amundi), Charles Cara (Absolute Strategy), Carmine De Franco (Ossiam), Giuliano De Rossi (Goldman Sachs), Marco Dion (Qube Research and Technologies), Kevin Endler (Acatis), Daniel Giamouridis (Bank of America Merrill Lynch), Alex Gracian (Resolute Investments), Farouk Jivraj (Barclays), Bryan Kelly (AQR), Petter Kolm, Alexei Kondratyev (Standard Chartered), Christos Koutsoyannis (Atlas Ridge Capital), Jörg Ladwein (Allianz Investment Management), Ke Lu, Jon Lukomnik (Sinclair Capital), Spyros Mesomeris (UBS), Matt Monach (Aberdeen Standard Investments), Andreas Neuhierl, Raghavendra Rau, Berkan Sesen, Maximilian Stroh (Invesco), Scott Taylor (AIG), Simon Taylor, Argyris Tsiaras, Nir Vulkan (Oxford MAN Institute), Markos Zachariadis, Riccardo Zecchinelli (Cabinet Office), and seminar participants at 2020 WBS Investment Challenge, Cambridge Judge Business School, Barclays Quantitative Investment Strategies (QIS) group, and the 13th Financial Risks International Forum are gratefully acknowledged. Söhnke Bartram gratefully acknowledges the warm hospitality of the Fudan University, the University of Cambridge, and the University of Oxford.

Table of Contents

1	Introduction	1
2	Trends in Artificial Intelligence.....	3
3	Portfolio Management	6
3.1	Alpha and Sigma	7
3.2	Portfolio Optimization.....	12
4	Trading.....	14
4.1	Algorithmic Trading.....	14
4.2	Transaction Cost Analysis	17
4.3	Trade Execution.....	19
5	Portfolio Risk Management.....	20
5.1	Market Risk.....	20
5.2	Credit Risk.....	23
6	Robo-Advisors.....	24
7	AI Risks and Challenges: What Can Go Wrong?.....	26
8	Conclusion	29
Appendix A: Basic Artificial Intelligence Concepts and Techniques.....		31
A.1	Artificial Intelligence and Machine Learning.....	31
A.1.1	Origin and Definition.....	31
A.1.2	Supervised Learning	32
A.1.3	Unsupervised Learning	33
A.1.4	Reinforcement Learning	33
A.2	Overview of Common Artificial Intelligence Techniques	33
A.2.1	Least Absolute Shrinkage and Selection Operator Regression.....	35
A.2.2	Artificial Neural Networks and Deep Learning.....	35
A.2.3	Decision Trees and Random Forests.....	36
A.2.4	Support Vector Machines.....	38
A.2.5	Cluster Analysis	39
A.2.6	Evolutionary (Genetic) Algorithms	39
A.2.7	Natural Language Processing.....	40
A.2.8	Comparisons of Artificial Intelligence Techniques	41
Appendix B: Trends and Patterns in Finance Research Using Artificial Intelligence		43
References		47

1 Introduction

Artificial intelligence (AI) is one of the hottest topics of current times, it has disrupted most industries in recent years, and the financial services sector has been no exception. With the advent of FinTech, which has had a particular emphasis on AI, the sector has experienced a revolution in some of its core practices. Asset management is probably the most affected area and is expected to suffer the highest number of job cuts in the near future (Buchanan, 2019). A sizeable proportion of asset management companies are now using AI and statistical models to run trading and investment platforms. The increased use of AI across a range of tasks in asset management calls for a more systematic examination of the various techniques and applications, as well as the concomitant opportunities and challenges that they bring to the sector.

This study provides a comprehensive overview of a wide range of existing and emerging applications of AI in asset management and sets out the key debates. We focus on three major areas of asset management in which AI plays a role: portfolio management, trading, and risk management. Portfolio management entails making asset allocation decisions to construct a portfolio with specific risk and return characteristics. AI techniques can contribute to this process by facilitating fundamental analysis through quantitative or textual data analysis and generating novel investment strategies. We also summarize the shortcomings of classical portfolio construction techniques and highlight how AI techniques improve the practice. In particular, AI can produce better asset return and risk estimates and solve portfolio optimization problems with complex constraints, yielding portfolios with better out-of-sample performance compared to traditional approaches.

Trading is another popular area for AI applications. Considering the growing speed and complexity of trades, AI techniques are becoming an essential part of trading practice. A particularly attractive feature of AI for the asset management industry is its ability to process large amounts of data for the generation of trading signals. Algorithms can be trained to automatically execute trades based on these signals, which has given rise to the industry of algorithmic (or algo) trading. In addition, AI techniques can reduce transaction costs by automatically analyzing the market and subsequently identifying the best time, size and venue of trades.

AI also has vast implications and potential for portfolio risk management. Since the 2008 financial crisis, risk management (and compliance) have been at the forefront of asset management practices. With increasing complexity of financial assets and global markets, traditional risk models may no longer be sufficient. At the same time, AI techniques that learn and evolve through the use of data can provide additional tools for monitoring risk. Specifically, AI assists risk managers in validating and back-testing risk models. AI approaches can also extract information from various sources of structured or unstructured data more efficiently and generate more accurate forecasts of bankruptcy and credit risk, market volatility, macroeconomic trends, financial crises, etc. than traditional techniques.

Finally, robo-advising has gained significant public interest in recent years. Robo-advisors are computer programs that provide digital financial investment advice based on mathematical rules or algorithms tailored to the needs and preferences of investors. The popularity of robo-advisors stems from their success in democratizing investment advisory services by making them cheaper and more accessible to unsophisticated individual investors. It is a particularly attractive tool for young and tech-savvy investor clienteles, such as generation Y (millennials). AI is the backbone of typical robo-advising algorithms, relying heavily on the applications of AI across all dimensions of asset management.

We also discuss a number of possible disadvantages of using AI in asset management. AI models are often opaque and complex, making it difficult for managers to monitor and scrutinize them. Their reliance on and sensitivity to data can introduce a considerable source of risk. AI models can be improperly trained as a result of using poor-quality or insufficient data. Ineffective human supervision might lead to systematic crashes, inability to identify inference errors, and lack of understanding of investment practices and attribution of performance by investors. Lastly, it is not clear whether the benefits associated with AI can justify its considerable development and implementation cost.

The remainder of the paper is organized as follows. Section 2 provides an overview of trends in AI and some of the key AI techniques frequently used in asset management. AI applications in portfolio management, trading, and risk management are discussed in Sections 3, 4, and 5, respectively. Section 6 covers the use

of AI in robo-advising, and Section 7 discusses some of the risks and concerns. Section 8 concludes with a summary of the main takeaways.

2 Trends in Artificial Intelligence

In recent years, there has been a surge in the popularity of AI in general and machine learning (ML) specifically, both in practice and academia. Consequently, the number of research papers published with the keywords “artificial intelligence” and “machine learning” has increased dramatically in the last five years (Figure 1). AI is a broader concept than ML, as it refers to the general use of computers to imitate human cognitive functions. ML is effectively a subset of AI, where machines have the ability to decide and perform actions based on past experiences. To date, AI applications in finance make mostly use of ML techniques, such as statistical learning, and thus the label artificial intelligence only applies in a very broad sense (e.g., Gu et al., 2018). Moreover, a large part of what is branded as AI (or ML) in finance is not new, but has existed in the form of statistical or econometric modelling techniques for a long time.

The recent hype about AI can be attributed to three developments that are not necessarily related to the science of AI itself (Giamouridis, 2017). First, computer processing and storage capacity have improved remarkably over the past decade, making it feasible to use some of the longstanding AI techniques. Second, there has been a substantial increase in the volume and the breadth of data that can be used to train AI models. Lastly, AI algorithms have been improved and have become widely accessible in many cases allowing them to be utilized without a necessity for expert knowledge in computer science. All these factors have also contributed to the popularity of AI and machine learning as research topics in social sciences.




Figure 1: Number of Published Research Papers by Topic Over Time

The figure presents the number of published papers with specific keywords by the year as reported by the Scopus database. The sample starts in 1996, ends in 2018, and includes papers having one of the following keywords in their abstract, title, or keyword section: “artificial intelligence” and “machine learning”.

While AI is a broad field that entails a range of approaches developed over the years, the recent interest in AI is almost entirely centered around machine learning, which is by far the most popular AI approach to date. ML is concerned with using data progressively to adapt the parameters of statistical, probabilistic, and other computing models. It essentially automates one or several stages of information processing. Although there is an extensive list of techniques that do this, most machine learning applications in asset management, and even in finance generally, rely on a number of major (classes of) techniques (Figure 2). These include artificial neural networks, cluster analysis, decision trees and random forests, evolutionary (genetic) algorithms, least absolute shrinkage and selection operator (LASSO), support vector machines, and natural language processing. Appendix A provides a detailed, more technical description of each of these techniques.

Artificial Neural Networks	<ul style="list-style-type: none"> • Non-linear regression model • Network of connected nodes that loosely model neurons in a brain • Receives a training set of input and desired output data pairs and is able to learn the relationship between them • Can then be used to predict the output of previously unseen inputs • Typical application: Forecasting
Decision Trees and Random Forests	<ul style="list-style-type: none"> • A decision tree classifies units based on their features • Classification is done by traversing a logical tree from root to leaves, at each branch moving left or right depending on the unit's features, such trees are human-readable • Constructed automatically based on training set of input and desired output pairs • Random forests simply average the outputs of several decision tree models in order to produce more reliable forecasts • Typical application: Classification and forecasting
Support Vector Machines	<ul style="list-style-type: none"> • Can be used for classification or regression • Can handle non-linear relationships by mapping the inputs to a higher-dimensional space • Faster to train than artificial neural networks • Typical application: Forecasting
LASSO	<ul style="list-style-type: none"> • Ordinary regression model with an additional penalty term that ensures choosing the smallest necessary subset of explanatory variables • Reduces spurious coefficient estimates to zero, which significantly enhances the out-of-sample performance of the model • Typical application: Forecasting
Cluster Analysis	<ul style="list-style-type: none"> • Clusters data into groups so that the units in each group have similar characteristics • The number of clusters can be defined by the user or determined automatically by the algorithm • Typical application: Asset classification
Evolutionary (Genetic) Algorithms	<ul style="list-style-type: none"> • Optimization technique capable of searching through large, complex, non-linear sets of solutions, identifying those that are preferred • Process inspired by natural evolution • Typical application: Variants of portfolio optimization that cannot be solved with classical optimization algorithms
Natural Language Processing	<ul style="list-style-type: none"> • Range of techniques used to process natural language data (e.g., textual, audio) • Particularly useful for extracting information from textual media (e.g., social media, websites, news articles) • Typical application: Automatic analysis of corporate annual reports and news articles

Figure 2: Summary of Key Artificial Intelligence/Machine Learning Techniques

The figure lists and describes major AI techniques commonly used in asset management.

Academic research interest in specific AI techniques has generally followed an upward trend over the past two decades, as illustrated by the number of published papers (Figure 3). Some of these techniques, such as evolutionary algorithms or neural networks, were established research topics long before machine learning gained popularity. On the other hand, support vector machines and natural language processing have gained interest more recently. Neural network, random forest, and natural language processing techniques have experienced the sharpest increase in their mention in published papers over the past five years characterized by the hype in machine learning research. Appendix B provides a more detailed view of the use of AI techniques in finance research based on analyzing all working papers posted on SSRN. The following sections discuss these techniques and their applications in the context of asset management.

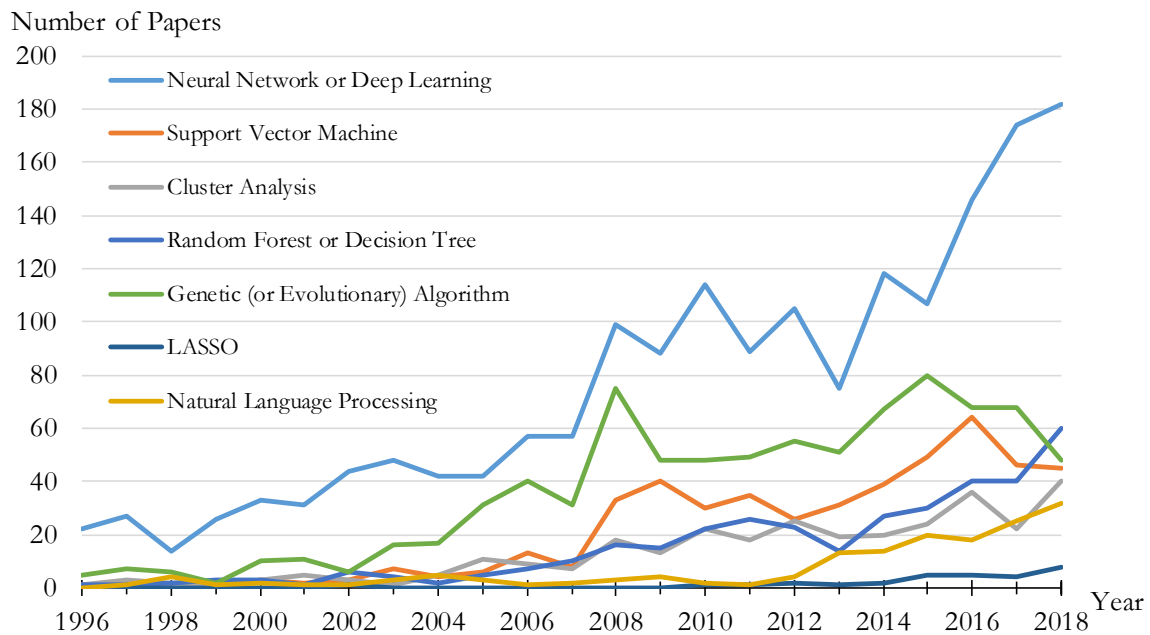


Figure 3: Numbers of Published Papers by Artificial Intelligence Technique Over Time

The figure presents the number of published papers with specific keywords by year. It is based on the number of published papers listed on Scopus starting in 1996 and ending in 2018. The papers have “finance” and/or “asset management” keywords together with at least one of the following keywords: “cluster analysis”, “genetic algorithm” or “evolutionary algorithm”, “lasso”, “natural language processing”, “neural network” or “deep learning”, “random forest” or “decision tree”, and “support vector machine”.

3 Portfolio Management

Artificial intelligence techniques can be used to perform sophisticated fundamental analysis, including the use

of text analysis, and to optimize asset allocations in financial portfolios. Amid various challenges of conventional portfolio optimization approaches, AI techniques often provide better estimates of returns and covariances that can then be used within traditional portfolio optimization frameworks. Moreover, AI can be utilized directly for asset allocation decisions to construct portfolios that more closely meet performance targets (Figure 4).

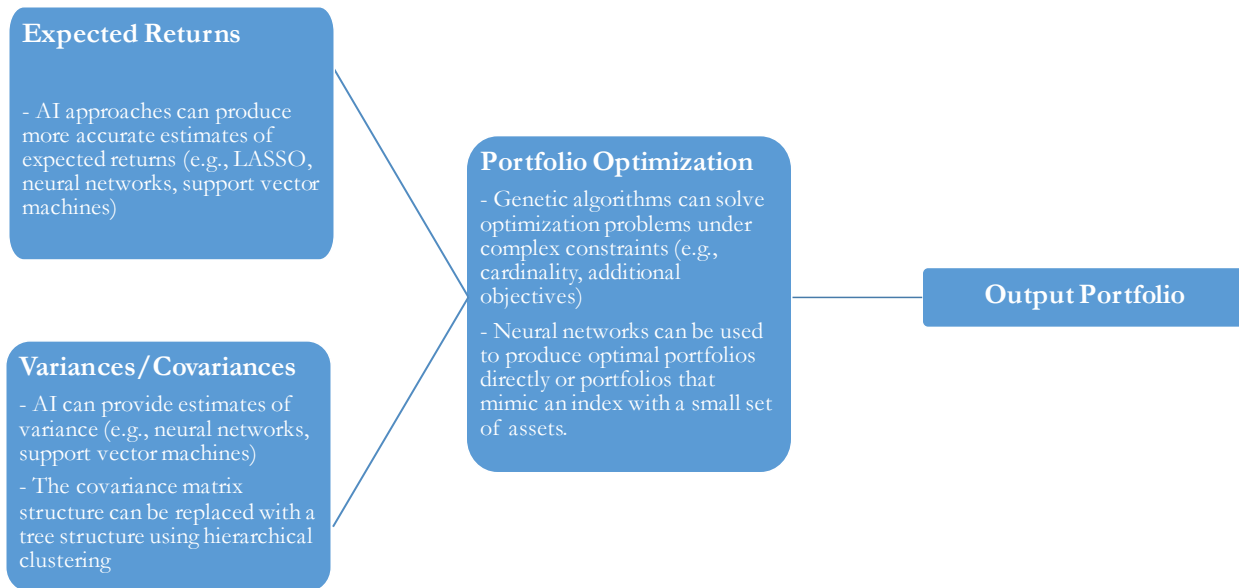


Figure 4: Artificial Intelligence in Portfolio Management

The figure presents a summary of how AI can be incorporated in portfolio construction. AI approaches can provide the inputs (i.e., expected returns, variance/covariance and asset views) and use them in asset allocation to meet targets of portfolio managers.

3.1 Alpha and Sigma

Fundamental analysis can be considered as the cornerstone of portfolio management and can be facilitated significantly by AI (Table 1). Arguably the most significant application of AI in fundamental analysis is textual analysis (Das, 2014; Kearney and Liu, 2014; Fisher et al., 2016). Natural language processing approaches are capable of extracting economically meaningful information from various sources of text, for example corporate annual reports (Azimi and Agrawal, 2019), news articles (Schumaker and Chen, 2006; Ke et al., 2019), or twitter messages (Sprenger et al., 2014). Unlike more traditional textual analysis techniques such as dictionary-based approaches that only extract information from individual words in the text, AI approaches can interpret the

context and the sentence structure as well.

LASSO regression is able to automatically select the factors with the highest explanatory power for future returns from a large set of return predictive signals documented in the literature (Feng et al., 2017; Freyberger et al., 2018). The LASSO framework can also be used to find lead-lag relationships between asset groups or markets. For example, one can investigate which domestic industry or market returns are the most significant predictors of returns among all other markets or industries (Rapach et al., 2013; 2019). More generalized versions of LASSO regression, known as “elastic nets”, complement the variable selection feature of LASSO by also ensuring that estimated coefficients are not disproportionately large (e.g., Gu et al., 2018). AI models can also be used to identify stocks expected to outperform or underperform using a range of economic or firm-level variables. The result from these analyses can then be incorporated by allocating more (less) weight to assets with high (low) alpha in the portfolio optimization process. Rather than training AI on historical data, it has also been successful to train them using actual experts’ buy or sell recommendations for each stock (Bew et al., 2019; Papaioannou and Giamouridis, forthcoming).

Across AI techniques available for return prediction, artificial neural networks have been found to perform best compared to ordinary least squares regression, elastic nets, LASSO regressions, random forests, and gradient boosted regression trees (Gu et al., 2018). In fact, the out-of-sample predictions of an artificial neural network with three hidden layers were almost 30% more accurate than those generated by a gradient boosted regression tree, which was the second best-performing technique among the six. It should be noted that these results might be highly task- and data-specific. Nevertheless, the success of neural networks in this case is largely attributed to their ability to capture complex nonlinear relationships. In addition, these models stand apart in that they are highly versatile, and a large number of functional forms and structures are available that allow neural networks to learn from data more efficiently. Recent studies have also introduced methods of interpreting neural networks statistically using confidence intervals and ranking the importance of input variables and interaction effects (Dixon and Polson, 2019).

Table 1: Artificial Intelligence and Fundamental Analysis

The table presents a list of frequently cited studies that use one or several major AI techniques (hybrid or ensemble approaches) for fundamental analysis.

Technique	Study	Sample / Data
Artificial Neural Networks	Atsalakis and Valavanis (2009)	No empirical work, surveys other studies
	Lam (2004)	Financial data for 364 S&P companies from 1985 to 1995
	Ballings et al. (2015)	Financial data for 5,767 listed European firms from 2009 to 2010
Cluster Analysis	Ballings et al. (2015)	Financial data for 5,767 listed European firms from 2009 to 2010
Decisions Trees	Ballings et al. (2015)	Financial data for 5,767 listed European firms from 2009 to 2010
	Bryzgalova et al. (2019)	Financial data for all US firms available on CRSP from 1964 to 2016
Genetic Algorithms	Hu et al. (2015)	No empirical work, surveys other studies
Hybrid / Ensemble	Li et al. (2014)	Stock data for all HKEx listed firms in year 2001
	Huang (2012)	Financial data for 200 stocks listed on the Taiwan Stock Exchange from 1996 to 2010
LASSO	Feng et al. (2017)	NYSE, AMEX, and NASDAQ stock data from 1976 to 2017
Natural Language Processing	Leung and Ton (2015)	Stock data for 2000 firms listed on ASX from 2003 to 2008
	Sprenger et al. (2014)	400,000 stock-related Twitter messages and S&P 500 stock prices for 2010
	Schumaker and Chen (2006)	9,211 financial news articles and 10,259,042 stock quotes for a 5-week period in 2005
Support Vector Machines	Han and Chen (2007)	Financial statement data for 251 stocks listed on in Shanghai and Shenzhen markets
	Fan and Palaniswami (2001)	Financial data for stocks listed on the Australian Stock Exchange from 1992 to 2000
	Ballings et al. (2015)	Financial data for 5,767 listed European firms from 2009 to 2010

It is not surprising, therefore, that neural networks are one of the most popular AI techniques for predicting stock returns (Vui et al., 2013; Abe and Nakayama, 2018), company fundamentals (Alberg and Lipton, 2018), and returns of other asset classes such as bonds (Bianchi et al., 2019). However, there is also evidence that support vector machines can be better at predicting the first two moments of asset returns compared to artificial neural networks provided that they are tuned appropriately (Huang et al., 2005; Chen et al., 2006; Arrieta-ibarra and Lobato, 2015). Consequently, a popular implementation consists of using the average prediction across a number of different AI techniques. This “ensemble” approach has been shown to produce better predictions than any individual AI technique (Rasekhschaffe and Jones, 2019; Borghi and De Rossi, forthcoming). Recent findings indicate that AI signals generate significant profits in short as well as long positions (0.78% abnormal returns per month for a long-only value-weighted portfolio) and that these profits remain statistically and economically significant even in the post-2001 period during which there is a global decay in abnormal returns (Avramov et al., 2019).

Modeling and predicting asset prices is a particularly challenging exercise when it comes to derivatives. This makes it difficult to construct optimal portfolios that include derivatives, since their prices and payoffs are not well defined and contingent on other assets. Most conventional derivative pricing approaches rely heavily on theoretical models, such as Black-Scholes, that are based on somewhat restrictive assumptions. This is, again, a realm where AI can play a role. For example, artificial neural networks can be used to price and hedge using non-parametric option pricing frameworks that perform better in delta hedging (Hutchinson et al., 1994) and forecasting future option prices (Yao et al., 2000) compared to the Black-Scholes model. Recent studies also extend the deep learning framework to price exotic (Becker et al., 2019a) and American-style (Becker et al., 2019b) options.

Lastly, AI can be used for estimating variance-covariance matrices amid its restrictive structure in the Markowitz framework. To illustrate, hierarchical cluster analysis can replace the covariance structure of asset returns with a tree structure (De Prado, 2016). This approach uses all the information contained in the covariance matrix but requires fewer estimates and thus leads to more stable and robust portfolio weights. Empirical

evidence using simulated return observations suggests that a minimum variance portfolio under this approach has a 31.3% higher Sharpe ratio than that under the classical Markowitz framework.

Ultimately, the jury is still out whether AI implementations are generally superior compared to more vanilla implementations in stock selection, factor investing or asset allocation. More evidence would be desirable to corroborate that the benefits of AI models, including their ability to capture nonlinearities, are worth the effort considering costs and potential data issues such as collinear variables. This will only become more important as time goes on because a large number of asset managers have recently started using AI, which can lead to the superior performance of AI-based investment strategies being arbitrated away in the near future. Moreover, some research advocating the use of AI in portfolio management looks at small samples of assets or emerging markets that lack liquidity and efficiency. At the same time, selecting relevant variables from the raw data and transforming them into appropriate formats for AI models to function properly, also known as “feature engineering”, constitutes an essential and time-consuming part of alpha research (Rasekhschaffe and Jones, 2019).

3.2 Portfolio Optimization

A portfolio manager’s decision entails allocating funds between a (large) set of assets such that the target portfolio satisfies an objective (e.g. mimicking an index, maximizing the Sharpe ratio, etc.) given certain constraints. The mean-variance framework of Markowitz (1952) is typically the theoretical backbone of this exercise. However, there are two main challenges when it comes to practice (Michaud and Michaud, 2008; Kolm et al., 2014). First, the optimal asset weights are highly sensitive to estimates of expected returns. Considering that estimates of future expected returns are often uncertain, the optimization exercise can yield unstable weights that perform poorly out of sample. In fact, the noise incorporated in return estimates can erode any diversification benefit. For example, DeMiguel et al. (2007) show that an equally-weighted portfolio has a higher out-of-sample Sharpe ratio than the optimal Markowitz portfolio and a range of other optimal portfolios.

Second, estimating the variance-covariance matrix, which is at the heart of Markowitz’s theory, requires large time series of data and the assumption of stable correlations between asset returns. Moreover, the matrix

becomes unstable when asset correlations increase, which are times when diversification is most important and yet more difficult to achieve (De Prado, 2016).

Artificial intelligence addresses these challenges in two ways. First, AI can produce more accurate return and risk estimates that can be used within traditional portfolio construction frameworks. Second, AI techniques can provide alternative portfolio construction approaches in order to generate more accurate portfolio weights and produce optimized portfolios with better out-of-sample performance. While there is limited empirical evidence, interest seems to be growing among both academics and practitioners.

In particular, artificial neural networks can be trained to make asset allocation decisions subject to complex constraints that are often not straightforward to integrate into the mean-variance framework. For example, a neural network can select portfolios according to a learning criterion that maximizes returns subject to value-at-risk constraints (Chapados and Bengio, 2001). Artificial neural networks can also solve complex multi-objective optimization problems. To illustrate, a neural network-based methodology can construct a mean-variance-skewness optimal portfolio in a fast and efficient manner (Yu et al., 2008). Furthermore, artificial neural networks can incorporate views about the future asset performance into the portfolio optimization using a Black and Litterman (1992) framework, generating higher out-of-sample Sharpe ratios than the market portfolio (Zimmermann et al., 2002).

Another popular artificial intelligence technique in portfolio construction is evolutionary algorithms that have the flexibility to accommodate more complex asset allocation problems. For example, evolutionary algorithms solve optimization problems under cardinality constraints (restricting the number of assets in the portfolio) and maximum or minimum holding thresholds (Branke et al., 2009). Evolutionary algorithms are also able to incorporate additional objectives. For example, one can incorporate model risk (i.e., the risk of failing to produce accurate estimates of asset returns and volatilities due to model mis-specifications) in the optimization problem to reduce forecasting error (Skolpadungket et al., 2016). Optimal portfolios using this approach have better realized Sharpe ratios by around 10% than those that consider only return and volatility in their objective functions.

The ability of artificial neural networks to capture non-linear relationships between assets without any prior knowledge about the underlying structure of the data can be useful in synthetic replication, that is, replicating a benchmark portfolio such as an index by holding a fraction of the constituents while minimizing the tracking error by matching some of the risk factors of the benchmark. For example, artificial neural networks can approximate the FTSE 100 index with only seven stocks (Lowe, 1994), resulting in lower transaction costs from portfolio rebalancing as well as reduced portfolio management and monitoring costs. This framework has promising out-of-sample performance and is flexible enough to generate target portfolios with other specified characteristics. For example, one can find the best strategy (i.e., the strategy with the lowest risk or cost) to construct a portfolio that outperforms a specific index by 1% on an annual basis (Heaton et al., 2017).

4 Trading

Algorithms can play a role in all stages of the trading process (Nutti et al., 2011). The trading process can be broken down into pretrade analysis, trade execution, and post-trade analysis (Figure 5). Pretrade analysis entails using data to analyze properties of financial assets with the objective of forecasting their future performance as well as the risks and costs involved in trading them. Insights from this analysis ultimately lead to the execution of trades. This can be a manual stage as asset managers might want to consider results from pretrade analyses together with risk assessments and client preferences. However, in high frequency or fully automated systems, pretrade analysis does not involve any human intervention. Trade execution implements trades while ensuring low transaction costs. Actual trading outcomes are evaluated during post-trade analysis in order to monitor performance and improve the trading system. Post-trade analysis often involves some form of human supervision or overlay. In contrast, pretrade analysis and trade execution are handled mostly by algorithms, since they require timely and complex analyses.

4.1 Algorithmic Trading

Artificial intelligence plays a role in trading primarily by facilitating algorithmic trading – defined as algorithms that automate one or more stages of the trading process. Algorithmic trading has experienced a growing presence in asset management thanks to three recent phenomena (Kirilenko and Lo, 2013). First, developments in computing power, data science, and telecommunication have led to structural changes in the way financial

markets operate. Computers are now capable of collecting and analyzing large amounts of data and execute trades in milliseconds without any human intervention. Second, breakthroughs in quantitative finance and machine learning have provided the necessary tools for computers to conduct insightful financial analysis faster and more efficiently than human beings. Third, the increasing speed, complexity and scale of financial markets together with the breadth of new structural products have made it difficult if not impossible for humans to keep track of markets and make real-time trading decisions, whereas complex AI techniques such as artificial neural networks can now be implemented in close to real time (Leshik and Cralle, 2011).

Strategies used in algorithmic trading are often based on technical analysis, which uses past stock and market data to predict future asset returns. While it is possible to perform fundamental analysis as well, algorithmic trades are often of a high frequency so that analyzing lower frequency data such as firm fundamentals is typically less effective. There is also evidence indicating that technical indicators dominate fundamental ones in generating profitable trading signals using AI (Borghi and De Rossi, forthcoming). Therefore, AI-based approaches have established a more active presence in technical analysis so far (Table 2).

Table 2: Artificial Intelligence and Technical Trading Rules

The table presents a list of frequently cited studies that use one or several major AI techniques (hybrid or ensemble approaches) to devise technical trading rules used in algorithmic trading.

Technique	Study	Sample / Data
Artificial Neural Networks	Dixon et al. (2016)	5-minute mid-prices for 43 CME listed commodity and FX futures from 1991 to 2014
	Choudhry et al. (2012)	JPY/USD, DM/USD, and USD/EUR exchange rates from 1998 and 1999
	Gradojevic and Yang (2006)	CAD/USD exchange rates from 1990 to 2000
	Atsalakis and Valavanis (2009)	No empirical work, surveys other studies
	Dunis et al. (2010)	Daily EUR/USD exchange rates from 1999 to 2007
	Fischer and Krauss (2018)	Daily stock data for the constituents of the S&P 500 from 1992 until 2015
Cluster Analysis	Liao and Chou (2013)	30 industrial indices from TAIEX, SSE/SZSE, and HSI from 2008 to 2011
Decisions Trees	Booth et al. (2014)	Stock data for 30 firms from the DAX stock index from 2000 to 2013
	Coqueret and Guida (2018)	Financial data for a sample of between 305 and 599 large US firms from 2002 to 2016
Genetic Algorithms	Hu et al. (2015)	No empirical work, surveys other studies
	Allen and Karjalainen (1999)	S&P 500 index daily prices from 1928 to 1995
	Manahov et al. (2014)	1-minute quote data for six major currency pairs from 2012 to 2013
	Berutich et al. (2016)	Stock data for 21 firms listed on the Spanish market from 2000 to 2013
Hybrid / Ensemble	Cheng et al. (2010)	TAIEX stock index data from 2000 to 2005
	Tan et al. (2011)	Stock data for 20+ firms from the US market from 1994 to 2006
	Tsai et al. (2011)	Stock data for a subset of the Taiwan stock market companies from 2002 to 2006
	Nuij et al. (2014)	Stock data for all FTSE 350 stocks from January to April 2007
	Geva and Zahavi (2014)	Stock data for 72 S&P 500 firms from 2006 to 2007
LASSO	Chinco et al. (2019)	1-minute returns of NYSE stocks from 2005 to 2012
Natural Language Processing	Renault (2017)	Dataset of stocks with messages published on StockTwits from 2012 to 2016
	Hagenau et al. (2013)	Stock data for a subset of German and British firms from 1997 to 2011

The main inputs to traditional forms of technical analysis are past price and trading volume data. Strategies that are based on prices often model trends, such as momentum or reversal, and cycles using historical data in order to forecast future returns. On the other hand, volume-based strategies predict future returns based on recent investor trading activity. Modern technical analysis also incorporates information from other sources, including fund flows, investor trades, and textual data from news articles or online sources. AI techniques using natural language processing can be particularly useful when it comes to these new unstructured sources of data.

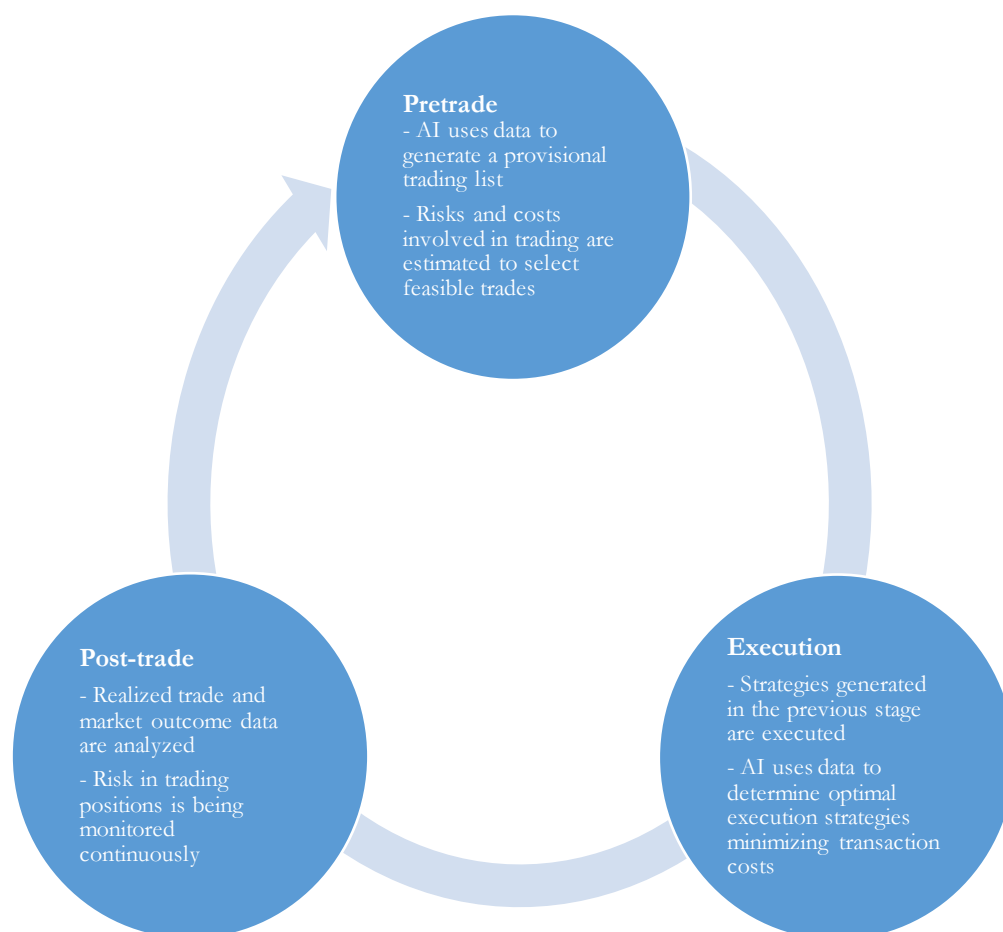


Figure 5: Algorithmic Trading With Artificial Intelligence

The figure presents the three stages of algorithmic trading and summarizes the applications of AI in each stage.

4.2 Transaction Cost Analysis

Analyzing transaction costs is an essential part of pretrade analysis that indicates whether the costs of trading

are small enough for a trading signal to generate profits net of implementation costs. Transaction costs have three main components: bid-ask spreads, market impact costs, and trading commissions. Among these three, market impact costs – defined as the adverse effect of a trade on market prices – are the only costs that are not observable before the trade is initiated. Nevertheless, it is crucial to have an estimate of market impact costs as they represent a significant portion of transaction costs: Market impact absorbs as much as two-thirds of trading gains made by systematic funds (The Financial Stability Board, 2017).

Artificial intelligence approaches complement traditional market impact models by providing additional insights. The non-parametric structure of AI techniques, together with their ability to capture nonlinear dynamics, are particularly useful for predicting market impact, and various AI techniques have been tested for this purpose. Performance weighted random forests are found to outperform linear regression, artificial neural networks, and support vector machines in predicting the market impact of a market order by 20% out of sample (Booth et al., 2015). On the other hand, support vector machines do not seem to perform particularly well when it comes to forecasting market impact, while artificial neural networks do well if they are properly defined and estimated (Park et al., 2016).

Although these non-parametric techniques perform well in estimating market impact, they have two major shortcomings. First, the majority of approaches do not have any economic intuition for the drivers of price impact. This makes them prone to capturing noise rather than relevant information. Second, these techniques cannot distinguish between permanent and temporary market impact, which would require additional variables including trade direction and liquidity (Farmer et al., 2006). To address these two issues, a parametric approach such as LASSO regression can be used alongside non-parametric techniques. With LASSO regression, the most informative variables capturing information related to the order book and other sources are selected to predict price impact. Empirical evidence indicates that trade sign, market order size, and liquidity based on best limit order prices are the most important variables for forecasting market impact (Zheng et al., 2013). A Bayesian network model is another approach for estimating market impact while providing intuition on the main drivers. Unlike most other machine learning techniques, this approach can also account for variables with data availability issues and model them as latent variables using Bayesian inference. Thanks to this feature, other

important variables can be identified (e.g., net order flow imbalance) and added to the model to improve the forecast (Briere et al., 2019).

Another useful application of AI consists of estimating the market impact of trades in assets that lack sufficient (or any) historical trading data, since it is almost impossible to use traditional approaches to estimate the market impact costs in this case. A cluster analysis approach can tackle this problem by identifying comparable assets with similar behavior and using their historical data instead. For example, cluster analysis can allocate bonds into clusters based on their duration, maturity, or value outstanding and measure their similarity according to these variables. Within each cluster, the information of other bonds is used for bonds without sufficient data. Bloomberg's liquidity assessment tool (LQA) notably uses this technique to provide liquidity information for various assets.

4.3 Trade Execution

Executing large trades often involves significant market impact costs. Therefore, such trades are typically broken up into a sequence of smaller orders, which are easier and cheaper to execute. This is known as the execution strategy that requires determining the timing and size of smaller orders using some form of execution model. The objective of such models is to minimize transaction costs while completing the transaction within a specified period of time. Classical modeling approaches for this problem use stochastic control techniques to determine optimal execution strategies (a methodology that goes back to Bertsimas and Lo, 1998). However, classical models often rely on restrictive assumptions regarding asset price dynamics and the functional form of market impact (Kearns and Nevmyvaka, 2013).

Artificial intelligence approaches facilitate trade execution modelling by actively learning from real market microstructure data when determining optimal execution strategies. Recent studies advocate reinforcement learning techniques, i.e. algorithms that receive vectors of microstructure and order book variables (bid-ask spread, volume imbalances between the buy and sell sides of limit order book, signed transaction volume, etc.) as input and return optimal execution strategies as output (e.g., Nevmyvaka et al., 2006; Kearns and Nevmyvaka,

2013; Hendricks and Wilcox, 2014; Kolm and Ritter, 2020). The algorithms essentially learn to map each combination of input variables, known as a “state”, to trading actions such that transaction costs are minimized (Kearns and Nevmyvaka, 2013).

The advantage of AI-based approaches is that they rely on data rather than normative assumptions to determine market impact costs, price movements and liquidity. This gives them the flexibility to adapt as market conditions change and new data becomes available. However, these models are often difficult to train and understand especially for large portfolios that benefit the most from a reduction in transaction costs. In addition, systematic execution strategies run the risk of cascading into a systemic event affecting the whole market. A famous precedent for this phenomenon is the so-called Flash Crash of 2010 (Kirilenko et al., 2017).

5 Portfolio Risk Management

Artificial intelligence also has applications in risk management, with regards to market risk but also for credit risk (The Financial Stability Board, 2017; Aziz and Dowling, 2019). Market risk refers to the likelihood of loss due to aggregate market fluctuation, while credit (or counterparty) risk is the risk of a counterparty not fulfilling its contractual obligations, which results in a loss in value (Figure 6). Although AI has broader uses in risk management, these two categories are the most important as far as asset management is concerned.

5.1 Market Risk

Market risk analysis involves modelling, assessing and forecasting risk factors that affect the investment portfolio. Artificial intelligence can play a role in this area in three ways: (i) making use of qualitative data for risk modelling, (ii) validating and back-testing risk models, and (iii) producing more accurate forecasts of aggregate financial or economic variables (Figure 6).

One area of application for AI in market risk management relates to extracting information from textual or image data sources. Textual data sources, including news articles, online posts, financial contracts, central bank minutes and statements, and social media, can contain valuable information for managing market risk (Groth and Muntermann, 2011). Satellite images are analyzed to predict sales at supermarkets or future crop harvests (Katona et al, 2018). The information provided by these sources are, in many cases, not captured by

other quantitative variables. For example, AI approaches using textual information have been shown to improve predictions of market crashes (Manela and Moreira, 2017), interest rates (Hong and Han, 2002), and other major macroeconomic outcomes (Cong et al., 2019) compared to information captured by other data sources. These approaches can also extract information from corporate disclosures with the aim of determining firms' systematic risk profiles (e.g., Groth and Muntermann, 2011; Bao and Datta, 2014; Cong et al., 2019). All these applications have triggered an interest among central banks to incorporate methods of AI-based text mining in macro-prudential analyses (Bholat et al., 2015). To date, empirical implementations and evidence in this area are scarce.

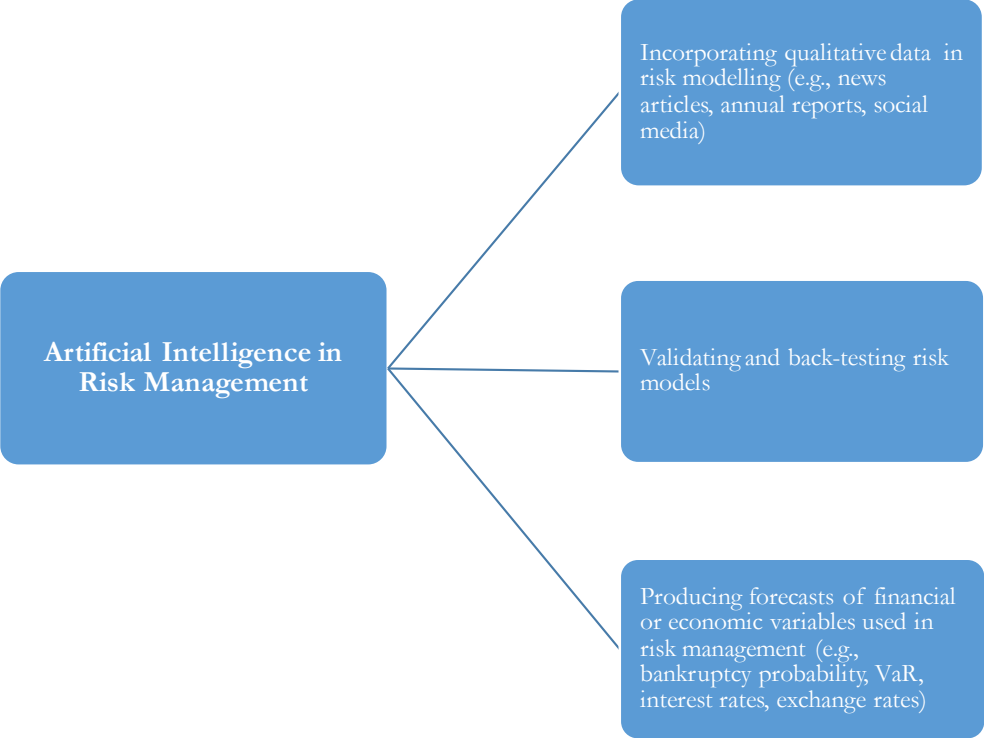


Figure 6: Artificial Intelligence Applications in Risk Management

The figure presents a summary of three areas, in which artificial intelligence can play a role in risk management.

AI can also help risk managers validate and back-test risk models (The Financial Stability Board, 2017). This is an important part of model risk management emphasized by regulators and financial supervisory institutions (Board of Governors of the Federal Reserve System, 2011). Unsupervised AI approaches can detect anomalies in risk model output by evaluating all projections generated by the model and automatically identifying any irregularities. Risk managers can also use supervised AI techniques to generate benchmark forecasts as part of model validation practice. Comparing model results and benchmark forecasts allows assessing whether the risk model is producing predictions that are significantly different from those generated by AI. A significant disagreement between AI forecasts and standard risk model outputs can highlight potential problems and trigger a more thorough investigation.

Depending on the exposure of the assets in a portfolio to the underlying risk factors, various financial or economic variables can affect its performance. Therefore, it is important to model future trends in these factors, especially macroeconomic variables (Elliott and Timmermann, 2008; Ahmed et al., 2010), and artificial neural networks are particularly popular in this context. For example, empirical evidence suggests that variants of artificial neural networks perform significantly better compared to linear autoregressive approaches in forecasting 47 monthly macroeconomic variables of the G7 economies (Terasvirta et al., 2005). However, using artificial neural networks entails the risk of producing implausible forecasts at long horizons. Nonetheless, artificial neural networks have been particularly successful in forecasting interest rates (e.g., Kim and Noh, 1997; Oh and Han, 2000) and exchange rates (e.g., Zhang and Hu, 1998; Kaashoek and van Dijk, 2002; Majhi et al., 2009).

Artificial neural networks can also be used to devise systematic risk factors. These models can capture nonlinearities and interactions of covariates including firm characteristics and macroeconomic variables (e.g., Bryzgalova et al., 2019, Chen et al., 2019; Gu et al., 2019; Feng et al., 2019). Such factors can better account for risk premia and distinguish between non-diversifiable and diversifiable (idiosyncratic) risk compared to conventional linear factors. LASSO regressions can also be useful in determining systematic factor structures. These models are able to select the most relevant systematic risk factors from a subset of factors or market indices (Giamouridis and Paterlini, 2019).

AI techniques, especially artificial neural networks and support vector machines, can also predict market volatility and financial crises. The ability of both methods to capture nonlinear dynamics results in their advantage over traditional generalized autoregressive conditional heteroscedasticity (GARCH) models. Artificial neural networks can predict market volatility either directly (Hamid and Iqbal, 2004) or in combination with a variant of GARCH (Donaldson and Kamstra, 1997; Fernandes et al., 2014). However, some researchers found support vector machines to be superior to artificial neural networks in this context (Chen et al., 2010). In addition to volatility modelling, artificial neural networks and support vector machines are used to predict financial crises. Models performing this forecasting task are often referred to as early warning systems. Almost all major financial institutions use a form of early warning system to monitor systemic risk. Artificial neural networks and support vector machines have been shown to predict currency crises (e.g., Lin et al., 2008; Sevim et al., 2014), banking crises (e.g., Celik and Karatepe, 2007; Ristolainen, 2018), and recessions generally (e.g., Yu et al., 2010; Ahn et al., 2011; Gogas et al., 2015) with reasonable accuracy. Nevertheless, crises are rare financial events so in the absence of a sufficient number of such events in the sample one could question the ability of AI models to accurately predict future instances.

5.2 Credit Risk

The objective of credit risk management is to ensure that the failure of any counterparty to meet its obligations does not have a negative effect on the portfolio beyond specific limits. Asset managers need to monitor the credit risk of the entire portfolio as well as individual positions and transactions. This practice involves modelling the solvency risk associated with institutions issuing financial products, including equities, bonds, swaps, and options. There is an extensive range of approaches for modelling solvency or bankruptcy risk. Multivariate discriminant analysis, logit and probit models are among the most common traditional methods used (Bellovary et al., 2007).

Credit risk modelling is one of the first areas of finance in which the application of artificial intelligence techniques was considered. The two most widely used techniques are artificial neural networks and support vector machines. In fact, artificial neural networks have become mainstream bankruptcy modelling techniques since the early 1990s (Tam, 1991). The popularity of artificial neural networks is largely due to their higher

success in forecasting bankruptcy and determining credit ratings compared to traditional techniques (e.g., Zhang et al., 1999; Tsai and Wu, 2008). However, more recent studies advocate the use of support vector machines (e.g., Auria and Moro, 2008; Ribeiro et al., 2012), since they yield slightly more accurate bankruptcy forecasts than artificial neural networks (Huang et al., 2004). Moreover, support vector machines are less likely to face some of the common issues of artificial neural networks such as overfitting. Artificial neural networks and support vector machines also perform particularly well at estimating loss given default (defined as the economic loss when default occurs), which the Basel II Accord requires financial institution to model in addition to the default probability for regulatory capital monitoring purposes (Loterman et al., 2012).

Beyond support vector machines and artificial neural networks, a wide range of other AI approaches including genetic algorithms (Varetto, 1998) can be used for credit risk modelling (Kumar and Ravi, 2007; Pena et al., 2011). Since each of the modelling techniques has its own specific advantages and disadvantages, an ensemble technique that uses various approaches separately and then combines the resulting predictions should be considered for achieving the best performance (Verikas et al., 2010).

6 Robo-Advisors

Robo-advisors are computer programs that provide customized advice to individual investors, assisting them in their investment activities. These programs have gained significant attention recently because of their success in reducing barriers to entry for retail investors. There is growing interest within academia to enhance robo-advisors using AI (Figure 7). The primary focus is on devising algorithms known as recommender systems that produce optimal portfolios catered to investors' risk appetites (e.g., Xue et al., 2018). However, robo-advising can integrate all types of AI applications into portfolio management, trading and risk management discussed in previous sections. By building on the success of AI in these fields, robo-advisors can not only produce portfolios with better out-of-sample performance for investors, but also rebalance portfolios automatically ensuring minimal transaction costs and managing the risks of the portfolio. This can be performed at a cheaper price and through a simplified interface, which ultimately makes investing both more beneficial and more accessible for retail investors than human advisors.

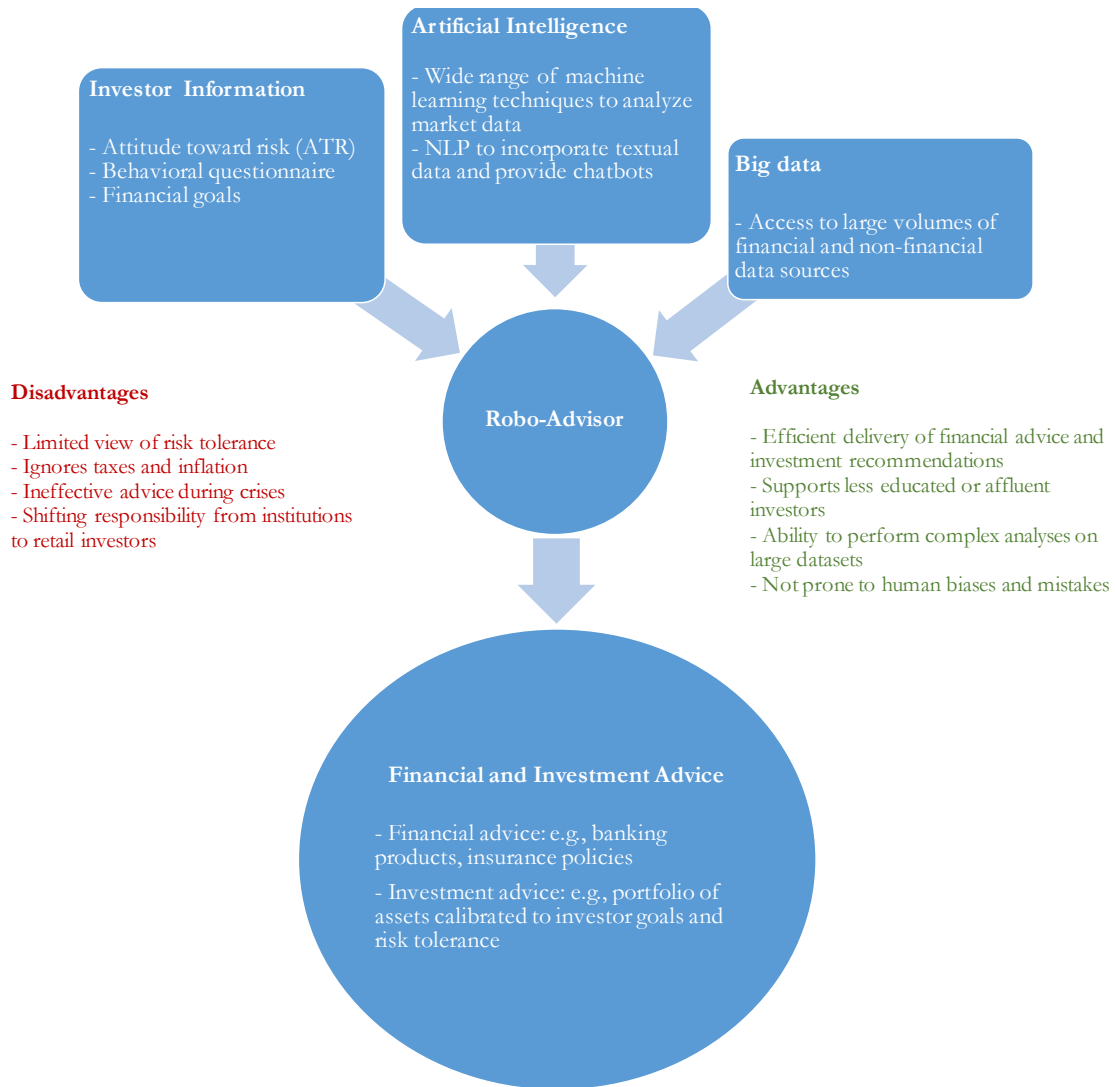


Figure 7: Robo-Advising With Artificial Intelligence

The figure illustrates the structure of robo-advisor systems that incorporate artificial intelligence. Advantages and disadvantages of these systems are also summarized below.

Robo-advisors are also less prone to behavioral biases, mistakes, and illegal practices. In fact, it has been shown that they appeal the most to investors that are afraid of being victimized by investment fraud (Brenner and Meyll, 2019). More sophisticated institutional investors can benefit from robo-advisors thanks to their ability to process a wide range of financial data more efficiently. While reducing behavioral biases when making investment decisions is beneficial to all types of investors (D’Acunto et al., 2017), less sophisticated investors particularly benefit from robo-advice in terms of enhancing portfolio performance, diversification,

and reducing volatility. At the same time, robo-advisors often encourage investors to trade more due to the integration of trade execution services. This can be both a benefit, in terms of encouraging investors to re-balance positions more often, and a pitfall, as it can lead to excessive trading that benefits robo-advising systems through commissions at the expense of investors. In order for an investor to be able to use robo-advisors and benefit from their advantages, a minimum level of technological understanding and financial sophistication is required.

Not all robo-advisors necessarily use new, sophisticated methods. An analysis of 219 international robo-advisors shows that Markowitz's portfolio theory is the most prevalent approach, though some systems do not disclose their techniques (Beketov et al., 2018). Particularly more sophisticated robo-advisors rely on proprietary algorithms and do not divulge the details their approach to analyze portfolios and make recommendations. Nevertheless, looking at some of the notable cases in the industry leaves no doubt that the most successful robo-advisors rely heavily on artificial intelligence to conduct investment and trading analyses (Sabbharwal, 2018). After all, robo-advising and FinTech in general derive most of their success from collecting and analyzing data, and AI is an integral part of this (Dhar and Stein, 2016).

7 AI Risks and Challenges: What Can Go Wrong?

While many studies of AI in finance highlight the advantages and benefits of AI for various finance applications, users of AI should also be aware of some of the actual or perceived risks and downsides associated with using AI in asset management. These issues are often related to complexity, opacity, and dependence on data integrity (Figure 8).

It is difficult, if not impossible, to understand how most AI models process data and make inference. As the complexity of the task or the algorithm grows, opacity can render human supervision ineffective and become an even more significant problem. This might have repercussions for asset managers in three ways. First, it is hard to predict how AI models would respond to major surprises or "black swan" events, which could lead to systematic crashes. Even in the absence of major events, AI algorithms may make the same errors at the same time, introducing the risk of cascading market crashes. Indeed, the considerable cost of producing AI algorithms has led to most asset management companies using the same tools and algorithms. This can

make AI-driven crashes much more likely than other cascading algorithmic crashes we have experienced. Cascading algorithmic crashes are not specific to AI systems and may arise from even simple widespread quantitative approaches such as value investing. What makes AI different, however, is that its opacity may prevent such risks from being properly modelled and monitored.



Figure 8: Artificial Intelligence Areas of Concern

The figure summarizes major potential sources of risk introduced by adopting AI in asset management.

Second, AI can take wrong decisions based on incorrect inferences capturing spurious or irrelevant patterns in the data. For example, artificial neural networks that are trained to pick stocks with high expected returns might select illiquid, distressed stocks (Avramov et al., 2019). Third, it can become more challenging to

attribute investment performance. For example, the widely used Barra Risk Factor Analysis which is based on linear factor models might not suit AI-based strategies that capture nonlinear relationships between characteristics and returns. Consequently, it can be difficult to explain to investors how and why the investment strategy failed in case of poor fund performance, undermining investors' trust in the fund or even in the industry. In an attempt to better understand the behaviour of AI models, some people approximate the prediction behavior of the AI model by constructing an additional, simpler and interpretable "surrogate model". Shapley values from game theory are a way to understand how much different feature values contribute to a prediction. A good overview on these and many other approaches to explaining AI models can be found in Molnar (2020).

Finally, the black box character of many AI systems raises the issue of responsibility and makes it challenging to regulate (Zetzsche et al., 2020).

Data quality and sufficiency can be another major source of concern. AI models, like any other empirical model, rely on the integrity and the availability of data. Poor data quality can easily trigger what is famously known as "garbage in, garbage out". This is a particularly important matter as AI outputs are often taken at face value. Therefore, it may not be as easy to identify data-related issues by evaluating the model outcomes. Furthermore, AI models require large amounts of data during the learning phase, often more than is available. This might lead to improper calibration caused by the input data's poor signal-to-noise ratio, especially when it comes to low-frequency financial data with plenty of missing observations. Imputation, a pre-processing step that substitutes missing observations by statistical values (e.g., Kofman and Sharpe, 2003) may help, but obviously only to a certain extent. It is argued that past data in general might not fully represent the future. This can become particularly problematic when the short time series of available financial data miss certain important extreme events in the past making it more likely for AI models to fail during crashes or crises (Patel and Lincoln, 2019). As a side effect, the growing presence of AI in the investment industry and the reliance of asset managers on it for day-to-day tasks might further increase their cybersecurity risk (Board of Governors of the Federal Reserve System, 2011).

Overall, it may not be clear whether the benefits of AI outweigh the considerable costs of investing in the software, hardware, human resources, and data required by AI systems. After all, the resources of asset

managers to develop and test new strategies are limited, so that investment in AI needs to be considered alongside mutually exclusive, competing research projects. Once the current AI hype vanishes, investors may become less keen to invest in AI-driven funds making it harder to break even on investments in AI infrastructure. Thus, asset managers will need to carefully consider both benefits and costs of AI (Patel and Lincoln, 2019; Buchanan, 2019), if only not to get cold feet in the next AI winter.

8 Conclusion

The use of artificial intelligence in asset management is an emerging field of interest among both academics and practitioners. AI has vast applications for portfolio management, trading, and risk management that enable the industry to be more efficient and compliant. It also serves at the heart of new practices and activities such as algorithmic trading and robo-advising. Nevertheless, AI is still far from replacing humans completely. Indeed, most of its operations within asset management are confined and controlled by some form of human supervision. Consequently, a better way to describe AI is as a collection of techniques that automate or facilitate (often small) parts of the practice of asset managers, from solving portfolio optimization problems with specific conditions to fully automated algorithmic trading systems.

Collectively, the success of AI in asset management is linked to its three key inherent capabilities. First, AI models are objective, highly successful in conducting repetitive tasks, and able to pick up patterns in high dimensional data, which may not be perceptible by humans. AI can also analyze data with minimal knowledge about its structure or the relation between input and output, including nonlinear relations. This feature is especially useful for forecasting and yields more accurate estimates as it does not rely on restrictive assumptions inherent in more traditional methods. Second, AI can extract information from unstructured data sources such as news articles, online posts, reports, or images. This results in the incorporation of a tremendous amount of information into financial analysis without manual processing and intervention. Third, AI algorithms, unlike other statistical techniques, are often designed to improve in terms of re-adjusting in accordance with the data. This means there is no need for manual reconfiguration or parameter re-estimation, which is essential for traditional models.

Finally, AI's greatest strength, its ability to process data with minimal theoretical knowledge or supervision, can also be its greatest weakness. Indeed, it is famously said that AI always generates a result even when there should not be one. This is especially a problem when data quality is poor, when the task is too complex for humans to monitor or understand, or when there is risk of cascading systemic failures due to several AI algorithms reacting to each other. It is important that asset managers bear such issues in mind as the role of AI becomes more pervasive and significant.

Appendices

Appendix A: Basic Artificial Intelligence Concepts and Techniques

A.1 Artificial Intelligence and Machine Learning

A.1.1 Origin and Definition

Artificial intelligence is widely believed to have started at the Dartmouth Summer Research Project on Artificial Intelligence, a workshop organized by John McCarthy in the summer of 1956 at Dartmouth College. Many prominent mathematicians and scientists, including Marvin Minsky and Claude Shannon, attended this six-week brainstorming workshop. The proposal for the workshop introduced the term artificial intelligence and stated the following objectives:

“The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.”

(McCarthy et al., 2006, p. 12)

In recent years, the original definition of AI and what it should encompass has evolved. Russell and Norvig (2009) distinguish the following four different dimensions, or schools of thought, that determine the objective of AI.

1. **Acting Humanly:** From this point of view, AI refers to the challenge of creating computers capable of performing tasks similar to humans. An example of this is the Turing Test, proposed by Alan Turing. The test poses a challenge in which a human interrogator presents questions and receives responses from either another human or a machine. The machine passes the test if the interrogator is unable to distinguish a human’s answers from a machine’s.
2. **Acting Rationally:** This dimension aims to build agents that act rationally, i.e., that aim to achieve the best outcome or, when there is uncertainty, the best expected outcome.
3. **Thinking Humanly:** This perspective refers to the replication of human thinking processes. The field of cognitive science is a major manifestation of this approach to AI. It uses computer programs and insights from experimental psychology to emulate the human mind.

4. **Thinking Rationally:** Thinking rationally refers to using rules for reaching logical conclusions based on premises assumed to be true.

Until a few decades ago, most research in AI fell into the category of thinking rationally represented by expert systems. Such systems have large knowledge bases and an inference mechanism that allows deducing new knowledge by logically deriving it through rules. For example, knowing that all men are mortal, and that Socrates was a man, an expert system could infer that Socrates was mortal. Expert systems were highly popular in the 1970s and 1980s, but the need for building large, complex knowledge bases and their deterministic nature have made them fall out of favor. The idea to have the machine learn through observations, i.e., machine learning, eventually turned out to be more applicable in practice, and it is the predominant AI technique behind most modern applications.

Problems studied under machine learning are of three main types: supervised learning, unsupervised learning, and reinforcement learning (Alpaydin, 2010; Murphy, 2012), each of which have common applications. While many of these techniques have been around for decades, the sudden surge in their popularity and application is due to the improvement in performance as the result of technological progress that has enabled computers to train machine learning models to a scale that was not possible even a few years ago.

A.1.2 Supervised Learning

Consider the problem of determining the sales price of a house based on a set of attributes, such as interior square footage, geographic location, and number of floors. In supervised learning, an algorithm establishes a mathematical relation between the feature data (square footage, location, and number of floors) and the response data (sales price). Rather than explicitly programming the model, a supervised learning algorithm is given a set of training data. It then adjusts its model in order to minimize the prediction error on the training data. Once the model has been established, it can be used to infer a response from features that have not been observed before. Often the training is iterative, i.e. a relation is first guessed randomly, and subsequently adjusted based on how erroneous the guess was. Over time, increasingly accurate relations are produced until, ideally, a “best” relation is found. Effectively a machine learns how to relate the feature data to the response data. Since a training set of correctly classified response data is used to guide the learning process, the learning

is deemed supervised. Supervised learning is currently the most common learning approach in practice.

There are two main applications for supervised learning: classification and regression. Predicting the sales prices of houses is an example of regression, since the response data is quantitative and continuous. In classification, one is interested in determining a response that falls into one of a few categories, such as whether a credit card transaction is fraudulent or not, based upon observed features such as the distance of the transaction from the cardholder's residence, the amount of the transaction, and the object purchased.

A.1.3 Unsupervised Learning

Unsupervised learning is used to identify structures in data without access to labels. The most popular example is clustering, i.e., the categorization of data into different groups where the elements of each group have similar characteristics. This is useful for example in marketing, where customers can be separated into different groups, and different marketing strategies can be developed for each group. Other applications are the detection of regularities (people who buy X also tend to buy Y), or the compression of data.

A.1.4 Reinforcement Learning

The premise of reinforcement learning is that an agent (e.g., a program, a robot, a control system) learns how to act successfully in an environment based on reward signals it receives in response to its actions. In each iteration, the agent can observe the state of the environment, decides how to act, and then receives a reward and information about the next state of the environment. Reinforcement learning is the core technology behind Google's AlphaZero, an algorithm that learned to beat the best human players in Go simply by playing against itself many times. These algorithms can also be used in finance to solve dynamic optimization problems including portfolio optimization and trading in the presence of transaction costs (Kolm and Ritter, 2020).

A.2 Overview of Common Artificial Intelligence Techniques

There are several artificial intelligence techniques that are widely used in asset management. These include artificial neural networks, cluster analysis, decision trees, evolutionary (genetic) algorithms, LASSO regression, support vector machines, and natural language processing. This section briefly characterizes these techniques, discusses their strengths and weaknesses, and areas of application (Table A.1).

Table A.1: Summary of Artificial Intelligence Techniques

The table summarizes the key characteristics of major artificial intelligence techniques and the branch of textual analysis approaches known as natural language processing. For each technique, the table details the key strengths, weaknesses, and areas of application.

Technique	Strengths	Weaknesses	Areas of Application
Artificial Neural Networks	<ul style="list-style-type: none"> - Complex and non- linear relationships - Incremental and transfer learning - Can generalize well 	<ul style="list-style-type: none"> - Data and computationally intensive - Predictions not explainable - Possible overfitting 	<ul style="list-style-type: none"> - Image processing and recognition - Speech recognition and synthesis - Forecasting
Cluster Analysis	<ul style="list-style-type: none"> - Labels unnecessary - Helps to understand data 	<ul style="list-style-type: none"> - Clusters may be intertwined - May require cluster count - Choosing attributes can be difficult 	<ul style="list-style-type: none"> - Data analysis - Anomaly detection - Recommendations
Decision Trees	<ul style="list-style-type: none"> - Classifications are explainable - Complex and non- linear relationships 	<ul style="list-style-type: none"> - Possible overfitting - Complex trees possible - Poor at predicting continuous variables 	<ul style="list-style-type: none"> - Decision making - Classification
Evolutionary (Genetic) Algorithms	<ul style="list-style-type: none"> - Ability to handle high-dimensional spaces - Finds novel solutions 	<ul style="list-style-type: none"> - Varies on initial conditions - Computationally intensive 	<ul style="list-style-type: none"> - Parameter optimization - Portfolio optimization
LASSO Regressions	<ul style="list-style-type: none"> - Identifies most relevant features - Flexible and fairly simple - Sparse solutions 	<ul style="list-style-type: none"> - Model can be unstable and hard to interpret - Performs poorly when independent variables are correlated 	<ul style="list-style-type: none"> - Forecasting and robust regression analysis
Natural Language Processing	<ul style="list-style-type: none"> - Analyze and generate text and speech - Find information in large textual datasets 	<ul style="list-style-type: none"> - Currently primitive and unable to fully understand text 	<ul style="list-style-type: none"> - Search engines and news filtering - Text classification and summarization
Support Vector Machines	<ul style="list-style-type: none"> - Structure of data can be unknown - Can generalize with less overfitting risk 	<ul style="list-style-type: none"> - Difficult to interpret - Kernel difficult to choose for non- linear classification 	<ul style="list-style-type: none"> - Classification - Regression

A.2.1 Least Absolute Shrinkage and Selection Operator Regression

Linear regression is a common and relatively simple way to fit a model to data in order to make predictions or estimate missing values. It seeks to find the coefficients of explanatory (or predictor) variables that contribute to the value of the dependent (or predicted) variable. To find the best model, the most common approach is to minimize the sum of squared errors, which are the difference between observed values and those predicted by the model. However, as model complexity increases with the number of regressors, the variability of predictions can also increase. Furthermore, with too many parameters, the model can overfit to the data, modelling noise rather than the underlying trend and leading to poor generalization to unseen data. The Least Absolute Shrinkage and Selection Operator (LASSO) method (Tibshirani, 1996; James et al., 2017) improves over standard linear or nonlinear models by additionally penalizing model complexity. By aiming to set some of the model's parameters to zero, LASSO regression automatically identifies the most relevant data features.


A.2.2 Artificial Neural Networks and Deep Learning

Artificial neural networks (ANNs) are inspired by biological brains (Aggarwal, 2018; Haykin, 2009). Similar to biological neurons, in each node of an ANN, the input signals are aggregated, processed, and the result forwarded to other nodes. ANNs are often arranged in feed-forward layers (Figure A.1), with input data applied to an input layer and further processed by a number of hidden layers before arriving at an output layer. ANNs with many hidden layers are called "Deep Neural Networks". ANNs are trained by altering the weights of the connections so that the errors between the predicted and desired data labels are minimized. Once trained, the ANN can be used to predict the output of previously unseen input data.

A fundamental concern among AI practitioners and researchers is the interpretability of trained ANNs. Presently, it is not easy to determine how a given neural network settles on its predictions. For some fields this is not an issue. However, in others, such as medical data analysis, doctors are reluctant to employ neural networks without a complete understanding of the mechanisms by which the network arrives at a prediction. ANNs also require fairly large datasets to train properly, and this may not be available for all assets or markets.

Figure A.1: Feed-Forward Artificial Neural Network

The figure illustrates a fully-connected feed-forward neural network. Output and input variables are linked through several layers of interconnected nodes.

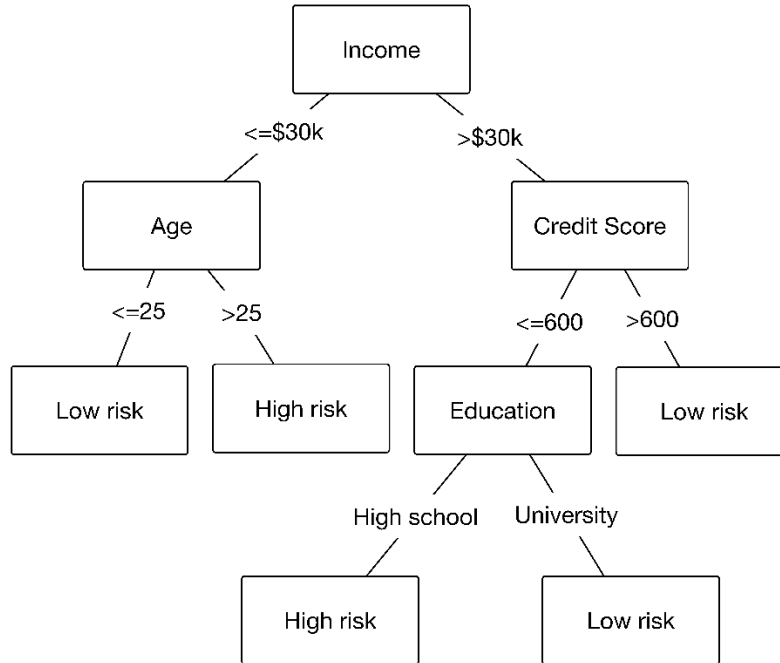


A.2.3 Decision Trees and Random Forests

A decision tree sequentially splits the data set into increasingly small subsets typically based on a single feature value. For the leaf nodes, the predicted class is determined (Figure A.2). Besides classification, decision trees can also be used for piecewise linear regression. Decision trees are a form of supervised learning, and the tree is constructed to replicate the labels in the training data as best possible.

Figure A.2: Decision Tree Classifier


A decision tree to classify credit applications into low and high risk. Starting from the top, applications move down the tree based on their characteristics. For example, an application of someone with an income less than or equal to \$30k and more than 25 years old would be classified as high risk.



A random forest (Breiman, 2001) is an ensemble of classification and regression methods that consists of many decision trees. Each decision tree makes predictions and contributes to the random forest's predictions via an averaging procedure. The assumption is that many decision tree models, each with a slightly different perspective, make better predictions than a single decision tree. Individual trees are generated from a distinct sampling of the training set while using a random subset of the attributes for nodes. Consensus voting on tree classifications determine the prediction (Figure A.3). A key advantage of decision trees over other AI techniques such as ANNs is that the rules according to which data is classified are human-readable and thus the reasons that lead to a particular classification can be easily traced.

Figure A.3: Random Forest Voting Scheme for Classification

The figure illustrates a random forest voting scheme that makes a prediction. The random forest combines the output of n decision trees and yields a final output that the majority of trees agree on.



A.2.4 Support Vector Machines

Support Vector Machines (SVMs) are supervised algorithms that are typically used for classification (Vapnik, 2000; James et al., 2017). These algorithms learn boundaries that partition the feature space into two or more classes. Once defined, the boundaries can be used to classify new data. SVMs are powerful, accurate tools for both classification and regression, and they are resistant to over-fitting their training data. However, they are computationally intensive and thus do not scale well to large data sets.

Figure A.4: Support Vector Machine Examples

The figure illustrates an example of using an SVM to separate data into two groups. The dotted lines are the support vectors. SVMs often transform data by adding one more dimension making it easier to classify data points. These transformations are called kernels. Panel (a) uses a linear kernel, while Panel (b) utilizes a non-linear one.

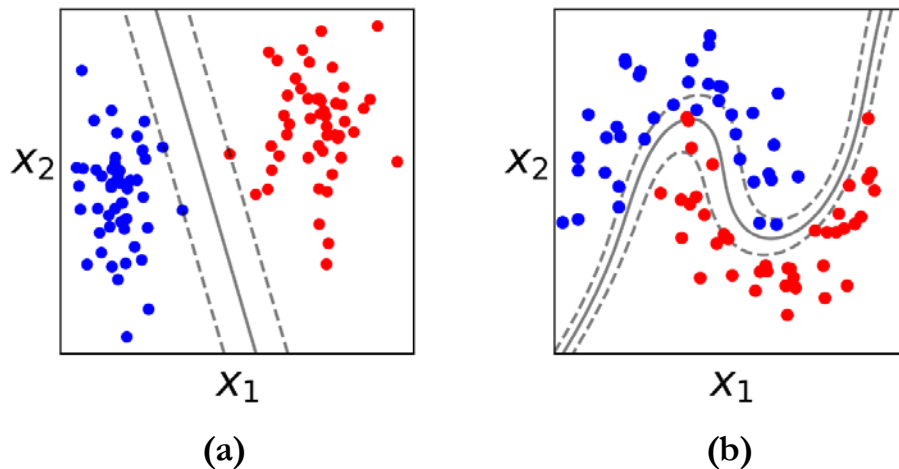


Figure A.4 (a) presents an example of linearly separable training data. The separation of the data into two classes is given by a line whose margin between the two classes of data points is as large as possible. In this example, the SVM has learned the linear boundary that most effectively separates the data. In practice, it is not always the case that training data are distributed in a way that permits a separation by a linear boundary. In such instances, kernel methods are employed to find more complex separation boundaries (Cortes and Vapnik, 1995). Figure A.4 (b) shows a non-linear boundary that was learned by an SVM with the radial basis function kernel.

A.2.5 Cluster Analysis

Cluster analysis is an unsupervised learning technique that seeks to partition data set into groups or clusters (Tan et al., 2018; Aggarwal and Reddy, 2014). Once the clusters are identified and labeled, the model can be used to classify new data. The same data may yield many different possible sets of clusters, depending on the number of clusters that are desired, how the data are distributed, and the clustering algorithm. Applications in asset management include cluster analysis of markets, companies, financial instruments, time series, and documents.

The most popular clustering algorithm is K-means clustering, which requires the user to specify the desired number of clusters, K . The method starts by randomly choosing some cluster centroids and allocating each data point to the closest centroid. It subsequently alternates between moving each centroid to the center of its data cluster and re-assigning the data points.

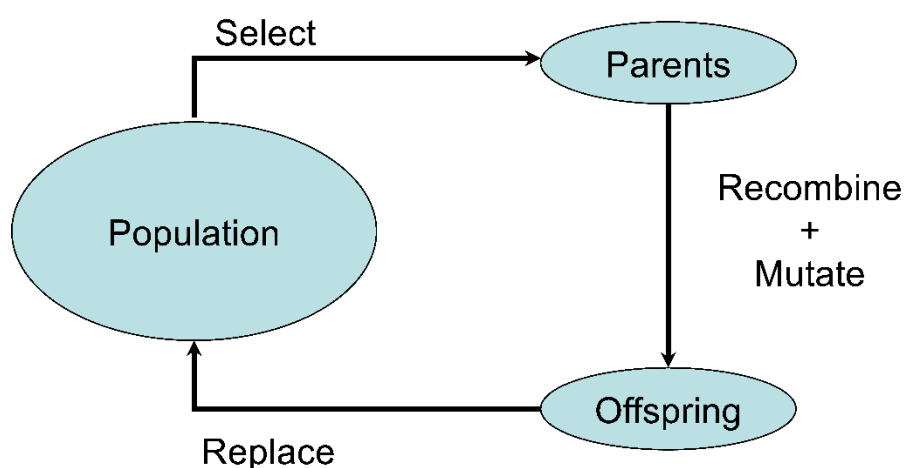
A.2.6 Evolutionary (Genetic) Algorithms

An evolutionary algorithm (often also called genetic algorithm) is an optimization algorithm based on Darwin's theory of evolution and natural selection (Eiben and Smith, 2015; Simon, 2013). An evolutionary algorithm starts with an initial population of candidate solutions, usually generated randomly. In each iteration, new candidate solutions are created by selecting a pair of good solutions in the population, merging the information from these two solutions into one (recombination) and introducing small random perturbations (mutation). The new solutions replace older, worse solutions in the population. An iterative process of varying existing good solutions and selectively keeping better ones, increasingly better solutions are found over time.

Because evolutionary algorithms do not require a mathematical formulation of the problem and do not make assumptions such as convexity or linearity about the objective function, they can also be applied to complex problems where other optimization algorithms fail. For example, evolutionary algorithms have been employed to solve mean-variance portfolio selection problems under cardinality constraints (restricting the number of assets in the portfolio).

Figure A.5: Iteration of an Evolutionary Algorithm

One iteration of an evolutionary algorithm: Good solutions are selected as parents, then used to create new solutions (offspring) by means of recombination and mutation, and finally the offspring replaces solutions in the population.



A.2.7 Natural Language Processing

Natural Language Processing (NLP) is a group of computational methods that can process or generate human natural language such as text or speech (Manning and Schütze, 1996; Mitkov, 2014). These methods include voice recognition, which converts spoken language to text; speech generation, which converts text to speech; natural language understanding, which extracts meaning from spoken or written text; and natural language generation, which produces natural language data from other data sources.

Owing to the ubiquity of social media and the countless natural language samples that it provides, the amount of literature on natural language processing has increased exponentially in recent years (Xing et al., 2018). Various support vector machines and neural network architectures such as deep learning or recurrent networks are frequently used in natural language processing.

In finance, much attention is focused on natural language financial forecasting, which extracts information such as sentiment from financial news or social media data and incorporates it into models that predict the movement of financial data. Most of the literature applies NLP to stock and foreign exchange rate prediction. This is due to the accessibility of information for these markets.

A.2.8 Comparisons of Artificial Intelligence Techniques

AI algorithms and techniques differ by the types of data on which they best operate, the kinds of predictions that they can make, the means by which they learn to fit their data, the computational power required for training, testing, and deployment, and the ease with which the algorithm can be scaled.

Data may be low-dimensional, as in the case of housing price data that depend on less than a dozen features, or high-dimensional, as in the case of an image consisting of millions of pixels that have values independent from all others. Data that reside in just a few dimensions allow using simpler techniques (LASSO, K-means clustering, K-nearest neighbors) for classification or regression. On the other hand, data that live in higher dimensions suffer from the curse of dimensionality: the more dimensions one has, the more samples one must obtain in order to make meaningful predictions, and the more computationally intensive it becomes to analyze and model the data. Autoencoder neural networks are an example of an algorithm that can project high-dimensional data into low-dimensions, so that the data can be modeled and predicted more effectively.

If the data are labeled, a wide variety of supervised learning algorithms can be used to learn a good mapping from input to predicted output. If data are unlabeled, unsupervised learning techniques can be employed, such as cluster analysis, to identify patterns in the data.

Some algorithms are versatile, whereas others have a narrow focus. For instance, K-means clustering is an algorithm that is designed to partition data into K subsets in an unsupervised fashion only. Neural networks, on the other hand, are more broadly applicable. They can be used for unsupervised and supervised learning, classification and regression, and on images, text, and time-series.

The interpretability of a mathematical model varies widely. Linear regression models, logistic regression models, decision trees and K-nearest neighbor classifiers are examples of algorithms with readily understood

learned behavior. In contrast, a precise statistical description of the learning and prediction decisions of neural networks is still an unresolved problem.

Some algorithms are difficult to deploy on large datasets due to their computational complexity. Specialized hardware or graphic cards may speed up computations, such as for neural networks, for example. Note that while training is sometimes very time-consuming, deploying a trained model for predictions on unseen data is usually fast.

The business logic of asset management determines the structure of data sets and the type of predictions to be obtained. Consequently, the business needs ultimately determine the choice of artificial intelligence algorithm.

Appendix B: Trends and Patterns in Finance Research Using Artificial Intelligence

Developments in academic research related to artificial intelligence can be analyzed by studying trends and patterns in working papers in finance that employ AI techniques. To capture the most recent trends and those related to the area of finance, we download all working papers that have AI-related keywords in their title, abstract, or listed keywords and that have been posted in the Financial Economics Network (FEN) on the Social Sciences Research Network (SSRN) between 1996 and 2018. The primary keywords are: “artificial intelligence”, “machine learning”, “cluster analysis”, “genetic algorithm” or “evolutionary algorithm”, “lasso”, “natural language processing”, “neural network” or “deep learning”, “random forest” or “decision tree”, and “support vector machine”.

There are 1,814 working papers on SSRN with at least one of our keywords. “Machine learning” is by far the most popular AI-related keyword and is used in 29% of all papers in our sample. Among the AI techniques, “neural network” (or “deep learning”) with 38% and “cluster analysis” with 16% of the papers are the two most popular keywords (Panel A of Figure B.1). The numbers of downloads magnify the relative popularity of neural networks even further considering that papers with “neural network/deep learning” keywords account for almost half of the total downloads of all papers related to the AI techniques (Panel B of Figure B.1). “Random forest” (or “decision tree”) with only 16% has the second highest proportion of downloads. Finally, “natural language processing” is associated with only 4% of the papers and 3% of downloads, which is the lowest among all techniques.

Interestingly, there are no working papers with a “machine learning” keyword until 2003. Since the terms “machine learning” and “artificial intelligence” have gained popularity more recently, we sum the number of papers for all AI-related keywords. In 1996, no working paper with any AI-related keyword was uploaded to SSRN. However, there is significant growth with 410 papers posted in 2018. The total number of papers on FEN also increases over the same period, but to a much smaller degree: Papers with AI-related keywords accounted for 3% of all papers in 2018.

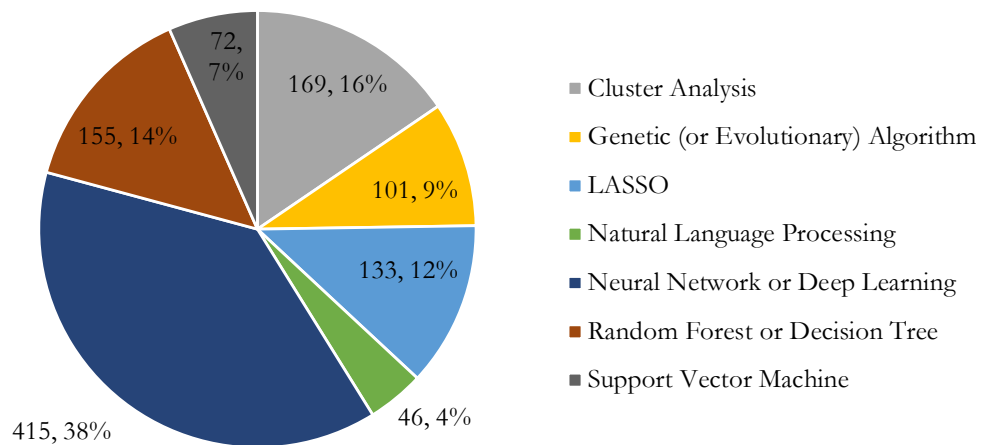
Among the AI techniques, the popularity of “neural network” is consistent over the years. However, “lasso” and “support vector machine” techniques seem to have gained popularity more recently. The same

applies to “natural language processing”: There are no papers with this keyword on SSRN until 2008, while 13 working papers with this keyword were uploaded in 2018.


Figure B.1: Number and Downloads of Finance Papers Using Artificial Intelligence

The figure shows the number of papers (Panel A) and downloads (Panel B) associated with various AI techniques for the sample of SSRN Financial Economics Network (FEN) working papers. The sample starts in 1996, ends in 2018, and includes papers having one of the following keywords in their abstract, title, or keyword section: “artificial intelligence”, “machine learning”, “cluster analysis”, “genetic algorithm” or “evolutionary algorithm”, “lasso”, “natural language processing”, “neural network” or “deep learning”, “random forest” or “decision tree”, and “support vector machine” .

Panel A: Number of Papers



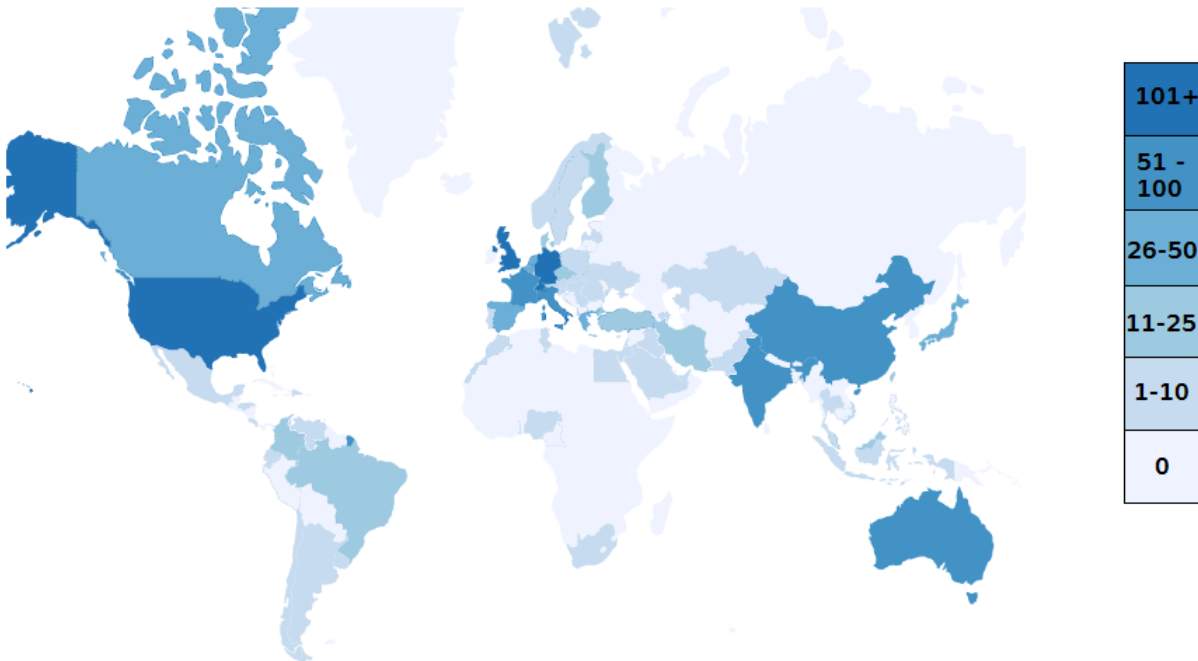
Panel B: Number of Downloads



Across countries, the United States, the United Kingdom, and Germany are the top three producers of papers with AI-related keywords (Figure B.2). These three countries account for more than half of all AI papers, followed by Switzerland, India, France, China, and Italy. Not surprisingly, these are also countries that show high productivity in terms of finance papers generally.

Figure B.2: Numbers of Papers Using Artificial Intelligence by Country

The figure shows a heat map based on the numbers of papers with AI-related keywords for each country. A paper belongs to a specific country if at least one of its authors is affiliated with institutions located in that country. The sample starts in 1996, ends in 2018, and includes papers having one of the following keywords in their abstract, title, or keyword section: “artificial intelligence”, “machine learning”, “cluster analysis”, “genetic algorithm” or “evolutionary algorithm”, “lasso”, “natural language processing”, “neural network” or “deep learning”, “random forest” or “decision tree”, and “support vector machine”.



Across institutions, Cornell University, Humboldt University of Berlin, University of Chicago, Stevens Institute of Technology, and ETH Zurich are the top 5 contributors to AI research papers (Table B.1). There is also a large number of non-academic institutions in the full list illustrating the interest in conducting research related to this topic in the finance industry and central banks.

Table B.1: Numbers of Papers by Institution

The table presents the number of SSRN Financial Economics Network (FEN) working papers by institution. A paper belongs to an institution if at least one of its authors has an affiliation with that institution. The sample starts in 1996, ends in 2018, and includes papers having one of the following keywords in their abstract, title, or keyword section: “artificial intelligence”, “machine learning”, “cluster analysis”, “genetic algorithm” or “evolutionary algorithm”, “lasso”, “natural language processing”, “neural network” or “deep learning”, “random forest” or “decision tree”, and “support vector machine”. Institutions with less than 15 papers are combined and reported as “Other Institutions”.

Institution / University	Sum of AI Papers
Independent	75
Affiliation not provided to SSRN	67
Cornell University	64
Humboldt University of Berlin	47
University of Chicago	42
Stevens Institute of Technology	41
ETH Zurich	39
New York University	30
Imperial College London	26
Harvard University	22
University of St. Thomas	21
London School of Economics & Political Science	19
Other Institutions	2937

References

Abe, Masaya, Hideki Nakayama, 2018, Deep learning for forecasting stock returns in the cross-section, *Advances in Knowledge Discovery and Data Mining*, 594–621 (Springer).

This paper examines the performance of deep learning for one-month ahead forecasting of stock returns in the cross-section in the Japanese stock market. It is found that deep neural networks (DNN) perform better than shallow neural networks and other ML techniques, such as Support Vector Regression (SVR) and Random Forests (RF). An ensemble of SVR, RF and DNN can provide a further small gain in performance.

Aggarwal, Charu C., 2018, *Neural Networks and Deep Learning* (Springer).

This textbook, aimed at graduate students, researchers, and practitioners, is devoted to studying neural networks or deep learning as it became known in machine learning literature. It starts with background material on neural network design and its relationship to traditional machine learning algorithms. The rest of the textbook provides a more detailed discussion of the fundamentals of neural networks (common neural architectures, shallow and deep neural networks, radial basis function networks, restricted Boltzmann machines) and the recent advances in the field (recurrent and convolutional neural networks, deep reinforcement learning, neural Turing machines, generative adversarial networks).

Aggarwal, Charu C., and Chandan K. Reddy, 2014, *Data Clustering: Algorithms and Applications* (Boca Raton: Taylor and Francis).

Data clustering is partitioning the data into multiple groups where each observation shares some characteristics with other data points within its group. Data clustering issues have been studied in different fields and this book attempts to bridge the gap by bringing the knowledge gained in machine learning and data mining literature together. The first part of the book discusses data clustering methods (probabilistic, distance-based, density-based, grid-based, and spectral clustering, clustering high-dimensional and big data). The second part is devoted to the application of data clustering in different fields and situations (clustering categorical, text, multimedia, time-series, biological, network, and uncertain data). Finally, the book focuses on the appropriateness of a particular clustering method through evaluation techniques available from different variations of cluster analysis (visual, supervised, ensemble and validation-based solutions).

Ahmed, Nesreen K., Amir F. Atiya, Neamat El Gayar, and Hisham El-Shishiny, 2010, An empirical comparison of machine learning models for time series forecasting, *Econometric Reviews* 29, 594–621.

The article conducts a comparative analysis of forecasting performance of several machine learning methods (multilayer perceptron, Bayesian neural networks, radial basis functions, generalized regression neural networks, K-nearest neighbor regression, CART regression trees, support vector regression, and Gaussian process) for time series data. Using M3 monthly time series data, the results show that multilayer perceptron and the Gaussian process regression outperform the rest of the machine learning methods considered. This indicates that the rarely used Gaussian process method may offer performance improvement in forecasting problems.

Ahn, Jae Joon, Kyong Joo Oh, Tae Yoon Kim, and Dong Ha Kim, 2011, Usefulness of support vector machine to develop an early warning system for financial crisis, *Expert Systems With Applications* 38, 2966–2973.

Existing studies on early warning system (EWS) for financial crises are based on a belief system that considers financial crises as self-fulfilling due to investor herding behavior. In this article, the authors take the alternative view that banking crises are a result of deteriorating long-term economic fundamentals. With this new hypothesis, the authors find that support vector machine is an effective technique for this classification problem.

Alberg, John, and Zachary C. Lipton, 2018, Improving factor-based quantitative investing by forecasting company fundamentals, Available at arXiv 1711.04837.

Fundamental financial data of publicly traded companies, such as operating income, revenue, assets, dividends, and debt, provide a picture of financial performance of companies. If the future fundamentals can be forecasted, then an investment strategy can build portfolios based on these predicted fundamentals that can outperform a standard factor approach. Using data on stocks for publicly traded US companies, the study develops and forecasts features from the fundamentals, and evaluates the out-of-sample forecasting performance using two classes of deep neural networks (multilayer perceptrons and recurrent neural network). This lookahead factor model with deep neural networks significantly outperforms the naive prediction (e.g., fundamentals remain unchanged). In addition, a simulated hypothetical stock portfolio using an industry-grade stock portfolio simulator shows a performance improvement over traditional factor models in terms of compounded annual returns.

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

In this paper, the authors examine the S&P 500 composite index from 1928 to 1995 to generate trading rules using genetic algorithms. Similar to past findings, trading rules created by genetic algorithms do not make money after controlling for transaction costs. Genetic algorithm-based trading rules can correctly identify periods of high return and low volatility and periods with opposite conditions.

Alpaydin, Ethem, 2010, *Introduction to Machine Learning*, second edition (The MIT Press).

This book is an introductory textbook on machine learning. It covers relevant algorithms and methods from computer science, neural computation, and statistics. The topics include, but are not limited to, supervised learning, Bayesian decision theory and Bayesian estimation, parametric and nonparametric methods, multivariate methods, dimensionality reduction, decision trees, linear discrimination, kernel machines, hidden Markov models, graphical models, combining multiple learners, reinforcement learning, design and analysis of machine learning experiments.

Arrieta-ibarra, Imanol, and Ignacio N. Lobato, 2015, Testing for predictability in financial returns using statistical learning procedures, *Journal of Time Series Analysis* 36, 672–686.

The study examines and compares the out-of-sample forecasting performance of three machine learning techniques (regression trees, neural networks, and support vector machines) to that of a traditional GARCH(1,1) model in predicting daily stock market returns. The authors conclude that, in terms of predictive accuracy, the three machine learning techniques do not outperform naive historical sample average returns when only past values of the series are considered. Adding other variables (interest rates, commodity prices or exchange rates) to the model may lead to model overfitting and insignificant results. However, support vector machines and neural networks offer potential improvements over a GARCH(1,1) model only for forecasting squared returns.

Atsalakis, George S., and Kimon P. Valavanis, 2009, Surveying stock market forecasting techniques – part ii: Soft computing methods, *Expert Systems with Applications* 36, 5932–5941.

This article provides a survey of the literature on stock return forecasting in different markets. It reveals that soft computing techniques have increasingly been used for predicting stock market behavior recently. Neural networks, fuzzy logic, and genetic algorithm are some examples of soft computing techniques. In the prior literature on stock return forecasting, soft computing techniques generally outperform traditional models in terms of forecast accuracy.

Auria, Laura, and Rouslan A. Moro, 2008, Support vector machines (SVM) as a technique for solvency analysis, Available at SSRN 1424949.

The authors propose the use of support vector machines (SVM) in predicting company solvency. Using data on German companies, the SVM technique outperforms alternative approaches, such as logistic regression and discriminant analysis, in out-of-sample accuracy.

Avramov, Doron, Si Cheng, Lior Metzker, 2019, Machine learning versus economic restrictions: Evidence from stock return predictability, Available at SSRN 3450322.

In this article, the researchers argue that ML methods, such as deep learning, often don't clear economic restrictions like value-weighting returns, excluding microcaps or distressed firms. When economic restrictions are imposed, value-weighted portfolio returns decline relative to equal-weighted portfolio returns. Upon further economic restrictions, the value-weighted portfolio returns attenuate further once microcaps, non-rated firms and distressed firms are excluded. In the presence of transaction costs, performance may further deteriorate. Nevertheless, ML methods are useful, especially in long positions and due to low downside risk.

Azimi, Mehran, and Anup Agrawal, 2019, Is positive sentiment in corporate annual reports informative? Evidence from deep learning, Available at SSRN 3258821.

The study uses deep learning for text classification of qualitative information found in corporate annual reports. The research findings indicate that, unlike prior literature, both positive and negative sentiments are important in predicting abnormal returns and abnormal trading volume around 10-K filing date. For instance, positive (negative) sentiment is associated with higher (lower) abnormal returns and lower (higher) abnormal trading volume around 10-K filing date.

Aziz, Saqib, and Michael Dowling, 2019, Machine learning and AI for risk management, *Disrupting Finance*, 33–50 (Springer).

This paper provides an overview of the application of machine learning and AI techniques in the risk management domain. It reviews the relevant literature and discusses the particular techniques used in the areas of credit risk, market risk, operational risk, and compliance.

Ballings, Michel, Dirk Van den Poel, Nathalie Hespels, and Ruben Gryp, 2015, Evaluating multiple classifiers for stock price direction prediction, *Expert Systems With Applications* 42, 7046–7056.

In this study, ensemble machine learning methods (Random Forest, AdaBoost, Kernel Factory) are compared to single classifiers (Neural Networks, Logistic Regression, Support Vector Machines (SVM), K-Nearest Neighbor) in their ability to predict the direction of stock price movements. Based on AUC (areas under the receiver operating characteristic curve) as a performance measure, the algorithms are ranked as follows: Random Forests, SVM, Kernel Factory, AdaBoost, Neural Networks, K-Nearest Neighbor, and Logistic Regression. The study concludes that ensemble models should not be neglected in predicting stock price direction.

Bao, Yang, and Anindya Datta, 2014, Simultaneously discovering and quantifying risk types from textual risk disclosures, *Management Science* 60, 1371–1391.

This study uses a latent dirichlet allocation topic model to simultaneously identify and measure the risk types from risk disclosure information in 10-K forms. Empirical results indicate that the presented textual analysis method outperforms all competing methods. Further, the proposed technique identifies that only one third of risk types are informative. Among the risk types that are found to be informative, only systematic and liquidity-related risks will increase investors' risk perceptions, while nonsystematic risk types will actually decrease them.

Becker, Sebastian, Patrick Cheridito, and Arnulf Jentzen, 2019a, Deep optimal stopping, *Journal of Machine Learning Research* 20, 1–25.

In this article, researchers develop an optimal stopping strategy using deep learning. They construct lower and upper bounds, point estimates, and confidence intervals for the price. The approach is then

tested in three different examples: the pricing of Bermudian max-call option, the pricing of a callable multi-barrier reverse convertible, and the problem of optimally stopping a fractional Brownian motion. The results show that in all three exercises the proposed method produces highly accurate results with short computing time.

Becker, Sebastian, Patrick Cheridito, and Arnulf Jentzen, 2019b, Pricing and hedging American-style options with deep learning, Available at arXiv 1912.11060.

In this article, the authors consider American-style options and develop a deep neural networks for optimal exercise behavior, pricing, and hedging. In the first stage, the neural network method is used to find the optimal stopping rule and derive a lower bound estimate of the price. In the next stage, it estimates an upper bound and confidence intervals for the price. Finally, the deep learning method constructs a dynamic hedging strategy. Exercises based on Bermudian max-call option show that the proposed method can achieve accurate price predictions and provide dynamic hedging strategies.

Beketov, Mikhail, Kevin Lehmann, and Manuel Wittke, 2018, Robo-advisors: quantitative methods inside the robots, *Journal of Asset Management* 19, 363–370.

This study explores the methods used in Robo-advisors around the world and finds modern portfolio theory as the predominant method used in the Robo-advisory systems. Further, the volume of assets managed by Robo-advisors is likely to be higher when they use more sophisticated systems.

Bellovary, Jodi L., Don E. Giacomino, and Michael D. Akers, 2007, A review of bankruptcy prediction studies: 1930 to present, *Journal of Financial Education* 1–42.

This study reviews the bankruptcy prediction literature published since the 1930s. Bankruptcy prediction studies in earlier years primarily utilized simple ratio analysis, while discriminant analysis has become more prevalent in the 1960s and 1970s. In the 1980s and 1990s, logit and neural networks methods gain more popularity in bankruptcy analysis. Comparisons of different methods reveal that multivariate discriminant analysis and neural networks deliver better bankruptcy prediction models. Additionally, the authors note that a model's accuracy may not always increase with the number of factors used in the model.

Bertsimas, Dimitris, and Andrew W. Lo, 1998, Optimal control of execution costs, *Journal of Financial Markets* 1, 1–50.

The performance of investments is affected by the presence of trading or execution costs, which include commissions, bid/ask spreads, opportunity costs. In this article, the authors present a dynamic trading strategy that seeks to minimize expected trading costs. The cost minimizing strategy is a function of time and several state variables so it can optimally adapt to changing market conditions and price movements. Compared to a naive trading strategy, the best-execution strategy reduces execution cost by 25% to 40%.

Berutich, Jose Manuel, Francisco Lopez, Francisco Luna, and David Quintana, 2016, Robust technical trading strategies using GP for algorithmic portfolio selection, *Expert Systems With Applications* 46, 307–315.

This study examines a genetic algorithm approach for generating profitable trading rules in the Spanish market. The main contribution of the paper is introducing a new random sampling method that improves the out-of-sample results of the genetic algorithm compared to traditional approaches. This method can also be generalized to cope with different types of markets. Empirical tests of the strategies generated using this approach confirm the superiority of the approach over a range of traditional alternatives.

Bew, David, Campbell R. Harvey, Anthony Ledford, Sam Radnor, and Andrew Sinclair, 2019, Modeling analysts' recommendations via Bayesian machine learning, *The Journal of Financial Data Science* 1, 75–98.

Individual analysts focus on only a small number of thousands of companies in making recommendation on stock performance. This paper proposes a novel approach for combining individual analyst's recommendations using Bayesian classifier combination (IBCC) that dynamically adjusts based on the length and quality of analyst's track record. The forecasts are obtained from the probabilistic IBCC model using data from the pan-European region from 2004 to 2013. The results indicate that the best outcome is achieved when analyst/broker recommendation agrees with forecasts from the IBCC model.

Bholat, David M, Stephen Hansen, Pedro Santos, and Cheryl Schonhardt-Bailey, 2015, Text mining for central banks, Available at SSRN 2624811.

Text mining can be used to extract meaning from texts. This handbook illustrates the value of text mining methods to central banks. In the first part, the handbook reviews the text mining literature and how it could be applied to central bank policymaking. In the second part, it explains the implementation of popular text mining techniques, such as Boolean and dictionary text mining, Latent Semantic Analysis, Latent Dirichlet Allocation, and Descending Hierarchical Classification.

Bianchi, Daniele, Matthias Büchner, and Andrea Tamoni, 2019, Bond risk premia with machine learning, Available at SSRN 3400941.

The paper compares performance of machine learning (ML) methods in predicting bond returns. Compared to classical approaches, such as dense and sparse linear regression frameworks, deep neural networks perform significantly better in forecasting excess bond returns. Empirical exercises using both one-year-ahead and one-month-ahead forecasts show that neural networks outperform other classical approaches largely due to their ability to capture non-linearities in the data.

Black, Fischer, and Robert Litterman, 1992, Global portfolio optimization, *Financial Analysts Journal*, 28–43.

In this article, the authors discuss the limitations of the modern portfolio management theory in global portfolio management. The main limitation of classic portfolio management theory is the difficulty of its implementation and unexpected poor behavior when implemented for global portfolios. The authors suggest combining Markowitz's mean-variance portfolio optimization model with the capital asset pricing model (CAPM) of Sharpe and Lintner in order to improve the portfolio allocation behavior.

Board of Governors of the Federal Reserve System, 2011, Supervisory guidance on model risk management, Available at: <https://www.occ.treas.gov/news-issuances/bulletins/2011/bulletin-2011-12a.pdf>.

Financial institutions have been using models in many aspects of banking. However, models also carry risks and costs. This paper presents an overview of model risk management, robust model development, model validation, and good practices in model development, implementation, and application.

Booth, Ash, Enrico Gerding, and Frank McGroarty, 2014, Automated trading with performance weighted random forests and seasonality, *Expert Systems With Applications* 41, 3651–3661.

This study presents an automated trading system based on performance weighted ensembles of random forests. The proposed method uses seasonal trends to improve the predictability of stock returns. Using data on the German Stock Index DAX, the proposed new technique shows promising results for out-of-sample forecasts. Specifically, it achieves better prediction accuracy as well as higher profitability compared to other ensemble methods.

Booth, Ash, Enrico Gerding, and Frank McGroarty, 2015, Performance-weighted ensembles of random forests for predicting price impact, *Quantitative Finance* 15, 1823–1835.

In this study, the researchers present a forecasting model based on an ensemble of random forests. The empirical results show that the proposed model outperforms other popular competing models –

liner regression, neural networks, and support vector regression – in out-of-sample forecasts. The ensemble of random forests improve forecast accuracy by at least 15% compared to the competing models.

Borghini, Riccardo, and Giuliano De Rossi, forthcoming, *The Artificial Intelligence Approach to Picking Stocks*, in *Machine Learning and Asset Management* (Springer).

The chapter explores machine learning (ML) techniques to select from a large set of firm characteristics to predict one-month-ahead stock returns using data from the US and European stocks. It uses a proprietary alpha model, a global stock selection model from Macquarie, as a benchmark for comparing the results from traditional machine learning techniques (OLS, GLM, LASSO, ridge, elastic net, neural networks, boosted trees and random forest, regression trees). The results indicate that an ensemble of ML methods could outperform even a proprietary model with a successful track record, such as Macquarie's Alpha Quant model, in terms of generating attractive risk-return ratios. However, ML tend to be less exposed to price momentum than an Alpha Model.

Branke, Juergen, Benedikt Scheckenbach, Michael Stein, Kalyanmoy Deb, and Hartmut Schmeck, 2009, Portfolio optimization with an envelope-based multi-objective evolutionary algorithm, *European Journal of Operational Research* 199, 684–693.

The mean-variance-based portfolio optimization problem with linear constraints can be solved efficiently using parametric quadratic programming. However, many real world constraints tend to be more complex and discrete, for example if the number of stocks in the portfolio is limited. This study combines parametric quadratic programming with a multi-objective evolutionary algorithm (MOEA) to find the best solution to the non-convex portfolio selection. As the results demonstrate, the proposed MOEA model significantly outperforms other existing evolutionary algorithms.

Breiman, Leo, 2001, Random forests, *Machine Learning* 45(1), 5–32.

In this seminal work, the author introduces Random Forests (RF). RF is a collection of tree predictors used in a classification problem. As the number of trees increase and due to law of large numbers, RF do not suffer from overfitting problems and are an effective forecasting tool used in machine learning.

Brenner, Lukas, and Tobias Meyll, 2019, Robo-Advisors: A substitute for human financial advice?, Available at SSRN 3414200.

With recent technological advances in financial services industry, the use of automated financial advisors, i.e. robo-advisors, has been increasing. The paper investigates the determinants of investor decisions to opt for robo-advisory services. Using a representative investor survey data, the paper shows that a significant substitution is taking place as investors are worried about potential conflicts of interest by human financial advisors.

Briere, Marie, Charles-Albert Lehalle, Tamara Nefedova, and Amine Raboun, 2019, Modelling transaction costs when trades may be crowded: A Bayesian network using partially observable orders imbalance, Available at SSRN 3420665.

This paper uses a Bayesian network model to forecast transaction costs measured as implementation shortfall. The Bayesian network approach can account for missing data by using the most probable values instead. This feature is particularly useful for integrating important latent variables such as the net order flow imbalance of investors to improve the transaction cost forecasts. The findings suggest that implementation shortfall forecasts are more accurate when the investors' order size is larger or when the stock volatility is lower.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2019, Forest through the trees: Building cross-sections of stock returns, Available at SSRN 3493458.

The paper describes a new method of building basis assets that capture the underlying information in the cross-sections of asset returns. The novel approach allows to capture the complex information of cross-sectional stock characteristics. The proposed method applies Asset-Pricing Trees to build portfolios that capture all relevant information, allow for non-linearities and interactions, act as building blocks for a stochastic discount factor, and provide test assets for asset pricing. The empirical results using monthly equity returns for all CRSP securities show that, unlike the proposed approach, conventional cross-sectional sorting methods fail to fully capture the relevant information on stock characteristics.

Buchanan, Bonnie, 2019, Artificial intelligence in finance, Available at: <http://doi.org/10.5281/zenodo.2612537>.

This literature review overviews AI, machine learning (ML) and deep learning (DL), and their applications in the financial services sector. The review first discusses how rapidly growing AI is changing financial services industry by using ML algorithms to analyze millions of data points in detecting fraudulent activity, by utilizing "robo-advising" that avoids human interaction and conflict of interest, and by using algorithmic trading to make fast trading decisions. The review then discusses differences between ML and econometrics and their applications. Specifically, ML and DL methods focus on predictive accuracy, while econometric methods are mainly concerned with inferential questions. ML and DL techniques have become powerful tools for out-of-sample forecasts and to identify the most useful predictors. Further, the article briefly reviews quantum computing and how it can process data at speeds impossible for traditional computers. Finally, as AI and ML are increasingly applied in the financial services industry and moving very fast, regulators are finding it difficult to keep pace with the new technology. The review then discusses the experience of some countries around the world.

Celik, Arzum Erken, and Yalcin Karatepe, 2007, Evaluating and forecasting banking crises through neural network models: An application for Turkish banking sector, *Expert Systems with Applications* 33, 809–815.

The authors of this study forecast the ratios of non-performing loans to total loans, capital to assets, profits to assets, and equity to assets for the Turkish banking sector. The results indicate that artificial neural networks (ANNs) with Taguchi's approach to determine the optimal values of ANN parameters are an effective tool in forecasting banking crises.

Chapados, Nicolas, and Yoshua Bengio, 2001, Cost functions and model combination for var-based asset allocation using neural networks, *IEEE Transactions on Neural Networks* 12, 890–906.

This study considers a value-at-risk (VaR), typically used in quantifying the risk associated with a portfolio, and extends it to the mean-VaR framework, similar in spirit to Markowitz' famous mean-variance model. Mean-VaR seeks to find the best portfolio allocation for a given VaR level. The authors build neural networks to optimize the portfolio of assets given the VaR constraint and transaction costs and show that its performance is comparable to that of a forecasting model based on the classical mean-variance portfolio management, but better than the performance of the benchmark market index. However, the proposed method relies on fewer model assumptions and takes into account transaction costs. Finally, both the proposed and the classical forecasting models can significantly outperform the performance of the benchmark market index when a forecast combination technique is applied.

Chen, Luyang, Markus Pelger, and Jason Zhu, 2019, Deep learning in asset pricing, Available at SSRN 3350138.

The article utilizes deep neural networks and no-arbitrage constraint to estimate individual stock returns. Specifically, it combines three different neural networks (feedwork network, LSTM network, generative adversarial network) with no-arbitrage condition to estimate the stochastic discount factor (SDF) that explains stock returns. The SDF factor portfolio uses time-varying weights for traded assets which are functions of firm-specific and macroeconomic variables. The proposed SDF factor asset pricing model outperforms other benchmark models out-of-sample.

Chen, Shiyi, Wolfgang K. Hardle, and Kiho Jeong, 2010, Forecasting volatility with support vector machine-based GARCH model, *Journal of Forecasting* 29, 406–433.

This study compares the volatility forecasts generated from the traditional volatility forecasting methods (moving average, GARCH, EGARCH, ANNGARCH) with the forecasts obtained using a support vector machine (SVM) or artificial neural network. Using a recursive forecasting framework, one-step ahead forecasts of the volatilities of the British exchange rate and NYSE index have been evaluated based on Diebold-Mariano forecast evaluation method. Out of all the models considered, SVM-GARCH models perform well in many cases, while a standard GARCH model performs well when the sample size is large and normally distributed. On the other hand, an EGARCH model better predicts the one-step ahead volatility when the data is highly skewed.

Chen, Wun-Hua, Jen-Ying Shih, and Soushan Wu, 2006, Comparison of support-vector machines and back propagation neural networks in forecasting the six major Asian stock markets, *International Journal of Electronic Finance* 1, 49–67.

Most of the existing research in the application of data mining techniques to financial time series forecasting focuses on the U.S. or European markets. This study examines the performance of Support Vector Machines (SVM) and back propagation neural networks in forecasting six Asian stock markets. Compared to naive AR(1) time series models, both of the proposed models perform well.

Cheng, Ching-Hsue, Tai-Liang Chen, and Liang-Ying Wei, 2010, A hybrid model based on rough sets theory and genetic algorithms for stock price forecasting, *Information Sciences* 180, 1610–1629.

Time series forecasting methods and artificial intelligence techniques used in forecasting stock returns are not easily understood by professional fund managers. Therefore, they often rely on subjective judgments in predicting stock prices based on some technical indicators. This paper develops a hybrid forecasting framework that utilizes the tools of rough set theory (RST) and genetic algorithms (GA). Empirical results show that the proposed hybrid techniques outperform the RST and GA in terms of forecast accuracy and profits.

Chinco, Alexander M., Adam D. Clark-Joseph, and Mao Ye, 2019, Sparse signals in the cross-section of returns, *The Journal of Finance* 74, 449–492.

Using the least absolute shrinkage and selection operator (LASSO), a penalized regression technique, the authors study the out-of-sample fit of one-minute ahead return forecasts. Evidence suggests that this improved performance is mainly due to identification of predictors that are unexpected, short-lived, and sparse. Further, the LASSO method is found to increase forecast-implied Sharpe ratio.

Choudhry, Taufiq, Frank McGroarty, Ke Peng, and Shiyun Wang, 2012, High-frequency exchange-rate prediction with an artificial neural network, *Intelligent Systems in Accounting, Finance and Management* 19, 170–178.

The authors study the effectiveness of artificial neural networks (ANN) for forecasting high-frequency exchange rates, ranging from one to a few minutes. They find that bid and ask prices are significant variables for exchange rate forecasting. The study concludes that high-frequency trading strategies using ANNs can generate positive profit beyond transaction costs.

Cong, Lin, Tengyuan Liang, and Xiao Zhang, 2019, Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information, Available at SSRN 3307057.

Increased access to unstructured data in the form of texts can be used to complement information obtained from the traditional structured data. However, extracting insight from textual data is more complex due to intricate language structure, high-dimensionality and lack of data-driven approaches to analyzing textual data. The authors develop an alternative method of analyzing textual data that is more effective than the existing methods. They use neural network models for natural language processing

to generate textual factors and show the procedure's application in analyzing financial and macroeconomic data.

Coqueret, Guillaume, and Tony Guida, 2018, Stock returns and the cross-section of characteristics: A tree-based approach, Available at SSRN 3169773.

This study uses regression trees that iteratively splits the sample into clusters to investigate the impact of 30 classical firm characteristics on future returns. Unlike linear regression models, tree-based regression models take into account conditional impact of firm attributes given the value of other attributes. The findings indicate that both technical indicators (for example, Relative Strength Index-RSI) and past performance indicators are important in stock pricing. Authors estimate that a portfolio built using a short-term RSI characteristic leads to an annual gain of 2.4% relative to a naive, equally-weighted portfolio.

Cortes, Corinna, and Vladimir Vapnik, 1995, Support-vector networks, *Machine Learning* 20(3), 273–297.

The support-vector network is a machine learning algorithm used in classification problems with two groups. This study generalizes the support-vector networks for separable training data and extends the algorithm to non-separable training data.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto Rossi, 2017, The promises and pitfalls of robo-advising, Available at SSRN 3122577.

This study investigates the robo-advising on the portfolio management of investors. Investors who use robo-advising tend to have more assets, trade more and have better risk-adjusted performance. The results show that among investors who adopt robo-advising diversification and volatility of stocks decrease when the number of stocks held are less than 5. However, the same is not true when the number of stocks initially held are more than 10. Thus, the study illustrates the possible heterogeneous impacts of robo-advising on investors' portfolio performance.

Das, Sanjiv Ranjan, 2014, Text and context: Language analytics in finance, *Foundations and Trends in Finance* 8, 145–261.

This monograph provides a comprehensive survey of techniques and applications of text analytics in the field of finance. Specifically, it discusses text extraction steps using the statistical package R, how to analyze the extracted text, classification of texts and words, and application of text-data in empirical studies.

De Prado, Marcos Lopez, 2016, Building diversified portfolios that outperform out of sample, *Journal of Portfolio Management* 42, 59.

The critical Line Algorithm (CLA) of Markowitz is based on quadratic optimization. However, the CLA procedure is unstable and produces different portfolios mainly because the algorithm requires the inversion of a covariance matrix. The authors introduce a hierarchical risk parity (HRP) method, an alternative algorithm that utilizes graph theory and machine learning, and that uses the information in the covariance matrix without requiring a matrix inversion. By replacing the covariance matrix with a tree structure, the HRP method allows to solve portfolio optimization problem even when the covariance matrix is singular. Monte Carlo experiments show that HRP produces lower out-of-sample variance than the CLA procedure, while CLA delivers minimum variance portfolios only in-sample.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2007, Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?, *The Review of Financial Studies* 22, 1915–1953.

The authors conduct an empirical exercise to evaluate the out-of-sample performance of optimal asset allocation models. In terms of Sharpe ratio, certainty equivalent return, or turnover, none of these

models fared better than the naive 1/N portfolio. The authors attribute this finding to model estimation error that more than offsets the gains from optimal portfolio diversification.

Dhar, Vasant, and Roger Stein, 2016, Fintech platforms and strategy, Available at SSRN 2892098.

Briefly, FinTech is described as financial sector innovations that have revolutionized most aspects of financial services. The article discusses the current status of FinTech in the Internet era, future directions of change, and possible strategies to pursue. The FinTech market is yet to achieve its highest potential in the coming decades as developing countries are entering the digital era and new platforms for financial services are still developing.

Dixon, Matthew, Diego Klabjan, and Jin Hoon Bang, 2016, Classification-based financial markets prediction using deep neural networks, *Algorithmic Finance* 1–11.

Deep neural networks (DNN) have been successfully applied to speech transcription and image detection. In this study, the authors describe the application of DNN to classification of direction of movement in financial data.

Dixon, Matthew and Nicholas G. Polson, 2019, Deep fundamental factor models, Available at arXiv 1711.04837.

The authors extend fundamental factor models using neural networks to allow for non-linearity, interaction effects, and non-parametric shocks. The proposed technique can be used under heteroskedastic errors, provides interpretability, and ranks the importance of factors (e.g., current enterprise value, price-to-book ratio, price-to-sales ratio, price-to-earnings ratio, log market cap) and interaction effects. The neural network factor model predicts S&P 500 stock returns with smaller out-of-sample mean-squared errors (MSE) than Generalized Linear Regressions. In addition, neural network factor models are found to produce information ratios that are three times higher than linear models.

Donaldson, R. Glen, and Mark Kamstra, 1997, An artificial neural network-GARCH model for international stock return volatility, *Journal of Empirical Finance* 4, 17–46.

The article examines the performance of stock return volatility forecasting models using daily returns data from London, New York, Tokyo, and Toronto. An artificial neural network GARCH model is found to generally outperform its traditional competing models – GARCH, EGARCH, and Sign-GARCH models – in both in-sample and out-of-sample periods.

Dunis, Christian L., Jason Laws, and Georgios Sermpinis, 2010, Modelling and trading the EUR/USD exchange rate at the ECB fixing, *The European Journal of Finance* 16, 541–560.

This paper explores the application of neural networks (NN) to the exchange rate forecasting problem. Unlike previous studies on neural networks that select the network's inputs in an ad-hoc basis, the latter NN techniques use autoregressive terms as inputs. The performance of four such NN techniques (a higher-order NN (HONN), a Psi Sigma Network, NN with multilayer perception or MLP) on the daily EUR/USD exchange rate are compared with traditional methods. In simple simulated exercises, the MLP technique significantly outperforms other competing models considered. For sophisticated trading strategies, the HONN technique achieves higher annualized returns than other NN techniques.

Eiben, Agoston, and Jim Smith, 2015, *Introduction to Evolutionary Computing*, Springer, second edition

This textbook provides a comprehensive introduction to evolutionary algorithms, from history to major design decisions to parameter tuning and control. Advanced topics include multi-objective optimization, optimization of dynamic and noisy functions, constraint handling, or hybridization with other techniques.

Elliott, Graham, and Allan Timmermann, 2008, Economic forecasting, *Journal of Economic Literature* 46, 3–56.

This study presents the forecasting problem in a unified framework, including forecasters' loss functions and types of economic forecasting methods. The authors then explore the various forecasting methodologies available to the researcher, such as optimal point forecasts, classical approach to forecasting, Bayesian approach to forecasting, density forecasts, and forecast combination techniques.

Fan, Alan, and Marimuthu Palaniswami, 2001, Stock selection using support vector machines, *Proceedings of the International Joint Conference on Neural Networks*, volume 3, 1793–1798, IEEE.

In this article, the authors choose stocks trading on the Australian Stock Exchange using support vector machine (SVM). Compared to the 71% return for the benchmark portfolio, an equally-weighted portfolio of stocks selected by SVM produced 208% return over a five year period.

Farmer, J. Doyne, Austin Gerig, Fabrizio Lillo, and Szabolcs Mike, 2006, Market efficiency and the long-memory of supply and demand: is price impact variable and permanent or fixed and temporary?, *Quantitative Finance* 6, 107–112.

The long memory of demand and supply means that future buying and selling behavior, and therefore price movements, should be predictable. However, in reality price movements are essentially uncorrelated. The aim of this article is to explain market efficiency given these two apparently contradictory facts. The authors revisit earlier studies on the topic and demonstrate that market efficiency is maintained by liquidity imbalances rather than mean-reverting price changes. To show this, they introduce transaction time into the model and demonstrate that liquidity covaries in time with the long-memory of supply and demand.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, 2017, Taming the factor zoo, Available at SSRN 2934020.

The authors present a LASSO (least absolute shrinkage and selection operator) model-selection method that allows to systematically evaluate the marginal gains from new factors out of many potential candidates to explain asset returns. A LASSO model selection technique allows to choose the best predictors, however, it cannot reliably tell if the selected variable is one of the variables that belongs to the true asset pricing model.

Feng, Guanhao, Nick Polson, and Jianeng Xu, 2019, Deep learning in characteristics-sorted factor models, Available at SSRN 3243683.

Unlike previous literature relating firm characteristics (inputs) to security returns (outputs), the paper introduces an intermediate channel involving risk factors (intermediate features). With the objective to reduce pricing error, this bottom-up approach trains a neural network that generates risk factors using firm characteristics to explain security returns. The paper provides an alternative dimension reduction framework on security sorting and factor generation.

Fernandes, Marcelo, Marcelo C. Medeiros, and Marcel Scharth, 2014, Modeling and predicting the CBOE market volatility index, *Journal of Banking and Finance* 40, 1–10.

The Chicago Board Options Exchange (CBOE) reports the volatility index VIX based on the 30-calendar day S&P 500 index option. This implied volatility is a key indicator of the overall market condition and is used in many trading strategies. This paper studies the time series properties of the VIX series. The time series analyses show that the VIX index is negatively correlated with S&P 500 index returns, but has a positive relationship with the S&P 500 index volume. Further, these two series do not appear to have a nonlinear relationship.

Fischer, Thomas, and Christopher Krauss, 2018, Deep learning with long short-term memory networks for financial market predictions, *European Journal of Operational Research* 270, 654–669.

In this study, the authors use long short-term memory (LSTM) networks in predicting out-of-sample movements for S&P 500 stocks. LSTM is a type of sequence learning method specifically designed to

learn long-term dependencies. LSTM clearly outperforms memory-free classification methods, such as deep neural nets and a logistic regression in terms of predictive accuracy. It also beats random forests except during the global financial crisis.

Fisher, Ingrid E., Margaret R. Garnsey, and Mark E. Hughes, 2016, Natural language processing in accounting, auditing and finance: A synthesis of the literature with a roadmap for future research, *Intelligent Systems in Accounting, Finance and Management* 23, 157–214.

In this article, the authors synthesize the literature that applies natural language processing (NLP) in accounting, auditing and finance. NLP is used to analyze textual data (e.g., corporate financial performance, management's assessment of firm performance, regulations, etc.) to analyze the data, fraud detection, make inferences and stock price predictions. The review of large literature in the three fields reveal that out of all the machine learning techniques applied in NLP, support vector machines (SVMs) are found to be the most popular tool, followed by naive Bayes (NB), hierarchical clustering, statistical methods, and term-frequency-inverse document frequency (TF-IDF) weighting.

Fletcher, Tristan, and John Shawe-Taylor, 2013, Multiple kernel learning with fisher kernels for high frequency currency prediction, *Computational Economics* 42, 217–240.

The authors show that constructing kernels based on EUR/USD exchange rate data from limited order book volumes, technical analysis, and from market microstructure models. Using these multiple kernels, support vector machines (SVM) are employed to predict the direction of price movements. SVMs based on multiple kernels significantly outperform individual SVMs and achieve 55 percent forecast accuracy, on average.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, 2018, Dissecting characteristics nonparametrically, Available at SSRN 3223630.

The paper proposes a nonparametric method to choose firm characteristics that have incremental predictive power for the cross-section of stock returns. Using adaptive LASSO (least absolute shrinkage and selection operator) nonparametric method for model selection and stock return forecast, it finds that out of 62 firm characteristics only 9-16 have incremental predictive power. Compared to linear model selection methods, LASSO performs relatively well in both model selection and expected return predictors.

Geva, Tomer, and Jacob Zahavi, 2014, Empirical evaluation of an automated intraday stock recommendation system incorporating both market data and textual news, *Decision Support Systems* 57, 212–223.

This paper examines the effectiveness of incorporating textual news data for stock return forecasting. To this end, various textual data types are added sequentially to numerical market data to study the incremental gain to forecasting performance. Overall, integration of textual data, especially more advanced ones, with market data and using neural networks significantly improves forecast accuracy.

Giamouridis, Daniel, 2017, Systematic investment strategies, *Financial Analysts Journal* 73, 10–14.

This editorial offers some thoughts on the latest research on systematic investment from both academic and practical perspectives. It also highlights areas where further research can improve the practice more considerably. Some of the topics highlighted in the paper include data science and machine learning, factor investing, market timing, and coordinated investing and holding.

Giamouridis, Daniel, Sandra Paterlini, 2010, Regular(ized) hedge fund clones, *Journal of Financial Research* 33, 223–247.

In this paper, the authors show the attractiveness of using constrained portfolio optimization methods for constructing efficient hedge fund clones. The findings indicate that the proposed hedge fund port-

folio is competitive relative to hedge fund clones obtained by commercially available proprietary construction techniques as well as benchmark hedge fund indices. The replicated hedge fund portfolio exhibits a very high correlation with the benchmark indices.

Gogas, Periklis, Theophilos Papadimitriou, Maria Matthaiou, and Efthymia Chrysanthidou, 2015, Yield curve and recession forecasting in a machine learning framework, *Computational Economics* 45, 635–645.

In this paper, the authors demonstrate the superior forecasting performance of support vector machines (SVM) in forecasting economic recessions in the U.S. Compared to logit and probit models, the SVM technique improves forecast accuracy by 67%–100% in forecasting recessions in the U.S.

Gradojevic, Nikola, and Jing Yang, 2006, Non-linear, non-parametric, non-fundamental exchange rate forecasting, *Journal of Forecasting* 25, 227–245.

This paper examines the exchange rate forecasting performance using artificial neural networks (ANN). Compared to naïve random walk or any linear time series model, ANN performs well in forecasting the Canadian-U.S. dollar exchange rate based on the root mean squared error criterion.

Groth, Sven S., and Jan Muntermann, 2011, An intraday market risk management approach based on textual analysis, *Decision Support Systems* 50, 680–691.

The paper studies the potential gains from using newly obtained qualitative text data to explain intraday market risk. Using naïve Bayes, k-nearest neighbor, neural network, and support vector machines for textual analysis, the authors conclude that these techniques are helpful in discovering intraday market exposure, with the latter model (SVM) being the best in terms of computational efficiency and classification accuracy.

Gu, Shihao, Bryan T. Kelly, and Dacheng Xiu, 2018, Empirical asset pricing via machine learning, Available at SSRN 3159577.

The authors show that machine learning techniques (trees and neural nets) may outperform traditional cross-section and time series models in predicting asset risk premia. The gain in predictive performance comes from the ability of machine learning methods to allow nonlinear predictive interactions that may be infeasible using traditional statistical methods.

Gu, Shihao, Bryan T. Kelly, and Dacheng Xiu, 2019, Autoencoder Asset Pricing Models, Available at SSRN 3335536.

The study proposes an autoencoder neural network to latent factor modeling for asset pricing. Unlike the linearity assumption found in other studies, the approach allows for non-linear relationships between factor exposures and asset characteristics. Incorporating the no-arbitrage condition into the machine learning framework, the new latent factor asset pricing model produces far smaller out-of-sample errors relative to other leading asset pricing modeling techniques, such as Fama-French, PCA, and linear conditioning methods.

Hagenau, Michael, Michael Liebmann, and Dirk Neumann, 2013, Automated news reading: Stock price prediction based on financial news using context-capturing features, *Decision Support Systems* 55, 685–697.

In this article, the authors describe the application of textual financial information to stock price prediction. Unlike other similar approaches, this study uses a feature selection that allows market feedback. By selecting the feedback-based relevant features, this approach reduces overfitting problems and significantly improves classification accuracy. As a result, stock return prediction accuracy increased by as much as 76%.

Hamid, Shaikh A., and Zahid Iqbal, 2004, Using neural networks for forecasting volatility of S&P 500 index futures prices, *Journal of Business Research* 57, 1116–1125.

The current study examines the performance of the popular neural networks technique in forecasting the volatility of S&P 500 Index futures prices. The results show that volatility forecasts from neural networks outperform implied volatility forecasts obtained from the Barone-Adesi and Whaley American futures options pricing model.

Han, Shuo, and Rung-Ching Chen, 2007, Using SVM with financial statement analysis for prediction of stocks, *Communications of the IIMA* 7, 8.

This article proposes applying support vector machines to financial statement analysis for stock market prediction. Using financial indices improves forecast accuracy and is more reliable than using technical indices.

Haykin, Simon, 2009, *Neural Networks and Learning Machines*, 3rd ed. (Pearson).

This comprehensive textbook presents an up-to-date treatment of neural networks. The book consists of six main parts that focus on supervised learning, kernel methods based on radial-basis function networks, regularization methods, unsupervised learning, reinforcement learning, and nonlinear feedback systems.

Heaton, James B., Nick G. Polson, and Jan H. Witte, 2017, Deep learning for finance: deep portfolios, *Applied Stochastic Models in Business and Industry* 33, 3–12.

Use of traditional methods from financial economic theory may not be feasible or difficult to apply to financial prediction and classification problems due to the large and complex nature of data. The authors describe deep learning hierarchical models that can improve performance in financial prediction problems and classification.

Hendricks, Dieter, and Diane Wilcox, 2014, A reinforcement learning extension to the Almgren-Chriss framework for optimal trade execution, *IEEE Conference on Computational Intelligence for Financial Engineering Economics (CIFER)*, 457–464.

In this article, the authors present a reinforcement learning technique to optimize the liquidation volume trajectory. The model is based on a standard Almgren-Chriss model, a popular trade execution strategy that assumes risk aversion to trade execution. The proposed technique is able to change the volume of trade using information on market spread and volume dynamics. Empirical tests show promising results. Specifically, on average, the new technique can improve post-trade implementation shortfall by up to 10.3% relative to the base model.

Hong, Taeho, and Ingoo Han, 2002, Knowledge-based data mining of news information on the internet using cognitive maps and neural networks, *Expert Systems with Applications* 23, 1–8.

The authors of this study develop a Knowledge-Based News Miner (KBNMiner) that integrates cognitive maps with neural networks. Unlike the time series models used in interest rate forecasting, the technique uses prior news on interest rates in predicting interest rates. Relative to neural network and random walk models, KBNMiner leads to an improved interest rate prediction.

Hu, Yong, Kang Liu, Xiangzhou Zhang, Lijun Su, E.W.T. Ngai, and Mei Liu, 2015, Application of evolutionary computation for rule discovery in stock algorithmic trading: A literature review, *Applied Soft Computing* 36, 534–551.

This study explores the literature on the application of evolutionary computation (EC) in stock algorithmic trading. The authors observed that most of the trading techniques considered in the surveyed studies do well in the downtrend, but poorly in the uptrend, which is likely due to the problems associated with the selection of factors and transaction costs.

Huang, Chien-Feng, 2012, A hybrid stock selection model using genetic algorithms and support vector regression, *Applied Soft Computing* 12, 807–818.

Stock selection problems are particularly challenging in the area of investment research. This paper examines a hybrid approach for stock selection using support vector regression (SVR) and genetic algorithms (GA). In the first stage, SVR chooses the top ranked stocks. In the second stage, GA is used to optimize the portfolio parameters and feature selection. As the results indicate, this hybrid approach significantly outperforms the benchmark model.

Huang, Wei, Yoshiteru Nakamori, and Shou-Yang Wang, 2005, Forecasting stock market movement direction with support vector machine, *Computers and Operations Research* 32, 2513–2522.

The authors examine the performance of support vector machines (SVM), a popular learning algorithm, in predicting the direction of the Nikkei 225 index. Relative to forecasts obtained from Linear Discriminant Analysis, Quadratic Discriminant Analysis, and Elman Backpropagation Neural Networks, forecasts from SVM are more accurate in predicting the direction of the stock index. The forecasting performance can further be improved by combining SVM with other classification techniques.

Huang, Zan, Hsinchun Chen, Chia-Jung Hsu, Wun-Hwa Chen, and Soushan Wu, 2004, Credit rating analysis with support vector machines and neural networks: a market comparative study, *Decision Support Systems* 37, 543–558.

This study applies support vector machines (SVM) to corporate credit rating analysis. Both the SVM and benchmark neural network techniques yield around 80% prediction accuracy for the U.S. and Taiwan market data. Additionally, authors implement input financial variable contribution analysis to help with the interpretation of neural network results.

Hutchinson, James M., Andrew W. Lo, and Tomaso Poggio, 1994, A nonparametric approach to pricing and hedging derivative securities via learning networks, *The Journal of Finance* 49, 851–889.

This article presents a nonparametric pricing approach for estimating pricing and hedging derivative securities. Unlike a parametric approach for estimating the arbitrage-based pricing formula, a nonparametric approach does not require that the dynamics of asset prices be known in advance. Using data on pricing and delta-hedging of S&P 500 futures options, the authors find that a network-pricing formula in the majority of cases outperforms the naive traditional Black-Scholes model. Additionally, compared to other popular models, such as ordinary least squares, radial basis function networks, multilayer perceptron network, and projection pursuit, the neural network pricing formula is computationally more efficient and accurate when the price dynamics of the underlying asset are unknown.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani, 2017, *An introduction to statistical learning with applications in R* (Springer).

This book is devoted to studying statistical learning, such as LASSO, classification and regression trees, boosting, and support vector machines. Unlike other textbooks in this new area, the book treats these topics in a less technical manner with a focus on their application using statistical software R.

Kaashoek, Johan F., and Herman K. Van Dijk, 2002, Neural network pruning applied to real exchange rate analysis, *Journal of Forecasting* 21, 559–577.

In this paper, the authors apply neural network pruning to exchange rate forecasting. As neural networks consist of many cells, the number of cells is pruned or reduced using basic descriptive procedures, such as multiple correlation coefficients, principal component analysis of residuals, and graphical analysis. Compared to the standard ARIMA models, the proposed method performs better in terms of its predictive accuracy and its ability to capture the long-term dynamics of exchange rates.

Katona, Zsolt, Marcus Painter, Panos Patatoukas and Jean Zeng, 2018, On the capital market consequences of alternative data: Evidence from outer space. Available at SSRN 3222741.

The paper examines the use of satellite images of parking lot traffic to predict the earnings of major US retailers before public disclosure.

Ke, Zheng Tracy, Bryan T. Kelly, and Dacheng Xiu, 2019, Predicting returns with text data. Available at SSRN 3489226.

The paper proposes a text mining methodology that extracts sentiment information from textual sources. This approach is simple, requires minimal computing power, and can be adapted to the dataset being used. A simple trading strategy that buys assets with positive recent news sentiment and sells assets with negative sentiment using this method generates higher out-of-sample abnormal returns than similar strategies based on sentiment scores from commercial vendors such as RavenPack.

Kearney, Colm, and Sha Liu, 2014, Textual sentiment in finance: A survey of methods and models, *International Review of Financial Analysis* 33, 171–185.

This article provides a comprehensive review of the text-based sentiment literature, qualitative information sources used in prior research, most frequently used textual data analysis methods and their empirical applications. Further, the paper identifies possible opportunities and directions for future research using text-based qualitative data.

Kearns, Michael, and Yuriy Nevmyvaka, 2013, Machine learning for market microstructure and high frequency trading, *High Frequency Trading: New Realities for Traders, Markets, and Regulators*. Risk Books

The authors explore the application of machine learning approaches to high frequency trading and microstructure data. They consider a few case studies used in solving different trading problems. Specifically, they review the application of machine learning methods in optimal trade execution problems and in predicting equity limit order book price movements.

Kercheval, Alec N., and Yuan Zhang, 2015, Modelling high-frequency limit order book dynamics with support vector machines, *Quantitative Finance* 15, 1315–1329.

In this paper, the authors develop a support vector machine-based learning model to examine the dynamics of information contained in a limit order book. A limit order book records high-frequency trading activity grouped by asks and bids, including time and type of transaction, order price and volume. The authors show the usefulness of features selected by the proposed method for forecasting short-term price dynamics.

Kim, Steven H., and Hyun Ju Noh, 1997, Predictability of interest rates using data mining tools: a comparative analysis of Korea and the US, *Expert Systems With Applications* 13, 85–95.

Forecasting techniques used in predicting interest rates have generally failed to improve over the simple random walk model. In this paper, the authors consider neural networks, case-based reasoning and their combination to forecast interest rates in the U.S. and Korea. Case-based reasoning uses accumulated past experiences in making decisions. Interestingly, these techniques are superior to the random walk model for the US market, but cannot outperform the random walk model for Korean market.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun, 2017, The flash crash: High-frequency trading in an electronic market, *The Journal of Finance* 72, 967–998.

This paper investigates intraday market intermediation in an electronic market around the time of large and temporary selling pressure that occurred on May 6, 2010. Known as the Flash Crash, the temporary decline in prices and increase in trading volume was triggered by a trader who initiated a large sell program to sell E-mini S&P 500 stock index futures. This kind of large temporary selling pressure may potentially lead to a market crash. Using empirical data on May 6, 2010 and three days before that date,

the study finds evidence that most intraday intermediaries did not change their trading behavior despite large and temporary selling pressure.

Kirilenko, Andrei A., and Andrew W. Lo, 2013, Moore's law versus Murphy's law: Algorithmic trading and its discontents, *Journal of Economic Perspectives* 27, 51–72.

In this paper, the authors survey algorithmic trading and its history, its major drivers, challenges and resulting unintended consequences for the financial system. The authors argue that while technological advance has reduced the financial transaction costs drastically, the current financial regulations need to be aligned with the requirements of the Digital Age.

Kofman, Paul, and Ian G. Sharpe, 2003, Using multiple imputation in the analysis of incomplete observations in finance, *Journal of Financial Econometrics* 1, 216–249

This paper examines the application of multiple imputation methods. When applied to two financial datasets involving severe data incompleteness, the imputation methods outperform the ad hoc approaches commonly used in the finance literature.

Kolm, Petter N., and Gordon Ritter, 2020, Modern perspectives on reinforcement learning in finance, Available at SSRN 3449401.

This paper provides an overview of reinforcement learning applications in finance. Reinforcement learning allows for solving dynamic optimization problems such as pricing and hedging of contingent claims, investment and portfolio allocation, buying and selling a portfolio of securities subject to transaction costs, market making, asset liability management, and optimization of tax consequences in a model-free way. The authors also highlight some of the challenges with using reinforcement learning.

Kolm, Petter N., Reha Tütüncü, and Frank J. Fabozzi, 2014, 60 Years of portfolio optimization: Practical challenges and current trends, *European Journal of Operational Research* 234, 356–371.

The authors review the approaches that aim at addressing some of the practical challenges of portfolio optimization. These challenges include accounting for transaction costs, portfolio management constraints, and the sensitivity to the estimates of expected returns and covariances. The paper also illustrates a number of developments in portfolio optimization including diversification methods, risk-parity portfolios, the mixing of different sources of alpha, and practical multi-period portfolio optimization.

Kumar, P. Ravi, and Vadlamani Ravi, 2007, Bankruptcy prediction in banks and finance via statistical and intelligent techniques – a review, *European Journal of Operational Research* 180, 1–28.

In this comprehensive review, the authors consider the application of statistical and intelligent techniques for bankruptcy prediction between 1985 and 2005. Historically, many researchers employ many different techniques to forecast bank and firm bankruptcy. These techniques include traditional statistical methods, neural networks, case-based reasoning, decision trees, evolutionary approaches, rough set based techniques, fuzzy logic, support vector machines, and various hybrid models.

Lam, Monica, 2004, Neural network techniques for financial performance prediction: integrating fundamental and technical analysis, *Decision Support Systems* 37, 567–581.

This study examines the effect of integrating technical and fundamental analysis to forecast the rate of return on common shareholders' equity. Empirical results show that a neural network incorporating financial statement and macroeconomic variables performs significantly better than the market return. However, it fails to outperform the maximum benchmark – average return of the top one-third of companies with the highest returns.

Leshik, Edward, and Jane Cralle, 2011, *An introduction to algorithmic trading: basic to advanced strategies* (John Wiley and Sons).

This book provides the reader with the background on algorithmic trading and describes current algorithmic trading strategies. It contains two broadly defined parts. Part I is an introduction to trading algorithms, current popular algorithms, and how to use and optimize these trading algorithms. The second part provides the reader with the tools and examples to learn the trading strategies developed by the authors.

Leung, Henry, and Thai Ton, 2015, The impact of internet stock message boards on cross-sectional returns of small-capitalization stocks, *Journal of Banking and Finance* 55, 37–55.

In this article, the authors study the effect of stock messages posted on the Australian online stock message board HotCopper of the Australian stock market. The empirical evidence suggests that both, the number of messages and sentiment, have a positive impact on the current return of underperforming small stocks and their trading volume. However, large stocks do not appear to be influenced by online message sentiment.

Li, Xiaodong, Xiaodi Huang, Xiaotie Deng, and Shanfeng Zhu, 2014, Enhancing quantitative intra-day stock return prediction by integrating both market news and stock prices information, *Neurocomputing* 142, 228–238.

Previous studies on stock prices and market news either consider market news or past stock prices as a predictor of future stock returns. In this article, the authors integrate these two sources of information using a kernel learning technique. The empirical evidence suggests that the proposed method outperforms three baseline models (market news, past stock prices, both market news and past stock prices).

Liao, Shu-Hsien, and Shan-Yuan Chou, 2013, Data mining investigation of co-movements on the Taiwan and China stock markets for future investment portfolio, *Expert Systems with Applications* 40, 1542–1554.

Taiwan and China signed an Economic Cooperation Framework Agreement (ECFA) in 2010. The aim of this paper is to investigate the co-movements in the two countries' stock markets after ECFA. To this end, 30 stock indexes are identified and their behavior analyzed using patterns, rules, and cluster analysis. The empirical results show that for the Taiwan stock market, electronics, financial and insurance, and semi-conductor stock indexes strongly move together with the TAIEX index. For the Hong Kong stock market, real estate, telecommunications, and financial services stock indexes move with the HIS index, and for the Shenzhen stock market, manufacturing, machinery, and electronics stock indexes move with the SZE index. The study also discusses the co-movement across these three stock markets.

Lin, Chin-Shien, Haider A. Khan, Ruei-Yuan Chang, and Ying-Chieh Wang, 2008, A new approach to modeling early warning systems for currency crises: Can a machine-learning fuzzy expert system predict currency crises effectively?, *Journal of International Money and Finance* 27, 1098–1121.

In this study, the authors develop a hybrid causal model to predict currency crises. This hybrid model, constructed by integrating neural networks with a fuzzy logic model, can predict currency crises better than other popular methods, such as neural networks and logistic regression.

Loterman, Gert, Iain Brown, David Martens, Christophe Mues, and Bart Baesens, 2012, Benchmarking regression algorithms for loss given default modeling, *International Journal of Forecasting* 28, 161–170.

The existing literature on credit risk models focuses mainly on the probability of default, while loss given default (LGD) remains an inadequately studied area. This study considers numerous regression methods to model the LGD problem. Using data on bank losses from several major international banks, the considered models can explain 4% to 43% of the variation in LGD. Out of 24 models

considered in the study, support vector machines and neural networks perform significantly better than traditional statistical techniques.

Lowe, David, 1994, Novel exploitation of neural network methods in financial markets, IEEE International Conference on Neural Networks, volume 6, 3623–3628, IEEE.

The author introduces feed forward networks for portfolio management. As the author demonstrates, neural network methods have the advantage of being able to approximate the nonlinearities in the data and to optimize the portfolio under constraints.

Majhi, Ritanjali, Ganapati Panda, and Gadadhar Sahoo, 2009, Efficient prediction of exchange rates with low complexity artificial neural network models, *Expert Systems With Applications* 36, 181–189.

This article proposes two artificial neural network (ANN) models, specifically functional link ANN (FLANN) and cascaded functional link ANN (CFLANN) to forecast currency exchange rates. Compared to the popular Widrow's least mean-square algorithm, both CFLANN and FLANN models are found to be superior in exchange rate forecasting, with the CFLANN model being the best among the three models considered.

Manahov, Viktor, Robert Hudson, and Bartosz Gebka, 2014, Does high frequency trading affect technical analysis and market efficiency? and if so, how?, *Journal of International Financial Markets, Institutions and Money* 28, 131–157.

About 40 percent of foreign exchange (FX) traders use technical analysis for their trading rules, while others may implicitly rely on the Efficient Market Hypothesis. In this article, the authors examine the relationship between technical analysis and high frequency trading (HFT) in FX markets. The authors develop a special adaptive form of Strongly Typed Genetic Programming (STGP) to forecast most frequently traded currency pairs using high frequency data. The results show that STGP significantly outperforms traditional econometric forecasting models. Importantly, excess returns are found to be both statistically and economically significant even when transaction costs are considered.

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

In this article, the authors propose a text-based measure of uncertainty – a news implied volatility (NVIX) measure – to study the relationship between uncertainty and expected returns. Using front-page articles of the Wall Street Journal starting in 1890, the constructed NVIX measure captures the time variation in risk premia. Specifically, a high NVIX measure indicates high future returns in normal times and rises just before rare disaster events (e.g., wars, concerns about government policy, natural disasters, etc.).

Manning, Christopher, and Hinrich Schütze, 1996, *Foundations of Statistical Natural Language Processing* (Cambridge, MA: MIT Press).

Statistical natural language processing (NLP) applies probabilistic and information theory, and linear algebra to characterize linguistic observations. This graduate-level textbook consists of four broad parts. In part 1 the essential mathematical and linguistic concepts are presented. Part 2 focuses on word-based statistics and inference. Part 3 is devoted to studying grammar-based statistical techniques. Finally, the fourth part of the book is devoted to the applications of statistical techniques to NLP.

Markowitz, Harry, 1952, Portfolio selection, *The Journal of Finance* 7, 77–91.

In this seminal theory paper, Markowitz describes the portfolio selection problem using expected returns and their variances. Also known as the mean-variance model, the model chooses an efficient portfolio with the maximum expected return for a given level of risk.

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, 2006, A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955, *AI magazine* 27, 12.

This article republished the seminal 1955 Dartmouth Summer Research Proposal on AI by four mathematicians. It is believed that the authors of this research proposal were the first to have used the term 'artificial intelligence'. The authors suggested conducting a 10 man study on AI during the summer of 1956.

Michaud, Richard O., and Robert O. Michaud, 2008, *Efficient asset management: a practical guide to stock portfolio optimization and asset allocation* (Oxford University Press).

This book discusses mean-variance (MV) portfolio optimization, and its limitations. It then reviews alternative portfolio management approaches from a statistical point of view. The range of topics covered include, but are not limited to, Markowitz efficiency, classic mean-variance optimization, traditional criticisms and alternatives, unbounded MV portfolio efficiency, MV efficiency with linear constraints, the resampled efficiency frontier and its properties, portfolio rebalancing and monitoring, input estimation, Bayes estimation and caveats, and avoiding optimization errors.

Mitkov, R., 2014, *The Oxford Handbook of Computational Linguistics*, 2nd ed. (New York: Oxford University Press).

This book is devoted to describing major concepts, methods, and applications in computational linguistics. It provides an overview of the field and describes a broad range of current techniques used in natural language processing (NLP). The book also provides a comprehensive survey of current applications of NLP.

Molnar, Christoph, 2020, *Interpretable Machine Learning – a Guide for Making Black Box Models Explainable*. <https://christophm.github.io/interpretable-ml-book/>

A free and comprehensive book on various approaches to facilitate the interpretation and understanding of machine learning predictions.

Murphy, Kevin P., 2012, *Machine Learning: A Probabilistic Perspective* (The MIT Press).

This advanced textbook on machine learning takes a particular focus on probability theory and distributions. The first part of the book is devoted to machine learning concepts and methods, probability theory and distributions, Bayesian and frequentist statistics. The second part of the book presents linear and logistic regression, generalized linear, mixture and latent linear models. The third part of the book discusses kernels, adaptive models (classification and regression trees, boosting, neural networks, ensemble learning), Markov and hidden Markov models, state space models, graphical models, variational inference, Monte Carlo inference and deep learning.

Nevmyvaka, Yuriy, Yi Feng, and Michael Kearns, 2006, Reinforcement learning for optimized trade execution, in *Proceedings of the 23rd International Conference on Machine learning*, 673–680, ACM.

Optimized trade execution is one of the important problems in the field of finance. In this study, reinforcement learning (RL) is applied to optimal trade execution and evaluated using NASDAQ market data. When market state variables are chosen carefully, RL can improve trade optimization relative to other baseline execution strategies.

Nuij, Wijnand, Viorel Milea, Frederik Hogenboom, Flavius Frasincar, and Uzay Kaymak, 2014, An automated framework for incorporating news into stock trading strategies, *IEEE Transactions on Knowledge and Data Engineering* 26, 823–835.

The authors introduce a framework that automatically incorporates news into stock trading strategies. Using genetic programming to find optimal trading strategies, the results indicate that optimal trading strategies include technical trading rules and in many cases news variables as additional input.

Nuti, Giuseppe, Mahnoosh Mirghaemi, Philip Treleaven, and Chaiyakorn Yingsaeree, 2011, Algorithmic trading, *Computer* 44, 61–69.

Trading in the financial sector uses automated systems that are fast and complex. This article provides an overview of trading algorithms and how the system works. In particular, it explains the trading objective, trading process, electronic trading execution, trading analysis, and provides some examples.

Oh, Kyong Jo, and Ingoo Han, 2000, Using change-point detection to support artificial neural networks for interest rate forecasting, *Expert Systems With Applications* 19, 105–115.

Because interest rates change due to the monetary policy of governments, this study proposes to identify intervals between these change points, and use this information in predicting interest rates. The authors use backpropagation neural network (BPN) to detect the change-point groups of interest rates and apply the BPN technique again to interest rate forecasting. The proposed BPN technique with change-point detection outperforms the simple BPN technique at a statistically significant level.

Papaioannou, Georgios V., and Daniel Giamouridis, forthcoming, Enhancing alpha signals from trade ideas data using supervised learning, in *Machine Learning and Asset Management* (Springer).

In this chapter the researchers use trade investment ideas along with supervised machine learning. Trade ideas are market experts' recommendations which are often used by institutional investors. Investment trade ideas are classified into two classes (success or failure) using supervised machine learning methods, specifically random forests and gradient boosting trees. In addition to stock characteristics, the study also uses characteristics of idea's contributor. The overall results demonstrate a performance improvement of more than 1% for long ideas and of more than 2% for short ideas.

Park, Saerom, Jaewook Lee, and Youngdoo Son, 2016, Predicting market impact costs using nonparametric machine learning models, *Plos One* 11, 1–13.

Transaction costs affect the profits of investment strategies. This paper seeks to more accurately predict market impact cost, which is the result of the difference between the initial price and the actual price after the transaction. Using data on the U.S. stock market, the authors apply nonparametric machine learning techniques (neural networks, Bayesian neural network, Gaussian process, support vector regression) to predicting market impact cost. The empirical results suggest that nonparametric machine learning models generally outperform their parametric counterparts.

Patel, Keyur, and Marshall Lincoln, 2019, It's not magic: Weighing the risks of AI in financial services, Available at: http://www.csfi.org/s/Magic_10-19_v12_Proof.pdf.

The report offers a detailed review of potential benefits and risks from the application of AI in financial services industry. The report is divided into three sections. Section one introduces AI in financial services industry. Section two discusses potential benefits of AI in financial services like improvements in security, compliance and risk management. The report devotes much attention to section three, which presents potential risks arising from the application of AI and machine learning (ML) in financial services. It identifies 12 key risks, including those related to ethical challenges related to increased use of AI and ML by the financial sector.

Pena, Tonatiuh, Serafi Martinez, and Bolanle Abudu, 2011, Bankruptcy prediction: A comparison of some statistical and machine learning techniques, in *Computational Methods in Economic Dynamics*, 109–131 (Springer).

The authors examine the accuracy of statistical and machine learning methods in predicting bank failures. They introduce Gaussian processes (GP) for classification and evaluate their performance relative to other statistical and machine learning techniques (logistic regression, discriminant analysis, least-squares support vector machines). The study finds that forecasts generated from different instances of GP classifiers can compete with the results generated from popular techniques.

Rasekhschaffe, Keywan Christian, and Robert C. Jones, 2019, Machine learning for stock selection, *Financial Analysts Journal* 75, 70–88.

Machine learning methods are gaining popularity among financial practitioners as they can better capture dynamic relationships between predictors and expected returns. However, given the noisy historical financial data, the risk of overfitting poses a real challenge. The paper discusses two main ways of overcoming the overfitting problem when using machine learning for predicting cross section of stock returns. Combination of different forecasts reduces noise. The paper recommends forecast combination along different dimensions, such as from different forecasting techniques, based on different training sets, and for different horizons. Similarly, feature engineering can help mitigate the overfitting problem by increasing the signal-to-noise ratio.

Rapach, David E., Jack K. Strauss, Jun Tu, and Guofu Zhou, 2019, Industry return predictability: A machine learning approach, *The Journal of Financial Data Science* 1, 9–28.

The study applies machine learning to predict industry returns. Specifically, it uses LASSO regression to fit sparse models that include lagged industry returns for 30 industries. The LASSO selected variables are then estimated using an OLS model to lessen the effect of downward bias in estimated coefficients from the LASSO model. In-sample and out-of-sample predictions of industry returns provide evidence for the relevance of information in lagged industry returns.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2013, International stock return predictability: what is the role of the United States?, *The Journal of Finance* 68, 1633–1662.

Stock return predictability has received significant attention in the literature. In this study, the authors introduce a new powerful predictor of stock returns in industrialized countries. They find that lagged U.S. market returns can dramatically improve stock return predictability in other industrialized countries, while lagged non-US returns are not good predictors of stock returns in the U.S. The contribution of lagged U.S. returns to stock return predictability in non-U.S. industrialized countries is explained through a news-diff model where shocks to U.S. stock returns are reflected in equity prices in other industrialized countries with a lag.

Renault, Thomas, 2017, Intraday online investor sentiment and return patterns in the US stock market, *Journal of Banking and Finance* 84, 25–40.

Using investor opinions and ideas about the stock market returns posted on the StockTwits online blog, the authors construct investor sentiment data to study its relationship with US stock returns. They provide evidence that investor sentiment is an important variable for forecasting intraday stock index returns.

Ribeiro, Bernardete, Catarina Silva, Ning Chen, Armando Vieira, and Joao Carvalho das Neves, 2012, Enhanced default risk models with SVM+, *Expert Systems with Applications* 39, 10140–10152.

Recent advances in bankruptcy prediction consider adding additional information, such as marketing reports, competitors landscape, economic environment, customers screening, industry trends, etc. This additional information can be incorporated in a support vector machine(SVM). Using data on French companies, the authors demonstrate that their adaptation produces a better bankruptcy prediction than a baseline SVM.

Ristolainen, Kim, 2018, Predicting banking crises with artificial neural networks: The role of nonlinearity and heterogeneity, *The Scandinavian Journal of Economics* 120, 31–62.

Early warning systems (EWSs) help to predict coming banking crises. Unlike traditional linear models, such as logistic regression, this study builds EWS using an artificial neural network (ANN) model. For regional as well as international data, the proposed ANN model outperforms logistic regression in predicting all banking crises 2 years in advance given the information about earlier crises.

Russell, Stuart, and Peter Norvig, 2009, *Artificial Intelligence: A Modern Approach*, third edition (Prentice Hall Press, Upper Saddle River, NJ, USA).

The textbook provides an in depth review of AI and covers introductory concepts as well as recent advances in the field. The range of topics includes, intelligent agents, problem-solving agents, search algorithms, logic agents, first-order logic and inference, planning, uncertainty in knowledge and reasoning, learning (e.g., learning from examples and learning probabilistic models), natural language processing and communication, and robotics.

Sabharwal, Chaman L., 2018, The rise of machine learning and robo-advisors in banking, *Journal of Banking Technology* 2, 28–43.

This article discusses the current use of machine learning and its role in the financial sector in the future. Robo-advisors utilized by the largest banks in the U.S. are examples that the financial sector has embraced machine learning for its banking services. Still, machine learning is yet to achieve its biggest impact in the finance industry.

Schumaker, Robert, and Hsinchun Chen, 2006, Textual analysis of stock market prediction using financial news articles, *AMCIS Proceedings* 185.

In this paper, the authors consider the impact of financial news on stock prices. Specifically, three textual document representations, namely Bag of Words, Noun Phrases, and Named Entities, obtained from news articles are considered. Using support vector machines (SVM), the experiment analyzes the impact of news articles on stock prices twenty minutes after a news article is published. There are two interesting results from this study. First, compared to linear models, the SVM finds a statistically significant impact of financial news on stock prices. Second, various textual analysis approaches yield different stock return prediction performance. Compared to the popular Bag of Words, Noun Phrase textual representation results in better prediction performance.

Sevim, Cuneyt, Asil Oztekin, Ozkan Bali, Serkan Gumus, and Erkam Guresen, 2014, Developing an early warning system to predict currency crises, *European Journal of Operational Research* 237, 1095–1104.

This article uses a Financial Pressure Index (FPI) – measuring a drastic deviation in the exchange rate and a drastic decrease in foreign exchange reserves – as the dependent variable and 32 macroeconomic variables as independent variables to predict currency crises in the Turkish economy. The three models considered in this article, artificial neural networks, logistic regression, and decision trees, are able to predict the 1994 and 2001 crises 12 months in advance and with 95% accuracy.

Simon, D., 2013, *Evolutionary Optimization Algorithms* (New Jersey: John Wiley & Sons, Inc.).

This applied textbook is devoted to studying evolutionary algorithms (EAs) for optimization. It is divided into five parts. The first part discusses types of optimization problems and algorithms. The second part reviews natural genetics and their history, and describes the use of artificial genetic algorithms for solving optimization problems. In part 3, the discussion is centered on related techniques, such as ant colony optimization, particle swarm optimization and differential evolution. Part 4 is devoted to special types of optimization problems (discrete, constrained, and multi-objective optimization problems) and problems associated with reducing computational costs of EAs. Finally, part 5 provides a practical guide on how to address problems (checking for bugs and problems in the code and software) and how to measure the performance of an algorithm against standard benchmark optimization problems.

Skolpadungket, Prisadarng, Keshav Dahal, and Napat Harnpornchai, 2016, Handling model risk in portfolio selection using multi-objective genetic algorithm, in *Artificial Intelligence in Financial Markets*, 285–310 (Springer).

The classical Markowitz (mean-variance) portfolio optimization model assumes that asset returns are normally distributed. In reality, means and volatilities of asset returns tend to vary, which requires forecasting these variables to construct an optimal portfolio. The authors present a solution to the portfolio optimization problem using a multi-objective genetic algorithm (MOGA) to account for inaccuracy inherent in forecasting models. This model risk can be reduced when an approximation of the Sharpe ratio error of the portfolio of assets is added as an additional objective to the portfolio optimization task.

Sprenger, Timm O., Philipp G. Sandner, Andranik Tumasjan, and Isabell M. Welp, 2014, News or noise? using twitter to identify and understand company-specific news flow, *Journal of Business Finance and Accounting* 41, 791–830.

Using 400,000 S&P 500 stock-related messages from Twitter messages, the authors compare company returns just before positive and negative news. Good news tend to have a larger information leakage on stock returns compared to the bad news.

Tam, Kar Yan, 1991, Neural network models and the prediction of bank bankruptcy, *Omega* 19, 429–445.

Bank bankruptcy first increased significantly in the 1980s. This article uses a neural network technique to predict bank failure. The experimental results show that the proposed method performs better than traditional statistical methods in terms of robustness, forecast accuracy, adaptability, and explanatory capability.

Tan, Pang-Ning, Michael Steinbach, Anuj Karpatne, and Vipin Kumar, 2018, Cluster Analysis: “Basic Concepts and Algorithms” in *Introduction to Data Mining* 2nd ed. (Pearson).

The authors provide an overview of cluster analysis and illustrate its application in different fields. In cluster analysis, data is partitioned into multiple groups or clusters that share some traits common within their group. The authors present different clustering techniques (e.g., K-means, hierarchical clustering), and discuss the strengths and weaknesses of various clustering methods.

Tan, Zhiyong, Chai Quek, and Philip Y. K. Cheng, 2011, Stock trading with cycles: A financial application of anfis and reinforcement learning, *Expert Systems With Applications* 38, 4741–4755.

This study develops a new non-arbitrage trading algorithm based on adaptive network fuzzy inference system (ANFIS) and reinforcement learning techniques. The proposed method predicts the changes in long-term movement of prices and is able to outperform trading algorithms such as DENFIS and RSPOP. Experimental trading outcomes using five U.S. stocks indicate that, on average, the total returns using the new framework is higher by about 50 percentage points.

Terasvirta, Timo, Dick Van Dijk, and Marcelo C. Medeiros, 2005, Linear models, smooth transition autoregressions, and neural networks for forecasting macroeconomic time series: A re-examination, *International Journal of Forecasting* 21, 755–774.

The authors consider three techniques, namely linear autoregressive, smooth transition autoregressive (STAR), and neural network time series models, to forecast macroeconomic variables. Using a dynamic model specification, the results indicate that the dynamic STAR model performs better than the linear autoregressive and several fixed STAR models in terms of forecast accuracy. Neural network models can produce more accurate forecasts when the forecast horizon is long and when the Bayesian regularization is applied to the model.

The Financial Stability Board, 2017, Artificial intelligence and machine learning in financial services, Available at: <http://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service>.

This report by the Financial Stability Board examines the implications arising from the application of AI and machine learning methods in financial services. The report presents some background on the

recent rise in the use of AI and machine learning in various finance applications and services, discusses possible effects on the financial system, and assesses risks to its stability. It concludes that in the absence of 'audibility' and interpretability, AI and machine learning techniques may pose macro-level risks and, therefore, should be monitored by micro-prudential supervisors.

Tibshirani, Robert, 1996, Regression shrinkage and selection via the LASSO, *Journal of the Royal Statistical Society. Series B (methodological)* 58(1), 267–288.

The author introduces a new method called LASSO (least absolute shrinkage and selection operator) in linear regression models. By adding a penalty term to the mean-squared error minimization problem, lasso shrinks some coefficients to zero and produces models that are interpretable. Compared to other variable selection techniques, such as subset selection and ridge regression, LASSO does best for small to moderate-sized numbers of moderate-sized effects.

Tsai, Chih-Fong, Yueh-Chiao Lin, David C. Yen, and Yan-Min Chen, 2011, Predicting stock returns by classifier ensembles, *Applied Soft Computing* 11, 2452–2459.

In this article, the authors compare the prediction performance of ensemble models of classification for stock returns. They find that relative to single classifiers, classifier ensembles perform well in terms of return on investment and prediction accuracy.

Tsai, Chih-Fong, and Jhen-Wei Wu, 2008, Using neural network ensembles for bankruptcy prediction and credit scoring, *Expert Systems With Applications* 34, 2639–2649.

The authors compare the forecast accuracy of multiple neural network classifiers with that of a single best neural network classifier for bankruptcy prediction and credit scoring problems. The empirical results show that multiple classifiers tend to perform worse than a single best neural network classifier. However, if type I or type II errors are considered, neither method outperforms the other in terms of prediction accuracy.

Vapnik, Vladimir, 2000, *The Nature of Statistical Learning Theory* (New York: Springer).

The book reviews statistical learning for small data samples that do not rely on a priori information. It is divided into three parts: the general theory of learning, support vector estimation, and statistical foundations of learning theory. The book is appropriate to be used as a graduate-level textbook on learning theory.

Varetto, Franco, 1998, Genetic algorithms applications in the analysis of insolvency risk, *Journal of Banking and Finance* 22, 1421–1439.

This paper studies insolvency risk and compares results obtained from traditional linear discriminant analysis (LDA) and genetic algorithms (GA). Using data on Italian companies from 1982–1995, the author find that LDA provides better prediction of insolvency. However, the GA produces results faster and with less contribution from the financial analyst, and can therefore serve as an effective tool in bankruptcy risk analysis.

Verikas, Antanas, Zivile Kalsyte, Marija Bacauskiene, and Adas Gelzinis, 2010, Hybrid and ensemble-based soft computing techniques in bankruptcy prediction: a survey, *Soft Computing* 14, 995–1010.

The authors present a survey of literature on hybrid and ensemble-based soft computing techniques used bankruptcy prediction studies. Because the literature on bankruptcy prediction rarely reports confidence intervals of prediction results and studies use vastly different data, comparisons of the obtained results are not feasible. Instead, the paper reviews specific techniques and their ensembles used in bankruptcy prediction problems.

Vui, Chang Sim, Gan Kim Soon, Chin Kim On, Rayner Alfred, and Patricia Anthony, 2013, A review of stock market prediction with artificial neural network, *International Conference on Control System, Computing and Engineering (ICCSCE)*, 477–482, IEEE.

Artificial neural networks (ANN) can be a useful technique for stock market prediction due to the non-linear and volatile nature of stock market dynamics. The authors provide a brief review of the literature on the application of various ANN approaches in predicting stock market returns.

Xing, Frank Z., Erik Cambria, and Roy E. Welsch, 2018, Natural language based financial forecasting: a survey, *Artificial Intelligence Review* 50, 49–73.

In recent years, the number of papers utilizing textual sentiment data for financial forecasting has been increasing. This paper reviews the natural language based financial forecasting (NLFF) literature. In particular, it discusses the history of NLP techniques, current text processing techniques, algorithms for predictive models.

Xue, Jingming, Qiang Liu, Miaomiao Li, Xinwang Liu, Yongkai Ye, Siqi Wang, and Jianping Yin, 2018, Incremental multiple kernel extreme learning machine and its application in robo-advisors, *Soft Computing* 22, 3507–3517.

Robo-advisors currently used in the finance industry provide investors with financial advice previously provided by finance sector employees. The authors suggest that the machine learning algorithm used by robo-advisors maybe less suitable when information is heterogeneous. They introduce an incremental multiple kernel extreme learning machine (IMK-ELM) model, which can initialize and simultaneously update the training dataset and combine information from different data sources. The proposed method is able to efficiently solve classification problems and can thus be used as an algorithm for robo-advisors.

Yao, Jingtao, Yili Li, and Chew Lim Tan, 2000, Option price forecasting using neural networks, *Omega* 28, 455–466.

The authors of this study use a neural network option pricing model and show that the proposed approach achieves different performance results depending on how the data is partitioned into groups based on moneyness. Using Japanese 225 futures data, the authors find that a neural network pricing model outperforms the Black-Scholes model when markets are volatile, while the latter does better when its theoretical assumption of constant volatility holds. The traditional Black-Scholes model performs better for at-the-money options. For in-the-money and out-of-the-money options, a neural network pricing model may be more appropriate when high risk and high return is the preferred strategy.

Yu, Lean, Shouyang Wang, and Kin Keung Lai, 2008, Neural network-based mean-variance-skewness model for portfolio selection, *Computers and Operations Research* 35, 34–46.

Researchers have been studying extensions of Markowitz's mean-variance model of portfolio optimization. Recent work in this area recommends incorporating higher order moments, such as the skewness, especially when asset returns are not normally distributed. This study presents a neural network-based mean-variance-skewness model of portfolio optimization. The authors integrate this model with investors' risk preferences, different forecasts, and trading strategies and show that the proposed algorithm is computationally fast and efficient in solving the triple trade-offs in the mean-variance-skewness portfolio optimization problem.

Yu, Lean, Shouyang Wang, Kin Keung Lai, and Fenghua Wen, 2010, A multiscale neural network learning paradigm for financial crisis forecasting, *Neurocomputing* 73, 716–725.

This study proposes a novel approach to predicting exchange rates. The authors develop a multiscale neural network learning algorithm for the exchange rates, which are decomposed into multiple inde-

pendent intrinsic mode components using the Hilbert-EMD (empirical mode decomposition) algorithm. The findings show that the new EMD-based multiscale neural network learning approach performs well in forecasting bank crises and is superior to other classification methods.

Zetsche, Dirk, Douglas Arner, Ross Buckley, and Brian W. Tang, 2020, *Artificial Intelligence in finance: Putting the human in the loop*, Available at SSRN 3531711.

This paper develops a regulatory roadmap for the use of AI in finance, in particular focusing on human responsibility and highlighting the necessity of human involvement. After describing various use cases of AI in finance, it discusses a range of potential issues, then considers the regulatory challenges and tools available. The key issues identified are increased information asymmetries, data dependencies, and system interdependencies leading to unexpected consequences.

Zhang, Gioqinang, and Michael Y Hu, 1998, Neural network forecasting of the British pound/US dollar exchange rate, *Omega* 26, 495–506.

Numerous studies successfully apply neural networks to exchange rate forecasting. In this study, the authors look at the impact of the number of parameters to be estimated on exchange rate forecasting because neural networks require the estimation of many parameters. Using the British pound/U.S. dollar exchange rate forecasting problem as an empirical exercise, the study finds that the number of input nodes and hidden nodes affect forecasting performance of neural networks.

Zhang, Guoqiang, Michael Y Hu, B Eddy Patuwo, and Daniel C Indro, 1999, Artificial neural networks in bankruptcy prediction: General framework and cross-validation analysis, *European Journal of Operational Research* 116, 16–32.

A comprehensive review of neural network methods used in bankruptcy prediction studies finds that these methods are superior to logistic regression models in bankruptcy forecasting and classification. The authors explain that the superior performance of neural networks is due to their link to Bayesian posterior probabilities.

Zheng, Ban, Eric Moulines, and Frederic Abergel, 2013, Price jump prediction in a limit order book, *Journal of Mathematical Finance* 3, 242–255.

In this paper, the authors study the empirical relationship between bid/ask limit order liquidity balance and trade direction and the usefulness of information contained in limit order books. The results show that bid/ask liquidity balance helps to predict the direction of the next trade. Further, it examines the intertrade price jump using logistic regression where the most informative features in limit order book are selected by the LASSO variable selection technique. Empirical results using data on French stocks show that trade sign, market order sign, and liquidity on the best bid/ask are all important factors for price jump prediction.

Zimmermann, Hans-Georg, Ralph Neuneier, and Ralph Grothmann, 2002, Active portfolio management based on error correction neural networks, in *Advances in Neural Information Processing Systems*, 1465–1472.

This study combines the Black-Litterman portfolio optimization model with neural networks to forecast excess returns. The forecasts of the expected return are based on error correction neural networks (ECNN) that utilize the last model's error. Using data from 21 financial markets in G7 countries, the proposed portfolio optimization model is shown to outperform a benchmark portfolio.