Decision analytics and machine learning in economic and financial systems

Qifeng Qiao¹ · Peter A. Beling²

Published online: 24 May 2016
© Springer Science+Business Media New York 2016

Decision analytics may be viewed as the combined use of predictive modeling techniques (e.g., forecasting and machine learning) and prescriptive decision frameworks (e.g., optimization and simulation) to create value in systems. Recent years have seen an explosion in decision analytics applications, driven by advances in machine learning methods and computational optimization and by massive increases in the data to which these techniques may be applied. Decision analytics has long been used in the domains of economic and financial systems, with credit scoring being an example of an early success, and the clear trend is to the development of ever more sophisticated methods and applications.

This issue of *Environment System & Decisions* is devoted to decision analytics in financial and economic systems and environments. The papers in the issue each bear an influence from theoretical developments and financial and economic applications of statistical methods, machine learning, and decision and risk modeling techniques. Before introducing the papers, we review some of the influential literature in this regard.

1 Forecasting in financial markets

Historical data and nonrandom price movement provide opportunities for technicians to develop automated algorithms to predict prices. This was noted by LeBaron et al.

(1992), who showed that much can be learned from a simulated stock market with simulated traders. Leigh et al. (2002) conducted experiments focused on stock market price forecasting, including conventional pattern recognition techniques of template matching to identify price–volume pattern market direction signals and neural network and genetic algorithm techniques to forecast price changes for the NYSE composite index. Those early papers were followed by a flood of research in the areas of data modeling tools, machine learning methods, high-performance computing, and big data management. The capabilities in these areas developed over the last decade have created an unprecedented opportunity for decision systems to offer valuable insights into complex problems in business world.

In the context of trading, the forecasting of trends in stock prices is a decision support process. Empirical results have shown that artificial neural networks (ANNs) perform better than linear regression since stock markets are dynamic and chaotic (Trippi 1995). White (1988) showed how to use neural networks to search for and decode nonlinear regularities in asset price movements. The experiment focused on the case of IBM common stock daily returns and showed that the finding evidence against efficient market with simple networks is not easy, but on the positive side, such simple networks are capable of extremely rich dynamic behavior. Price prediction using ANNs is commonly done with backpropagation, a training algorithm in which steepest descent gradient is used to learn optimal network parameters. Gradient search is not a robust global optimization method, as the process can converge to a local optimum. In an attempt to develop more robust approaches to neural network training, researchers have investigated genetic algorithms, simulated annealing, and other methods that embed mechanisms for escaping local minima (see, e.g., Chen et al. 1991;

✉ Peter A. Beling
pb3a@virginia.edu

¹ NetApp, 495 East Java Drive, Sunnyvale, CA 94089, USA

² Department of Systems and Information Engineering,
University of Virginia, Charlottesville, VA 22904, USA

McInerney and Dhawan (1992). ANNs are studied in Trippi and DeSieno (1992) in the context of an attempt to build a trading system for S&P 500 index futures that can outperform passive investment in the index. Grudnitski and Osburn (1993) investigated ANN's ability to forecast the S&P 500 index and Gold futures, taking as input data historical prices as well as open interest patterns that aim to characterize the beliefs of the traders in the corresponding market. Baba and Kozaki (1992) and Armano et al. (2005), among many others, continued to combine technical indicators and ANNs to forecast stock indices. Along similar lines, Tan and Yao (2000) applied ANNs in foreign exchange rates, forecasting relationships between the dollar and five other major currencies. The results illustrated the difficulty, expected from finance theory, in making profit using a predictive model in an efficient market.

Some studies, such as Wood and Dasgupta (1996) and Zemke (1999), have focused on predicting the up and down directions of market indices. Zemke (1999) found that a nearest neighbor method outperformed naive Bayes classifiers and a genetic algorithm that evolved classification rules. Chen et al. (2003) modeled and predicted the direction of return on market index of the Taiwan Stock Exchange, showing that probabilistic neural network-based investment strategies obtain higher returns than other strategies based on random walks and generalized methods of moments with Kalman filter. Instead of formulating the direction prediction problem as a binary classification problem, Tilakaratne et al. (2008) developed a neural network algorithm to predict trading signals in three classes: buy, hold, and sell, when the distribution of these signals is imbalanced.

Tio et al. (2005) and Hasbrouck (2007), among many others, have applied machine learning methods to model limit order book dynamics. Fletcher and Shawe-Taylor (2012) applied support vector machines (SVMs) to predict the direction of price movement of a currency. Kercheval and Zhang (2015) developed a SVM-based machine learning framework to model the dynamics of high-frequency limit order books in financial equity markets and automate real-time prediction of metrics such as mid-price movement and price spread crossing.

2 Trading strategy recognition

Algorithmic traders design their trading algorithms and systems to obtain consistent returns under different market conditions. The most commonly used measure of effectiveness for trading systems is the Sharpe ratio (Sharpe 1994). Mullei and Beling (1998) used genetic algorithms to learn technical decision rules that then are the basis for

trading decisions. Similarly, Allen and Karjalainen (1999) evolved rules to determine the days that are likely to give a positive return with low volatility. Moody (2002) proposed an adaptive algorithm called recurrent reinforcement learning to discover investment policies and test in real applications intra-daily currency trading and a monthly asset allocation system for the S&P 500 stock index. Other work, such as that by Chung et al. (2004), Kim and Shin (2007), and Lee et al. (2010), focused on enhancing the trading system by employing machine learning methods to predict stock price patterns.

In the past decade, high-frequency trading strategies have attracted significant attention from investors, regulators, and policy makers. Although many HFT strategies exist today, they are largely unknown to public. Recently, researchers have begun to shed some light on the general characteristics of these strategies. Several illustrative HFT strategies include (1) acting as an informal or formal market-maker, (2) high-frequency relative-value trading, and (3) directional trading on news releases, order flow, or other high-frequency signals (Jones 2012). Research on specific HFT trading strategies falls within the larger theme of machine learning-based approaches to identifying and characterizing behavior in algorithmic trading. Studies in this vein examine the activities of different types of algorithms and traders and attempt to recognize their behavioral differences. Hendershott and Riordan (2013) summarized the role of algorithmic traders in liquidity supply and demand in the 30 Deutsche Aktien Index stock and studied the classification problem of distinguishing algorithmic traders from humans. Kirilenko et al. (2011) classified traders into several categories, including HFTs, opportunistic traders, fundamental traders, and noise traders. Hayes et al. (2013) used supervised and unsupervised learning to reverse engineer fund allocation strategies for groups of human participants in simulated trading competition. As documented by Yang et al. (2012) and Hayes et al. (2012), algorithmic trading strategies can be monitored by market operators and regulators to prevent unfair trading practices and improve the health of the financial markets. Qiao and Beling (2013) proposed a general approach to behavior recognition in sequential decision problems that is based on Markov decision process (MDP) models and Gaussian process inverse reinforcement learning (cf. Qiao and Beling 2011). Yang et al. (2015) applied that approach to algorithmic trading, modeling the trading strategies in terms of an MDP and then learning trader behavior in the space of reward functions learned through inverse reinforcement learning. Empirical results on E-Mini S&P 500 futures market show that the machine learning-based approach provides significant and consistent improvement on previous rule-based classification methods.

3 Financial crisis and risk modeling

Models of the cause and development of financial crises have been topics of much research, especially in the years since the 2008. The risk of a banking crisis suggests the importance of identifying banks with potential problems before they face liquidity or solvency crises. Bank failures may result from poor management practices, expanded risk-taking, interest rate volatility, inadequate accounting accounts, and increased competition (Miletic 2008). Ecer (2013) compared the ability of artificial neural networks and support vector machines in predicting bank failures. A notable component of the literature on financial crises is comprised of decision analytic methods that bear on consumer and corporate credit risk, since unexpected levels of credit default or other mis-pricing of risk can result in substantial shocks to the financial system.

Machine learning techniques have been studied widely as tools for default and bankruptcy prediction in both the consumer and corporate spaces, as both prediction problems can be formulated as a binary classification problem that assigns a good or bad risk label to a new observation. Zhong et al. (2014), to take an example, compared the learning effectiveness of backpropagation, extreme learning machines, support vector machines (SVMs), and ANNs for corporate credit ratings, and showed that ANNs and SVMs have demonstrated remarkable performance. Other studies focused on the impact of feature selection in credit scoring models (e.g., Liu and Schumann 2005) and bankruptcy prediction (e.g., Tsai 2009; Xu and Wang 2009). Classifier ensembles and hybrid classifiers have been used to improve the single classifier's performance in the prediction of financial crises through majority voting ensembles (Li and Sun 2009); k -NN ensembles (Paleologoa et al. 2010; Twala 2010); logistic regression ensembles (Twala 2010); SVM ensembles (Nanni and Lumini 2009; Paleologoa et al. 2010); and neural network ensembles (West 2000; West et al. 2005; Kim and Kang 2010). Hybrid classifiers are widely applied because, relative to other methods, they consume less computational resources and are less demanding in terms of complexity of choosing classifiers and training data sets.

Gestel et al. (2006) applied least-squares support vector machine classifiers within a Bayesian evidence framework to infer and analyze the creditworthiness of potential corporate clients. The inferred posterior class probabilities of bankruptcy were used to analyze the sensitivity of the classifier output and assist in the credit assignment decision-making process. A hybrid model of neural network and adaptive boosting method was studied by Xu et al. (2015). Sexton and McMurtrey (2006) and Setiono et al. (2009) studied a genetic algorithm-based neural network algorithm for credit card screening and found that the

neural network rule extraction is very effective in discovering knowledge and is particularly appropriate in applications that require comprehensibility and accuracy. A two-stage hybrid model based on artificial neural networks and multivariate adaptive regression splines (MARs) was proposed for credit scoring by Lee and Chen (2005). The output of MARs is used as the input variables of the designed neural networks.

Consumer credit outstanding exceeded \$14 trillion in the fourth quarter of 2015; the opportunities and risk exposures in that space are equally outsized. The complicated nature and large number of decisions involved in the consumer lending business make it necessary to use algorithms to automate risk assessment and management for individuals. It is common for lending institutions to create their models based on the information from the credit files collected by credit bureau agencies and private information regarding borrowers' previous behavior. Credit scoring has been successfully used in practice for more than four decades. Desai et al. (1996) compared a number of alternative to traditional credit scoring methods, highlighting an increasing interest in the application of machine learning techniques in the assessment of individual credit risk. To date, a rich variety of learning methods have been studied in this area. Ong et al. (2005), for example, concluded that genetic algorithms outperform the models based on ANNs, decision trees, and logistic regression. Martens et al. (2007) proposed a rule-extraction method for SVMs to generate consumer credit models with human interpretable rules and little loss in accuracy. Amir et al. (2010) used generalized classification and regression trees to construct credit risk management models.

4 In this issue

The well-known parameter *Beta* provides a measure of the volatility of an asset relative to the market, a quantity that is needed in many investment decisions. Baker et al. (2016) study methods for estimating Beta and ultimately suggest a new method with desirable characteristics. Following the same vein, Laird-Smith et al. (2016) consider an alternative to Beta that provides a more realistic and stable estimator for market-related risk and return.

The central banks of nations often use short-term interest rates to drive monetary policy and influence investment decisions. Shaw et al. (2016) use stochastic process models to explore how jumps in market prices might predict central bank announcements.

Historically, stock performance models have been based on a narrow set of observables, such as trading activity (e.g., buy and sell orders), prices themselves, and the

fundamentals of companies. Today, through new media sources and data sets, it is possible to form direct estimates of the relationships between people's opinions of a stock and its price. Mo et al. (2016) study a feedback hypothesis that news sentiment drives trading activity and that trading activity generates publicity, influencing the news.

The advent of algorithmic trading has vastly expanded the quantity and complexity of the data that stock market regulators must consider in operational and policy-level decision making. Paddrik et al. (2016) consider a variety of visualization tools to support regulatory decisions in markets dominated by electronic trading.

Three papers in the issue bear on decision analytics in economic and industrial environments and systems. Danielle et al. (2016) consider economic planning under environmental risk and uncertainty, developing a fuzzy optimization model that takes inputs from experts and provides support to decision makers. In manufacturing environments, operational and financial decisions are often driven by key performance indicators (KPIs), which are productivity and economic metrics. Collins et al. (2016) present a new approach and insight into the improvement of these important decision analytic tools. Hopkins (2016) considers a fundamental economic decision problem, cost-benefit analysis, studying the characteristics of the problem in several cases studies in environmental project planning.

The papers in this special issue of *Environment Systems & Decisions* span a range of decision problems, analytic approaches, and domains. We hope this collection serves to illustrate the rich nature of decision analytics research in economic and financial systems.

References

- Allen F, Karjalainen R (1999) Using genetic algorithms to find technical trading rules. *J Financ Econ* 51:245–271
- Amir K, Kim A, Lo A (2010) Consumer credit-risk models via machine-learning algorithms. *J Bank Finance* 34:2767–2787
- Armano G, Marchesi M, Murru A (2005) A hybrid genetic-neural architecture for stock indexes forecasting. *Comput Intell Econ Finance* 170:3–33
- Baba N, Kozaki M (1992) An intelligent forecasting system of stock price using neural networks. In: International joint conference on neural networks
- Baker C, Rajaratnam K, Flint EJ (2016) Beta estimates of shares on the jse top 40 in the context of reference-day risk. *Environ Syst Decis* 37(1). doi:10.1007/s10669-016-9595-4
- Chen AS, Leung MT, Daouk H (2003) Application of neural networks to an emerging financial market: forecasting and trading the taiwan stock index. *Comput Oper Res* 30:901–923
- Chen S, Cowan C, Grant P (1991) Orthogonal least squares learning algorithm for radial basis function. *IEEE Trans Neural Netw* 2:302–309
- Chung F, Fu T, Ng VTY, Luk RWP (2004) An evolutionary approach to pattern-based time series segmentation. *IEEE Trans Evol Comput* 8(5):471–489
- Collins AJ, Hester P, Ezell B, Horst J (2016) An improvement selection methodology for key performance indicators. *Environ Syst Decis* 37(1). doi:10.1007/s10669-016-9591-8
- Danielle K, Yu S, Aviso KB, Promentilla MA, Santos J, Tan R (2016) A weighted fuzzy linear programming model in economic input–output analysis: an application to risk management of energy system disruptions. *Environ Syst Decis* 37(1). doi:10.1007/s10669-016-9599-0
- Desai V, Crook J, Overstreet G (1996) A comparison of neural networks and linear scoring models in the credit union environment. *Eur J Oper Res* 95(1):24–37
- Ecer F (2013) Comparing the bank failure prediction performance of neural networks and support vector machines: the Turkish case. *Econ Res* 26(3):81–89
- Fletcher T, Shawe-Taylor J (2012) Multiple kernel learning with fisher kernels for high frequency currency prediction. *Comput Econ* 42(2):217–240
- Gestel TV, Baesens B, Suykens JA, den Poel DV, Baestaens DE, Willekens M (2006) Bayesian kernel based classification for financial distress detection. *Eur J Oper Res* 172(3):979–1003
- Grudnitski G, Osburn L (1993) Forecastings and gold futures prices: an application of neural networks. *J Futures Markets* 13:631–643
- Hasbrouck J (2007) Empirical market microstructure: the institutions, economics, and econometrics of securities trading. Oxford University Press, New York
- Hayes R, Paddrik M, Todd A, Yang S, Beling PA, Scherer W (2012) Agent based model of the e-mini future: Application for policy making. In: Proceedings of the Winter simulation conference, Winter simulation conference, WSC '12, pp 111:1–111:12
- Hayes RL, Beling PA, Scherer WT (2013) Action-based feature representation for reverse engineering trading strategies. *Environ Syst Decis* 33(3):413–426
- Hendershott T, Riordan R (2013) Algorithmic trading and the market for liquidity. *J Financ Quant Anal* 48:1001–1024
- Hopkins E (2016) Program efficiency analysis within the environmental landscape. *Environ Syst Decis* 37(1). doi:10.1007/s10669-016-9594-5
- Jones C (2012) What do we know about high-frequency trading? Columbia Business School Research Paper No. 13-11
- Kercheval AN, Zhang Y (2015) Modelling high-frequency limit order book dynamics with support vector machines. *Quant Finance* 15:1315–1329
- Kim H, Shin K (2007) A hybrid approach based on neural networks and genetic algorithms for detecting temporal patterns in stock markets. *Appl Soft Comput* 7:569–576
- Kim MJ, Kang DK (2010) Ensemble with neural networks for bankruptcy prediction. *Expert Syst Appl* 37:3373–3379
- Kirilenko A, Kyle A, Samadi M, Tuzun T (2011) The flash crash: the impact of high frequency trading on an electronic market. Available at SSRN. <http://ssrn.com/abstract=1686004>
- Laird-Smith J, Meyer K, Rajaratnam K (2016) A study of total beta specification through symmetric regression: the case of the Johannesburg Stock Exchange. *Environ Syst Decis* 37(1). doi:10.1007/s10669-016-9596-3
- LeBaron B, Brock WA, Lakonishok J (1992) Simple technical trading rules and the stochastic properties of stock returns. *J Finance* 47:1731–1764
- Lee SJ, Ahn JJ, Oh KJ, Kim TY (2010) Using rough set to support investment strategies of real-time trading in futures market. *Appl Intell* 32(3):364–377
- Lee TS, Chen IF (2005) A two-stage hybrid credit scoring model using artificial neural networks and multivariate adaptive regression splines. *Expert Syst Appl* 28:743–752
- Leigh W, Purvis R, Ragusa JM (2002) Forecasting the nyse composite index with technical analysis, pattern recognizer, neural network,

- and genetic algorithm: a case study in romantic decision support. *Decis Support Syst* 32:361–377
- Li H, Sun J (2009) Majority voting combination of multiple case-based reasoning for financial distress prediction. *Expert Syst Appl* 36(3):4363–4373
- Liu Y, Schumann M (2005) Data mining feature selection for credit scoring models. *J Oper Res Soc* 56(9):1099–1108
- Martens D, Baesens B, Gestel TV, Vanthienen J (2007) Comprehensive credit scoring models using rule extraction from support vector machines. *Eur J Oper Res* 183:1466–1476
- McInerney M, Dhawan AP (1992) Use of genetic algorithms with back-propagation in training of feed-forward neural networks. In: *IEEE International joint conference on neural networks*, pp 203–208
- Miletic I (2008) Supervision of banks and the stability of the banking system. *Econ Res* 21(3):43–56
- Mo S, Liu A, Yang SY (2016) News sentiment to market impact and its feedback effect. *Environ Syst Decis* 37(1). doi:[10.1007/s10669-016-9590-9](https://doi.org/10.1007/s10669-016-9590-9)
- Moody J (2002) Learning to trade via direct reinforcement. *IEEE Trans Neural Netw* 12:875–889
- Mullei S, Beling PA (1998) Hybrid evolutionary algorithms for a multiobjective financial problem. In: *Systems, man, and cybernetics. 1998 IEEE international conference on*, vol 4, pp 3925–3930
- Nanni L, Lumini A (2009) An experimental comparison of ensemble of classifiers for bankruptcy prediction and credit scoring. *Expert Syst Appl* 36:3028–3033
- Ong C, Huang J, Tzeng G (2005) Building credit scoring systems using genetic programming. *Expert Syst Appl* 29:41–47
- Paddrik ME, Haynes R, Todd AE, Scherer WT, Beling PA (2016) Visual analysis to support regulators in electronic order book markets. *Environ Syst Decis* 37(1). doi:[10.1007/s10669-016-9597-2](https://doi.org/10.1007/s10669-016-9597-2)
- Paleologoa G, Elisseeffb A, Antoninib G (2010) Subagging for credit scoring models. *Eur J Oper Res* 201(2):490–499
- Qiao Q, Beling PA (2011) Inverse reinforcement learning with gaussian process. In: *Proceedings of the 2011 American control conference*, pp 113–118. doi:[10.1109/ACC.2011.5990948](https://doi.org/10.1109/ACC.2011.5990948)
- Qiao Q, Beling PA (2013) Recognition of agents based on observation of their sequential behavior. *Mach Learn Knowl Discov Databases* 8188:33–48
- Setiono R, Baesens B, Mues C (2009) A note on knowledge discovery using neural networks and its application to credit card screening. *Eur J Oper Res* 192:326–332
- Sexton RS, McMurtrey S (2006) Knowledge discovery using a neural network simultaneous optimization algorithm on a real world classification problem. *Eur J Oper Res* 168:1009–1018
- Sharpe WF (1994) The sharpe ratio. *J Portf Manag* 21:49–58
- Shaw F, Murphy F, O'Brien FG (2016) Jumps in euribor and the effect of ECB monetary policy announcements. *Environ Syst Decis* 37(1). doi:[10.1007/s10669-016-9600-y](https://doi.org/10.1007/s10669-016-9600-y)
- Tan C, Yao T (2000) A case study on using neural networks to perform technical forecasting of forex. *Neurocomputing* 34:79–98
- Tilakaratne CD, Mammadov MA, Morris SA (2008) *AI 2008: advances in artificial intelligence: 21st Australasian joint conference on artificial intelligence Auckland, New Zealand. Proceedings.* Springer, Berlin, chap Predicting Trading Signals of Stock Market Indices Using Neural Networks, pp 522–531
- Tio P, Nikolaev N, Yao X (2005) Volatility forecasting with sparse bayesian kernel models. In: *4th International conference on computational intelligence in economics and finance*, pp 1052–1058
- Trippi R, DeSieno D (1992) Trading equity index futures with a neural network. *J Portf Manag* 19(1):27–33
- Trippi RR (1995) *Artificial intelligence in finance and investing: state-of-the-art technologies for securities selection and portfolio management*, 1st edn. McGraw-Hill, Inc., New York, NY
- Tsai CF (2009) Feature selection in bankruptcy prediction. *Knowl Based Syst* 22:120–127
- Twala B (2010) Multiple classifier application to credit risk assessment. *Expert Syst Appl* 37(3):3326–3336
- West D (2000) Neural network credit scoring models. *Comput Oper Res* 27(11):1131–1152
- West D, Dellana S, Qian J (2005) Neural network ensemble strategies for financial decision applications. *Comput Oper Res* 32(10):2543–2559
- White H (1988) Economic prediction using neural networks: the case of ibm daily stock returns. In: *IEEE international conference on in neural networks*
- Wood D, Dasgupta B (1996) Classifying trend movements in the MSCI U.S.A. capital market index—a comparison of regression, arima and neural network methods. *Comput OR* 23(6):611–622
- Xu X, Wang Y (2009) Financial failure prediction using efficiency as a predictor. *Expert Syst Appl* 36(1):366–373
- Xu X, Wang G, Ding S, Jiang X, Zhao Z (2015) A new method for constructing granular neural networks based on rule extraction and extreme learning machine. *Pattern Recognit Lett* 67(Part 2):138–144
- Yang S, Qiao Q, Beling PA, Scherer WT, Kirilenko AA (2015) Gaussian process-based algorithmic trading strategy identification. *Quant Finance* 15:1683–1073
- Yang SY, Paddrik ME, Hayes R, Todd A, Kirilenko A, Beling PA, Scherer WT (2012) Behavior based learning in identifying high frequency trading strategies. In: *Proceedings of the 2012 IEEE conference on computational intelligence for financial engineering & economics, CIFE 2012.* New York City, NY, pp 1–8
- Zemke S (1999) Nonlinear index prediction. *Phys A Stat Mech Appl* 269:177–183
- Zhong H, Miao C, Shen Z, Feng Y (2014) Comparing the learning effectiveness of BP, ELM, I-ELM, and SVM for corporate credit ratings. *Neurocomputing* 128:285–295