Expert Systems With Applications 83 (2017) 187–205

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies

Eunsuk Chong[a], Chulwoo Han[b], Frank C. Park[a,*]

[a] Robotics Laboratory, Seoul National University, Seoul 08826, Korea
[b] Durham University Business School, Mill Hill Lane, Durham DH1 3LB, UK

A R T I C L E I N F O

Article history:
Received 15 November 2016
Revised 15 April 2017
Accepted 16 April 2017
Available online 22 April 2017

Keywords:
Stock market prediction
Deep learning
Multilayer neural network
Covariance estimation

A B S T R A C T

We offer a systematic analysis of the use of deep learning networks for stock market analysis and prediction. Its ability to extract features from a large set of raw data without relying on prior knowledge of predictors makes deep learning potentially attractive for stock market prediction at high frequencies. Deep learning algorithms vary considerably in the choice of network structure, activation function, and other model parameters, and their performance is known to depend heavily on the method of data representation. Our study attempts to provides a comprehensive and objective assessment of both the advantages and drawbacks of deep learning algorithms for stock market analysis and prediction. Using high-frequency intraday stock returns as input data, we examine the effects of three unsupervised feature extraction methods—principal component analysis, autoencoder, and the restricted Boltzmann machine—on the network's overall ability to predict future market behavior. Empirical results suggest that deep neural networks can extract additional information from the residuals of the autoregressive model and improve prediction performance; the same cannot be said when the autoregressive model is applied to the residuals of the network. Covariance estimation is also noticeably improved when the predictive network is applied to covariance-based market structure analysis. Our study offers practical insights and potentially useful directions for further investigation into how deep learning networks can be effectively used for stock market analysis and prediction.

1. Introduction

Research on the predictability of stock markets has a long history in financial economics (*e.g.*, Ang & Bekaert, 2007; Bacchetta, Mertens, & Van Wincoop, 2009; Bondt & Thaler, 1985; Bradley, 1950; Campbell & Hamao, 1992; Campbell & Thompson, 2008; Campbell, 2012; Granger & Morgenstern, 1970). While opinions differ on the efficiency of markets, many widely accepted empirical studies show that financial markets are to some extent predictable (Bollerslev, Marrone, Xu, & Zhou, 2014; Ferreira & Santa-Clara, 2011; Kim, Shamsuddin, & Lim, 2011; Phan, Sharma, & Narayan, 2015). Among methods for stock return prediction, econometric or statistical methods based on the analysis of past market movements have been the most widely adopted (Agrawal, Chourasia, & Mittra, 2013). These approaches employ various linear and nonlinear methods to predict stock returns, *e.g.*, autoregressive models and artificial neural networks (ANN) (Adebiyi, Adewumi, & Ayo, 2014; Armano, Marchesi, & Murru, 2005; Atsalakis & Valavanis, 2009; Bogullu, Dagli, & Enke, 2002; Cao, Leggio, & Schniederjans, 2005; Chen, Leung, & Daouk, 2003; Enke & Mehdiyev, 2014; Guresen, Kayakutlu, & Daim, 2011a; Kara, Boyacioglu, & Baykan, 2011; Kazem, Sharifi, Hussain, Saberi, & Hussain, 2013; Khashei & Bijari, 2011; Kim & Enke, 2016a; 2016b; Monfared & Enke, 2014; Rather, Agarwal, & Sastry, 2015; Thawornwong & Enke, 2004; Ticknor, 2013; Tsai & Hsiao, 2010; Wang, Wang, Zhang, & Guo, 2011; Yeh, Huang, & Lee, 2011; Zhu, Wang, Xu, & Li, 2008). While there is uniform agreement that stock returns behave nonlinearly, many empirical studies show that for the most part nonlinear models do not necessarily outperform linear models: *e.g.*, Lee, Sehwan, and Jongdae (2007), Lee, Chi, Yoo, and Jin (2008), Agrawal et al. (2013), and Adebiyi et al. (2014) propose linear models that outperform or perform as well as nonlinear models, whereas Thawornwong and Enke (2004), Cao et al. (2005), Enke and Mehdiyev (2013), and Rather et al. (2015) find nonliner models outperfrom linear models. Table 1 provides a summary of recent works relevant to our research. For more exhaustive and detailed reviews, we refer

* Corresponding author.
E-mail addresses: bear3498@snu.ac.kr (E. Chong), chulwoo.han@durham.ac.uk (C. Han), fcp@snu.ac.kr (F.C. Park).

http://dx.doi.org/10.1016/j.eswa.2017.04.030

Table 1
A summery of recent studies on stock market prediction.

Authors (Year)	Data type (Num. of input features × lagged times)	Target output	Num. of samples (Training: Validation: Test)	Sampling period (Frequency)	Method	Performance measure
Enke and Mehdiyev (2013)	US S&P 500 index (20 × 1)	Stock price	361	Jan-1980 to Jan-2010 (daily)	Feature selection+fuzzy clustering+fuzzy NN	RMSE
Niaki and Hoseinzade (2013)	Korea KOSPI200 index (27 × 1)	Market direction (up or down)	3650 (8:1:1)	1-Mar-1994 to 30-Jun-2008 (daily)	Feature selection+ANN	statistical tests
Cervelló-Royo et al. (2015)	US Dow Jones index (1 × 10)	Market trend (bull/bear-flag)	91,307	22-May-2000 to 29-Nov-2013 (15-min)	Template matching	trading simulation
Patel, Shah, Thakkar, and Kotecha (2015)	India CNX and BSE indices (10 × 1)	Stock price	2393*	Jan-2003 to Dec-2012 (daily)	SVR+ {ANN, RF, SVR}	MAPE, MAE, rRMSE, MSE
T.-l. Chen and Chen (2016)	Taiwan TAIEX[a] and US NASDAQ[b] indices (27 × 20)	Market trend (bull-flag)	3818[a, *] 3412[b, *] (7:0:1)	7-Jan-1989 to 24-Mar-2004 (daily)	Dimension reduction+template matching	trading simulation
Chiang, Enke, Wu, and Wang (2016)	World 22 stock market indices ({3∼5} × 1)	Trading signal (stock price)	756 (2:0:1)	Jan-2008 to Dec-2010 (daily)	Particle swarm optimization +ANN	trading simulation
Chourmouziadis and Chatzoglou (2016)	Greece ASE general index (8 × 1)	Portfolio composition (cash:stock)	3907*	15-Nov-1996 to 5-Jun-2012 (daily)	Fuzzy system	trading simulation
Qiu, Song, and Akagi (2016)	Japan Nikkei 225 index (71 × 1)	Stock return	237 (7:0:3)	Nov-1993 to Jul-2013 (monthly)	ANN+{genetic algorithm, simulated annealing}	MSE
Arévalo, Niño, Hernández, and Sandoval (2016)	US Apple stock (3× {2∼15}+2)	Stock price	19,109 (17:0:3)	2-Sep-2008 to 7-Nov-2008 (1-minute)	Deep NN	MSE, directional accuracy
Zhong and Enke (2017)	US SPDR S&P 500 ETF (SPY) (60 × 1)	Market direction (up or down)	2518 (14:3:3)	1-Jun-2003 to 31-May-2013 (daily)	Dimension reduction+ANN	trading simulation, statistical tests
Our study	Korea KOSPI 38 stock returns (38 × 10)	Stock return	73,041 (3:1:1)	4-Jan-2010 to 30-Dec-2014 (5-minute)	Data representation+deep NN	NMSE, RMSE, MAE, MI

NN: neural network, SVR: support vector regression, RF: random forest, rRMSE: relative RMSE, NMSE: normalized MSE, MI: mutual information.
* In some studies the number of samples is not explicitly provided. We have calculated the number of samples based on each country's business days.

the reader to Atsalakis and Valavanis (2009), Guresen, Kayakutlu, and Daim (2011b), and Cavalcante, Brasileiro, Souza, Nobrega, and Oliveira (2016).

Recently there has been a resurgence of interest in artificial neural networks, in large part to its spectacular successes in image classification, natural language processing, and various time-series problems (CireşAn, Meier, Masci, & Schmidhuber, 2012; Hinton & Salakhutdinov, 2006; Lee, Pham, Largman, & Ng, 2009). Underlying this progress is the development of a feature learning framework, known as deep learning (LeCun, Bengio, & Hinton, 2015), whose basic structure is best described as a multi-layer neural network, and whose success can be attributed to a combination of increased computational power, availability of large datasets, and more sophisticated algorithms (Bengio, Lamblin, Popovici, Larochelle et al., 2007; Deng & Yu, 2014; Hinton, Osindero, & Teh, 2006; Salakhutdinov & Hinton, 2009; Srivastava, Hinton, Krizhevsky, Sutskever, & Salakhutdinov, 2014).

There has been growing interest in whether deep learning can be effectively applied to problems in finance, but the literature (at least in the public domain) still remains limited.[1] With the increasing availability of high-frequency trading data and the less-than-satisfactory performance of existing models, comprehensive studies that objectively examine the suitability of deep learning to stock market prediction and analysis seem opportune. The ability to extract abstract features from data, and to identify hidden nonlinear relationships without relying on econometric assumptions and human expertise, makes deep learning particularly attractive as an alternative to existing models and approaches.

ANNs require a careful selection of the input variables and network parameters such as the learning rate, number of hidden layers, and number of nodes in each layer in order to achieve satisfactory results (Hussain, Knowles, Lisboa, & El-Deredy, 2008). It is also important to reduce dimensionality to improve learning efficiency. On the other hand, deep learning automatically extracts features from data and requires minimal human intervention during feature selection. Therefore, our approach does not require expertise on predictors such as macroeconomic variables and enables us to use a large set of raw-level data as input. Ignoring the factors that are known to predict the returns, our approach may not be able to outperform existing models based on carefully chosen predictors. However, considering the fast growth of deep learning algorithms, we believe our research will serve as a milestone for the future research in this direction. Chourmouziadis and Chatzoglou (2016) also predict that deep learning will play a key role in financial time series forecasting.

We conjecture that, due to correlation, past stock returns affect not only its own future returns but also the future returns of other stocks, and use 380 dimensional lagged stock returns (38 stocks and 10 lagged returns) as input data letting deep learning extract features. This large input dataset makes deep learning a particularly suitable choice for our research.

[1] There exist a few studies that apply deep learning to identification of the relationship between past news events and stock market movements (Ding, Zhang, Liu, & Duan, 2015; Yoshihara, Fujikawa, Seki, & Uehara, 2014), but, to our knowledge, there is no study that apply deep learning to extract information from the stock return time series.