RECENT TRENDS IN MACHINE LEARNING WITH A FOCUS ON APPLICATIONS IN FINANCE

Jan Kalina – Aleš Neoral

Abstract

Machine learning methods penetrate to applications in the analysis of financial data, particularly to supervised learning tasks including regression or classification. Other approaches, such as reinforcement learning or automated machine learning, are not so well known in the context of finance yet. In this paper, we discuss the advantages of an automated data analysis, which is beneficial especially if a larger number of datasets should be analyzed under a time pressure. Important types of learning include reinforcement learning, automated machine learning, or metalearning. This paper overviews their principles and recalls some of their inspiring applications. We include a discussion of the importance of the concept of information and of the search for the most relevant information in the field of mathematical finance. We come to the conclusion that a statistical interpretation of the results of the automatic machine learning remains crucial for a proper understanding of the knowledge acquired by the analysis of the given (financial) data.

Key words: statistical learning, automated machine learning, metalearning, financial data analysis, stock market investing

JEL Code: C55, C45, G24

Introduction

Machine learning represents a subfield of artificial intelligence focused on algorithms and techniques allowing to learn, i.e. to extract knowledge from information in the form given data. The importance of numerous available machine learning tools for data analysis keeps increasing in economics and finance especially (but not only) with the advent of Big Data (Athey and Imbens, 2019). Machine learning methods for modeling of financial data include shallow or deep regression neural networks or time series prediction methods (Charpentier et al., 2021). Their applications naturally require an estimation of parameters in the considered models exploiting available tools of statistical learning theory (Kalina, 2013). Machine learning applications have been successful in various domains thanks to the power of deep learning (Deng et al., 2017), but also thanks to the development of computational technologies and massive possibilities for paralelizing the computations.

According to the type of learning, machine learning algorithms can be divided to several categories, which include supervised learning, unsupervised learning (clustering), combination of both these approaches, reinforcement learning, transfer learning, etc. If it is necessary to perform a quick analysis of a larger number of datasets, then it is desirable to perform the analysis as much automated as possible (Dixon et al., 2020). Therefore, broad classes of methods as well as whole systems for automated machine learning have been proposed and implemented. As an example, an automatic method selection (also known as automated selection problem, ASP) for robust regression over a database of 643 real datasets was recently presented (Kalina et al., 2021). This motivated the current work devoted to various types of machine learning.

The overview of available tools for automated machine learning represents the aim of this paper, while recent examples of economic applications will be presented as well. Section 2 presents and discusses supervised learning and Section 3 reinforcement learning. Section 4 is devoted to principles of metalearning. Section 5 discusses stock market investment decisions to be based on information and Section 6 is focused on the potential of machine learning for such investment decision.

1 Supervised learning

Supervised learning is a general concept encompassing regression and classification, i.e. two statistical tasks with a (continuous or categorical) response variable with an enormous number of applications. The importance of supervised learning methods for the analysis of data in economics and finance is well acknowledged and has been recently recalled and overviewed (Athey and Imbens, 2019). Remarkable examples of supervised learning methods include credit risk analysis, where the aim is to recommend whether a given individual/company should get a loan, or regression modeling and predictions for financial markets. Recommender systems as another example offer recommendations of products (e.g. services or advertisements) based on given data within e-business e.g. in e-shops.

One of common problems in supervised learning is overfitting, i.e. exploiting too specific features from given data, which cannot be generalized (extended) to new data generated from the same model under the same assumptions. An overfitted regression curve is too influenced by a random noise. A contrary of overfitting is underfitting, i.e. the phenomenon when the information from given data is not fully exploited. A general aim of supervised learning is to distinguish between the signal (relevant information) and noise, i.e. to ignore the random contamination. Such task can never be solved ideally due to the tradeoff between bias and variance. A common solution is to use cross-validation to assess a good (optimal) generalization ability of the given method beyond the training data. This can be at the same time used for finding optimal hyperparameters of a given supervised learning method. To explain the difference between parameters and hyperparameters, values of parameters are found within the supervised learning task itself, while values of hyperparameters are chosen before and are not modified during the learning process.

2 Reinforcement learning

Reinforcement learning is inspired by neuroscience (neurocomputing models) and particular by models describing the functioning of (not only human) brains. Without any given response variable, the learning is based primarily on rewarding in case of correct performance. The algorithms (like e.g. agents in robotics) are rewarded if they perform their tasks appropriately. While the quite general ideas of reinforcement learning are applicable to various tasks of economics and finance (Charpentier et al., 2021), let us now mention a couple of recent inspiring applications.

A mixture of high-frequency financial data including stock data as well as commodity future contracts related to silver and sugar was modeled in (Deng et al., 2017). Deep neural networks were used for the modeling and reinforcement learning was used as the learning approach for their training. The method turned out to perform in a robust way, allowing to outperform human traders experienced in financial asset trading. A predictive modeling of energy demand based on several predictors using a nonlinear regression model was presented in (Yin and Chao, 2018); supervised variable selection was performed in the paper to determine the most relevant predictors for a given time instance.

Within reinforcement learning, an agent attempts to maximize the total reward obtain from the environment that adopts such strategy. Typically, agents are allowed to start a random exploration of their environment and to perform some actions. They are either rewarded or punished according to their strategy. Their strategy is modified before the start of the next iteration, so the agent has a vision to obtain a better reward. The environment may be formally described as a Markovian decision process (MDP). This seems reasonable, because numerous algorithms of reinforcement learning exploit techniques of dynamic programming.

Formally, an MDP is completely characterized by the set (S, A, P, R), where

- S is a finite set of states of the environment and states of the agent,
- A is a finite set of actions of the agent,
- $P_a(s, s')$ is a transition function describing the probability that if the agent performs the action a in the state s, then his state becomes s',
- $R_a(s, s')$ is a function describing the reward given to the agent, if he performs the action a in the state s and his state becomes s'.

The behaviour of the agent is described by means of the strategy $\pi : S \times A \rightarrow (0,1)$, where $\pi(s, a)$ is the probability of performing the given action in the state s. The aim is to find such strategy π, which maximizes the total reward given to the agent; this is given by the sum

$$\sum_{t=0}^{\infty} \gamma^t R_a(s_t, s_{t+1}), \tag{1}$$

where $a = \pi(s_t)$ is the action performed by the agent in step t and where $\gamma < 1$ is a discount factor ensuring the sum to be finite. A scheme of reinforcement learning is shown in Figure 1.

Fig. 1: A scheme explaining the concept of the reinforcement learning

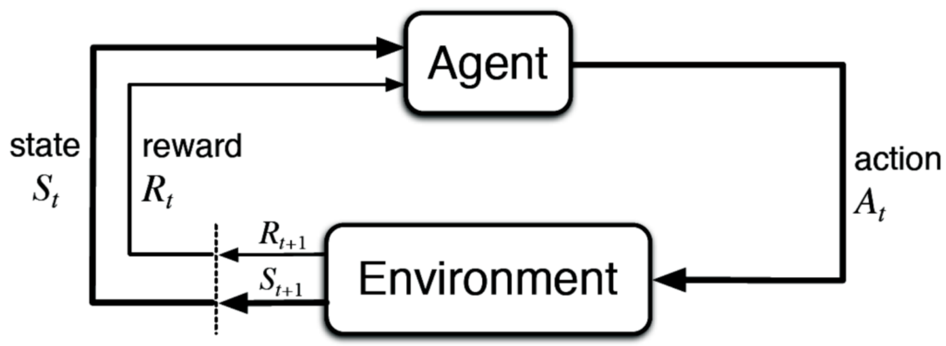

Source: own figure

3 Automated machine learning (AutoML)

Automated machine learning commonly abbreviated as AutoML is motivated by the desire to simplify and automate the process of data analysis (or supervised learning) as much as possible (Kotthoff et al., 2017). This seems especially appealing for the analysis of economic data, although we are not aware of any application of AutoML to financial data. Let us therefore recall an interesting application related to geographical analysis. In the paper (Dakin et al., 2020), the crime rate in a small geographical area in Northern England was investigated. The authors measured the relationship between the number of crime incidents and selected variables. To perform the analysis by means of supervised learning, the authors included an AutoML analysis of individual variables based on more than 60 000 images of the area; this included an automatic identification whether there are street lights or trees on a particular place, ease of access of individual houses, etc. The deployment of an automatic procedure saved very much time with pre-processing the data prior to their analysis by means of regression modeling.

Often, AutoML systems represent black boxes not allowing the user to understand and intepret the results due to their being fully automatic. This saves much time, because it is the data pre-processing (including transformations of data in diverse formats) that is usually the most demanding. These manipulations are arranged to a so-called workflow pipeline, which typically starts with data collection and contains cleaning and pre-processing the data, feature extraction, model selection, optimization of hyperparameters, and finally predictions and model evaluation. Let us now give an overview of popular AutoML systems. While the development of AutoML systems contains many open problems, the developers currently seem to focus on improving the model selection. Our list of systems is focused only on tools for supervised learning.

TPOT (Tree-based Pipeline Optimization Tool). The system TPOT starts with a simple pipeline that is successively modified and extended by means of genetic programming techniques. These include evolution strategies such as mutation, crossover or selection, and help the methods to become more suitable for the given data.

Auto-WEKA. This system was the first Auto-ML system allowing to combine the selection of the best algorithm with searching for the optimal values for hyperparameters (Kotthoff et al., 2017). Such combination is one of possible approaches to the common problem known as the Combined Algorithms Selection and Hyperparameter Optimization (CASH). Particularly, the optimization of hyperparameters exploits a Bayesian optimization approach here, which turns out to be very effective. From the point of view of using Auto-WEKA for data analysis, its available graphical user interface (GUI) is very user-friendly for data analysts without any programming knowledge.

Auto-SKLearn. This is currently the most commonly used AutoML system, although it is useful for model selection only over a limited range of (only standard) machine learning methods. The user may set an upper limit for the time necessary for the computations. In the pipeline of Auto-SKLearn, three basic pillars are used:

- Firstly, automatic selection algorithms find a set of potentially good models;
- Secondly, Bayesian optimization is used to find the optimal values of hyperparameters for given algorithms;

- Finally, a combination of models (ensemble of the models) is performed.

H2O and AutoKeras. While deep learning tools are becoming important for the analysis of financial data, these systems were constructed primarily for an automatic selection of the most suitable neural network for given data. These systems use techniques for an automatic proposal of neural networks and methods for combining neural network models.

Cloud systems. All above-mentioned AutoML systems belong to a category of software locally operated on a computer. Companies like Google or DataRobot started however to offer AutoML systems in their giant computational clouds. Such approaches have received intensive attention, as some IT companies aim at implementing cloud technologies to their systems for software engineers, who need to perform data mining quickly but do not have experience with machine learning.

As methods of AutoML become more available, the authors of this paper still hold the opinion that the abilities of financial mathematicians will not be replaced in the near future by artificial intelligence. The user of AutoML systems is not required to understand the (statistical) background of the methods, but if the user is aware of principles of the classification method used in the system, the interpretation of the results may be much improved. More importantly, there is not a very good experience with feature selection performance of AutoML systems so that it is still advisable to leave this part of the data analysis to a human expert. We can say that AutoML systems in supervised as well as reinforcement learning are still under development. Also from the point of view of research, there is a need for new general methods of dimensionality reduction; particularly, there have not been many robust methods (resistant to the presence of outliers in the data) proposed for the task (Kalina, 2014).

4 Metalearning

Metalearning represents a very broad class of automatic approaches and can be only vaguely characterized as a methodology based on a transition to a higher (in some sense new) level. Such indeterminate definition was presented in (Brazdil and Giraud-Carrier, 2018). The process of learning, i.e. extracting knowledge applicable to the primary level, is then performed on the higher level (or levels). Figures 2 shows a simplistic scheme revealing that metalearning starts with a basic algorithm applied to available datasets, and proceeds with a meta-algorithm (on the higher level) allowing to extract the information from them in a form applicable to a new given dataset. Also available methods of automatic selection of the optimal algorithm or method belong to metalearning as approaches with an increasing popularity (Kalina et al., 2021).

To give a recent example, energy demand was modeled by means of several algorithms of statistics (e.g. ARIMA) or machine learning (e.g. support vector regression) based on a set of predictors in (Arjmand et al., 2020). To improve the performance of the supervised learning approach, the authors used metalearning to exploit exogenous variables, which include meteorological characteristics from different geographical areas; this helps to find the best predictors for the response (e.g. for energy demand).

Fig. 2: Structure of metalearning

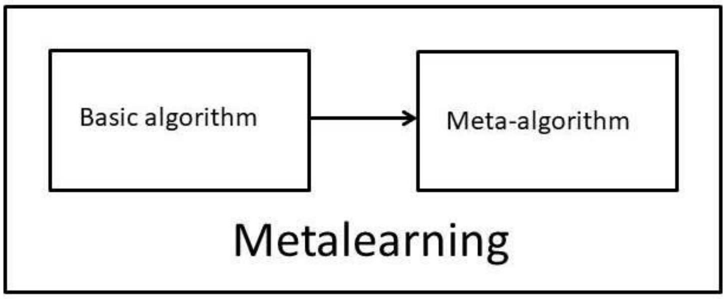

Source: own figure

5 Stock market investment decision are based on information

It is well known that dominant market have influence on less important ones, or that a time series of stock prices may have influence on the time series of cryptocurrency returns. In both situations we speak about information flow, i.e. transfer of information (about the performance of the stock market) with causal effects on financial returns in another market. Information flow has been subject of theoretical studies in the context of information theory and/or cybernetics. Available attempts to evaluate (quantify) the effect of one stock market on another by means of measures of information theory focused on the use of transfer entropy as a measure allowing to evaluate the transfer of information. Transer entropy has a better interpretation compared to correlation coefficients (Dixon et al., 2020), where the latter do not imply causal relationships.

The so-called active investment is an investment strategy characterized by an active search for perspective stocks that are likely to outperform available combinations of stocks (stock market indexes). To overcome the performance of the market is a difficult task and requires to predict future performance under risk (uncertainty) in a better way compared to the majority of other investors. The useful information for active investors may include newly available information about a given company or about the development of the whole economy. Active investors have advantage if able to combine multi-modal information coming from various sources; at the same time, automatic information extraction from free (narrative) texts, which represents a task for natural language processing, is beneficial here. A special case of active investing is the evidence-based investing (factor investing), which requires to take various macroeconomic and microeconomic factors into consideration.

Stock market investment practices summarized e.g. in the paper (Sharma et al., 2018) confirm that stock market investing can be characterized as a competition, where the success is determined in the availability or better quality of information. This practical experience contrasts with theoretical ideas of the economic doctrine, where e.g. Friedrich August von Hayek (1899-1992) perceived the availability of information as an assumption. Hayek took access to information to be granted for the general public and considered information to belong to non-excludable goods.

6 Machine learning tasks related to stock market investing

An important task within portfolio management strategies is to perform dimensionality reduction (variable selection, stock selection) aimed at finding a diverse portfolio of stocks. Statistically speaking, focusing on the most relevant information can decrease the risk (uncertainty). Nevertheless, common stock market indexes contain many stocks and their primary aim is not sparsity (i.e. construction only over a small set of stocks). Instead, the portfolio construction is performed by optimization (rather than by statistical methods) formulated in the framework of the risk-return trade-off. A typical procedure is to maximize the return under the constraint of keeping variability under control, thus also avoiding a selection of strongly correlated stocks.

Other possible applications of machine learning (neural networks) to finance include stock return predictions (possibly using Bayesian methods allowing to incorporate prior expectations) or unsupervised learning (clustering) applied to stocks with the aim to determine existing but unknown homogeneous groups of stocks. Machine learning also allows to automate the processes of decisions related to buying and selling, and thus allows to implement a fully automated algorithmic trading (Jeong and Kim, 2019). Also decision support systems for supporting investment decisions belong to artificial intelligence tools; implementing decision support systems is already relatively straightforward, if the user can rely on realistic predictions of future returns.

Conclusion

This paper recalls and discusses some important directions within the rapidly growing field of automated machine learning. While supervised learning methods have already established their position within analysis of financial data, other approaches presented in this paper seem to have been studied especially in computer science journals, without a more vivid reception in the financial mathematics literature. Still, as the literature research of this paper indicates, the existing financial applications of automated machine learning are very promising. We would like to stress that a statistical interpretation of the results of the automatic machine learning (although difficult) remains crucial for a proper understanding of the result. This is definitely true if financial data are analyzed. Therefore, we would like to focus on investigating robust methods for high-dimensional data in the future, allowing also an automatic detection of outliers in contaminated financial data.

Acknowledgment

The work was supported by the project GA21-19311S ("Information Flow and Equilibrium in Financial Markets") of the Czech Science Foundation. The authors would like to thank Petra Vidnerová for discussion.

References

Arjmand, A., Samizadeh, R., Saryazdi, M.D. (2020). Meta-Learning in Multivariate Load Demand Forecasting with Exogenous Meta-Features. *Energy Efficiency*, 13, pp. 871–887.

Athey, S., Imbens, G.W. (2019). Machine Learning Methods that Economists Should Know About. *Annual Review of Economics*, 11, pp. 685–725.

Brazdil, P., Giraud-Carrier, G. (2018). Metalearning and Algorithm Selection: Progress, State of the Art and Introduction to the 2018 Special Issue. *Machine Learning*, 107, pp. 1–14.

Charpentier, A., Élie, R., Remlinger, C. (2021). Reinforcement Learning in Economics and Finance. *Computational Economics*. Available from: arxiv.org/abs/2003.10014.

Dakin, K., Xie, W., Parkinson, S., Khan, S., Monchuk, L., Pease, K. (2020). Built Environment Attributes and Crime: An Automated Machine Learning Approach. *Crime Science*, 9, Article 12.

Deng, Y., Bao, F., Kong, Y., Ren, Z., Dai, Q. (2017). Deep Direct Reinforcement Learning for Financial Signal Representation and Trading. *IEEE Transactions on Neural Networks and Learning Systems*, 28, pp. 653–664.

Dixon, M.F., Halperin, I., Bilokon, P. (2020). *Machine Learning in Finance: From Theory to Practice*. Cham: Springer.

Jeong, G., Kim, H.Y. (2019). Improving Financial Trading Decisions using Deep Q-Learning: Predicting the Number of Shares, Action Strategies, and Transfer Learning. *Expert Systems with Applications*, 117, 125–138.

Kalina, J. (2013). Highly Robust Methods in Data Mining. *Serbian Journal of Management*, 8, 9–24.

Kalina, J. (2014). On Robust Information Extraction from High-Dimensional Data. *Serbian Journal of Management*, 9, 131–144.

Kalina, J., Neoral, A., Vidnerová, P. (2021). Effective Automatic Method Selection for Nonlinear Regression Modeling. *International Journal of Neural Systems*, 31, Article 2150020.

Kotthoff, L., Thornton, C., Hoos, H.H., Hutter, F., Leyton-Brown, K. (2017). Auto-WEKA 2.0: Automatic Model Selection and Hyperparameter Optimization in WEKA. *Journal of Machine Learning Research*, 18, 826–830.

Sharma, M., Sharma, S., Singh, G. (2018). Performance Analysis of Statistical and Supervised Learning Techniques in Stock Data Mining. *Data*, 3, Article 54.

Yin, P.Y., Chao, C.H. (2018). Automatic Selection of Fittest Energy Demand Predictors based on Cyber Swarm Optimization and Reinforcement Learning. *Applied Soft Computing*, 71, 152–164.

Contact

Jan Kalina
Charles University, Faculty of Mathematics and Physics
Sokolovská 83, 186 75 Prague 8, Czech Republic
& The Czech Academy of Sciences, Institute of Computer Science
Pod Vodárenskou věží 2, 182 07, Praha 8, Czech Republic
kalina@cs.cas.cz

Aleš Neoral
The Czech Academy of Sciences, Institute of Computer Science
Pod Vodárenskou věží 2, 182 07, Praha 8, Czech Republic
neoralales@gmail.com