“Machine learning for robo-advisors: testing for neurons specialization”

AUTHORS

Roman Semko <https://orcid.org/0000-0002-3777-4064>
<http://www.researcherid.com/rid/N-9128-2018>

ARTICLE INFO

Roman Semko (2019). Machine learning for robo-advisors: testing for neurons specialization. *Investment Management and Financial Innovations*, 16(4), 205-214. doi:[10.21511/imfi.16\(4\).2019.18](https://doi.org/10.21511/imfi.16(4).2019.18)

DOI

[http://dx.doi.org/10.21511/imfi.16\(4\).2019.18](http://dx.doi.org/10.21511/imfi.16(4).2019.18)

RELEASED ON

Monday, 09 December 2019

RECEIVED ON

Tuesday, 01 October 2019

ACCEPTED ON

Monday, 18 November 2019

LICENSE

This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

JOURNAL

"Investment Management and Financial Innovations"

ISSN PRINT

1810-4967

ISSN ONLINE

1812-9358

PUBLISHER

LLC “Consulting Publishing Company “Business Perspectives”

FOUNDER

LLC “Consulting Publishing Company “Business Perspectives”

NUMBER OF REFERENCES

22

NUMBER OF FIGURES

3

NUMBER OF TABLES

1

© The author(s) 2023. This publication is an open access article.

BUSINESS PERSPECTIVES

LLC "CPC "Business Perspectives"
Hryhorii Skovoroda lane, 10,
Sumy, 40022, Ukraine

www.businessperspectives.org

Received on: 1st of October, 2019
Accepted on: 18th of November, 2019

© Roman Semko, 2019

Roman Semko, Ph.D., Associate
Professor, National University of
Kyiv-Mohyla Academy, Kyiv, Ukraine.

This is an Open Access article, distributed under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license, which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

Roman Semko (Ukraine)

MACHINE LEARNING FOR ROBO-ADVISORS: TESTING FOR NEURONS SPECIALIZATION

Abstract

The rise of robo-advisor wealth management services, which constitute a key element of fintech revolution, unveils the question whether they can dominate human-based advice, namely how to address the client's behavioral biases in an automated way. One approach to it would be the application of machine learning tools during client profiling. However, trained neural network is often considered as a black box, which may raise concerns from the customers and regulators in terms of model validity, transparency, and related risks. In order to address these issues and shed more light on how neurons work, especially to figure out how they perform computation at intermediate layers, this paper visualizes and estimates the neurons' sensitivity to different input parameters. Before it, the comprehensive review of the most popular optimization algorithms is presented and based on them respective data set is generated to train convolutional neural network. It was found that selected hidden units to some extent are not only specializing in the reaction to such features as, for example, risk, return or risk-aversion level but also they are learning more complex concepts like Sharpe ratio. These findings should help to understand robo-advisor mechanics deeper, which finally will provide more room to improve and significantly innovate the automated wealth management process and make it more transparent.

Keywords

wealth management, robo-advisor, fintech, machine learning, neural network, portfolio optimization

JEL Classification

G11, G23, O33, C45

INTRODUCTION

Great Recession of 2008 and fintech revolution sparked the massive development of robo-advisor (RA) services, which broke into the market by posing several fundamental questions to the wealth management industry and providing disruptive solutions to them. First, one of the key underlying drivers of the crises – lack of transparency – was expected to be addressed by automated advice, which could be better tracked and analyzed than subjective human conversations (Government Accountability Office, 2011). Second, RAs, which rely on passive investment strategies, could outperform the active ones. In particular, according to SPIVA Scorecard (S&P Indices Versus Active Funds), 65% of active US equity funds underperformed S&P Index 1500 in 2008 (Soe, Liu, & Preston, 2019). One of the items significantly contributing to the difference between active and passive investment returns is fees and costs, namely, when human advisors charge 1-2% of assets under management (AuM), RA management fee usually ranges from 0 to 50 basis points. Needless to say, legendary investor Warren Buffett issued a challenge to hedge fund industry that it would not be possible to construct hedge fund portfolio that will outperform S&P 500 Index fund during the next decade (nytimes.com, 2008). According to Morningstar (McDevitt & Schramm, 2019), 2009–2018 was a period of massive flow of money from active to passive funds and, citing SPIVA, US equity funds outperformed S&P 1500 only dur-

ing 2009, 2010, and 2013, while failed to beat the index during the remaining years. Given this performance results, Buffett ultimately won a bet (Eavis & Grocer, 2018).

Total AuM in the RA segment reached almost one trillion US dollars with 46 million users globally served by hundreds of financial companies (statista.com, 2019). RA services are rapidly developing with their strengths and weaknesses summarized by Jung, Glaser, and Köpplin (2019), and Fisch, Laboure, and Turner (2019). One of the key weaknesses of RA is its limitation to assess the client profile (e.g., willingness and ability to take risk), to provide personalized portfolio, which could be also summarized to some extent as inability to address the client's behavioral biases. For example, robot will usually suggest selling inherited income-generating property, even though there could be present emotional relationship to it. In these circumstances, human advisor will reveal potential endowment bias and will try to adapt rather than moderate it by taking into account client's standard of leaving risk and other factors.

All in all, to capture and process such unstructured client's information, RAs need to pay attention to the opportunity of integrating AI and machine learning tools. That could be done via replacing standard quadratic optimization problem with neural network that will consume different types of data. Given that the client profiles and suggested portfolios are hardly available due to the compliance and security reasons, we first construct portfolios of hypothetical clients using Modern Portfolio Theory. Then, having input-output parameters, we will train neural network and evaluate its neurons for their specialization in response to risk, return, or other metrics.

1. LITERATURE REVIEW

There is a vast literature investigating the RAs methodology and the main ideas behind them. In this review, we will briefly analyze what typical methods are now being used for portfolio construction by leading companies and then will switch to portfolio selection using the neural networks.

Beketov, Lehmann, and Wittke (2018) conducted comprehensive analysis of the methods utilized by RAs. In particular, they investigated web pages of 219 existing companies globally to figure out what the most popular techniques are used for core portfolio construction. It is stated that typical workflow of RA consists of five steps: starting from universe specification and client profiling, through portfolio optimization, and ending up with monitoring, rebalancing, and performance reporting. Asset allocation is usually handled via Modern Portfolio Theory (MPT) and, in case of high-scale RAs (by AuM), it is augmented with Black-Litterman approach.

Leading RAs like Wealthfront, Betterment, or Schwab Intelligent Portfolios are providing their clients the possibility to conduct automated tax-loss harvesting by augmenting classical opti-

mization algorithm with additional tax-related problem.

It is clear that 65-year old MPT (Markowitz, 1952, 1959) cannot be a persuasive approach for portfolio construction. While in the short- to medium-term, business will try to modify or update MPT, in the longer-term digitized world, radically new methodology may become dominant. For example, the application of artificial neural networks (ANNs) can potentially address the existing shortcomings of RAs, like ignorance of client's emotional and cognitive biases.

There are numerous AI approaches that are in scope of wealth management advisors, namely genetic algorithms, ANNs, fuzzy logic, and so on (see for more details Gonzalez-Carrasco, Colomo-Palacios, Lopez-Cuadrado, Garcia-Crespo, & Ruiz-Mezcua, 2012). Nevertheless, they are more complementary than competing. For example, while neural net mostly describes mapping function between inputs and outputs, and genetic algorithm is used mainly for optimization, the former method can still exploit the latter to derive its parameters. Similarly, ANN and fuzzy logic can be combined into one pipeline to process the unstructured and ambiguous information. For the sake of simplicity and given the latest tendencies,

we focus on classical ANNs since they provide the user with all facet of techniques to process textual and numerical information in an effective way and are well-suited for the extension with the abovementioned methods.

One of the seminal papers on portfolio selection based on ANNs was published by Fernandez and Gomez in 2007. The authors developed heuristic method based on Hopfield-type ANN to solve portfolio optimization problem with cardinality and boundary constraints, which, in fact, is mixed quadratic and integer programming task and is very computationally intensive. Ko and Lin (2008) introduced multi-layer resource-allocation ANN (MLALNN) to guarantee that portfolio weights summation constrain is hold. To achieve the goal, weights are dynamically adjusted and modified via learning rate, which is, in contrast to conventional ANN, variable rather than fixed so that summation of outputs is always maintained at one. The authors reported that return on investment of the buy-and-hold portfolio constructed based on proposed MLALNN is significantly greater (by almost 14% on average) than return of Taiwan Stock Exchange index in all sliding windows covering 2000–2005 period.

The other big part of the research is focused on the introduction of AI methods to portfolio construction implicitly. For example, by forecasting such input parameters as risk and returns via Long Short-Term Memory (LSTM) ANN combined with Principal Component Analysis (PCA) and then reusing MPT (e.g., Obeidat, Shapiro, Lemay, MacPherson, & Bolic, 2018; Snihovyi, Kobets, & Ivanov, 2019). Although the authors report that such approach can generate superior return for a given level of risk, we believe that this is pure forecasting task rather than portfolio construction: it is still important for wealth management but has limited impact on MPT itself and cannot directly address client's behavioral biases and specific needs.

Even though during the last several years ANNs have demonstrated remarkable results in solving numerous problems, as Yosinski, Clune, Nguyen, Fuchs, and Lipson (2015) state, they still often are treated as black boxes that are hardly interpretable, especially when speaking about hidden

layers. Often, ANN is thought to be a highly distributed system but if the opposite is true, that is neuron specialization is present, then we could expect ANN is functioning like a natural system and in that case it is expected to be more effective (for example, more diverse neurons population can prevent from converging to suboptimal solutions). Knowing that some neuron is focused on financial feature like return allows not only to better understand and improve ANN architecture but also to analyze more complex cases, e.g., when specialization in risk, return, and risk-free rate could mean that activation happens when Sharpe ratio is changing (by analogy to image recognition when specialization in shelf and text detection is a sign of library photo). Finally, visualization and interpretation of hidden layers themselves means that they represent concentrated structured information. In case of image recognition, it could be some synthetic figure recreated from the figures used for ANN training. If ANN was trained on secret data, only having its weights can allow the analyst to reproduce to some extent selected original features, which will result in security breach. Of course, confidentiality and preservation of client information is highly important for wealth management industry.

There are numerous methods and tools (e.g., Yosinski et al., 2015) that can be utilized to visualize and learn ANN. For the sake of simplicity, we will apply straightforward approach that determines what neurons are mostly activated by particular feature(s) during forward propagation.

Unsolved parts of the problem. To the best of our knowledge, there are no studies that try to visualize and interpret hidden layers (neurons) of ANN which is used for portfolio construction.

Research goal and questions. Specifically, we will try to test the hypothesis whether neurons specialization is present, that is whether some ANN elements react to selected feature(s) in a more significant way.

This paper is organized as follows. We begin with state-of-the-art RA methodology review, which will help us to generate data for ANN training. Then, we will define ANN and describe the approach for visualizing neurons specialization. In

section 3, the data used in the experiment will be described. Results discussion and analysis is provided in the next section 4. Finally, we will conclude and define further steps for investigation.

2. METHODOLOGY

Before constructing ANN, detailed description of the methods that are typically used by modern RAs will be provided. Having implemented this methodology, we will generate portfolios for different client profiles. This information will be the output layer of the future ANN and will be used for its training.

As it has been already said, MPT is the key concept for portfolio construction. However, due to the presence of selected drawbacks (resulting portfolios are unintuitive, concentrated, and extremely sensitive to input parameters), it is often augmented with Black-Litterman approach as demonstrated on Figure 1. First, reverse optimization is run to get equilibrium returns and then we mix them with analyst predictions, which in the end are fed into classical MPT optimization problem. In parallel, tax-loss harvesting can also be implemented to boost portfolio performance.

RAs can set up different MPT optimization problems when constructing portfolio for the clients. It

could try to maximize return for a given level of risk or solve its dual problem of minimizing risk for a given level of return. However, more common approach is to maximize risk-adjusted utility function subject to a set of constraints:

$$\max_w w^T r - \frac{RA}{2} w^T \Sigma w, \quad (1)$$

subject to

$$w^T \mathbf{1} = 1 \quad (2)$$

$$lower\ bound \leq w \leq upper\ bound'$$

where w is the $N \cdot 1$ vector of portfolio weights, r is the $N \cdot 1$ vector of expected returns, Σ is the $N \cdot N$ variance-covariance matrix of returns, RA is the risk-aversion level defined on the scale from 0.5 (ready to accept large risks) toward 10 (very risk-averse individual) with 0.5 increment, *lower bound* and *upper bound* are vectors setting lower and upper limits for each security holdings, e.g. usually 0 or 0.05 for lower bound and 0.35 for upper bound.

Due the already mentioned limitations of standard MPT, large RAs usually augment it with Black-Litterman approach (Black & Litterman, 1990; Black & Litterman, 1991). First, reverse optimization problem is solved resulting in the so-called

Figure 1. Optimization algorithm

equilibrium returns, that is returns that should be expected within MPT given current asset market structure:

$$r^{eq} = S^{gp} \Sigma w^{gp} + r^{rf}, \quad (3)$$

where S^{gp} is the Sharpe ratio of global portfolio calculated as weighted average of Sharp ratios of its constituents, w^{gp} is the vector of weights of global portfolio constituents calculated based on market capitalizations, and r^{rf} is the risk-free rate of return.

Second, these equilibrium returns are blended with analyst forecasts (see, for example, Yang, Ye, Wei, & Bao, 2017 for more details) using Bayes theorem:

$$r^{blended} = \left[(\tau \Sigma)^{-1} + P^T \Omega^{-1} P \right]^{-1} \times \left[(\tau \Sigma)^{-1} r^{eq} + P^T \Omega^{-1} Q \right], \quad (4)$$

where Q is the vector representing analyst views on the expected returns, matrix P contains the information about available analyst returns (potentially, it allows to model not only absolute views like return of asset i will be equal to some number but also relative ones like asset i will outperform asset j by some specific percentage), Ω represents the uncertainty in the views and is calculated as $\tau P \Sigma P^T$, τ is the constant that scales up prior analyst forecasts uncertainty.

Leading RAs like Wealthfront or Betterment often allow their clients to execute automated tax-loss harvesting. This is done via augmenting MPT optimization problem with a new one that aims at minimizing tracking error (direct indexing) and maximizing tax alpha in the following way (Malkiel, 2016):

$$\max TaxAlpha - TrakingError^2 \quad (5)$$

subject to the following set of constraints:

1. No short positions: portfolio is usually assumed to be long-only.
2. To achieve plausible level of diversification, portfolio weights must be less than some reasonably large upper bound, e.g., 0.35 or 35%.

3. The difference between portfolio and benchmark weights lies within short interval limits, e.g., $[-0.01; 0.01]$.
4. Wash sales are not allowed, that is investors cannot apply tax deductions when they sell a stock at loss and buy the same or similar security during 30 days window.

The above problem can be further specified in greater detail as follows:

$$\max_w \alpha^T (w - w_0) - \frac{\lambda}{2} (w - b)^T \Sigma (w - b), \quad (6)$$

subject to

$$\begin{aligned} w^T \mathbf{1} &= 1 \\ 0 &\leq w \leq \text{upper bound} \\ \text{lower limit} &\leq w - b \leq \text{upper limit} \\ \text{no wash sales} \end{aligned} \quad (7)$$

where w_0 is the $N \cdot 1$ vector of current portfolio weights, b is the $N \cdot 1$ vector of benchmark portfolio weights, w is the $N \cdot 1$ vector of candidate portfolio weights, λ is the scalar that plays the role of regularization parameter discriminating between tax alpha and tracking error and is set based on the company appetite toward tax optimization and the desire to follow the selected strategy, Σ is the $N \cdot N$ covariance matrix of returns, α is the $N \cdot 1$ vector containing security level tax alphas (one-sided):

$$\alpha = \frac{STCL \cdot STTR + LTCL \cdot LTTR}{Beginning\ value}, \quad (8)$$

where $STCL$ and $LTCL$ are the short-term and long-term capital losses, respectively, multiplied by the short-term and long-term tax rates. Similarly, tax alpha can be defined also as the excess after-tax return minus excess pre-tax return, where excess means additional return over portfolio benchmark.

3. DATA DESCRIPTION

We use similar set of assets in the investment universe as it is done by Wealthfront (research. wealthfront.com, 2019), namely such asset classes

as U.S. stocks, foreign stocks, emerging markets, bonds, and alternative investments are represented by seven ETFs: Vanguard Total Stock Market ETF (VTI), Vanguard FTSE Developed Markets ETF (VEA), Vanguard FTSE Emerging Markets ETF (VWO), Vanguard Dividend Appreciation ETF (VIG), Energy Select Sector SPDR ETF (XLE), iShares National Muni Bond (MUB), and Schwab Barclays Capital U.S. TIPS (SCHP). Daily quotes representing 2010–2017 period are transformed into monthly returns and variance-covariance matrix. Maximum allowed allocation is 35%, while minimum is set to 0% for SCHP and 5% for all other ETFs. We assumed that management fee is 0.15%, expense ratio varies from 0.03 to 0.25%. Capital gain are taxed at 25%, dividends at 25%, and SCHP income at 40%. Capital market assumptions (e.g., correlation matrix) are also similar to Wealthfront (research.wealthfront.com, 2019). Current market capitalizations are snapshot as of end of 2016 and constitute net assets of above ETFs. Black-Litterman scalar τ is set at conventional level of 0.025. Inflation rate is 2% and risk-free rate is approximated by USD LIBOR.

Technically, clients are asked seven multiple-choice questions about their age, family status, income, investment amount, behavior in case of stock market downturn, and primary goal of investment, which are then transformed into risk-aversion level varying from 0.5 to 10 with 0.5 step, that is 20 different investor types are possible. Using five-year time horizon for optimization, we ended up with 396 business days that together with 20 risk-aversion types per each time period amounted to 7,920 observations that will be consumed by the ANN.

4. RESULTS ANALYSIS AND DISCUSSION

140 input features, m , (above described variables like historical and forecasted returns, variance-covariance matrices, fees, risk-aversion level and so on) generate asset distribution for 7 instruments. Due to the low variability in selected features (e.g. management fee is constant), we have reduced data dimensionality by applying principal component analysis using Python

scikit-learn module, namely, the first loading vector (PC direction) is given by:

$$p_{(1)} = \arg \max_{\|p=1\|} \{p^T D^T D p\}, \quad (9)$$

where D is the column-wise demeaned matrix of original data matrix $X_{N \times m}$, and the remaining loading vectors (given previous $k-1$) are defined as:

$$p_{(k)} = \arg \max_{\substack{\|p=1\| \\ p^T p_{(j)}=0, j=1, \dots, k-1}} \{p^T D^T D p\}. \quad (10)$$

ANN architecture is shown on Figure 2. There is one hidden layer that consists of 20 neurons each activated by element-wise sigmoid function:

$$a_i^{(0)} = \frac{1}{1 + e^{-z_i^{(0)}}}, \quad (11)$$

where $z_i^{(0)}$ is the i element of vector $(x_k - M)P\Theta_0^T$ with x_k being the k row (observation) of matrix X , M being the column-wise mean vector of X , P being the loading matrix and Θ_0^T being the transposed coefficients matrix.

Output layer activation function is softmax, namely at each output unit it is:

$$a_i^{(1)} = \frac{e^{z_i^{(1)}}}{\sum_{j=1}^7 e^{z_j^{(1)}}}, \quad (12)$$

where $z_i^{(1)}$ is the i element of $A^{(0)}\Theta_1^T$ with Θ_1^T being the transposed coefficients matrix (both $(x_k - M)P$ and $A^{(0)}$ are augmented with bias units).

We replaced standard sigmoid function with softmax at output layer since it guarantees that sum of all portfolio weights equals to 1.

Cost function, defined as mean squared error, is minimized using gradient descent with back propagation algorithm; learning rate is set at 0.02.

To test for neurons specialization, we can assess their sensitivity to the changes to underlying feature(s) by calculating the elasticities. They will show by how much neuron's activation value

Figure 2. ANN architecture

changes when, for example, risk-aversion level or historical returns change by one percent. Total differential of $a_i^{(0)}$ results in the following:

$$\begin{aligned}
 da_i^{(0)} &= \sum_{j=1}^m \frac{\partial a_i^{(0)}}{\partial x_{kj}} dx_{kj} = \\
 &= \sum_{j=1}^m \frac{\partial a_i^{(0)}}{\partial x_{kj}} \frac{x_{kj}}{a_i^{(0)}} \frac{a_i^{(0)}}{x_{kj}} dx_{kj} = \\
 &= \sum_{j=1}^m \varepsilon_{ij} \frac{dx_{kj}}{x_{kj}} a_i^{(0)} \Rightarrow \frac{da_i^{(0)}}{a_i^{(0)}} = \\
 &= \sum_{j=1}^m \varepsilon_{ij} \frac{dx_{kj}}{x_{kj}}.
 \end{aligned} \quad (13)$$

It means, if we perturb one or several features by 1%, that is $dx_{kj} / x_{kj} = 0.01$, the percentage change of neurons activation value would be equal to the sum of respective elasticities, ε_{ij} , with

$$\varepsilon_{ij} = \frac{e^{(x_k - M)P(\Theta_{0i})^T}}{\left(1 + e^{(x_k - M)P(\Theta_{0i})^T}\right)^2} c_j (\Theta_{0i})^T, \quad (14)$$

where Θ_{0i} is the i row of matrix Θ_0 and c_j is the j row of matrix P .

To formalize this type of sensitivity analysis, we have calculated elasticities for the average observation (means of all features) and provide their absolute values in Table A1 in Appendix. We also highlighted cells that represent 90th percentile in each column. For example, when we change historical re-

turns by 1%, neuron 13 and 20 react the most. Then, if we check each row of the table, we can see the cases when specific neuron plays the most important role. For instance, neuron 1 is the most important for forecasted covariances. Even though its elasticity may be larger in other cases like response to risk-aversion level, its role for the latter is not the key one since neuron 13 and 15 are the most elastic to risk aversion. We can say that neuron 1 (and similarly neuron 11) is specializing in forecasted covariances. Analyzing further, neuron 2 specializes in taxes, neurons 5 and 20 in returns, neurons 8 and 13 in covariances, and neuron 15 in risk-aversion level. There are also more complex responses, when neurons play the most important role for several different features. This could be interpreted as they are capturing more complex patterns. For example, neuron 7 is specializing in historical covariances, risk-free rate, and taxes. In addition, neuron 7 is responding relatively strong to the changes in historical returns. All in all, we may conclude that this neuron is specializing in calculating the complex feature that could be similar to Sharpe ratio.

Finally, we can visualize the effect of specific feature on hidden layer units (neurons) to see how they are reacting to the changing input parameters. On Figure 3, we have shown ANN pipeline.

The first quadrant represents the initial normalized set of features (one feature is one point). If the value of some input parameter changes, then corresponding color will also change (feature numerical value is used to select the color). In this

Figure 3. Testing for neurons specialization: response to risk-aversion level

example, we are stressing risk-aversion level. The next quadrant represents principal components. Then, we have vector of neurons and the darker is their color, the more they have deviated from the baseline scenario. Finally, we have seven out-

put neurons representing portfolio weights. As we can see from example on the figure, neurons 13 and 15 are responding the most to the changes in risk-aversion level (in line with elasticities analysis).

CONCLUSION

Robo-advisor wealth management services have entered the market during the last decade. Possessing a set of advantages over human like low fees, greater transparency and accessibility, automated solutions cannot provide personalized portfolios that will incorporate client-specific circumstances and their behavioral biases. Natural approach to address this drawback is to apply ANN for portfolio optimization. However, the client, advisor, and regulators may want to know what lies behind this ANN black box. One way to analyze it is to test neurons for their specialization when responding to different features. By constructing ANN that should replicate the behavior of classical RA, we have trained it and, then, calculated elasticity of hidden layer neurons to different features. It has been found that selected neurons are more sensitive and playing larger role for specific feature(s), e.g., some neurons may be more focused on risk, return, risk aversion, taxes or other feature. There are also hidden units that demonstrate more complex behavior responding strongly to multiple input parameters. We may conclude that they have learned how to calculate such metrics as, for example, Sharpe ratio, which takes as inputs risk, return, and risk-free rate.

In summary, this analysis could help to understand ANN inner mechanism deeper, add more clarity to the clients and regulators, and suggest financial advisors further steps to improve the model. For the future research, having more textual data that describe client profile and their behavioral biases, recursive ANN can be constructed to generate more tailored financial portfolios. Testing and visualizing such complex models could be a challenging but, at the same time, rewarding task for all stakeholders leading to more disruptive innovations in the fintech industry.

REFERENCES

1. Beketov, M., Lehmann, K., & Wittke, M. (2018). Robo Advisors: Quantitative Methods Inside the Robots. *Journal of Asset Management*, 19(6), 363-370. <https://doi.org/10.1057/s41260-018-0092-9>
2. Black, F., & Litterman, R. (1990). Asset Allocation: Combining Investor Views with Market Equilibrium. *The Journal of Fixed*

- Income*, 1(2), 7-18. <https://doi.org/10.3905/jfi.1991.408013>
3. Black, F., & Litterman, R. (1991, October). Global Asset Allocation with Equities, Bonds, and Currencies. *Goldman, Sachs & Company Fixed Income Research*. Retrieved from https://faculty.fuqua.duke.edu/~charvey/Teaching/BA453_2006/Black_Litterman_GAA_1991.pdf
 4. Eavis, P., & Grocer, S. (2018, February, 24). Buffett's Annual Letter: Berkshire Records \$29 Billion Gain From Tax Law. *The New York Times*. Retrieved from <https://www.nytimes.com>
 5. Fernandez, A., & Gomez, S. (2007). Portfolio Selection Using Neural Networks. *Computers & Operations Research*, 34(4), 1177-1191. <https://doi.org/10.1016/j.cor.2005.06.017>
 6. Fisch, J. E., Laboure, M., & Turner, J. (2019). The Emergence of the Robo-Advisor. In J. Agnew & O. S. Mitchell (Eds.), *The Disruptive Impact of FinTech on Retirement Systems* (pp. 13-37). Oxford, UK: Oxford University Press. Retrieved from <https://books.google.com.ua/books?id=xJCqDwAAQBAJ&lpg=PA35&ots=VIQhPjcQ0H&dq=%22The%20Emergence%20of%20the%20Robo-Advisor%22%20agnew&pg=PR8#v=onepage&q=%22The%20Emergence%20of%20the%20Robo-Advisor%22%20agnew&f=false>
 7. Gonzalez-Carrasco, I., Colomo-Palacios, R., Lopez-Cuadrado, J. L., Garcia-Crespo, A., & Ruiz-Mezcua, B. (2012). PB-ADVISOR: A Private Banking Multi-Investment Portfolio Advisor. *Information Sciences: An International Journal*, 206, 63-82. <https://doi.org/10.1016/j.ins.2012.04.008>
 8. Government Accountability Office (2011). *401(K) Plans: Improved Regulation Could Better Protect Participants from Conflicts of Interest*. GAO-11-119. Retrieved from <http://www.gao.gov/products/GAO-11-119>
 9. Jung, D., Glaser, F., & Köpplin, W. (2019). Robo-Advisory: Opportunities and Risks for the Future of Financial Advisory. In V. Nissen (Ed.), *Advances in Consulting Research: Recent Findings and Practical Cases* (pp. 405-427). Cham: Springer. Retrieved from https://books.google.com.ua/books?id=VSFzDwAAQBAJ&dq=nissen+Robo-Advisory:+Opportunities+and+Risks+for+the+Future+of+Financial+Advisory&source=gbs_navlinks_s
 10. Ko, P.-C., & Lin, P.-C. (2008). Resource Allocation Neural Network in Portfolio Selection. *Expert Systems with Applications*, 35(1-2), 330-337. <https://doi.org/10.1016/j.eswa.2007.07.031>
 11. Malkiel, B. G. (2016). Automated Investment Services. *Journal of Investment Management*, 14(4), 1-13. Retrieved from <https://www.princeton.edu/~bmalkiel/Automated%20Investment%20Services.pdf>
 12. Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, 7(1), 77-91. <https://doi.org/10.1111/j.1540-6261.1952.tb01525.x>
 13. Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investment*. New York: John Wiley & Sons. Retrieved from <https://www.jstor.org/stable/j.ctt1bh4c8h>
 14. McDevitt, K., & Schramm, M. (2019). 2018 U.S. Fund Flows Trends in 5 Charts. *Morningstar Research Services LLC*. Retrieved from <https://www.morningstar.com/blog/2019/01/28/us-fund-flows-trends.html>
 15. Nytimes.com (2008). Buffett lays down bet against hedge funds. *The New York Times*. Retrieved from <https://www.nytimes.com/2008/06/10/business/worldbusiness/10iht-invest11.1.13598388.html>
 16. Obeidat, S., Shapiro, D., Lemay, M., MacPherson, M. K., & Bolic, M. (2018). Adaptive Portfolio Asset Allocation Optimization with Deep Learning. *International Journal on Advances in Intelligent Systems*, 11(1 & 2), 25-34. Retrieved from http://www.iari-journals.org/intelligent_systems/intsys_v11_n12_2018_paged.pdf
 17. Research.wealthfront.com (2019). *Wealthfront Investment Methodology White Paper*. Retrieved September 23, 2019, from <https://research.wealthfront.com/whitepapers/investment-methodology/>
 18. Snihovyi, O., Kobets, V., & Ivanov, O. (2019). Implementation of Robo-Advisor Services for Different Risk Attitude Investment Decisions Using Machine Learning Techniques. In V. Ermolayev, M. Suarez-Figueroa, V. Yakovyna, H. Mayr, M. Nikitchenko, & A. Spivakovsky (Eds.), *Information and Communication Technologies in Education, Research, and Industrial Applications, ICTERI 2018. Communications in Computer and Information Science* (pp. 298-321). Cham: Springer. http://dx.doi.org/10.1007/978-3-030-13929-2_15
 19. Soe, A. M., Liu, B., & Preston, H. (2019). SPIVA® U.S. Scorecard. *S&P Dow Jones Indices, Year-End (2018)*. Retrieved from https://www.spglobal.com/_assets/documents/corporate/us-spiva-report-11-march-2019.pdf
 20. Statista.com (2019). *Robo-Advisors: Worldwide*. Retrieved September, 23, 2019, from <https://www.statista.com/outlook/337/100/robo-advisors/worldwide>
 21. Yang, J., Ye, H., Wei, Y., & Bao, L. (2017). How Robo-Advisors Manage Investment Portfolios. *Cutter Business Technology Journal*, 30(4), 14-26. Retrieved from <https://www.cutter.com/article/how-robo-advisors-manage-investment-portfolios-495656>
 22. Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., & Lipson, H. (2015). Understanding Neural Networks Through Deep Visualization. *Deep Learning Workshop, 31st International Conference on Machine Learning*, Lille, France. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.713.1915&rep=rep1&type=pdf>

APPENDIX A

Table A1. Neurons elasticities (in percentages) to 1% change in the feature

Neurons	Historical Returns	Forecasted returns	Historical Covariances	Forecasted Covariances	Risk-aversion Level	All Returns	All Covariances	Risk-free Rate	Taxes	Fees and Expenses	Specialization
N1	0.0132435	0.0033566	0.0051144	0.0050126	0.1241446	0.0098869	0.0101270	0.0002527	0.0515649	0.0000707	Forecasted covariances
N2	0.0018526	0.0117196	0.0052579	0.0021841	0.5263215	0.0135721	0.0074420	0.0001412	0.1105370	0.0003043	Taxes
N3	0.0073050	0.0095397	0.0079786	0.0016616	0.0787880	0.0168447	0.0063171	0.0008105	0.0284286	0.0000588	–
N4	0.0037037	0.0107450	0.0019668	0.0011974	0.6239927	0.0144487	0.0031641	0.0001301	0.0080038	0.0001115	–
N5	0.0136474	0.0190280	0.0019162	0.0042242	0.9062565	0.0326754	0.0061404	0.0001631	0.0683331	0.0000309	Returns
N6	0.0125937	0.0018789	0.0038075	0.0009899	0.8483424	0.0107148	0.0028176	0.0001306	0.0416002	0.0003300	–
N7	0.0120643	0.0139654	0.0097170	0.0012209	1.3726878	0.0019016	0.0109379	0.0013460	0.1335377	0.0002741	Sharpe ratio
N8	0.0039322	0.0050702	0.0097110	0.0046176	0.8159057	0.0011380	0.0143286	0.0001236	0.0435010	0.0001368	Covariances
N9	0.0040413	0.0068318	0.0049388	0.0014391	1.0126733	0.0108730	0.0063780	0.0004838	0.0283192	0.0001841	–
N10	0.0001743	0.0056508	0.0091159	0.0020772	0.9494889	0.0058250	0.0111932	0.0006506	0.0437663	0.0006004	Fees and expenses
N11	0.0027085	0.0038217	0.0049591	0.0046404	0.1189715	0.0011132	0.0095995	0.0000515	0.0236769	0.0002512	Forecasted covariances
N12	0.0059555	0.0004079	0.0006233	0.0002303	1.2801005	0.0063634	0.0008536	0.0001865	0.0695873	0.0000791	–
N13	0.0140876	0.0027842	0.0045284	0.0000219	1.9510858	0.0113034	0.0045065	0.0002390	0.0804461	0.0000764	Complex response
N14	0.0079171	0.0006152	0.0065650	0.0028565	0.7110904	0.0085324	0.0094215	0.0001224	0.0254911	0.0000647	–
N15	0.0028952	0.0085530	0.0026335	0.0017717	2.3410687	0.0114481	0.0008618	0.0006295	0.0400031	0.0001283	Risk-aversion level
N16	0.0020534	0.0156718	0.0003146	0.0027643	0.6898017	0.0177252	0.0024497	0.0009796	0.0559447	0.0004090	Complex response
N17	0.0085993	0.0055258	0.0101336	0.0012501	0.2071328	0.0141251	0.0113837	0.0004651	0.0743825	0.0002037	Covariances
N18	0.0035603	0.0025285	0.0014566	0.0037463	1.5064833	0.0010318	0.0052029	0.0002850	0.0232804	0.0003279	–
N19	0.0055317	0.0036139	0.0075069	0.0020168	0.1646872	0.0091456	0.0095237	0.0006817	0.0751061	0.0000413	–
N20	0.0165030	0.0025377	0.0026074	0.0000762	0.1679000	0.0190406	0.0025312	0.0006319	0.0025548	0.0000226	Returns