



SALES TIME SERIES ANALYTICS USING DEEP Q-LEARNING

Bohdan M. Pavlyshenko

Ivan Franko National University of Lviv, Ukraine
b.pavlyshenko@gmail.com

Paper history:

Received 10 January 2020
Received in revised form 02 June 2020
Accepted 08 July 2020
Available online 27 September 2020

Keywords:

Sales;
Time Series;
deep Q-learning;
Reinforcement Learning;
Machine Learning.

Abstract: The article describes the use of deep Q-learning models in the problems of sales time series analytics. In contrast to supervised machine learning, which is a kind of passive learning, where historical data are used, Q-learning is a kind of active learning aimed at maximizing a reward by optimal sequence of actions. Model free Q-learning approach to optimal pricing strategies and supply-demand problems is considered in the work. The main idea of the study is to show that using deep Q-learning approach in time series analytics causes the sequence of actions to be optimized by maximizing the reward function when the environment for learning agent interaction can be modeled using the parametric model and in the case of using the model which is based on the historical data. In the pricing optimizing case study environment was modeled using sales dependence on extra price and randomly simulated demand. In the pricing optimizing case study, the environment was modeled using sales dependence on extra price and randomly simulated demand. In the supply-demand case study, it was proposed to use historical demand time series for environment modeling, agent states were represented by promo actions, previous demand values and weekly seasonality features. Obtained results show that using deep Q-learning, we can optimize the decision making process for price optimization and supply-demand problems. Environment modeling using parametric models and historical data can be used for the cold start of learning agent. On the next steps, after the cold start, the trained agent can be used in real business environment.

*Copyright © Research Institute for Intelligent Computer Systems, 2020.
All rights reserved.*

1. INTRODUCTION

Sales time series analytics is an important part of modern business intelligence. We can mention classical and popular time series models – Holt-Winters, ARIMA, SARIMA, SARIMAX, GARCH, etc. Different time series models and approaches can be found in [1-5]. In [6] we studied the use of machine-learning models for sales predictive analytics. We considered the main approaches and case studies of using machine learning for sales forecasting and the effect of machine-learning generalization. In this paper, we also considered a stacking approach to building regression ensemble of single models. In [7], we studied the logistic regression in the problem of detecting manufacturing failures. For the logistic regression, we considered a generalized linear model, machine learning and Bayesian models. In [8], we studied

stacking approaches to time series forecasting and logistic regression with highly imbalanced data. The use of regression approaches to sales forecasting can often give us better results compared to time series methods. One of the main assumptions of regression methods is that the patterns in the historical data will be repeated in future. Supervised machine learning can be considered as a kind of passive learning using historical observations. Time series forecasting using Machine Learning gives us insights and allows us to make the right business decisions. Reinforcement learning allows us to find sequences of optimized actions directly without historical data. In this approach, we have an environment and a learning agent which interacts with environment. As a result of each interaction, learning agent receives reward. The goal of Reinforcement Learning is to find such sequence of actions which will maximize

an average cumulative reward on the episodes of agent-environment interactions. There are policy based and policy free approaches. Policy can be described by parameterized distribution function for states and actions. The parameters of these distributions can be found using policy gradient approach where on each of the iterations we calculate the gradient of objective function. Policy free approach can be Q-learning which is based on Bellman equation [9, 10, 11]. On each of the iterations, we upgrade the Q-table where rows represent states and columns represent actions. In the case of continuous action, space Q-table can be approximated by Neural Network using DQN approach [10, 11].

2. RELATED WORK

The main principles of reinforcement learning can be found in [9]. In [10, 11] deep Q-learning approaches were studied. In [12], real-time dynamic pricing in a non stationary environment has been considered. In the article, the problem of establishing a pricing policy that maximizes the revenue for selling a given inventory by a fixed deadline has been highlighted. In [13] reinforcement learning for fair dynamic pricing has been considered. In [14], Reinforcement Learning algorithm that can tune parameters of a seller's dynamic pricing policy in a gradient direction even when the seller's environment is not fully observable was proposed. Different approaches to dynamic pricing are considered in [15]. In the paper [16] adaptive inventory-control models for a supply chain consisting of one supplier and multiple retailers were proposed. The paper [17] describes the use of reinforcement learning techniques in the problem of determining dynamic prices in an electronic retail market. The papers [18, 19] consider the use of reinforcement learning for financial analytics. The paper [20] formulates an autonomous data-driven approach to identify a parsimonious structure for the NN so as to reduce the prediction error and enhance the modeling accuracy. The Reinforcement Learning based Dimension and Delay Estimator (RLDDE) was proposed. The paper [19] presents a model-free Reinforcement Learning framework to provide a deep machine learning solution to the portfolio management problem. In [21], Deep Reinforcement Learning in Large Discrete Action Spaces was considered.

3. DEEP Q-LEARNING APPROACHES

The goal of Q-learning is to maximize cumulative future reward [9, 10, 11]. To train the Q-learning network, the gradient descent algorithm is

often used. To eliminate influence between sequence data and non-stationary distribution, the replay mechanism [22] can be used. This approach consists in random sampling of previous data which represent states and actions. It makes it possible to average distributions of data which describe previous agent's behaviour. The goal of the agent lies in choosing a sequence action strategy which maximizes future rewards [11]. An optimal action-value function can be considered as:

$$Q^*(s, a) = \mathbb{E}_{s' \sim \varepsilon} \left[r + \gamma \max_{a'} Q(s', a') \mid s, a \right] \quad (1)$$

where r is a reward, s is a state, a is an action, s', a' are possible states and actions on the next time step. To estimate the function $Q^*(s, a)$ approximately, Bellman equation can be used in the iteration process:

$$Q_{i+1}(s, a) = \mathbb{E} \left[r + \gamma \max_{a'} Q(s', a') \mid s, a \right] \quad (2)$$

The main problem of such an approach is that there is no generalization of revealed patterns in the agent-environment interaction due to the fact that $Q(s, a)$ is estimated on each separate step. To improve generalization, one can use an approximation function for $Q^*(s, a)$. Neural network can be used for this purpose. Parameters of such deep Q-network can be found using the gradient methods minimizing the loss function

$$L_i(\theta_i) = \mathbb{E}_{s, a \sim \rho} \left[(y_i - Q(s, a; \theta_i))^2 \right] \quad (3)$$

$$y_i = \mathbb{E}_{s' \sim \varepsilon} \left[r + \gamma \max_{a'} Q(s', a', \theta_{i-1}) \mid s, a \right] \quad (4)$$

where ρ – is the behaviour distribution, θ is the weights of Q-network.

Experience replay approach [22] is effectively used in deep Q-network [11, 10]. In this approach, agent actions and states are stored in the replay memory at each time step as tuples $e_t = (s_t, a_t, r_t, s_{t+1})$. Tuples e_t are stored in the data set $D_t = \{e_1, \dots, e_t\}$. On each step for Q-learning updates, the algorithm gets samples e_t from the replay memory by uniform random sampling $e_t \sim U(D)$ [11]. On the next experience replay step, Q-learning updates the weights θ . Then, the agent selects an optimal action, using ε -greedy policy

[11]. Data mini batches are formed on each of the iterations for updating the weights for Q-network. Mini batches are chosen randomly. Such an approach provides generalization of approximation of given data for agent-environment interaction. Due to the experience replay approach, behavioural distribution is averaged on many previous states and actions of learning agent that provides the convergence of iteration process. One of widely used approaches to Q-network consists in treating agent states as input parameters when outputs are Q -values for each separate agent action [11]. In such an approach, Q -value of each action is calculated during one Q-network step forward.

In this work, we consider two cases of using Q-learning in sales time series analytics. One case is an optimal pricing strategy and the second case is supply and demand problems which appear in retail domain areas. In our numerical experiments, we used the algorithms which were based on model architecture described in [10, 11] and the approaches to Q-learning agent implementations from [23, 24, 25]. For environment modeling, we used the parametric models in the case of price optimization and historical time series of demand in the case of supply-demand problem. For the analysis of supply-demand problem, we used the store sales historical data from 'Rossmann Store Sales' Kaggle competition [26]. These data describe sales in Rossmann stores. The calculations were conducted in the Python environment using the main packages *pandas*, *sklearn*, *numpy*, *keras*, *matplotlib*, *seaborn*. To conduct the analysis, *Jupyter Notebook* was used.

4. OPTIMAL PRICING STRATEGY USING Q-LEARNING

Let us consider a simple case for the pricing strategy. The question is what strategy can be applied to maximize profit for some time period. Generally, the strategy can involve many factors and means. In the simplest model, it consists of discrete extra price values which are the fraction of cost price. Dependence between sales and extra price can be considered as

$$F_{Sales} = \frac{a}{(1 + b \cdot \exp(c \cdot (Price_m \cdot (1 + Price_e) - d)))}, \quad (5)$$

where $Price_m$ is an marginal price, $Price_e$ is an extra price, a, b, c, d are the parameters for function F_{Sales} . The function F_{Sales} (5) describes relative decreasing of sales with an increasing extra price. Extra price is considered in relative parts of

marginal price. The reward function for Q-learning can be considered as

$$Reward = Demand \cdot F_{Sales} \cdot Price_e \quad (6)$$

We created a simple model with normalized variables. The Figure 1 shows sales vs extra prices ($Price_m=1$, $a=1$, $b=1$, $c=15$, $d=1.5$). We can observe that for high extra price we get small sales. The real parameters for the logistic curve can be found by the gradient method using historical data. The Figure 2 shows profit versus extra prices. One can see that there is no profit with low and high extra price.




Figure 1 – Sales vs extra prices




Figure 2 – Profit versus extra prices

The challenge is to find the optimal extra price. We can tabulate the objective function using simulated demand time series and find the optimal value for extra price. But in real cases, profit versus extra prices can have a complicated functional dependence, including the dependence on many qualitative factors which are included into complex multilevel pricing strategy. To find which pricing strategy is optimal, we can use a Q-learning approach. This simple model can be used for the

cold start of learning agent before the interaction with real business environment.

Let us consider the parameters for numerical modeling. For the Q-values approximation we used 2-layers feed forward neural network with 32 neurons in each layer. Output neural network dimension is equal to the number of possible actions, which was chosen equal to 8. For the actions, we took the list of the following values for normalized extra price [0, 0.15, 0.25, 0.5, 0.75, 0.85, 1, 1.5]. The extra price is considered in relative parts of the marginal price. As time steps we consider days. The number of time steps in the each episode was equal to 7, the batch size was 32, the number of learning iterations was 50, the optimizer is Adam, the learning rate for neural network was 0.001, the epsilon decay was 0.97. For approximation of the function (5), we used the following parameters: $Price_m = 1, a=1, b=1, c=7, d=1.7$.

In this case study, we used DQN approach with epsilongreedy exploration-exploitation trade off. Epsilon describes the probability of random action. On each iteration, epsilon decreases. The Figure 4 shows epsilon vs time dependence. For the numerical experiment, we simulated demand with random uniform distribution of relative units. As a result of numerical experiment we received trained DQN model which can output optimal actions that maximize future cumulative reward. Figure 3 shows simulated demand time series. Figure 5 shows mean reward over the episodes. One can observe that the reward goes up with iterations. It means that the learning agent improves the way it interacts with the environment. Figure 6 shows the frequencies of actions and we can see one dominated action which corresponds to the optimal pricing strategy for this simple model. Figure 7 shows the actions with time and we can see a big dispersion of actions at the beginning of interaction process due to the domination of exploration type of interactions. We can observe that on the next steps after exploration time period, one action dominates which can be considered as a found optimal pricing strategy.




Figure 3 – Simulated demand time series




Figure 4 – Epsilon vs time dependence




Figure 5 – Mean reward over episodes




Figure 6 – Frequencies of actions




Figure 7 – Actions vs time

5. SUPPLY AND DEMAND CASE STUDY USING QLEARNING

Let us consider a case study of using Q-learning for supply and demand problems. We can also use historical data for the initial start of Q-learning algorithm along with parametric modeling of the environment. Such an approach makes it possible to conduct the cold start for Q-learning agents. In this case study, we took normalized demand time series with seasonality and promo action factor. The challenge is to find optimal discrete actions in the supply-demand problem. Products can be supplied by batches with discrete amount. In the model, we took into account the expenses for product processing which are related to logistic, storing and other expenses. The reward on each step can be considered as

$$\text{Reward} = \text{SalesProfit} - \text{ProcessCost} \quad (7)$$

In this case study, we have more features for state representation. They are the product demand the day before, expected promo action, day of week. Let us

consider the parameters for modeling supply-demand problem. For modeling environment, we used the historical time series for demand. For each episode, the demand for 150 days was taken. This time series was simulated given the sales time series from ‘Rossmann Store Sales’ Kaggle competition [26]. To eliminate overfitting, random lag for the first point of demand time series was applied. The lag was calculated by uniform random distribution for integer values ranging from 0 to 25. For each episode, we used the time period of 150 days, the number of actions was 7. We used the feed forward neural network with 2 layers with 64 neurons in each layer. The batch size was 32, learning rate was 0.001, epsilon decay was 0.995, gamma coefficient for Bellman equation was 0.3. To calculate the reward, we used the following parameters: the value of price profit was 1, pack unit was 0.05, price support was 0.5. For state features we used a promo binary feature, previous day sales, weekly seasonality features which were 7 binary features for each weekday which were obtained by one hot encoding of the weekday categorical feature. We also set up generating stop episode flag when the reward became lower than specified reward. Applying this rule accelerates the process of Q-learning, since in the cases of low current reward values learning process was stopped and new learning iteration was started. These actions were considered as discrete values for supply. For numerical modeling, 7 actions were chosen with the following values [0, 2, 4, 6, 8, 10, 12], which are a discrete number of packs. The amount of simulated product in the pack was 0.025 of relative units. Figure 8 shows simulated demand time series for arbitrary chosen episode. For numerical experiment, demand is simulated in some relative units. Figure 9 shows calculated mean reward over episodes.




Figure 8 – Simulated demand time series




Figure 9 – Mean reward over episodes

We can observe that the learning agent optimizes action sequences over the episodes. In Figure 10, we can see that several actions dominate in comparison with the previous case study for price strategy optimization where only one action dominated. Here we have more features for state representation, so the learning agent chose different optimal actions for different states. Figure 11 shows the heatmap for action frequencies against weekday. We can see that for different days we have different dominated actions which also depend on promo action which can take place in different days. Figure 12 shows the time series of demand, supply, stock and shortage. Dynamics of supply, shortage and stock depend on reward function which can be formed by profit, expenses on logistic, product processing and lost profit.

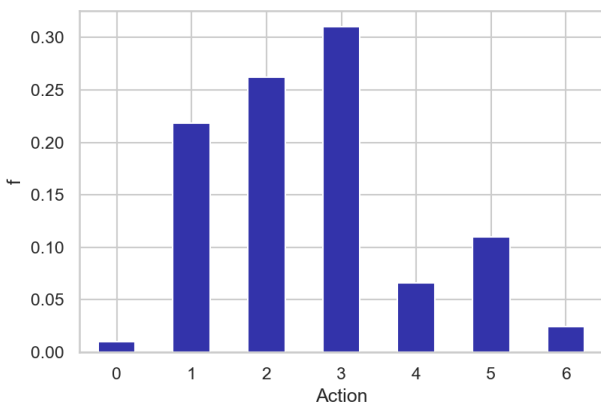


Figure 10 – Frequencies of actions for supply-demand problem




Figure 11 – Heatmap for action frequencies vs weekday

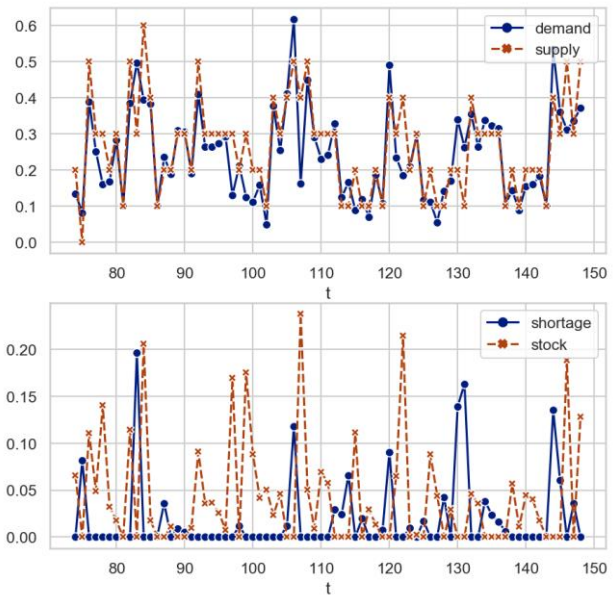


Figure 12 – Time series of demand, supply, stock and shortage

6. CONCLUSIONS

The article describes the use of deep learning models for Q-learning in the problems of sales time series analytics. In contrast to supervised machine learning, which is a kind of passive learning using historical data, Q-learning is a kind of active learning aimed at maximizing the reward by optimal sequence of actions. Model free Q-learning approach for optimal pricing strategies and supply-demand problems was considered in the work. It was shown that using deep Q-learning approach, the sequence of actions can be optimized by maximizing the reward function. In the pricing optimizing case study, the environment was modeled using sales dependence on extra price and randomly simulated demand. In the supply-demand case study, it was proposed to use historical demand time series for

environment modeling, agent states were represented by promo actions, previous demand values and weekly seasonality features. Obtained results show that using deep Q-learning, we can optimize the decision making process for price optimization and supply-demand problems. Environment modeling with parametric models and historical data can be used for the cold start of learning agent. On the next steps, after the cold start, the trained agent can be used in real business environment. So, using Q-learning we can build a decision-making algorithm, which proposes qualitative decisions. This algorithm can start with the data simulated by parametric expert model or with the model based on historical data and then it can work with real business environment and do adaptation to the changes in business processes. In more complicated cases, we can take into account the price of spoiled products in case we deal with perishable products and other expenses.

7. REFERENCES

- [1] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, John Wiley & Sons, 2015.
- [2] P. Doganis, A. Alexandridis, P. Patrinos, and H. Sarimveis, "Time series sales forecasting for short shelf-life food products based on artificial neural networks and evolutionary computing," *Journal of Food Engineering*, vol. 75, no. 2, pp. 196-204, 2006.
- [3] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, OTexts, 2018.
- [4] R. S. Tsay, *Analysis of Financial Time Series*, vol. 543. John Wiley & Sons, 2005.
- [5] W.W. Wei, *Time Series Analysis, The Oxford Handbook of Quantitative Methods in Psychology*, Volume 2; Oxford University Press: Oxford, UK, 2006.
- [6] B. M. Pavlyshenko, "Machine-learning models for sales time series forecasting," *Data*, vol. 4, no. 1, paper 15, pp. 1-11, 2019.
- [7] B. Pavlyshenko, "Machine learning, linear and Bayesian models for logistic regression in failure detection problems," *Proceedings of the IEEE International Conference on Big Data (Big Data)*, 2016, pp. 2046-2050.
- [8] B. Pavlyshenko, "Using stacking approaches for machine learning models," *dings of the 2018 IEEE Second International Conference on Data Stream Mining & Processing (DSMP)*, 2018, pp. 255-258.
- [9] R. S. Sutton, A. G. Barto, et al., *Introduction to Reinforcement Learning*, vol. 2. MIT press Cambridge, 1998.
- [10] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [11] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [12] R. Rana and F. S. Oliveira, "Real-time dynamic pricing in a non-stationary environment using model-free reinforcement learning," *Omega*, vol. 47, pp. 116-126, 2014.
- [13] R. Maestre, J. Duque, A. Rubio, and J. Arévalo, "Reinforcement learning for fair dynamic pricing," *Proceedings of the SAI Intelligent Systems Conference*, Springer, 2018, pp. 120-135.
- [14] D. Vengerov, "A gradient-based reinforcement learning approach to dynamic pricing in partially observable environments," *Technical Report*, Sun Microsystems, Inc., 2007.
- [15] A. V. den Boer, "Dynamic pricing and learning: historical origins, current research, and new directions," *Surveys in operations research and management science*, vol. 20, no. 1, pp. 1-18, 2015.
- [16] C. O. Kim, J. Jun, J. Baek, R. Smith, and Y.-D. Kim, "Adaptive inventory control models for supply chain management," *The International Journal of Advanced Manufacturing Technology*, vol. 26, no. 9-10, pp. 1184-1192, 2005.
- [17] C. Raju, Y. Narahari, and K. Ravikumar, "Reinforcement learning applications in dynamic pricing of retail markets," *Proceedings of the IEEE International Conference on E-Commerce, CEC 2003*, 2003, pp. 339-346.
- [18] C. Y. Huang, "Financial trading as a game: A deep reinforcement learning approach," *arXiv preprint arXiv:1807.02787*, 2018.
- [19] Z. Jiang, D. Xu, and J. Liang, "A deep reinforcement learning framework for the financial portfolio management problem," *arXiv preprint arXiv:1706.10059*, 2017.
- [20] F. Liu, C. Quek, and G. S. Ng, "Neural network model for time series prediction by reinforcement learning," *Proceedings of the 2005 IEEE International Joint Conference on Neural Networks*, 2005, vol. 2, pp. 809-814.
- [21] G. Dulac-Arnold, R. Evans, H. van Hasselt, P. Sunehag, T. Lillicrap, J. Hunt, T. Mann, T. Weber, T. Degris, and B. Coppin, "Deep reinforcement learning in large discrete action

- spaces,” arXiv preprint arXiv:1512.07679, 2015.
- [22] L.-J. Lin, “Reinforcement learning for robots using neural networks,” tech. rep., Carnegie-Mellon Univ Pittsburgh PA School of Computer Science, 1993.
- [23] “Github Repository,” [Online]. Available at: <https://github.com/dennybritz/reinforcement-learning>. resource, accessed 5 December 2019.
- [24] “Github Repository,” [Online]. Available at: <https://github.com/keon/deep-q-learning>, accessed 5 December 2019.
- [25] “Github Repository,” [Online]. Available at: <https://github.com/rlcode/reinforcement-learning>, accessed 5 December 2019.
- [26] “Rossmann Store Sales’, Kaggle.Com,” [Online]. Available at: <http://www.kaggle.com/c/rossmann-store-sales>, accessed 5 December 2019.



Bohdan Pavlyshenko, PhD (physics), a postdoctoral researcher at Electronics and Computer Technologies Faculty at Ivan Franko National University of Lviv. His current scientific interest lies in the area of quantitative linguistics, machine learning, reinforcement learning, predictive analytics, computer vision, social network mining, time series analytics, numeric modeling, risk assessment, reliability theory, financial modeling. He has practical experience in retail and supply chain analytics, customer's behavior analytics, fraud detection.