

Feature selection method using multi-agent reinforcement learning based on guide agents

Minwoo Kim

Gachon University

Jinhee Bae

Gachon University

Bohyun Wang

Gachon University

Joon S. Lim (✉ jslim@gachon.ac.kr)

Gachon University

Article

Keywords: Feature Selection, Guide Agents, Main Agents, Multi-Agent, Reinforcement Learning (RL), Rewards

Posted Date: June 23rd, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1732607/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

In this study, we propose a method to automatically find features from a dataset that are effective for classification or prediction, using a new method called multi-agent reinforcement learning and a guide agent. Each feature of the dataset has one each of main and guide agents, and these agents decide whether to select a feature. Main agents select the optimal features, and guide agents present the criteria for judging the main agents' actions. After obtaining the main and guide rewards for the features selected by the agents, the main agent that behaves differently from the guide agent, by calculating the learning reward delivered to the main agents, updates their Q-values. The behavior comparison between them helps the main agent decide whether its own behavior is correct, without using other algorithms. After performing this process for each episode, the features are finally selected. The feature selection method proposed in this study uses multiple agents, which reduces the number of actions that each agent can perform and finds optimal features effectively and quickly. Finally, comparative experimental results on multiple datasets show that the proposed method can select effective features for classification and increase the classification accuracy.

Introduction

Numerous methods and technologies have emerged, and as the cost of collecting data through these methods and technologies decreases, the number of variables that can be used increases exponentially. To successfully train a machine learning model, it is important to select features that do not adversely affect the prediction of the model and minimize the variation in judgment [1]. In addition, these optimized features positively affect the interpretation of the results and achieve high accuracy [2][3].

Feature selection in the field of machine learning has been studied for an exceptionally long time [4]. Feature selection methods include minimum redundancy maximum relevance (mRMR) [5], Relief [6], and a method using a heuristic algorithm, such as the genetic algorithm [7]. Although feature selection methods are mostly fast and efficient, it is difficult to find the best feature subset. In addition, to find the optimal features, either the sequential forward search method (SFS), which adds features individually starting with the one with the highest importance, or the sequential backward search (SBS) method, which removes the low-importance features individually from the entire data set, can be used. It should be found using [8]. These methods are not a big problem when the number of features in a dataset is small, but when this number is large, considerable computation time and computing power are required. Although feature selection methods using heuristic optimization algorithms, such as genetic algorithms, work flexibly on multiple datasets, it is difficult to make detailed selections because features are judged based on a combination of features rather than individually [9][10][11].

Reinforcement learning (RL) is one of the best techniques to make optimal decisions, and many studies have recently been conducted on this topic [12]. Similar to the method using the heuristic optimization algorithm, the newly proposed feature selection methods using RL operate flexibly on any dataset and can be adjusted for each feature, enabling a detailed judgment. Feature selection methods using RL, such

as the heuristic optimization algorithm, work well for multiple datasets, but they are inefficient because they require significant computation time and computing power to search a large number of states [13].

The objective of this study was to develop an effective and efficient feature selection method using RL algorithms. If no additional algorithm advice is provided or there are no criteria to evaluate the agent's own behavior, the agent is unaware of the purpose of the problem to be solved, and there is a high probability that the user will behave unexpectedly. Therefore, we proposed an agent, called a guide agent, and designed it to evaluate its own behavior. Because these guide agents randomly set their behavior, they can be flexibly operated on any dataset, unlike external algorithms. Figure 1 shows the flow of the multi-agent reinforcement learning feature selection method using a guide agent (MARFS-GA) proposed in this study. For a given dataset, two groups of agents, called the main agents and the guide agents, perform their respective actions to select features. Two rewards are calculated from the selected features using a classification algorithm. The algorithm proposed in this study uses these two rewards to obtain a learning reward and trains the main agent responsible for feature selection. Using the proposed algorithm, experiments were conducted on six datasets, and the superiority of the proposed feature selection method was verified through experimental results.

Related Works

In 2019, Liu proposed a multi-agent reinforcement learning framework for the feature selection problem. In this framework, by assigning an agent to each feature, feature selection is performed using an RL framework. Then, the state is obtained using a statistical description, autoencoder, and graph convolutional network (GCN) for the selected feature subset, so that the algorithm can sufficiently understand the learning progress. It achieved the highest accuracy of 0.8731 for the forest cover type dataset using the above framework [14]. In 2020, Fan et al. proposed an algorithm that balances efficiency and effectiveness by combining feature selection methods using RL and other feature selection methods, such as K-best, decision-tree-based ranking, and mRMR. For each agent of the RL feature selection algorithm, the aforementioned K-best, decision tree-based ranking, and mRMR, are designated as external trainers provide advice to the agents. Using this method, they achieved accuracies of 0.80 ~ 0.81, 0.93, 0.91, and 0.98 for the forest cover type, Spambase, insurance company benchmark, and Musk datasets, respectively [15]. A paper published by Khurana et al. in 2018 presented a framework using RL to automate feature engineering. They presented a method for the exploration of a transformation graph that systematically and compressively captures the space of a given option. After listing a subset of a given dataset through a transformation graph, the agent explores the space of the listed feature choices. This method demonstrated an accuracy of 0.961 for the Spambase dataset [16]. A paper published by Rasoul et al. in 2021 formalized the state space as a Markov decision process and used the temporal difference algorithm to find a subset of the best features. It achieved an accuracy of 0.7629 on the breast cancer Wisconsin dataset [17]. Finally, the method presented by Fard et al. in 2013 considers the state space of a feature as a Markov decision process, and then proceeds with an optimal graph search. Each state is evaluated using the lattice of the feature sets. Finally, the optimal features are selected using a

method based on filters and wrappers. This method demonstrated an accuracy of 0.7306 for the colon cancer dataset [18].

Methods

Definitions

Agents (Agt): One agent is assigned to each feature. For example, assuming 30 features, 30 agents are created. The agents assigned to each feature determine whether to select a feature. There are two types of agents: main and guide agents.

Main Agents (Agt_m): The main agents are agents that perform feature selection. They perform the main action and learn by receiving a reward based on the action.

Guide Agents (Agt_g): Guide agents are those that are not involved in feature selection. They do not learn even if they receive a reward for their actions, and all their actions are randomly determined.

Action (A): Each agent decides whether to select a feature. The action of the i^{th} agent A_i is either zero or one. Zero means that no feature is selected, and one means that a feature is selected. In this study, there are two types of actions: the main action, which is the action performed by the main agents, and the guide action, which is performed by the guide agents.

Main Actions (A_m): The main action is performed by the main agents. It is an indicator for selecting a feature and has a Q-value that indicates the value of each action.

Guide Actions (A_g): Guide actions are performed by guide agents. Guide behavior is an index for evaluating whether the main agents are correct when performing the main behavior. Unlike the main action, a guide action does not have a Q-value that indicates the value of each action.

Environment (E): Environment refers to the environment in which the agents operate. In this study, the entire dataset is defined as the environment.

State (S): In RL, the state refers to the situation after the agents have performed actions in the environment. In this study, the state is defined as a subset of features selected based on the actions of the main agents for each episode.

Reward (R): The reward is based on the actions of the agents. In the case of classification problems, the reward is defined as the classification accuracy of the selected features. In this experiment, there were three types of rewards: main, guide, and learning.

Main Reward (R_m): This is the reward for the characteristics selected through the main actions performed by the main agents. It is defined as the classification accuracy of the features selected through the main actions. It is not used for learning the value of the main agents' actions.

Guide Reward (R_g): This is a reward for the characteristics selected through guide actions performed by the guide agents. It is defined as the classification accuracy of the features selected through the guide actions. Similar to the main reward, it is not used as a reward for learning the value of the main agents' actions.

Train Reward (R_t): This reward is used to learn the values of the main agents' actions. Learning reward is defined as the difference between the main and guide rewards.

Experiments

Framework Overview: Fig. 2 shows the overall framework of the experiment. The main and guide agents perform their respective actions. The main and guide actions performed by them interact with the environment. Through interaction, the characteristics selected by actions, state, and rewards (main reward & guide reward) based on the state are obtained. Then, train reward is calculated using the derived reward, and the main agents selected by the training strategy are trained with the train reward value.

Initializing Main and Guide Agents: Fig. 3 illustrates the creation of the main and guide agents. When the algorithm starts, the number of main and guide agents created each is equal to the number of features in the data set. For example, if there are 30 features in one dataset, 30 main agents and 30 guide agents are created. The agents in each group perform actions on the same dataset. The main agent performs feature selection, and the guide agent is used by the main agent to evaluate its own behavior. Therefore, the main agents have a Q-value indicating the value of their actions, and the guide agents do not have a Q-value because they are only evaluation targets.

Actions of main agents: When the episode starts, the main agents decide their actions according to the epsilon-greedy policy. This policy conducts exploration with epsilon (ϵ) probability based on the epsilon (ϵ) value, and exploits it with $1-\epsilon$ probability [21]. Here, the epsilon (ϵ) value has a real value between 0 and 1. Further exploration reveals that all the actions of the main agents are randomly determined. In the case of exploitation, the actions are selected based on the Q-value of each main agent. The Q-value value indicates how good each action is when the action is performed, and it exists for each of the two actions (Select: 1 / Deselect: 0) that each main agent can take. The agent performs the action with the larger Q-value. In this case, the action of the i^{th} main agent can be expressed as follows: A_m^i is the behavior of the i^{th} main agent and $\text{rand}(0 \sim 1)$ is a random real value between 0 and 1. $Q(1)$ and $Q(0)$ are the Q-values of the feature-selecting and non-selecting actions, respectively.

$$A_m^i = \begin{cases} 0 \text{ or } 1, & \text{if } \epsilon > \text{rand}(0 \sim 1) \\ \begin{cases} 1, & \text{if } Q(1) \geq Q(0) \\ 0, & \text{if } Q(1) < Q(0) \end{cases} & \text{if } \epsilon \leq \text{rand}(0 \sim 1) \end{cases}$$

1

When one episode ends, the epsilon value is reduced by the epsilon decay rate (EDR) value. The EDR has a real value between 0 and 1. Thus, the ratio of utilization between exploration and exploitation gradually increases. This is expressed as:

$$\epsilon = \epsilon \times \text{EDR}$$

2

Actions of guide agents: When the episode starts, guide agents, unlike the main agents, randomly determine all their actions. There is no Q-value for the actions of the guide agents, and the actions performed at time t do not have any effect on the actions performed at time $t + 1$. The action of the i^{th} guide agent is expressed as follows:

$$A_g^i = 0 \text{ or } 1$$

3

Reward calculation method: The main and guide rewards are defined as the accuracy of features selected through the main actions performed by the main agents and guide actions performed by the guide agents, respectively. The learning reward is defined as the difference between the main and guide rewards; when the main reward is higher, it becomes a positive reward, and when the guide reward is higher, it becomes a negative reward. The reward calculation method is expressed as follows:

$$R_t = R_m - R_g$$

4

where R_t is the learning reward used for learning the values of the main agent, R_m is the main reward calculated by the action of the main agent, and R_g is the guide reward calculated by the action of the guide agent.

Training strategy for updating the Q-value of main agents: After one episode, not all the main agents learn. If the main action is performed by the main agent at the same location it is in, and the guide action performed by the guide agent, are different from each other, the main agent corresponding to the location learns. For example, if the third main and guide actions are the same, the third main agent is not rewarded with a train reward. Conversely, if the fourth main and guide actions are different from each other, the 4th main agent will be rewarded with the train reward for the action it took. Figure 3 shows a schematic of the training strategy.

There are two reasons for learning when different actions are performed. First, the difference between the accuracies of the main and guide actions is that different actions occur at the same location; hence, only the agent at that location learns. Second, if a train reward is given to all main agents, detailed adjustment of the main agents' actions become impossible. This is because if the accuracy of the main action

continues to be higher than that of the guide action, all main agents will continue to receive rewards as it cannot filter out less important features such as noise or noise.

When the main agents that receive the reward are determined through the above process, each main agent multiplies the Q-value of the action they execute and adds them to the learning rate (α) value. At this time, the α value is responsible for adjusting the amount of train reward reflected, and has a real value between 0 and 1. For the action of the i^{th} main agent, the Q-value update expression through the train reward can be expressed as:

$$Q(A_m^i) = Q(A_m^i) + (\alpha \times R_t)$$

5

where $Q(A_m^i)$ is the Q-value for the action of the i^{th} main agent, α is the learning rate used to designate the degree of learning, and R_t is the reward for the action of the main agent.

Algorithm

The above process is expressed as Algorithm 1.

Algorithm 1. Multi-Agent Reinforcement Learning Feature Selection Method with Guide Agents (MARFS-GA)

Input: Number of Features: N , Epsilon value: ϵ , Learning Rate: α , Number of Episodes: E , Epsilon Decay Rate

1. **for** $i = 1$ to N **do** # make Agents
2. make Agt_m^i, Agt_g^i
3. **End for**
4. **for** $i = 1$ to N **do** # initialize Q-values
5. Q-value of Agt_m^i [random value, random value]
6. **End for**
7. **for** $E = 1$ to Number of Episode **do**
8. **for** $i = 1$ to N **do**
9. **if** random value $> \epsilon$ **do**
10. Agt_m^i do Exploitation
11. **else do**
12. Agt_m^i do Exploration
13. **End for**
14. **for** $i = 1$ to N **do**
15. **if** random value $> \epsilon$ **do**
16. Agt_g^i do Random Action
17. **End for**
18. Selected Features by $A_m = \text{Select Features}(A_m)$
19. Selected Features by $A_g = \text{Select Features}(A_g)$
20. $R_m = \text{get Reward}(\text{Selected Features by } A_m)$
21. $R_g = \text{get Reward}(\text{Selected Features by } A_g)$
22. $R_t = R_m - R_g$ # Classification Problem
23. **for** i to N **do**
24. **if** $A_m^i \neq A_g^i$

Algorithm 1. Multi-Agent Reinforcement Learning Feature Selection Method with Guide Agents (MARFS-GA)

25. Q-value of $A_m^i + = \alpha \times R_t$

26. **End for**

27. $\epsilon = \epsilon \times \text{Epsilon Decay Rate}$

28. **End for**

29. **Return** Selected Feature Subset

Results

Datasets

Breast cancer Wisconsin: This is a breast cancer dataset provided by the University of Wisconsin comprising 212 negative samples and 357 positive samples, for a total of 569. The total number of features is 32, including the radius, texture, and perimeter [19].

Forest cover type: This dataset categorizes seven forest cover types and consists of 581,012 samples. It has 54 features including elevation, quantitative, slope, and soil types [19].

Spambase: This is a short message service (SMS) dataset with a label indicating whether a message is spam or non-spam. It contains 4,601 samples, consisting of 1,813 Spam messages and 2,788 non-spam messages. There are 57 features in this dataset, and each feature consists of information about the number of specific characters appearing in a message [19].

The insurance company benchmark: This dataset is labeled according to whether a customer has purchased Caravan insurance. There are 9,822 samples, consisting of 586 subscriber data and 9,236 non-subscriber data. This dataset has 86 features, which are composed of product usage data and socio-demographic data derived from zip area codes [19].

Musk: This a dataset with a Musk or non-Musk label for each data. It has 6,598 samples that consist of 1,017 Musk data and 5,581 non-Musk data, and 168 features. The first two attributes of the Musk dataset were excluded from this experiment because they refer to the names of molecules and forms [19].

Colon cancer: This is a dataset provided by Princeton University. It has 62 samples of which 40 contain normal data and 22 contain colorectal cancer patient data. It has 2,000 features, and each feature represents gene information [20].

Experiment Settings

The composition of the variables in this experiment were as follows: The initial epsilon (ϵ) value was 0.5, and the EDR was 0.9995. The number of main and guide agents generated depends on the number of features in the dataset being tested. The learning rate was 0.01 and the maximum number of episodes was 10,000. For the initial main agent action, the Q-value corresponding to 0 (Deselect) was randomly initialized to a real value between 0 and 1, and the Q-value corresponding to 1 (Select) was randomly initialized to a real value between 0 and 0.05. This was intended to allow a small number of features to be initially selected, and gradually increase the number.

In this experiment, an artificial neural network was used as the classifier. It consists of an input layer, two hidden layers, and an output layer. The two hidden layers consist of five and two nodes, respectively. Furthermore, the rectified linear unit function was used as the activation function and the learning rate was set to 0.01. Epochs were assigned a maximum score of 10.

In addition, to compare and verify the performance of the proposed method, feature selection was performed using the mRMR, Relief, and genetic algorithm. They were implemented and directly tested. In the case of the genetic algorithm, only a basic concept was used without introducing a special method, and the fitness value of the genetic algorithm was defined as the classification accuracy. In this case, a classifier for obtaining classification accuracy had the same configuration as that of the aforementioned classifier.

Figure 4 shows the number of features selected for each episode in each dataset, and Fig. 5 shows the classification accuracy for each episode. For each experiment, 10,000 episodes were performed, and the values plotted on the graph are the average values of the number and accuracy of the features derived per 100 episodes.

It is evident that the experiments conducted with the guide agent increase the classification accuracy compared with the experiments that do not. In the experiments wherein the guide agent is applied, the classification accuracy increases as the episode progresses, but in the ones without the guide agent, there is either no significant increase compared to the initially derived accuracy or the accuracy is constant. It is not always maintained and appears to rise or fall continuously. In the experiments that did not apply the guide agent, all agents received the initial result values as a reward without any strategy. Therefore, it is evident that this is because the Q-value for one action in the beginning increases considerably and the opportunity to take another action decreases. In addition, the continuous change in accuracy can be attributed to a behavioral method of agents, called exploration. Exploration causes an agent to perform various actions, and it appears that the change in accuracy is significantly caused by such an exploration action. Therefore, the accuracy of the selected characteristics can be seen to vary greatly when there are significant changes in the types of selected characteristics because this greatly affects the change in the behavior of the entire agent.

Through various experiments, it was deduced that the method in which the guide agent is applied does not change the number of selected features and the accuracy constantly increases. By applying the guide agent, the main agent can reliably determine whether their actions are correct, and by learning only a

small number of features through the proposed learning strategy, large fluctuations can be avoided. However, in the case of the forest cover type dataset, the selection number of features changes significantly compared to other datasets. This is when many of the features of the forest cover type dataset have meaningless values. These features are filtered out during the learning process, while other features are selected; hence, the amount of change is large compared to other datasets.

Table 1 lists the experimental results for each dataset, where each numerical value represents the classification accuracy. The numbers on the left in parentheses indicate the total number of features in the dataset, and the numbers on the right indicate the number of features finally selected by applying our algorithm. Evidently, compared with other experiments, our proposed method achieves a slightly better performance. In other methods, there were no data on the number of selected features; therefore, a comparison of the number of features could not be made. The results of six experiments show that our proposed feature selection method enables efficient feature selection for classification, regardless of the number of features in the dataset.

Table 2 lists the results of our proposed method, and the feature selection methods commonly used in other studies. The results obtained after applying each feature selection method are that of the features showing the greatest accuracy.

In the case of the Wisconsin breast cancer dataset, there was no significant difference in accuracy between the four feature selection methods and our proposed method; however, the mRMR method achieved the highest accuracy. A satisfactory result of 0.9404 was obtained even when the experiment was conducted without feature selection. All features of the Wisconsin breast cancer dataset have some significant characteristics. Moreover, from the results in Table 2, it can be observed that the linear data analysis method is the most efficient for this dataset.

For the forest cover type dataset, our algorithm achieved the best result, with an accuracy of 0.8802. In the case of this dataset, the accuracy improved by approximately 0.17 or more when the feature selection method was used. However, it is judged that the linear analysis method, that is, the selection method based on a combination of features, does not show very good performance because the improved classification accuracy is not very high.

For the Spambase dataset, our algorithm achieved the best results, with an accuracy of 0.963. Compared with other experimental results, the Genetic Algorithm and proposed method achieved high accuracy. From the results of these experiments, in the case of this dataset, the feature selection method based on a combination of features achieved better results than the linear analysis method.

In the case of the insurance company benchmark dataset, our algorithm achieved the second-highest result with an accuracy of 0.943. The feature selection method using Relief achieved the best accuracy. A similar method, mRMR, also achieved high accuracy. The genetic algorithm's accuracy was relatively low. Therefore, in the case of the insurance company benchmark dataset, there seems to be linear relationship between features and labels.

For the Musk dataset, the proposed method achieved the highest accuracy of 0.984. Compared with when no method was used, it achieved a significant accuracy improvement when feature selection was performed. It can be observed that the Musk dataset contains many meaningless features. Compared with the other algorithms, the genetic algorithm and our proposed method achieved higher classification accuracy. However, in the case of the experiment using mRMR, it can be assumed that there is linearity between each characteristic and the label of the Musk dataset, as it shows high accuracy.

The experiment with the colon cancer dataset showed the greatest difference compared to the other dataset experiments. Feature selection using the linear analysis method achieved a slightly higher accuracy than when no method was used. However, the feature selection methods based on combinations, such as the genetic algorithm and our algorithm, showed superior performance. In addition, the difference in the results obtained using our proposed feature selection method and the genetic algorithm was 0.0364; a difference of approximately 3.64%.

Conclusion

Our proposed feature selection algorithm selects features that are effective for classification in most datasets. Existing feature selection methods also performed well for most datasets, but our proposed algorithm showed surprisingly better results, especially for datasets where it is difficult to determine a specific relationship, such as the colon cancer dataset. Widely used feature selection methods, such as mRMR and Relief, define the relationship between data by using formulas. Therefore, it is difficult to apply them to data that show high performance when several features are applied together, although the relationship between them is weak. In addition, the feature extraction method using the genetic algorithm, which is used as an alternative, is similar to the direction we are pursuing, but has a limitation in that it focuses on the combination of features. In this study, the proposed algorithm can learn more flexibly by learning each feature individually or a combination of features. To design the method proposed in this paper, we focused on emphasizing individual features rather than focusing on a combination of features. Because it is tailored, a detailed selection is not performed well. Although most of the same features were selected whenever the experiment was repeated, problems such as features mixed with noise continued to occur. In addition, although all guide agents' actions are randomly assigned, only two types of actions can be taken. Therefore, if one action occurs continuously, there is a high possibility that the action of the main agent corresponding to that location will be concentrated on one side. Therefore, in the future, we plan to develop a learning method for each agent to select features in more detail.

Declarations

Acknowledgments

This study was supported by the Gachon University Research Fund of 2019 (GCU-2019-2019-0766). The research was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the Ministry of Education(2020R111A1A01066599).

Data Availability

The Breast Cancer Wisconsin, Spambase, The Insurance Company Benchmark and Musk datasets generated and/or analyzed during the current study are available in the 'UCI Machine Learning repository(<http://archive.ics.uci.edu/ml>)' and the Colon Cancer dataset are available in the 'Princeton University Gene Expression Project(<http://genomics-pubs.princeton.edu/oncology>)'

References

1. Guyon, I. & Elisseeff, A. An introduction to variable and feature selection. *J. Mach Learn Res.* 1157–1182 (2003).
2. Bousquet, O., Boucheron, S. & Lugosi, G. Introduction to statistical learning theory in *Summer school on machine learning* (eds. Bousquet O., von Luxburg U., Rätsch G.) 169–207 (Springer, Berlin, Heidelberg, 2003).
3. Bermingham, M. L. et al. Application of high-dimensional feature selection: evaluation for genomic prediction in man. *Sci Rep.* **5**(1), 1–12 (2015).
4. Cai, J., Luo, J., Wang, S. & Yang, S. Feature selection in machine learning: A new perspective. *Neurocomputing* **300**, 70–79 (2018).
5. Peng, H., Long, F. & Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy." *IEEE Trans. Pattern Anal. Mach. Intell.* **27**(8), 1226–1238 (2005).
6. Kira, K. & Rendell, L. A. A practical approach to feature selection in *Machine learning proceedings 1992* (eds. Sleeman, D. and Edwards, P.) 249 – 25 (Morgan Kaufmann, 1992).
7. Mirjalili, S. Genetic algorithm in *Evolutionary algorithms and neural networks* 43–55 (Springer, Cham, 2019).
8. Pudil, P., Novovičová, J. & Kittler, J. Floating search methods in feature selection in *Pattern Recognit. Lett.* **15**(11), 1119–1125 (1994).
9. Babatunde, O. H., Armstrong, L., Leng, J. & Diepeveen, D. A genetic algorithm-based feature selection. *Asian J. Comput. Inf. Syst.* (2014).
10. Frohlich, H., Chapelle, O. & Scholkopf, B. Feature selection for support vector machines by means of genetic algorithm. *Proceedings. 15th IEEE International Conference on Tools with Artificial Intelligence, IEEE.* 142–148 (2003).
11. Tan, F., Fu, X., Zhang, Y. & Bourgeois, A. G. A genetic algorithm-based method for feature subset selection. *Soft Comput.* **12**(2), 111–120 (2008).
12. Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction.* (MIT press, 2018).
13. Tesauro, G. et al. Managing Power Consumption and Performance of Computing Systems Using Reinforcement Learning in *Advances in Neural Information Processing Systems* (eds. Platt, J., Koller D., Singer Y., and Roweis, S.) **7**, 1–8 (Curran Associates, Inc., 2007).

14. Liu, K. et al. Automating feature subspace exploration via multi-agent reinforcement learning. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 207–215 (2019).
15. Fan, W. et al. Autofs: Automated feature selection via diversity-aware interactive reinforcement learning. *2020 IEEE International Conference on Data Mining (ICDM), IEEE*. 1008–1013 (2020).
16. Khurana, U., Samulowitz, H. & Turaga, D. Feature engineering for predictive modeling using reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*. **32(1)** (2018).
17. Rasoul, S., Adewole, S. & Akakpo, A. Feature Selection Using Reinforcement Learning. Preprint at arXiv:2101.09460v1 (2021).
18. Fard, S. M. H., Hamzeh, A. & Hashemi, S. Using reinforcement learning to find an optimal set of features. *Comput. Math. Appl.* **66(10)**, 1892–1904 (2013).
19. Asuncion, A. & Newman, D. UCI machine learning repository. Available at <http://archive.ics.uci.edu/ml> (2007).
20. Alon, U. et al. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *Proceedings of the National Academy of Sciences*. **96(12)**, 6745–6750 (1999)
21. Tokic, M. & Palm, G. Value-difference based exploration: adaptive control between epsilon-greedy and softmax in *Annual conference on artificial intelligence* 335–346 (Springer, Berlin, Heidelberg, 2011).

Figures

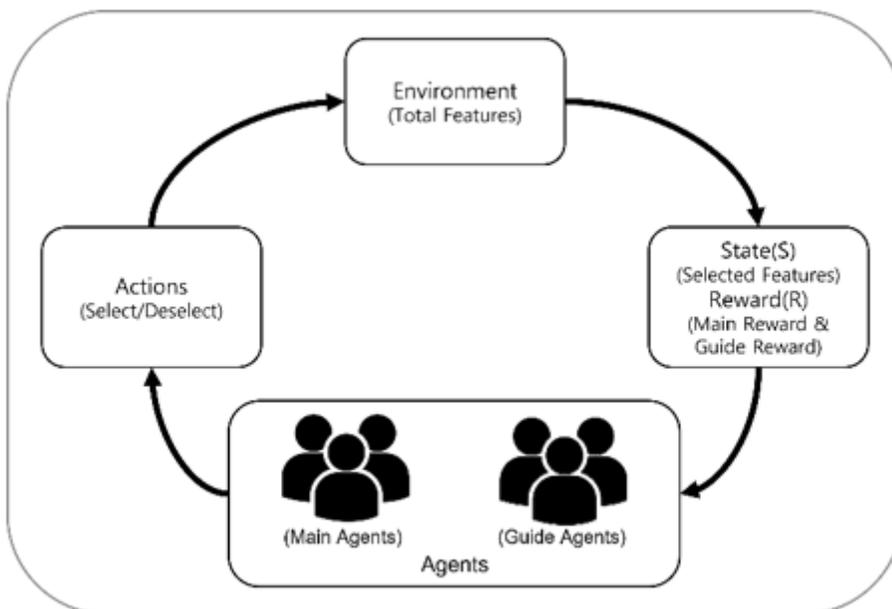


Figure 1

Illustration of Multi-Agent Reinforcement Learning Feature Selection Method with Guide Agents (MARFS-GA)

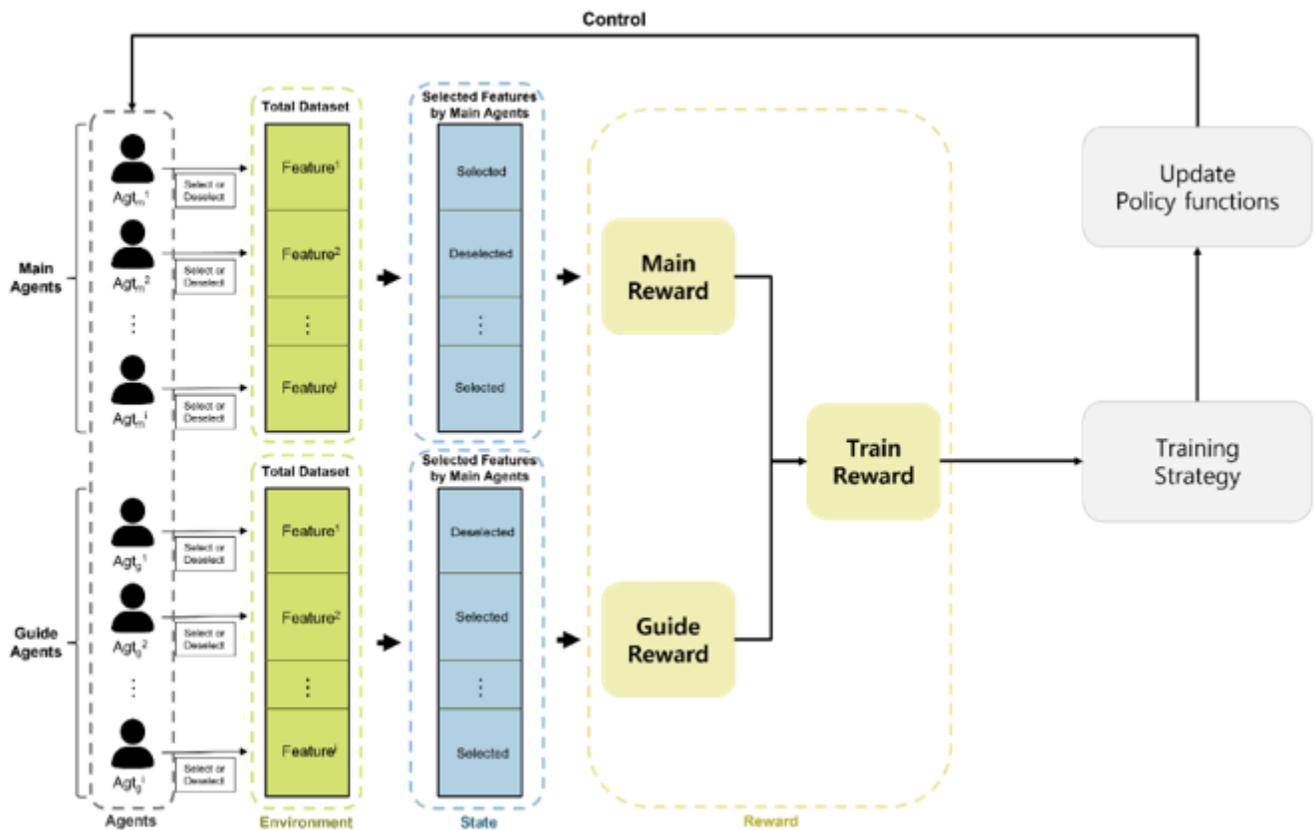


Figure 2

Framework Overview

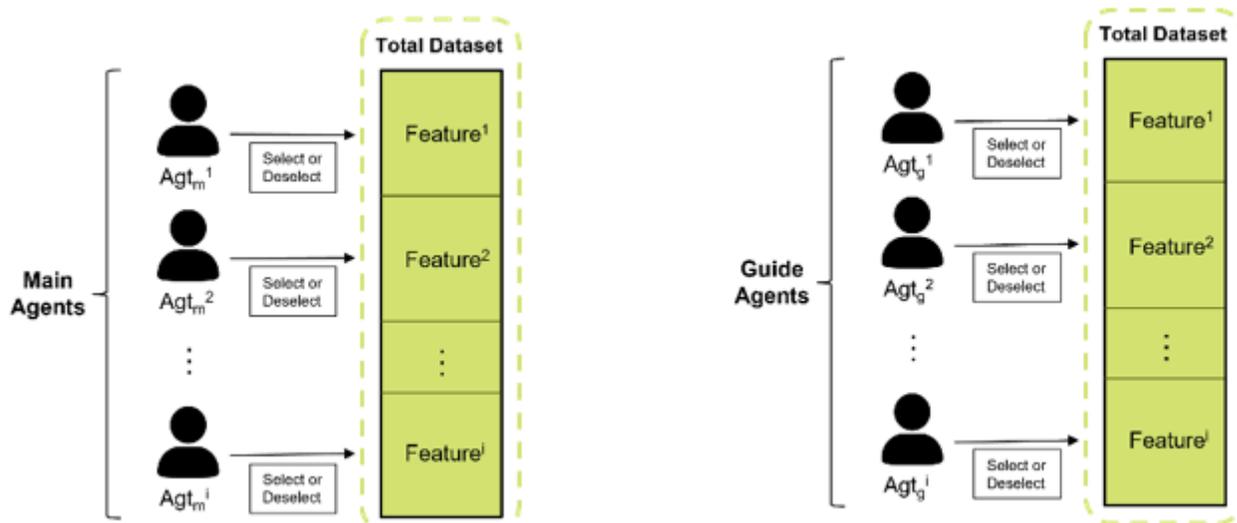


Figure 3

Initializing Main Agents and Guide Agents

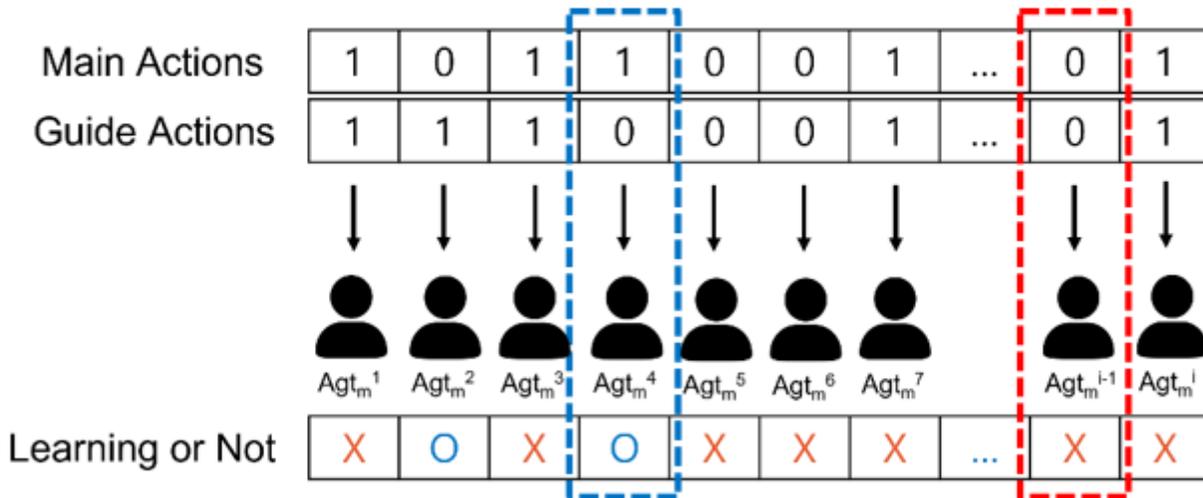


Figure 4

Training Strategy



Figure 5

Number of selected features per episode, for each dataset



Figure 6

Accuracy of selected features per episode, for each dataset