Automating Feature Subspace Exploration via Multi-Agent Reinforcement Learning

Kunpeng Liu, Yanjie Fu,
Pengfei Wang, Le Wu, Rui Bo, Xiaolin Li

Outline

- **Background and Motivation**
- Problem Statement
- Methodology
- Evaluation
- Conclusion

Feature Selection

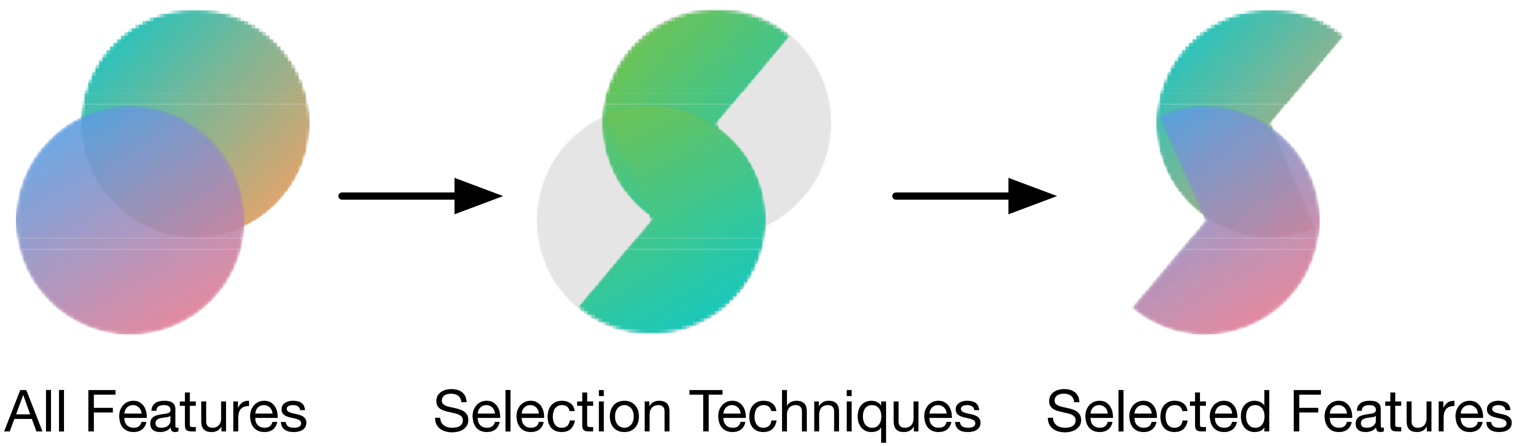

Feature Selection as An Exploration Process

Feature selection: An iterative exploration process to find an optimal / near optimal subset of features

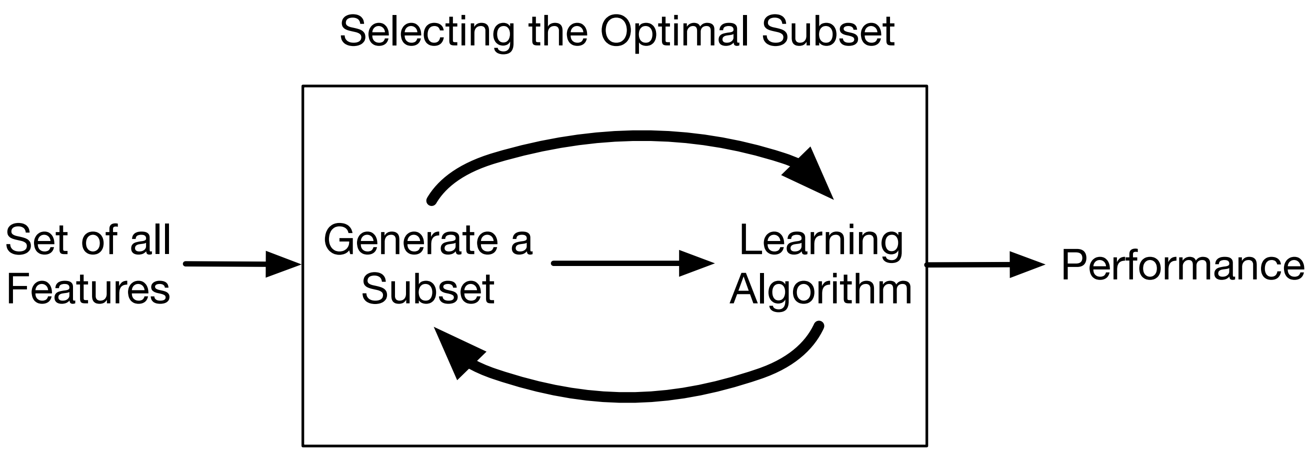

Reinforcement Learning as A Tool of Exploration

- Reinforcement learning: **exploration + exploitation**

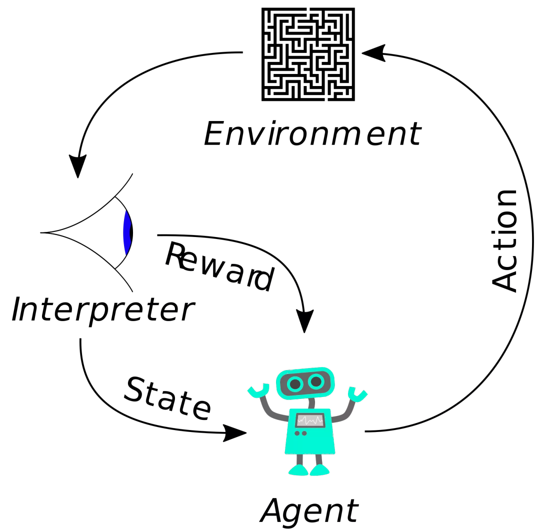

- Applications

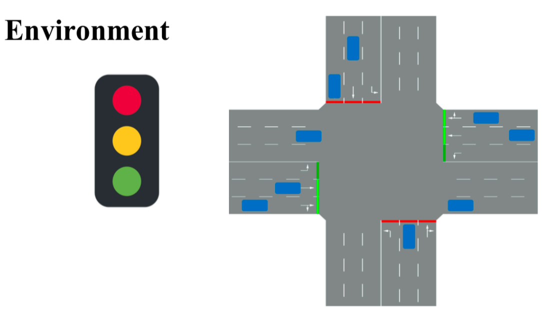

Traffic light control via RL

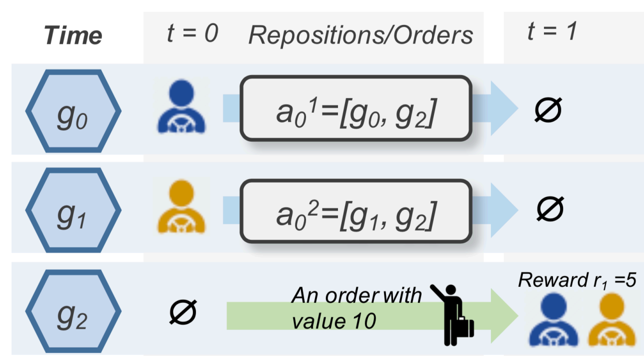

Taxi fleet management via RL

Automated Feature Subspace Exploration

Inspiration: Can reinforcement learning help to solve/improve feature selection?

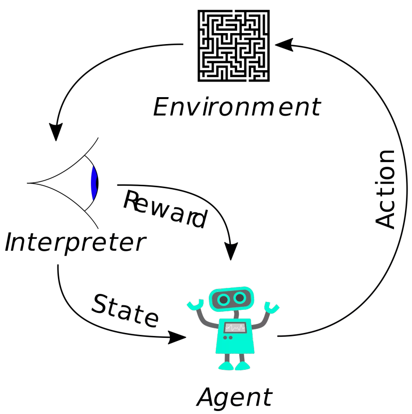

Reinforcement Learning
(Exploration Tool)

+

Selecting the Optimal Subset

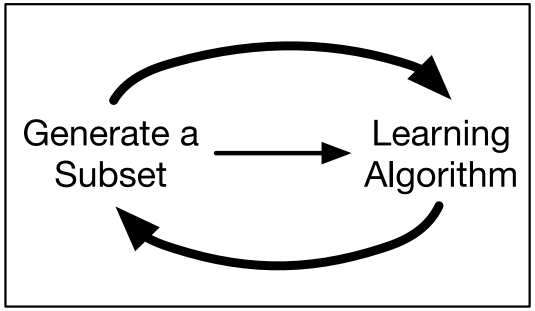

Feature Selection
(Exploration Problem)

= ?

Question 1 of 4

How can we reformulate the feature selection problem into a reinforcement learning task?

Feature Selection as A Multi-Agent Reinforcement Learning Task (1)

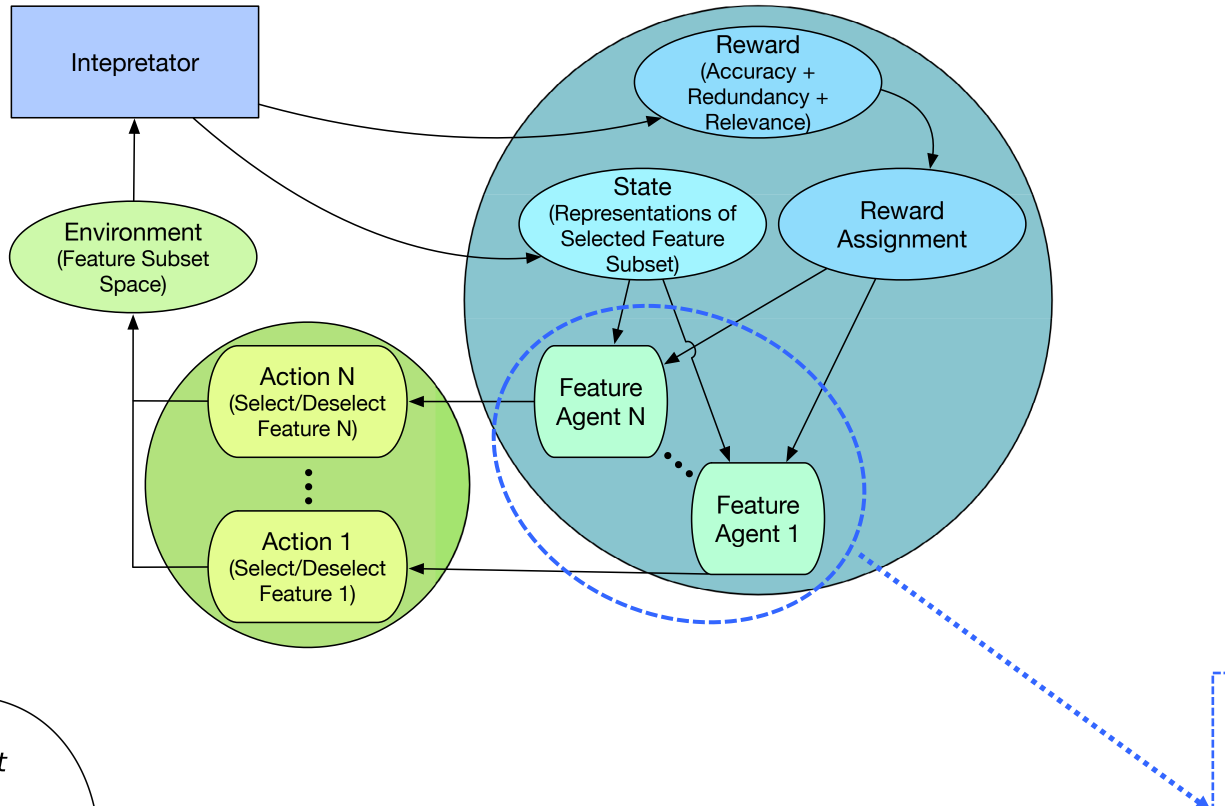

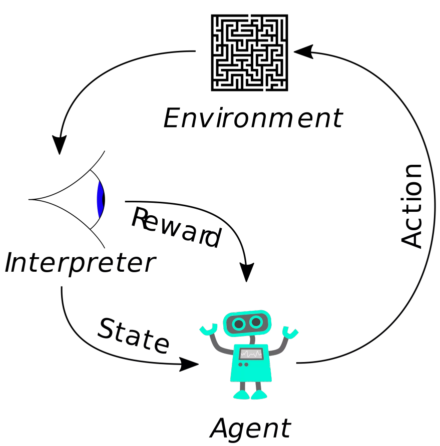

Agent: Each feature is controlled by a corresponding feature agent.

Feature Selection as A Multi-Agent Reinforcement Learning Task (2)

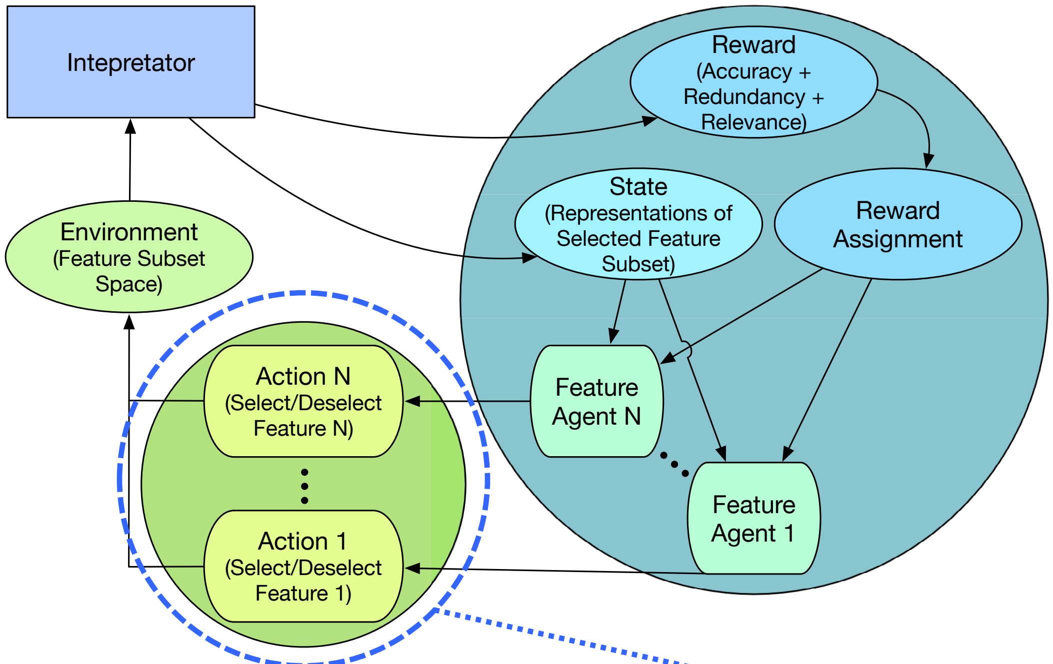

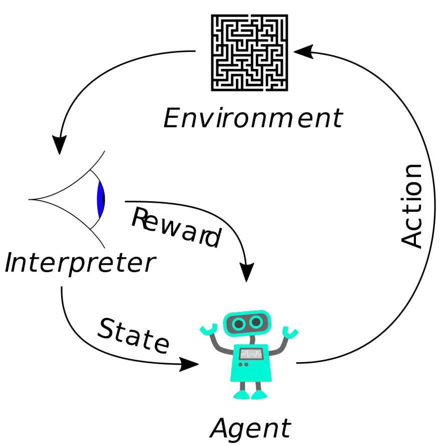

Action: Select or deselect the corresponding feature.

Feature Selection as A Multi-Agent Reinforcement Learning Task (3)

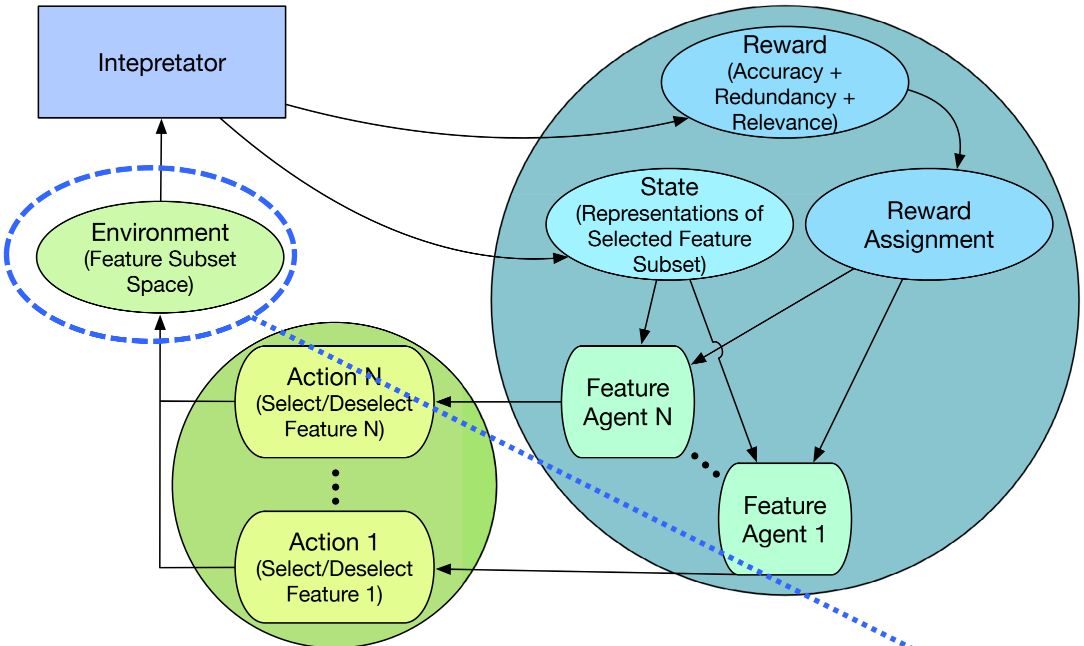

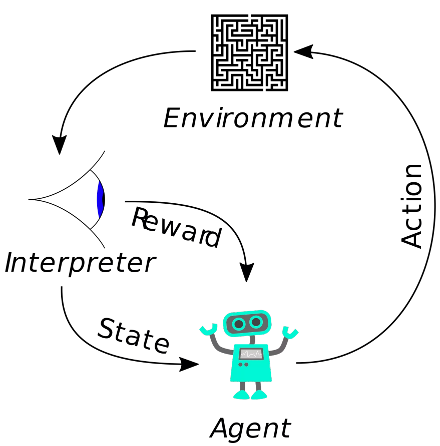

Environment:
Feature subset space: {{f4}, {f1,f2}, {f1,f3,f5}…}.

Feature Selection as A Multi-Agent Reinforcement Learning Task (4)

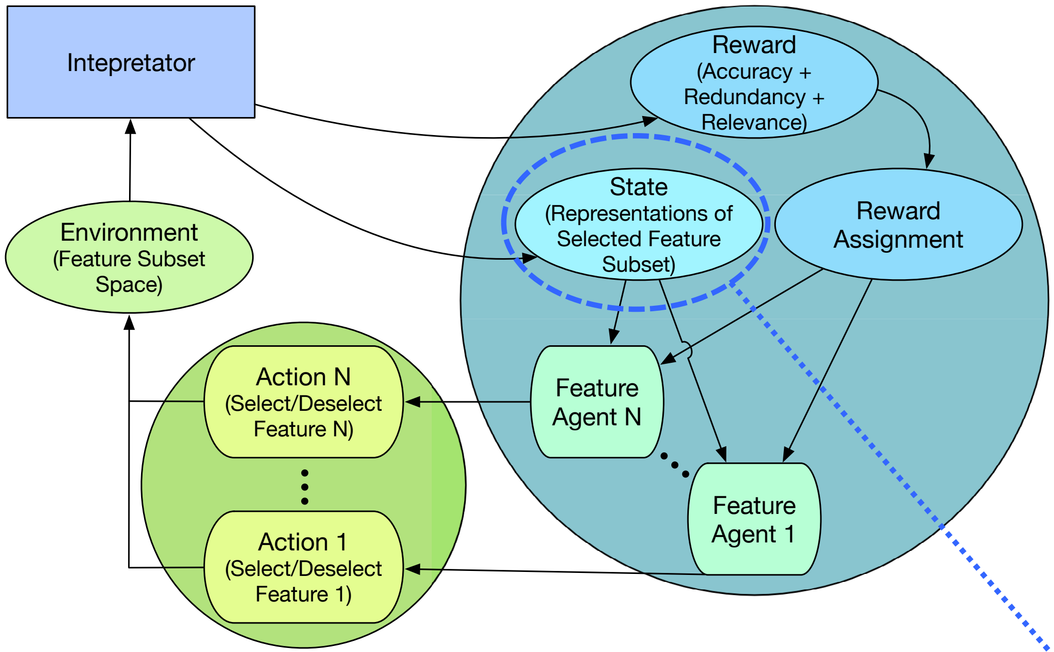

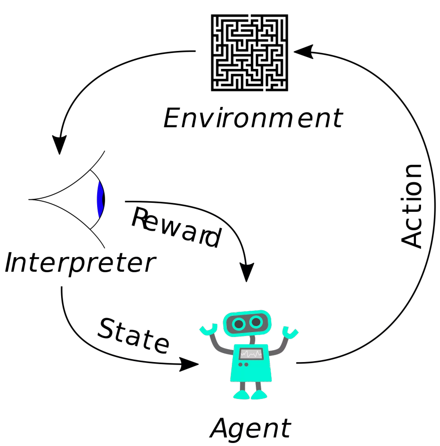

State: Representation of selected feature subset. E.g., S({f1,f3,f5}).

Feature Selection as A Multi-Agent Reinforcement Learning Task (5)

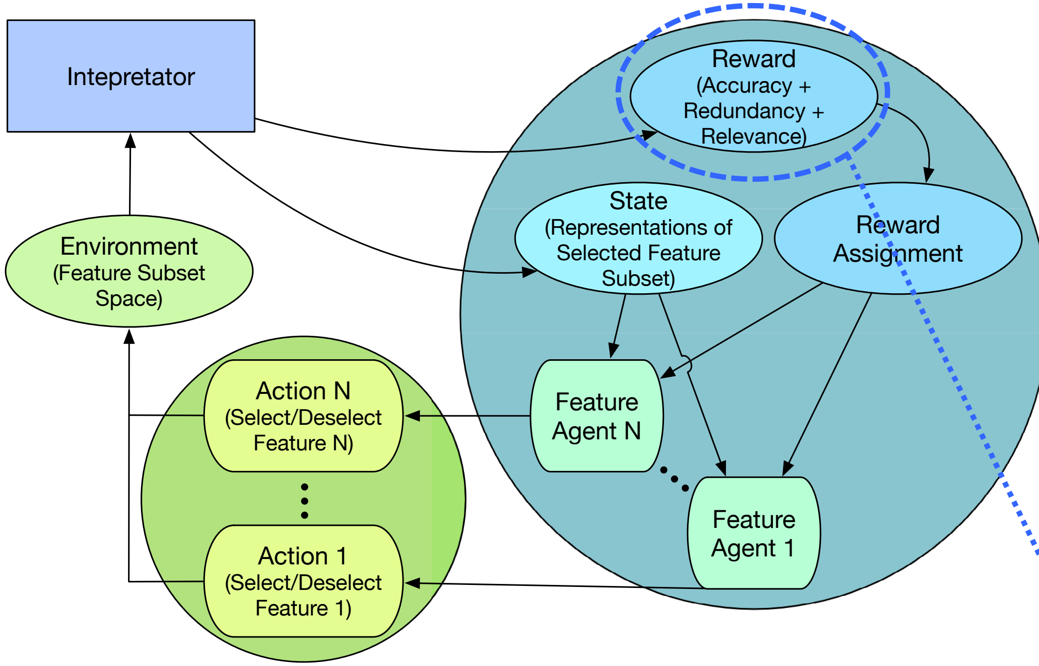

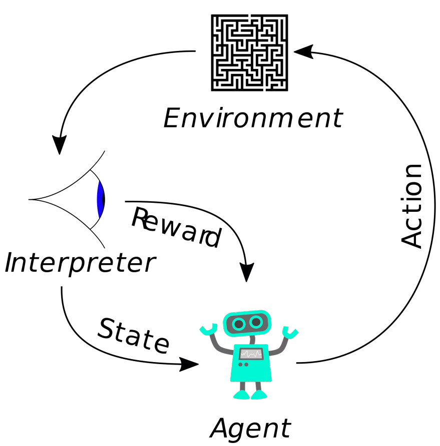

Overall Reward: Weighted sum of prediction accuracy, redundancy and relevance of selected feature subset.

Feature Selection as A Multi-Agent Reinforcement Learning Task (6)

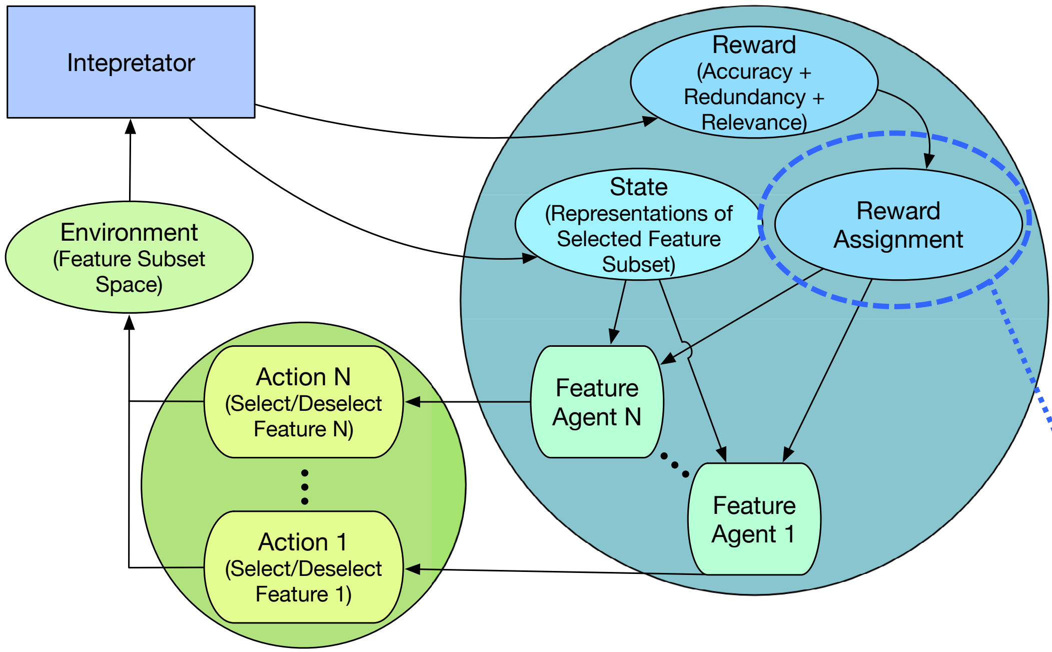

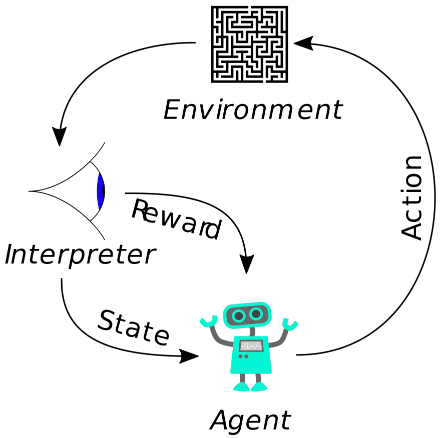

Reward Scheme: Assign overall reward to each agent.

Question 2 of 4

How can we design the assignment strategy?

Participating & Non-participating Agents

- For the current k_{th} iteration
- Participating feature agents:
  - **Select** action (k_{th} iteration) & **Select** action($(k-1)_{th}$ iteration)
  - **Select** action (k_{th} iteration) & **Deselect** action($(k-1)_{th}$ iteration)
  - **Deselect** action (k_{th} iteration) & **Select** action($(k-1)_{th}$ iteration)
- Non-participating agents:
  - **Deselect** action (k_{th} iteration) & **Deselect** action($(k-1)_{th}$ iteration)

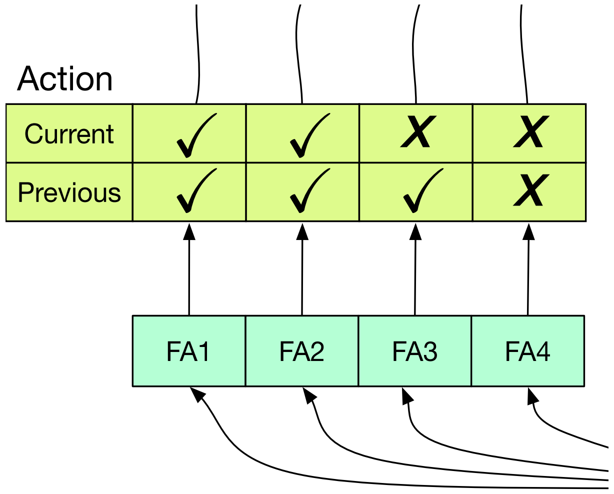

Reward Assignment Strategy

- Participating agents
  - Equally share the overall reward.
- non-participating agents:
  - 0 reward.

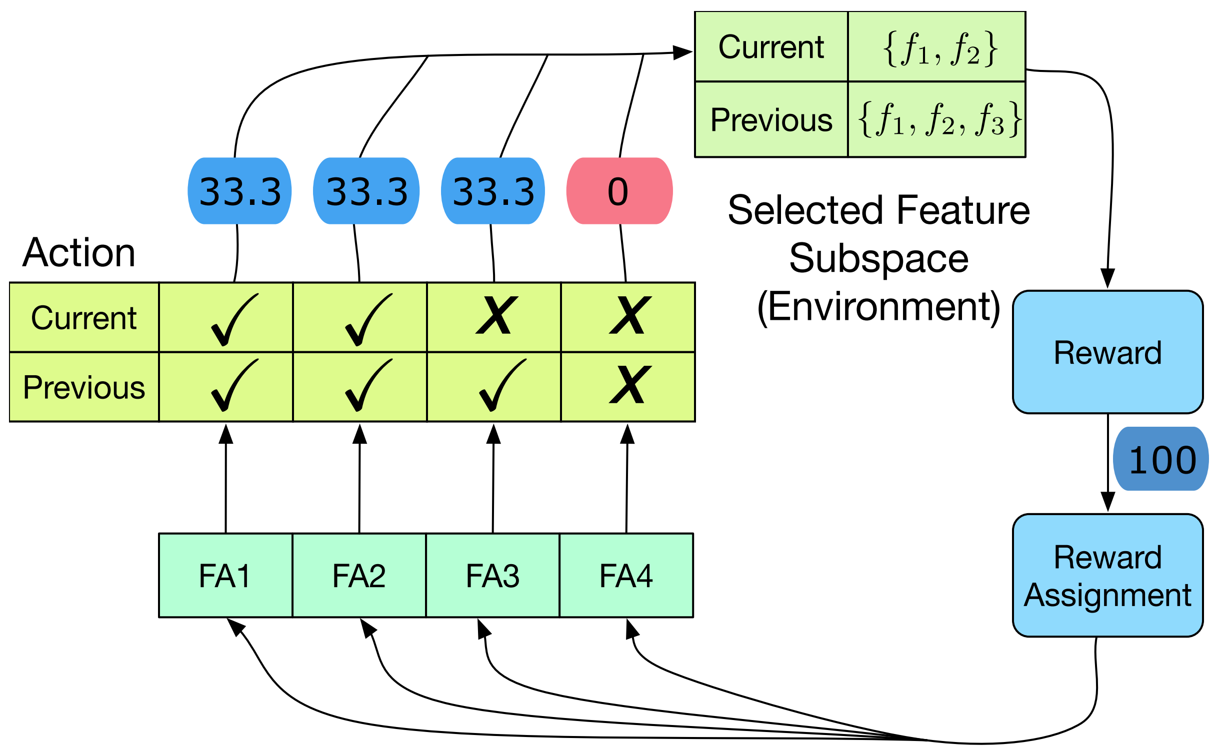

Question 3 of 4

How can we better quantify the state representation?

Three State Representation Methods

- Meta descriptive statistics.
- Auto-encoder based representation.
- Dynamic-graph based Graph Convolutional Network(GCN).

Meta Descriptive Statistics

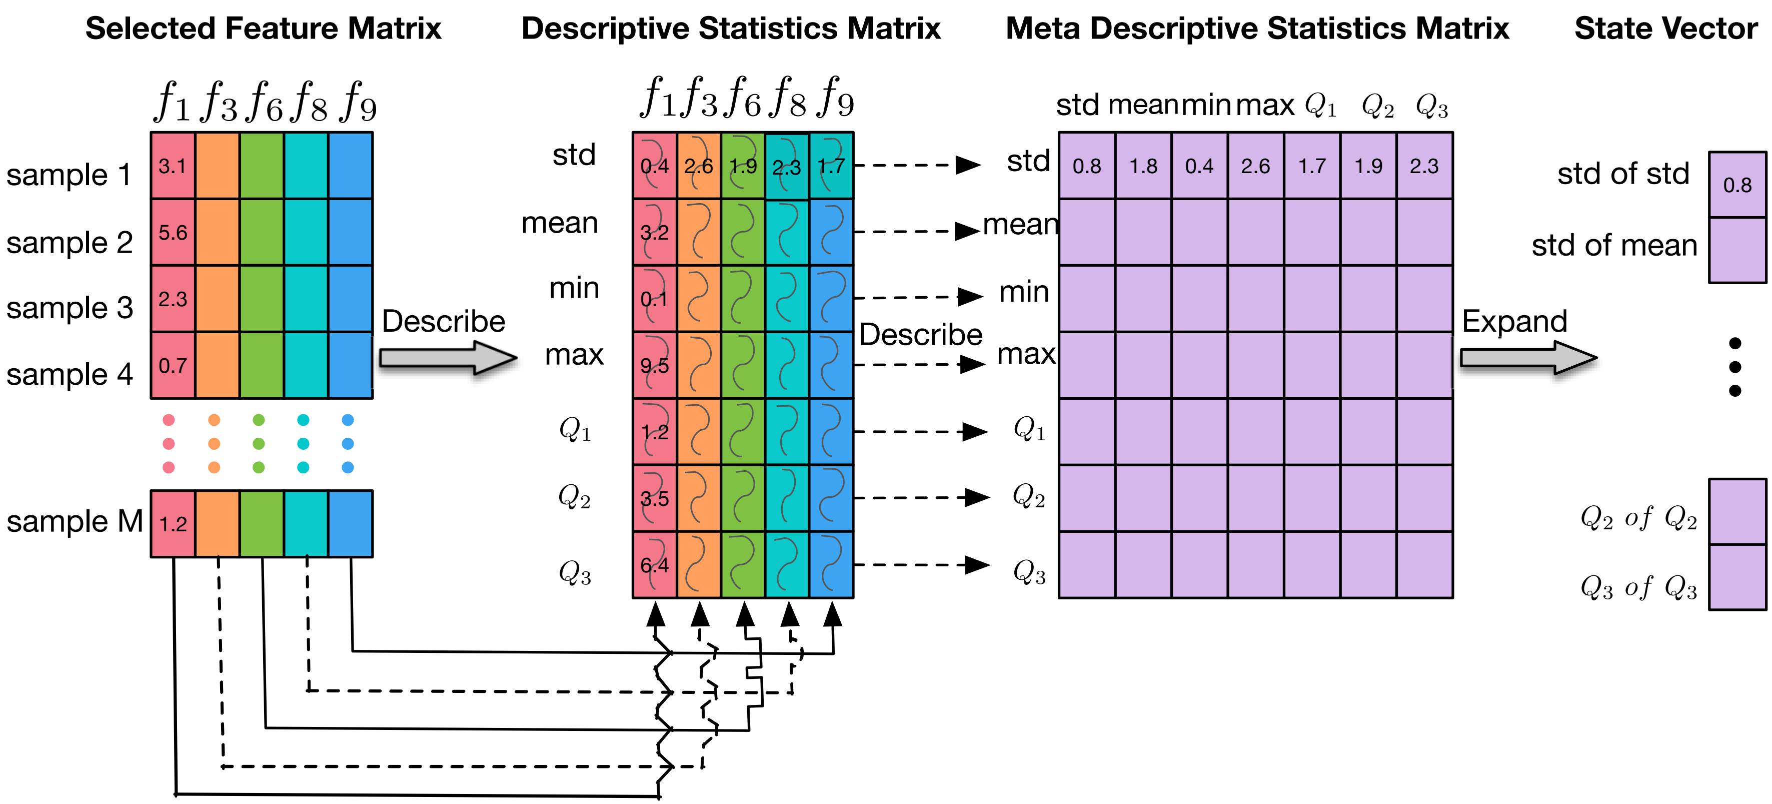

Step 1: Draw statistics column-wisely.
Step 2: Draw statistics row-wisely.
Step 3: Expand the statistics matrix.

Auto-Encoder Based Representation

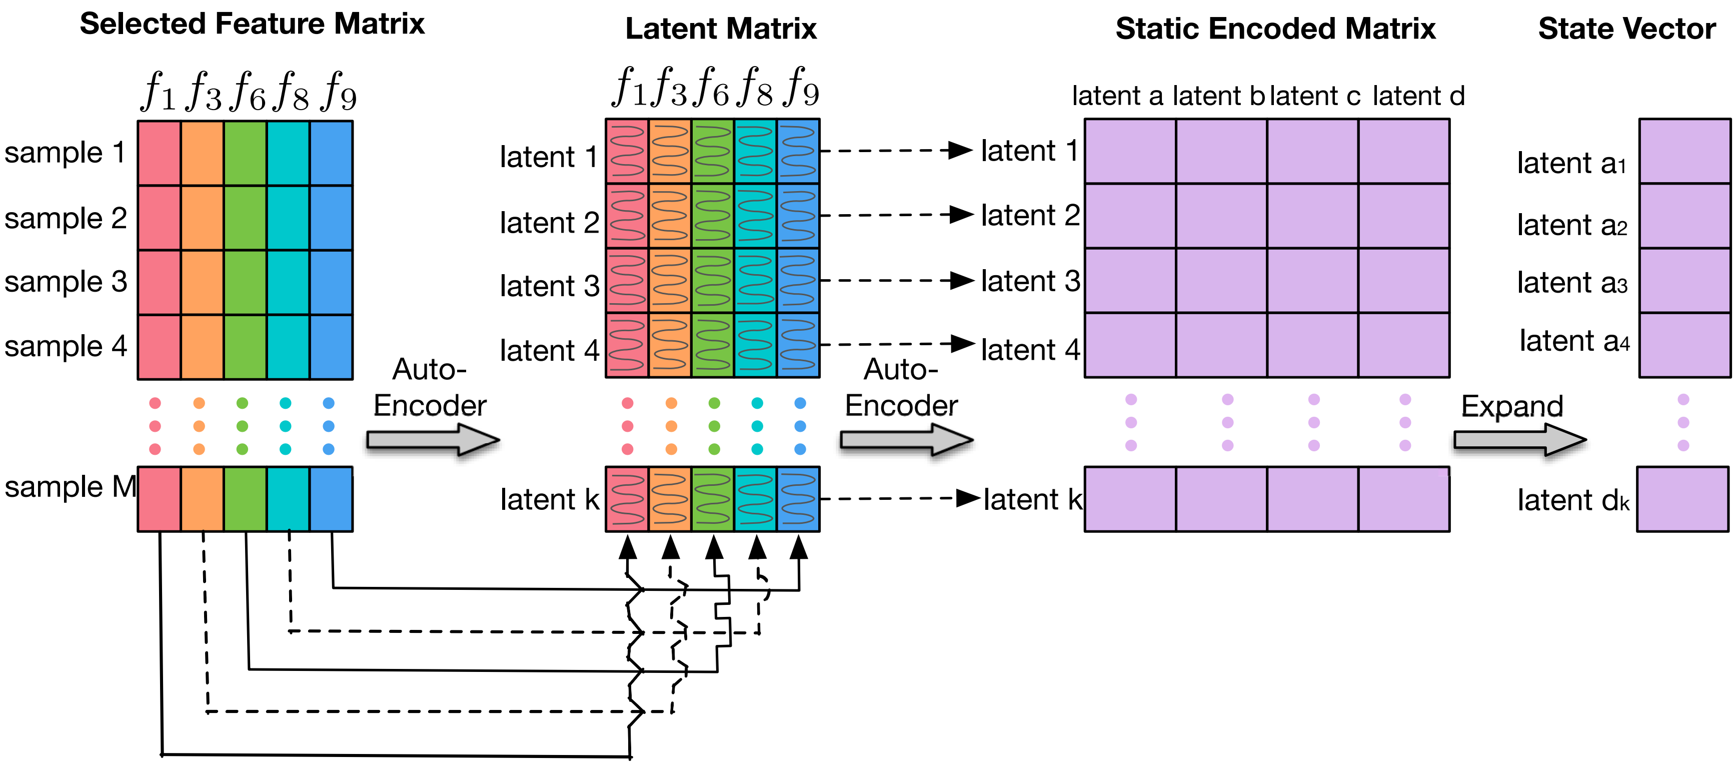

Step 1: Encode column-wisely.
Step 2: Encode row-wisely.
Step 3: Expand the encoded matrix.

Dynamic-Graph Based GCN

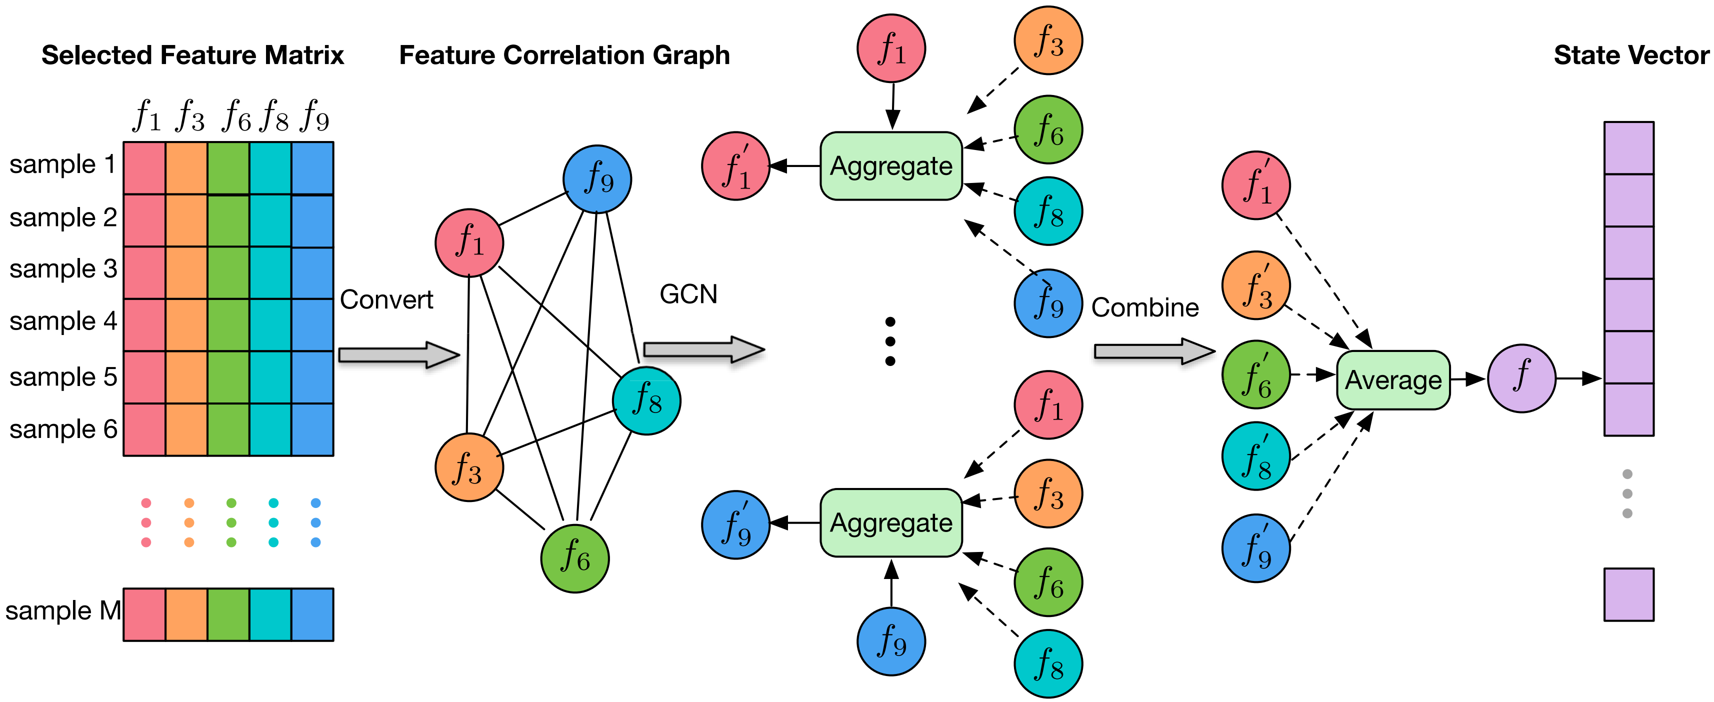

Step 1: Draw a fully-connected graph.
Step 2: Update each node's representation.
Step 3: Aggregate all nodes' representations.

Question 4 of 4

How can we improve the training efficiency of DQN in MARL?

GMM Based Sampling for Acceleration

- Improve quality of training data in Experience Replay.

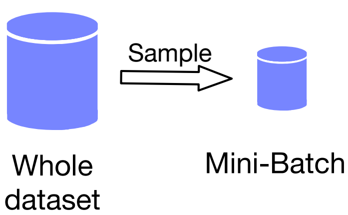

Conventional sampling strategy.

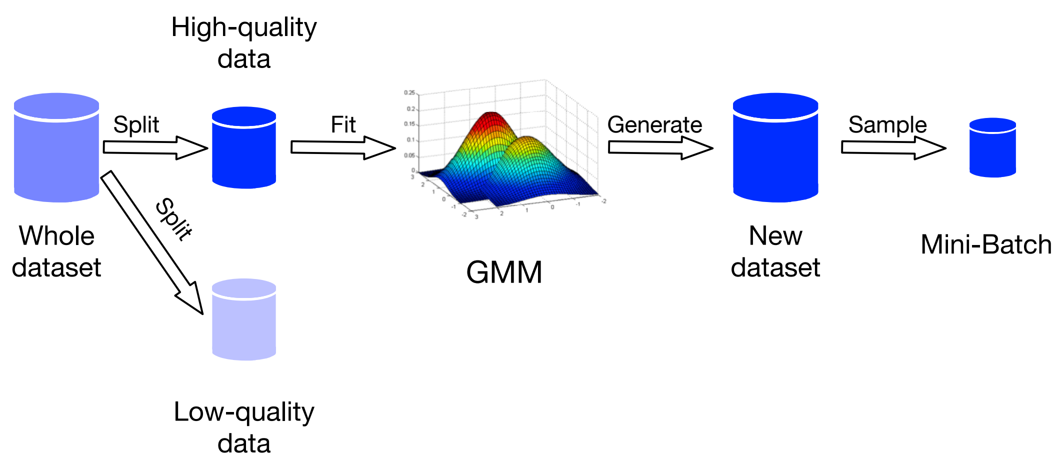

GMM based sampling strategy.

Summary

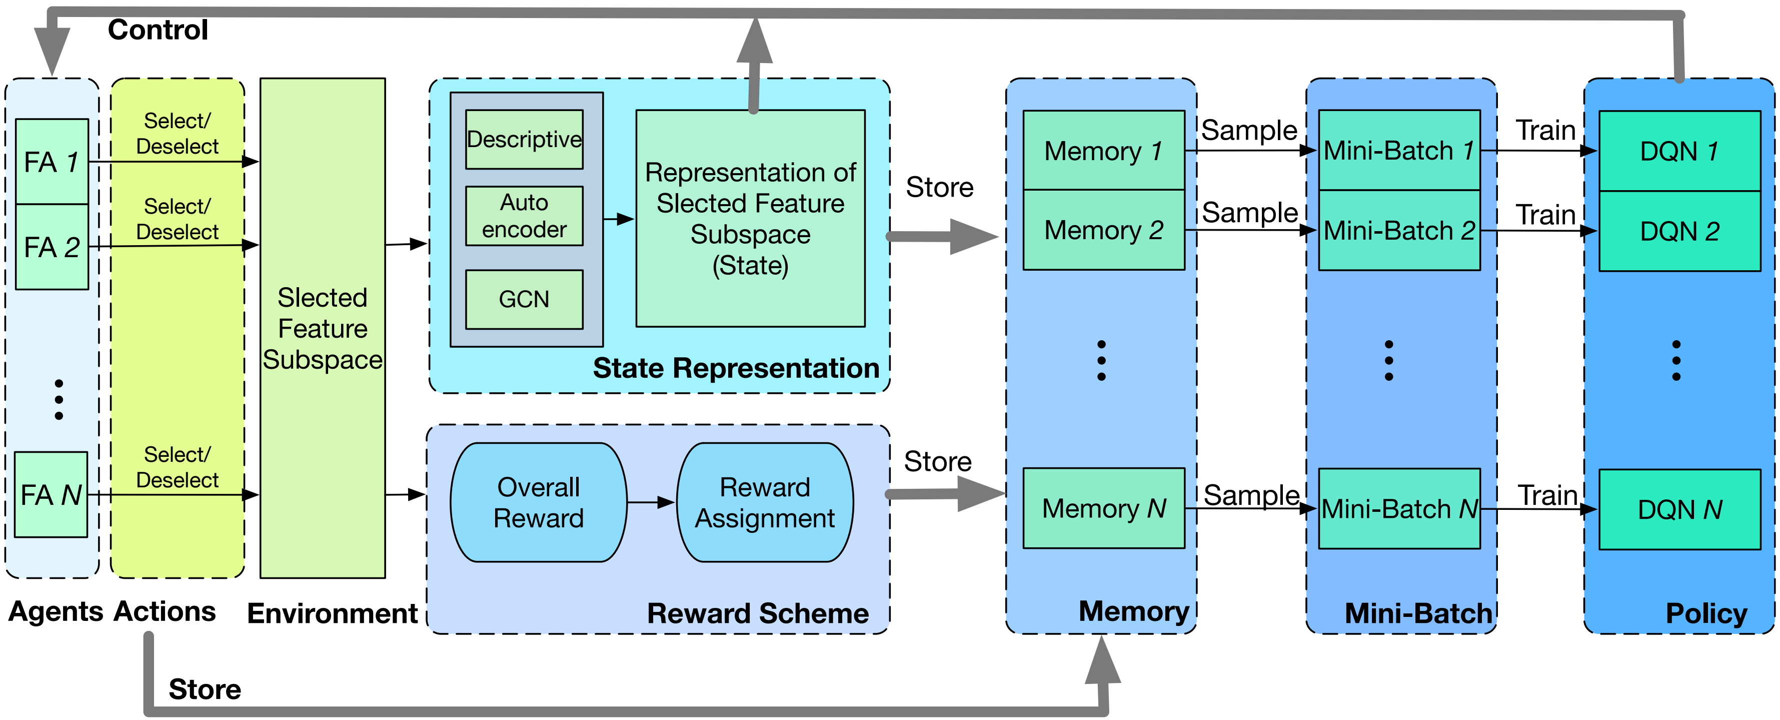
Control
FA 1
FA 2
FA N
Select/ Deselect
Select/ Deselect
Select/ Deselect
Slected Feature Subspace
Descriptive
Auto encoder
GCN
Representation of Slected Feature Subspace (State)
State Representation
Overall Reward
Reward Assignment
Reward Scheme
Store
Store
Memory 1
Memory 2
Memory N
Memory
Sample
Sample
Sample
Mini-Batch 1
Mini-Batch 2
Mini-Batch N
Mini-Batch
Train
Train
Train
DQN 1
DQN 2
DQN N
Policy
Agents
Actions
Environment
Store

Experimental Setup

- Experimental Data
  - The experiments are carried on a publicly available dataset with 15120 samples and 54 features.
  - https://www.kaggle.com/c/forest-cover-type-prediction/data.
- Predictive Task
  - The task is to classify the forest cover types into 7 classes.
- Experimental Questions
  - Can our study improve feature selection performance?
  - How do different reward quantification methods impact the performance of our method?
  - How do different state representation methods impact the performance of our method?
  - Can GMM sampling strategy improve exploration efficiency?

Performances over Different Classifiers and Feature Selectors

		Predictors				
		RF	LASSO	DT	SVM	XGBoost
Algorithms	K-Best	0.7943	0.8246	0.8125	0.8324	0.8076
	mRMR	0.8042	0.8124	0.8096	0.8175	0.8239
	LASSO	0.8426	**0.8513**	0.8241	0.8131	0.8434
	RFE	0.8213	0.8236	0.8453	0.8257	0.8348
	GFS	0.8423	0.8318	0.8350	0.8346	0.8302
	SARLFS	0.8321	0.8295	0.8401	0.8427	0.8450
	MARLFS	**0.8690**	0.8424	**0.8583**	**0.8542**	**0.8731**

Our method: **MARLFS**.
For the accuracies, the **higher**, the **better**.

Baselines: **K-Best** Selection, **mRMR**, **LASSO**, Recursive Feature Elimination (**RFE**), Genetic Feature Selection (**GFS**) and Single-Agent Reinforcement Learning Feature Selection (**SARLFS**).
Evaluation Metrics: overall accuracy

Performances over Different Reward Functions

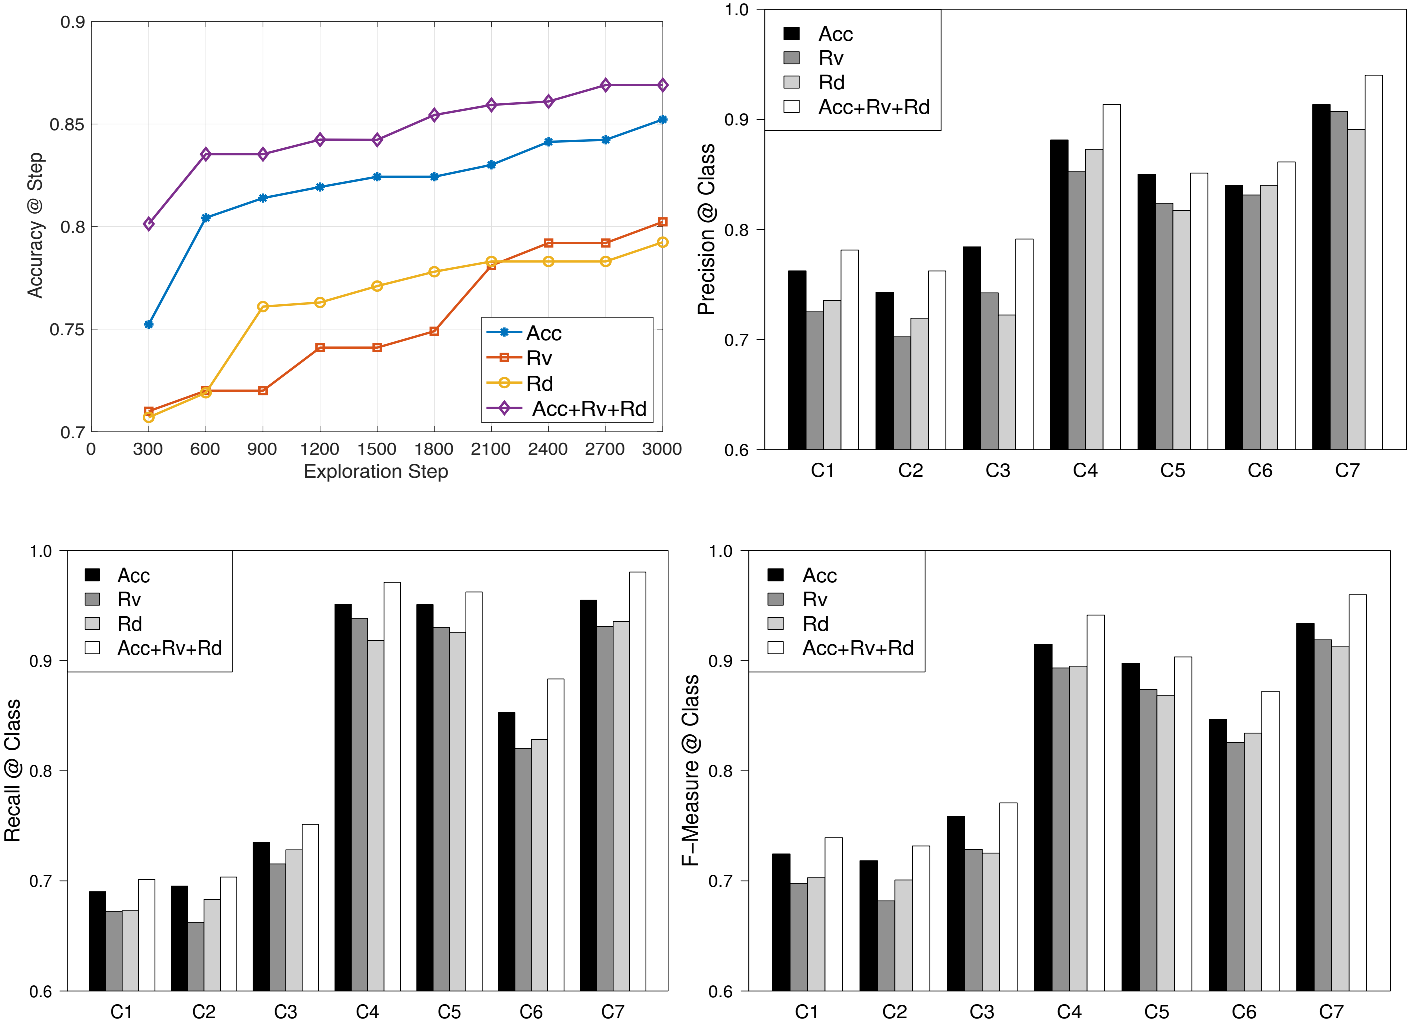

Variants: **ACC** (accuracy), **RV** (relevance), **RD** (redundancy), **ACC+RV+RD**.

Evaluation Metrics: overall accuracy, precision, recall and F-measure.

For accuracies and bars, the **higher**, the **better**.

Performances over Different State Representation Methods

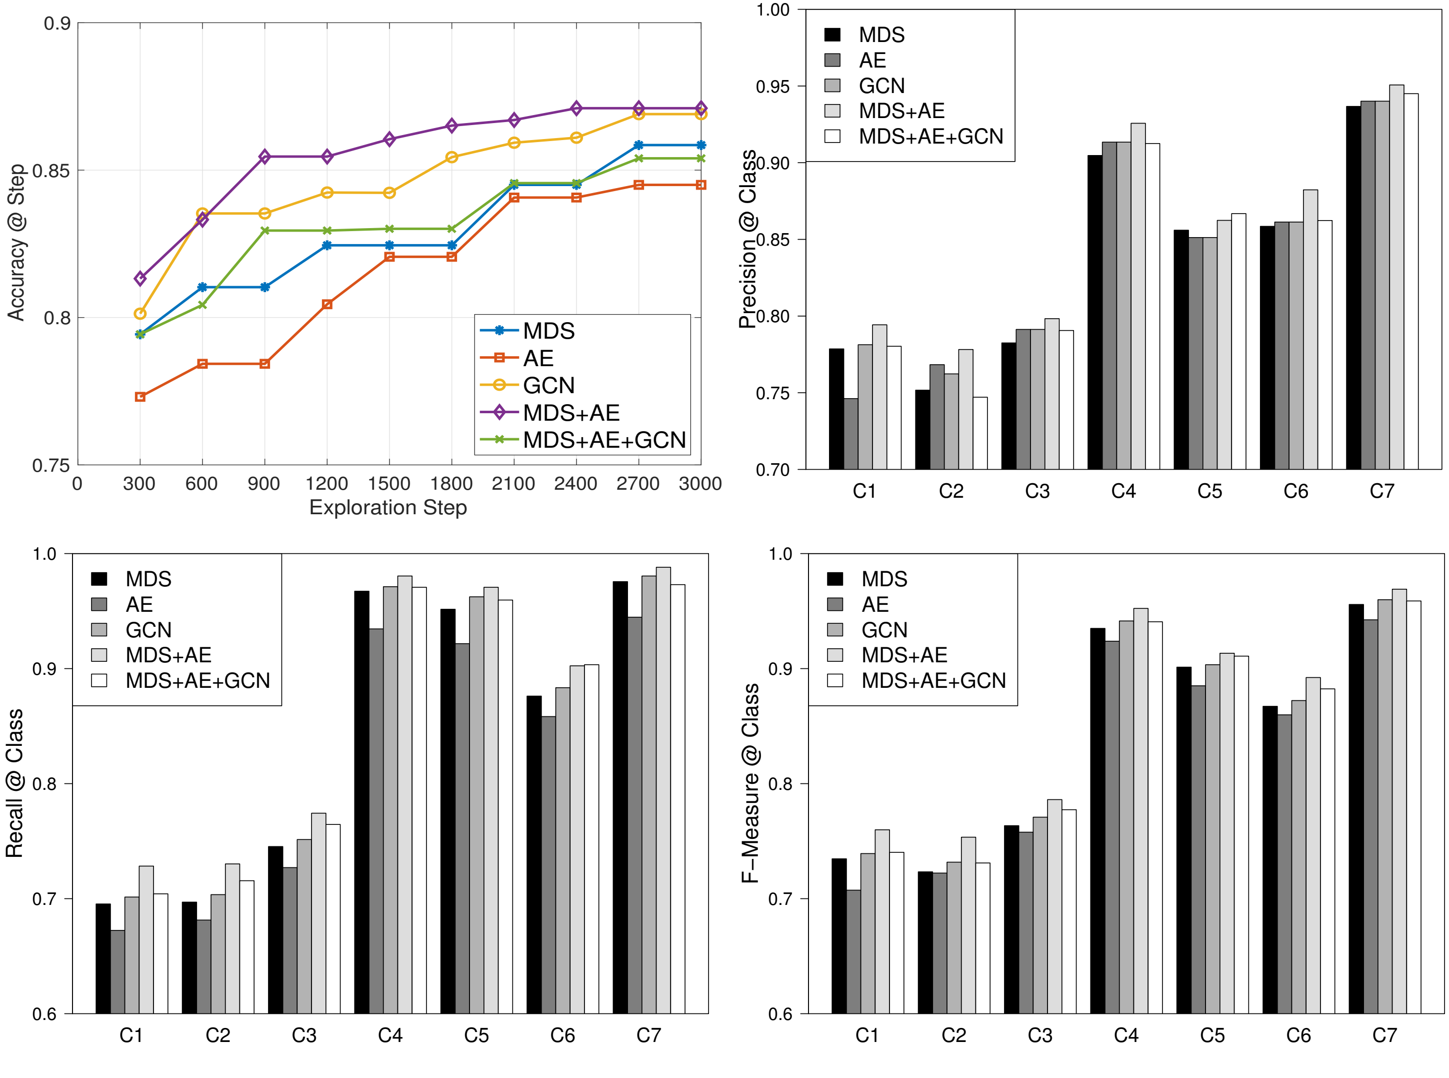

Variants:

MDS(meta descriptive statistics), **AE** (autoencoder based representation), **GCN** (GCN based representation)**,** **MDS+AE** (combination of MDS and AE), **MDS+AE+GCN** (combination of MDS, AE and GCN)

Evaluation Metrics:

overall accuracy, precision, recall and F-measure.

For accuracies and bars, the **higher**, the **better**.

Performance over Different GMM Sampling Strategies

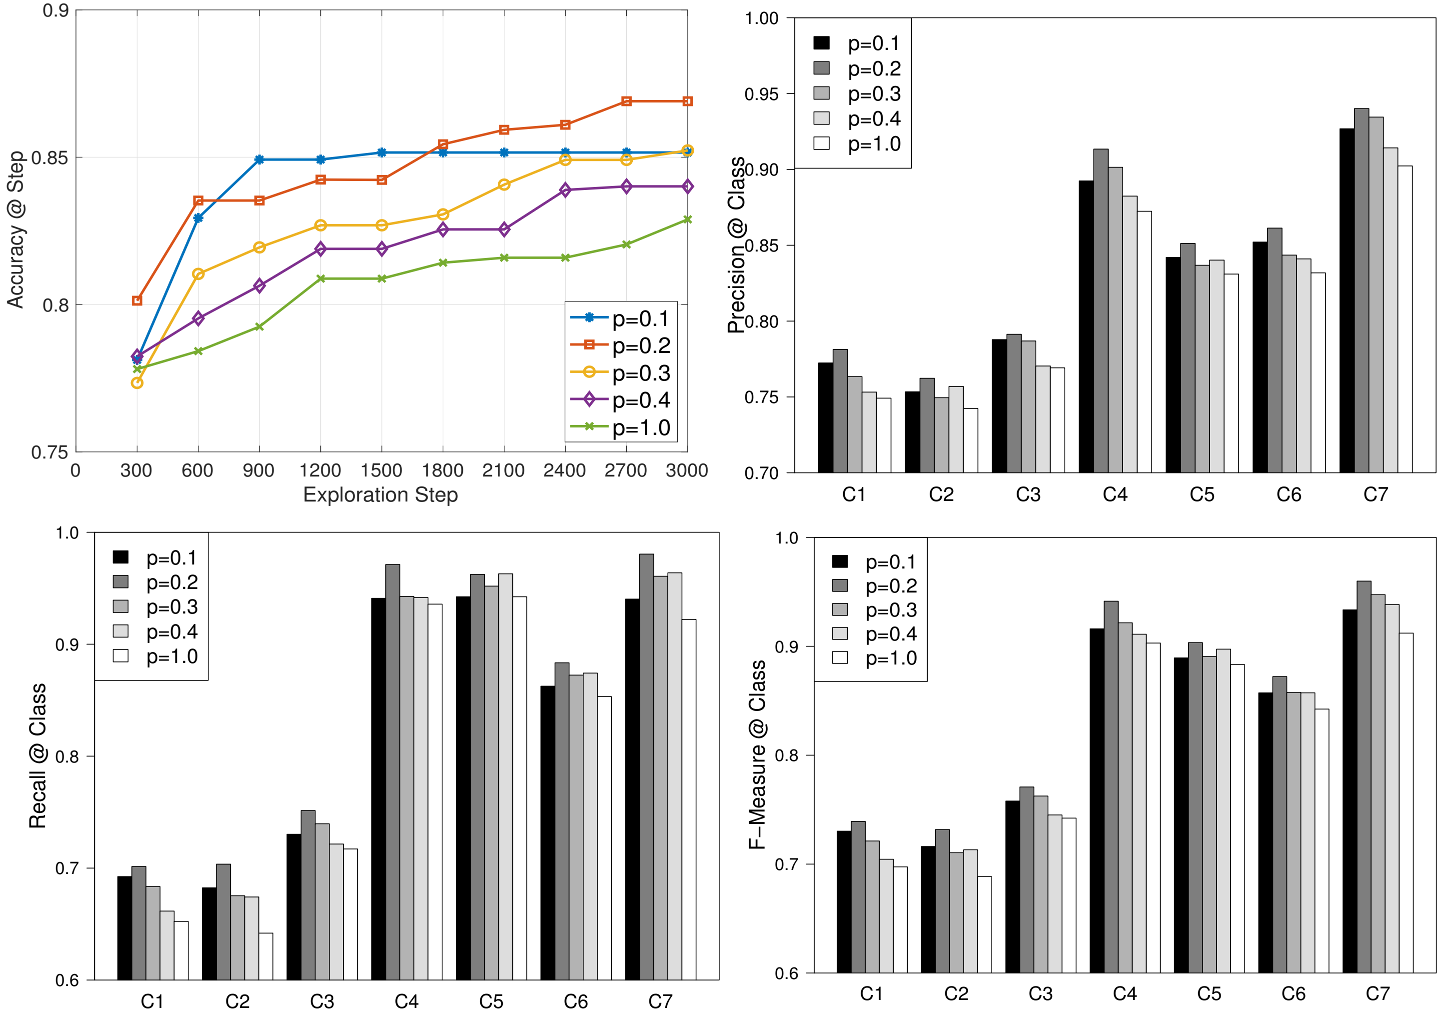

Variants:
proportion of high-quality varies from 0.1 to 1.0
Evaluation Metrics:
overall accuracy, precision, recall and F-measure.

For accuracies and bars, the **higher**, the **better**.

Conclusions

- Understandings
  - Feature selection is a space exploration process.
  - Feature selection can be improved by multi-agent reinforcement learning framework.
- Techniques
  - We propose three **state representation** methods.
  - We propose GMM-based **sampling** strategy.
  - We design **reward** quantification and assignment.

Thank you!