Informative Path Planning for Mobile Sensing with Reinforcement Learning

+

A Reinforcement Learning Framework for Efficient Informative Sensing

Vincent Mak

2024-02-29

Agenda

1) Informative Path Planning Problem
2) Reinforcement Learning
3) Proposed Solution
4) Improvement
5) Results
6) Conclusion

Informative Path Planning Problem

Goal: Find the best path to collect maximum amount of reward from starting vector to terminal vector with limited budgets

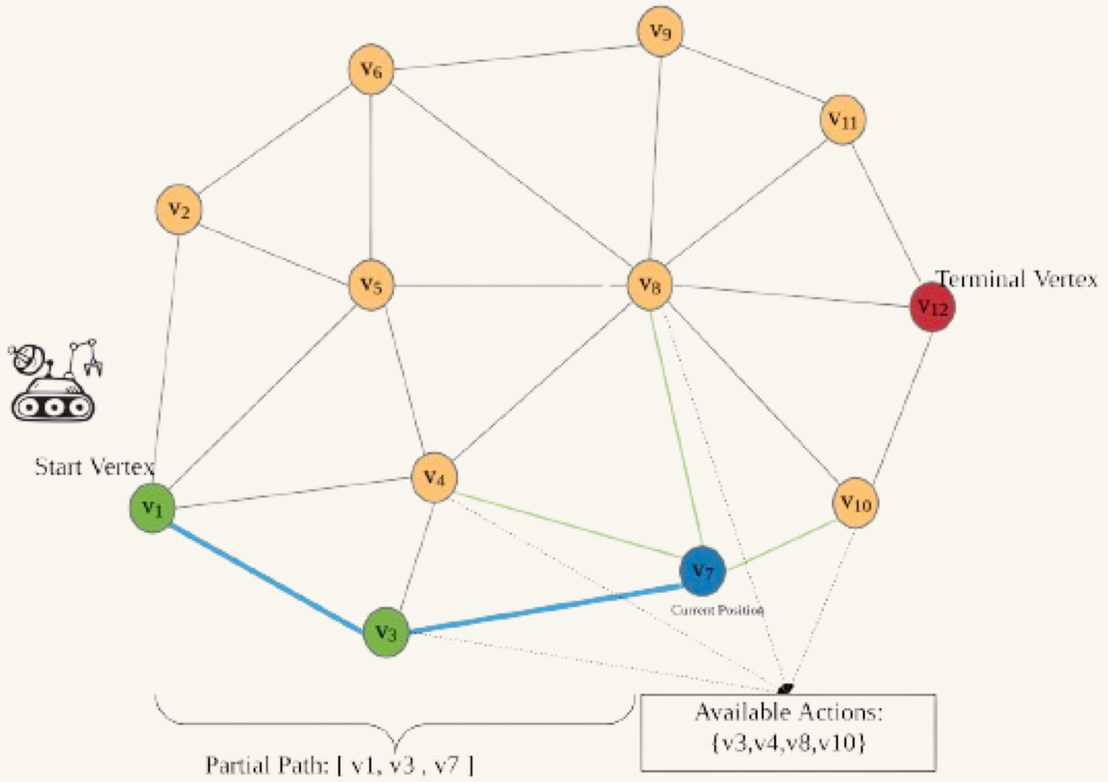

Reinforcement Learning

Existing approaches (eg. control policy and heuristic algorithms) suffer from optimization and complexity

Reinforcement learning

- Perform close to optimal with low execution time
- Utilize Markov Decision Process
- Stochastic policy and deterministic policy
- Value-based approach and policy-base approach

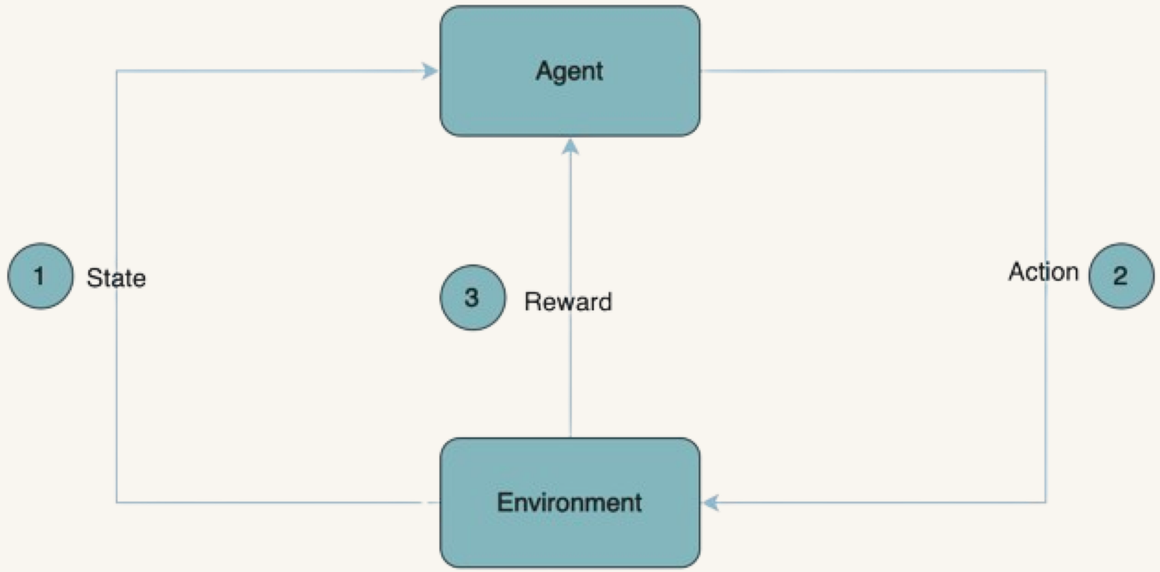

Proposed Solution

Recurrent Neural Network (RNN) as the Q-value approximator

Epsilon Greedy Algorithm to explore the solution space

- Exploration and Exploitation

Experience buffer is used for learning

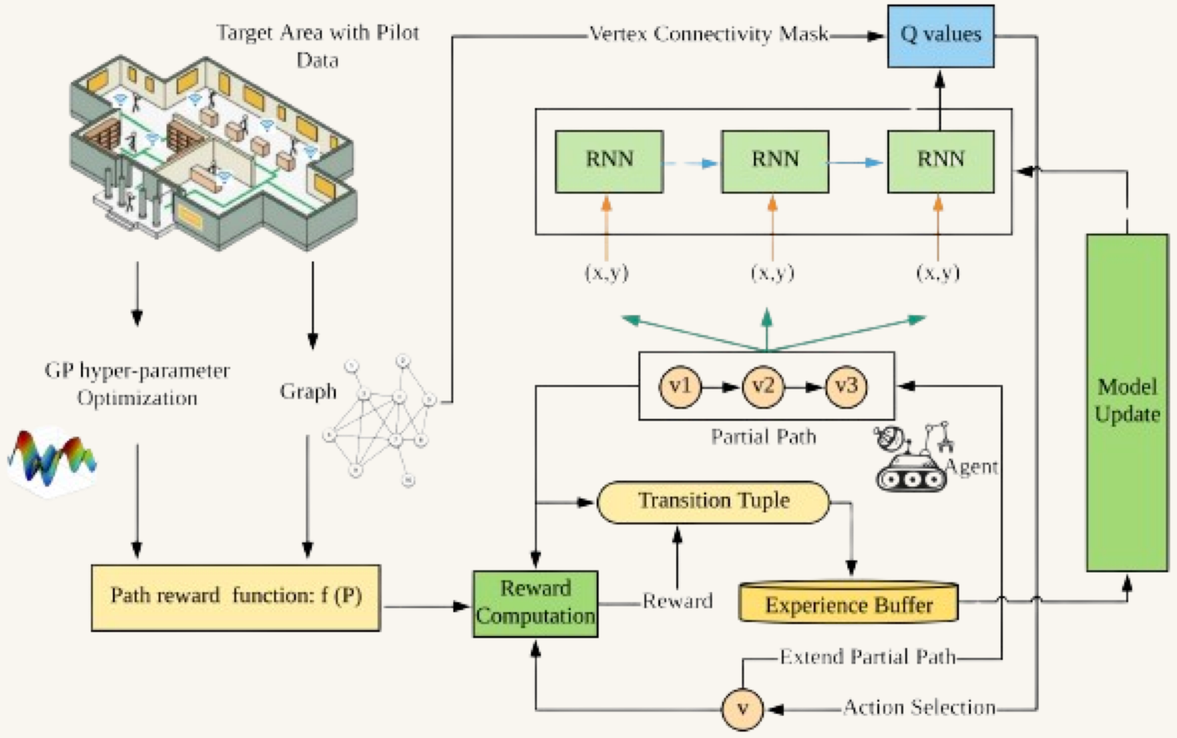

Improvement

Introduced state encoding

Implement 3 RL models and compare the performance

- Q-Learning
- Policy Gradient
- Actor-Critic

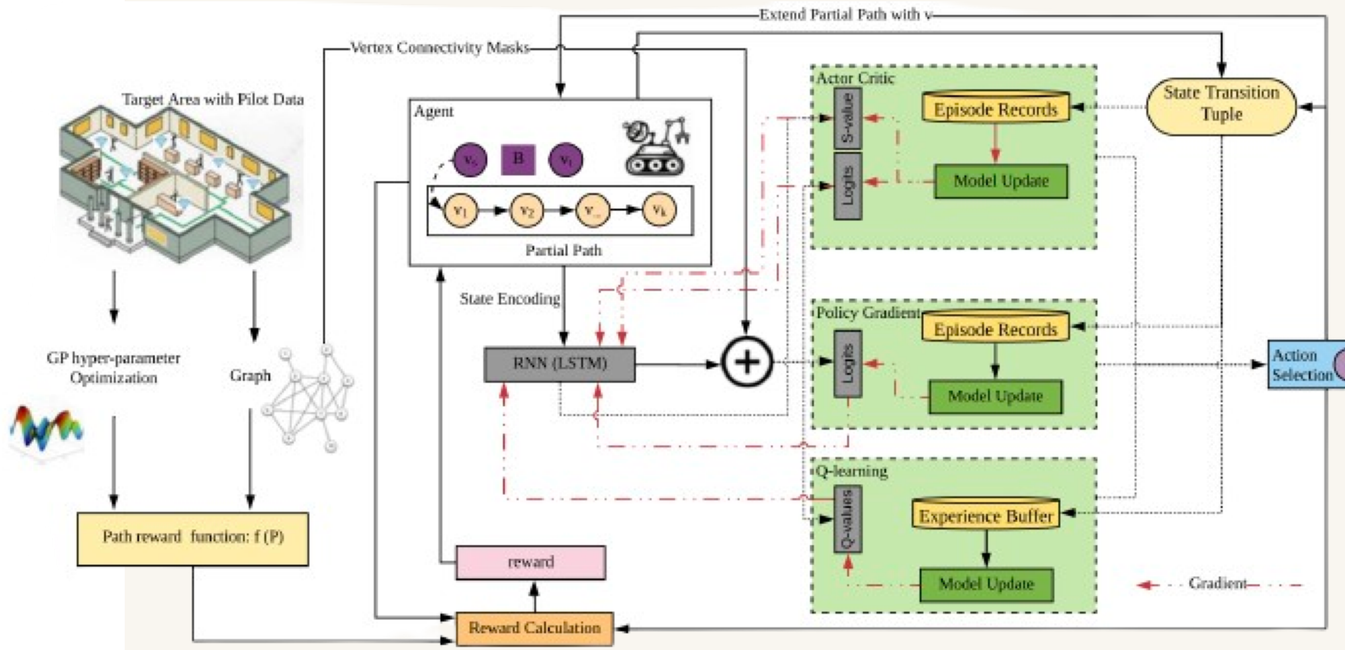

Research Result

A2C has a faster convergence speed in both Areas

Reinforce algorithm shows a similar trend to Q-learning in Area One

Reinforce algorithm has a competitive performance compared with A2C

Both reinforce algorithm and A2A outperform Q-learning

Reinforce algorithm = Policy Gradient

A2A = Actor-Critic

Area One

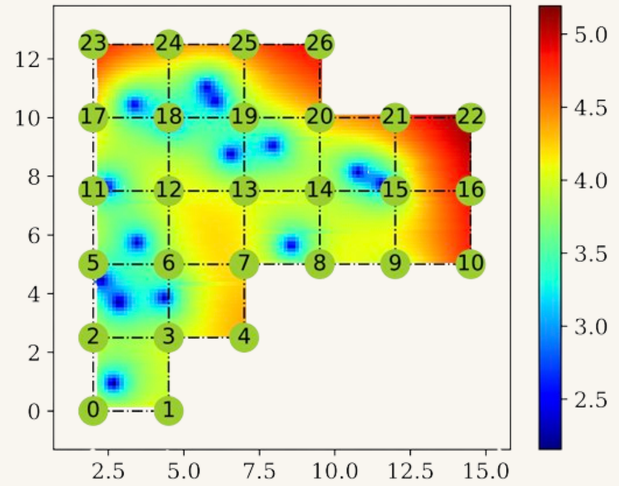

Area Two

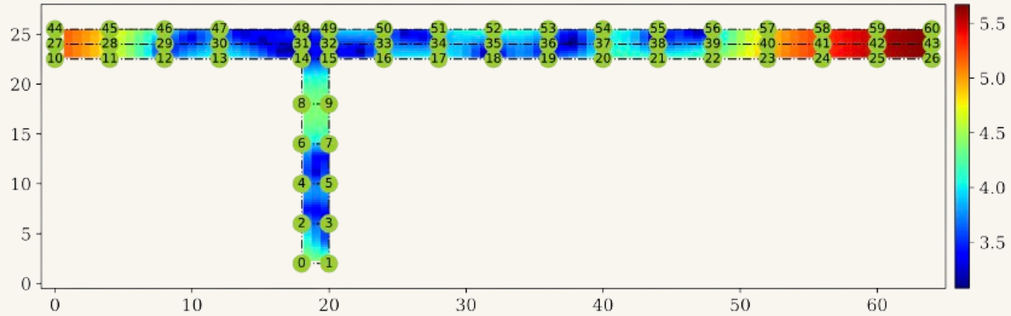

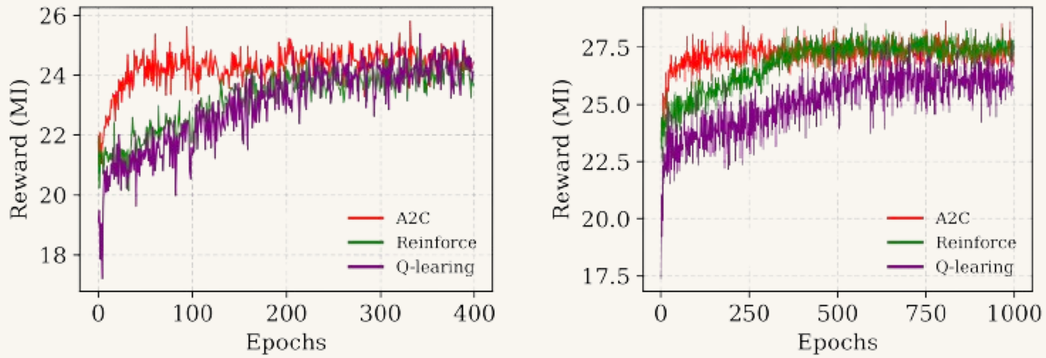

(a) Area One

(b) Area Two

Research Result

The rewards from the GA and RG consistently increase with higher budgets.

Rewards from the Greedy Algorithm remains unchanged with increasing budgets.

RL achieved the best performance in 15 cases out of 20 test cases

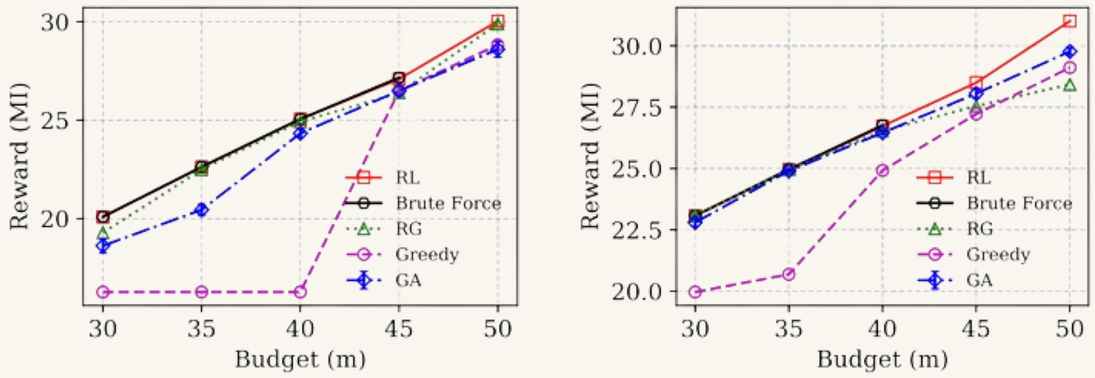

(a) tour (b) non-tour

Fig. 1

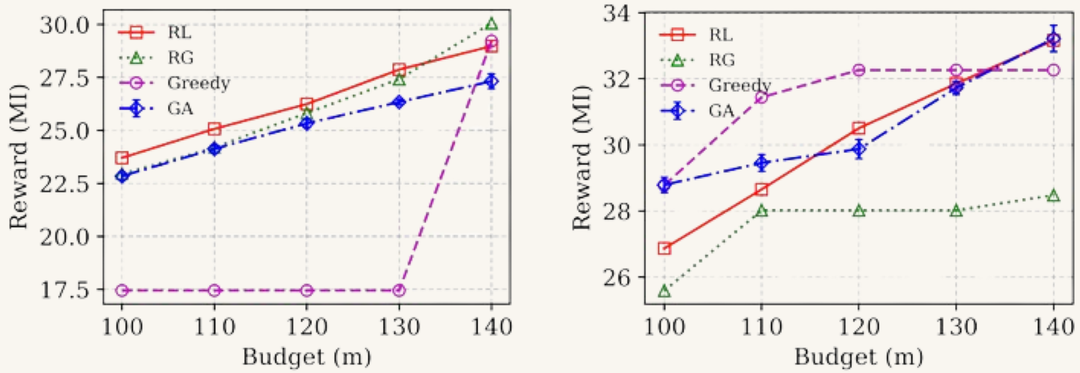

(a) tour (b) non-tour

Fig. 2

Tour: Agent is required to return to the starting vector

Non-Tour: the terminal location is different from the start location

Conclusion

Reinforcement learning can provide result outperforms other algorithms, and is much more efficient

A2A appears to be the best RL framework among all three RL approaches

Thank you