Ensemble Method for Reinforcement Learning Algorithms Based on Hierarchy

D. A. Kozlov[1]
djoade100@gmail.com

V. V. Myasnikov[1]
vmyas@geosamara.ru

Introduction

Reinforcement learning (RL) is a popular machine learning approach for solving sequential decision-making problems. One of the challenges of RL is that the performance of a single algorithm may be highly dependent on the choice of hyperparameters or the specific problem being solved. One way to address this challenge is to use an ensemble of algorithms that can work together to improve performance. In this poster, we present our study on improving RL performance using an ensemble of algorithms.

Background

Our study is based on the proposed method that combines the Robust Ensemble of Deep Q-Networks (REDQ) algorithm with the Soft Actor-Critic (SAC) algorithm. The ensemble output is controlled by the DQN control algorithm. We implemented this algorithm using the PyTorch RL framework and found that it performed better than, or at the worst case, at the level of the best algorithm in the ensemble.

Method

Our proposed method involves using an ensemble of algorithms that are specialized for different tasks. We explored the use of a larger number of algorithms in the ensemble, including the use of the same algorithm with different hyperparameters. This approach required fewer interactions with the environment than traditional methods for hyperparameter selection. We also considered using a hierarchical approach that would include multiple control algorithms, with one top-level algorithm responsible for task recognition and selecting the appropriate ensemble of algorithms.

Results

Our study showed that the implemented algorithm improved RL performance compared to using a single algorithm. In addition, the proposed method allowed for more efficient hyperparameter selection, reducing the number of interactions with the environment required. The hierarchical approach also shows promise for more complex environments.

Fig 1. Environments from the dm-control package. Cheetah (left), Humanoid (middle), Walker (right)

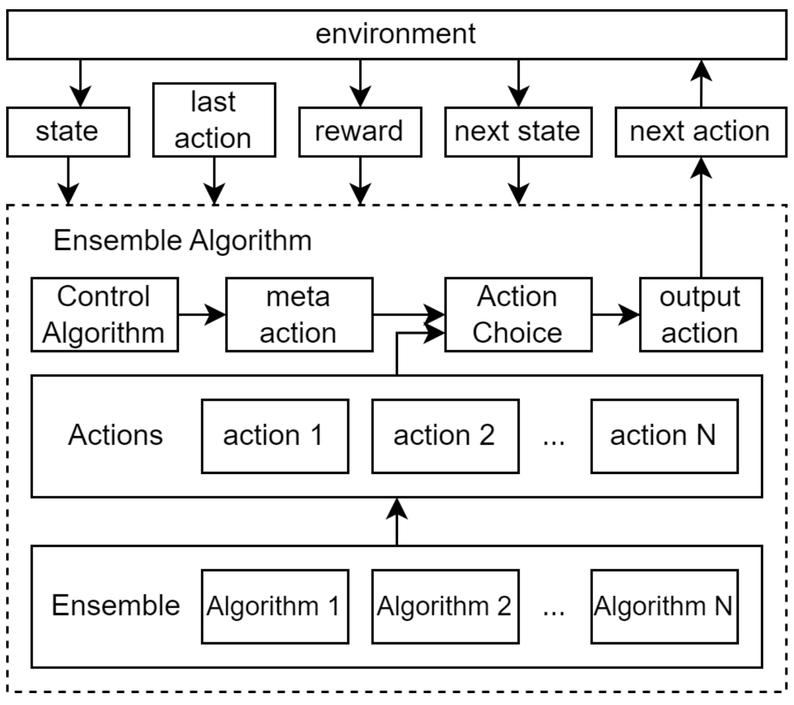

Fig 2. Proposed method

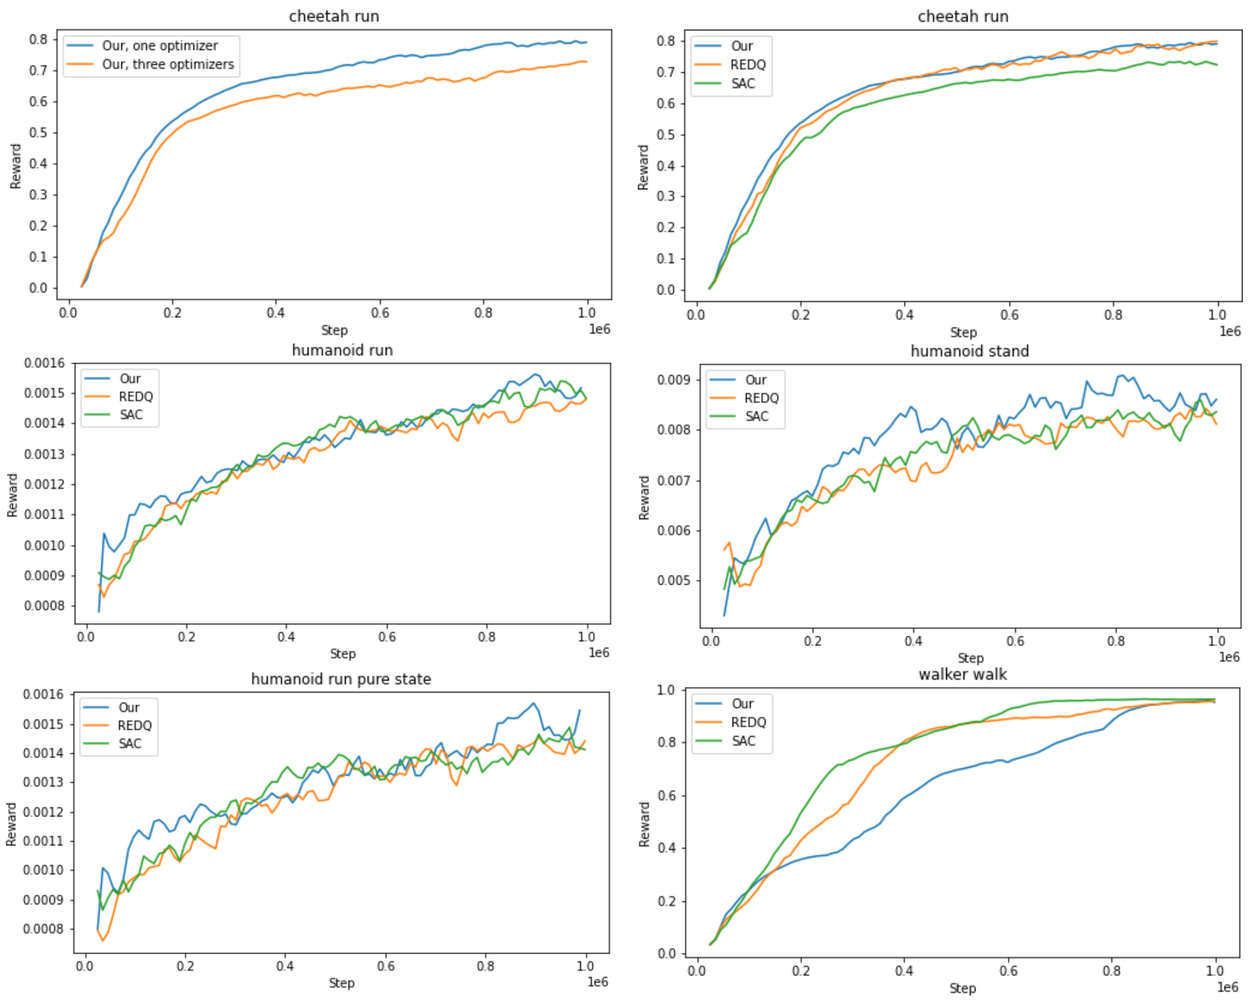

Fig 3. Comparison of the implementation of the proposed method with the SAC and REDQ algorithms

Discussion

Our study suggests that using an ensemble of algorithms can be an effective way to improve RL performance. However, further research is needed to explore the use of larger ensembles and more complex hierarchical approaches. In addition, the choice of environment is an important consideration, as real environments may be more suitable for this approach due to the presence of a large number of subtasks. Further work could focus on developing methods for clustering input data to identify subtasks and create specialized ensembles for each one.

Conclusion

Our study demonstrates the potential of using ensembles of algorithms for improving RL performance. By combining multiple specialized algorithms, we were able to achieve better results than using a single algorithm. Our proposed method also allowed for more efficient hyperparameter selection and shows promise for more complex environments.

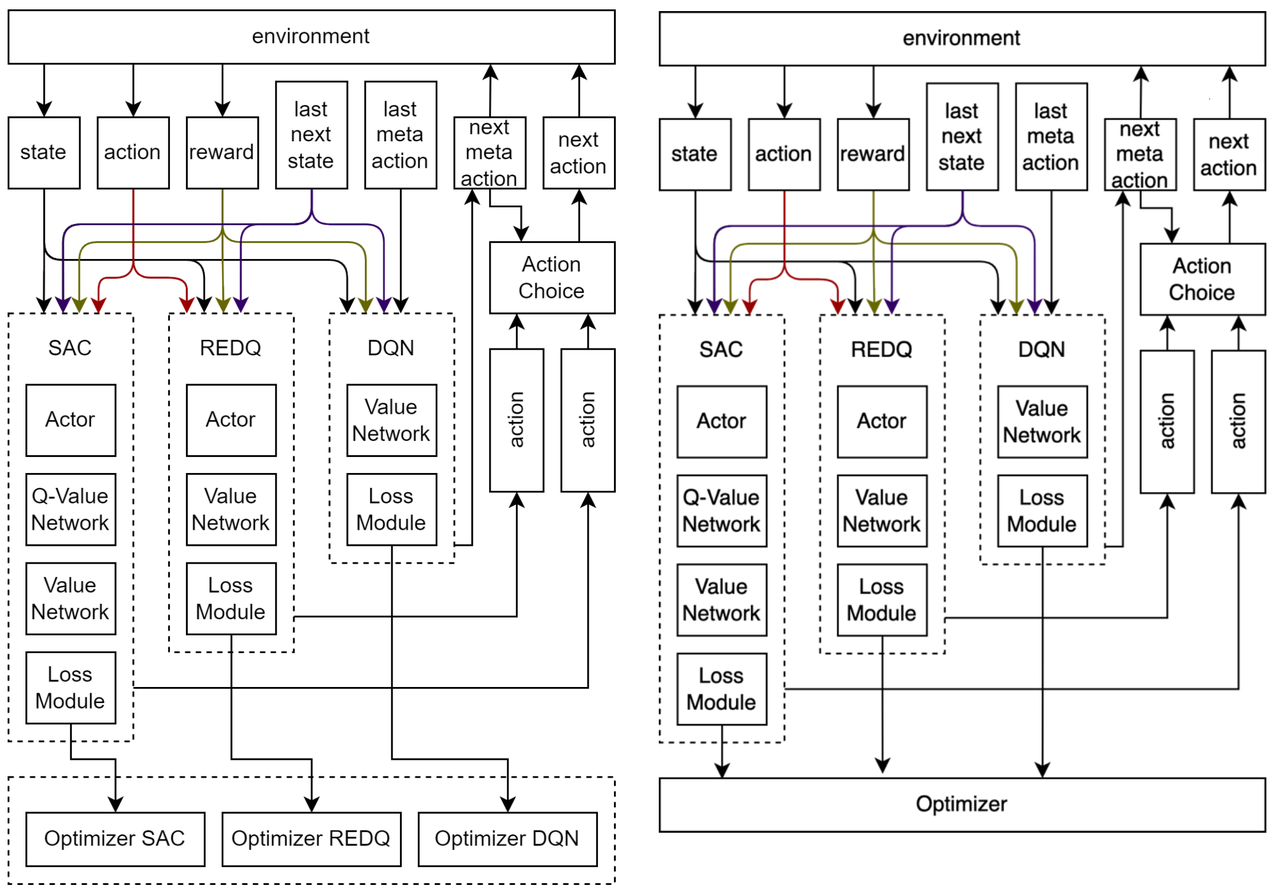

Fig 4. Implemented algorithm. Version with individually optimized neural network parameters on the left. Version with jointly optimized neural network parameters on the right

Acknowledgement

The work was supported by the Russian Science Foundation grant No. 21-11-00321, https://rscf.ru/en/project/21-11-00321.

[1]SAMARA NATIONAL RESEARCH UNIVERSITY
34, Moskovskoye shosse, Samara, 443086, Russia
www.ssau.ru

SAMARA UNIVERSITY