Open source project: Mava

Asbjorn Lorenzen

December 4, 2024

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

Mava's Architecture and Usefulness

Key Components:

- Centralized Training, Decentralized Execution (CTDE): Efficient training with global information while enabling independent agent execution.
- Modularity: Plug-and-play support for various MARL algorithms.
- Scalability: Optimized for large-scale systems with distributed training support.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00

Why Mava?

- Easy development of MARL algorithms.
- Built for research

Counterfactual Multi-Agent Policy Gradients (COMA)

What is COMA?

- ► A MARL algorithm designed for cooperative tasks.
- Uses a centralized critic to compute counterfactual baselines.
- Addresses the credit assignment problem by isolating each agent's impact.

Why Extend Mava with COMA?

- COMA is one of the most cited MARL algorithms
- Mava is built for rapid prototyping and research. Implementing a new algorithm puts this to the test

Project Implementation Plan

Steps Using Mava:

- Centralized Critic: Extend Mava's critic to compute counterfactual baselines.
- Training Integration: Add COMA's training loop to Mava's modular architecture.
- Reward and Baseline: Implement shared rewards and counterfactual calculations.
- Performance Evaluation: Test COMA with Mava in collaborative tasks. Compare with MADDPG and QMIX.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00

Conclusion

Conclusion:

- Mava is a robust platform for MARL research.
- Extending Mava with COMA is useful for exploring cooperative MARL scenarios.

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ □ のへぐ