NRI: FND: Connected and Continuous Multi-Policy Decision Making

Edwin Olson ebolson@umich.edu

Challenge

How can teams of robots coordinate over long time scales with poor communication?



Solution

- Plan in policy space, rather than in action space.
- Each timestep, use Monte-Carlo simulations to dynamically select which policy to use from a small set of policies.
- Fully decentralized; will still make progress towards goal even with zero communications.
- Agents compute their own choice, based on their local model of the other agents
- Sharing of information, not plans, leads to implicit coordination.



Coverage

Versus picking a single policy (shown in different colors), dynamically switching between policies results in faster task completion. This holds true across different communication levels, and across different environments, including brand-new environments (with the exception of the bottom right plot).

University of Michigan

Adversarial Search

Scientific Impact

This method is promising for its ease-of-implementation, high performance, and good generalization across environments.

Broader Impact

Applications: Robot teams for search and rescue; agriculture; autonomous driving.

Versions of these algorithms are gaining commercial interest.



Our approach is computationally efficient, compared to a tree search baseline.

Award ID#: 1830615

The PI has a financial interest in a company that may have rights to foreground or background technology described here.





