how do we improve teams of robots cooperate on long-horizon problems?

Tag!

N robots versus 1 human

Problem Properties

- Partial observability
- Stochastic actions and communications
- Adversarial human ➔ long planning horizons

Solution Properties

- Decentralized
- Degrade gracefully with comms

Idea: dynamically switch between 2+ strategies (MPDM).

Spread out and look for the human

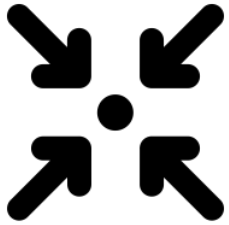

Converge upon a suspected location

know when to use which strategy by evaluating with Monte Carlo roll-outs.

It works!

(but there's lots of interesting wrinkles)

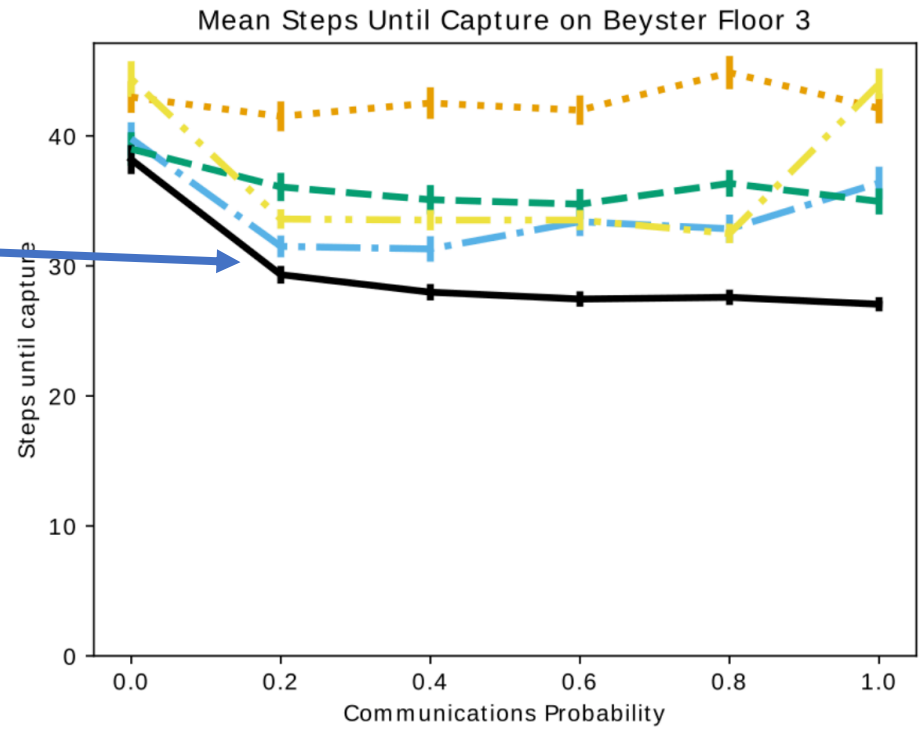

Come talk with me at Poster Session 2, #46