

2018 Workshop on Autonomy for Future NASA Science Missions

October 10-11, 2018



Multi-Arm Bandits for Automated Decision Making

Steve McGuire, Nisar Ahmed: CU Boulder

P.M. Furlong, Terry Fong: NASA Ames Research Center

Research Vision: Autonomy Failure Recovery

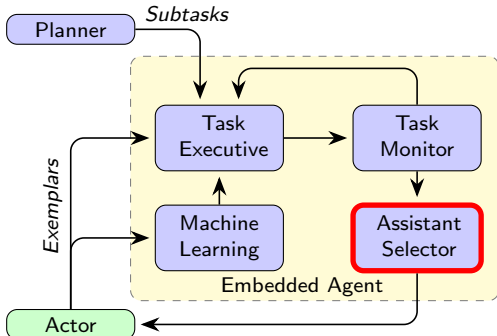
Scenario: *Autonomy isn't perfect and will eventually fail. Then what?*

Problem: How does a robot in need of help decide who to ask?



Mission: Joint Exploration

Model of an Ideal Learning Robot



Key Research Challenges:

- ▶ Individual Human Modeling
- ▶ Signals of Opportunity
- ▶ Long-Term Operations

Plan of Action


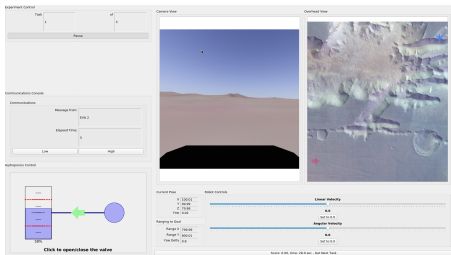
Method: Reinforcement Learning via Contextual Multi-Arm Bandits

Implementation Levels:

- ▶ Assessment in simulation: *Failure is Not an Option: Policy Learning for Adaptive Recovery in Space Operations*, IEEE RA-L, Jan 2018
Signals of opportunity are useful!
- ▶ Assessment in human trials: *Everybody Needs Somebody Sometimes: Validation of Adaptive Recovery in Robotic Space Operations*, IEEE RA-L, In review
Works with real humans!
- ▶ Practical Problems: **In progress**
 - ▶ Initialization
 - ▶ Dynamic Human Models

Future Impact:

- ▶ More agile missions
- ▶ Collaboration with less structure
- ▶ Enabling for multi-robot ops
 - ▶ Better use of human assets
 - ▶ Leverage onboard autonomy



Human-in-the-loop experiments showing improvement when contextual information's impact is learned (red, lower is better)