# Coordinating Human-Robot Teams in Uncertain Environments

**Christopher Amato**, *Northeastern University* & **Laurel D. Riek**, *UC San Diego*

Northeastern University

UC San Diego JACOBS SCHOOL OF ENGINEERING Computer Science and Engineering

## Project overview

**Goal: create and solve realistic models for coordinating teams of humans and robots in uncertain environments**

1. Re-conceptualize multi-human teamwork that include dynamic, stochastic environments
   - See **Intention Modelling for Teaming under Uncertainty** [4][5]
2. Develop realistic (POMDP) models of human-robot teamwork with uncertainty and partial observability
   - In progress

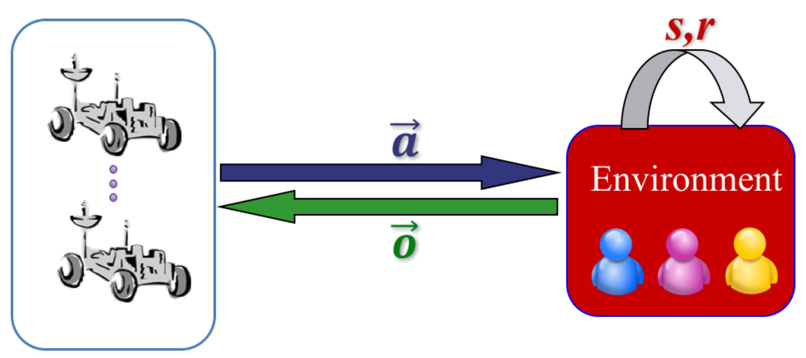


3. Create scalable techniques for planning and learning in these models
   - See **Bayesian Reinforcement Learning (RL) for POMDPs** [1] and **Hierarchical Deep Multi-Agent Reinforcement Learning (MARL)** [2][3]
4. Test in simulation and emergency department (ED) settings
   - See **Situating Robots in the Emergency Department** [6][7]

## Bayesian Reinforcement Learning (RL) for POMDPs

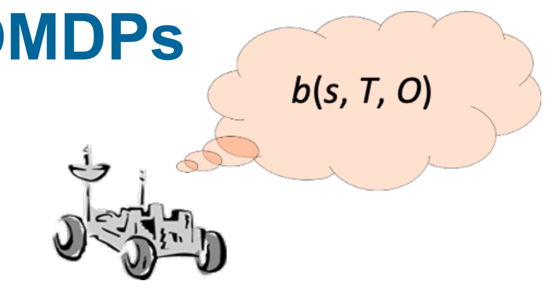


### Developed scalable Bayesian RL methods for POMDPs

- Bayesian RL can optimally balance exploration and exploitation
- Ideal for online learning—optimally sample efficient!
- Can be computationally challenging, but developed factored, sample-based methods
- These methods outperform previous methods, allowing learning in large POMDPs

**POMDPs**
- S, a set of states
- A, a set of actions
- T, the state transition model: $\Pr(s'|s,a)$
- R, the reward model: $R(s,a)$
- O, a set of observations
- Z, the observation model: $\Pr(o|s',a)$

### Bayesian RL for POMDPs (e.g., [Ross et al. JMLR 11])

- Explicitly consider uncertainty over possible trans. and obs. models
- Can start with prior over models and update based on observations
- Can now have belief over state *and models*

### Scalable solution methods

- Developed method for learning factored model and solutions [1]
- Sampling method for scalable particle filtering and updating [1]

### Results

- Our methods can learn quickly in all situations
- Drastically outperform previous methods
- Can learn efficiently by learning factored model and solution together

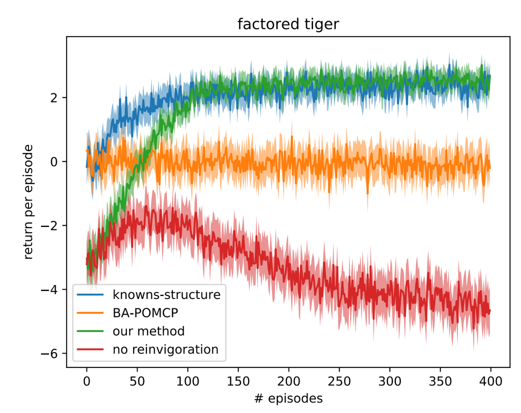


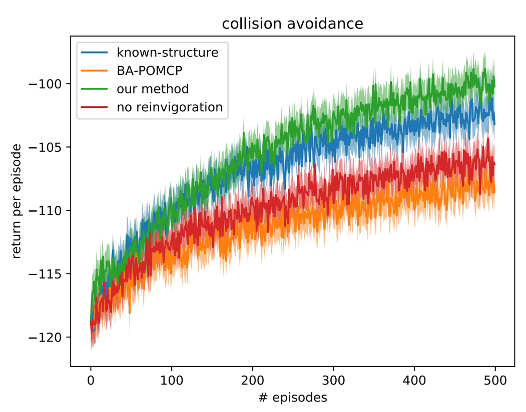


## Hierarchical Deep Multi-Agent Reinforcement Learning (MARL)

### Developed deep MARL methods that can learn in asynchronous, hierarchical settings

- Robots will execute at different task levels (e.g., high-level *macro-actions* and low-level controls)
- This is result in asynchronous execution for the agents
- Current deep MARL methods cannot solve this problem

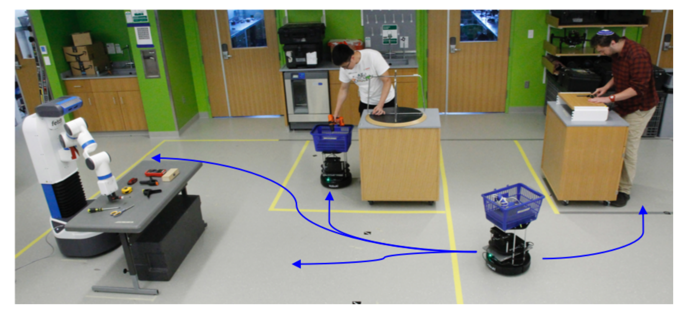

### Scalable solution methods

- Developed methods for multi-agent versions of POMDPs
  - *Decentralized learning for decentralized execution* [2]
    - Online learning without communication
  - *Centralized learning for centralized execution* [2]
    - Online learning with full communication
  - *Centralized learning for decentralized execution* [3]
    - Communication during learning, but not execution

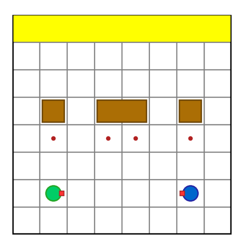

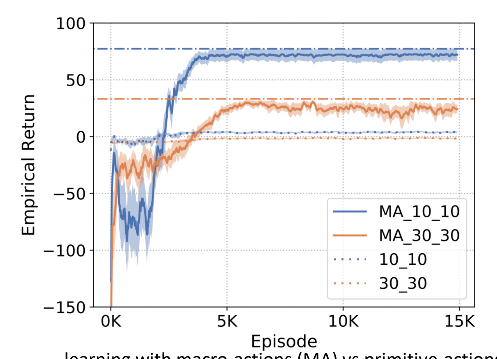


learning with macro-actions (MA) vs primitive-actions

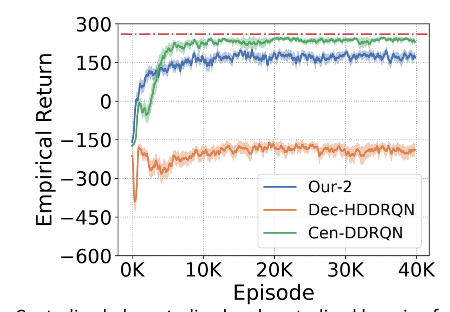


Centralized, decentralized and centralized learning for decentralized execution (Our-2)

### Results

- Experiments in simulation and hardware
- All methods can learn much faster and converge to higher solutions than using primitive (non-hierarchical) actions
- Centralized learning for decentralized execution methods can approach fully centralized solutions
- These methods require much more data than the Bayesian ones

## Intention Modelling for Teaming under Uncertainty

### Created new deep learning methods for non-visual activity modelling [4]

- Can detect both fine and gross motor movements, is immune to occlusion and avoids privacy concerns of visual sensing, augments linear and angular velocity w/ muscle activity data from wearable (Myo)
- Our methods outperformed the state-of-the-art classifiers by 28%, sEMG+Inertial yielded significantly higher classification accuracy than inertial alone
- Wearables are well-suited to activity recognition in uncertain environments

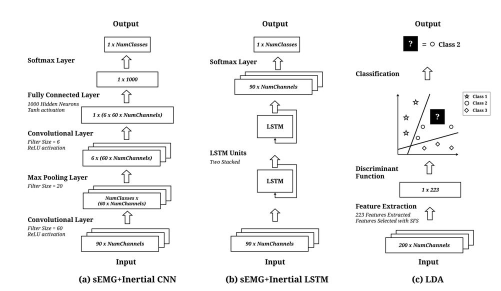


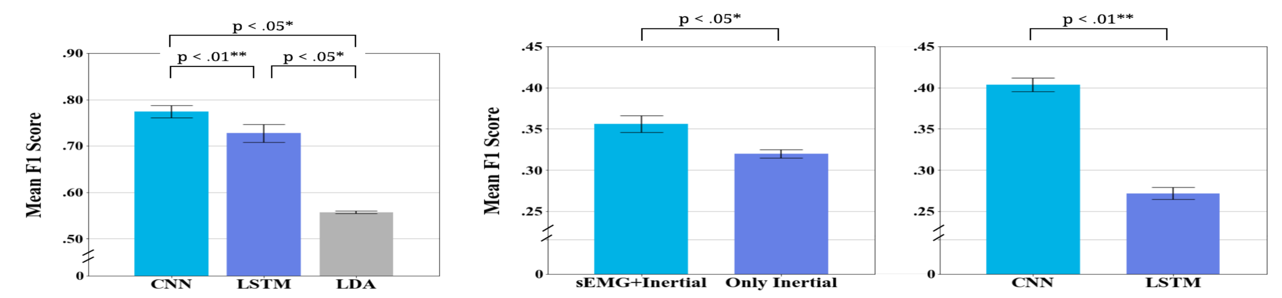


### Designed new approaches for multimodal contextualized activity recognition [5]

- Created new multimodal dataset of gross and fine motor tasks (EMG/Inertial/MoCap), compared multiple activity recognition approaches for recognition suitability, employed early fusion
- Results suggest complementary strengths of each sensor type – task type should be taken into account when engaging in sensor selection

Mean F1 scores obtained for each data modality on each dataset using different classifiers. Across the datasets and sensors, we averaged the F1 scores from every trial. A higher F1 score is better.

| | SVM | | | LDA | | | KNN | | |
|---|---|---|---|---|---|---|---|---|---|
| | Vicon | Myo | Vicon+Myo | Vicon | Myo | Vicon+Myo | Vicon | Myo | Vicon+Myo |
| **Automotive** (Gross motion) | **.79** | .42 | .43 | **.76** | .48 | .49 | **.88** | .58 | .59 |
| **Block** (Fine-grained motion) | .09 | **.37** | .36 | .23 | **.39** | .36 | .32 | **.43** | .43 |

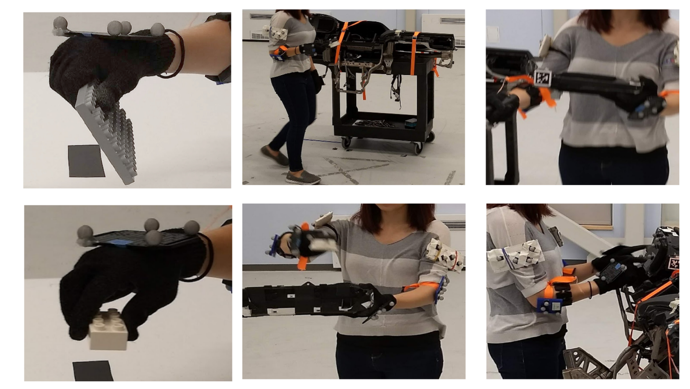

## Situating Robots in the Emergency Department

### Characterized Workflow and Task Representation for Situating Robots in the ED [5]; Designed Acuity-Aware Social Navigation Algorithms [6]

- The ED is an uncertain environment in which mistakes can be deadly and providers are over burdened.
- Well-designed & contextualized robots could relieve providers of non-value added tasks and enable them to spend more time on patient care. e.g., delivery robots.
- We used domain knowledge to characterize staff workflow and patient experience, identify key considerations for robots in the ED, inc.: safety, physical and behavioural attributes, usability, and training
- We introduced a task representation [5] and new acuity-aware social navigation algorithm [6] which incorporates both patient criticality and staff workflow.

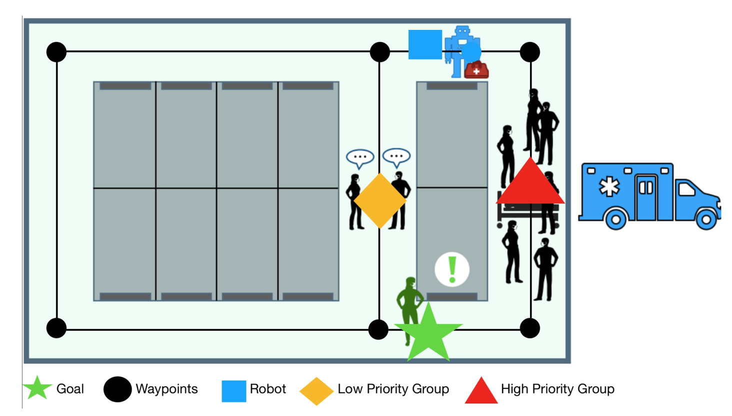


[1] **Bayesian Reinforcement Learning in Factored POMDPs.** Sammie Katt, Frans A. Oliehoek and Christopher Amato. In the *Proceedings of the Eighteenth International Conference on Autonomous Agents and Multi-Agent System* (AAMAS-19), May 2019

[2] **Macro-Action-Based Deep Multi-Agent Reinforcement Learning.** Yuchen Xiao, Joshua Hoffman and Christopher Amato. In *Proceedings of the 2019 Conference on Robot Learning* (CoRL-19), October 2019.

[3] **Multi-Robot Deep Reinforcement Learning with Macro-Actions.** Yuchen Xiao, Joshua Hoffman, Tian Xia, Christopher Amato. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation* (ICRA-20), May 2020.

[4] **Frank, A., Kubota, A., and Riek, L.D. (2019). "Wearable activity recognition for robust human-robot teaming in safety-critical environments via hybrid neural networks".** Andrea Frank, Alyssa Kubota, Laurel D. Riek. *IEEE International Conference Intelligent Robots and Systems* (IROS-19), 2019

[5] **Activity recognition in manufacturing: The roles of motion capture and sEMG+inertial wearables in detecting fine vs. gross motion.** Alyssa Kubota, Tariq Iqbal, Julie A. Shah, and Laurel D. Riek. *IEEE International Conference on Robotics and Automation* (ICRA-19), 2019.

[6] **Situating Robots in the Emergency Department.** Angelique Taylor, Sachiko Matsumoto, and Laurel D. Riek. *AAAI Spring Symposium on Applied AI in Healthcare: Safety, Community, and the Environment* (AAAI-20) , 2020

[7] **Acuity-Aware Social Navigation in Emergency Medicine.** Angelique Taylor, Sachiko Matsumoto, Wesley Xiao, and Laurel D. Riek. *IEEE International Conference Intelligent Robots and Systems* (IROS-20), *In review*