ROB 537:
Learning-Based Control

Week 5, Lecture 2
Project Update, Anatomy of an Abstract

Announcements:
HW 3 Due on 10/30
Midterm Exam: 11/6

Oregon State University

How to do a project in 10 weeks

- Five steps:

– Team selection, topic discussion	week 1
– Project proposal (Topic selection)	week 3
– Paper Intro & background	week 5
– Paper draft	week 9
– Final paper	week 11 (finals week)

Let's discuss next steps

Oregon State University

Paper: Intro & Background

- Introduction
  - What is the interesting problem: big picture?
  - What is the difficulty?
  - What is the significance of solving this problem
  - What do you intend to do ?
  - What is the contribution of this paper?
  - What is coming in the next paragraphs?
- Background
  - Specifics about the problem
  - Key background needed to understand/solve the problem
  - General approaches to the problem
  - Related work addressing these problems
- References
  - Most should be here

Oregon State University

Paper: Preliminary

- Abstract
  - 1 paragraph "ad" for the paper
  - Each section below should have 1-2 sentence summaries
- Introduction
- Background
- Method
  - What is your solution
  - Describe algorithm/theory
- Results
  - Describe set of experiments you will conduct
  - Give preliminary results
- References
  - Full list

Oregon State University

Paper: Final

- Abstract
  - 1 paragraph "ad" for the paper
  - Each section below should have 1-2 sentence summaries
  - 1-2 sentence of key results
- Introduction
- Background
- Method
- Results
  - Describe set of experiments you will conduct
  - Give detailed results
  - Provide Analysis on the results
- Discussion/conclusion
  - Key contributions of paper
  - Key insight
  - Future work
- References

Project Discussion

- Grading:

–	Project proposal (Topic selection)	5
–	Paper intro & background	10
–	Paper draft	15
–	Final paper	50
–	Presentation	20
–	Team performance	+/- 10%

Team Performance

	TM 1	TM 2	TM 3
Professionalism			
Technical direction			
Coding contribution			
Writing contribution			
Presentation contribution			
Percentage effort			

- First five rows are scored out of 3 (3 is best)
- Last one is summary contribution and needs to sum to 100.
- All members need to agree to one table per team
  - This is due with preliminary and final paper

Oregon State University

Team Performance: Example

	TM 1	TM 2	TM 3
Professionalism	2	3	2
Technical direction	3	1	2
Coding contribution	2	3	1
Writing contribution	3	2	2
Presentation contribution	3	2	2
Percentage effort	**40**	**35**	**25**

Oregon State University

Team Performance: What happens if …

Percentage effort	60	20	20

Percentage effort	10	45	45

Percentage effort	33	33	33

Oregon State University

Anatomy of an Abstract

Oregon State University

Abstract: Option 1: Define Idea/Solution First

In this paper, we introduce …

Provide context: The problem … is important, critical …

However, the problem current methods don't work

How does our approach address this? The key feature of this new approach is …

This does x, y, and z to enable solving this difficult problem

Our results show …

Oregon State University

Abstract 1: Conference (Solution Driven)

This paper introduces a new reward shaping approach to achieving effective multiagent coordination in domains that require tight coupling among agents.

Tightly coupled tasks are difficult for many coordination algorithms because they require agents to stumble upon complex joint actions involving precise timing and teaming to fulfill the team objective.

Our new reward framework based on the idea of *counterfactuals* addresses this issue by incentivizing actions that are potentially useful, but were not previously incentivized because other agents in the team had not yet discovered the complementary actions.

Our results show that the proposed approach improves performance by 166%, and speeds up convergence by a factor of four over robots using the global objective directly (Reference).

Oregon State University

Abstract: Option 2: Define Problem First

The problem of … is important, critical …

However, the problem is difficult and current methods don't work

In this paper, we …

The key feature of this new approach is …

This does x, y, and z to enable solving this difficult problem

Our results show

Oregon State University

Abstract 2: Conference (Problem driven)

Autonomous multi-robot teams can improve both the speed and effectiveness of robotic exploration in complex domains.

However, use of multi-robot systems presents additional challenges when the robots' actions are tightly coupled, and robots need to stumble upon complex joint actions involving precise timing and teaming to fulfill the team objective.

In this paper, we present a new reward shaping framework based on the idea of *counterfactuals* to address the coordination problem in tightly coupled domains. This approach eliminates the need for robots to stumble upon the correct joint action at the same time, allowing robots to learn complex tasks independently.

Our results show that the proposed approach improves performance by 166%, and speeds up convergence by a factor of four over robots using the global objective directly (Reference).

Oregon State University

Abstract 3: Magazine – For a General Audience

Teams of artificially intelligent planetary rovers have tremendous potential for space exploration, allowing for reduced cost, increased flexibility and increased reliability.

However, having these multiple autonomous devices acting simultaneously leads to a problem of coordination: to achieve the best results, the they should work together. This is not a simple task. Due to the large distances and harsh environments, a rover must be able to perform a wide variety of tasks with a wide variety of potential teammates in uncertain and unsafe environments.

Directly coding all the necessary rules that can reliably handle all of this coordination and uncertainty is problematic.

Instead, this article examines tackling this problem through the use of coordinated reinforcement learning: rather than being programmed what to do, the rovers iteratively learn through trial and error to take take actions that lead to high overall system return. To allow for coordination, yet allow each agent to learn and act independently, we employ state-of-the-art reward shaping techniques.

The article uses visualization techniques to break down complex performance indicators into an accessible form, and identifies key future research directions.

Oregon State University

Abstract 4: Journal – Application

Intelligent air traffic flow management is one of the fundamental challenges facing the Federal Aviation Administration (FAA) today. FAA estimates put weather, routing decisions and airport condition induced delays at 1,682,700 hours in 2007 (Ref), resulting in a staggering economic loss of over $41 Billion (Ref).

New solutions to the flow management are needed to accommodate the threefold increase in air traffic anticipated over the next two decades. Indeed, this is a complex problem where the interactions of changing conditions (e.g., weather), conflicting priorities (e.g., different airlines), limited resources (e.g., air traffic controllers) and heavy volume (e.g., over 40,000 flights over the US airspace) demand an adaptive and robust solution.

In this paper we explore a multiagent algorithm where agents use reinforcement learning to reduce congestion through local actions.

Each agent is associated with a fix (a specific location in 2D space) and has one of three actions: setting separation between airplanes, ordering ground delays or performing reroutes. We simulate air traffic using FACET which is an air traffic flow simulator developed at NASA and used extensively by the FAA and industry.

Our FACET simulations on both artificial and real historical data from the Chicago and New York airspaces show that agents receiving personalized rewards reduce congestion by up to 80% over agents receiving a global reward and by up to 90% over a current industry approach (Monte Carlo estimation).

Oregon State University