

Module 1

Course Overview

CS 886 Sequential Decision Making
and Reinforcement Learning
University of Waterloo


Modeling Dimensions for Decision Making

- Uncertainty: **Deterministic** vs **stochastic** proc.
- Time: **One shot** vs **sequential** process
- Observability: **full** vs **partial** observability
- Learning: **Complete** vs **incomplete** model
- # of agents: **One** vs **multiple** agents
- Variable type: **Discrete** vs **continuous**

Roadmap

- Start with Markov Decision Process
 - **Sequential** and **stochastic** process
 - **Single** agent in a **complete** model of a **fully observable** and **discrete** process
- Gradually relax assumptions
 - **Continuous** variables
 - **Partial observability**
 - **Incomplete** model (learning)
 - **Multiple** agents

Roadmap



Legend:
MDP: Markov Decision Process
PO: Partially Observable
RL: Reinforcement Learning
MA: Multi-agent
SG: Stochastic Games

Relaxations:
→ : Partial observability
→ : Incomplete model
→ : Multiple agents

Computational Tricks

- Structure:
 - Additive and multiplicative decompositions
- Approximation:
 - Monte Carlo techniques