

ADVANCED REINFORCEMENT LEARNING

JULY 28–29, 2022 | Instructors: Pulkit Agrawal, Cathy Wu | professional.mit.edu/arl

Note: All times are US Eastern Daylight Time. Schedule is subject to change.

DAY 1: THURSDAY, JULY 28 (9:00am–7:30pm)	
9:00–10:00 AM	WELCOME <ul style="list-style-type: none"> • Meet and greet • Problem description from class attendees
10:00–11:00 AM	SESSION 1: OVERVIEW OF DEEP RL AND ITS LIMITATIONS <ul style="list-style-type: none"> • Summary of popular algorithms like DQN, PPO, A3C • Motivation for advanced topics <ul style="list-style-type: none"> - Guidance from experts: Learning from demonstrations - Use of existing datasets: Offline RL - Increasing data efficiency via model learning - Multi-task RL - Curriculum learning - Advanced exploration and exploitation
11:00–11:30 AM	BREAK
11:30 AM–1:00 PM	SESSION 2: LEARNING FROM DEMONSTRATIONS <ul style="list-style-type: none"> • Learning from experts <ul style="list-style-type: none"> - Behavior cloning - DAGGER - Augmenting behavior cloning with RL - Observational learning • Practical issues and solutions: Causal confusion, inability to go beyond expert performance, etc.
1:00–2:00 PM	LUNCH
2:00–3:00 PM	SESSION 3: OFFLINE RL <ul style="list-style-type: none"> • Introduction to offline RL • State-of-the art in offline RL • Off-policy evaluation • Applications of offline RL
3:00–3:30 PM	SESSION 4: HANDS-ON SESSION <ul style="list-style-type: none"> • Go through Jupyter Notebook with code • Run algorithms and modify them to analyze what matters
3:30–4:00 PM	COFFEE BREAK
4:00–5:00 PM	SESSION 5: MODEL-BASED RL <ul style="list-style-type: none"> • Inverse models • Cross-entropy method • When to use a model and when to be model-free?
5:00–5:30 PM	SESSION 6: HANDS-ON SESSION ON MODEL-BASED RL
5:30–5:45 PM	BREAK
5:45–6:30 PM	SESSION 7: MONTE CARLO TREE SEARCH <ul style="list-style-type: none"> • Applications to go • Applications to other areas
6:30–7:30 PM	RECEPTION

ADVANCED REINFORCEMENT LEARNING

JULY 28–29, 2022 | Instructors: Pulkit Agrawal, Cathy Wu | professional.mit.edu/arl

Note: All times are US Eastern Daylight Time. Schedule is subject to change.

DAY 2: FRIDAY, JULY 29 (9:30am–7:00pm)	
9:30–11:00 AM	SESSION 8: MULTI-TASK / GOAL-CONDITIONED RL <ul style="list-style-type: none"> • Inverse models • Connections between inverse models and RL • Hindsight experience replay • Approaches to hierarchical RL
11:00–11:30 AM	BREAK
11:30 AM–12:30 PM	SESSION 9: EXPLORATION & CURRICULUM LEARNING <ul style="list-style-type: none"> • Curriculum learning • Advanced exploration algorithms
12:30–1:30 PM	LUNCH
1:30–3:00 PM	SESSION 10: MISCELLANEOUS TOPICS <ul style="list-style-type: none"> • Multi-agent RL • Generalization in RL
2:00–3:00 PM	SESSION 3: OFFLINE RL <ul style="list-style-type: none"> • Introduction to offline RL • State-of-the art in offline RL • Off-policy evaluation • Applications of offline RL
3:00–3:30 PM	COFFEE BREAK
3:30–4:45 PM	SESSION 11: APPLICATIONS OF ADVANCED RL ALGORITHMS OR A TOPIC OF PARTICIPANT'S CHOICE
4:45–5:00 PM	BREAK
5:00–6:30 PM	SESSION 12: DISCUSSION OF PARTICIPANT APPLICATIONS <ul style="list-style-type: none"> • Miscellaneous assumptions and roadmap for RL <ul style="list-style-type: none"> - The reset requirement in RL - Tradeoffs in using different RL algorithms - Evolutionary methods and its relationship to reinforce • Discussion on topic of participant's choice • Participants can talk about their application and we will formulate as a group approaches to solve it
6:30–7:00 PM	SESSION 13: AMA