Northwestern

Northwestern University – Kellogg School of Management
2022 RL Workshop
August 15 – 18, 2022
James L. Allen Center, Barr Forum Classroom 200

DAY 1 – Monday, August 15
9:00am - 10:30am: Background on MDPs, value iteration + policy iteration, online tabular algorithms
10:30am - 10:45am: *Break*
10:45am - 12:00pm: Online tabular algorithms: intro to deep RL
12:00pm - 1:00pm: *Lunch, Park Dining Room*
1:00pm - 3:00pm: Finish deep RL, lab session: Developing simulator for OR models with Open AI Gym
                      API
3:00pm - 3:30pm: *Break*
3:30pm - 5:00pm: Lab session: Implementing online tabular algorithms for WindyGridWorld and
                      ambulance routing on graph
5:15pm – 6:00pm: *Cocktail Reception, Park Dining Room*
6:00pm: *Dinner, Park Dining Room*

DAY 2 – Tuesday, August 16
9:00am - 10:30am: Nonparametric RL with adaptive discretization
10:30am - 10:45am: *Break*
10:45am - 12:00pm: RL in MDPs with exogenous inputs
12:00pm - 1:00pm: *Lunch, Park Dining Room*
1:00pm - 3:00pm: Lab session: tree-based discretization techniques for ambulance routing on metric
                      Space
3:00pm - 3:30pm: *Break*
3:30pm - 5:00pm: Lab session: hindsight planning in revenue management problems

DAY 3 – Wednesday, August 17
9:00am - 10:15am: Deep RL and its applications, challenges of processing network controls, and policy-gradient
                      algorithms
10:15am - 10:30am: *Break*
10:30am - 11:00am: Intro into neural networks
11:00am - 12:00pm: Deep actor-critic algorithms for processing network controls
12:00pm - 1:00pm: *Lunch, Park Dining Room*
1:00pm - 2:00pm: Implementation of REINFORCE
2:00pm - 3:00pm: Implementation of a value neural network
3:00pm - 3:20pm: *Break*
3:20pm – 5:00pm: Implementation of an actor-critic algorithm

Day 4 - Thursday, August 18
9:00am – 9:30am: Introduction (inventory optimization models, and RL approaches that have been applied
                      to them)
9:30am - 10:30am: The newsvendor problem as a multi-armed bandit
10:30am - 10:45am: *Break*
10:45am - 12:00pm: The base-stock optimization problem as an MDP
12:00pm - 1:00pm: *Lunch, Park Dining Room*
1:00pm - 2:30pm: The beer game as a deep RL problem
2:30pm - 2:45pm: *Coffee break*
2:45pm - 4:00pm: Finish coding deep RL for beer game