Reinforcement Learning Study Guidelines for REU Students

Objective: The objective of this study plan is to provide a comprehensive introduction to Reinforcement Learning (RL) concepts and techniques and prepare students to be able to apply these techniques to solve real-world problems. The plan will cover foundational RL concepts, explore various RL algorithms, and highlight recent advancements in the field. By the end of this study plan, REU students will have gained a deep understanding of RL and its applications and be equipped with the necessary skills to pursue research in RL.

Timeline	Agenda
Week 1 **Introduction to Reinforcement Learning and Markov Decision Processes**	● Learn about the basic concepts of Reinforcement Learning (RL), such as agents, environments, states, actions, rewards, and policies. ● Understand the difference between Model-Free and Model-Based RL. ● Understand the difference between Off-Policy and On-Policy. ● Learn about Markov Decision Processes (MDPs) and their mathematical framework for modeling RL problems. ● Understand the concepts of value functions, Bellman equations, and the Bellman optimality principle. ● Read chapters 1, 2, and 3 from Sutton's RL book. ● Work on chapter 1, 2, and 3 Implementation (Check Python Implementation).
Week 2 **Policy Iteration and Dynamic Programming**	● Learn about Policy Iteration and its mathematical framework for solving RL problems. ● Understand the concepts of dynamic programming, value iteration, and policy evaluation. ● Read chapter 4 from Sutton's RL book. ● Work on chapter 4 Implementation (Check Python Implementation)
Week 3 **Monte Carlo Methods and Temporal Difference**	● Learn about Monte Carlo methods for RL and their advantages over Dynamic Programming methods. ● Understand the concepts of first visit and every-visit Monte Carlo algorithms. ● Learn about Temporal Difference (TD) learning and its advantages over Monte Carlo methods.

	• Understand the concepts of TD (0) and TD (lambda) algorithms. • Read chapter 5 and 6 from Sutton's RL book. • Work on chapter 5 and 6 Implementation (Check Python Implementation)
Week 4 **Function Approximation and Policy Gradients**	• Learn about the challenges of RL in high-dimensional state and action spaces. • Understand the concepts of function approximation and feature engineering. • Learn about Policy Gradients (PG) methods and their advantages over Q-learning. • Understand the concepts of policy gradients, advantage functions, and the REINFORCE algorithm. • Read Paper Policy Gradient Methods for Reinforcement Learning with Function Approximation • Read chapters 9, 10, 11 and 12 from Sutton's RL book. • Work on chapters 9, 10, 11 and 12 Implementation (Check Python Implementation).
Week 5 **Deep Q-networks**	• Read the original paper on DQNs: Human-level control through deep reinforcement learning • Understand the challenges of DQNs, such as stability and exploration-exploitation trade-offs. • Implement a DQN algorithm for playing Atari games, using PyTorch and OpenAI Gym • Programming Example: Implementing DQN with PyTorch and OpenAI Gym
Week 6 **Actor-Critic Methods**	• Learn about Actor-Critic methods and their combination of value-based and policy-based RL. • Understand the concepts of Actor-Critic architectures, TD error, and the A2C algorithm. • Implement an A2C algorithm for playing Pong, using PyTorch and OpenAI Gym • Programming Example: Implementing A2C with PyTorch and OpenAI Gym

Week 7 **Fast Learning and Batch RL**	● Understand the concept of Fast Learning, Experience Replay, Sample Efficiency, Prioritized Experience Replay, Policy Distillation, and Batch RL. ● Watch Stanford's Fast Learning and Batch RL lectures - CS234: Reinforcement Learning Winter 2021 ● Read Paper Prioritized Experience Replay ● Read Paper Off-Policy Deep Reinforcement Learning without Exploration
Week 8 **Multi-Agent RL**	● Learn about Multi-Agent RL (MARL) and its challenges and applications in games, robotics, and social dilemmas. ● Understand the concepts of MARL, independent and joint learning, and the MADDPG algorithm. ● Implement a MADDPG algorithm for playing a simple game with two agents, using PyTorch and OpenAI Gym ● Programming Example: Implementing MADDPG with PyTorch and OpenAI Gym.
Week 9 **Recent Advances in RL**	● Learn about recent advances in RL research, such as Deep RL, Meta RL, Imitation Learning, Inverse RL, and Hierarchical RL ● Read the selected RL papers (Check next page).
Week 10 **Work on RL Projects**	● Learn to design a project with problem formulation, algorithm selection, parameter selection, and representation design---fits together into a complete solution, and how to make appropriate choices when deploying RL in the real world. ● A Complete Reinforcement Learning System (Capstone) | Coursera ● Work on listed Robotics Projects.

Resources

- Reinforcement Learning: An Introduction" by Richard S. Sutton and Andrew G. Barto.
- Python Implementation of Reinforcement Learning: An Introduction
- Spinning Up in Deep RL!
- Key Papers in Deep RL — Spinning Up documentation
- Recent RL Papers from Conference

RL Libraries

- OpenAI Gym: A toolkit for developing and comparing reinforcement learning algorithms.
- Stable Baselines 3 is a learning library based on the Gym API. It is designed to cater to complete beginners in the field who want to start learning things quickly.
- RL Baselines3 Zoo builds upon SB3, containing optimal hyperparameters for Gym environments as well as code to easily find new ones.
- Tianshou is a learning library that's geared towards very experienced users and is designed to allow for ease in complex algorithm modifications.
- RLlib is a learning library that allows for distributed training and inferencing and supports an extraordinarily large number of features throughout the reinforcement learning space.
- PettingZoo is like Gym, but for environments with multiple agents.

RL Courses

- Reinforcement Learning Specialization
- DeepMind Reinforcement Learning Lecture Series 2021
- UC Berkley CS 294 Deep Reinforcement Learning, Spring 2017
- Stanford CS234: Reinforcement Learning | Winter 2019 - YouTube

Selected Papers to Read

- Playing Atari with Deep Reinforcement Learning
- Learning to Adapt in Dynamic, Real-World Environments through Meta-Reinforcement Learning
- Apprenticeship Learning via Inverse Reinforcement Learning
- Deep Reinforcement Learning from Self-Play in Imperfect-Information Games
- Proximal Policy Optimization Algorithms
- Deep Reinforcement Learning with Double Q-learning
- Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor
- Smith & Gasser, The Development of Embodied Cognition: Six Lessons from Babies
- Silver, Huang et al., Mastering the Game of Go with Deep Neural Networks and Tree Search
- Bojarski et al., *End to End Learning for Self-Driving Cars*
- Ross et al., Learning Monocular Reactive UAV Control in Cluttered Natural Environments
- Abbeel et al., Apprenticeship Learning via Inverse Reinforcement Learning
- Ziebart et al., Maximum Entropy Inverse Reinforcement Learning
- Lillicrap et al., Continuous control with deep reinforcement learning
- Lake et al., Building Machines That Learn and Think Like People
- Finn et al., Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks
- Silver et al., Mastering the Game of Go without Human Knowledge

Articles to Read

- Math Behind Reinforcement Learning, the Easy Way | by Ziad SALLOUM | Towards Data Science
- Hands-On Reinforcement Learning Course: Part 1 | by Pau Labarta Bajo | Towards Data Science
- States, Observation, and Action Spaces in Reinforcement Learning | by #Cban2020 | The Startup | Medium
- A hands-on introduction to deep reinforcement learning using Unity ML-Agents

Applying RL in Robotics

- Reinforcement Learning with ROS and Gazebo - Artificial Intelligence Research
- Benchmarking Reinforcement Learning Algorithms on Real-World Robots
- A ROS2-based framework for TurtleBot3 DRL autonomous navigation
- ML-Agents with Unity
- High Dimensional Planning and Learning for Off-Road Driving
- Deep Reinforcement Learning for Autonomous Driving: A Survey
- Using TurtleBot in Deep Reinforcement Learning

Important

If you are not familiar with Machine Learning, please consider building a foundation first, before diving into Reinforcement Learning.

- Machine Learning Specialization
- Mathematics for Machine Learning and Data Science Specialization