Reinforcement Learning

Austin Wu

May 29, 2025

Contents

Chapter 1

An Introduction to RL

1.1 Preface

While I am in now way an expert and honestly straight up inexperienced in reinforcement learning, I think a lot of the modern material can have poorly explained overviews on what is happening and leave people scratching their heads at the bigger picture even if they understand the smaller fine points. I wanted to create an article that helps alleviate this problem. The material in this article is heavily structured off of the class I took last semester at Cornell, Wen Sun's 4789: Introduction to Reinforcement Learning

1.2 The thing we care about

At some level, all of reinforcement learning comes back to the idea of a Markov Decision Process, or MDP. Specifically we care about optimizing an MDP. A Markov decision process contains multiple things. Recall that

Definition 1.1: Markov Decision Process

ecall an MDP is defined as such

$$MDP = \{A, S, \pi, mu, r\} \tag{1.1}$$

Where A is the action space, S is the state space, π is a policy, μ is the initial state, and r is the reward function.

Now this being said, we actually need to create a distinction between

Definition 1.2: Finite Horizon Markov Decision Process

i

Definition 1.3: Infinite Horization Markov Decision Process

i

Chapter 2

Convergence of RL Systems with Finite Policy Action Space

PSYCH you thought I would actually fill this out im so lazy lol

Chapter 3

Convergence of RL Systems with Parameterized Policy Space

Chapter 4

Fine Tuning RL Models

Chapter 5

Imitation Learning