A Code-Driven Introduction to Reinforcement Learning

Dr. Phil Winder, CEO

https://WinderResearch.com

Winder
DATA SCIENCE | CLOUD NATIVE

O'REILLY®

Reinforcement Learning

Industrial Applications of Intelligent Agents

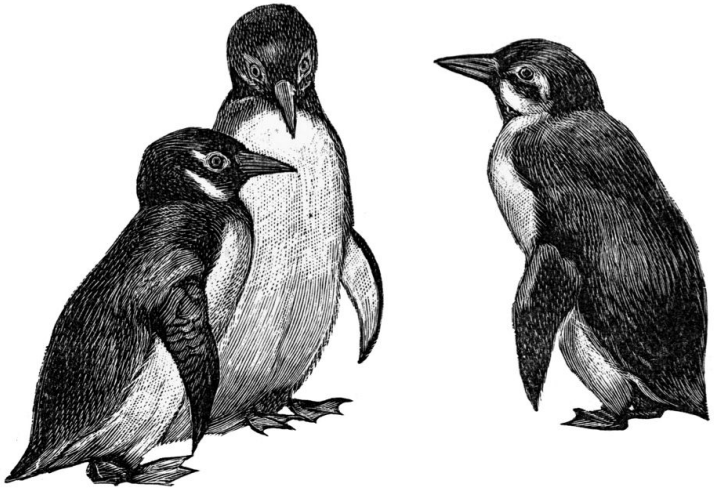

Phil Winder, Ph.D.

rl-book.com

Release date: December 1st, 2020

https://WinderResearch.com

1. What is RL?
2. Coding the Simulation
3. Coding the RL solution
4. Next steps

https://WinderResearch.com

Winder
DATA SCIENCE | CLOUD NATIVE

What is RL?

https://WinderResearch.com

"Backwards Bike" - Destin Sandlin - SmarterEveryDay

https://WinderResearch.com

"Backwards Bike" - Destin Sandlin - SmarterEveryDay

https://WinderResearch.com

"Backwards Bike" - Destin Sandlin - SmarterEveryDay

https://WinderResearch.com

"Backwards Bike" - Destin Sandlin - SmarterEveryDay

https://WinderResearch.com

"Backwards Bike" - Destin Sandlin - SmarterEveryDay

https://WinderResearch.com

What do I need?

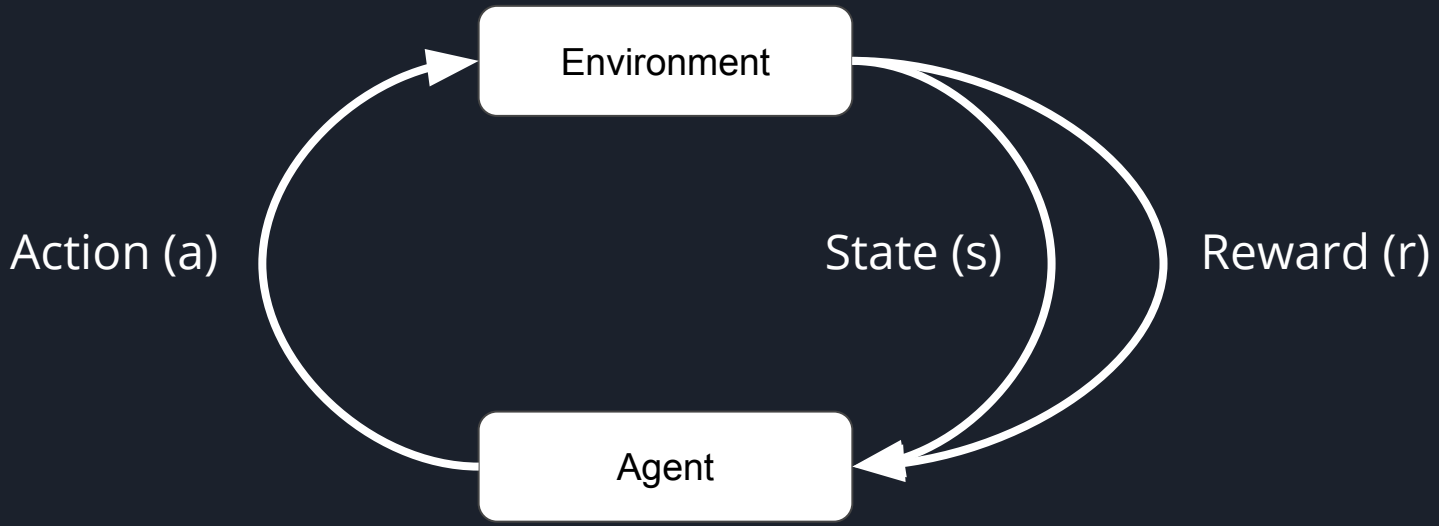

https://WinderResearch.com

Coding the MDP

https://WinderResearch.com

Accompanying Notebook

Go here in your browser:

https://tinyurl.com/rl-code-notebook

Then click on your preferred notebook service…

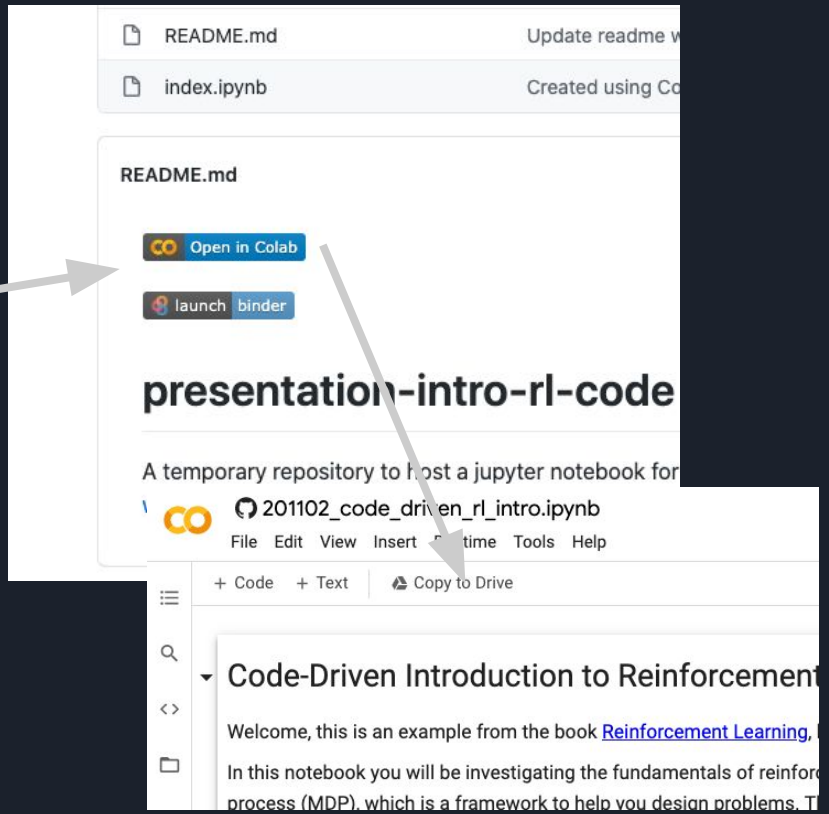

Source: https://rl-book.com/learn/

https://WinderResearch.com

Coding the RL Solution

https://WinderResearch.com

Next Steps

https://WinderResearch.com

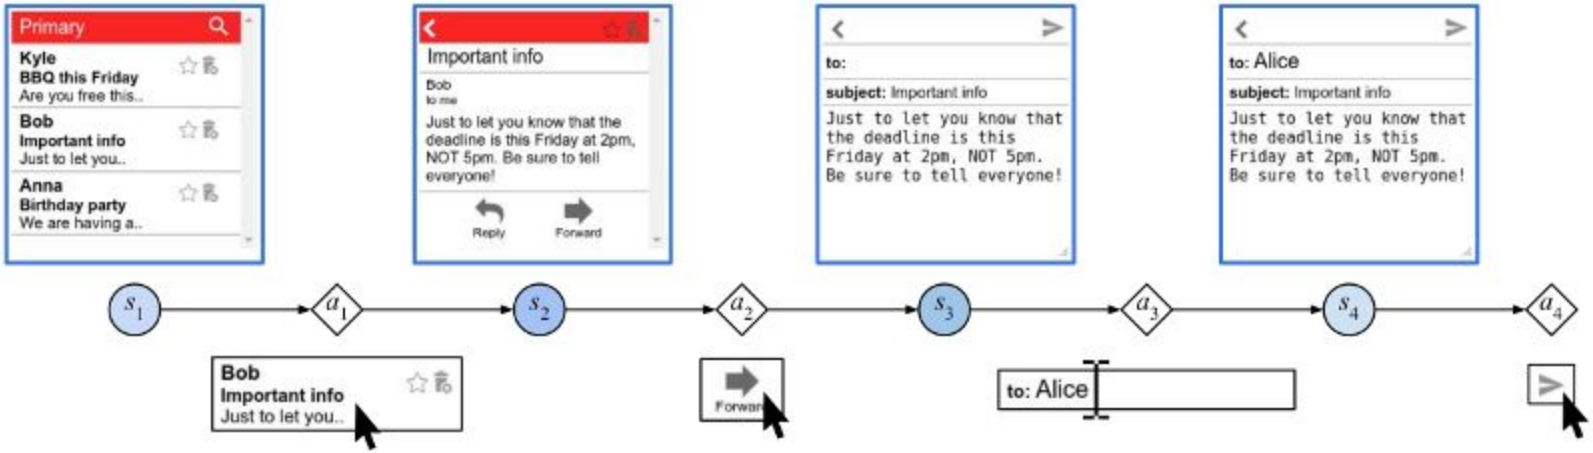

Liu et al. 2018

https://WinderResearch.com

Winder
DATA SCIENCE | CLOUD NATIVE

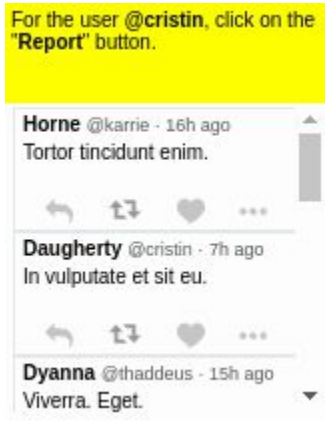

https://WinderResearch.com

Winder
DATA SCIENCE | CLOUD NATIVE

Liu et al. 2018

O'REILLY®

Reinforcement Learning

Industrial Applications of Intelligent Agents

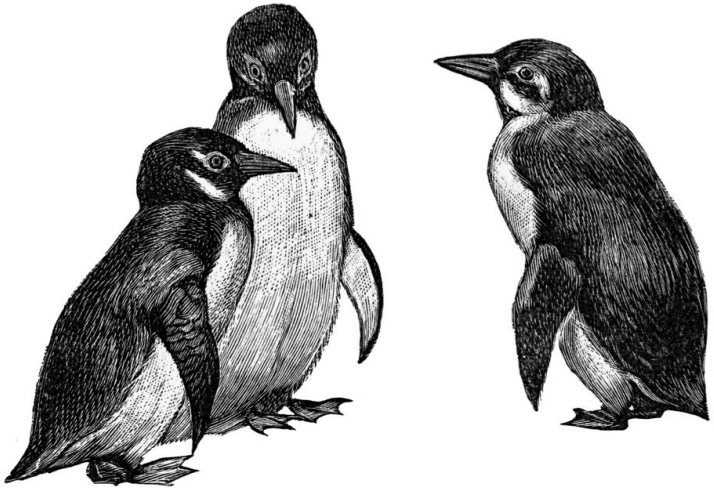

Phil Winder, Ph.D.

rl-book.com

phil@WinderResearch.com

DrPhilWinder

WinderResearch.com

https://WinderResearch.com

Title
A code-driven introduction to reinforcement learning

Abstract
Reinforcement learning (RL) is lined up to become the hottest new artificial intelligence paradigm in the next few years. Building upon machine learning, reinforcement learning has the potential to automate strategic-level thinking in industry.

In this presentation I present a code-driven introduction to RL, where you will explore a fundamental framework called the Markov decision process (MDP) and learn how to build an RL algorithm to solve it.

First I show you how to create a simple “GridWorld” simulation of the MDP, from the ground up, to help demonstrate why and how RL works. Then I derive a simple RL algorithm that’s capable of solving your simulation. Finally I will provide actionable next steps to show you how to take this learning and apply it to industry.

This presentation includes a Jupyter notebook that you can tinker with during the presentation. Full instructions will be provided. Although this presentation is suitable for beginners, you will benefit if you have some exposure to data science and machine learning.

https://WinderResearch.com