

Kavan Mehta

Mrs. Secord

ISM 1

19 January 2022

Research Assessment #14

Date: 19 January 2022

Subject: Reinforcement Learning Overview

MLA citation(s):

“What Is Reinforcement Learning? – Overview of How It Works.” *Synopsys*, Synposys, www.synopsys.com/ai/what-is-reinforcement-learning.html#:~:text=The%20Reinforcement%20Learning%20problem%20involves,maximization%20of%20expected%20cumulative%20reward.

Assessment:

In this semester, I am currently in the process of selecting a mentor and exploring final product ideas in order to create a successful machine learning project in ISM. Hence, I went to explore one of the most intriguing fields of machine learning, reinforcement learning. I was hoping to find out more about what reinforcement learning exactly is, how we can implement it, and what are the benefits and use cases of reinforcement learning. To earn this new knowledge about reinforcement learning, I found the article, “What is Reinforcement Learning?,” which went over how reinforcement learning works, what the specific benefits of reinforcement learning are, and explores some examples and challenges of reinforcement learning.

After understanding reinforcement learning through this article, I was able to understand the differences between reinforcement learning and supervised/unsupervised learning,

benefits/examples of reinforcement learning, and what reinforcement learning actually is. I learned that reinforcement learning is all based on an agent learning based on past mistakes to maximize its success or reward over a long period of time (“What is Reinforcement Learning?” 3). In addition to this, reinforcement learning models mainly consist of 4 main parts: “The agent or learner, the environment the agent interacts with, the policy that the agent follows to take actions, and the reward signal that the agent observes upon taking actions” (“What is Reinforcement Learning?” 3). There are also specific challenges in reinforcement learning such as difficulty obtaining a variety of new experiences to an extent of measurable success, difficult to interpret as observers, and the long-term efficiency is usually very delayed in terms of time (“What is Reinforcement Learning?” 8). This helped connect with my knowledge of machine learning models and helped me get a better understanding of reinforcement learning and has helped me gain a better idea of how I could possibly use it. Some examples that helped me gain more inspiration were how we can use reinforcement learning in robotics to find the shortest path to a certain point using trial and error reinforcement learning and how reinforcement learning could also be used are strategic games such as the AlphaGo model that is the best at the Chinese game Go (“What is Reinforcement Learning?” 6). By using an extremely large variety of situations, reinforcement learning models could possibly find the most optimized solutions to get the best accuracy at a specific task. Furthermore, since reinforcement learning is not affected by individual biases and focuses on an entire situation instead, they are much better able to adapt to new situations and predict accurately to gain a “reward” or win. This could be very useful for my final project as it could help create a consistent model that is even more optimized than neural networks in some cases. This is another reason why I want to continue exploring reinforcement learning for applications in my final project.

I am currently developing ideas for my final project that is likely to focus on computer vision or natural language processing and thus am considering if I should also explore reinforcement learning to create my model. I will continue to explore any other ideas and since I do like reinforcement learning, I need to consider the feasibility of learning and determine if reinforcement learning would be a great model to use for my final project. I will look at some examples of projects for inspiration and research some machine learning problems to determine my final project idea in the future.