

**SDS PODCAST**

**EPISODE 440**

**FIVE MINUTE**

**FRIDAY:**

**MUZERO: LEARNING**

**WITHOUT RULES**



(00:05):

This is Five-Minute Friday on MuZero: Learning Without Rules.

(00:19):

On last week's Five-Minute-Friday episode, I introduced the concept of artificial general intelligence (or AGI, for short), a theoretical algorithm that one day could have all of the intellectual capacities of a human being. I also introduced the company DeepMind and the landmark deep reinforcement learning algorithms they developed over the past decade, each one a stepping stone on the road to creating AGI. If any of AGI, DeepMind, or deep reinforcement learning are unfamiliar terms to you, you might want to check out last week's Five Minute Friday episode to brush up.

(00:58):

Last week's coverage of DeepMind's deep reinforcement learning advances bring me now to MuZero, an algorithm that David Silver and his DeepMind research team published on in the final days of 2020 in the journal Nature, arguably the most prestigious academic science journal.

(01:15):

Like the Deep Q-Learning and AlphaZero algorithms before it, DeepMind's MuZero doesn't require data from human gameplay to learn how to master a game; it learns through trial and error how to maximize its performance and win. The aspect that makes MuZero exceptional relative to its predecessors is that it does not even need to be provided with the rules of the game -- it figures out the rules via trial and error from scratch too!

(01:43):

With this additional flexibility, MuZero can excel on a broader range of challenges than its predecessors. Like its Deep Q-Learning predecessor,



MuZero can master Atari video games; however, this single MuZero algorithm can master board games with perfect information such as chess, Go, and shogi -- board games that were previously the stomping grounds of AlphaZero. So now we have one algorithm, MuZero, that can simultaneously master video games as well as board games, two game categories that require fundamentally different learning styles to master.

(02:22):

Thus, MuZero is a significant advancement toward artificial general intelligence because it excels at a broader range of tasks than any algorithm before it. It is all the more remarkable that it can do so without any initial knowledge of what the rules are for any given problem it sets out to solve.

(02:42):

Given its breadth of applicability, MuZero will no doubt serve as a starting point for even more broadly applicable algorithms in the future. An immediate task it's being considered for is devising new compression algorithms for video, allowing YouTube -- another Alphabet group company that sits alongside DeepMind and Google -- to save tons of storage space and bandwidth in its data centers, and therefore also saving oodles of cash.

(03:11):

Pretty cool, right? The pace of progress in AI is breathtaking, and methodically, step by step, the human race is making advances toward the theorized AGI benchmark, where a single algorithm can command the same breadth of abilities as a human brain.

(03:27):

If you'd like to learn more about the fascinating world of deep reinforcement learning, the data science approach that underlies all of the algorithms I discussed today -- from Deep Q-Learning to AlphaGo, and from AlphaZero



to MuZero -- you can check out Chapter 4 of my book *Deep Learning Illustrated* for a broad overview of how deep reinforcement learning works and how it relates to the concepts of artificial intelligence, including artificial general intelligence.

(03:52):

Further, in Chapter 13 of my book, I explain all of the essential theory and hands-on Python code you need to design your own deep reinforcement learning algorithm, specifically a Deep Q-Learning network like the one DeepMind used to master dozens of Atari video games.

(04:08):

I also have a free 72-minute video on YouTube called *Deep Q Learning Networks*, which is my most popular YouTube video, and like Chapter 14 of my book, in the video I step you through the essential theory and a hands-on code demo in Python.

(04:23):

For a video series that goes beyond Deep Q-Learning into even more advanced deep reinforcement learning techniques, check out SuperDataScience's *Artificial Intelligence A-Z*, which is available via Udemy.

(04:36):

Or, if you prefer books to videos, you can read *Foundations of Deep Reinforcement Learning* by friends of mine, Laura Graesser and Wah Loon Keng. This book is a fabulously rich, hands-on introduction to the gamut of deep reinforcement learning applications all the way through to the most sophisticated approaches.



(04:55):

All right, I hope you learned something new today, perhaps you even became inspired to learn more about the leading approaches behind artificial intelligence today, and regardless are looking forward to another exciting episode of the SuperDataScience podcast soon.