

After AIs mastered Go and Super Mario, scientists have taught them how to 'play' experiments at NSLS-II

March 25 2021



NSLS-II scientists, Daniel Olds (left) and Phillip Maffettone (right), are ready to let their AI agent level up the rate of discovery at NSLS-II's PDF beamline. Credit: Brookhaven National Laboratory



Inspired by the mastery of artificial intelligence (AI) over games like Go and Super Mario, scientists at the National Synchrotron Light Source II (NSLS-II) trained an AI agent—an autonomous computational program that observes and acts—how to conduct research experiments at superhuman levels by using the same approach. The Brookhaven team published their findings in the journal Machine Learning: Science and Technology and implemented the AI agent as part of the research capabilities at NSLS-II.

As a U.S. Department of Energy (DOE) Office of Science User Facility located at DOE's Brookhaven National Laboratory, NSLS-II enables scientific studies by more than 2000 researchers each year, offering access to the facility's ultrabright X-rays. Scientists from all over the world come to the facility to advance their research in areas such as batteries, microelectronics, and drug development. However, time at NSLS-II's experimental stations—called beamlines—is hard to get because nearly three times as many researchers would like to use them as any one station can handle in a day—despite the facility's 24/7 operations.

"Since time at our facility is a precious resource, it is our responsibility to be good stewards of that; this means we need to find ways to use this resource more efficiently so that we can enable more science," said Daniel Olds, beamline scientist at NSLS-II and corresponding author of the study. "One bottleneck is us, the humans who are measuring the samples. We come up with an initial strategy, but adjust it on the fly during the measurement to ensure everything is running smoothly. But we can't watch the measurement all the time because we also need to eat, sleep and do more than just run the experiment."

"This is why we taught an AI agent to conduct scientific experiments as if they were video games. This allows a robot to run the experiment, while we—humans—are not there. It enables round-the-clock, fully



remote, hands-off experimentation with roughly twice the efficiency that humans can achieve," added Phillip Maffettone, research associate at NSLS-II and first author on the study.

According to the researchers, they didn't even have to give the AI agent the rules of the 'game' to run the experiment. Instead, the team used a method called "reinforcement learning" to train an AI agent on how to run a successful scientific experiment, and then tested their agent on simulated research data from the Pair Distribution Function (PDF) beamline at NSLS-II.

Beamline Experiments: A Boss Level Challenge

Reinforcement learning is one strategy of training an AI agent to master an ability. The idea of reinforcement learning is that the AI agent perceives an environment—a world—and can influence it by performing actions. Depending on how the AI agent interacts with the world, it may receive a reward or a penalty, reflecting if this specific interaction is a good choice or a poor one. The trick is that the AI agent retains the memory of its interactions with the world, so that it can learn from the experience for when it tries again. In this way, the AI agent figures out how to master a task by collecting the most rewards.

"Reinforcement learning really lends itself to teaching AI agents how to play video games. It is most successful with games that have a simple concept—like collecting as many coins as possible—but also have hidden layers, like secret tunnels containing more coins. Beamline experiments follow a similar idea: the basic concept is simple, but there are hidden secrets we want to uncover. Basically, for an AI agent to run our beamline, we needed to turn our beamline into a video game," said Olds.

Maffettone added, "The comparison to a video game works well for the



beamline. In both cases, the AI agent acts in a world with clear rules. In the world of Super Mario, the AI agent can choose to move Mario up, down, left, right; while at the beamline, the actions would be the motions of the sample or the detector and deciding when to take data. The real challenge is to simulate the environment correctly—a video game like Super Mario is already a simulated world and you can just let the AI agent play it a million times to learn it. So, for us, the question was how can we simulate a beamline in such a way that the AI agent can play a million experiments without actually running them" said Maffettone.

The team "gamified" the beamline by building a virtual version of it that simulated the measurements the real beamline can do. They used millions of data sets that the AI agent could gather while "playing" to run experiments on the virtual beamline.

"Training these AIs is very different than most of the programming we do at beamlines. You aren't telling the agents explicitly what to do, but you are trying to figure out a reward structure that gets them to behave the way you want. It's a bit like teaching a kid how to play video games for the first time. You don't want to tell them every move they should make, you want them to begin inferring the strategies themselves." Olds said.





The researchers' goal is to translate science into the language of a video game so that AIs can learn how to 'play' the game and eventually master it, thus improving research outcomes. Credit: Brookhaven National Laboratory

Once the beamline was simulated and the AI agent learned how to conduct research experiments using the virtual beamline, it was time to test the AI's capability of dealing with many unknown samples.

"The most common experiments at our beamline involve everything from one to hundreds of samples that are often variations of the same material or similar materials—but we don't know enough about the samples to understand how we can measure them the best way. So, as humans, we would need to go through them all, one by one, take a snapshot measurement and then, based on that work, come up with a good strategy. Now, we just let the pre-trained AI agent work it out," said Olds.

In their simulated research scenarios, the AI agent was able to measure unknown samples with up to twice the efficiency of humans under strongly constrained circumstances, such as limited measurement time.

"We didn't have to program in a scientist's logic of how to run an experiment, it figured these strategies out by itself through repetitive playing." Olds said.

Materials Discovery: Loading New Game

With the AI agent ready for action, it was time for the team to figure out



how it could run a real experiment by moving the actual components of the beamline. For this challenge, the scientists teamed up with NSLS-II's Data Science and System Integration (DSSI) program to create the backend infrastructure. They developed a program called Blueskyadaptive, which acts as a generic interface between AI tools and Bluesky—the software suite that runs all of NSLS-II's beamlines. This interface laid the necessary groundwork to use similar AI tools at any of the other 28 beamlines at NSLS-II.

"Our agent can now not only be used for one type of sample, or one type of measurement—it's very adaptable. We are able to adjust it or extend it as needed. Now that the pipeline exists, it would take me 45 minutes talking to the person and 15 minutes at my keyboard to adjust the agent to their needs," Maffettone said.

The team expects to run the first real experiments using the AI agent this spring and is actively collaborating with other beamlines at NSLS-II to make the tool accessible for other measurements.

"Using our instruments' time more efficiently is like running an engine more efficiently—we are making more discoveries per year happen. We hope that our new tool will enable a new transformative approach to increase our output as user facility with the same resources."

The team who made this advancement possible also consists of Joshua K. Lynch, Thomas A. Caswell and Stuart I. Campbell from the NSLS-II's DSSI program and Clara E. Cook from the University at Buffalo.

More information: Phillip M Maffettone et al, Gaming the beamlines—employing reinforcement learning to maximize scientific outcomes at large-scale user facilities, *Machine Learning: Science and Technology* (2021). DOI: 10.1088/2632-2153/abc9fc



Provided by Brookhaven National Laboratory

Citation: After AIs mastered Go and Super Mario, scientists have taught them how to 'play' experiments at NSLS-II (2021, March 25) retrieved 5 May 2024 from <u>https://techxplore.com/news/2021-03-ais-mastered-super-mario-scientists.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.