

Exploring reinforcement learning to control nuclear fusion reactions

September 9 2022, by Aaron Aupperlee



Ian Char works in the control room of the DIII-D National Fusion Facility.
Credit: Jeff Schneider

A student in Carnegie Mellon University's School of Computer Science (SCS) has used reinforcement learning to help control nuclear fusion

reactions, a significant step toward harnessing the immense power produced in nuclear fusion as a source of clean, abundant energy.

Ian Char, a doctoral candidate in the Machine Learning Department, used [reinforcement learning](#) to control the hydrogen plasma of the tokamak machine at the DIII-D National Fusion Facility in San Diego. He was the first CMU researcher to run an experiment on the sought-after machines, the first to use [reinforcement](#) learning to affect the rotation of a tokamak plasma, and the first person to try reinforcement learning on the largest operating tokamak machine in the United States. Char collaborated with the Princeton Plasma Physics Laboratory (PPPL) on the work.

"Reinforcement learning affected the plasma's pressure and its rotation," Char said. "And that's really our big first here."

Nuclear [fusion](#) happens when hydrogen nuclei smash, or fuse, together. This process releases a tremendous amount of energy but remains challenging to maintain at levels necessary for putting electricity on the grid. Hydrogen nuclei will only fuse under extremely high temperatures and pressure such as those found at the center of the sun, where [nuclear fusion](#) occurs naturally. Physicists have also achieved nuclear fusion in thermonuclear weapons, but these are not useful as energy sources.

Another method to produce nuclear fusion uses magnetic fields to contain a plasma of hydrogen at the required temperature and pressure to fuse the nuclei. This process happens inside a tokamak—a massive machine that uses magnetic fields to confine the hydrogen plasma in a donut shape called a torus. Containing the plasma and maintaining its shape require hundreds of micromanipulations to the magnetic fields and blasts of additional hydrogen particles.

There are few large-scale tokamaks operating in the world that can

facilitate this type of research, and time to run experiments on them is coveted. The DIII-D National Fusion Facility is the only one operating in the United States.

DeepMind, an artificial intelligence subsidiary of Alphabet, Google's parent company, was the first to use reinforcement learning to control the [magnetic field](#) containing the fusion reaction. The lab successfully kept the plasma steady and sculpted it into different shapes. DeepMind ran its experiment on the Variable Configuration Tokamak (TCV) in Lausanne, Switzerland, and published its findings in February in *Nature*.

Char was the first to run a similar reinforcement learning experiment at DIII-D. Reinforcement learning uses data from past attempts to achieve an optimal outcome. During Char's experiment, reinforcement learning algorithms examined historic and real-time data to vary and control the speed of the plasma's rotation in search of optimal stability.

The plasma donut rotates when additional hydrogen particles are shot into it. Varying the speed of these shot particles can potentially stabilize the plasma and make it easier to contain. Char used two learning algorithms for his experiment. In one, he used data from the tokamak collected over several years to train it on how the plasma reacts. The second algorithm observes the condition of the plasma and then decides at what rate and direction to shoot in the additional particles to affect its speed.

"The short-term goal is to give the physicists the tools to cause this differential rotation so they can do the experiments to make this plasma more stable," said Jeff Schneider, a research professor in the Robotics Institute and Char's Ph.D. adviser. "Longer term, this work shows a path to using reinforcement learning to control other parts of the plasma state and ultimately achieve the temperatures and pressures long enough to have a power plant. That would mean limitless, clean energy for

everyone."

Char pitched the project to DIII-D, which is a U.S. Department of Energy Office of Science User Facility managed by General Atomics, last year and was granted a three-hour slot to run his algorithms on June 28. Seated in the control room of the massive DIII-D facility and surrounded by operators, Char loaded his algorithms.

Char demonstrated his algorithms could control the speed of the plasma's rotation. This was the first time reinforcement learning was used to control the rotation. Some problems crept up during the control session and more testing is needed. Char returned to DIII-D at the end of August to continue his work.

"Ian showed a tremendous ability to digest the fusion device-specific control issues and [plasma](#) physics that underlines it," said Egemen Kolemen, an associate professor in Princeton University's Mechanical and Aerospace Engineering Department and one of Char's collaborators at PPPL. "It is a great achievement to apply the theory he learned at CMU to a real fusion problem and lead an experiment on a national fusion facility. That work normally requires years of [plasma physics](#) and engineering training."

More information: Jonas Degraeve et al, Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* (2022). [DOI: 10.1038/s41586-021-04301-9](https://doi.org/10.1038/s41586-021-04301-9)

Provided by Carnegie Mellon University

Citation: Exploring reinforcement learning to control nuclear fusion reactions (2022, September 9) retrieved 19 April 2024 from

<https://techxplore.com/news/2022-09-exploring-nuclear-fusion-reactions.html>

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.