Reinforcement Learning to Increase Nuclear Fusion Efficiency

Joseph Z.X. Koh[1] Sergey Issinski[1] Eric Fan[2] Amanda K. Brown[1] Curtis P. Berlinguette[1,3-5]

[1]Department of Chemistry, University of British Columbia, 2036 Main Mall, Vancouver, British Columbia V6T 1Z1, Canada, [2]Department of Electrical and Computer Engineering, University of Alberta, 9211 116 Street Edmonton, Alberta T6G 1H9, Canada, [3]Department of Chemical and Biological Engineering, University of British Columbia, 2360 East Mall, Vancouver, British Columbia V6T 1Z3, Canada, [4]Stewart Blusson Quantum Matter Institute, University of British Columbia, 2355 East Mall, Vancouver, British Columbia V6T 1Z4, Canada, [5]Canadian Institute for Advanced Research (CIFAR), 661 University Avenue, Toronto, Ontario, M5G

Correspondence to: Curtis P. Berlinguette curtis.berlinguette@ubc.ca

Deuterium–deuterium (D–D) fusion in a metal lattice

Nuclear fusion occurs when atomic nuclei collide to form a heavier nucleus, releasing energy. Deuterium (D), a stable isotope of hydrogen, is commonly used as a fusion fuel. Increases in fuel density increase the probability of nuclear collisions. For example, magnetic and inertial confinement fusion typically confine deuterium in a plasma at densities of approximately $10^{20}\,\mathrm{m}^{-3}$ and $10^{31}\,\mathrm{m}^{-3}$, respectively [1]. In contrast, solid metal lattices, such as palladium, can reach deuterium densities on the order of $10^{28}\,\mathrm{m}^{-3}$ and provide a more stable medium than a plasma for storing D fuel [2]. These advantages make lattice confinement fusion an attractive platform to study fusion in condensed matter.

Thunderbird Reactor electrochemically loads D into palladium to increase fusion rates

We built a bench-sized apparatus, the Thunderbird Reactor, to test the hypothesis that electrochemically loading a palladium target with deuterium can enhance nuclear fusion rates [2]. In the reactor, a plasma thruster delivers deuterium to one side of the palladium target through plasma-immersion ion implantation. D–D fusion events are quantified with a neutron detection system that measures neutrons as a proxy for these fusion reactions. We showed that fusion rates increase by 15% when an electrochemical cell supplies additional deuterium to the opposite side of the target [2].

Many parameters can impact fusion efficiency in the Thunderbird Reactor

We now seek to increase fusion efficiency, $\frac{reaction\ yield}{input\ energy}$, by maximizing reaction yield while minimizing input energy. The Thunderbird Reactor has many parameters that can be controlled in pursuit of increasing fusion efficiency, such as plasma pressure and ion density, target material and geometry, plasma-generation input power, and target voltage.

Development of a reinforcement learning system to increase fusion efficiency

We are developing a reinforcement learning model to discover reactor parameters that improve fusion efficiency. The model defines fusion efficiency as the ratio of neutron yield, a direct proxy for fusion rate, to the total energy input to the fuel in the target. In this model, an agent receives a reward when a particular combination of parameters and their values increases fusion efficiency. The agent then updates the reactor state accordingly, for example by adjusting target voltage, plasma pressure, and plasma-generation input power.

We also plan to leverage Ada, the Berlinguette Group's self-driving lab, to accelerate exploration of materials for the Thunderbird Reactor's metal target in pursuit of higher fusion efficiency.

Acknowledgments

Thistledown Foundation.
Canadian Natural Sciences and Engineering Research Council (RGPIN-2018-06748).
Canadian Foundation for Innovation (229288)
Canadian Institute for Advanced Research (BSE-BERL-162173).
Thunderbird team members, past and present, including Kuo-Yi Chen, Faye Yu, and Jackson Williams.

References

[1] Meschini et al., Review of commercial nuclear fusion projects. Front. Energy Res., 2023, 11, 1157394.
[2] Chen et al., Electrochemical loading enhances deuterium fusion rates in a metal target. Nature, 2025, 644, 640–645.