Magnetic Control of the TCV tokamak through Deep Reinforcement Learning

F. Felici[1], J. Degrave[2], J. Buchli[2], M. Neunert[2], B. Tracey[2], F. Carpanese[1,2],T. Ewalds[2], R. Hafner[2], A. Abdolmaleki[2], D. de las Casas[2], C. Donner[2], L. Fritz[2], C. Galperti[1], A. Huber[2], J. Keeling[2], M. Tsimpoukelli[2], J. Kay[2], A. Merle[1], J-M. Moret[1], S. Noury[2], F. Pesamosca[1], D. Pfau[2], O. Sauter[1], C. Sommariva[1], S. Coda[1], B. Duval[1], A. Fasoli[1], P. Kohli[2], K. Kavukcuoglu[2], D. Hassabis[2], M. Riedmiller[2]

[1] *École Polytechnique Fédérale (EPFL), Swiss Plasma Center (SPC), Lausanne, Switzerland*

[2] *DeepMind, London, UK*

Maintaining a stable tokamak discharge in the desired equilibrium requires control of the plasma position and shape using poloidal field coils. Traditionally, shape and position control have been handled by linear controllers designed using control engineering methods. In this work [1], we demonstrate the first successful application of Deep Reinforcement Learning (RL) to the magnetic control of tokamaks. The Deep RL algorithm learns magnetic controllers for a target TCV plasma equilibrium solely by interacting with a free-boundary evolution model, which simulates the plasma equilibrium evolution coupled to the circuit equations for external conductors. These controllers generate and maintain a plasma of the desired shape by receiving only magnetic measurements, without the need for equilibrium reconstruction, and actuating directly the coil power supply voltages in real-time. A variety of controllers have been implemented and successfully tested in TCV experiments, achieving several different target shapes including elongated plasmas, negative triangularity and snowflake configurations. We additionally demonstrate sustained ‘droplets’ on TCV where two separate plasmas are maintained simultaneously within the vessel using this controller. This control design method has the potential to reduce the effort required to obtain a new target shape in a tokamak. Moreover, this work represents the first use of reinforcement learning for feedback control on a tokamak and paves the way for combining physics models and machine learning for improving other aspects of control of fusion plasmas.

References

[1] J. Degrave, F. Felici, J. Buchli et al. Magnetic control of tokamak plasmas through deep reinforcement learning. Nature 602, 414–419 (2022). https://doi.org/10.1038/s41586-021-04301-9