ACTIVE TEARING MODE AVOIDANCE WITH MACHINE LEARNING CONTROLLERS

[1]A. ROTHSTEIN, [2]J. SEO, [3]S.K. KIM, [1]A. JALALVAND, [1]H.J. FARRE-KAGA, [3]R. SHOUSHA, [1]R. CONLIN, [1]J. ABBATE, [3]K. ERICKSON, [1]J. WAI, [1,3]E. KOLEMEN

[1]Princeton University, Princeton, NJ, USA
[2]Chung-ang University, Seoul, South Korea
[3]Princeton Plasma Physics Laboratory, Princeton, NJ, USA

Email: arothstein@princeton.edu

We present two novel machine-learning (ML) based controllers that actively avoid tearing modes (TMs): the first controller is the first reinforcement learning (RL) controller deployed on DIII-D's ITER baseline scenario (IBS) and the second controller combined ML predictions with actuator sharing to balance multiple control objectives. We developed the first-of-its-kind RL controller to actively track a stable operating path through a dynamic plasma shot while maintaining high plasma confinement by utilizing an offline training environment with a multimodal dynamic model to predict the likelihood of TMs. A follow-up controller directly used the TM likelihood within the time-evolving shot to divide the electron cyclotron heating (ECH) power between active TM control and off-axis current drive for scenario optimization. Both newly deployed controllers demonstrated improved plasma confinement by active TM avoidance and are promising options for advanced TM control in future tokamak reactors.

DIII-D IBS has been observed to be highly susceptible to TMs that will degrade plasma confinement or in the worst cases, disrupt the plasma for total loss in confinement. In DIII-D's IBS and other higher performance scenarios close to TM stability limits, active feedback control schemes are highly desirable for their ability to monitor and respond to changes in the plasma's TM stability to avoid onset of TMs.

Active control of TMs with Reinforcement Learning: Taking advantage of the flexibility provided by the offline training environment to predict TMs [1], we trained multiple RL controllers differing levels of cautiousness to balance TM avoidance with β_N optimization [2]. The more cautious RL controllers were rewarded for low TM prediction metrics while the less cautious controllers allow less stable plasma to higher β_N values.

When deployed in experiment, the RL controller was able to maintain the "Tearability," the predicted TM stability, below a fixed threshold of 0.5 that was set at training time by making various adjustments to plasma shape and neutral beam power as seen in Figure 1. These feedback control adjustments took advantage of the nonlinear effects of multi-actuator control to balance TM stability and high plasma pressures. In this experiment, the RL controller was able to sustain the β_N values more efficiently with less NBI power than in the reference.

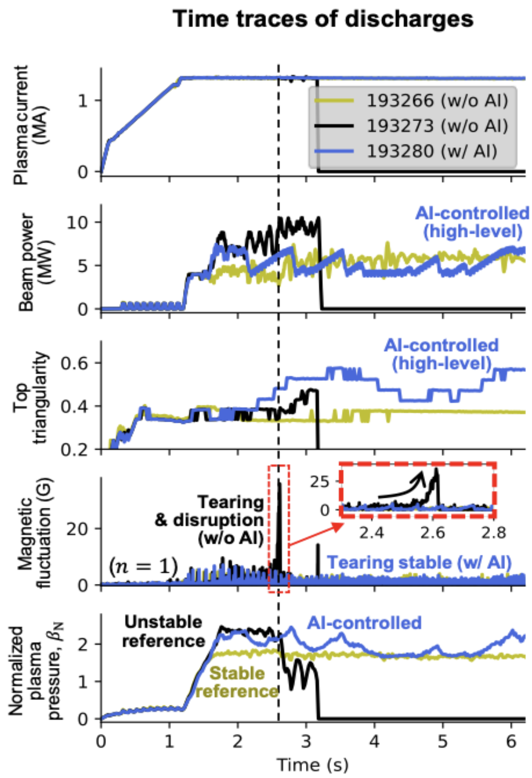

Figure 1. Time traces of various plasma parameters showing AI improvement (blue) over stable (green) and unstable (black) references. Beam power and top triangularity were the control actuators under feedback control. The magnetic fluctuations show there are no TMs in our controlled shot, and we see higher β_N than the stable reference shot.

Actuator sharing of ECH for TM control and scenario development: We deployed a second TM controller on DIII-D with the capability of sharing ECH between multiple goals [3] to align with ITER's planned ECH control system that will allow gyrotrons to flexibly switch between tasks within a single plasma shot. In our experiment the two goals were pre-emptive TM suppression and advanced scenario optimization where we used q=2 tracking for the former and broad off-axis current drive for the latter to maintain DIII-D's non-inductive advanced scenario. Taking advantage of TM warning times greater than 100ms [4], the controller pre-emptively steered ECH to the 2/1 rational surface to pre-emptively suppress TM growth and give more robust access to the high performance noninductive scenario.

With this control approach, we see in Figure 2 how our controller improved the reference scenario to avoid TMs that typically plague this scenario as seen in Figure 2c and 2d. Using preset thresholds in Figure 2a, each gyrotron task was assigned to TM suppression when TM probability crossed the threshold or off-axis CD when below the threshold and the effect on EC current drive is seen in Figure 2b. In these experiments, all plasma actuators were held constant except for the feedback-controlled ECH deposition locations. By balancing TM suppression and scenario optimization tasks, we pre-emptively suppressed TMs and maintained higher β_N values.

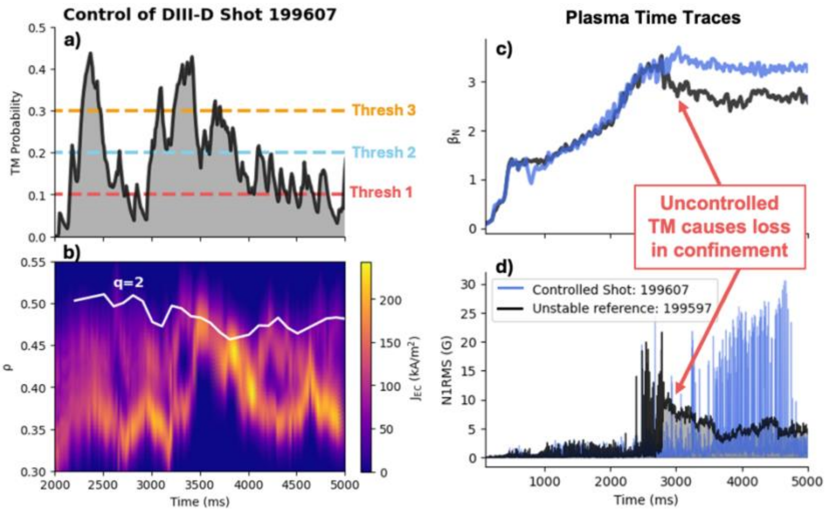

Figure 2. a) ML-based TM probability prediction with the preset threshold value for each gyrotron. b) EC current drive deposition profiles over time along with real-time q=2 location used for TM suppression target. c) β_N time traces from a shot with feedback control compared to one without. d) N1RMS signal to where saturation of TM is clear in the uncontrolled shot.

We have demonstrated the effectiveness of ML controllers in actively preventing TMs and gaining access to higher-performance plasma regimes. By leveraging the unique capabilities of these controllers—specifically multi-actuator feedback control and ECH actuator sharing—we successfully stabilized previously unstable reference shots by actively adjusting the plasma state. These innovative controllers have achieved high-performance plasmas, highlighting the significant potential of ML feedback controllers for future fusion reactors.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award DE-FC02-04ER54698. Additionally, this material is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2039656 and by the U.S. Department of Energy, under Awards DE-SC0015480.

REFERENCES

[1] J. Seo et. al. Multimodal prediction of tearing instabilities in a tokamak, IJCNN (2023) 1-8.

[2] J. Seo et. al. Avoiding fusion plasma tearing instability with deep reinforcement learning, Nature **626** 8000 (2024) 746-751

[3] A. Rothstein et. al. Pre-emptive tearing mode stabilization with multi-tasking ECCD, Nucl. Fusion (in preparation)

[4] H.J. Farre-Kaga et. al. Physics analysis of machine learning tearing mode predictions, PoP (in preparation)

DISCLAIMER