Princeton and DIII-D: AI to avoid tearing instabilities

IMPLEMENTED 2024–2025 · ENERGY AND UTILITIES · US · BAND D · PARTLY PROVED
VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

A Princeton-led team trained a reinforcement learning controller to keep a tokamak plasma away from tearing instabilities and ran it in experiments on DIII-D in San Diego in 2024. A 2025 preprint describes a wider AI control system now deployed on the machine.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

Frontier case: AI used for a scientific, medical or engineering breakthrough. The transformation is potential and unproven; MarLuc grades what has changed in practice so far.

AI used: Deep reinforcement learning controller with a neural network instability predictor.

Implemented: 2024–2025. Grade: band D, Task support, Partly proved.

Headline result, as reported: In DIII-D experiments the AI controller avoided tearing instabilities in all but two controlled discharges, in a single session (Nature, 2024)

1 · The case

The facts, as reported by the 3 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Princeton University and Princeton Plasma Physics Laboratory, with the DIII-D National Fusion Facility. Research team at Princeton University and Princeton Plasma Physics Laboratory (PPPL), a US Department of Energy laboratory, with a co-author at Chung-Ang University in Seoul. Experiments ran on DIII-D, a DOE user facility located at General Atomics in San Diego. Canvas block: 6. Data and technology. Also 4. Operations and process.

Business challenge

The challenge. Tearing instability is the leading cause of plasma disruptions in tokamaks, per the paper, and is hard to forecast. The authors say it is a particular obstacle in the ITER baseline scenario, and that disruptions can damage the machine's components.

How it was addressed. The team trained a neural network on past DIII-D data to estimate the chance of a future tearing instability, then used it as a simulated environment for reinforcement learning. The resulting controller adjusts neutral beam power and plasma shape to keep that chance below a set threshold. It was tested in DIII-D discharges, where the machine's existing control system turned its commands into coil currents and individual beam powers. In November 2025 members of the group published a preprint on a control framework, PACMAN, deployed on DIII-D for end-to-end machine learning control experiments. [1] [2] [3]

Who implemented it

Princeton University and Princeton Plasma Physics Laboratory (University or research lab): Built and ran the models and experiments.
DIII-D National Fusion Facility, General Atomics (University or research lab): Hosted the tokamak experiments.
The work was supported by the US Department of Energy, per the paper.

Technology (SEAM pillar E)

SEAM types: T5 Custom / fine-tuned models, T8 Predictive / ML models. Deployment: Not reported. Vendors and products: Tearability prediction model, Reinforcement learning controller, PACMAN control framework.

Where it was built in: The controller sets beam power and plasma triangularity targets; DIII-D's plasma control system derives coil currents and individual beam powers.

What changed

Before: preprogrammed control that did not avoid tearing at the targeted conditions. After, in experiments: an AI controller adapts beam power and shape to the predicted risk. Routine use on DIII-D or any power plant: Not reported. [1] [2] [3]

What it cost

Not reported

What it saved

Not reported. The paper's results are in avoided disruptions in one experiment session, not in time or money. [1] [2] [3]

Governance and control

Researchers arbitrate between what the AI wants and what the tokamak can accommodate, by a PPPL scientist's account, and the controller's actuator ranges were limited to the experiment's conditions. Accountable owner and independent assurance: Not reported.

S2 · Accountable owner · Not reported
Not reported.

A1 · Safe use · Not reported
Not reported.

A2 · Human in the loop · Claimed
A PPPL scientist says staff smooth the model's requested changes and arbitrate between what the AI wants and what the tokamak can accommodate. [2]

A3 · Monitoring and assurance · Not reported
Not reported.

A4 · Vendor AI · Not reported
Not reported.

A5 · Agent controls · Claimed
The paper says the controller's high-level commands go through DIII-D's plasma control system and the controllable ranges were limited to match the experiment. [1]

A6 · Standards and regulation · Not reported
Not reported.

What worked, what did not, what is unproven

WORKED

– The model forecast tearing instability up to 300 ms ahead, enough time for the controller to change beam power and shape, per Princeton.

– In the session the paper reports, all shots but one disrupted before the AI controller was applied and two disrupted after.

– The controller kept predicted tearing likelihood below its threshold at low safety factor and low torque while holding high-confinement mode.

DID NOT

– One run with a conservative setting failed to avoid a large instability because the beam power hit its preset lower bound.

– A second controlled discharge also disrupted.

– The first author said the group has strong evidence for DIII-D but needs more data for other situations.

UNPROVEN

– Whether the controller works on other tokamaks, with other heating systems and in other plasma scenarios.

– Whether it holds up in ITER-scale devices; the authors say further verification is needed when actuators and conditions change.

– Independent replication: the data are available only on reasonable request to the author.

Timeline

Feb 2024 Nature paper published: reinforcement learning controller tested in experiments on DIII-D [1]

Feb 2024 Princeton and PPPL describe the work as a proof of concept and say more data are needed for other situations [2]

Nov 2025 Preprint describes a control system, PACMAN, deployed on DIII-D for machine learning control experiments, including a state-machine tearing mode predictor-controller [3]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: A. Grade: D, Task support. Proof: Partly proved. Level of change: Task support.

Task support (Level 1, band D), claimed as transformational (band A). One research session showed a working controller. Fusion is not yet commercial, so the transformation is potential and unproven.

Graded D, claimed A. The controller changed what operators could do in one experiment session on a research tokamak, and the group says an AI control system is now deployed there. Fusion power is not yet commercial and no reactor relies on this controller. A transformation of how fusion machines are run is potential and unproven.

The SEAM view and lessons learnt

The design pairs a learned model of risk with a controller that works within a conventional control system, which keeps the final actuation familiar to operators. The group is open about its limits: one session, one machine, one instability. The lesson for any organisation is to report failed runs next to successful ones and to treat one session as a pilot, not a rollout.

1. Train controllers in a learned model, then test them inside the existing control system.
2. Report failed runs and their causes alongside successes.
3. Treat a single-session result as a pilot and plan replication before relying on it.

Evidence

Score 2 of 4. Partly independent: a peer-reviewed paper and a university release, both by the research team; the 2025 work is a preprint. No independent replication.

Evaluation method. Peer-reviewed paper (Nature, 2024) with a university press release relayed by ScienceDaily and a later preprint; results come from one experiment session by the authors.

Status. Ongoing: the group describes further experiments at DIII-D and plans to extend the controller to other tokamaks and more instabilities.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/princeton-diii-d-tearing-instability-control

[1] Nature: Avoiding fusion plasma tearing instability with deep reinforcement learning
https://www.nature.com/articles/s41586-024-07024-9

[2] ScienceDaily, relaying Princeton University and PPPL: Engineers use AI to wrangle fusion power for the grid
https://sciencedaily.com/releases/2024/02/240221160337.htm

[3] arXiv: Enabling Integrated AI Control on DIII-D: A Control System Design with State-of-the-art Experiments (preprint)
https://arxiv.org/abs/2511.08818

Princeton University and Princeton Plasma Physics Laboratory, with the DIII-D National Fusion Facility website: https://www.pppl.gov

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 27 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/princeton-diii-d-tearing-instability-control

VERSION HISTORY

v1.0 15 November 2026 First published