AI-accelerated fusion: advancements and open challenges

MIT Plasma Science & Fusion Center

Cristina Rea
crea@psfc.mit.edu

SULI lecture June 4th, 2026

Special thanks to:

Andrew Maris, Allen Wang & GL Trevisan, Y Wei, Z Keith, S Benjamin, A Kumar the PSFC Disruptions team, General Atomics collaborators, and the SPC-EPFL team

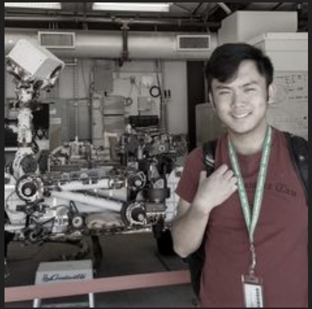

Work partly funded by DE-FC02-04ER54698, DE-SC0014264, DE-SC0024368, and CFS

Outline

1. Machine Learning and data-driven research in fusion
2. Problem setting: disruptions, instability prevention, and controllability
3. Explainable and adaptive Machine Learning for disruption prevention
4. Conclusions

Data drives fusion experiments' design, simulation, analysis, control and optimization

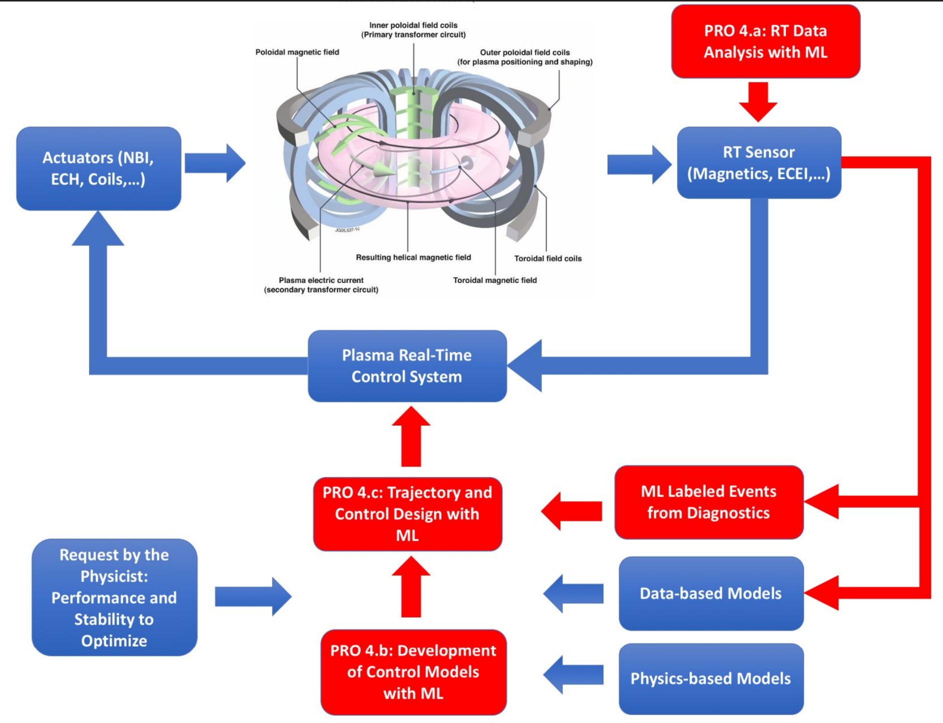

Adapted from D. Humphreys et al. "Advancing Fusion With Machine Learning" DOE Workshop (2020)

Focus on MFE applications

ML-applied research at the PSFC aimed at advancing scientific discovery via:

- Database development & data analysis
- Accelerating simulations via surrogates
- Predictive modeling (sim2real)
  - Optimizing operations
  - Trajectory planning and control

PSFC is the 1st IAEA Collaborating Centre in AI in Fusion and Plasma Science

IAEA Coordinated Research Project "AI for Fusion":
5yr agreement bridging MFE and IFE research teams internationally.

- **Open Science**, e.g. data/models openly accessible
- **FAIR** access to fusion data
  - Fusion Data Lake architecture being designed by IAEA
- Develop research products to enable **cross-pollination**
  - Plasma forecasting competitions & hackathons
- **Education**, **capacity building** in ML-applied fusion science

https://nucleus.iaea.org/sites/ai4atoms/ai4fusion/SitePages/AI4F.aspx

Upcoming event:

- **Second Workshop on AI for Accelerating Fusion Energy and Plasma Science** (Fall 2026) https://conferences.iaea.org/event/454/

Key challenges for trustworthy ML in fusion research

- **Benchmarking** → trusted performance comparison
- **Domain/data shift** → robust out-of-distribution prediction
- **Interpretability** → confidence in predictions and model behavior
- **Labeled data + metadata** → scalable training and evaluation
- **UQ, verification, and validation** → rigorous reliability assessment
- **Open, FAIR data and models** → reproducibility and faster progress

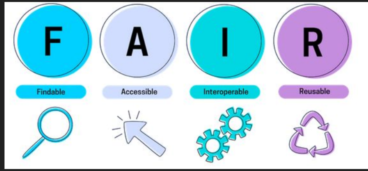

M. Wilkinson, *et al.* The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* **3**, 160018 (2016)

→ *Data availability – labeled, curated (FAIR) data – is the foundation for reliable ML in fusion*

Open source software and public datasets are crucial to develop robust ML solutions

DisruptionPy: an interoperable, open-source Python library for data access across different fusion devices. Enables assembly of big databases for ML applications.

https://github.com/MIT-PSFC/disruption-py

Trevisan (2026), J. Open Source Softw. 11 (119)
10.21105/joss.09364

C-Mod Open Density Limit Dataset: a curated C-Mod dataset related to density limits.

https://github.com/MIT-PSFC/open_density_limit_database

Tearing Physics Suite: a numerically robust nonlinear tokamak tearing analysis tool for large-scale database generation.

https://github.com/MIT-PSFC/tearing-physics-suite

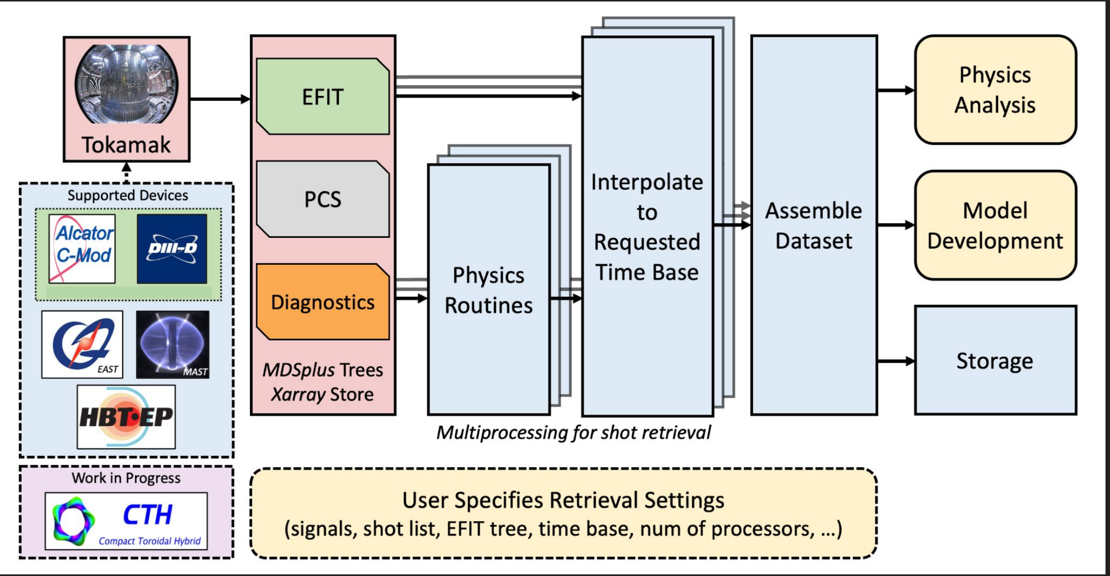

DisruptionPy: interoperable library for data retrieval and database development across different devices

- ✓ Modular structure + version controlled
- ✓ Adopted automated unit testing, linting, and deployment workflows
- ✓ Available for Alcator C-Mod, DIII-D, EAST, and now supporting MAST open data!

★ Trevisan, Rea and the MIT Disruptions team
https://github.com/MIT-PSFC/disruption-py
- 2025 Zenodo DOI: 10.5281/zenodo.13935223
- 2026 JOSS 11 (119) DOI: 10.21105/joss.09364

Machine	Methods	Features	Shots
C-Mod	22	62	35,000
DIII-D	20	59	75,000
EAST	31	82	40,000
HBT-EP	12	35	35,000
MAST	4	20	10,000

- ✓ Adopted multi-backend support to allow access to both *MDSplus* and *xarray* databases.
- ✓ Recomputing equilibria with finer time resolution to support analysis of faster timescale physics
- ❑ Continue expanding support to new devices e.g. CTH & HBT-EP
- ❑ Integrate physics and ML-driven event labeling workflows into data pipeline, and adopt *MLOps* best practices
- ❑ **Release curated datasets by end of 2026**

Data-driven learning for disruption prevention and performance optimization

Overview of ML workflow

- Disruption Prediction (DPRF)
- Disruption Precursors' modeling
- Time-to-event frameworks

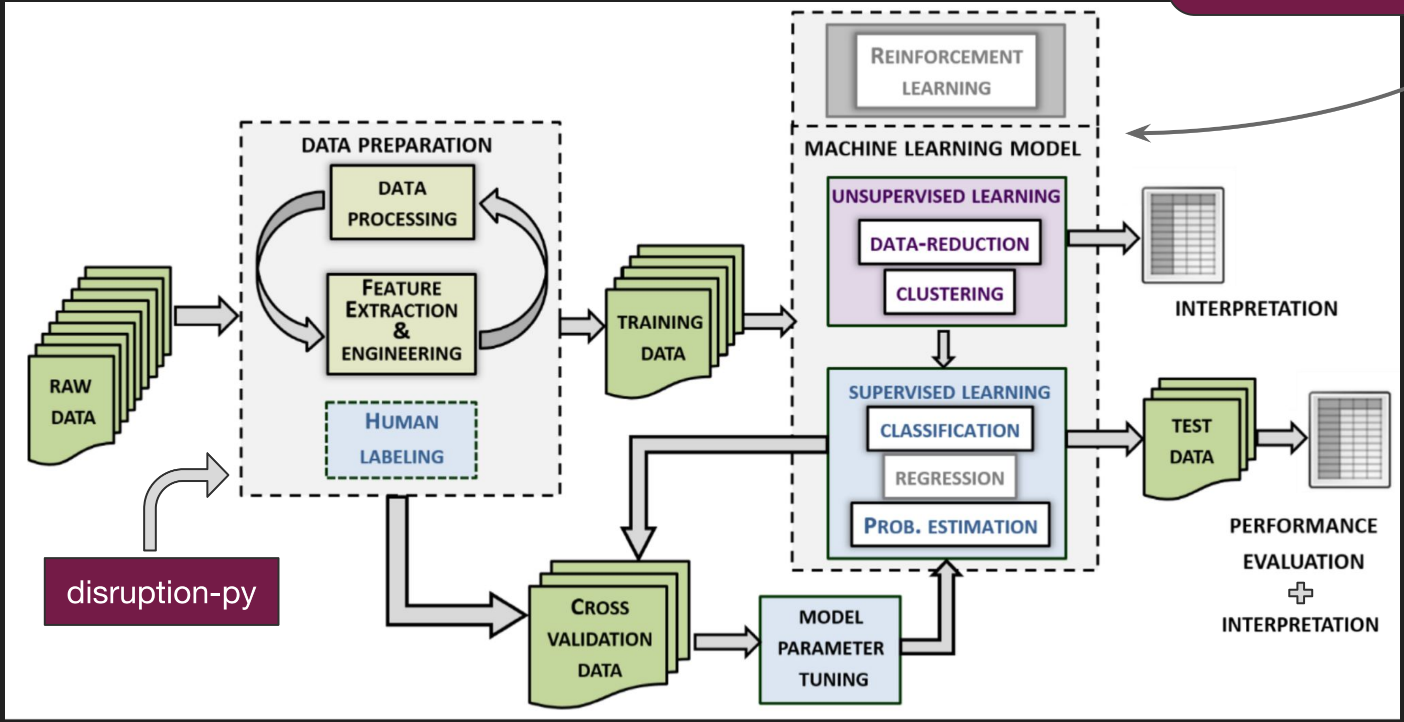

Adapted from A. Pau et al, Nuclear Fusion, 59(10):106017, 2019

Outline

1. Machine Learning and data-driven research in fusion
2. Problem setting: disruptions, instability prevention, and controllability
3. Explainable and adaptive Machine Learning for disruption prevention
4. Conclusions

Plasma pushed close to operational limits → instabilities onset or control faults: disruptions

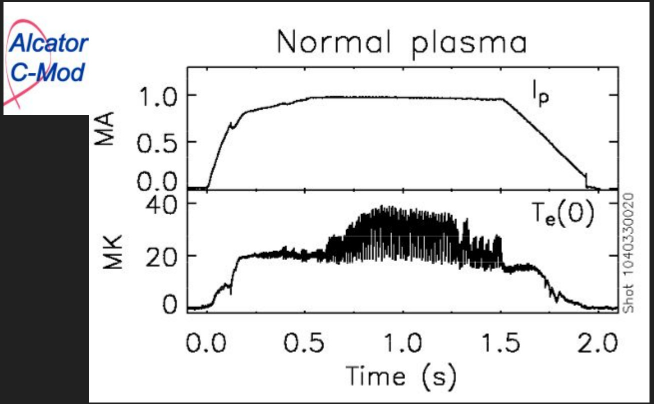

Disruptive plasma:

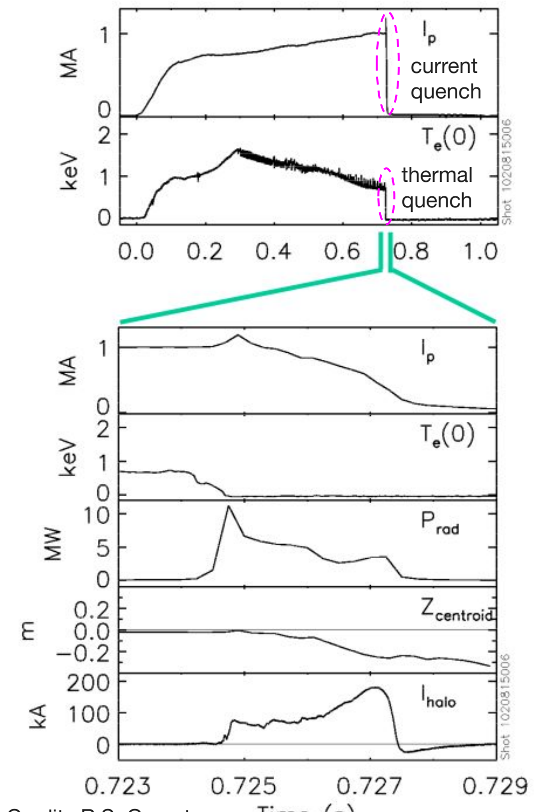

Credits R.S. Granetz

Final loss of control evolving on timescales of milliseconds:

- Fast drop I_p leads to loss of confining poloidal field
- Fast I_p transient leads to large induced voltages/currents/forces
- Rapid thermal losses leads to surface damage

What a disruption looks like:

Visible camera view of RE beam hitting Alcator C-Mod first wall. Courtesy R.A. Tinguely

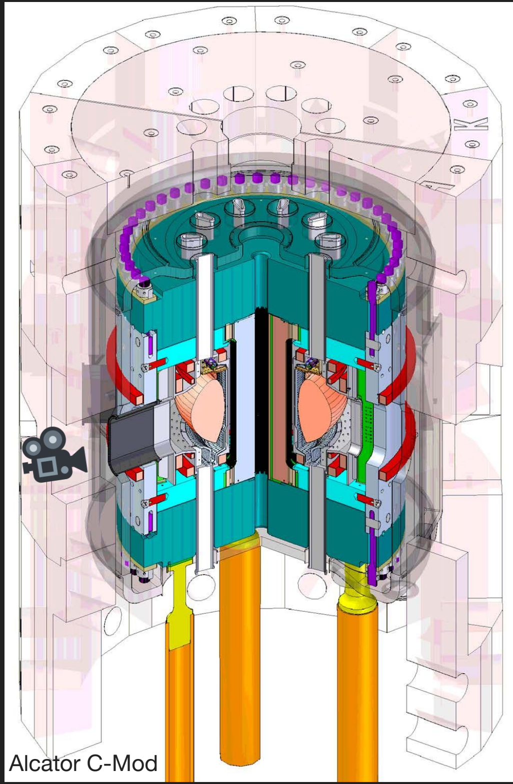

Rapid energy loss during disruptions causes strong electric field → runaway electron beam

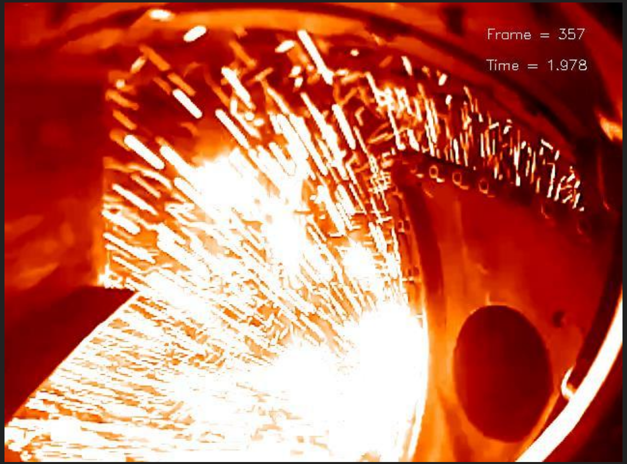

Visible camera view of RE beam hitting Alcator C-Mod first wall. Courtesy R.A. Tinguely

melting!

JET runaway electrons damage.
https://www.iter.org/newsline/-/2234

Accept the damage and live with it.

Mitigate the damage by injecting massive gas (or shattered pellets or passive coils).

Avoid altogether by detecting precursors & steer plasma away from disruptive boundary.

Next gen devices disruption rate < 1% !

Precursor's prediction and detection crucial to avoid disruptions, but challenging task!

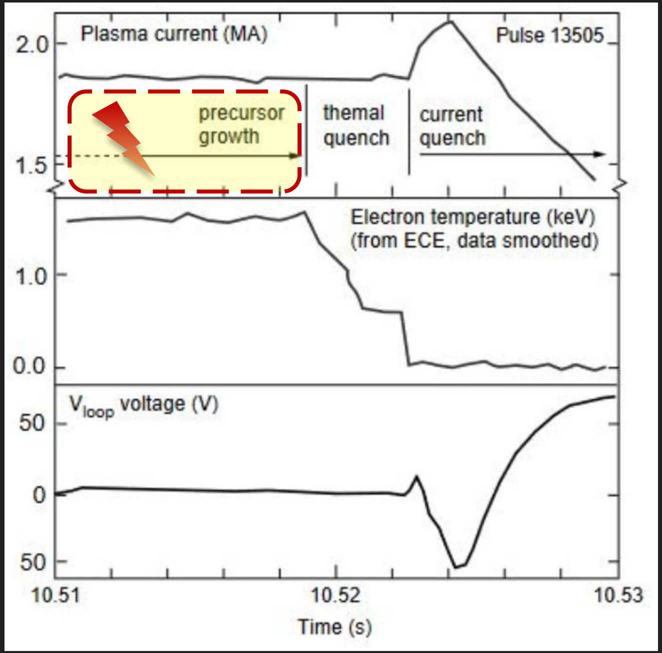

ITER Physics Expert Group on Disruptions, Plasma Control, and MHD (1999) Nucl. Fusion 39 2251

Precursor's prediction and detection crucial to avoid disruptions, but challenging task!

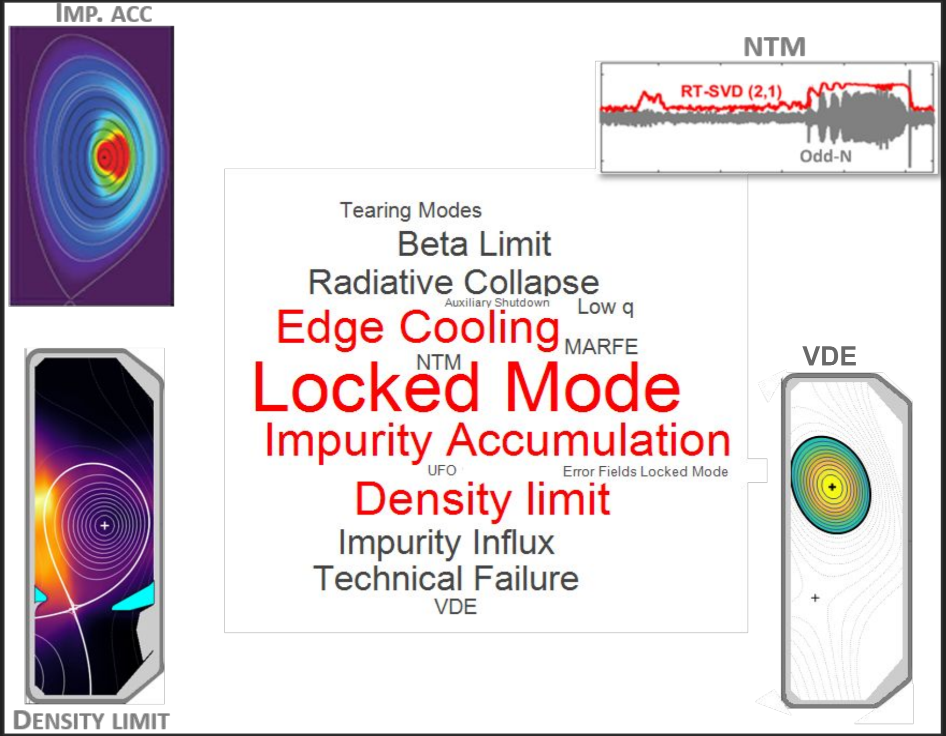

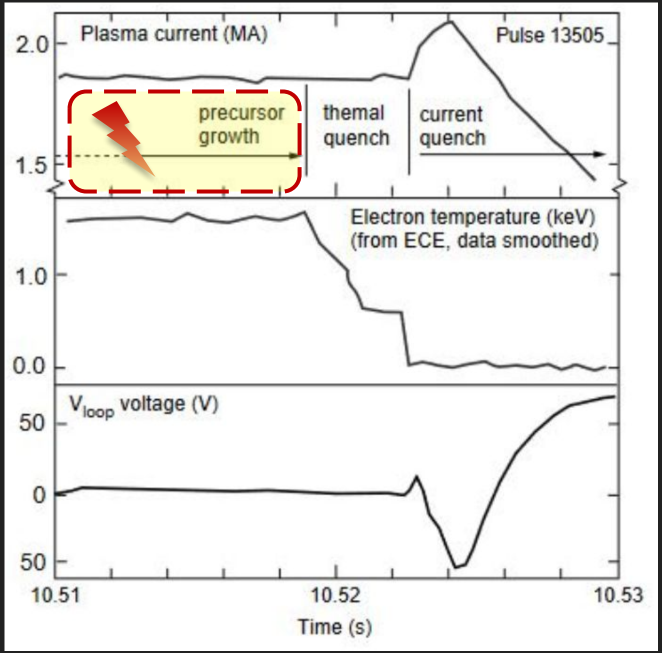

ITER Physics Expert Group on Disruptions, Plasma Control, and MHD (1999) Nucl. Fusion 39 2251

Plasma Science & Fusion Center

Courtesy: A. Pau, PSFC seminar, April 2023

De Vries et al. NF 51 (2011) 053018 "Survey of disruption causes at JET"

No first-principle model to capture full pre-disruptive chain of events

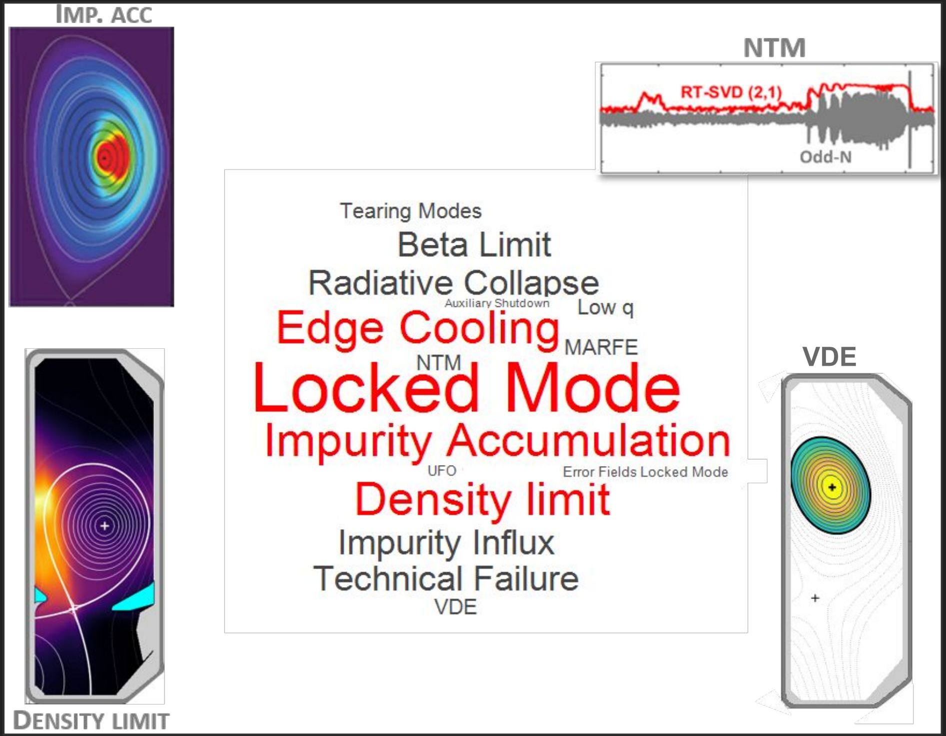

- Statistical studies on disruption frequency (and type) not available across different tokamaks.
- Notable efforts on event analyses:
  KSTAR, MAST, NSTX/-U → **DECAF**[1]
  JET, TCV → **DEFUSE**[2]

  [1]Sabbagh et al PoP 30 032506 (2023)
  [2]Pau et al 29th IAEA-FEC (2023)
- Wealth of experimental data from different tokamaks enables "**learn from data**" paradigm → ML
- The timely identification of **precursors** allows the plasma control system (PCS) to take proper **avoidance** action → extend discharge lifetime → make fusion!

Plasma Science & Fusion Center

Courtesy: A. Pau, PSFC seminar, April 2023
De Vries et al. NF 51 (2011) 053018 "Survey of disruption causes at JET"

Outline

1. Machine Learning and data-driven research in fusion
2. Problem setting: disruptions, instability prevention, and controllability
3. Explainable and adaptive Machine Learning for disruption prevention
4. Conclusions

Large portfolio of ML/AI applications in fusion – broader literature available

- Explainable and adaptive Machine Learning for disruption prevention
  - Data labeling and augmentation
  - DPRF + precursors' modeling
    - Density limit
    - Confinement regime identification
    - UFO disruptions
  - Deep Learning modeling
    - Scenario-adaptive experiments
    - Multi-task learning
    - Transformers and CCNN
  - Time-to-event frameworks
    - Survival analysis
    - Neural State Space Models + RL

- Montes, Rea et al 2021 Nuclear Fusion 61 026022
- Rath, ..., Rea et al, J. Plasma Phys. (2022), 88, 895880502

- Rea et al, Nucl. Fusion 59 (2019) 096016
- Rea et al 2021, 28th IAEA FEC, EX/P1-25
- Benjamin, Rea et al. (2024) PPCF 66 075016
- Benjamin, Rea et al. (2026) Nucl. Fusion 66 036049 doi:10.1088/1741-4326/ae44ad
- Maris, Rea et al 2025 Nucl. Fusion 65 016051 *doi:10.1088/1741-4326/ad90f0*

- J.X. Zhu, Rea et al, 2021 Nucl. Fusion 61 026007
- J.X. Zhu, Rea et al, 2021 Nucl. Fusion 114005
- J.X. Zhu, Rea et al, 2023 Nucl. Fusion 63 046009
- Spangher, Rea, Eni+MIT partnership 2025 JoFE https://doi.org/10.1007/s10894-025-00495-2

- Tinguely, Montes, Rea et al 2019 PPCF 61
- Keith, Rea et al 2024 JoFE doi:10.1007/s10894-024-00413-y
- Boyer, Rea, Clement, Nucl. Fusion 62 (2022) 026005 doi:10.1088/1741-4326/ac359e
- Wang, Rea et al, 2025, Nature Comm. Physics *https://www.nature.com/articles/s42005-025-02146-6*
- Wang, Pau, Rea et al, 2025, Nature Comms. https://doi.org/10.1038/s41467-025-63917-x

Outline

1. Machine Learning and data-driven research in fusion
2. Problem setting: disruptions, instability prevention, and controllability
3. Explainable and adaptive Machine Learning for disruption prevention
    - Disruption Prediction (DPRF)
    - Disruption Precursors' modeling
    - Time-to-event frameworks
4. Conclusions

Available in real-time at DIII-D

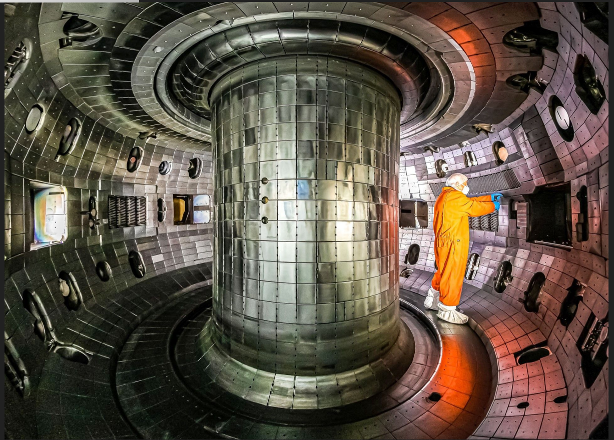

Explainable ML for disruption prediction and stability boundaries identification in real-time

DPRF

- **D**isruption **P**rediction via **R**andom **F**orest (DPRF) computes probability of an impending disruption, while interpreting its drivers in real-time.
  - Available on DIII-D (and EAST).

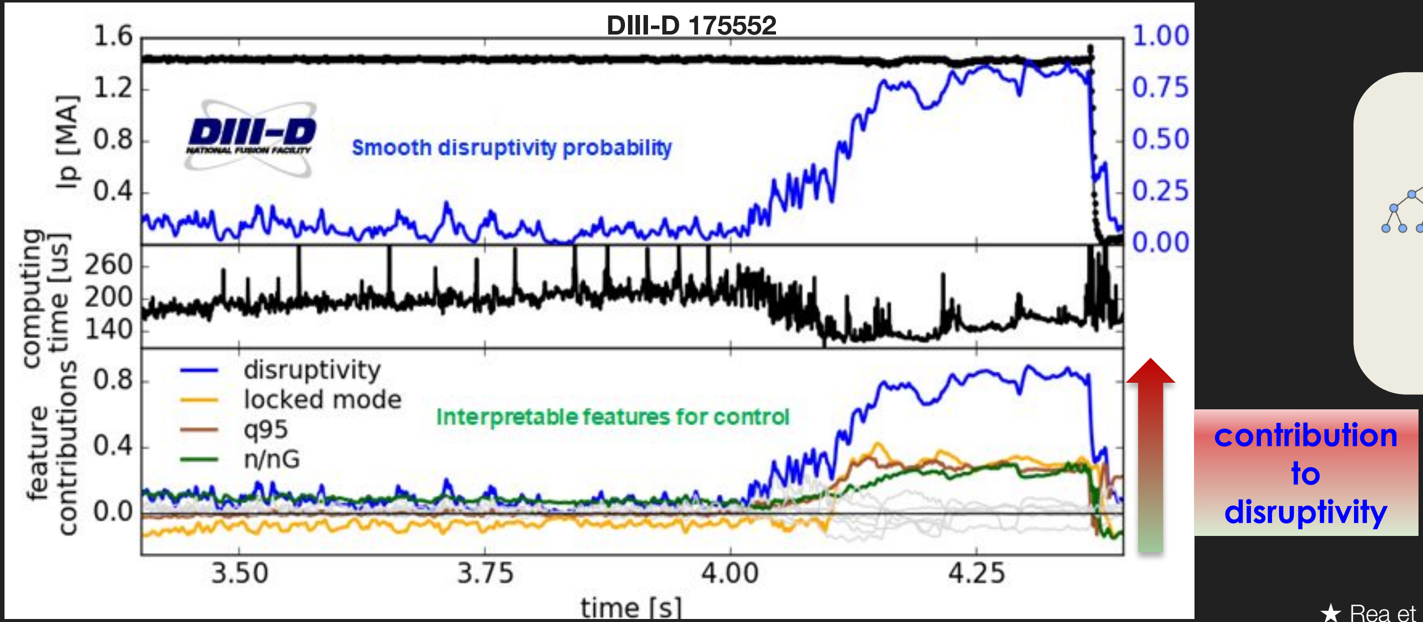

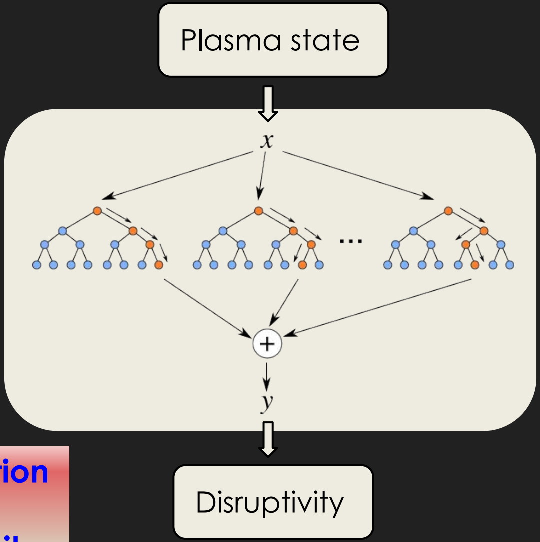

★ Rea et al, Nucl. Fusion 59 (2019) 096016
★ Rea et al, 2021 IAEA EX/P1–25,
J. Barr et al, Nucl. Fusion 61 (2021) 126019

Real-time implementation and optimization for asynchronous avoidance and emergency response

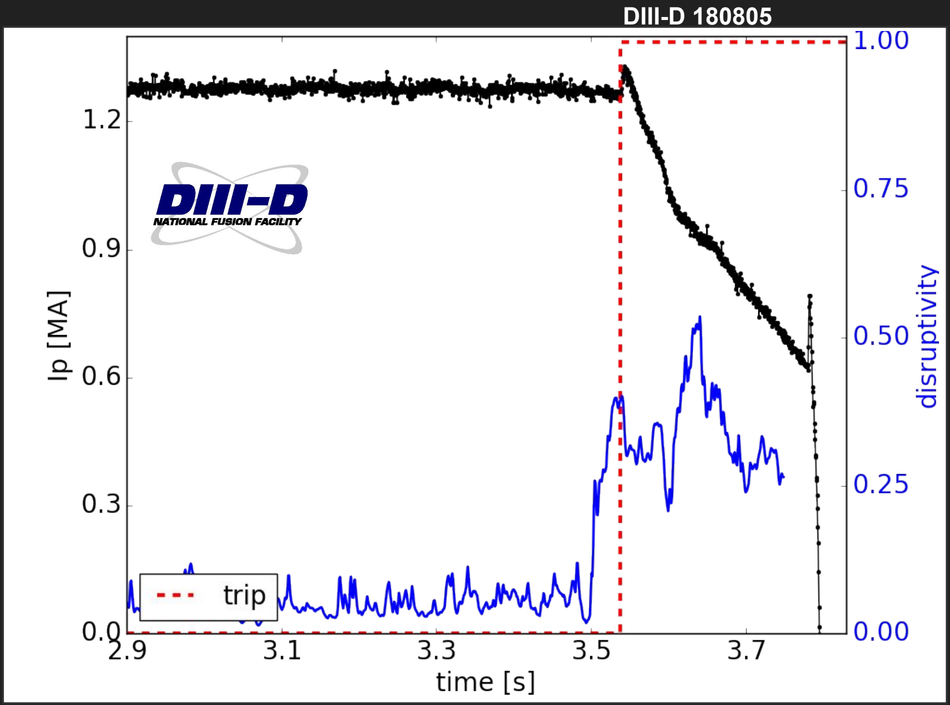

★ Rea et al, 2021 IAEA EX/P1–25

Real-time model evaluation and feature contribution computation: < 200 μs

Successful ONFR integration:

- Fast shutdown triggered by preset disruptivity threshold.
- MGI response and ECH trigger in closed loop.

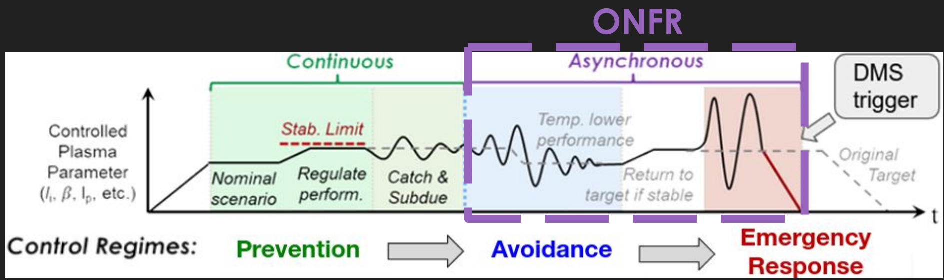

Off-Normal Fault Response → **Asynchronous and Emergency response**.

Eidietis et al., 2018 Nucl. Fusion 58 056023

DPRF + Proximity Controller to continuously regulate plasma stability and performance

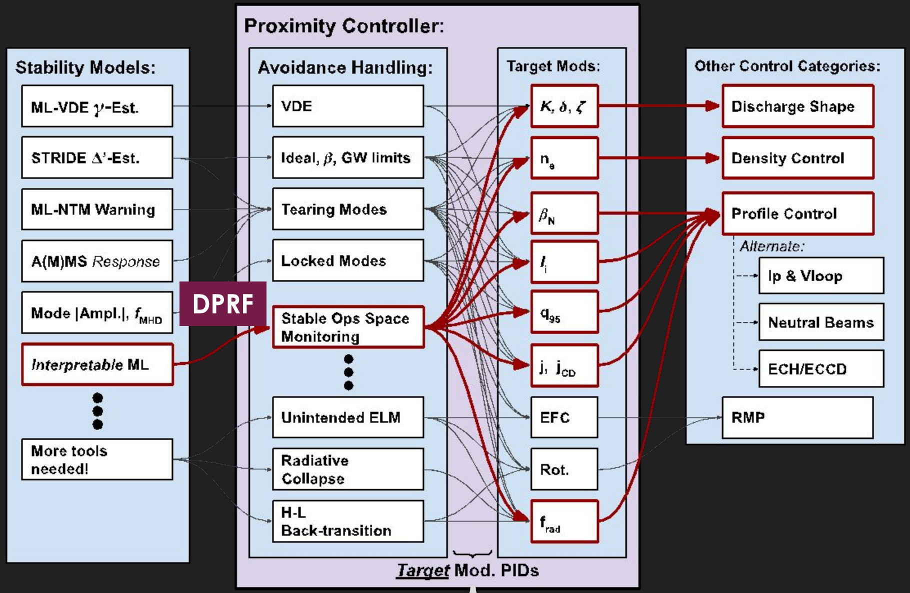

$$\Delta\kappa = PID\left[f_{danger} * f_{\kappa,contrib} * sign\left(\frac{d\kappa}{dt}\right)\frac{\Delta\kappa_{target}}{\Delta f_{\kappa,contrib}}\right]$$

- Disruptivity measures proximity to unstable operating space
- Feature contributions can be mapped onto controllable plasma parameters to regulate stability

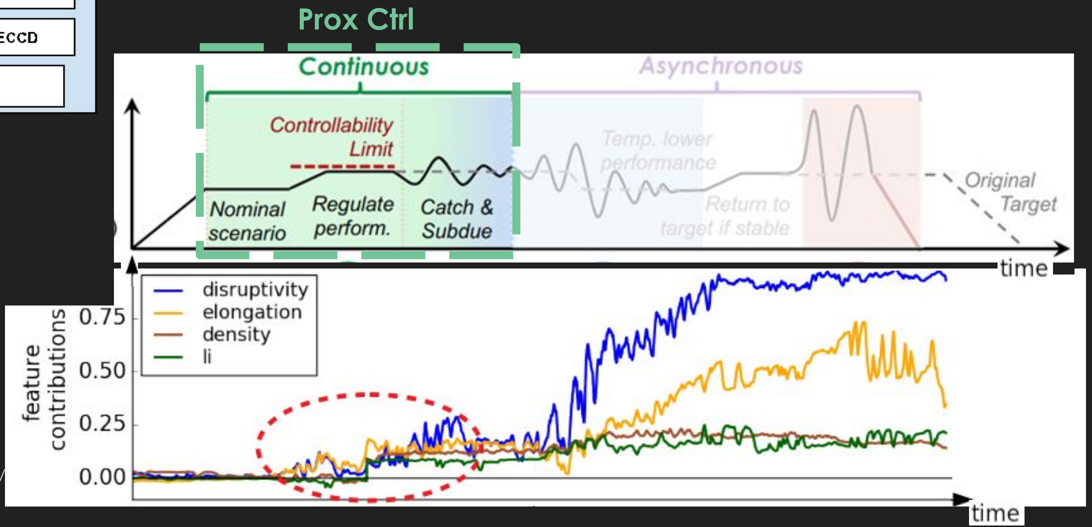

Barr et al, Nucl. Fusion 61 (2021) 126019

★ Rea et al 2021, 28th IAEA FEC, EX/P1-25

Rea /

DPRF + ProxCtrl tested in exp ~5 weeks ago

- Retrained DPRF with high betap / SVR scenario data. DisruptionPy enabled the quick retraining.
- PID control of target actuators (beams, gas) using feature contributions obtained from SHAP analysis on DPRF.
- DIII-D experiment March 24, 2026 as part of the 2025 AI Task Force. Data analysis is currently in progress.

Experiment led by Y. Wei, A. Kumar, J. Barr
Manuscript under preparation, 2026

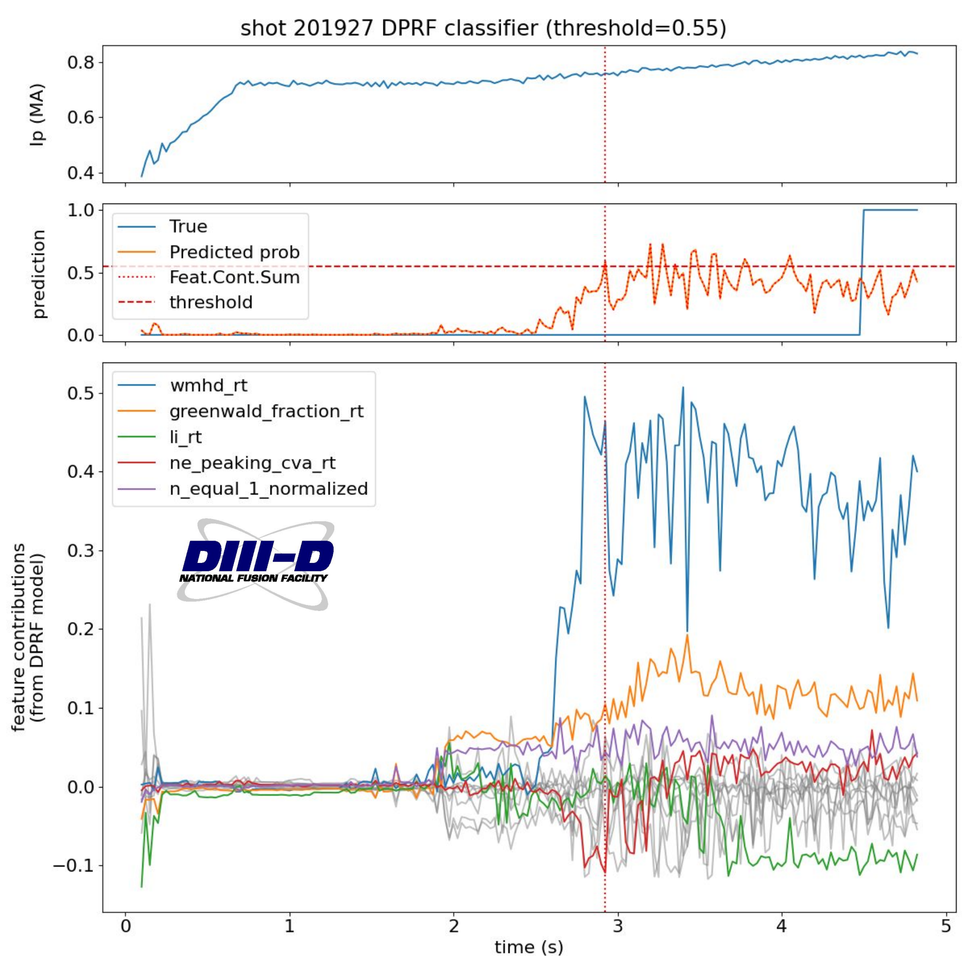

Analysis of dominant features & contributions on target shot 201927

DPRF: "gray box" + post-hoc explanation for generic disruption prediction

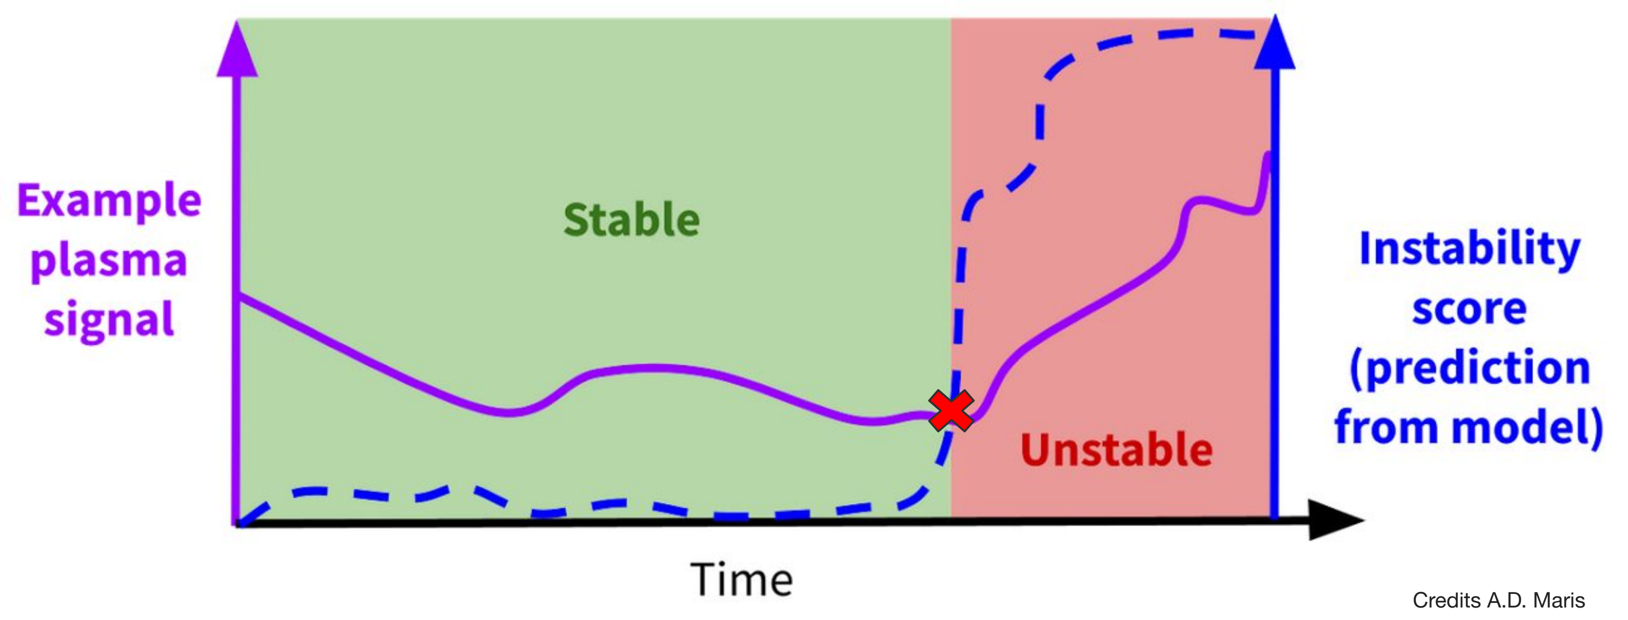

Shift focus on interpretable models for precursors' onset

Symbolic ML for disruption-free operations in regimes close to density-limit

density limit

Density is critical design parameter, fusion power density ∝ n^2

Power plant design points usually close to empirical density limit

$$\frac{n}{n_G} < 1, \text{ where } n_G = \frac{I_p}{\pi a^2}$$

Greenwald et al., Nucl. Fusion (1988)

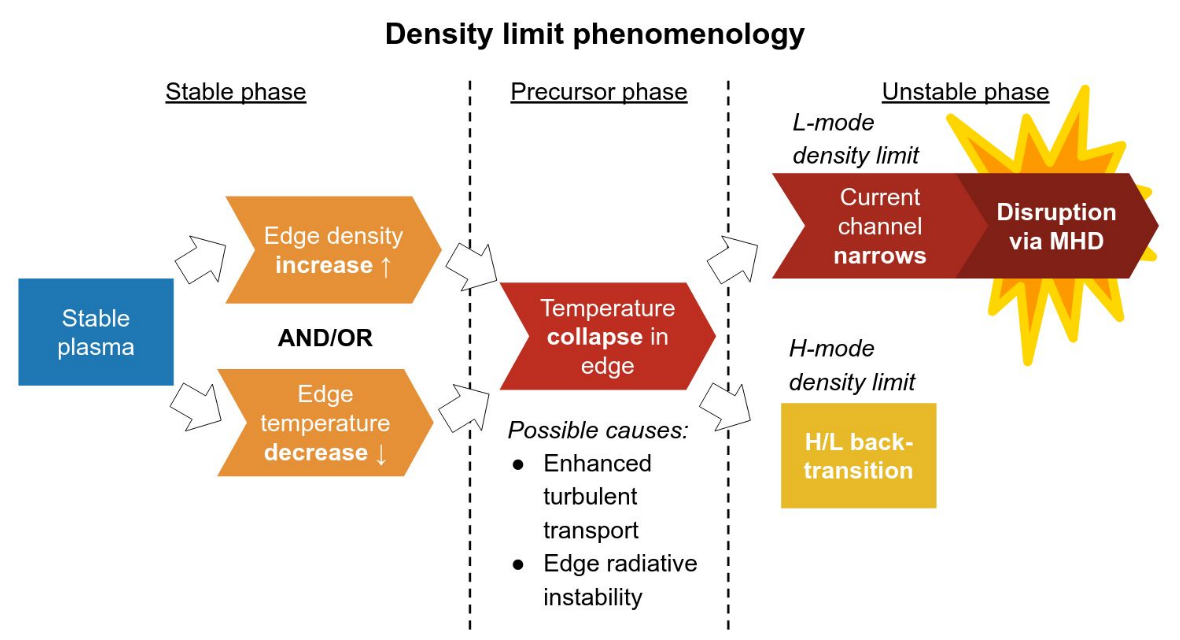

Broad literature & multi-machine data available.

Manz et al 2023 Nucl. Fusion 63 076026;
Zanca et al 2019 Nucl. Fusion 59 126011;
Giacomin et al 2022 Phys. Rev. Lett. 128 185003;
Singh and Diamond 2022 Plasma Phys. Control. Fusion 64 084004;
Stroth et al 2022 Nucl. Fusion 62 076008;
Brown and Goldston 2021 Nucl. Mater. Energy 27 101002;
Bernert et al Plasma Phys. Control. Fusion 57 (2015) 014038;
Maraschek et al Plasma Phys. Control. Fusion 60 (2018) 014047;
...

Interpretable tools needed to understand & avoid density limit events.

★ Maris, Rea et al 2025 Nucl. Fusion 65 016051
doi:10.1088/1741-4326/ad90f0

Recursive feature selection + ML-driven boundary id enables interpretable stability metric

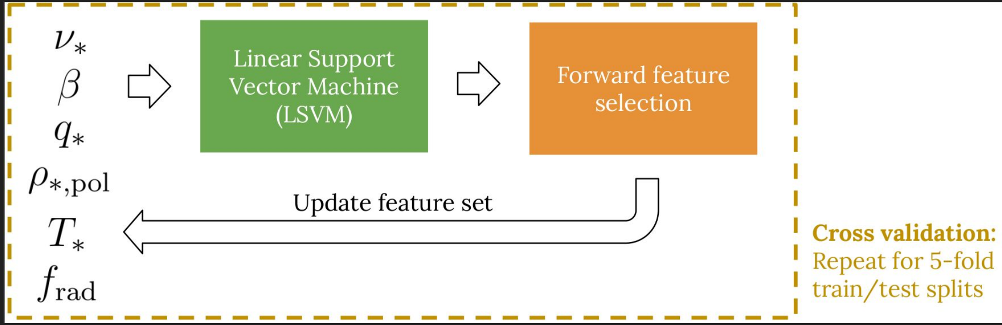

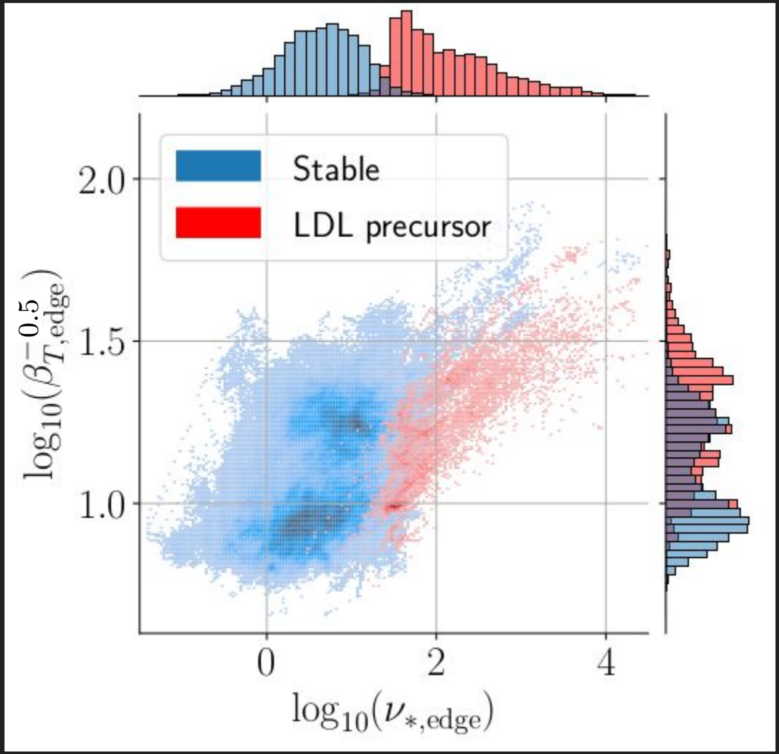

- L-mode density limit (LDL) database developed using:
  Alcator C-Mod, DIII-D, AUG, TCV
- LDL precursor risk metric: $\boxed{\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.5}}$

★ Maris, Rea et al 2025 Nucl. Fusion 65 016051
doi:10.1088/1741-4326/ad90f0

Recursive feature selection + ML-driven boundary id enables interpretable stability metric

- L-mode density limit (LDL) database developed using:
  Alcator C-Mod, DIII-D, AUG, TCV

- LDL precursor risk metric: $\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.5}$

→ yields 4x fewer false positives than Greenwald ✅

- H-Mode analyses expanded to 5 tokamaks (348 DL events)

Greenwald fraction **Risk metric:** $\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.5}$

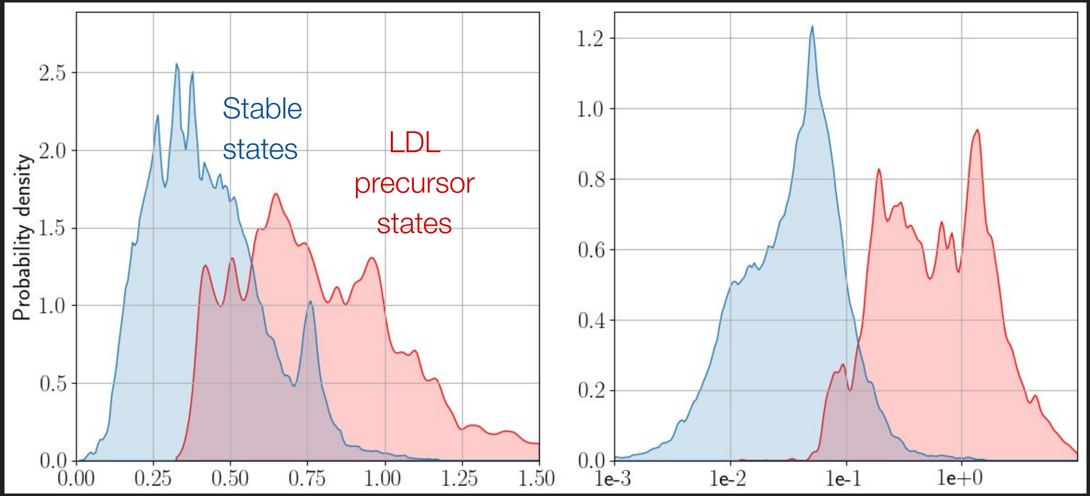

★ Maris, Rea et al 2025 Nucl. Fusion 65 016051
[doi:10.1088/1741-4326/ad90f0](https://doi.org/10.1088/1741-4326/ad90f0)

LSVM+sequential feature selection → Risk metric integrated in real-time DIII-D PCS

density limit

samples: ~4000 stable shots (1.5e6 timeslices), ~90 H-mode density limit events (768 HDL time slices),
Features: ~20
Size full database: 2.5 GB

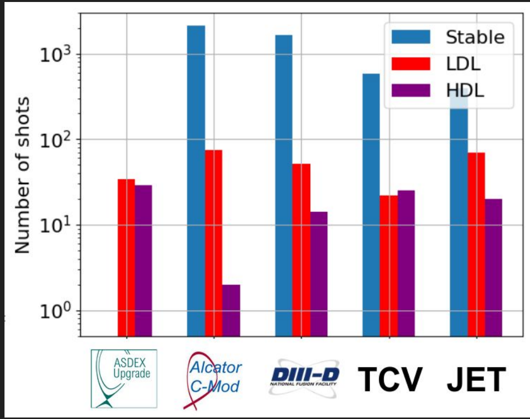

Greenwald fraction

Risk metric: $\dfrac{\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.9}q_*^{1.1}}{\rho_{*,\mathrm{edge}}^{0.8}}$

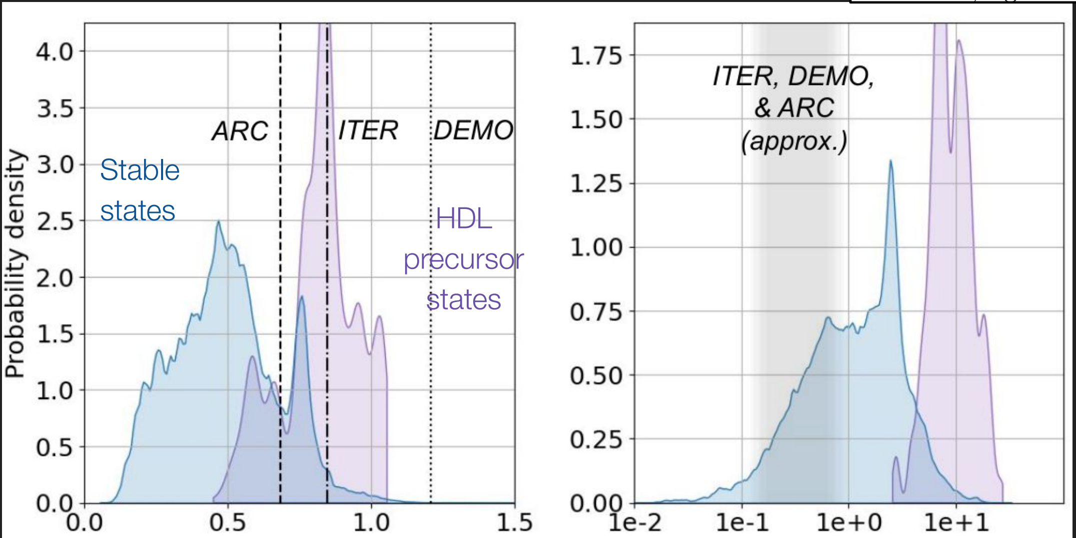

→ yields 2x fewer false positives than Greenwald ✅

★ Maris, Rea et al 2026 NF 66 056002 "Real-time avoidance of the L-mode and H-mode density limit via machine-learned stability metrics" doi:10.1088/1741-4326/ae4efe

Real-time trajectory steering accomplished via Proximity Controller integration

- Goal: get as close to original targets as possible while extending discharge lifetime

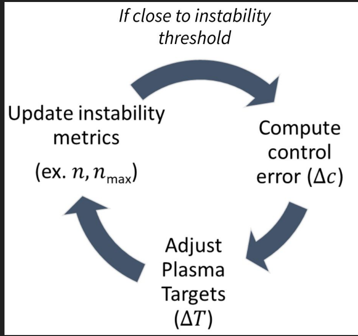

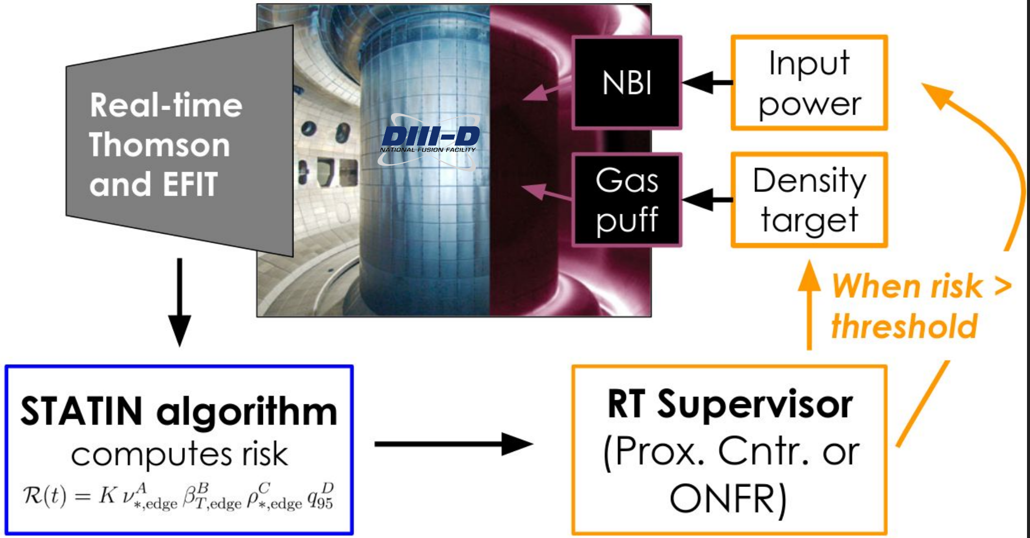

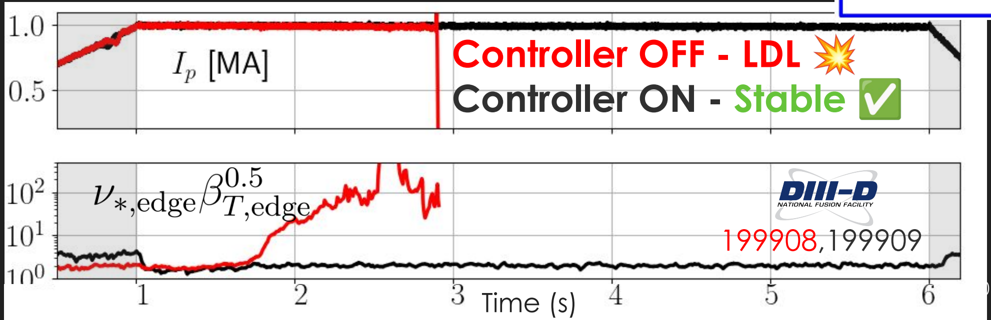

Reproducible results, varying actuators' response

★ Maris, Rea et al 2026 NF 66 056002 doi:10.1088/1741-4326/ae4efe

Interpretable ML leverages multi-machine database to extract new density limit (DL) scalings

L-mode DL

$$\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.49}$$

H-mode DL

$$\frac{\nu_{*,\mathrm{edge}}\beta_{T,\mathrm{edge}}^{0.88}q_*^{1.09}}{\rho_{*,\mathrm{edge}}^{0.80}}$$

≥2x fewer false positives than Greenwald fraction ✅

- **Collisionality scalings utilized for real-time DL avoidance**

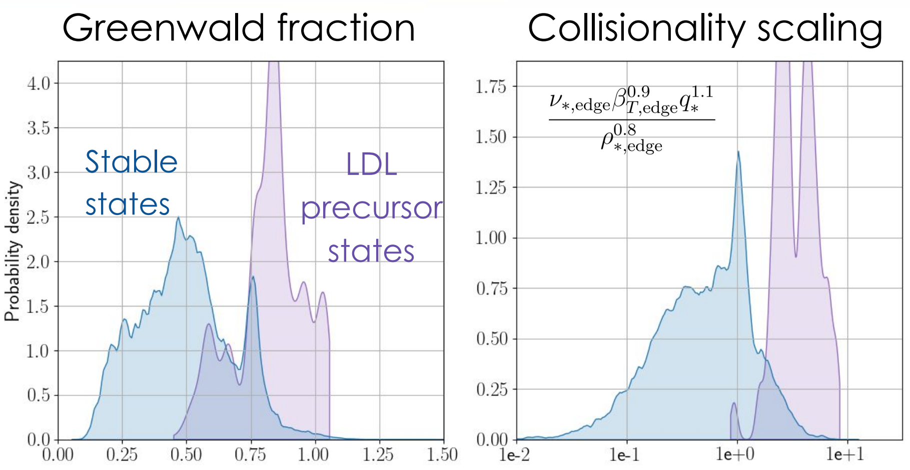

Control metric:
$$\frac{\nu_* q_{95}\beta_T^{0.8}}{\rho_*^{0.6}}$$

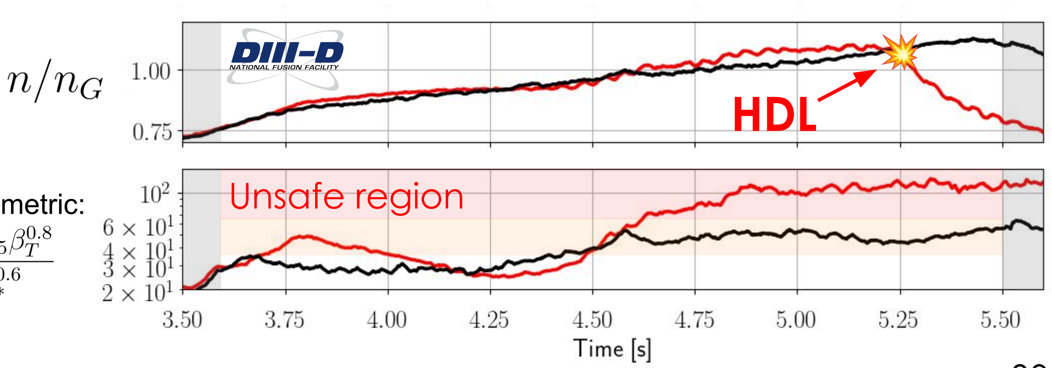

DIII-D NATIONAL FUSION FACILITY

Outline

1. Machine Learning and data-driven research in fusion
2. Problem setting: disruptions, instability prevention, and controllability
3. Explainable and adaptive Machine Learning for disruption prevention
   - Disruption Prediction (DPRF)
   - Disruption Precursors' modeling
   - Time-to-event frameworks or rather, *trajectory design**
4. Conclusions

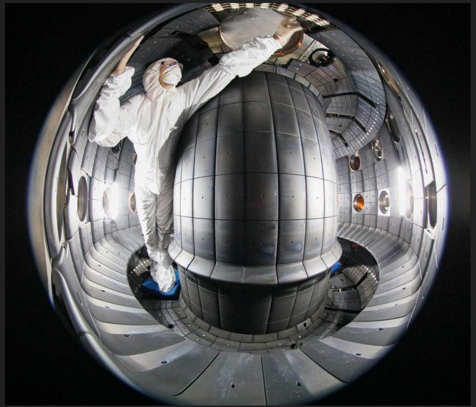

TCV experiments

**slides freely adapted from Allen Wang's thesis defense (Jan 2026)*

Plasma Science & Fusion Center

Original Motivation & Overview: Ramp-down Control + Trajectory Optimization

- **Problem:** rampdowns can push the plasma closer to instability limits
- Some limits (e.g. Greenwald fraction) are related to kinetic dynamics (temperature and density)
- Motivates simulation of plasma temperature and density evolution in response to control knobs (extremely challenging from a pure physics perspective)

Wang et al., 2025 Nature Comms.

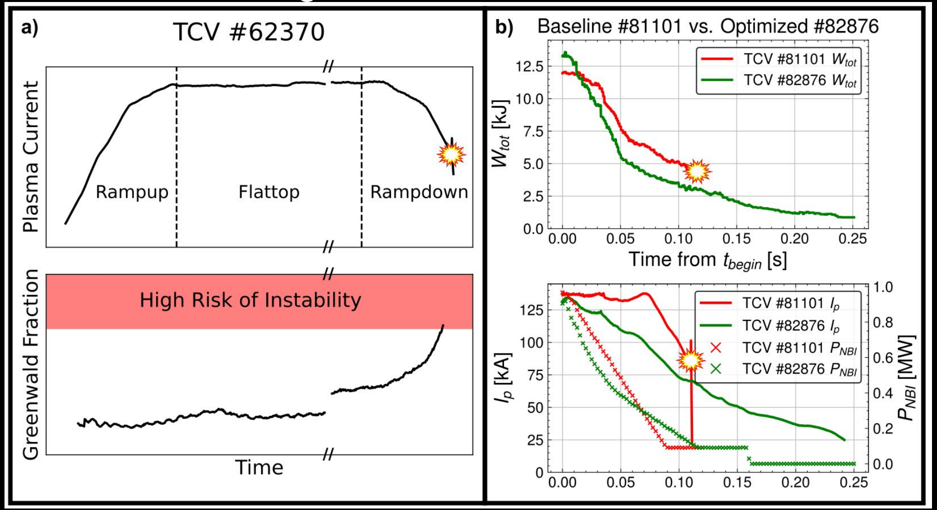

- **SciML kinetic dynamics model** → for rampdown trajectory optimization of a high-performance scenario (betaN~2, Greenwald Fraction ~ 0.8-1)
  - Only trained for rampdowns
  - Also predicts density dynamics and vertical instability growth rates

Neural State-Space Models

- For control, it's important to have actions and observations in the model structure
- Classically, often done with linear state-space models
  - **Many** applications to real-world systems, including aircraft
- Modern ML frameworks allow us to program NSSMs with arbitrary combinations of physics + networks

Linear State-Space Model

state actions

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{a}$$

observations

$$\mathbf{o} = C\mathbf{x} + D\mathbf{a}$$

Nonlinear State-Space Model or Neural State-Space Model (NSSM)

$$\dot{\mathbf{x}} = f_\theta(\mathbf{x}, \mathbf{a})$$

$$\mathbf{o} = O_\theta(\mathbf{x}, \mathbf{a})$$

SciML + JAX enable time-dependent model training

- NN f_θ used as system of differential equations integrated forward in time with a differential equations solver

- Adjoint back-propagation + automatic differentiation to determine the gradient of loss with respect to the network parameters θ

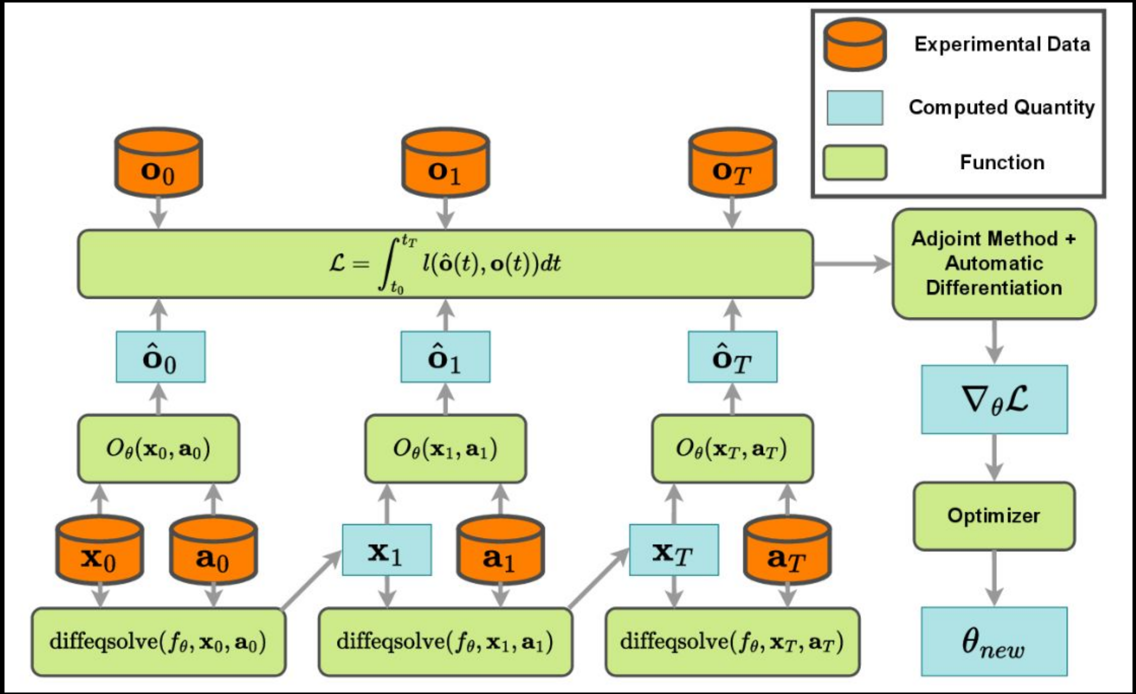

Plasma Science & Fusion Center

From Observables to Actuators and Considered Constraints

W_{tot}	Plasma total stored energy
$\langle n_e \rangle$	Line-averaged electron density
$\langle n_{g,frac} \rangle$	Greenwald fraction
β_p	Plasma poloidal beta
ι_{95}	Rotational transform at the 95% flux surface
γ_{vgr}	Vertical Growth Rate
$T_e(\rho)$	Electron temperature profile (not predict-first)
$n_e(\rho)$	Electron density profile (not predict-first)

dI_p/dt	Plasma current ramp-rate [MA/s]
dP_{NBI}/dt	NBI heating ramp-rate [MW/s]
$d\kappa/dt$	Elongation ramp-rate [1/s]
da_{minor}/dt	Minor radius ramp-rate [m/s]

Plasma Constraints

Greenwald Fraction

$$f_{GW} \propto \frac{\bar{n}_{e,20} a^2_{minor}}{I_p}$$

Decreasing plasma current worsens the problem, and density dynamics are slow

Rotational Transform

$$\iota_{95} \propto \frac{I_p}{a^2_{minor}(1+\kappa^2)}$$

Need to avoid kink instability limits while changing shape + Ip

Vertical Instability Growth Rate Need to manage as we change shapes + plasma current

Poloidal Beta Loosely correlated with NTM drive

The target goal:

Can we manage this scenario (140kA, baseline high performance), with minimal data?

Can the model extrapolate a bit to this scenario with zero data?

~400 shots
~2GB

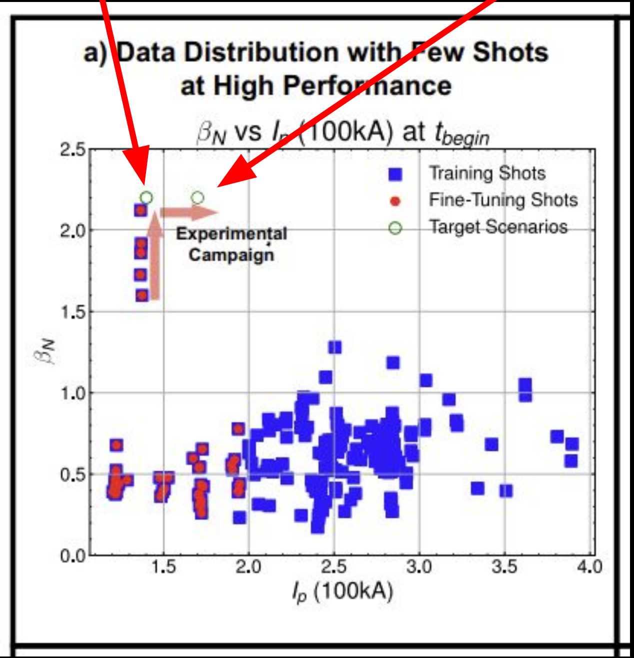

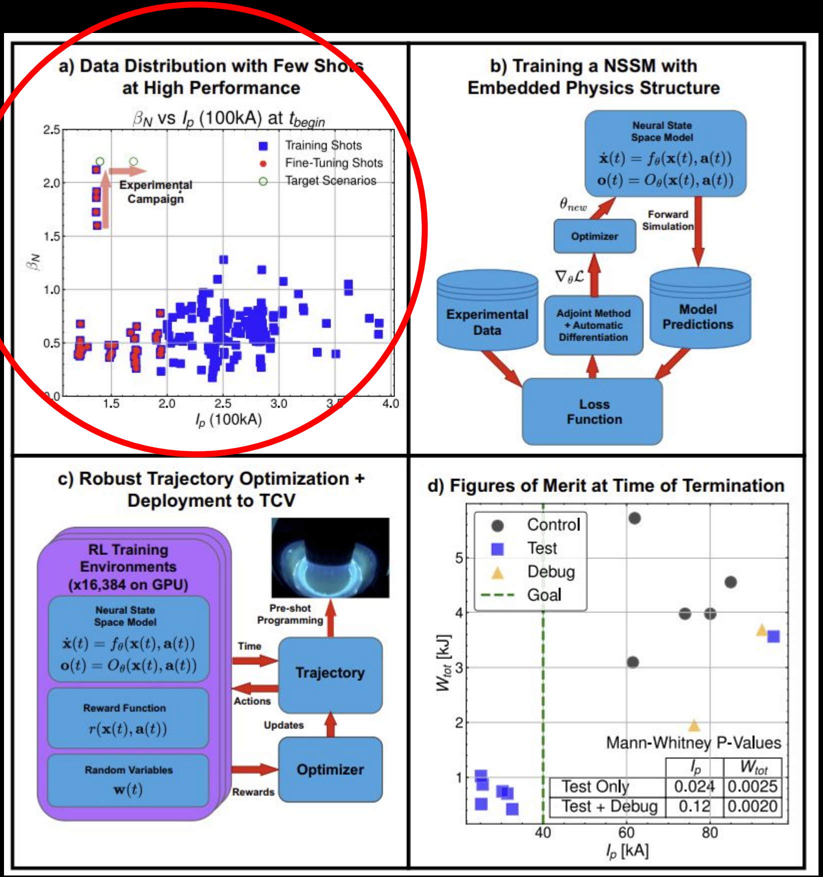

	I_p	W_{tot}
Test Only	0.024	0.0025
Test + Debug	0.12	0.0020

b) We train a NSSM with embedded physics structure: 0D energy and particle balance equations, which is a model that blends simple physics principles, power laws, and NNs

$$\frac{dW_{tot}}{dt} = -\underbrace{\frac{W_{tot}}{\tau_{E,pred}}}_{\text{Transport}} + \underbrace{I_p^2 NN_{ohm,rad,0}(\mathbf{x},\mathbf{a})}_{\text{Ohmic Heating}} - \underbrace{\bar{n}_{e,20} V NN_{ohm,rad,1}(\mathbf{x},\mathbf{a})}_{\text{Radiated Power}} + \underbrace{P_{NBI} + P_{ECRH}}_{\text{Aux. Heating}}$$

$$\frac{d(\bar{n}_{e,20}V)}{dt} = -\underbrace{\frac{\bar{n}_{e,20}}{\tau_{n,pred}}}_{\text{Transport}} + \underbrace{c_{NBI}P_{NBI}}_{\text{NBI Fueling}} + \underbrace{c_{gas,0}\,\sigma(c_{gas,1}V_{gas} + c_{gas,2})}_{\text{Gas Valve Fueling}} + \underbrace{NN_{wall}(\mathbf{x},\mathbf{a})\exp^{-c_{wall}g_{HFS}}}_{\text{Wall Effects}}$$

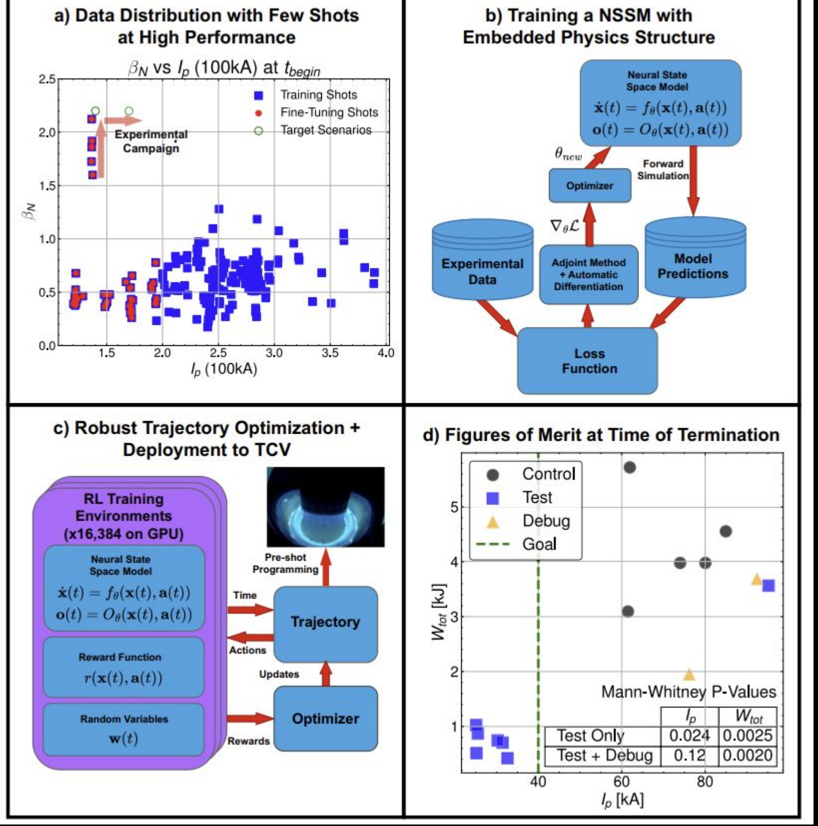

c) "RL gym" allows id of optimal trajectory ~ RL Policy that Only Observes Time

Statistical test provides degree of confidence there was improvement*

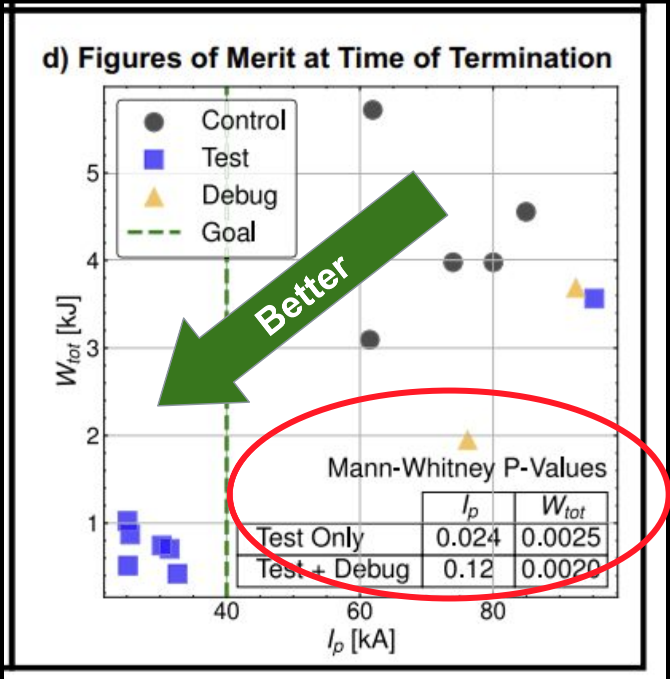

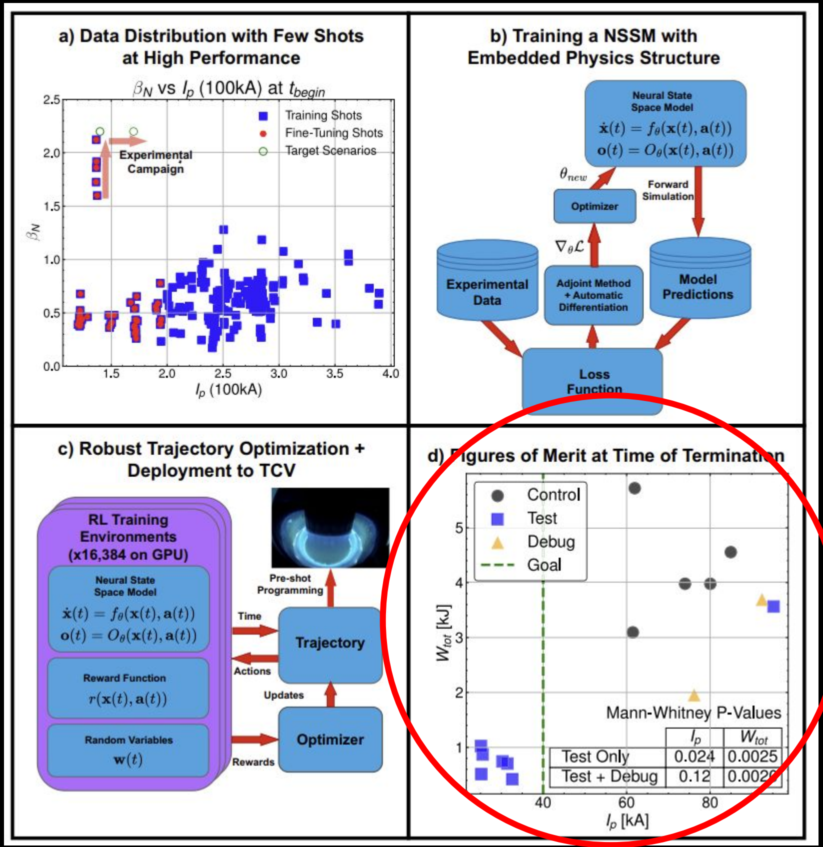

*Tokamaks are stochastic w/ distributional drift, some skepticism is important in the absence of overwhelming statistics

Continuous Model Validation + Updates for Incremental Improvements

- Continuously updated dynamics model + trajectories to new data
- Trajectories became faster and non-disruptive over time

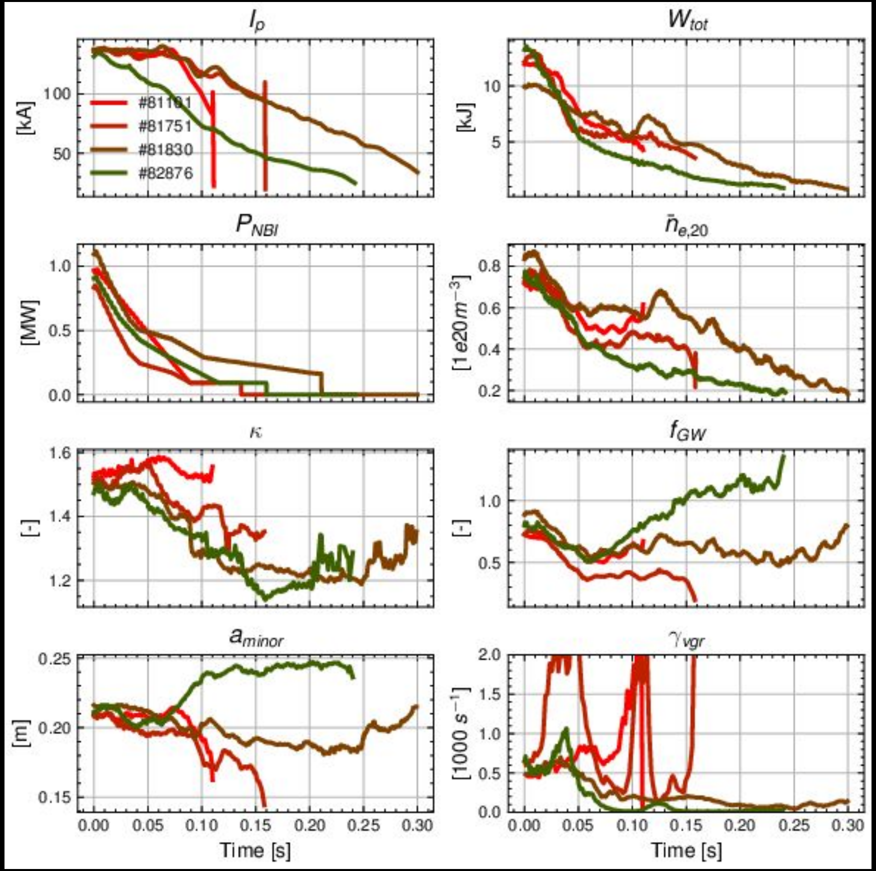

Predict-First

Can the model design a non-disruptive trajectory and make a priori predictions slightly out-of-distribution?

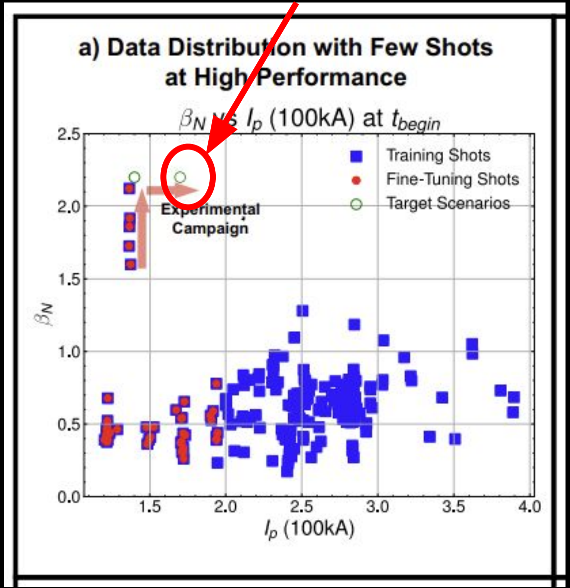

Yes!

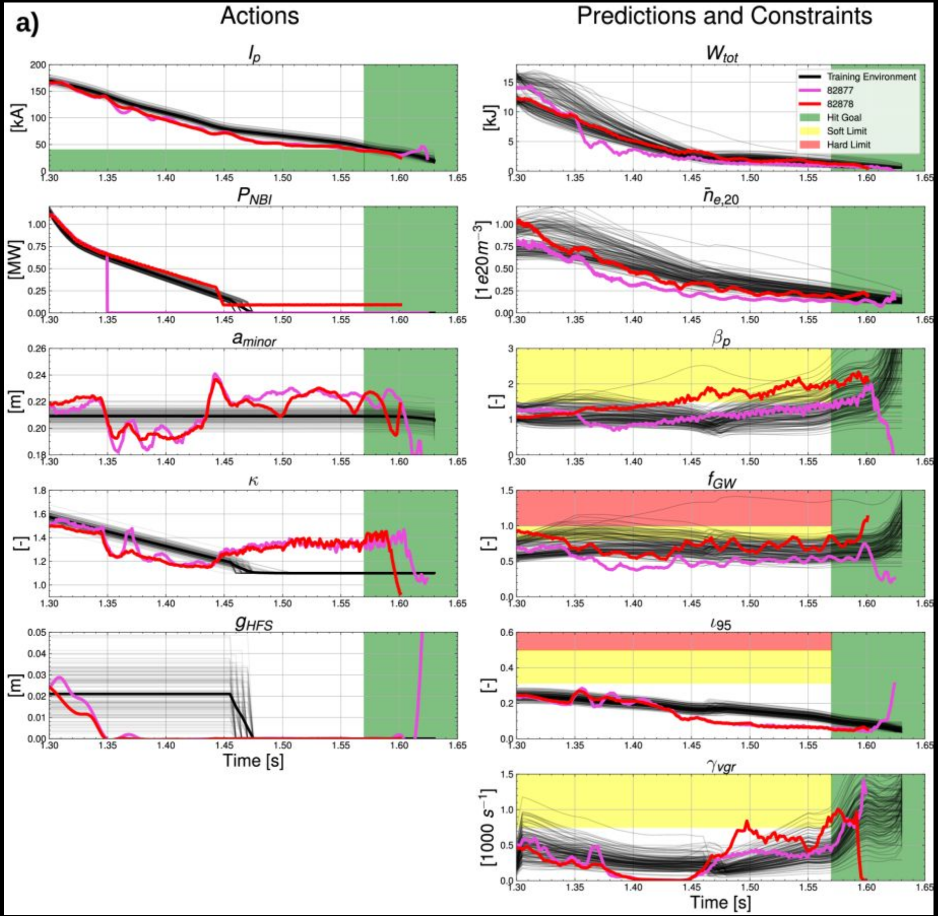

In summary: SciML for Kinetic Dynamics (NSSM) shown useful for rampdown trajectory optimization

- Demonstrated that a RL-based approach to rampdown robust trajectory optimization can help manage disruptivity
  - Accounting for distributional uncertainty has been previously assumed to be infeasible
  - Multi-actuator trajectory optimization with success on experiment is very rare
- Biggest limitation: we did not address tungsten
  - TCV is a carbon wall machine
  - Tungsten is a major challenge for rampdowns on high performance machines
- Developed approach may be promising for full-pulse trajectory optimization
- Framework components now open-sourced!
  https://github.com/MIT-PSFC/POPSIM-Public

Outline

Incomplete physics understanding of observable fusion instabilities prompts AI/ML solutions

- Interpretable/explainable ML = trustworthy solutions for safety-critical systems

- Continuous Model Validation + Updates are Critical
  - Even if the physics is known, models need to be continuously validated (also in operations) and updated if necessary

Incomplete physics understanding of observable fusion instabilities prompts AI/ML solutions

- Interpretable/explainable ML = trustworthy solutions for safety-critical systems

- Continuous Model Validation + Updates are Critical
  - Even if the physics is known, models need to be continuously validated (also in operations) and updated if necessary

→ Multi-machine database study + interpretable ML workflow to identify boundary functional form
  - Collisionality-based scaling for density limit events in L/H-mode
  - DIII-D avoidance experiments + new risk metric

→ Experimental demonstration of robust rampdown trajectory optimization at TCV
  - Improved high performance scenario, close to stability limits and with minimal data
  - Successful a priori trajectory design + prediction of a scenario slightly out of distribution

Thank you! < crea@psfc.mit.edu >

Backup slides