Control and machine learning algorithms for disruption avoidance and proximity control

Dan Boyer mboyer@pppl.gov

IAEA TM on Plasma Disruptions and their Mitigation July 20, 2022

C. Rea [MIT], I. Char, V. Mehta, J. Schneider [CMU], J. Wai, J. Abbate, R. Conlin, E. Kolemen, R. Shousha, Y. Fu , A. Jalalvand [Princeton U.], S. Sabbagh [CU], J. Barr [GA]

Disruption mitigation vs. avoidance

- **Mitigation:** minimize damage of a disruption
  - Reliable detection of imminent disruption and deployment of a strategy to limit impact of the event
    - Collision sensor deploying airbags in a vehicle
- **Avoidance:** modify scenario to avoid onset of a disruption
  - **Discrete changes:** give up on target scenario, or give up on shot and initiate 'safe landing'
    - Drive a different route to avoid dangerous icy patch
    - Return home due to treacherous conditions
  - **Continuous changes:** actively modify scenario in the minimum possible way that maintains robust operation
    - Steer around a slow vehicle
    - Swerve around a dangerous driver

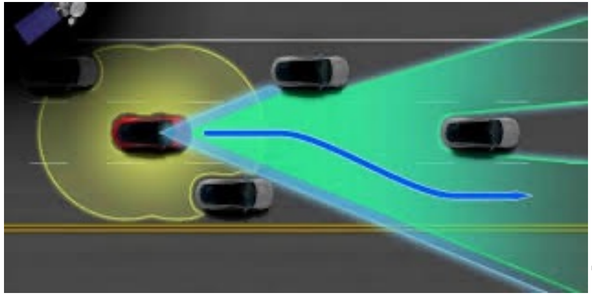

Proximity control

- Ensure the operating point remains robustly within the safe operating space
  - Avoid crossing thresholds that trigger mitigation/discrete changes to the system
  - Remain sufficiently far away to avoid triggering mitigation system on noise
- Due to limited actuation and time-varying behavior/constraints, important to consider predictive control strategies

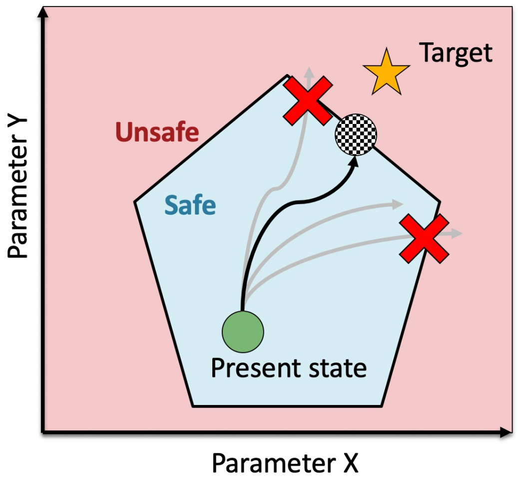

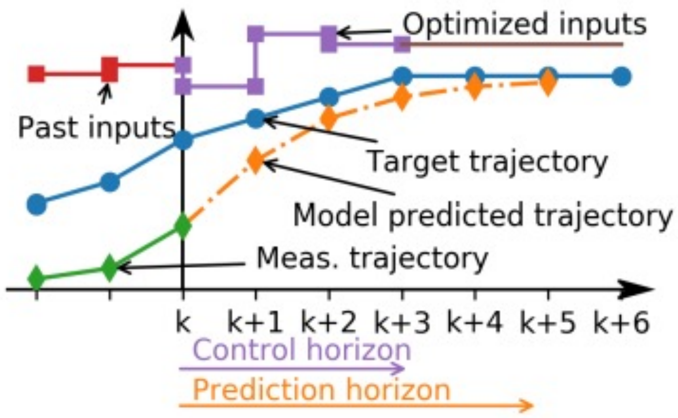

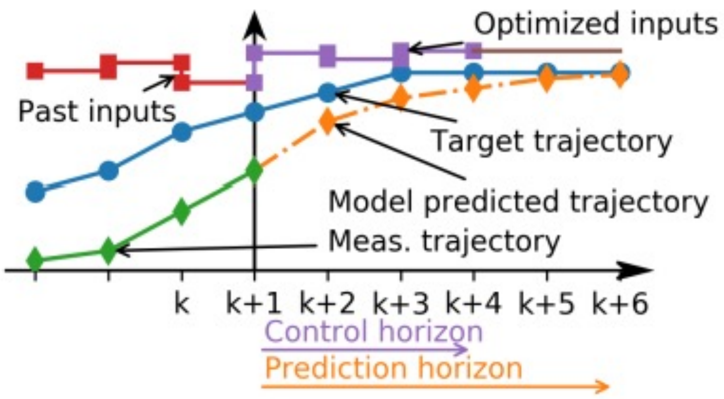

Requirements for Model Predictive Control

- Fast (to allow long enough look-ahead) + accurate (enough) predictive model
  - Approaches
    - System identification
    - Empirical neural networks
    - Neural networks for physics model acceleration
- Real-time capable approximation of limits
  - Numerically, prefer linear, convex sets of constraints
    - Safe operating region identification algorithm

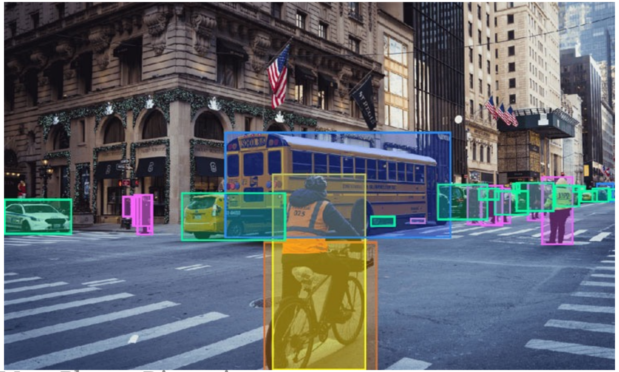

Example using TRANSP simulations of KSTAR

Data → Model → Constraints → Targets → Closed loop control

Boyer et al., 2020 Nuclear Fusion 60 (9), 096007

Example using TRANSP simulations of KSTAR

Data → Model → Constraints → Targets → Closed loop control

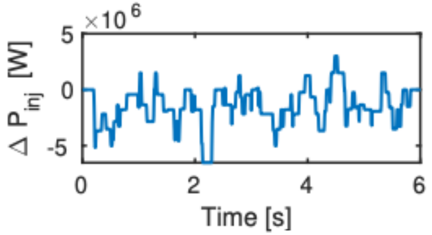

Beam power

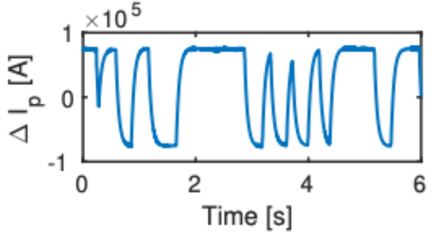

Plasma current

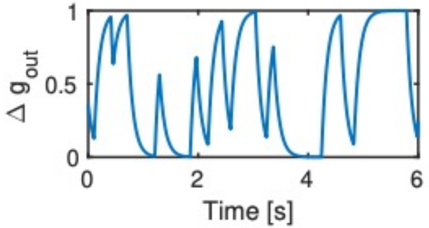

Outer gap

- Actuator modulations applied to a reference predictive simulation provide data on response of variables of interest
  - Stored energy
  - Internal inductance
  - Loop voltage

Example using TRANSP simulations of KSTAR

- System identification Matlab toolbox enables fitting a state-space model to the simulated response data

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$$
$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t)$$

TRANSP
Identified model

Internal inductance

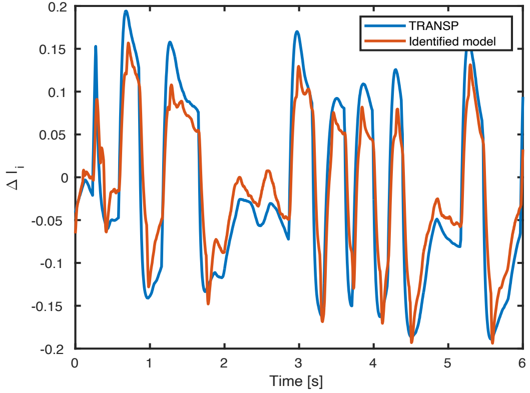

Stored energy

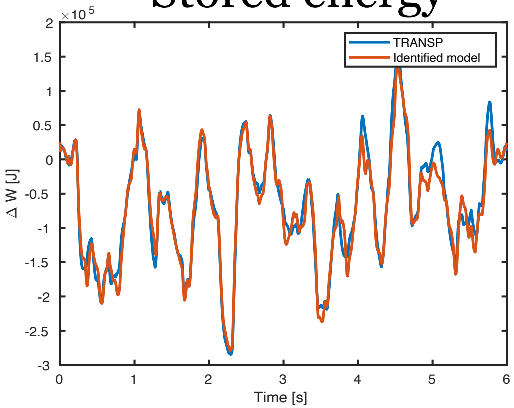

Loop voltage

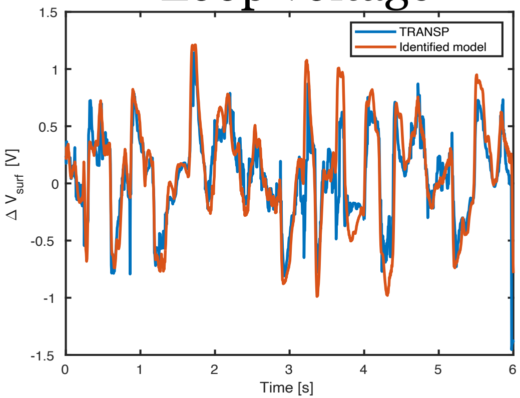

Example using TRANSP simulations of KSTAR

Data → Model → Constraints → Targets → Closed loop control

Ad hoc constraints chosen for demonstration

Simulated beam fault: changing actuator constraint

Time-varying maximum allowable stored energy

Internal inductance dependent limit on elongation (VDE limit)

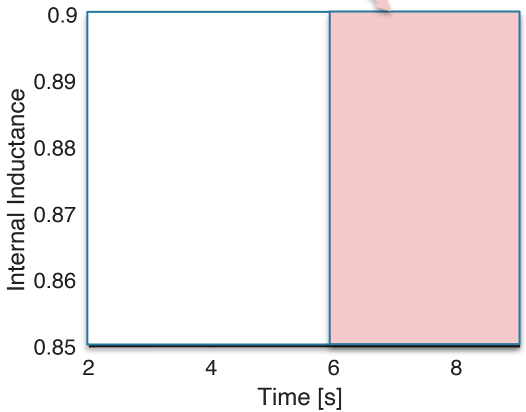

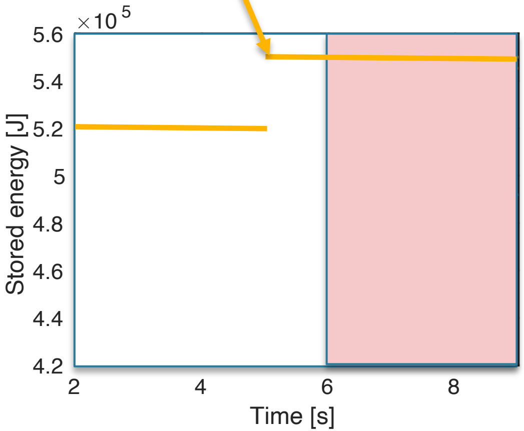

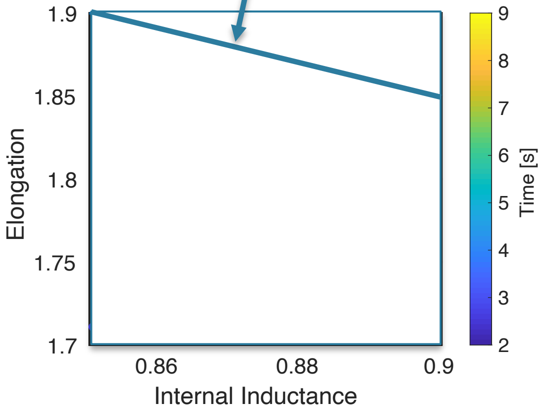

Example using TRANSP simulations of KSTAR

Data → Model → Constraints → Targets → Closed loop control

Time varying internal inductance target

Time varying stored energy target.
Step change initially violates imposed constraints.

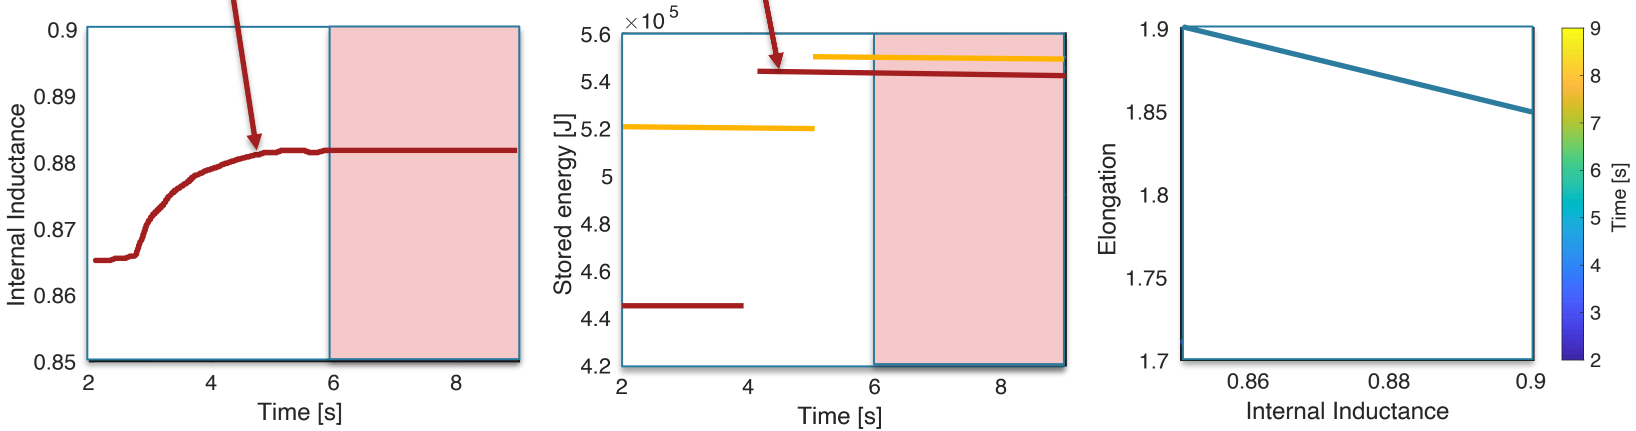

Example using TRANSP simulations of KSTAR

Data → Model → Constraints → Targets → Closed loop control

Feedback off

Beam fault

Controller limits stored energy until constraint is relaxed.

Shape change is used to make up for loss of beam. Controller respects elongation limit to avoid VDE

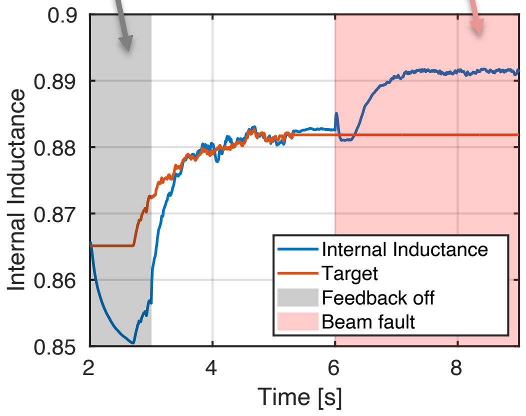

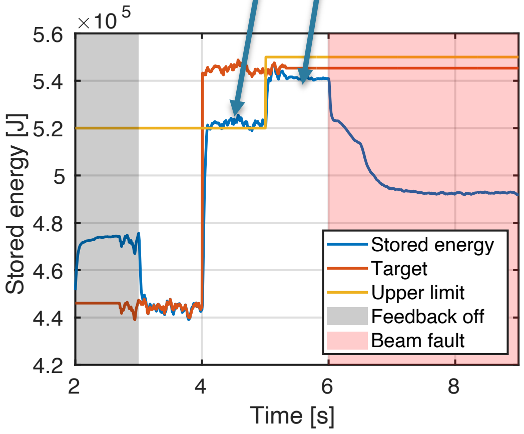

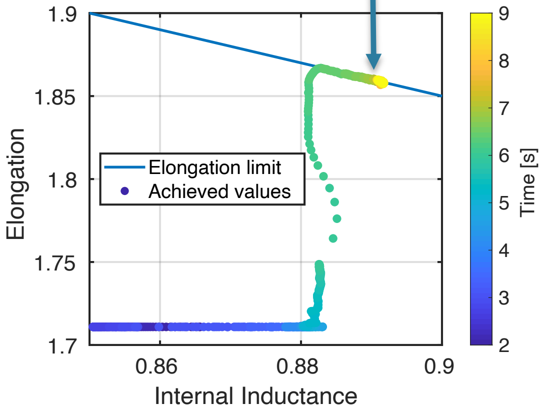

Opportunities for improvement addressed by recent results and ongoing activities

- The identified model is linearized around a reference TRANSP simulation and depends only on TRANSP predictions
  - **Improvements:** Empirical models and/or accelerated nonlinear physics models
- The constraints tested were ad hoc
  - **Improvements:** Base limits on state-of-the-art disruption predictors

Opportunities for improvement addressed by recent results and ongoing activities

- The identified model is linearized around a reference TRANSP simulation and depends only on TRANSP predictions
  - **Improvements:** Empirical models and/or accelerated nonlinear physics models
- The constraints tested were ad hoc
  - **Improvements:** Base limits on state-of-the-art disruption predictors

Empirically trained neural networks enable predictions over a wide operating space and close matching to experiments

- Several successful approaches on DIII-D
  - **Convolutional+recurrent neural networks:** J. Abbate et al 2021 Nucl. Fusion 61 046027LS.2021.3085 504.
  - **Reservoir computing:** Jalalvand, et al., IEEE Transactions on Neural Networks and Learning Systems doi: 10.1109/TNN
  - **Ian Char's work on reinforcement learning** (Carnegie Mellon University)

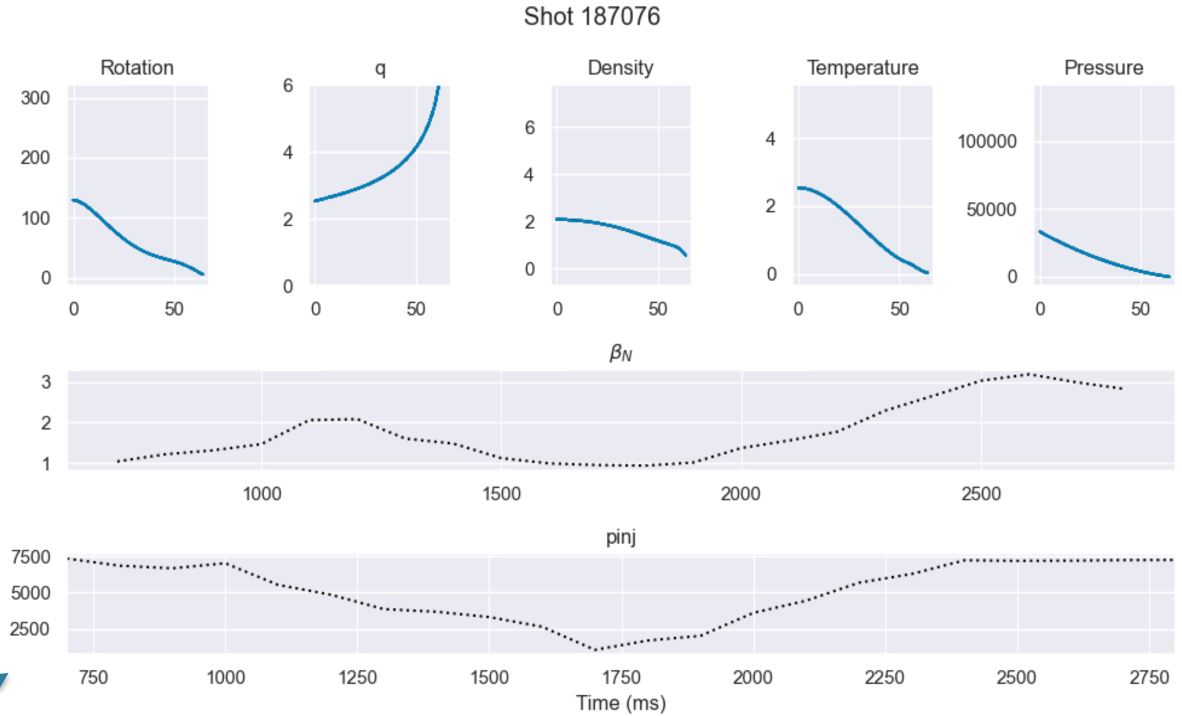

- **However:** models may not make useful predictions outside of distribution of experimental data
  - Especially important for new machines, upgrades, or new scenarios
  - **Embedding physics knowledge could help!**

Machine learning enables faster versions of physics models for optimization and control

- **NUBEAM** is a Monte Carlo code that calculates the effect of neutral beams on the plasma (heating, current drive, torque)
  - Often takes >30% of TRANSP calculation time
- Machine learning approaches enable the development of NubeamNet

Boyer et al., 2019 Nuclear Fusion 59 (5), 056008

Execution time

NubeamNet calculation of beam effects takes ~100 microseconds per sample as opposed to minutes for NUBEAM

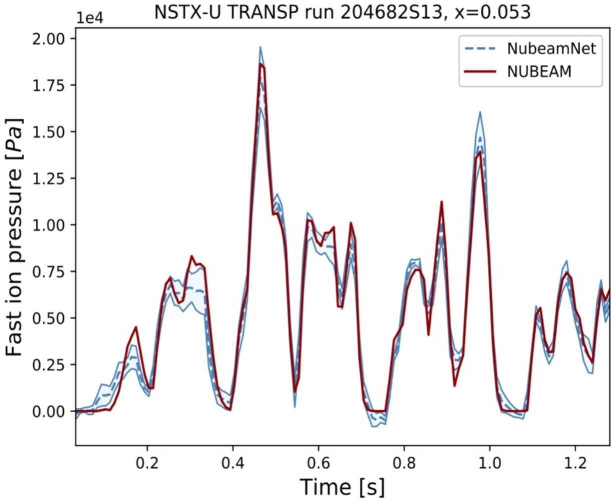

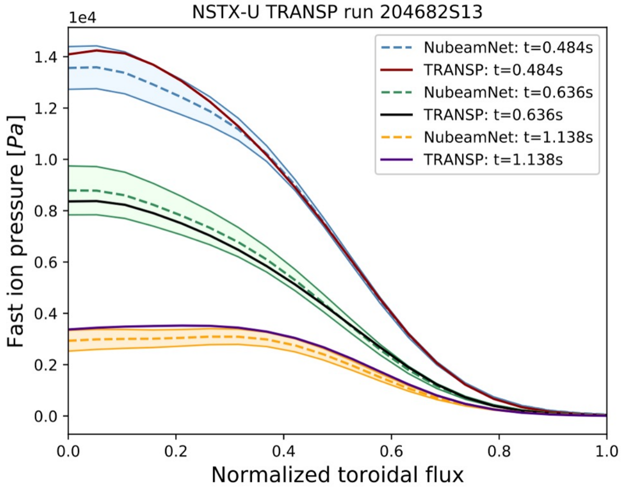

Similar approach developed for equilibrium modeling

- Neural network trained to predict flux from coil currents and profiles
  - Trained on reconstructions, could use modeled equilibria
  - Fast equilibrium prediction for feedforward coil current design
- Separate model predicts linearized response of plasma to changes in currents
  - Provides plasma-modified mutual inductance for conductor modeling
  - Expected to enable fast nonlinear modeling of equilibrium evolution

Wai, J, Boyer, M.D., Kolemen, E, 2022 Nucl. Fusion 62 086042

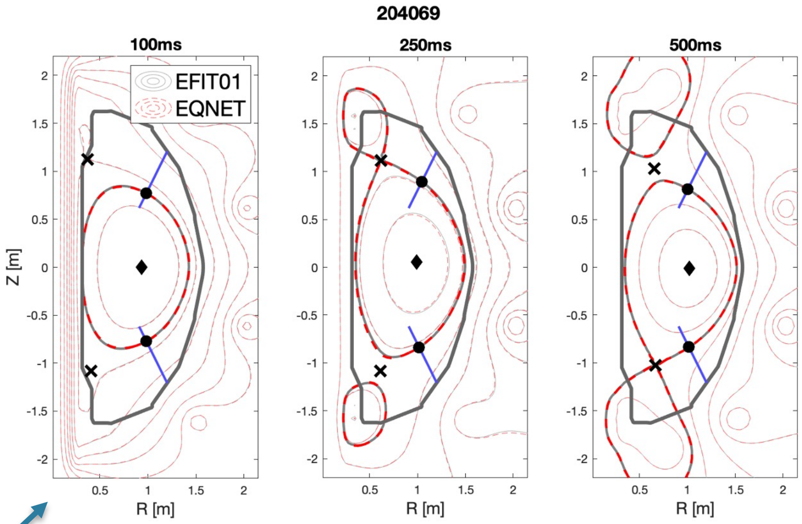

Comparison of equilibrium predictions from EFIT01 and neural network

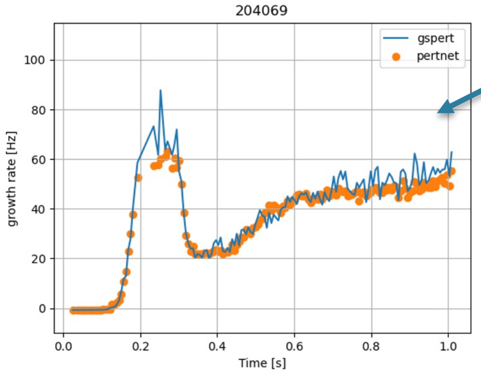

Predicted vertical growth rate matches non-rigid calculation

Combining accelerated physics models with empirical models to optimize accuracy and generalizability

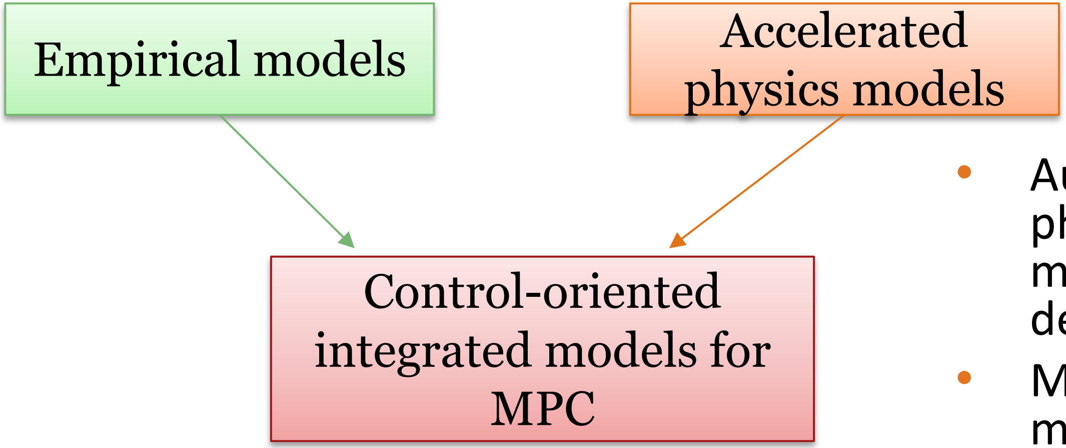

- Augment well-validated, accelerated physics surrogate models with empirical models for phenomena not well-described by available physics
- Modular approach to include new models or data as they become available

Mark Boyer Proceedings of the 2nd Conference on Learning for Dynamics and Control, PMLR 120:698-707, 2020.

V. Mehta et al., "Neural Dynamical Systems: Balancing Structure and Flexibility in Physical Prediction," 2021 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 3735-3742.

Opportunities for improvement addressed by recent results and ongoing activities

- The identified model is linearized around a reference TRANSP simulation and depends only on TRANSP predictions
  - **Improvements:** Empirical models and/or accelerated nonlinear physics models
- The constraints tested were ad hoc
  - **Improvements:** Base limits on state-of-the-art disruption predictors

Machine learning has emerged as a powerful disruption prediction/detection approach

- Random forests[1], deep learning[2], SVMs[3], etc.
  - Predict **disruptivity** and select a **threshold** based on confusion matrix
- Efforts have so far focused largely on **triggering shutdown/mitigation**
  - Can we use models like these to support **moving from prediction to control?**

[1] Rea et al., PPCF 2018.

[2] Kates-Harbeck, et al., Nature 2019.

[3] Cannas, et al., et al., FED 2007.

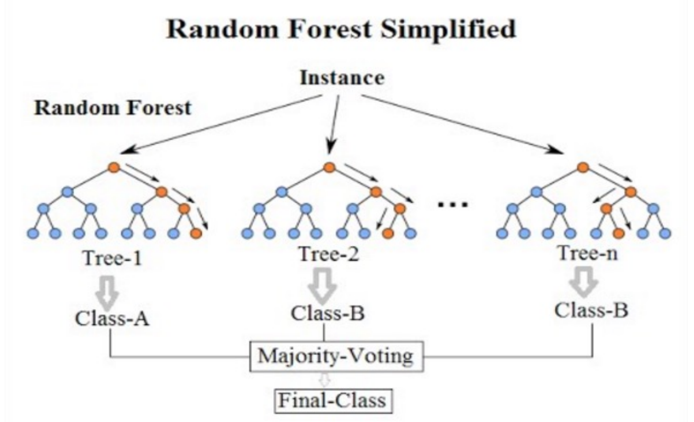

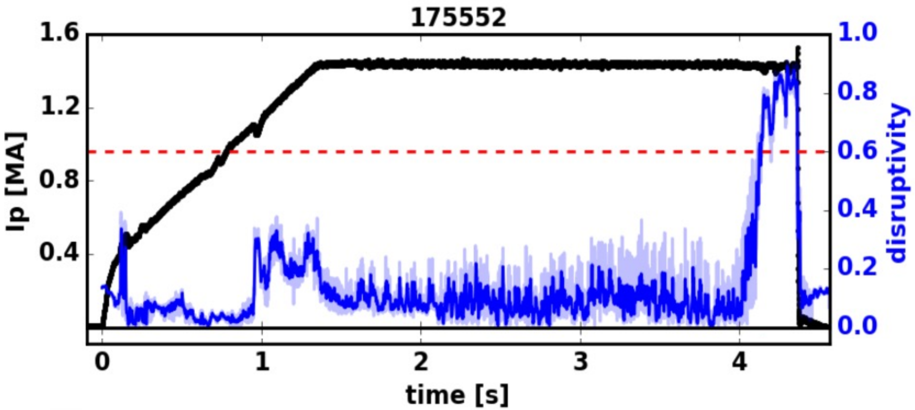

- Feedback adjustment of beam power was able to keep 'tearability' at a safe level
  - Empirically tuned feedback parameters
  - Relied on assumption that tearability is proportional to beam power

Fu, Y., et al. PoP 27, 022501 (2020)

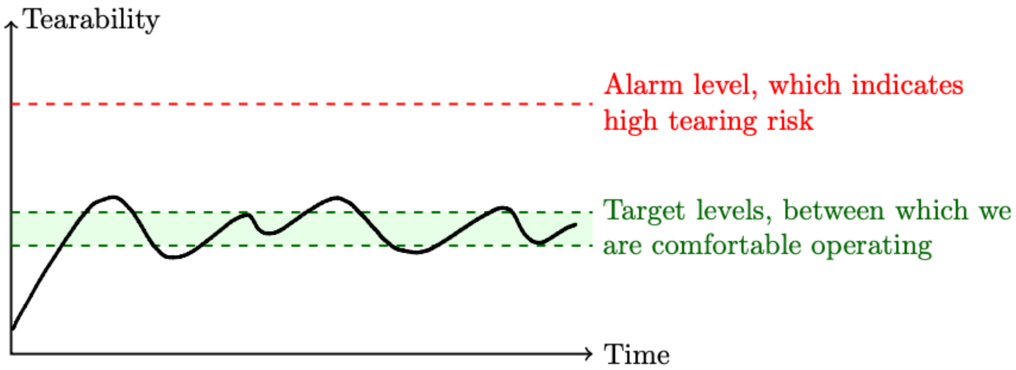

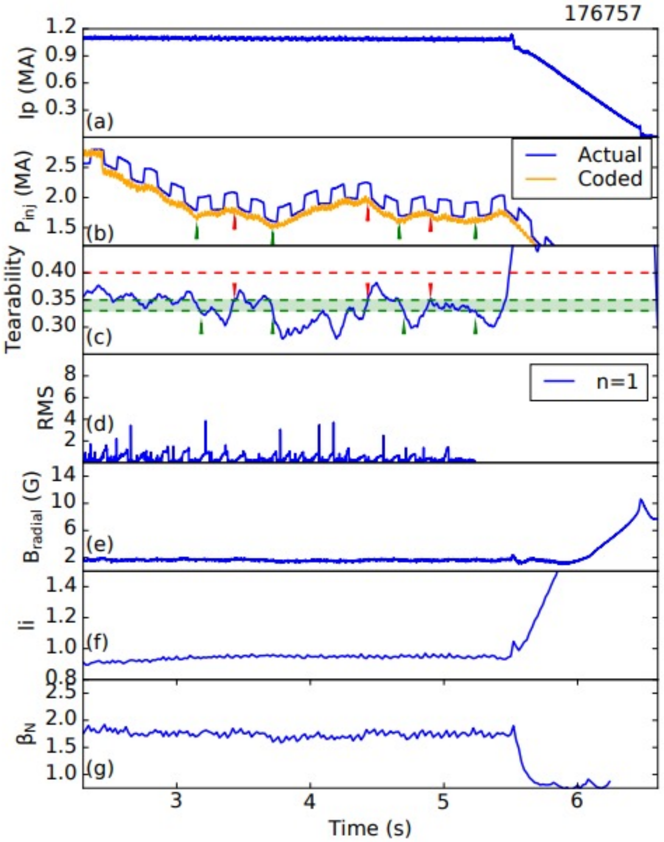

By sampling the local operating space, a map of disruptivity predictions from the model can be produced

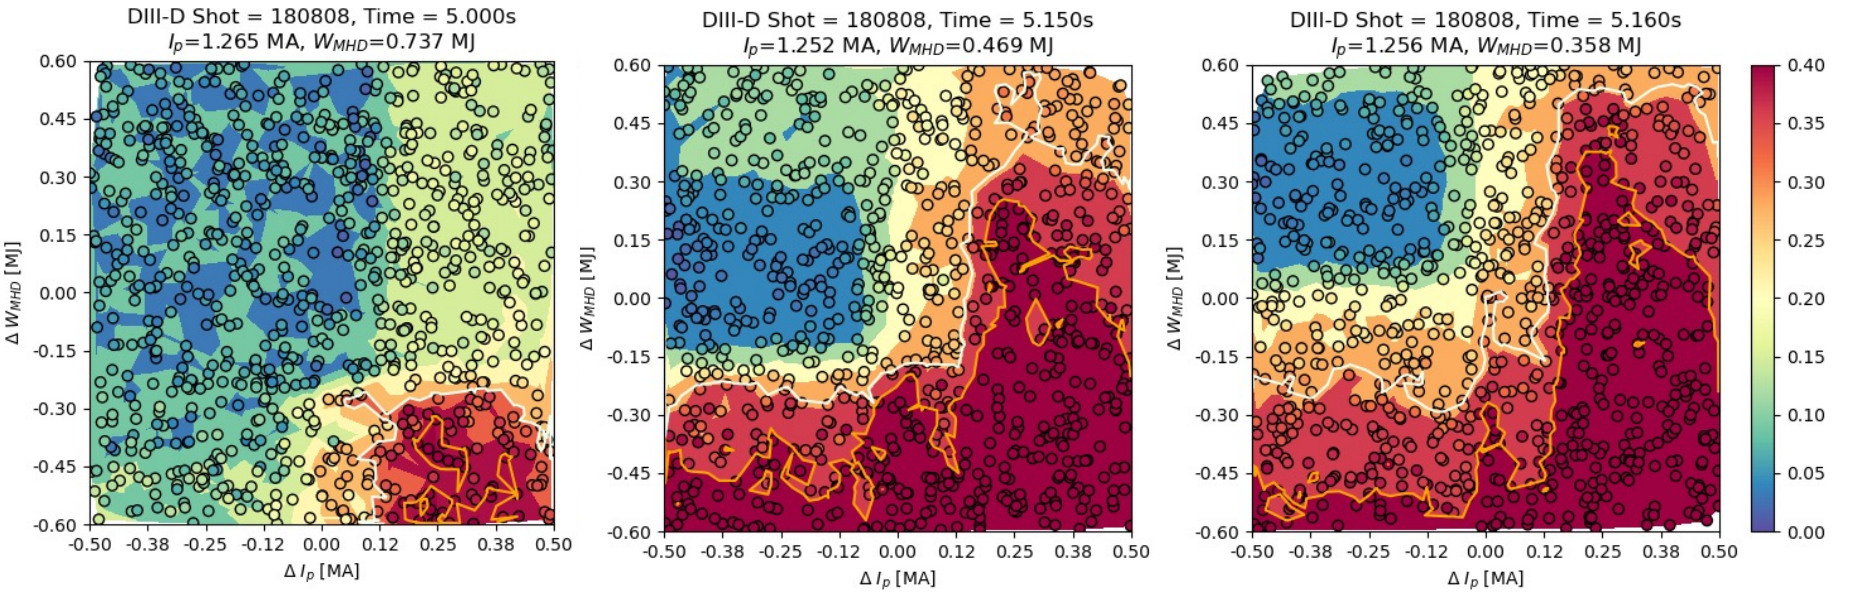

Disruptivity vs. distance from disruption boundary

- Distance to disruption boundary varies significantly for the same values of disruptivity
- Minimum safe distance to disruption boundary is directly related to control performance (noise, disturbances, etc.)

Points: Evaluation of DPRF in Wmhd-Ip space
Color: Interpolation of disruptivity
Orange line: Disruptivity threshold (0.40)
White line: 75% of disruptivity threshold

Real-time safe operating region identification (SORI)

- New algorithm can convert **disruption prediction models** into an estimate of the **safe operating space**
  - Uses genetic optimization to identify constraints that bound points around current operating point that are predicted to be safe
- Will enable monitoring of **disruption proximity** and **active optimization** of controllable parameters to **avoid disruption**
- Real-time relevant **execution times achieved with GPUs** for parallelization
- Worked with C. Rea [MIT] to train a new version of DPRF and apply proposed algorithm to DIII-D data

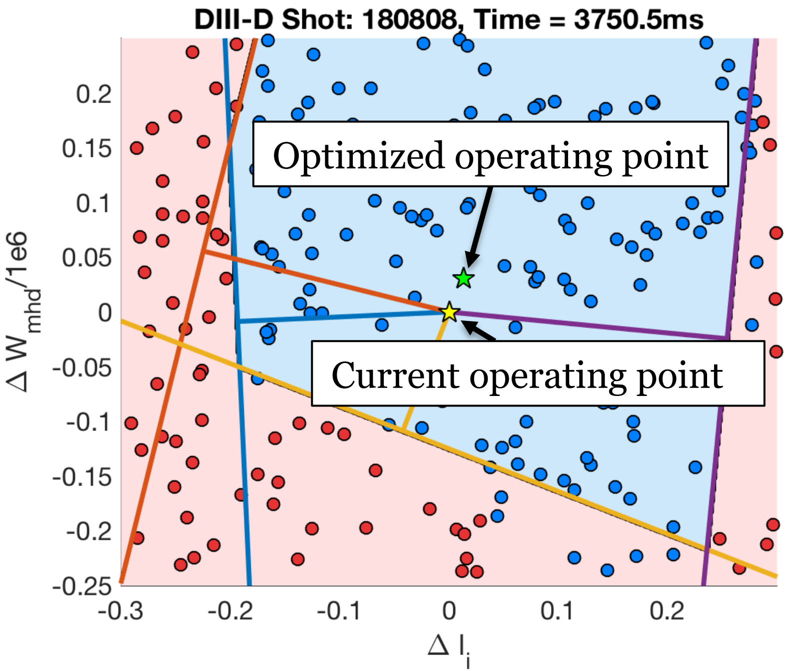

Boyer et al 2021 *Nuclear Fusion*
https://doi.org/10.1088/1741-4326/ac359e

GPU implemented genetic algorithm finds optimal set of constraints (largest safe operating space)

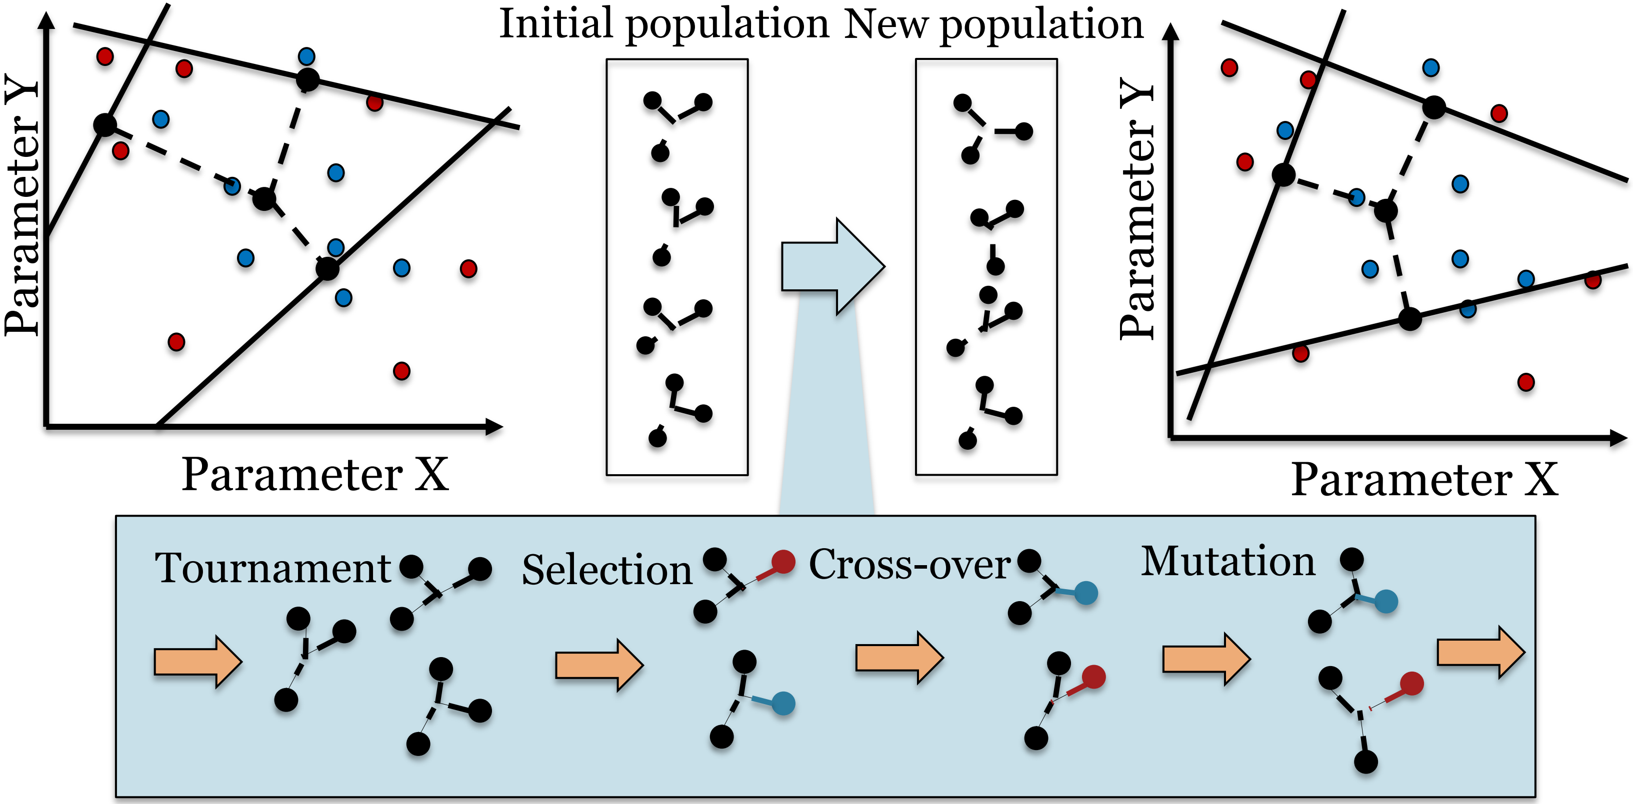

900 points, 3 constraints, 40 individuals, 25 generations = **~7ms per cycle**

SORI algorithm identifies set of convex linear constraints separating safe/unsafe points in the local operating space

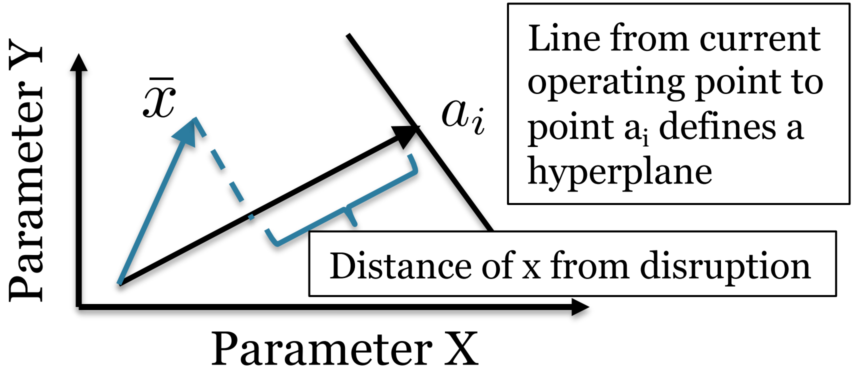

- Each constraint takes the form

$$a_i \bar{x} < \|a_i\|_2^2$$

where $\bar{x}$ is a point in **normalized** (by standard deviation of tracking/estimation error), **relative** (to current operating point) operating space.

SORI algorithm identifies set of convex linear constraints separating safe/unsafe points in the local operating space

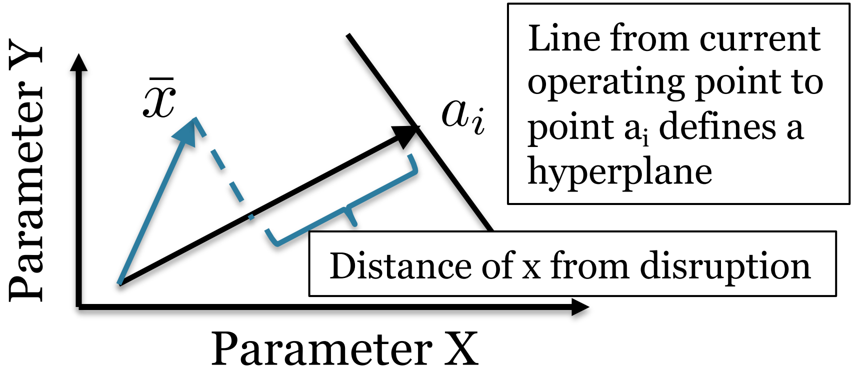

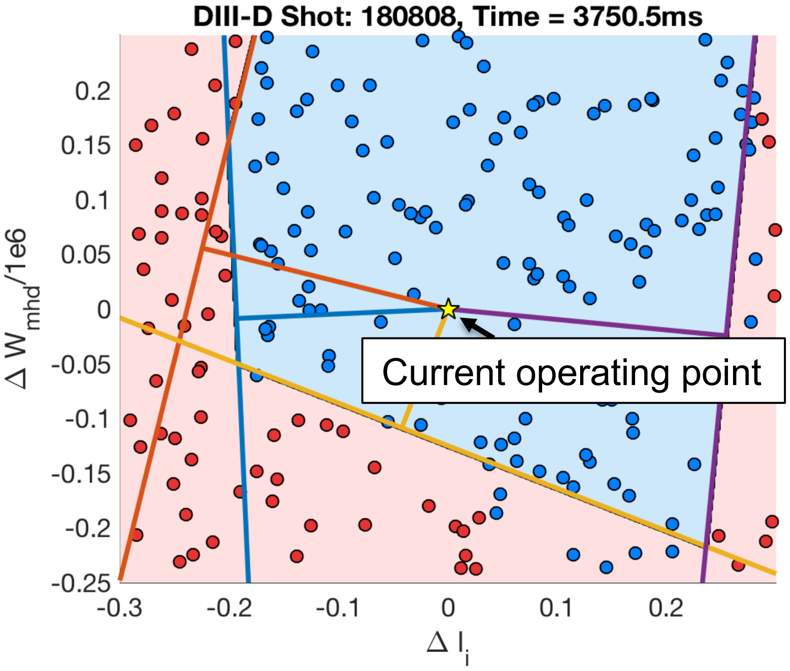

- Each constraint takes the form

$$a_i \bar{x} < \|a_i\|_2^2$$

where $\bar{x}$ is a point in **normalized** (by standard deviation of tracking/estimation error), **relative** (to current operating point) operating space.

Normalized distance of operating point from disruptive boundary: $\min_i \|a_i\|_2$

Example: SORI identifies evolving safe operating region for DIII-D shot 180808 showing approach to disruption

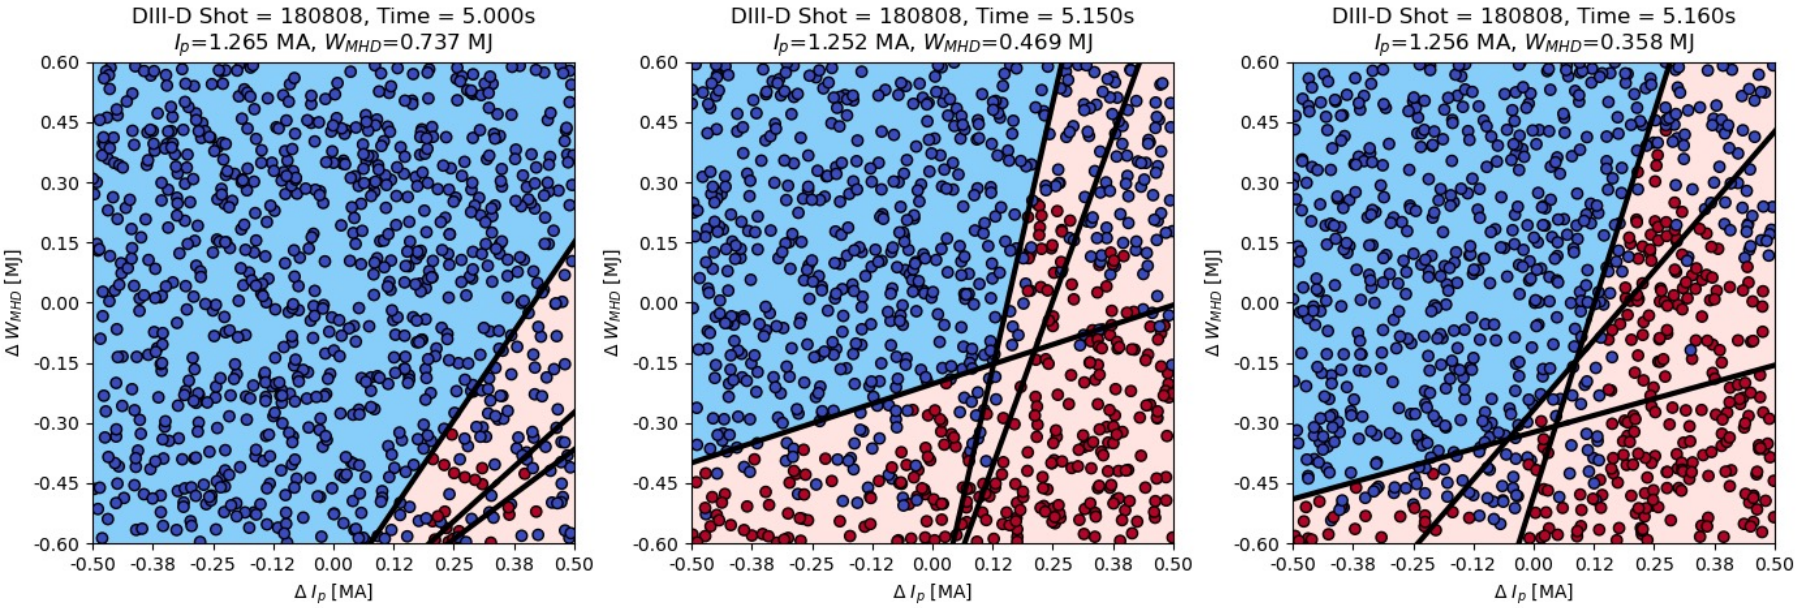

- Indicates decreasing plasma current and/or increasing stored energy could avoid/delay disruption in this case -> **interpretable result tied to specific actions**

Identified constraints enable adjusting targets to maintain a prescribed safety margin

- A safety margin m can be added to the constraint:

$$a_i \bar{x} < \|a_i\|_2 \left(\|a_i\|_2 - m\right)$$

- Optimization of targets, penalizing violation of safety margin:

$$\min_{\bar{x}} \quad \frac{1}{2} \bar{x}^t Q \bar{x} + w_p P(\bar{x})$$

Q weights deviation in active variables.
P penalized constraint violation.
Wp sets weight on penalty

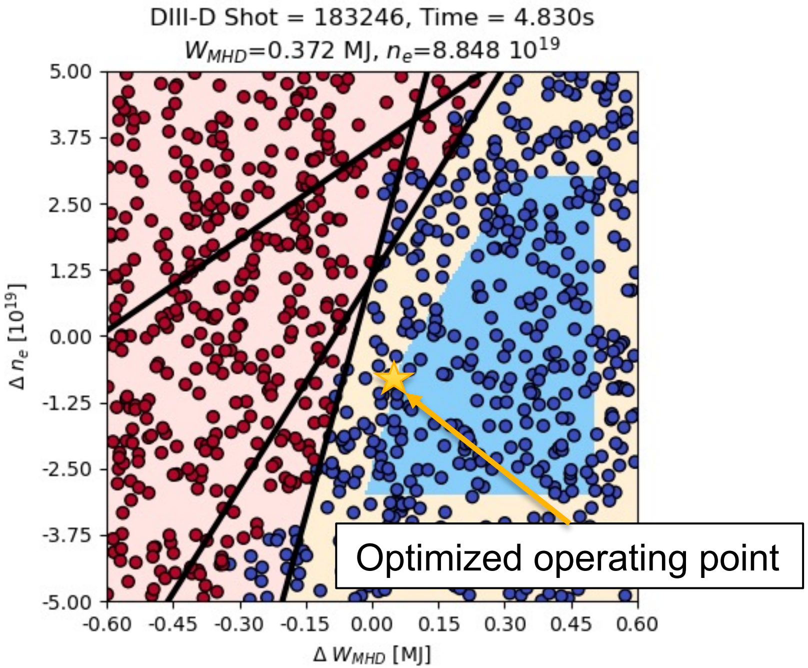

Many possible cost functions, this one tries to find the point closest to the current operating point that maintains the specified safety margin.

Example: Optimization of targets indicates how to maintain a safety margin in disruptive shot 183246 with small adjustments in density and stored energy

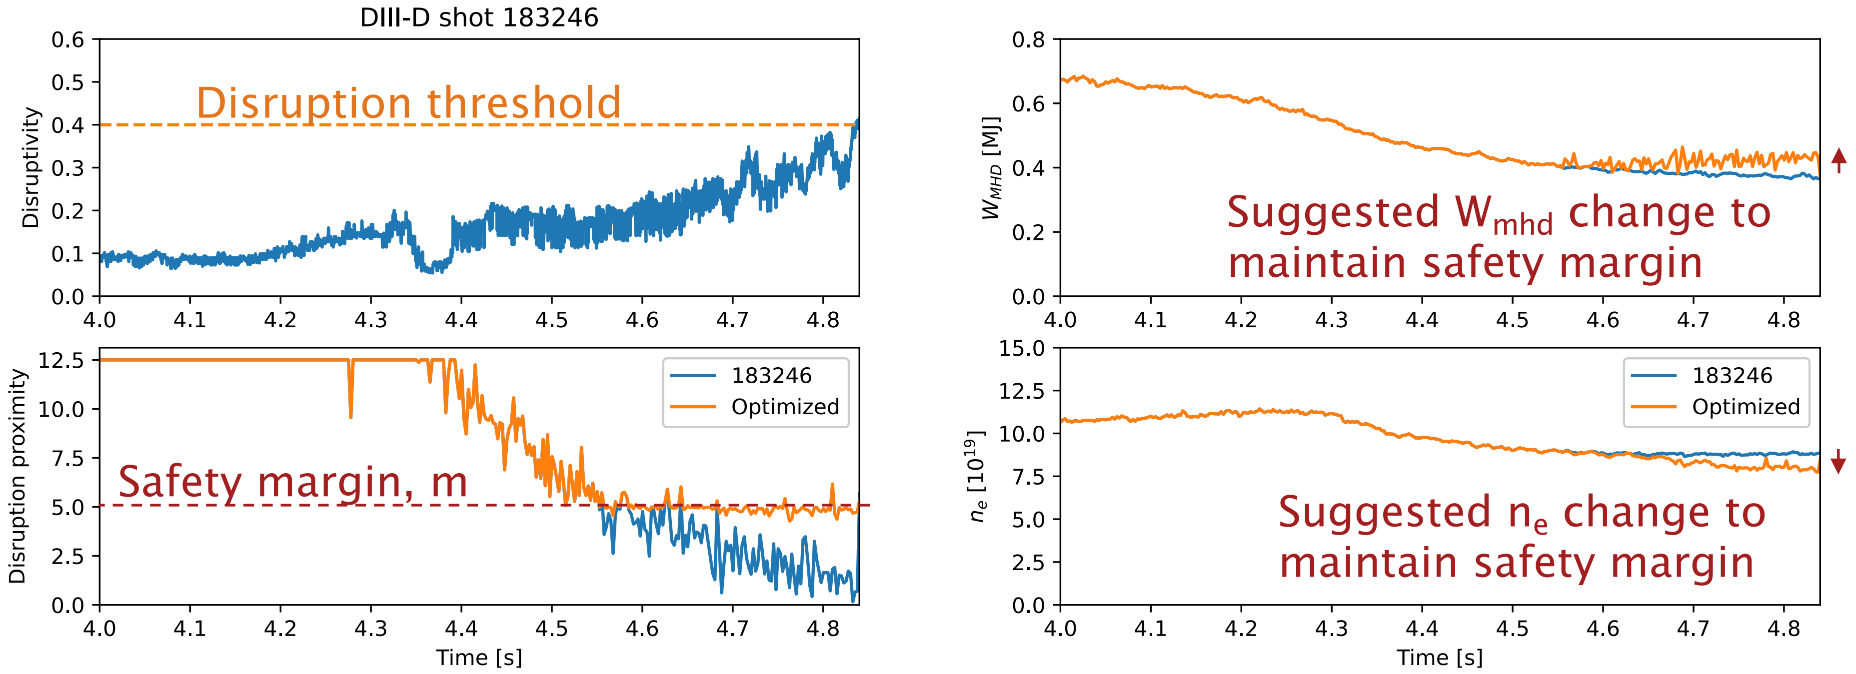

ML algorithms are enabling model predictive control development for disruption avoidance

- Predictive proximity control to ensure robust operation requires:
  - Accelerated modeling
    - Combination of accelerated physics models and empirical models
  - Identification of constraints
    - Important to control proximity vs. disruptivity
    - Important to develop causal understanding of disruptions
- Experimental validation of these approaches underway at DIII-D

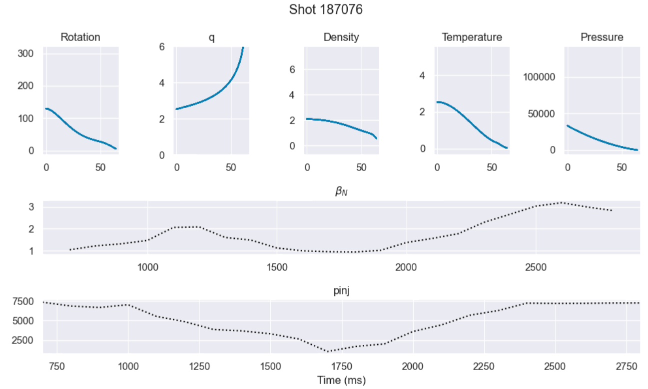

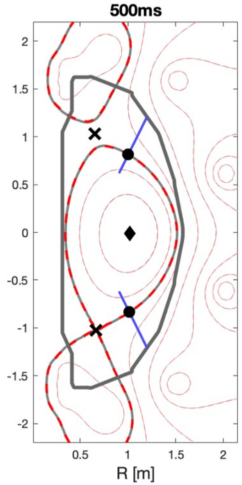

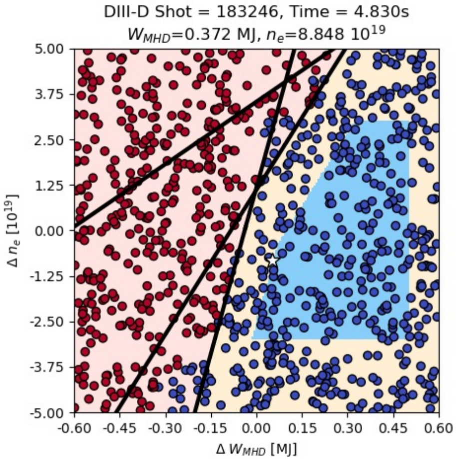

Thank you!

It is important to distinguish between disruption detection and prediction: symptoms vs. causes

- **Symptoms of COVID-19:** cough, fever, loss of smell/taste
  - Recognition of the symptoms enables treatment
    - quarantine, hospitalization, etc.
- **Causes of COVID-19:** contact with carriers' respiratory droplets
  - Control of the causes enables avoidance: avoid close contact and wear masks
- **Plasma example:**
  - Large vertical position/velocity is a symptom of VDE, not the cause
    - Could trigger mitigation
  - Vertical growth rate exceeding the controllability limits of the power supplies is the cause.
    - Reducing elongation can reduce the growth rate.

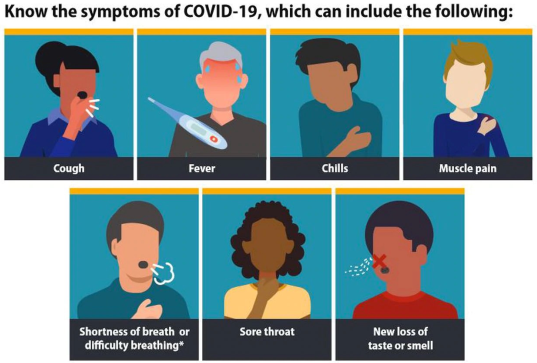

Implications for choice of disruption prediction model

- For **detection**, all inputs should be considered
  - **Goal:** detect disruptions as early as possible with high reliability
- For **avoidance**, inputs should be restricted
  - Controllable, predictable, or at least weakly correlated
  - **Goal:** predict likelihood of disruption based on a set of inputs we can affect in a reasonably predictable way

- **Active variables**
  - Variables we **want to change** to actively avoid disruptions
  - Or have **significant impact on disruptivity**
  - Can be **controlled within some range** (sets range of scan)
- **Context variables**
  - Variables we **don't want to (or can't) change**
  - Or have **less impact on disruptivity**
  - Define **expected range** based on combination of
    - Measurement/estimation error
    - Tracking error
    - Projection/trending
    - Future pre-programmed targets
- **Choice of active variables may change throughout a discharge**

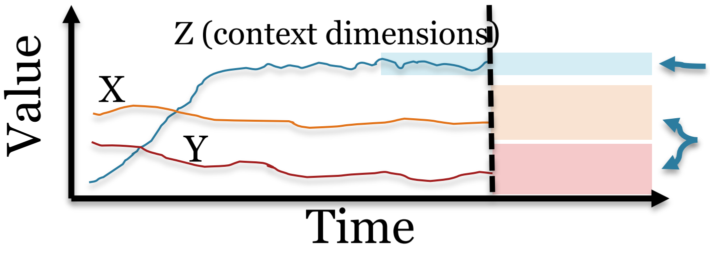

Assuming Z stays in this range

How does disruptivity change as we scan this range?