Science and Technology Facilities Council

Hartree Centre

AI for Magnetic Confinement Fusion

A. Agnello for the Fusion Computing Lab

HAMLET-Physics, Aug.19th-21st 2024

The STFC Hartree Centre

High-performance computing, data analytics and artificial intelligence research facility located at the Sci-Tech Daresbury research and innovation campus

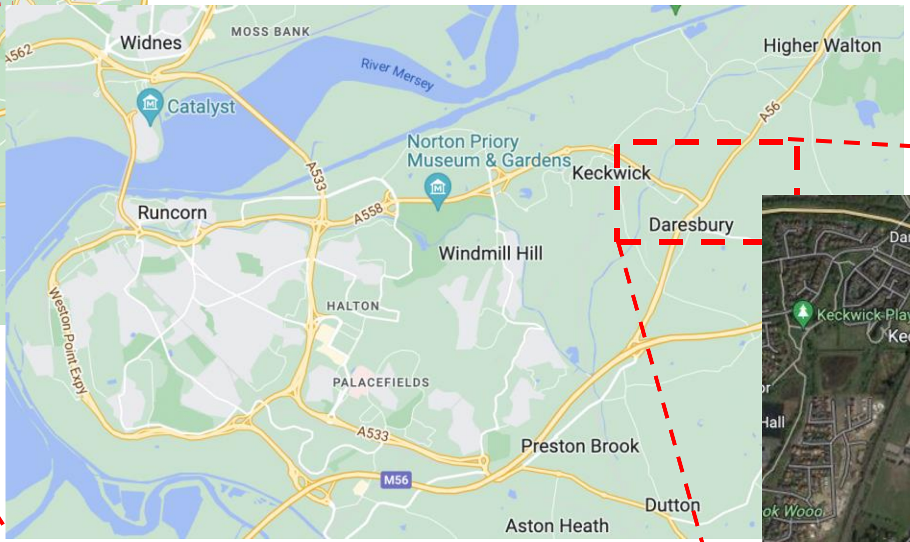

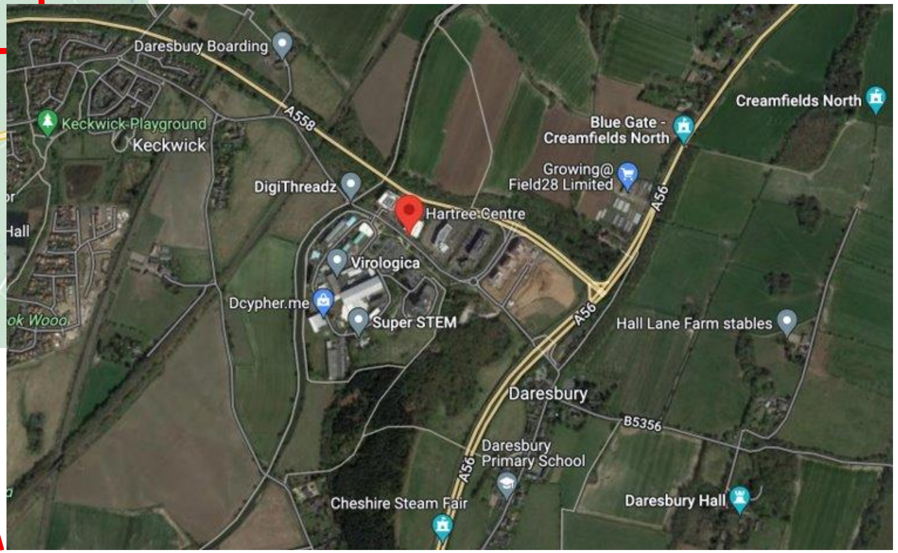

AI activities at the Hartree Centre

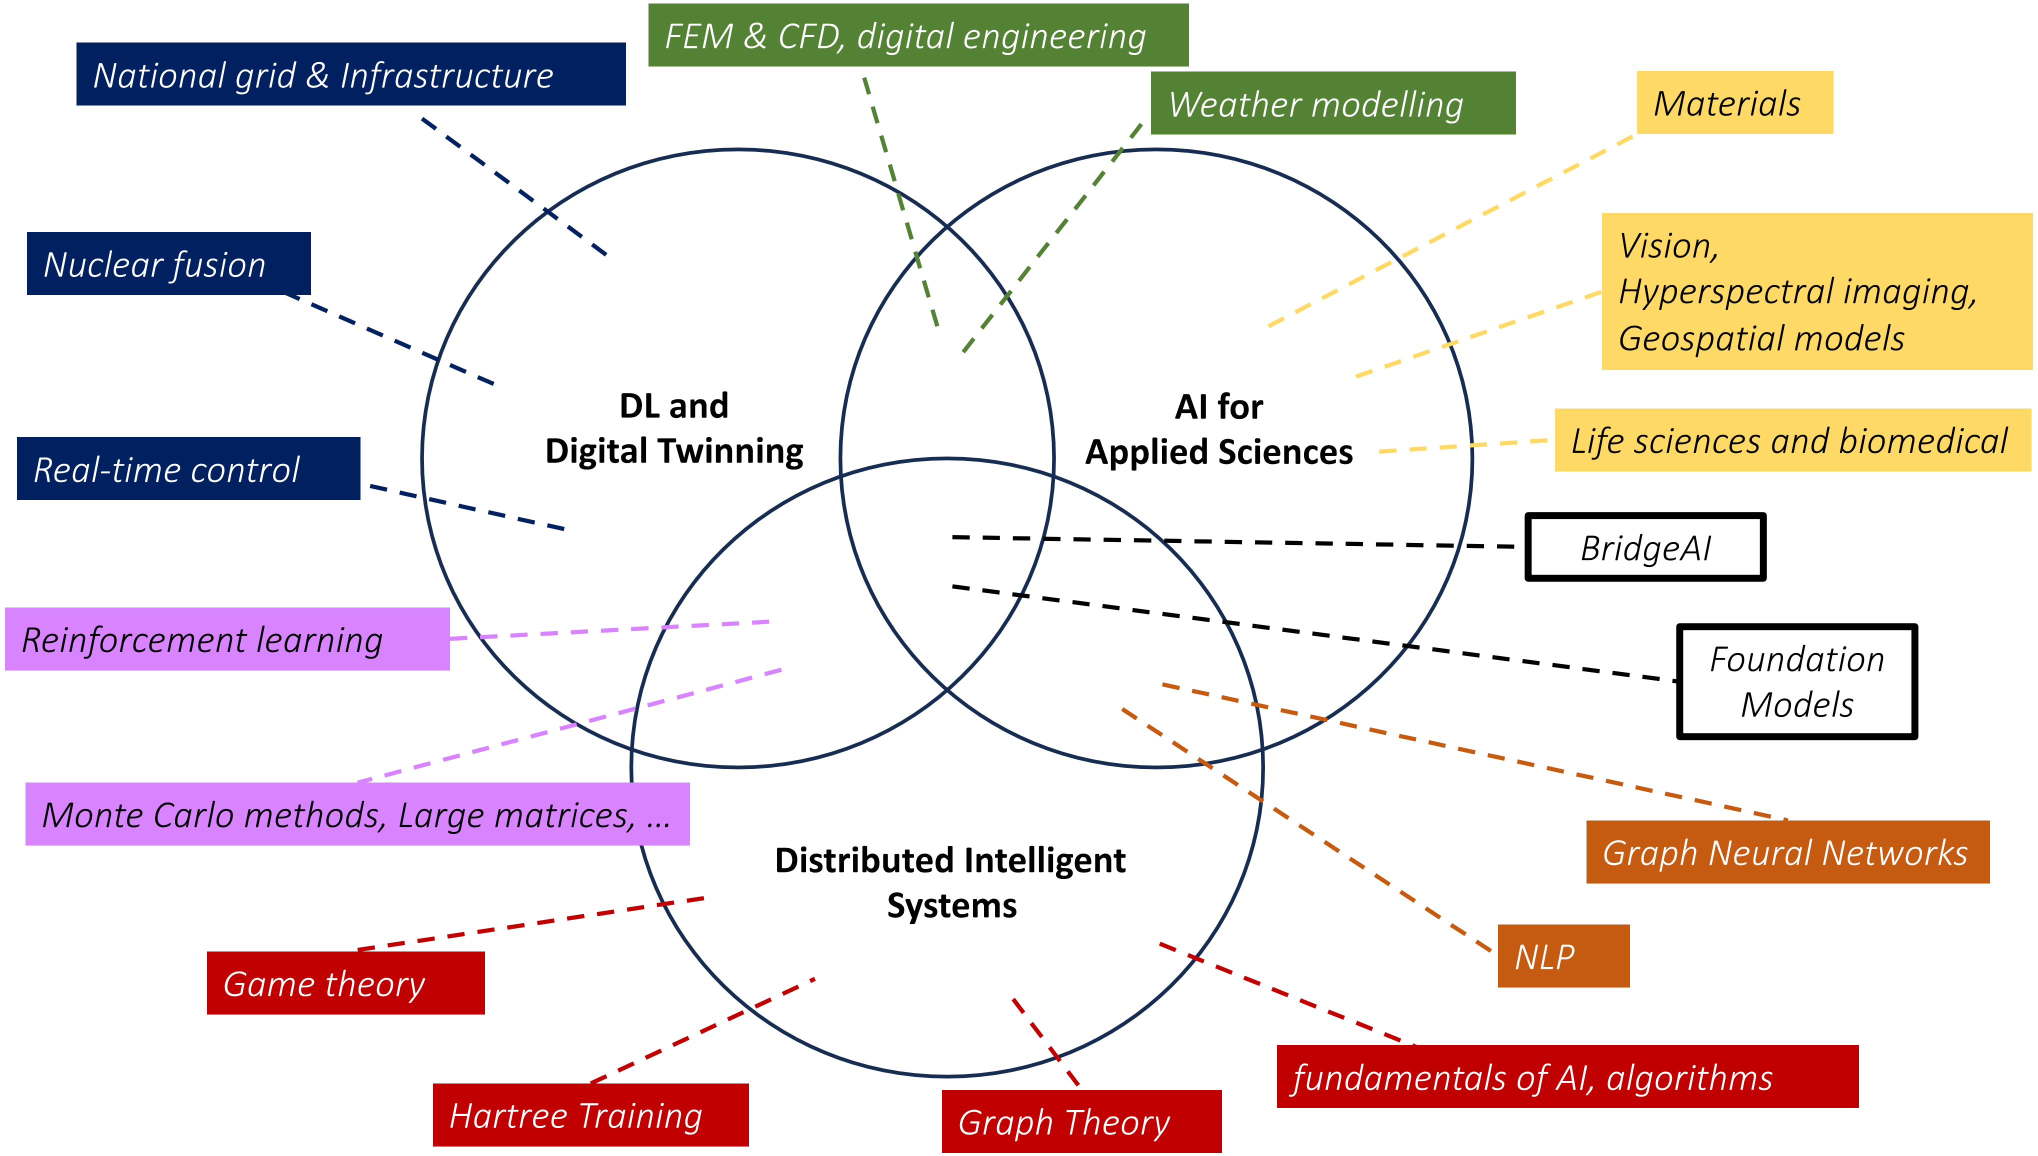

Science and
Technology
Facilities Council
Hartree Centre

The Fusion Computing Lab

A collaboration with digiLab and the UKAEA

How to design
the fusion power plant of the future?

Iterative cycle
design – build – improve
is unfeasible

Can we get a full *in silico* replica
of a nuclear fusion reactor?

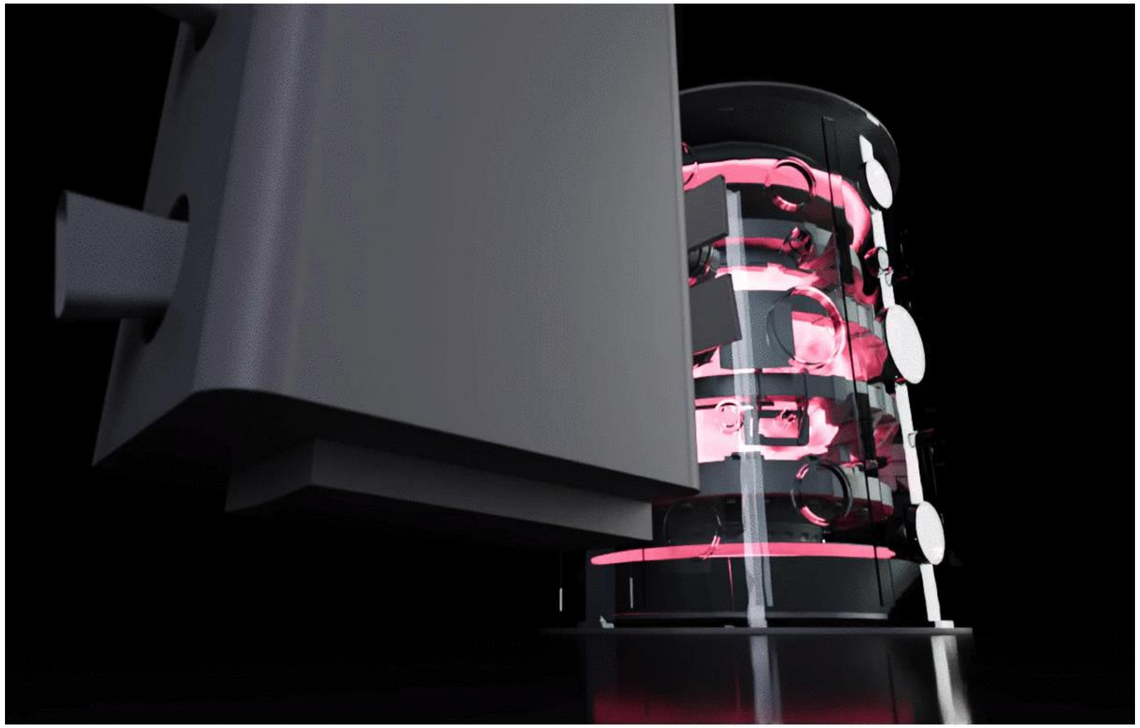

Science and Technology Facilities Council
Hartree Centre

digiLab

UK Atomic Energy Authority

Fusion Computing Lab

4-year collaborative agreement,
aiming to continue beyond 2027

Around 60 individuals (mostly researchers + operational support, communications, etc.),

Evenly split between STFC and UKAEA

Secondments of UKAEA researchers at Hartree

Aim: to explore and implement advanced computing methods in nuclear fusion research, and workflows towards the development of a reactor digital twin

Agreements also with academic institutions (incl Univ. Manchester and University College London), US National Labs (ESCAPE Project), and SMEs

Fusion Computing Lab

Five "Work Streams"

WS-2:
Fast and Actionable Emulators
using AI to replace expensive simulations with accurate surrogate models,
and to optimize the simulations to be run

WS-1:
Digital Thread co-integration
bringing it all together to build a digital twin

WS-3:
Plasma Real-Time Control
AI-powered algorithms
to keep the plasma confined and safeguard the surrounding structures

WS-4:
Platform Architecture Exploitation
high-performance computing SW and HW solutions

WS-0:
Project Coordination, Training, Communication

WS-5:
Uncertainty Quantification
guarantee safe operations and understand sensitivity to design changes

Science and Technology Facilities Council
Hartree Centre

digiLab

UK Atomic Energy Authority

AI for Magnetic Confinement Fusion

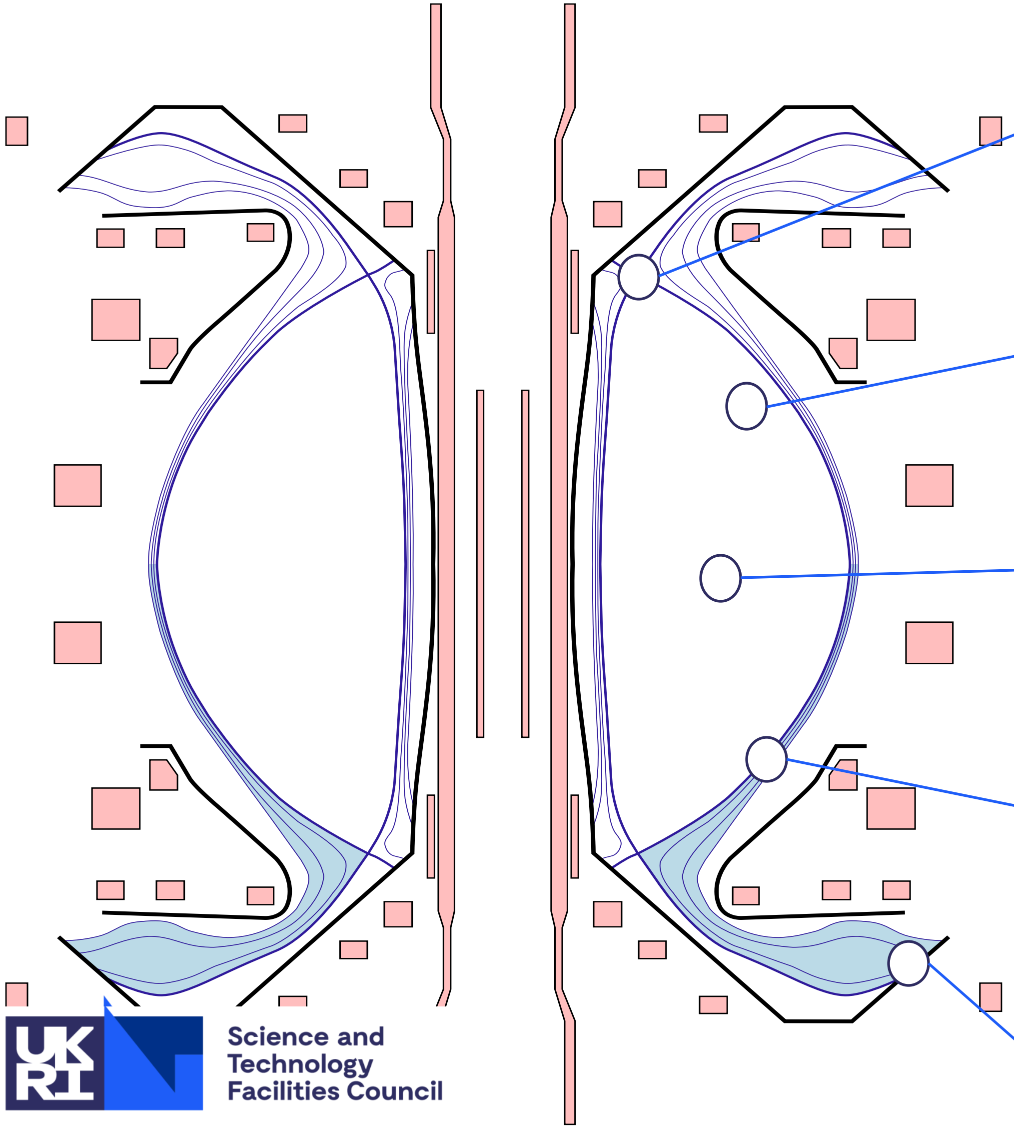

Plasma shape control
How do we keep the plasma where we want it to be?

Disruptions
Can we predict/avoid/mitigate plasma instabilities?

Core Turbulence and Gyrokinetics
How much energy is transported inside the “burning” plasma?

Scrape-off Layer
How much energy is transported away from the “burning” plasma?

Divertor Detachment
How do we cool the exhausts so not to damage the tokamak?

Science and Technology Facilities Council
Hartree Centre

Surrogate Modeling

Science and
Technology
Facilities Council

Hartree Centre

Plasma Shape Control

Real time magnetic control of 2D shape of plasma in the poloidal plane.

High frequency control of actuator coil currents magnetically coupled to the plasma.

Generally tackled using linear control techniques.





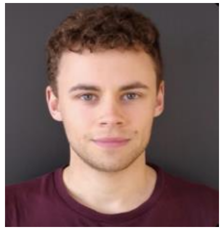


Hartree AI researchers:
Nicola Amorisco, Adriano Agnello,
George Holt, Abbie Keats,
Alasdair Ross, Aran Garrod

UKAEA Collaborators:
Stan Pamela, James Buchanan,
Graham McArdle, Charlie Vincent,
Kamran Pentland, etc.

example MAST-U equilibrium

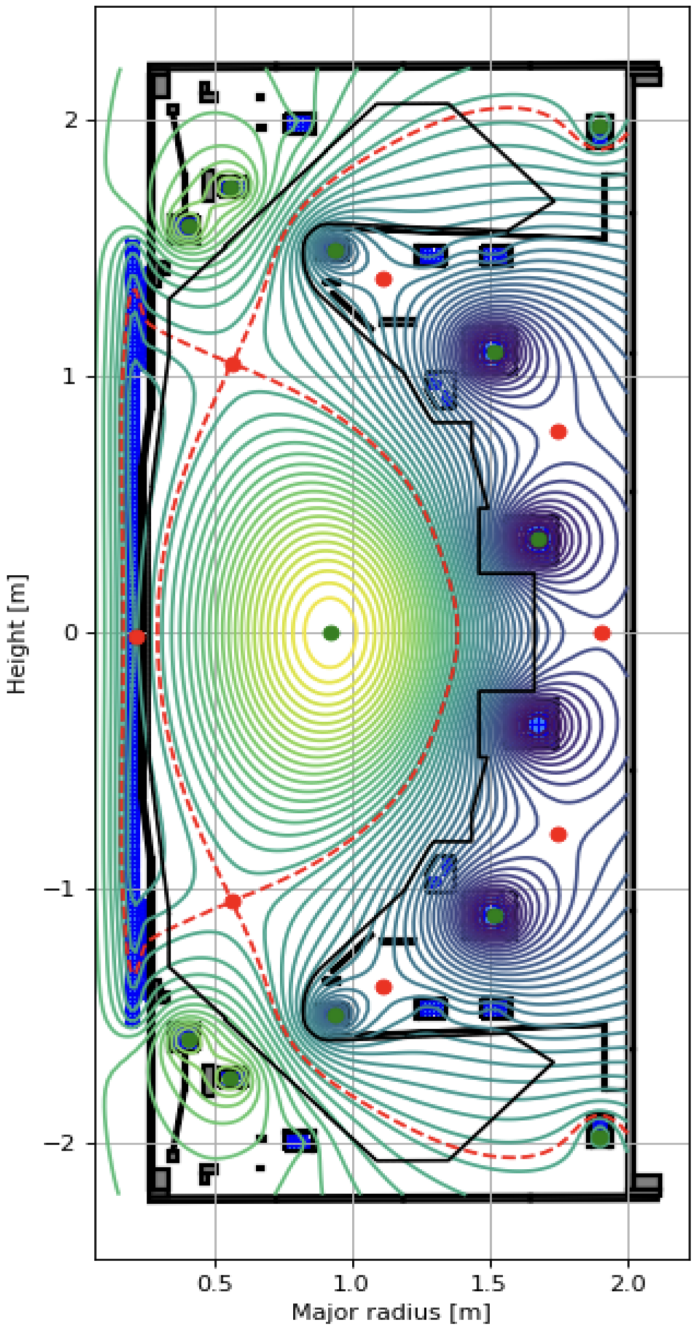

Science and Technology Facilities Council

Hartree Centre

(Ohmic) Plasma Shape Control

The plasma is confined by currents in the coils and in the plasma itself.

The "solenoid sweep" and other drives are preprogrammed to keep the plasma current up, but they can also alter its shape.

Probes around the tokamak give us noisy and incomplete information on the plasma. Is it departing from what we designed? And, how do we bring it back where it should be?

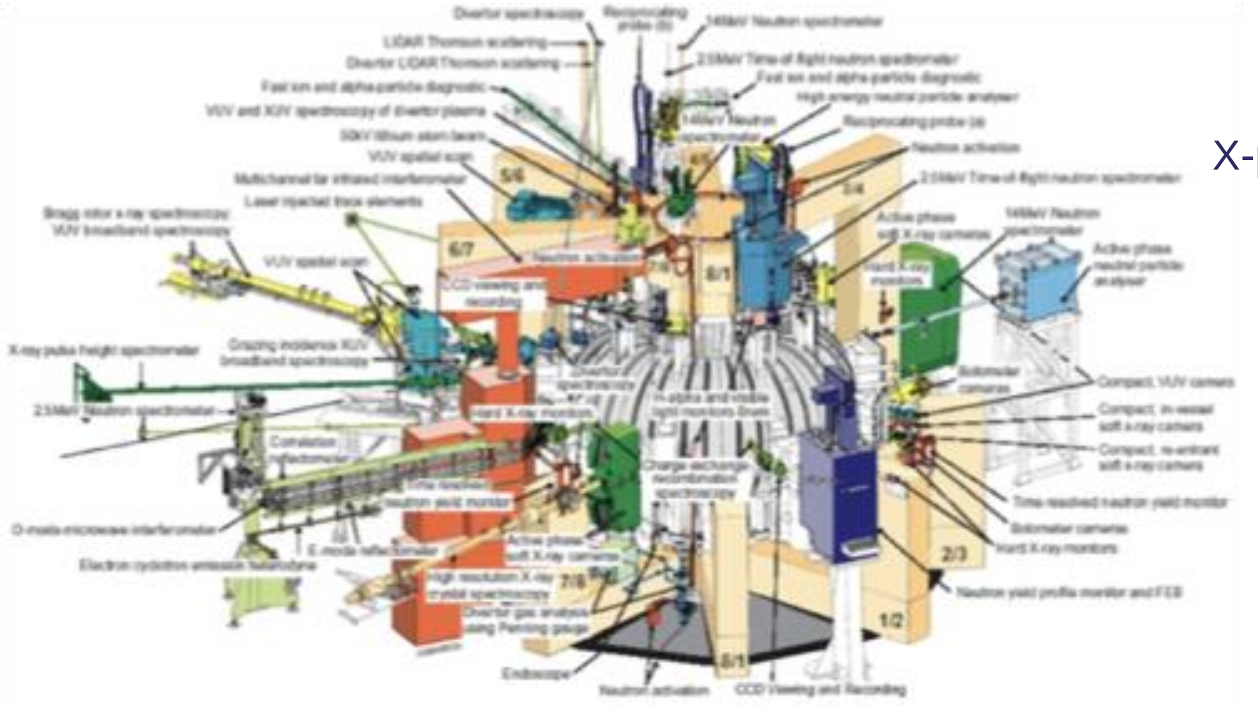

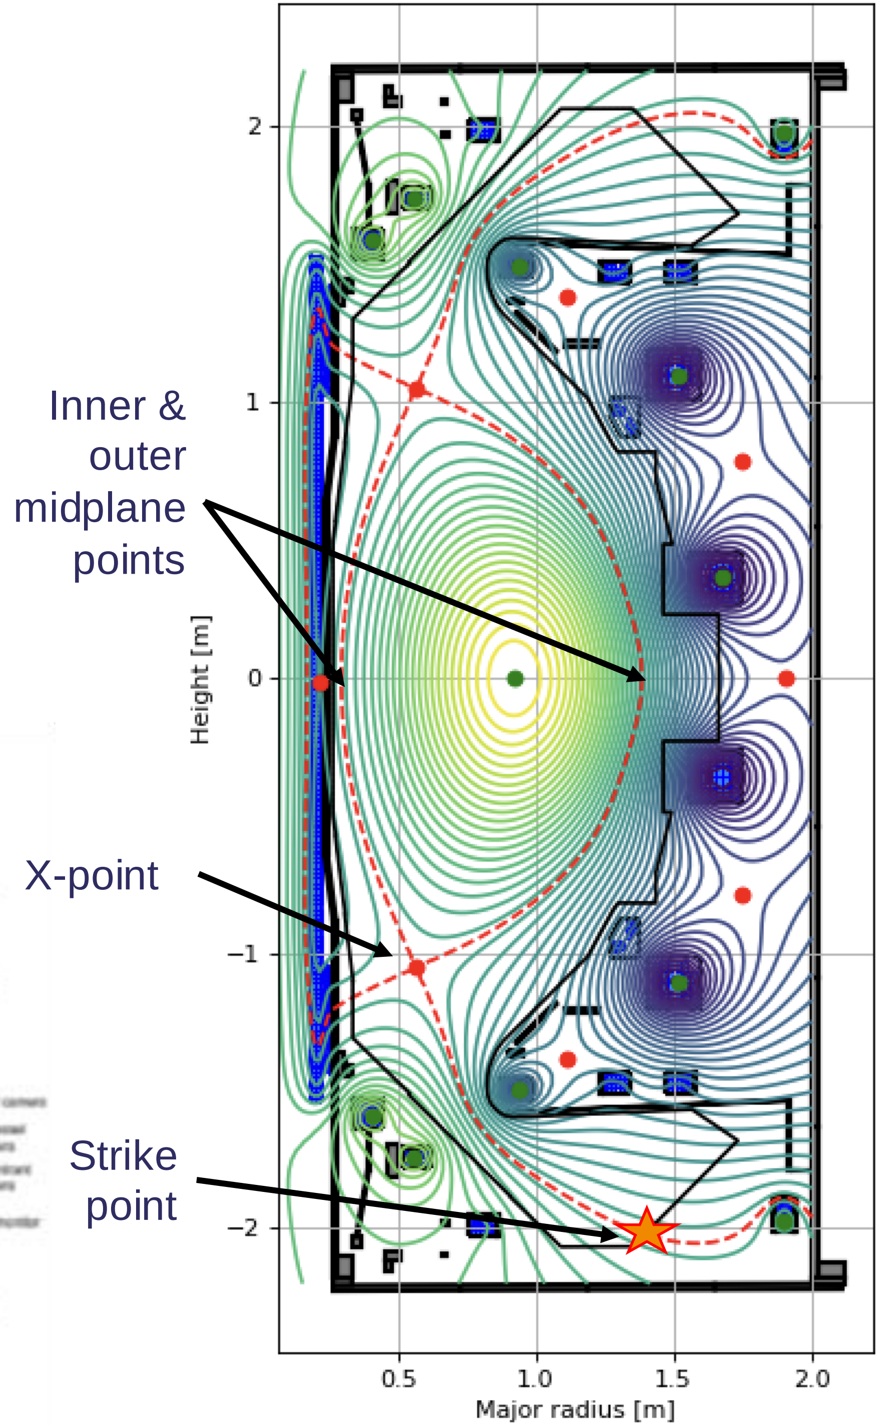

Science and Technology Facilities Council

Hartree Centre

‘Classical’ Plasma Shape Control

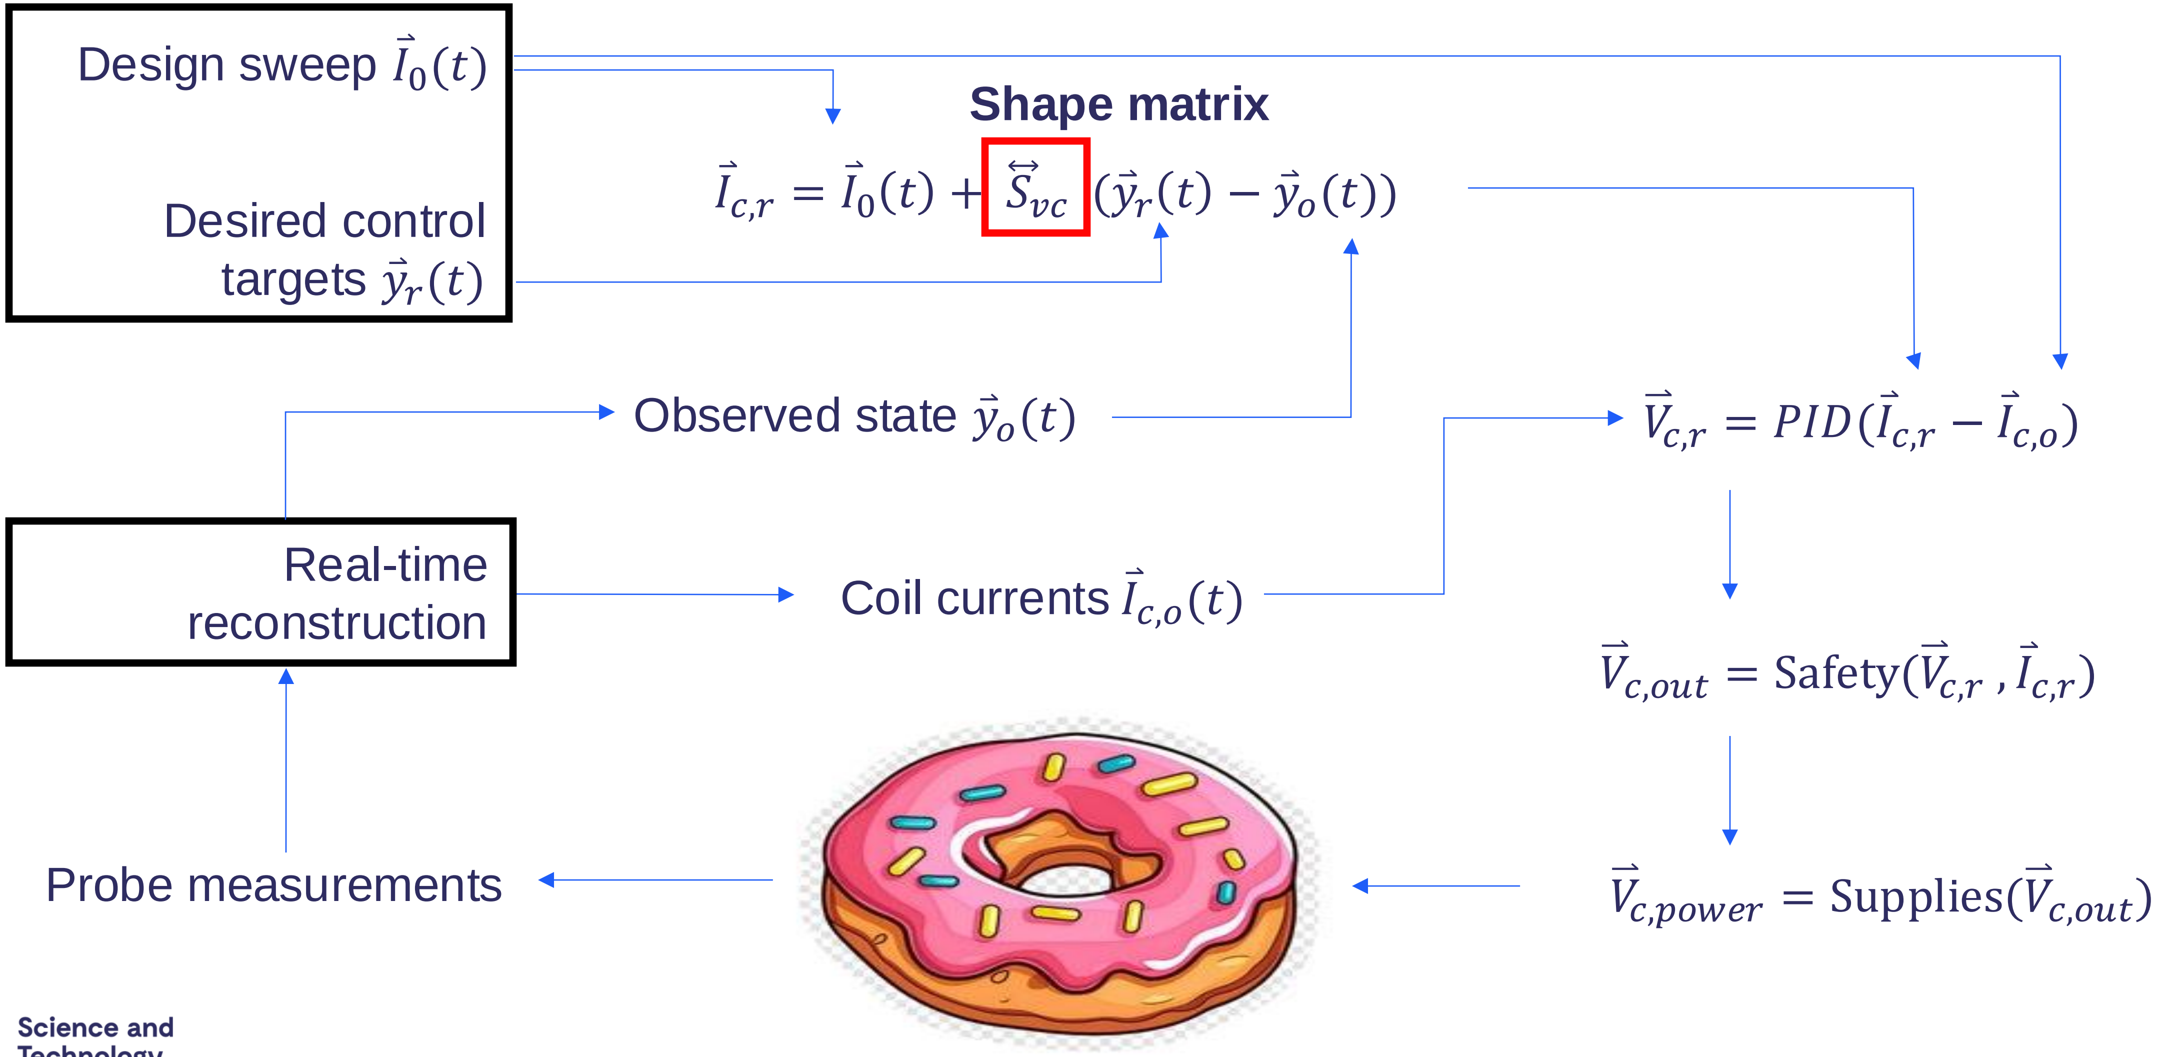

with the Shape matrix pre-calculated along the desired scenario.

AI-supported Plasma Shape Control

'Classical control'

AI / ML supported control

Wall/Passive currents are neglected

Dynamic
Wall/Passive currents are NOT neglected

Current framework

Shape matrix emulation

Reinforcement Learning

- GS surrogate: shape target emulator
- Accurate Shape matrices at any point in the experiment
- Ensemble averages

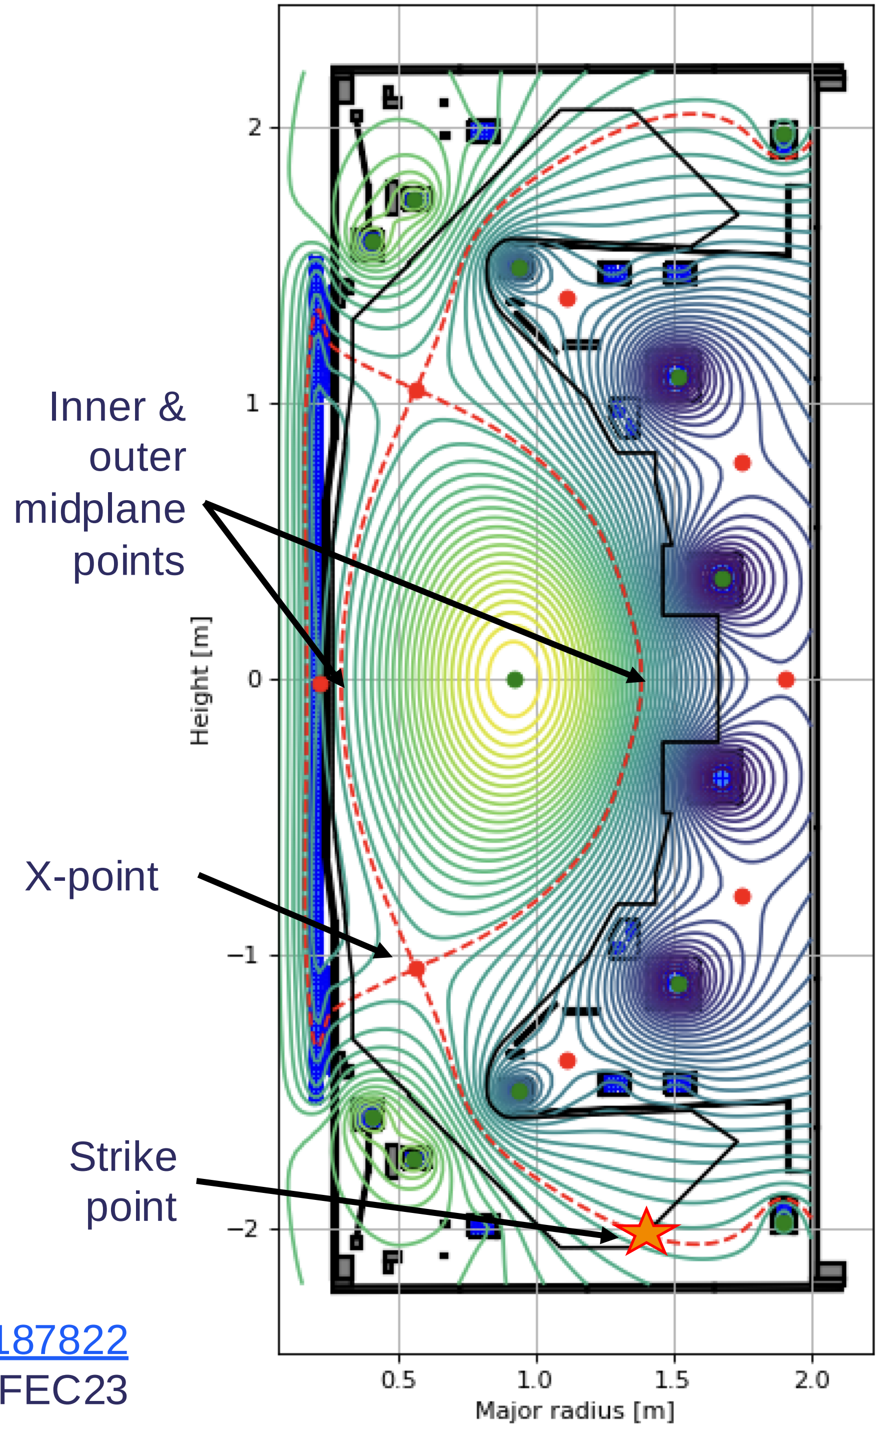

Science and Technology Facilities Council
Hartree Centre

Agnello et al. (2024), *Phys.Plasmas*, doi:10.1063/5.0187822
Amorisco et al. (2023), ICDDPS-4 and IAEA-FEC23

FreeGSNKE: FreeGS Newton-Krylov Evolutive

Fully Python non-linear solver for the evolutive equilibrium problem. Extends Ben Dudson's FreeGS.

1. **Static GS solver**:
    - forward-solve of Grad-Shafranov eq., Newton-Krylov method.
2. **Linear dynamics:**
    - Automated normal mode decomposition of passive structure model
    - Linear stability analysis, linear growth-rate of vertical instability
3. **Non-linear dynamics:**
    - NK-based solver of fully non-linear problem
    - Prescribed time evolving profiles, parameterized by $p_a(\mathrm{t})$ or $\beta_p(\mathrm{t})$, evolving $\alpha_m(t), \alpha_n(t)$

FreeGSNKE-RL library for RL experiments and training.

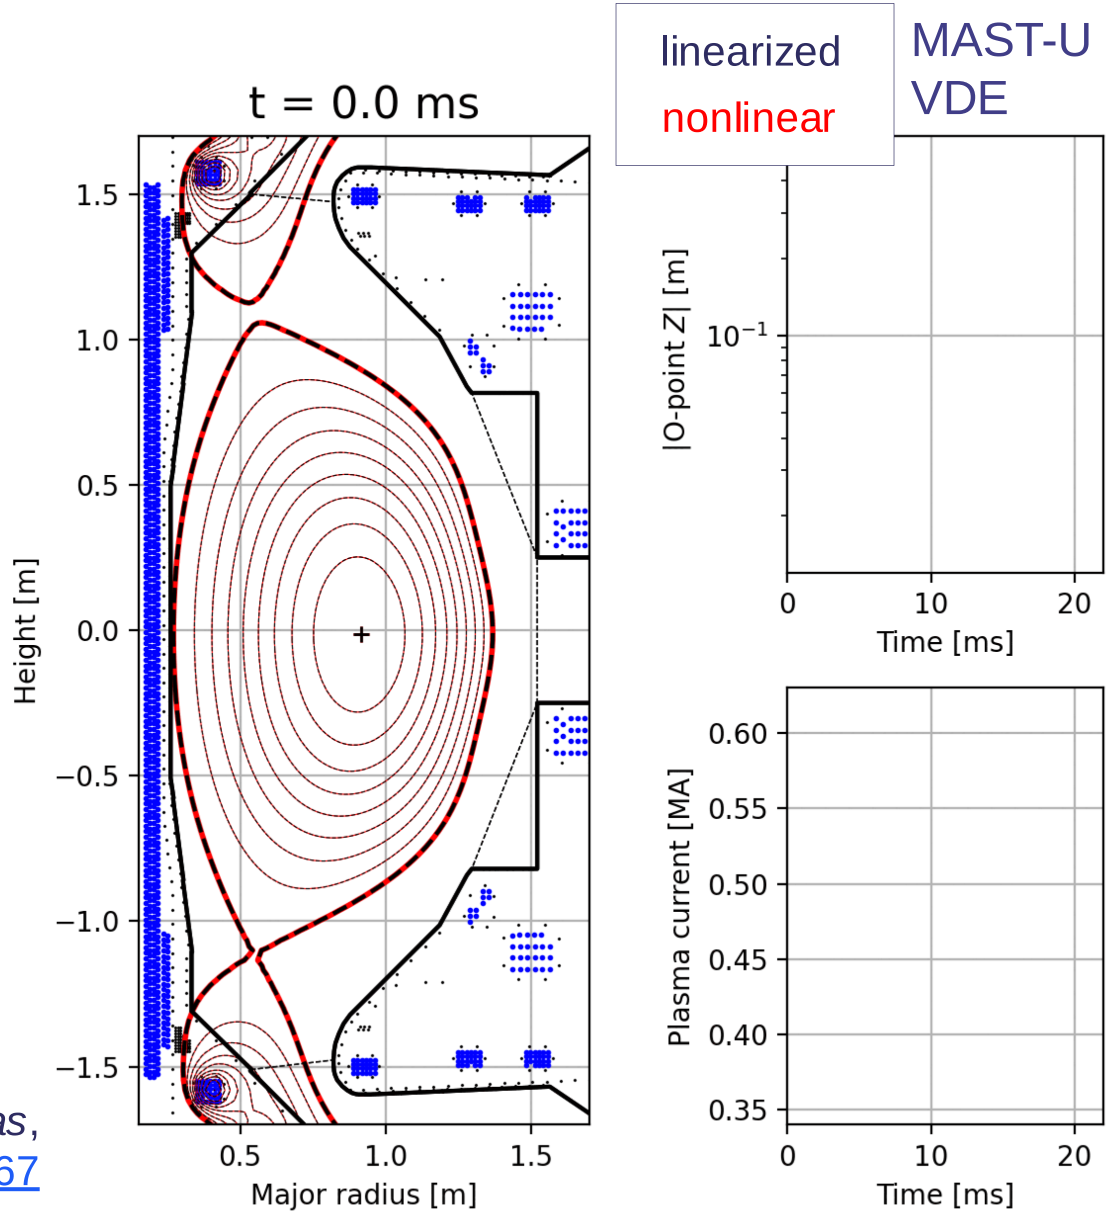

Amorisco et al. (2024), *Phys.Plasmas*,
doi:[10.1063/5.0188467](https://doi.org/10.1063/5.0188467)

Science and Technology Facilities Council

Hartree Centre

Divertor Detachment


Hartree: George Holt, Abbie Keats
Collaborators: Stan Pamela[1], Mike Kryjak[2], Ben Dudson[3], Lorenzo Zanisi[1]
[1]UKAEA, [2]Univ. York, [3]LLNL

Background

- Tokamak plasma exhaust is extremely energetic
- There are no materials that would withstand unmitigated deposition of the exhaust
- Advanced divertor configurations are being designed and tested to reduce the energy load on exhaust components

Problems we are addressing

- Scrape-off layer and divertor simulation is computationally expensive but can be massively sped up with machine learning
- Current control policies are based on linear theory but can potentially be improved with nonlinear policy development

Led by Hartree Centre and UKAEA, contributions from Lawrence Livermore National Laboratory and University of York

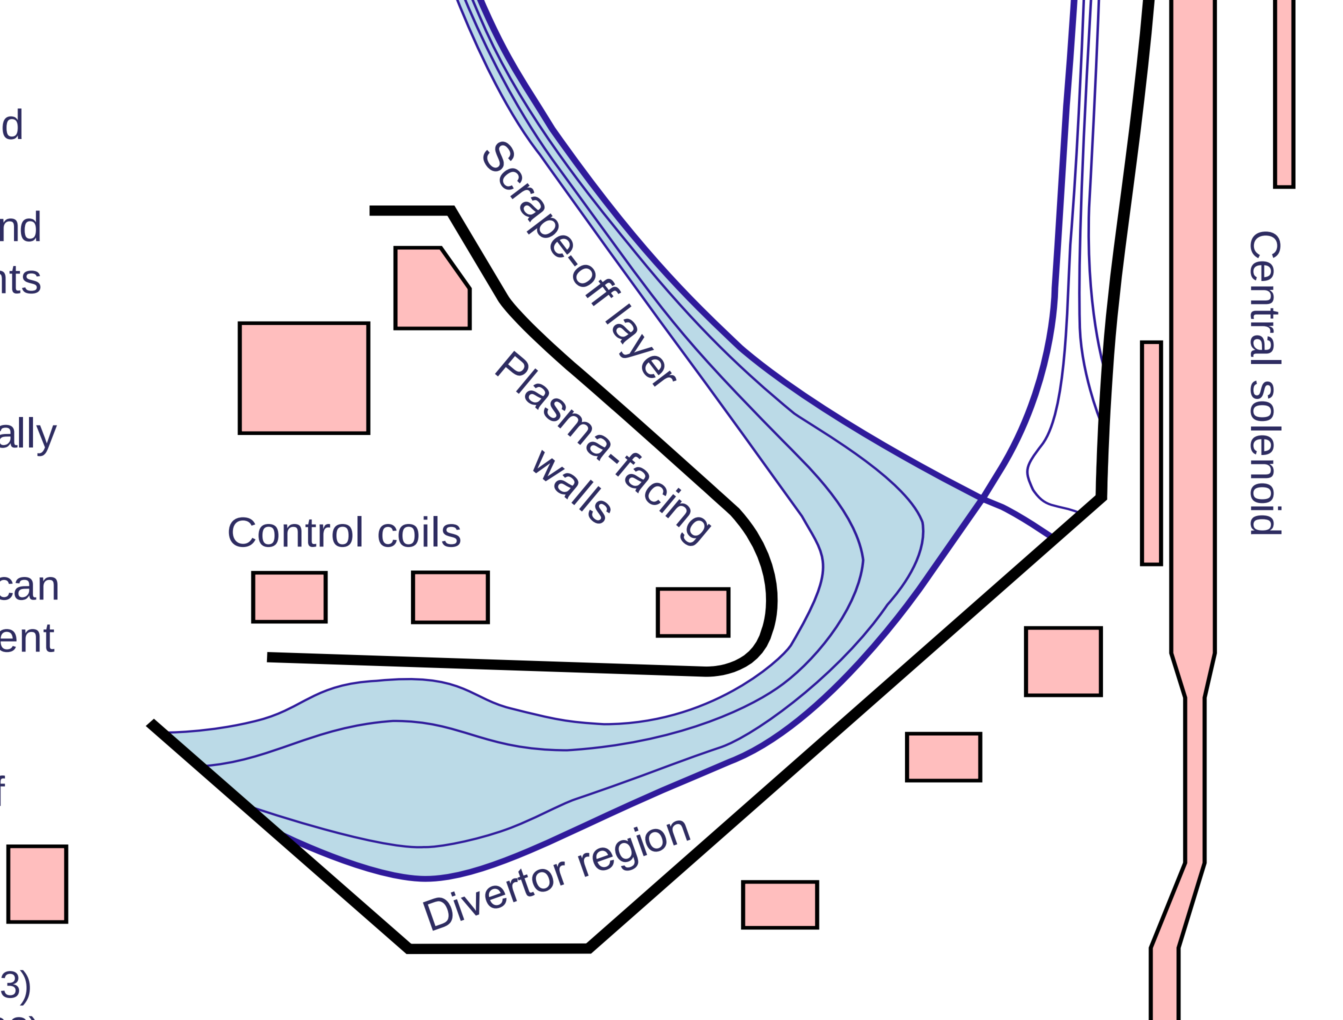

Quarter cross-section of the MAST-U tokamak

UKRI Science and Technology Facilities Council
Hartree Centre

G. K. Holt, et al, ICDDPS-4 (2023)
G. K. Holt, et al., IAEA-FEC (2023)
A. Keats, et al., ICDDPS-5 (2024)
G. K. Holt, et al., under review (2024)

Divertor Detachment

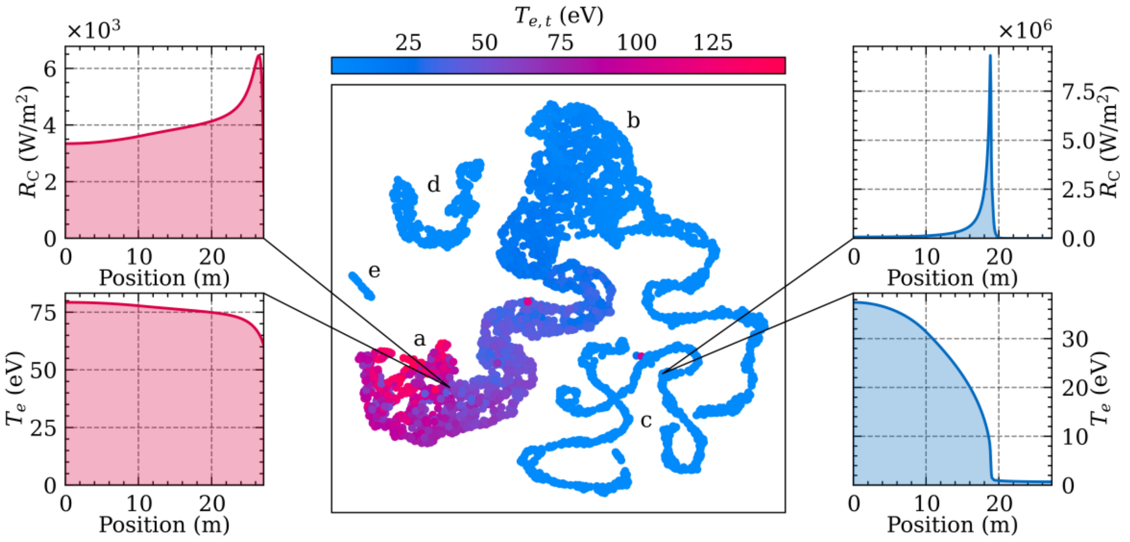

Top: UMAP target visualisation. Bottom-left: job placement diagram. Bottom-right: search space initialization schematic.

Data set creation

Automation → scaling → HPC exploitation

Automation

- Simulation input generation
- Convergence testing
- Diagnostic cleanup

Scaling

- Trade-off between efficiency and wall time is problem dependent

HPC exploitation

- Search space initialisation
- Simulation batching
- Array jobs

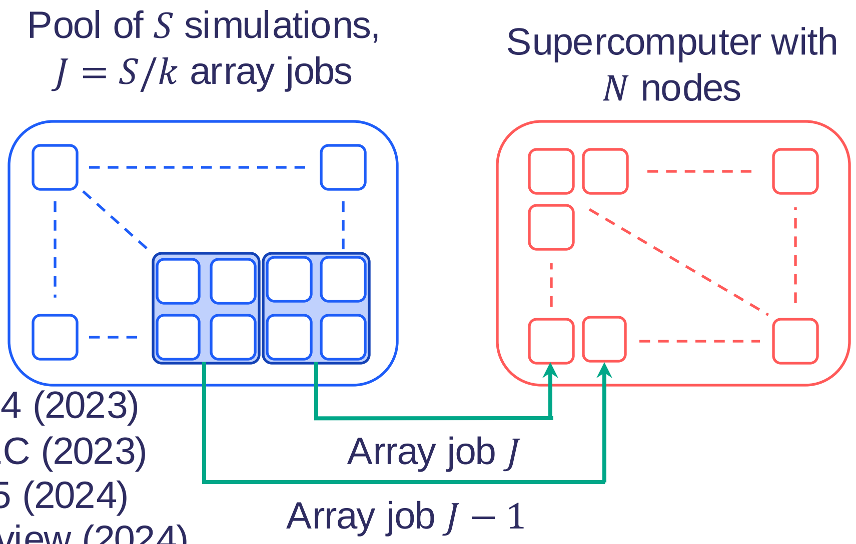

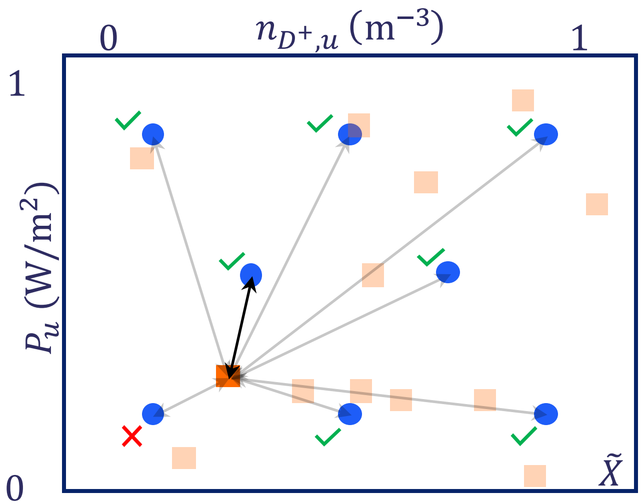

Science and Technology Facilities Council

Hartree Centre

G. K. Holt, et al, ICDDPS-4 (2023)
G. K. Holt, et al., IAEA-FEC (2023)
A. Keats, et al., ICDDPS-5 (2024)
G. K. Holt, et al., under review (2024)

Divertor Detachment

Neural network training, results and interpretability

Rigorous hyperparameter optimisation

- Tree-structured Parzen estimator for trial selection
- Asynchronous hyperband scheduler for culling
- Automated experiments, run to convergence

Model performance

- R^2: 0.98
- Time-to-solution reduced from ~1 day to ~1 ms

SHAP analysis

- Shapley additive explanations for global and local model interpretability

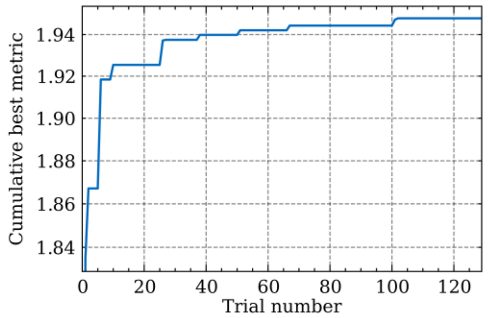

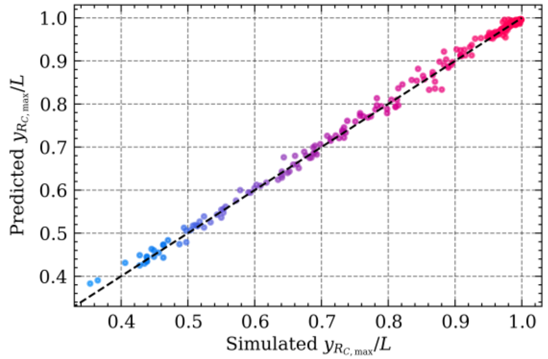

Top-left: hyperparameter tuning experiment progress. Top-right: trained model calibration plot. Right: Local SHAP interpretations. Bottom: Global SHAP interpretations.

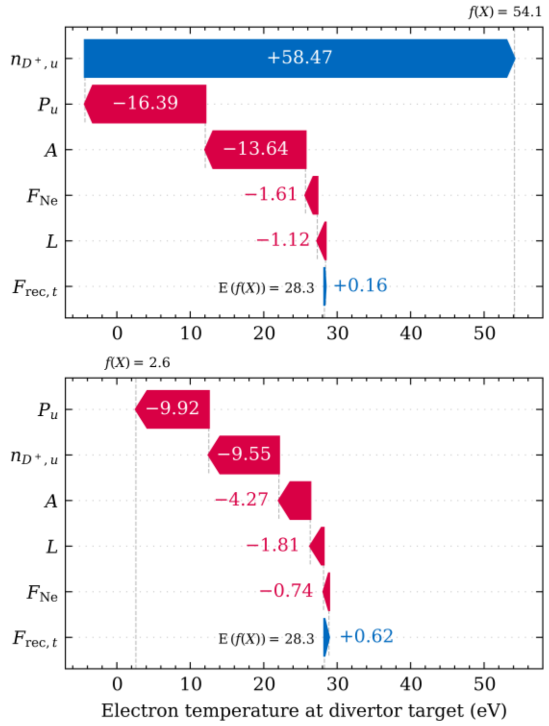

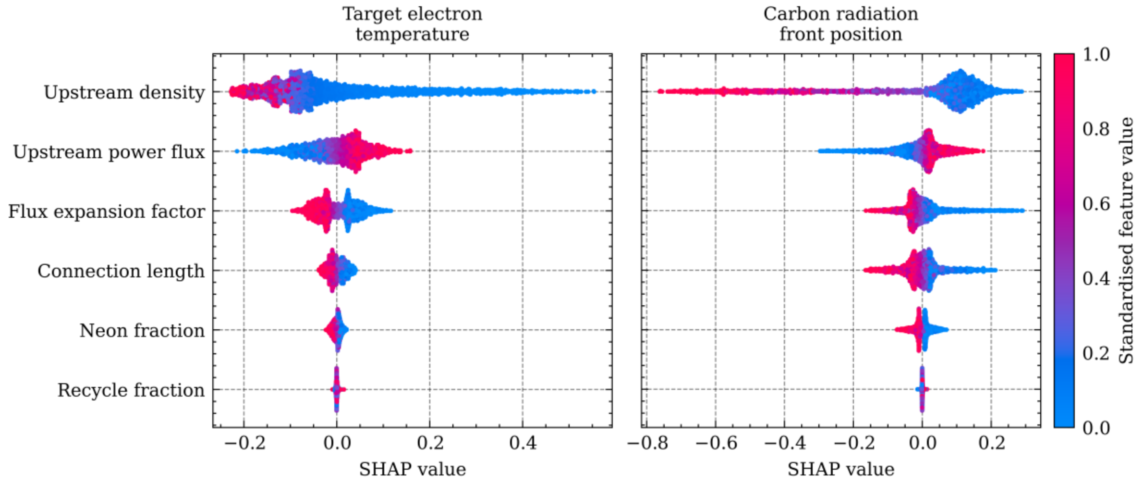

G. K. Holt, et al, ICDDPS-4 (2023)
G. K. Holt, et al., IAEA-FEC (2023)
A. Keats, et al., ICDDPS-5 (2024)
G. K. Holt, et al., under review (2024)

Science and Technology Facilities Council

Hartree Centre

Accelerating Gyrokinetic simulations

Transport in a plasma is governed by the phase space distribution *f*(*x*,*v*) of particles
Focus on slab ITG turbulence.

Model:	Difficulty:	Caveats:
Fluid	easy	No info on *f*(*x*,*v*)
(quasi)linear gyrokinetic	Hard, but doable	Can miss important physics
Nonlinear gyrokinetic	Very hard, expensive	Provides *f*(*x*,*v*) info **But**: how do we sample the parameter space fast and efficiently?

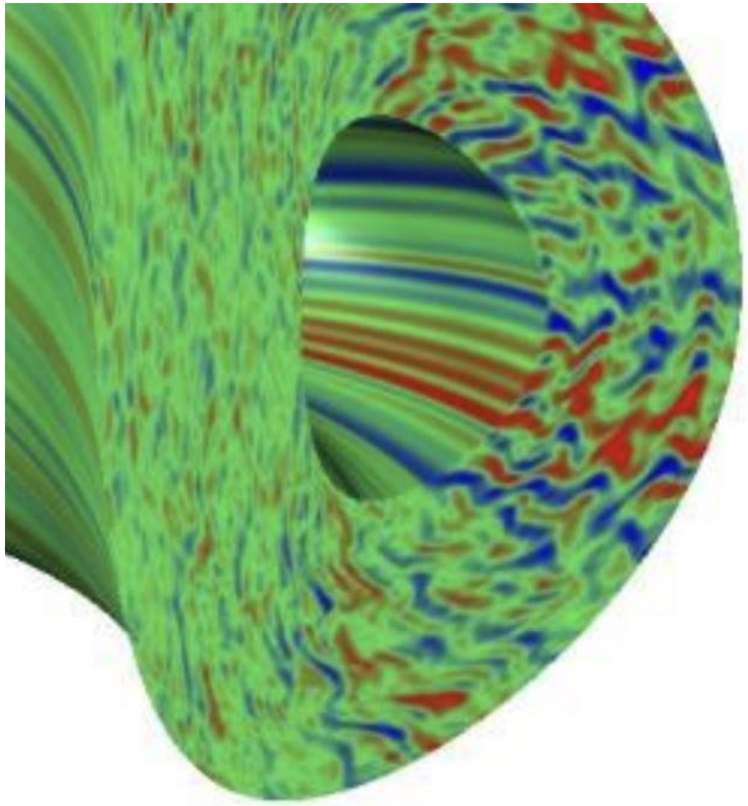

(Candy & Waltz, GA 2003)

Accelerating Gyrokinetic simulations

Closures for nonlinear slab ITG turbulence:
higher-order velocity moments as simple functions of lower-order moments.

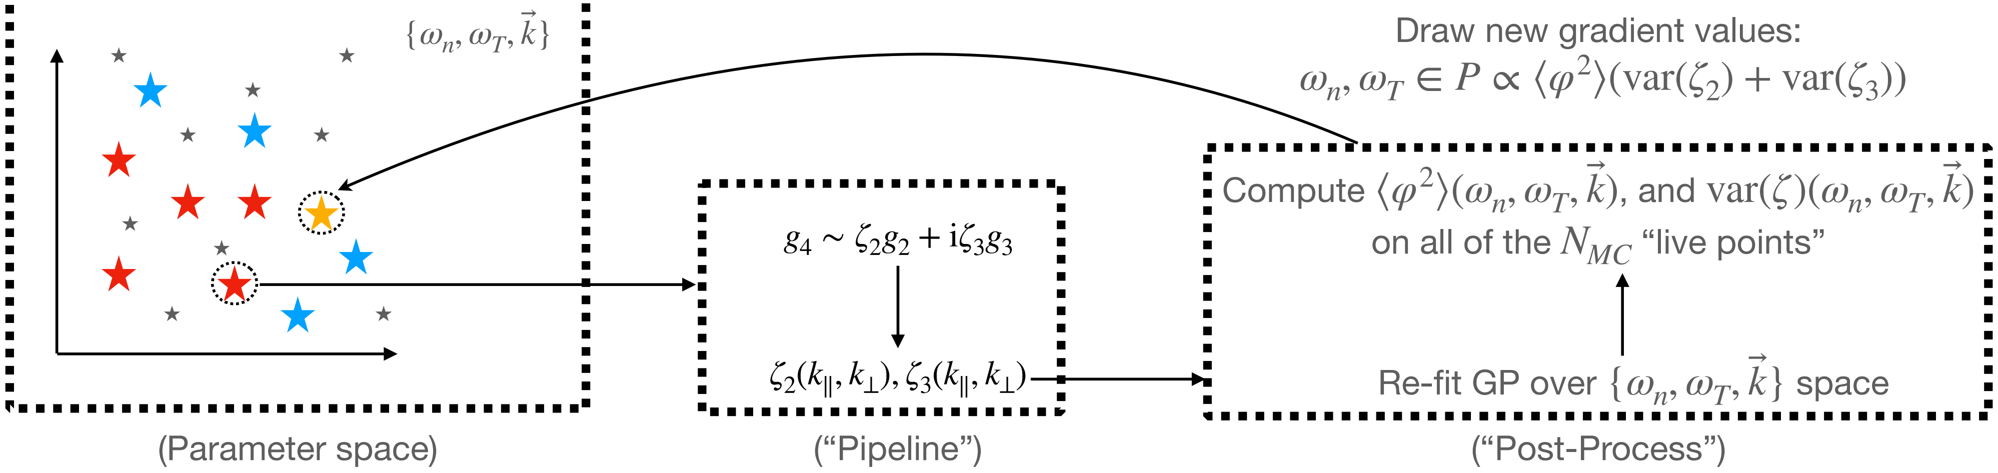

Simpler than full-geometry problem, but gives good insight.
It also shows some behaviour that was not caught in "paper and pen" linear-Landau closures.

Science and Technology Facilities Council
Hartree Centre

https://github.com/JosephThomasParker/SpectroGK

Hartree: Adriano Agnello
Collaborators: JT Parker[1,2], James Buchanan[1], William Hornsby[1],
[1],Irish Centre for HE Computing, [2]UKAEA

Active Learning on JET

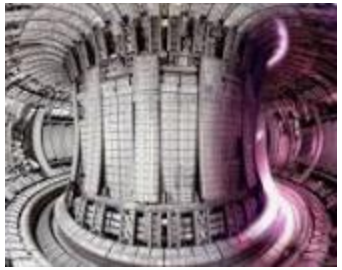

JET

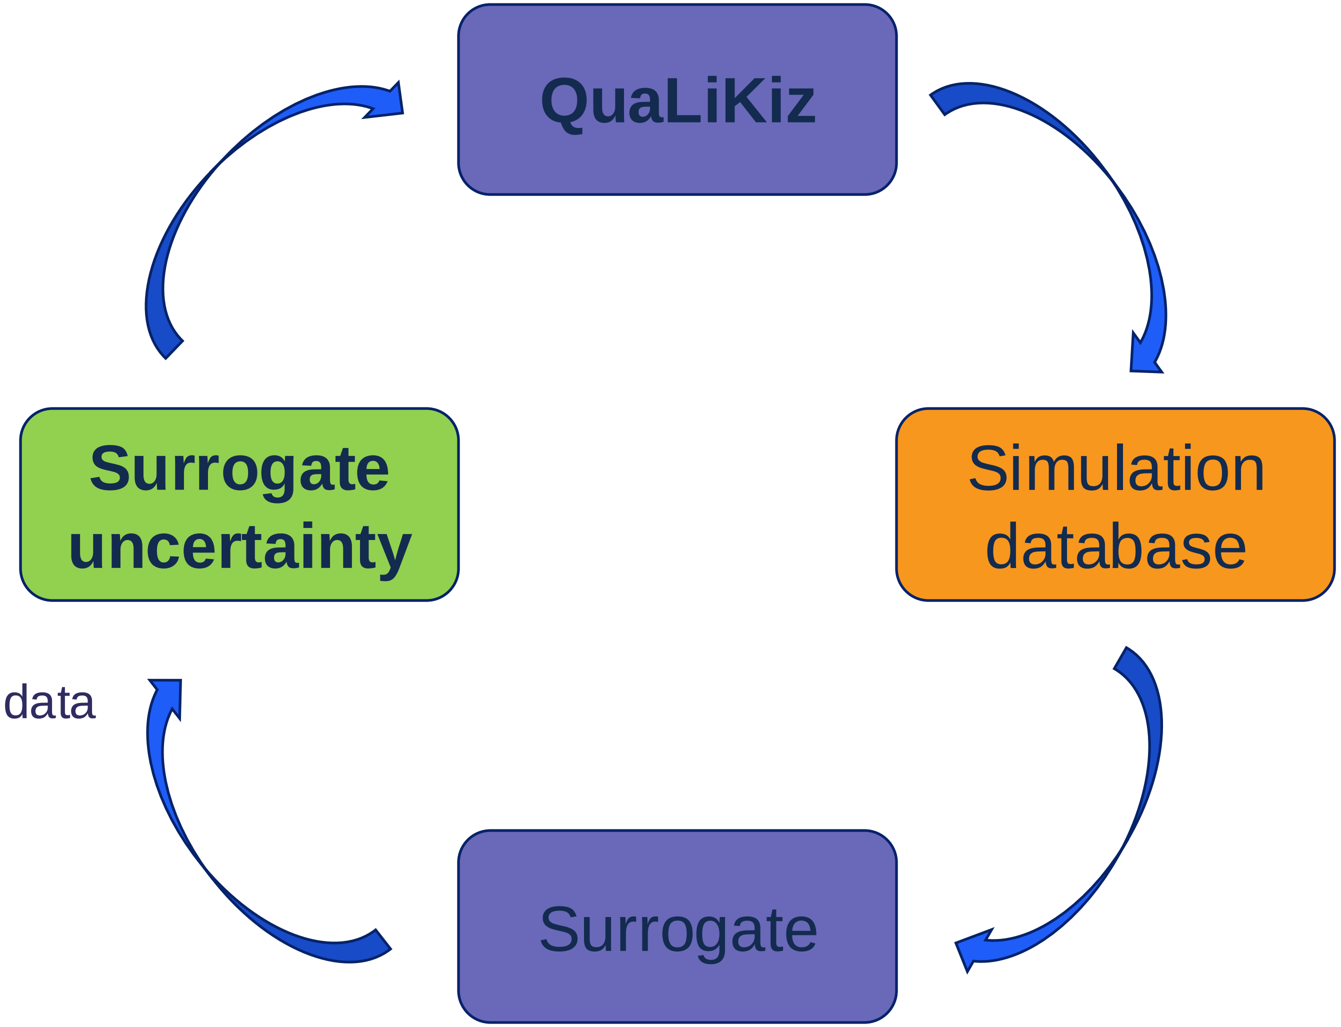

JET experimental 15D historical data
Fixed design space
Quasi-linear simulations

Zanisi et al. (2024)

Active Deep Ensembles for Plasma Turbulence

Two Deep Ensembles
(Lakshminarayanan et al. 2017)

- Classification
  - Critical gradient estimation
  - Prevents sampling in stable region
- Regression
  - ITG turbulence flux estimation

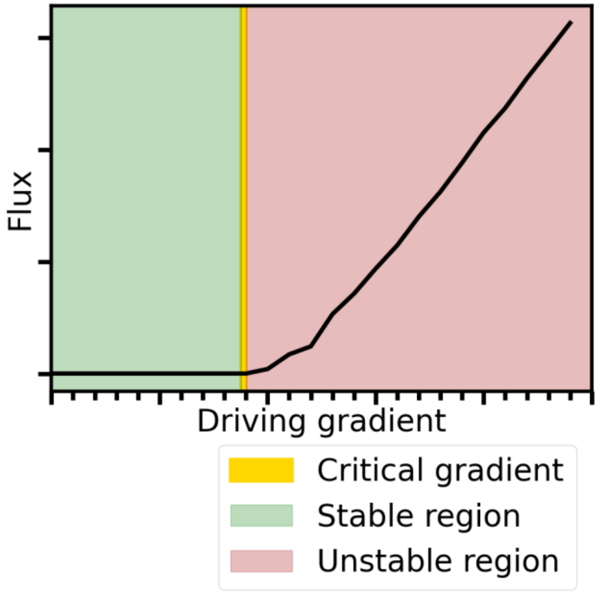

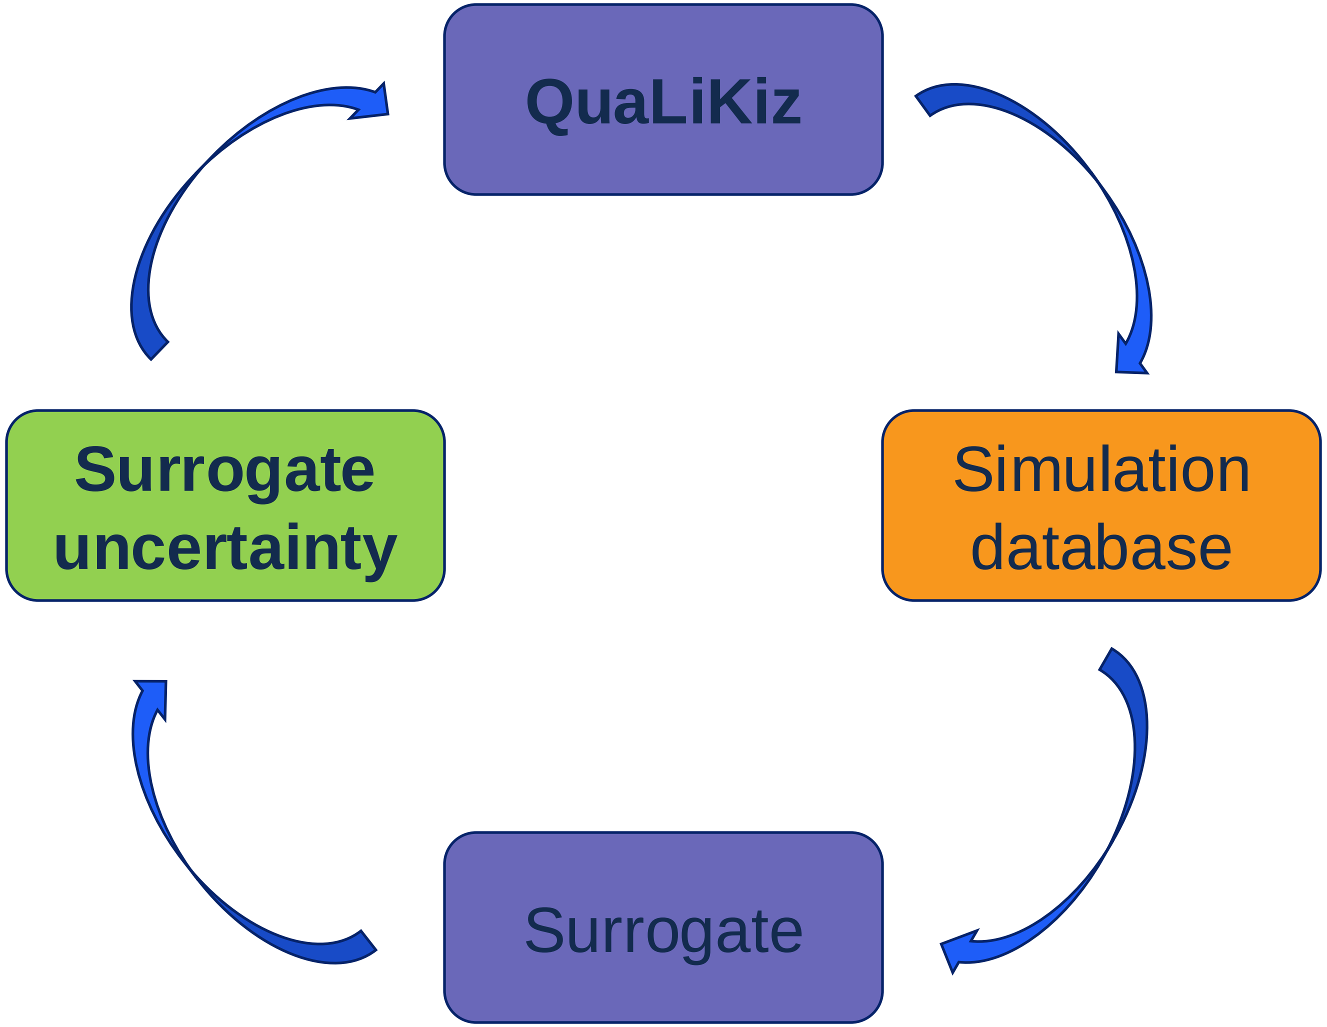

Zanisi et al. 2024

ADEPT Results on JET

Regressor performance
(ion particle flux)

R2

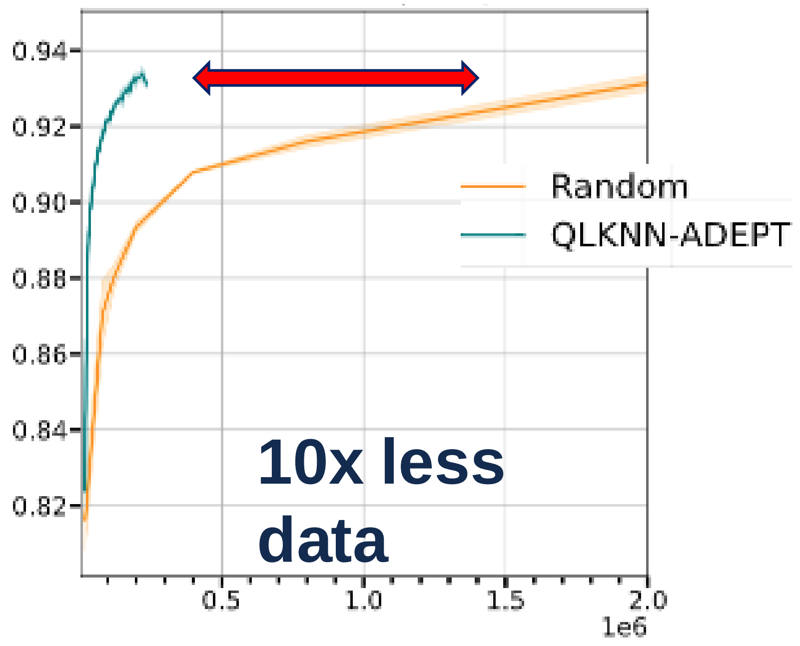

Number of training examples

Classifier performance

F1

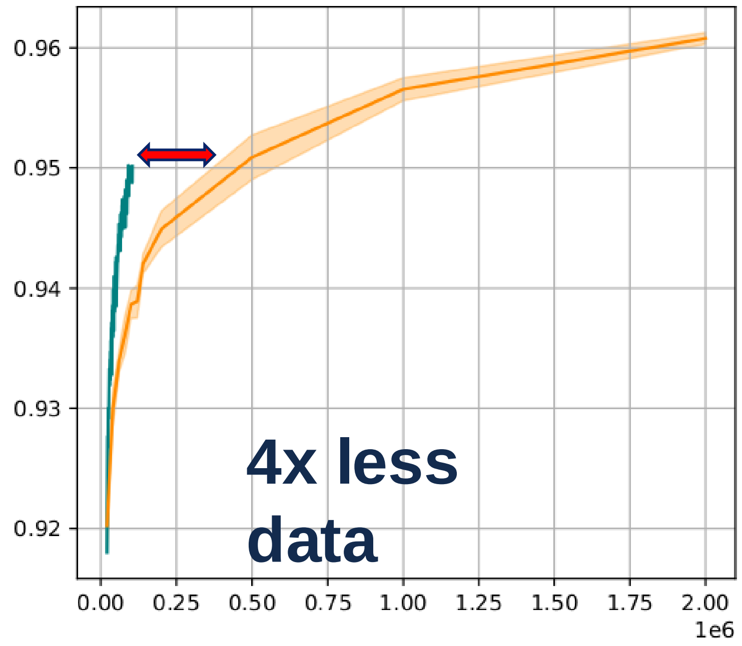

Number of training examples

Science and
Technology
Facilities Council
Hartree Centre

Zanisi et al. (2024) *Nuclear Fusion*
doi:10.1088/1741-4326/ad240d
Siddle et al. in prep.

StyleGAN for SOL Turbulence

Using Generative Adversarial Networks as deconvolution operator for Large Eddy Simulations

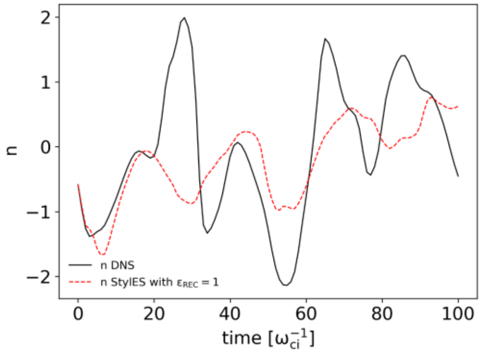

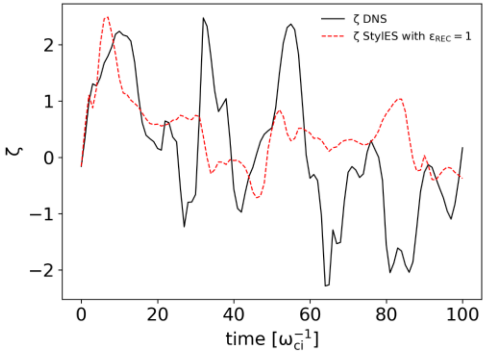

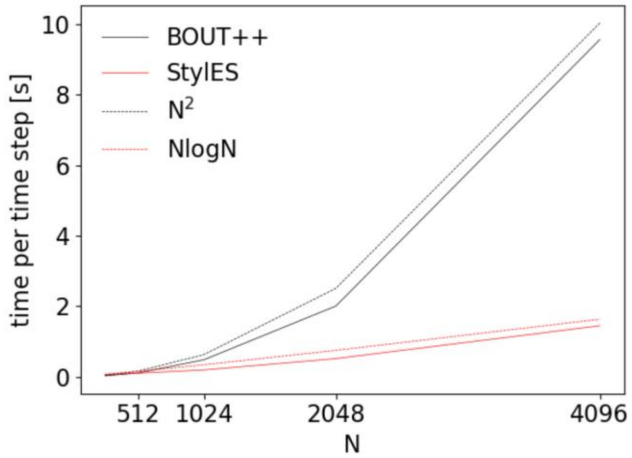

Speedup ~10x
and complexity NlogN
vs N^2 of BOUT++

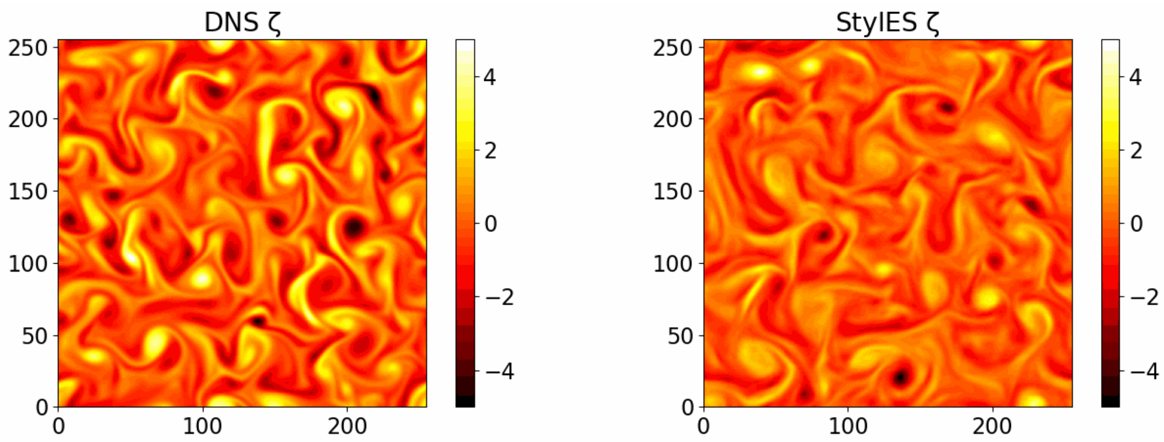

Castagna et al. (2024), doi:10.1063/5.0189945 *Physics of Plasmas*

Unsupervised Disruption prediction

Predict impending plasma disruptions based on plant diagnostics in MAST database

- **Unsupervised approach** vs literature work based on manually labelled data
- **Uncertainty-aware prediction** for robust inference and advance warnings for mitigation
- **Unsupervised pre-training** based on β-Variational Auto-Encoders
- Training tailored to ensure robustness to missing data
- RNN baseline prediction and customised transformers underway

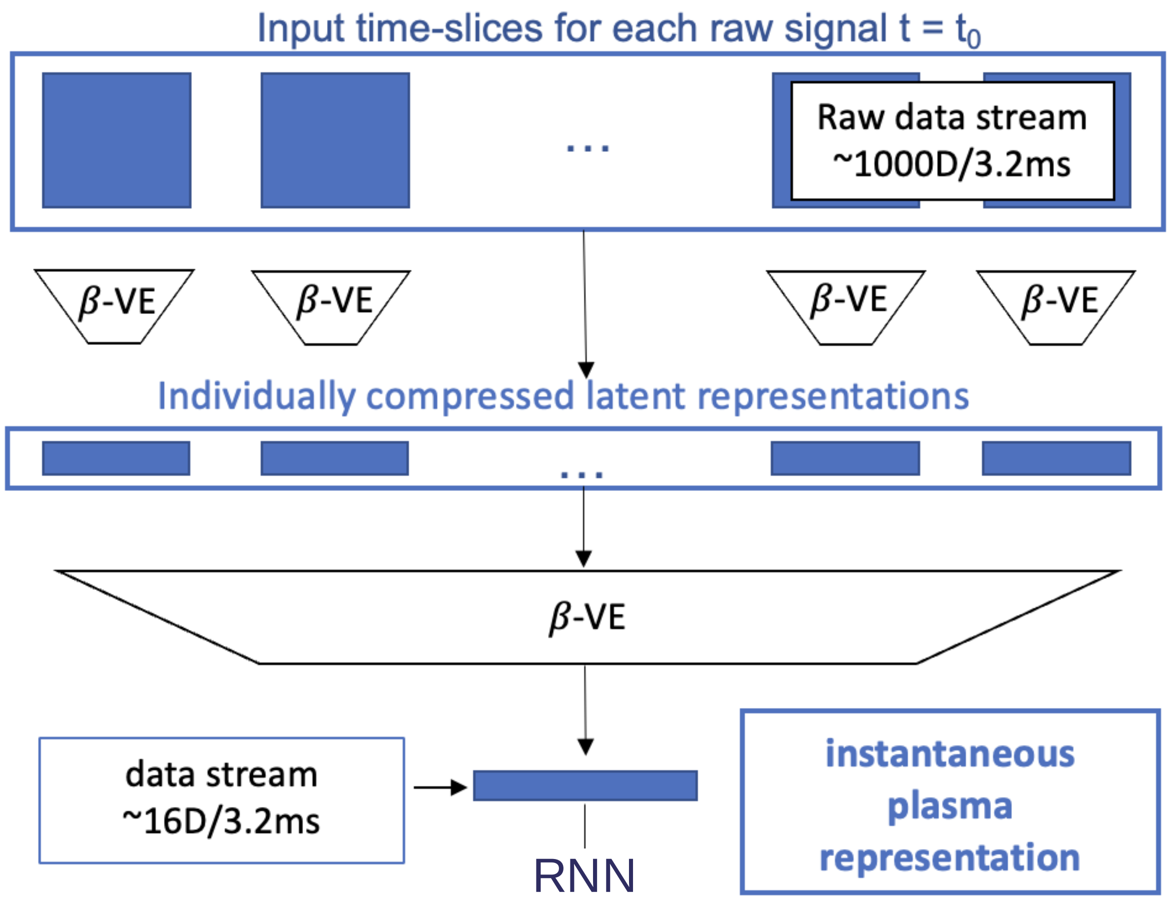

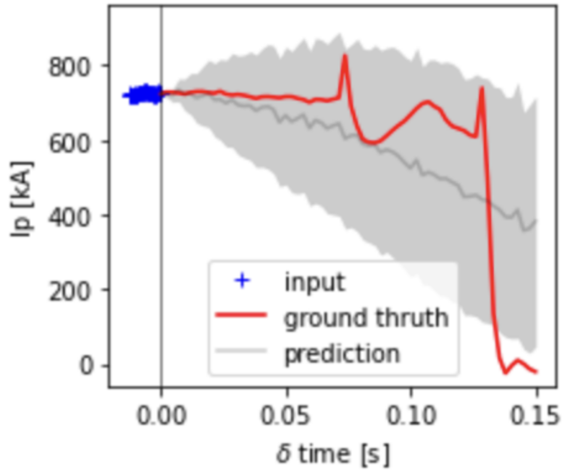

Plasma current prediction with uncertainty

Science and Technology Facilities Council
Hartree Centre


Hartree: N. Amorisco, Wonny Lee
UKAEA: Stan Pamela

AI for magnetic confinement fusion

Plasma shape control

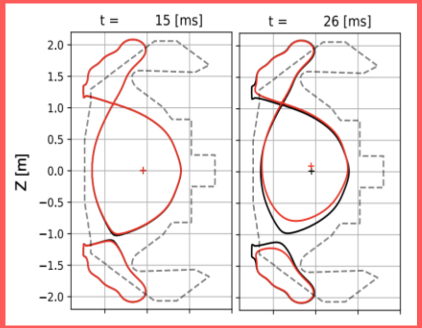

- Evolutive Grad-Shafranov solver
- Shape matrix emulation
- Reinforcement learning for control
- GPU support with JAX

Detachment control

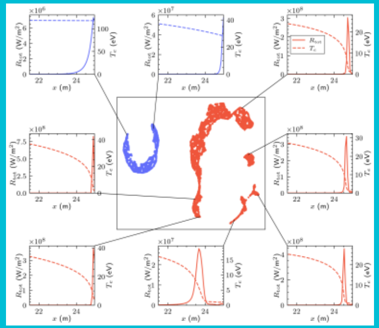

- Emulators of Hermes-3 & SD1D
- Efficiency with active learning
- Interpretability with SHAP
- Reinforcement learning for control

Gyrokinetics

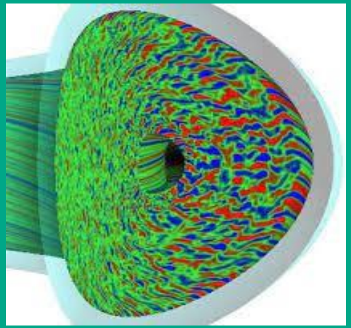

- Nonlinear slab gyrokinetic simulations of ITG turbulence
- Data-driven closures discovery
- Efficiency with active learning

Disruption prediction

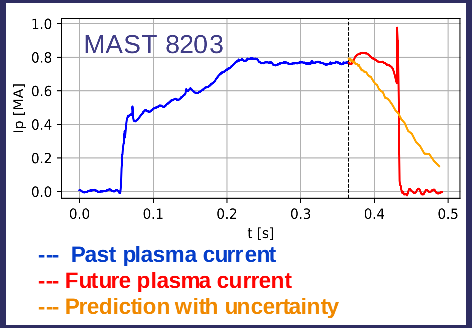

- Unsupervised approach eliminates need for manual labeling
- Stack of VAEs for pre-training
- Uncertainty-aware prediction

SOL Turbulence

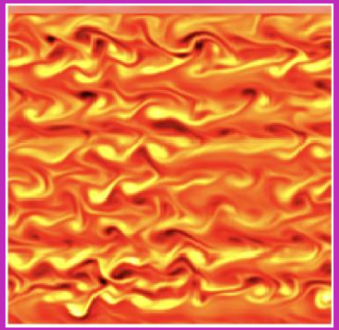

- StyleGAN as deconvolution operator for large eddy simulations
- Integrated into BOUT++