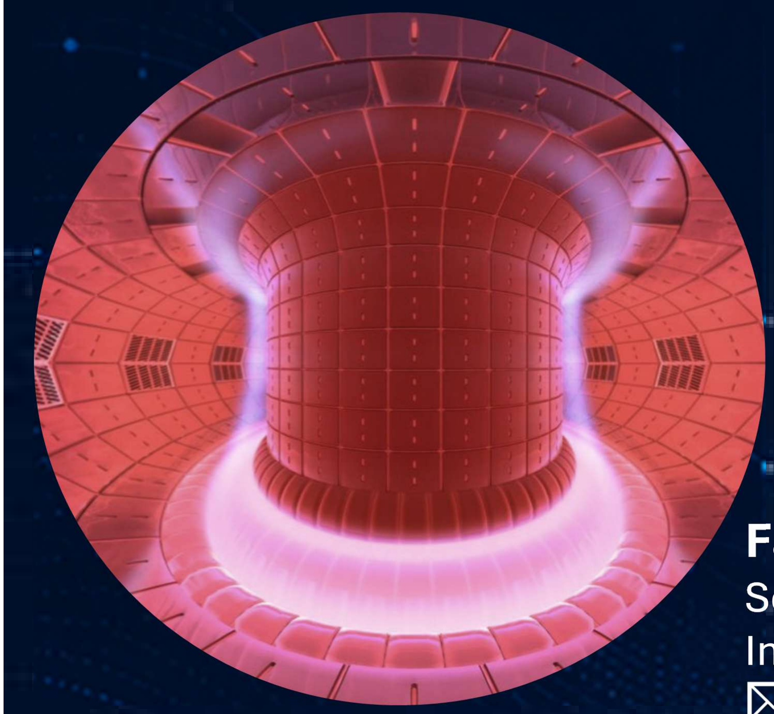

Understanding turbulence-induced transport in fusion plasmas using neuro-symbolic AI

Toward physics-constrained discovery and modelling of plasma dynamics from data

Farbod Faraji, PhD
Schmidt AI in Science Research Fellow
Imperial College London
f.faraji20@imperial.ac.uk

European Physical Society Plasma Physics Conference 2026
Edinburgh International Conference Centre, 29 June - 3 July

Turbulence-induced transport: a key bottleneck in fusion plasmas

Plasma turbulence and the resulting transport of particles and energy remain major long-standing problems in fusion science.

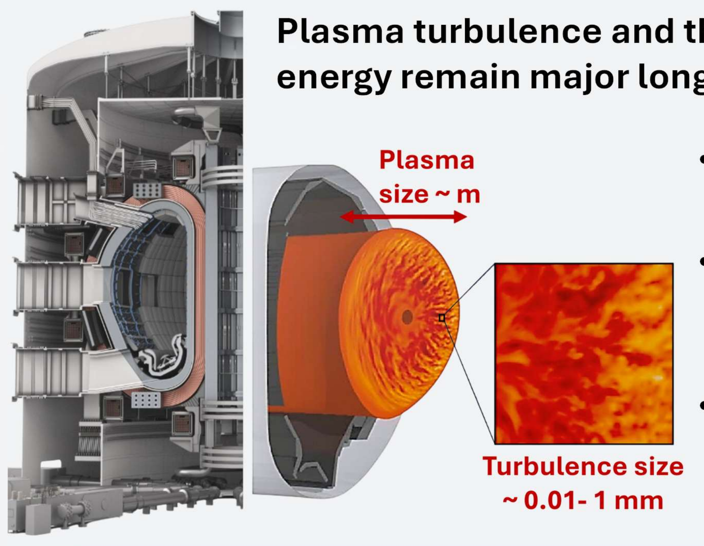

- Plasma turbulence evolves across wide spatiotemporal scales.
- Turbulence and transport influence confinement, pedestal dynamics, transient events, and divertor heat loads.
- Predicting and controlling these processes remains limited by the lack of high-fidelity, computationally efficient models.

Key open questions

- How does turbulence evolve in space and time?
- How do coherent structures emerge?
- How do these structures interact with plasma species to drive transport?
- How does transport change across plasma conditions and operating regimes?

Why today's first-principles models remain insufficient

No existing modelling framework fully satisfies the modelling needs of fusion plasmas at realistic device scales.

Kinetic models

- Capture wave-particle interactions, non-Maxwellian features, and non-local effects that drive transport.
- **But** become prohibitively expensive in realistic geometry, dimensionality, and operational time horizons.

Fluid / hybrid models

- More computationally efficient and tractable for larger-scale studies.
- **But** often miss key kinetic effects and rely on closure relations that are regime-specific or incomplete.

Why current AI/ML is not enough for discovery

AI/ML can derive fast surrogate models and complement sparse diagnostics. Yet current approaches remain limited for scientific discovery.

Data-driven and neural models

- Methods such as OPT-DMD [1,2], Neural ODEs [3], and multiscale architectures [4] can extend predictive capability.
- But they often rely on assumptions about system behaviour, offer limited prediction horizons, or remain poor in generalizability.
- As networks become deep and highly parametrized, potential gains in performance can come at the cost of transparency and explainability.

Physics-informed models (e.g., PINNs [5])

- Embed governing equations or residuals into training.
- But often rely on soft constraints and become limited when the full system of equations is unknown.

Symbolic regression (e.g., SINDy [6], Evolutionary [7])

- Infers interpretable governing relations from data.
- But often depends on predefined term libraries or expert-crafted constraints, which may bias discovery.

Neuro-symbolic AI

Neuro-symbolic AI [8] offers a promising route to combine the strengths of neural learning and symbolic reasoning.

- **Neural learning**
  extracts compact latent dynamics from high-dimensional plasma data.
- **Symbolic discovery**
  expresses those dynamics through interpretable mathematical expressions.
- **Logic-based reasoning**
  guides discovery under physical constraints.

Aim: Not only **better prediction**, but **interpretable, verifiably physics-grounded, and more generalizable** models of plasma dynamics and turbulence-induced transport.

Why neuro-symbolic AI can overcome current ML limitations

- **Generalizability**
  physics-consistent reasoning can guide cross-regime learning.
- **Data efficiency**
  models respect known physics, reducing data demands.
- **Autonomy**
  discovery does not rely solely on existing expert knowledge or biased priors.
- **Transparency**
  symbolic expressions offer interpretability into learned dynamics.

Physics-aware, reasoning-enabled dynamics discovery architecture

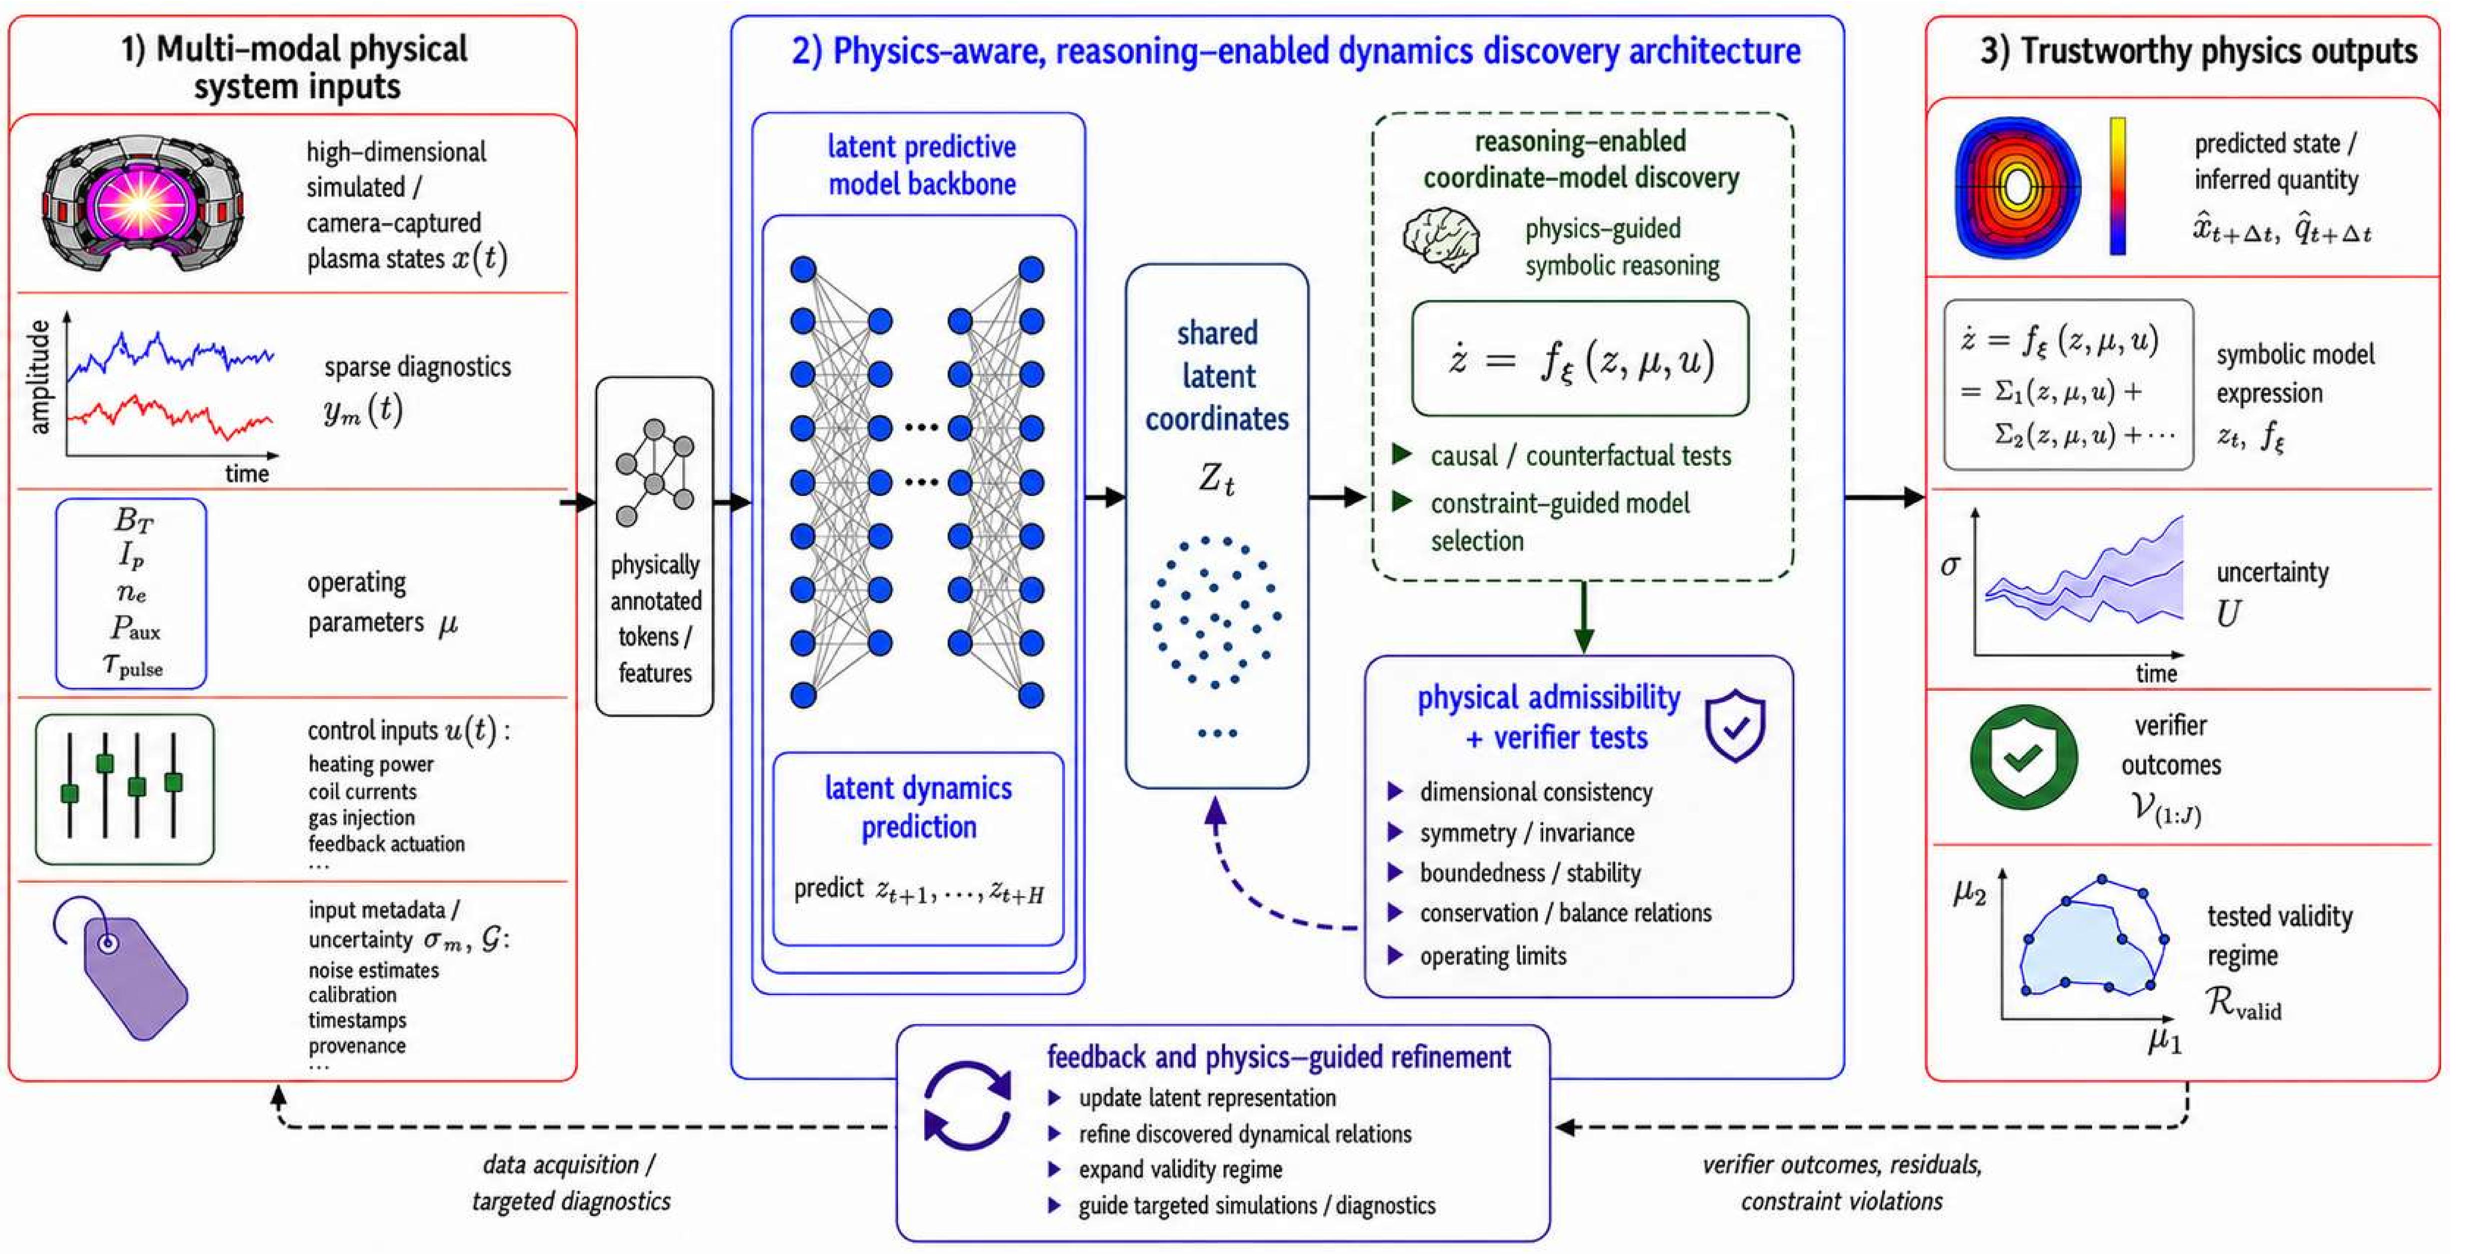

The proposed architecture connects multi-modal plasma data, latent predictive learning, symbolic dynamics discovery, logic-driven physical constraints, and validity-aware outputs.

From learned dynamics to transport insight

- **Latent coordinates** identify compact patterns in turbulent plasma evolution.
- **Symbolic discovery** proposes candidate reduced dynamics for those patterns.
- **Logic-driven** physical constraints reject inadmissible expressions.
- **Parameter and control variations** test whether candidate mechanisms remain consistent across regimes.
- Accepted models can expose **interpretable links** between turbulence, coherent structures, and transport.

Key challenges for fusion-relevant neuro-symbolic discovery

1. **Representing physics knowledge**
   - How should physical knowledge be represented in machine-readable form?
2. **Encoding real physical constraints**
   - Physics constraints are often global, relational, or transformation-based.
   - How should conservation, symmetry, invariance, dimensional consistency, and stability constraints be formalized?
   - How do we test and guarantee whether a discovered model truly satisfies the encoded constraints?
3. **Learning causally and physically meaningful coordinates**
   - How can latent variables correspond to physically interpretable and causally related turbulent structures or transport-relevant mechanisms?
4. **Benchmarking performance**
   - How should discovered models be cross-validated across withheld regimes, sparse diagnostics, noise, and known transport quantities?

Near-term demonstration pathway

Canonical plasma test cases
Verify discovery against known theory and instability / transport behaviours.

Parametric simulation families
Test whether discovered coordinate-dynamics models track systematic changes in plasma evolution and transport as prescribed physical inputs are varied.

Fusion-relevant simulated data
Demonstrate interpretable forecasts and expose interaction pathways between turbulent evolution, coherent structures, and transport.

Acknowledgment:
This research has been generously supported by Schmidt Sciences via an Eric and Wendy Schmidt AI in Science Research Fellowship.

Schmidt Sciences

References:

For the tokamak rendering and turbulence simulation

- Image credit @turbosquid.
- A. E. White et al., *Front. Nucl. Eng.* **3**:1380108 (2024).

[1] T. Askham, J. N. Kutz, *SIAM J. App. Dyn. Sys.* **17**:1, 380-416 (2018).
[2] F. Faraji et al., *J. Phys. D: Appl. Phys.* **57** 065202 (2024).
[3] R.T. Chen, *Advances in neural information processing systems*, 31 (2018).
[4] F. Faraji, M. Reza, *Phys. Scr.* **100** 076015 (2025).
[5] M. Raissi et al., *J. Comp. Phys.*, Vol. **378**, pp. 686-707 (2019).
[6] S. L. Brunton et al., PNAS **113**(15) 3932-3937 (2016).
[7] M. Schmidt, H. Lipson, *Science* **324**, 81-85 (2009).
[8] A. d'Avila Garcez, L.C. Lamb, *Artif. Intell. Rev.* **56** 12387-12406 (2023).

IMPERIAL
Department of Computing