Reduced-order Data-driven Surrogate Modeling of MAST Discharges via Selective Dense State-Space Models

Sergio A. Angelini[1], Alex Skillen[1], Małgorzata J. Zimoń[2], Mingfei Sun[1], Wei Pan[1]

[1] *The University of Manchester, Manchester, UK*

[2] *IBM Research Europe, Daresbury, UK*

ABSTRACT

We present a reduced-order, data-driven surrogate model for MAST plasma evolution. The model predicts the compact performance state

$$x_t = [I_p(t), W_n(t)],$$

where I_p is the plasma current and W_n is the neutron emission rate, used here as a global reactivity proxy. The dynamics are conditioned on experimentally available actuator signals: PF coil voltages, gas injection and neutral beam injection power, with line-averaged density treated as an exogenous plasma-state input. To focus on transport and confinement behaviour, disruptive and mitigation-dominated phases are excluded. We train a Selective Dense State-Space Model (SD–SSM), whose input-dependent selector interpolates between a dictionary of locally linear transition operators. The resulting model reproduces multi-step I_p and W_n trajectories and provides interpretable internal structure: selector weights segment discharges into operational regimes, while the spectra of the learned operators yield characteristic data-driven timescales. This provides a compact and interpretable route towards scenario evaluation and control-oriented trajectory modelling in spherical tokamaks.

INTRODUCTION

Trajectory optimisation and scenario planning will remain essential even after high-gain fusion operation is achieved, since reactor-relevant plasmas must be steered reliably through actuator-constrained operating spaces. First-principles simulations are indispensable, but remain too expensive and incomplete for rapid trajectory evaluation. Reduced-order, data-driven surrogates offer a complementary route: they can exploit the growing archive of experimental data while remaining fast enough for control-oriented studies.

This work uses curated MAST campaign data to learn a compact dynamical model of plasma evolution. Rather than reconstructing the full plasma state, we target the global variables I_p and W_n, and condition their evolution on available controls. The aim is not to model disruptions, but to learn stable ramp-up and quasi-stationary dynamics under realistic actuator histories.

DATA AND LEARNING FORMULATION

The dataset consists of MAST M8 and M9 discharges. Shots involving pellet injection are removed, while discharges with disruption mitigation or disruption onset are truncated 30 ms before the event. Additional thresholding on time derivatives of I_p and W_n removes pathological trajectories, leaving 794 viable runs. Overlapping windows of length L are then extracted and all channels are standardised using training-set statistics.

The model input is

$$u_t = [I_p, W_n, V_s, V_1, V_2, V_3, G_{\text{in}}, G_{\text{out}}, P_{\text{NBI}}, n_e], \tag{1}$$

where only I_p and W_n are treated autoregressively. Coil voltages, gas injection, NBI power and line-averaged density are provided as measured context variables.

SELECTIVE DENSE STATE-SPACE MODEL

Standard state-space models evolve a latent state through a fixed linear operator,

$$x_t = Ax_{t-1} + Bu_t, \tag{2}$$

$$y_t = Cx_t + Du_t. \tag{3}$$

This is efficient, but too restrictive when plasma dynamics change across operating regimes. Selective SSMs replace the fixed transition with an input-dependent one,

$$x_t = A(u_t)x_{t-1} + b(u_t)u_{t-1}, \tag{4}$$

$$y_t = c(x_t) + d(u_t). \tag{5}$$

In the dense variant used here, $A(u_t)$ is represented through a learned dictionary of dense transition operators. A selector network assigns time-dependent weights to these operators, allowing the effective dynamics to adapt between ramp-up, heating and quasi-stationary phases. This provides both expressive sequence modelling and a direct object for physical interpretation.

TRAINING AND VALIDATION

Training uses an $80/20$ split of the windowed dataset. To obtain stable long-horizon rollouts, we combine teacher forcing with autoregressive training. The model is first conditioned on a context window of length H and is then rolled forward for K steps, feeding its own predicted state back into the next prediction while retaining the measured actuator schedule.

The model outputs a Gaussian predictive distribution for each target channel. With predicted mean $\hat{\mu}$ and variance σ^2, the empirical negative log-likelihood is

$$\mathscr{L} = \mathbb{E}\left[\frac{(\hat{\mu} - y)^2}{\sigma^2} + \log(\sigma^2)\right]. \tag{6}$$

This penalises inaccurate predictions while learning calibrated uncertainty. A small amount of Gaussian noise is added to recycled states during training, improving robustness to rollout error accumulation.

We ablate autoregressive horizons $K \in \{20, 40, 60, 80, 100\}$ and context lengths $H \in \{10, 20, 30\}$. Performance is evaluated on fully autoregressive test rollouts using scaled RMSE over time for both I_p and W_n.

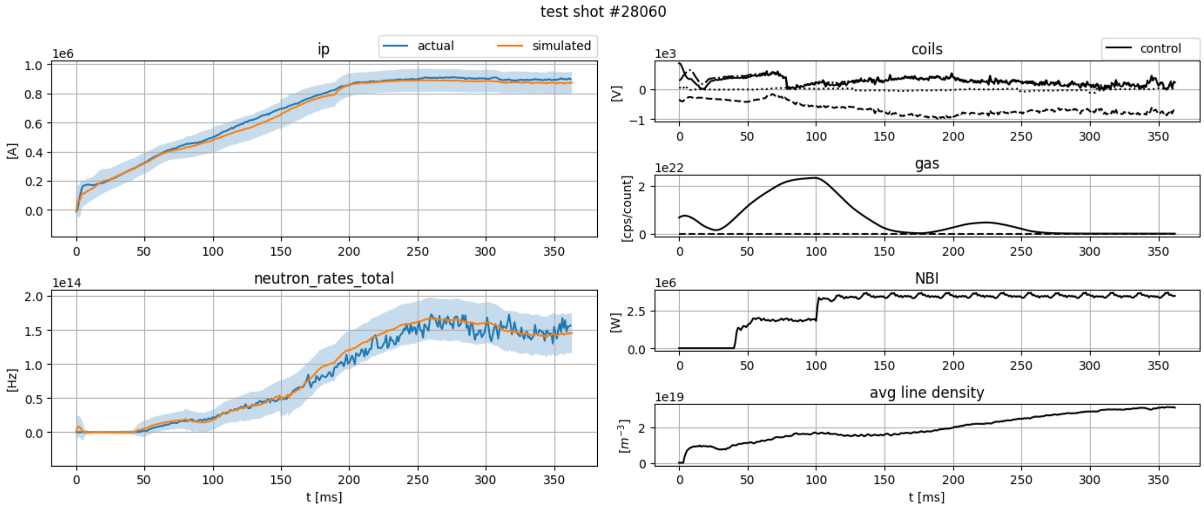

Figure 1: Example autoregressive rollout. Left: generated I_p and W_n trajectories with 90% confidence intervals. Right: measured context and control signals used to condition the rollout.

RESULTS AND OUTLOOK

Short autoregressive horizons lead to unstable long-term prediction: the $K = 20$ model rapidly accumulates rollout error because it is never exposed to the relevant plasma evolution timescales during training. Models trained with longer horizons are substantially more stable, with the best overall performance obtained for $K = 80$ and $H = 30$. In the retrained best model, the scaled RMSE converges to approximately 0.20 of the natural channel variability, and late-time I_p predictions remain stable rather than diverging.

The selector trajectories provide an interpretable segmentation of the discharge, separating transient ramp-up, heating and quasi-stationary phases. This demonstrates that the SD–SSM is not only a predictive surrogate, but also a compact system-identification tool. Future work will extend the interpretability analysis by ablating the context predictors, trying to further characterise the internal representation of the model.

References

[1] A. J. H. Donné, Phil. Trans. R. Soc. A **377**, 20170432 (2019).

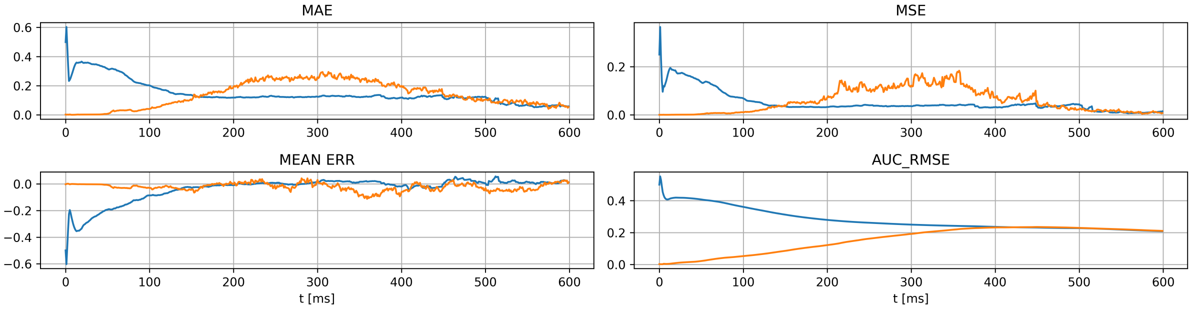

Figure 2: Performance of the retrained best model with $K = 80$ and $H = 30$. Early uncertainty is dominated by transient current dynamics, while later rollout error remains bounded.

[2] I. T. Chapman et al., Nucl. Fusion **55**, 104008 (2015).

[3] A. Gu, K. Goel and C. Ré, ICLR (2022).

[4] A. Orvieto et al., ICML (2023).

[5] A. Gu and T. Dao, arXiv:2312.00752 (2024).

[6] B. Terzić et al., arXiv preprint (2025).

[7] A. J. Thornton et al., J. Nucl. Mater. **415**, S836–S840 (2011).