Machine learning methods for modelling local, linear gyrokinetic simulations of MAST-U pedestal turbulence

Anna Niemelä[1], Daniel Jordan[1], Aaro Järvinen[1], Amanda Bruncrona[1], Adam Kit[1], Lorenzo Frassinetti[2], David Hatch[3], Samuli Saarelma[4], EUROfusion Tokamak Exploitation team[*], and the MAST Upgrade Team[**]

[1]*VTT Technical Research Centre of Finland, Espoo, Finland*
[2]*KTH Royal Institute of Technology, Stockholm, Sweden*
[3]*Institute for Fusion Studies, University of Texas at Austin, Austin, Texas, USA*
[4]*UK Atomic Energy Authority, Abingdon, UK*
[*]*See the author list of E. Joffrin et al 2024 Nucl. Fusion **64** 112019*
[**]*See the author list of J. Harrison et al 2019 Nucl. Fusion 59 112011*

In H-mode tokamak plasmas, a steep edge pressure gradient forms the pedestal, which strongly affects the overall confinement and fusion performance [1]. Pedestal structure is limited by both MHD instabilities and gyrokinetic (GK) microinstabilities, such as ion and electron temperature gradient modes (ITG, ETG), micro-tearing modes (MTM), and kinetic ballooning modes (KBM), which drive transport at ion and electron scales [1-3]. While high-fidelity GK codes, such as GENE, can capture these effects, their computational cost limits their routine use in integrated pedestal models [4,5]. As a result, most pedestal models, such as EPED or Europed, rely on reduced physics assumptions to bypass the need for detailed transport analysis [4-6].

This work investigates the use of machine learning (ML) to develop surrogate models for GK features of spherical tokamak pedestals. We present a proof-of-principle surrogate trained on local linear GENE simulations in the parameter space surrounding MAST-U discharge #49108. The aim is to enable fast, accurate predictions of GK stability for future integration into pedestal modeling workflows. This requires a dataset that spans a high-dimensional, physically realistic parameter space. In this work, the approach is to sample plasma profiles, which are then used to construct the equilibrium by solving the axisymmetric 2D Grad-Shafranov equation. Each sampled case is defined by prescribing the plasma shape, toroidal current, and magnetic field, all matched to the representative MAST-U discharge. The separatrix temperature is fixed at 50 eV, consistent with MAST-U edge conditions.

Rather than using experimental profiles directly, we use them to define sampling bounds, enabling the generation of physically plausible, generalizable inputs (Fig. 1). This approach supports scalable data generation and allows the surrogate model to be applied beyond existing experimental regimes. The sampled parameters include pedestal top electron temperature and density, separatrix density, and the pedestal widths of both temperature and density profiles. Bounds for these parameters are derived from experimental data from MAST-U discharge #49108, using Thomson scattering measurements. The data is mapped onto normalized poloidal flux coordinates, and only time slices within 75% to 95% of the inter-ELM period are used, corresponding to intervals when the pedestal is nearly saturated and thus more representative of its quasi-steady-state structure [1].

For data efficiency, static sparse grids are used for populating the training database [7]. For each combination of sampled values, a modified hyperbolic tangent (mtanh) parameterization is used to construct smooth, continuous profiles to represent experimental profile shapes [5]. Each sampled profile pair, together with the fixed plasma shape and magnetic configuration from MAST-U #49108, is passed to the HELENA Grad-Shafranov solver, which computes the equilibrium including the self-consistent bootstrap current that dominates the pedestal region [8]. For each case, HELENA is run iteratively to adjust the temperature profile until a target core plasma beta is achieved. 126 self-consistent equilibria were generated with this procedure (Fig. 1).

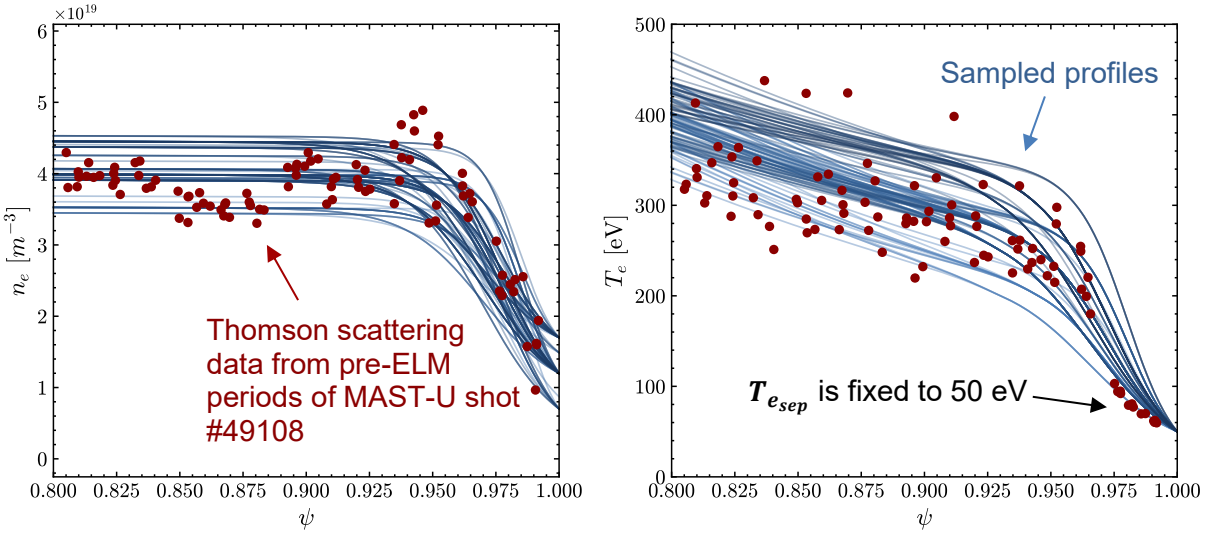

Figure 1. Sampled temperature and density profiles. Thomson scattering measurements of electron density and temperature from MAST-U discharge #49108, taken during 75–95% of the inter-ELM period, are overlaid on the samples.

To prepare the equilibrium data from HELENA for gyrokinetic analysis, the output is post-processed using CHEASE, which is used primarily to smoothly map the magnetic flux surfaces to a R–Z grid and to generate an EQDSK file, which provides magnetic geometry input for GENE [9]. Alongside this, an ITERDB file is constructed to specify the plasma profiles on a ρ_{tor} grid. For the ITERDB file, rotation values are estimated using the radial force balance [10].

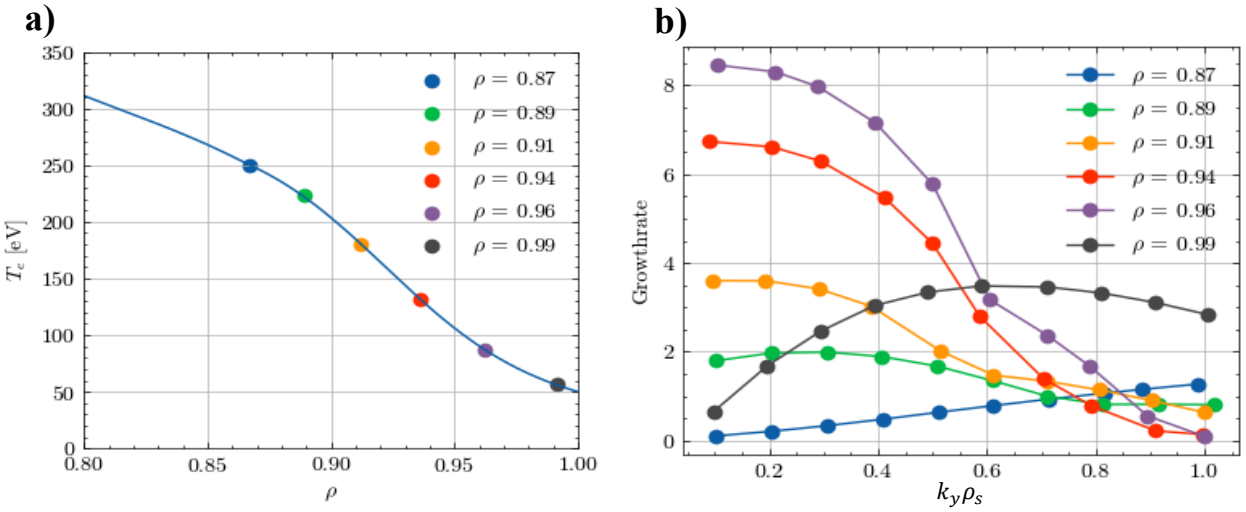

Figure 2. (a) Temperature profile of a sample. Six radial positions from pedestal top to 95% of the pedestal region are chosen for gyrokinetic simulations. (b) Growthrate over binormal wavenumber as outputted by GENE, for each radial position.

GENE is run using the tracer-EFIT geometry [4]. To capture the radial variation of stability across the pedestal, six radial positions are selected for each equilibrium. These positions are uniformly spaced from the pedestal top of the broader profile (either temperature or density) to a location corresponding to 95% of the pedestal width (Fig. 2a).

At each of the six radial locations, local linear GENE simulations are done over a uniform grid of 10 normalized binormal wavenumbers ($k_y\rho_s$) at the ion-scale (Fig. 2b). Each GENE run outputs linear instability features, such as the mode frequency, growth rate, and estimates of ion/electron particle and heat fluxes. While linear simulations do not capture saturation mechanisms, and thus cannot predict absolute flux levels, prior work has shown that ratios of the diffusivities, calculated from the fluxes, can be used to identify the dominant underlying microinstability [11].

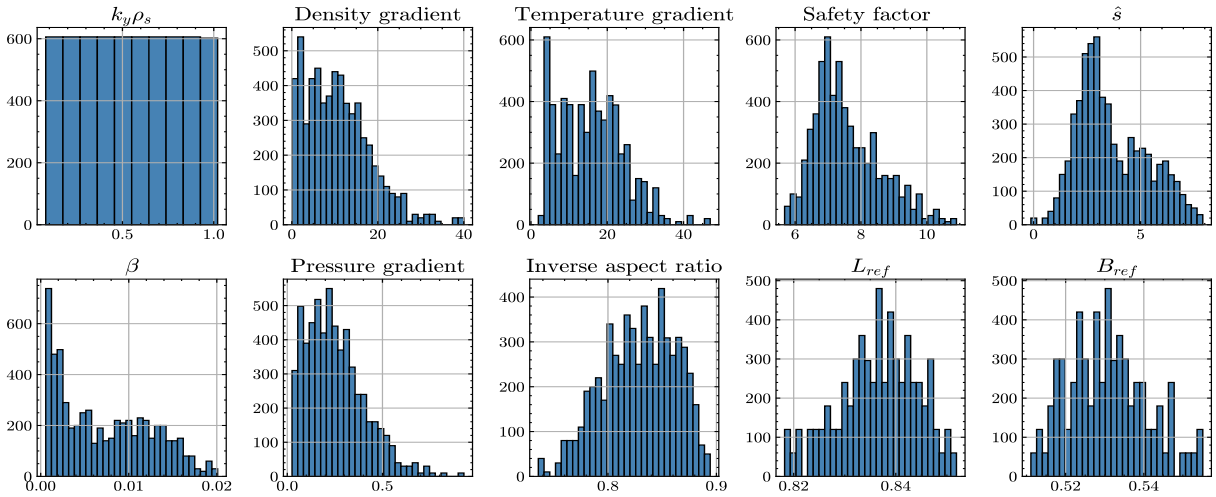

Figure 3. Distribution of GENE input parameters of the training dataset.

The full dataset generated by GENE comprises approximately 7,500 local linear simulations, corresponding to six radial locations and ten binormal wavenumbers per equilibrium for 126 unique pedestal profiles. To train and evaluate a machine learning model, the dataset is split into a training set (80%) and a test set (20%), with 101 profiles (~6,000 simulations) used for training and 25 profiles (~1,500 simulations) held out for testing. The distribution of the training inputs used in the ML model, are shown in Figure 3. The target outputs are the linear growth rate, the ion-to-electron heat diffusivity ratio $\frac{\chi_i}{\chi_e}$, and the electron particle to heat diffusivity ratio $\frac{D_e}{\chi_e}$. The surrogate model is a fully connected neural network with 10 input features, hidden layers with 64, 32, and 16 neurons, and a 3-neuron output layer corresponding to the targets. The model is trained with the Adam optimizer and mean squared error loss.

Initial results show good predictive performance for the linear growth rate, with low test error and strong correlation between predictions and GENE outputs (Fig. 4). However, the model currently exhibits higher errors and reduced reliability for the diffusivity ratios. Present hypotheses suggest that the mode transitions complicate the regression landscape. These diffusivity targets remain an area for further refinement.

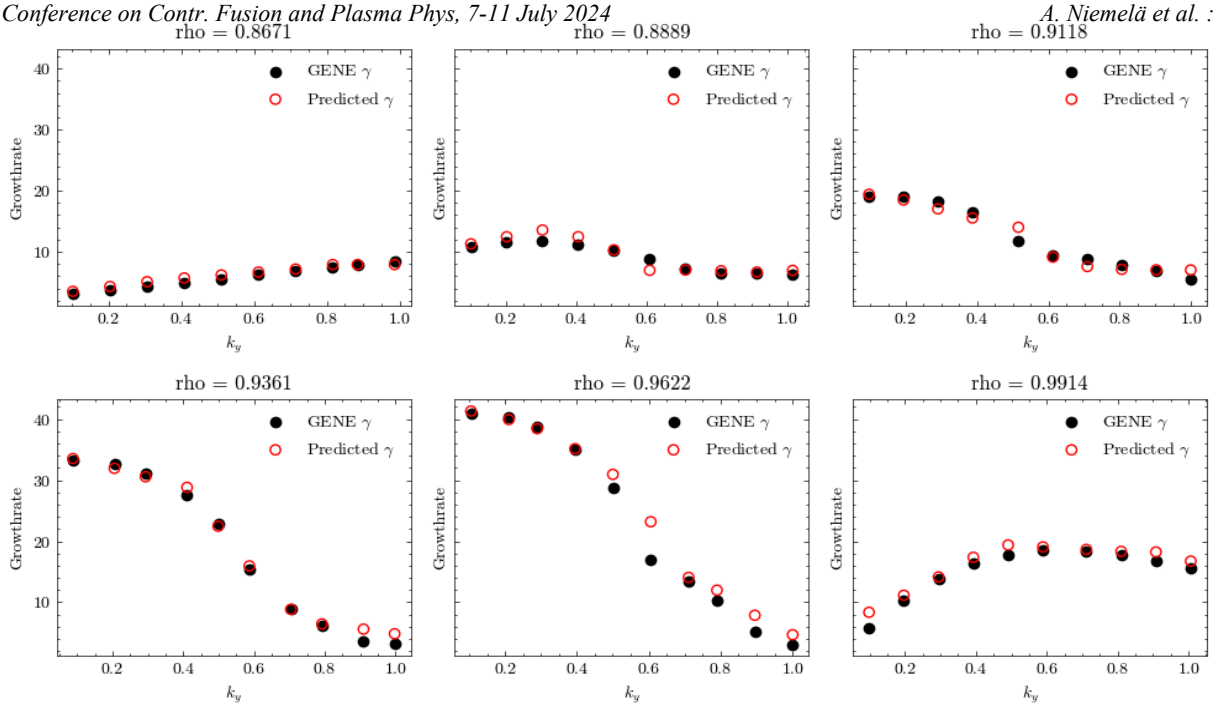

Figure 4. Comparison between growthrate outputs of GENE and the fully connected neural network for each radial location. Sample is in the test set, and therefore not used in the training of the model.

Future work will focus on exploring alternative models, such as convolutional or ensemble-based approaches. Some previously listed assumptions will be relaxed, to enable sampling in a wider operational space of MAST-U and JET, and for expanding the dataset. Data generation will also be expanded to electron scales.

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 —EUROfusion). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. The work of Aaro Järvinen, Adam Kit, and Amanda Bruncrona has been partially supported by the Research Council of Finland grant no. 355460, and the work of Daniel Jordan and Anna Niemelä by the Research Council of Finland grant no. 358941. The authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

References

[1] P.-Y. Li, et al. Nucl. Fusion 64 (2024) 016040.
[2] D Dickinson, et al. Plasma Phys. Control. Fusion 55 (2013) 074006.
[3] J.F. Parisi, et al. Nucl. Fusion 64 (2024) 054002.
[4] F. Jenko, et al. Phys. Plasmas 7 (2000) 1904-1910.
[5] P.B. Snyder, et al. Phys. Plasmas 49 (2009) 085035.
[6] S. Saarelma, et al. Phys. Plasmas 26 (2019) 072501.
[7] D. H. Jordan, et al. This conference.
[8] G.T.A. Huysmans, et al. Proc. Int. Conf. Comp. Phys. (1991) 371.
[9] H. Luetjens, et al. Comp. Phys. Commun. 97 (1996) 219.
[10] M. Landreman and D.R. Ernst, Plasma Phys. Control. Fusion 54 (2012) 115006.
[11] M. Kotchenreuter, et al. Nucl. Fusion 59 (2019) 096001