

CONFERENCE PRE-PRINT

DEVELOPING MACHINE LEARNING FACILITATED PEDESTAL MODELS

A.E. JÄRVINEN, A.M. BRUNCRONA, D. JORDAN, A. KIT, A. NIEMELÄ
VTT Technical Research Centre of Finland
Espoo, Finland
Email: aaro.jarvinen@vtt.fi

M. DUNNE
Max-Planck-Institut für Plasmaphysik
Garching, Germany

L. FRASSINETTI, H. NYSTRÖM
Division of Fusion Plasma Physics, KTH Royal Institute of Technology
Stockholm, Sweden

D.R. HATCH, J. SCHMIDT, C. STEPHENS
Institute for Fusion Studies, University of Texas at Austin
Austin, Texas, USA

L. LEPPIN
Oden Institute, University of Texas at Austin
Austin, Texas, USA

V. MENKOVSKI
Eindhoven University of Technology, Mathematics and Computer Science
Eindhoven, the Netherlands

Y.R.J. POELS
Swiss Plasma Center, EPFL
Lausanne, Switzerland

S. SAARELMA, L. ZANISI
UK Atomic Energy Authority
Abingdon, UK

S. WIESEN
Dutch Institute for Fundamental Energy Research – DIFFER
Eindhoven, the Netherlands

THE ASDEX UPGRADE TEAM*

See the author list of ‘Overview of ASDEX upgrade results in view of ITER and DEMO’ by H Zohm et al. Nuclear Fusion 64 (2024) 112001.

JET CONTRIBUTORS

See the author list of ‘Overview of T and D-T results in JET with ITER-like wall’ by CF Maggi et al. Nuclear 64 (2024) 112012.

Abstract

This conference manuscript provides an overview of recent activities in a project developing machine learning (ML) facilitated pedestal models. The project is divided to three branches, consisting of surrogate modelling techniques for pedestal magnetohydrodynamics, development of reduced pedestal transport models with ML methods, as well as data-driven methods to learn corrections for the remaining gap between numerical predictions and experimental observations. A proof-of-principle model for accelerating pedestal MHD stability evaluations has been recently published, and the next step activities to go beyond this proof-of-principle are detailed. First proof-of-principle models are emerging from the part of the project developing surrogate models for local, linear pedestal gyrokinetic evaluations based on GENE simulations for JET and MAST-U. The data-driven models are proceeding from purely observations-based models to models that combine both physics models and experimental observations for a combined representation.

1. INTRODUCTION

Pedestal optimization is central to establishing sustained high performance in conventional tokamak scenarios [1]. The key point of tension is in the typically conflicting demands of pedestal performance and the necessary power exhaust measures to avoid overheating of the divertor components. A core-edge integrated pedestal, connected to an exhaust solution and a self-consistent core, is a multiscale, multiphysics system that is scientifically and computationally challenging to simulate with integrated models of various fidelities (Fig. 1).




Fig. 1. Cartoon illustration of the mutually linked core, pedestal, and SOL plasmas.

To alleviate the scientific and computational complexity of pedestal plasmas, the scientific community has developed reduced models for projecting performance between scenarios and devices [2 – 5]. The common structure of these models is such that a set of pedestal profiles is generated using reduced transport models with assumed edge transport barrier (ETB) widths, and that the overall magnetohydrodynamic (MHD) stability envelope is computed with a linear MHD stability solver, such as MISHKA or ELITE [6, 7] (Fig. 2). The predicted pedestal profile is that for which the transport and MHD constraints meet. Meanwhile, to push the envelope of predictive capability of these models, the research community is actively investigating the key physics areas contributing to pedestal performance, including transport, the role of resistive MHD, neutral fuelling, and separatrix boundary conditions [8 – 13].




Fig. 2. Cartoon illustration of the standard pedestal prediction workflow

The development of data science methods has opened a pathway to bridge the gap between model complexity and computational throughput [14]. This conference contribution presents an overview of an active project focused on

developing machine learning (ML) facilitated pedestal models for tokamaks. The project consists of three parallel branches:

- ML surrogate models for pedestal MHD stability evaluations;
- Reduced pedestal transport models based on surrogate models for linear gyrokinetic (GK) simulations;
- Methods to learn reduced dimensionality representations for the remaining gap between simulations and observed reality, *sim2real*, based on large databases of experimental observations.

The first branch is focused on accelerating linear MHD stability evaluations for pedestals. As can be seen in Figure 2, assuming a reduced transport model that can generate a set of pedestal profiles for assumed ETB widths without a significant computational demand, the computational bottleneck is in the linear MHD stability evaluations that need to be conducted for these profiles. Through the first branch of this project, the aim is to develop ML accelerated surrogates for ideal and resistive pedestal MHD stability evaluations to accelerate this part of the workflow.

The second branch is investigating the transport aspects of the pedestal, aiming to build higher fidelity reduced models to augment the approach beyond the standard ballooning critical pedestal assumption or models based on experimental observations [2, 3]. This approach is focused first on developing an ML surrogate for local, linear pedestal GK with the training data generated using the GENE code [15]. The plan is to then use the trained surrogate to develop a quasilinear transport model for the pedestal region.

The third branch is focused on developing methods to learn features that are not represented by the surrogated numerical models by leveraging the information encompassed by large database of experimental observations, such as the JET pedestal database (JETPDB) [16]. The work in this branch is focused on generative artificial intelligence (GenAI) algorithms [17, and references therein]. These algorithms aim to learn a joint distribution between the relevant input, latent, and observed parameters of the system. Having a representation for the joint distribution enables conditional prediction of the unknown features of a system, given the known information from the system. In the context of this work, the vision is for this approach to provide an elegant method to fill the *sim2real* gap with experimental observations.

2. ENCHANTED-SURROGATES

To facilitate large-scale simulation data generation for surrogate model development, a unified framework, *Enchanted-surrogates*, has been developed in collaboration with other projects (Fig. 3) [18, 19]. The key challenges related to large-scale simulation execution, sampling strategies and database handling are common to any task requiring more than several thousand executions of a computationally demanding forward model. The aim of *Enchanted-surrogates* is to be a common framework that can address many of these challenges and the software is available Open Source at [18]. A modular class structure is implemented, such that common *Executors* and *Samplers* classes are model agnostic and intended to be applicable to any forward model, and that the model specific features are implemented in the *Parsers* and *Runners* classes. To avoid escalation of the dependencies within the main software package, the code specific implementations in *Parsers* and *Runners* are presently implemented as plugins that can be installed as needed.

Executors: A naïve implementation of a sampling strategy in a surrogate model development task is likely to lead to a workflow generating several thousand of independent simulation tasks. However, the usual Slurm Workload Manager system, as is common to many high-performance computing (HPC) environments, is not intended for handling several thousand independent job submissions from a single user. The main purpose of the *Executors* is to provide an interface layer that can internally conduct the simulation task orchestration such that the underlying computer system is not faced with thousands of small tasks. Instead, the computer system receives significantly larger, but fewer computing tasks that consist of many simulation tasks internally scheduled by the *Executors*. The primary implementation of *Executors* uses Dask as the interface to the underlying HPC system [20].

Samplers: A common need for simulation database generation is also the selection of sampling strategy, such as Grid or Random Sampling. Since these do not depend on the actual forward model that is sampled, they are implemented in a generic class that can be applied for any of the models implemented through the *Parsers* and *Runners* classes. The present active development is focused on smart sampling schemes, such as active learning (see Zanisi et al. [19]) and Bayesian Optimization [21].

Parsers: A parser for a numerical model implements the code input and output file parsing features. As Executor and Sampler generate a sample and a directory for the simulation, the parser has features to write the input file that is necessary for the simulation to run. This is the part of the framework where most simulation code specific implementation is needed as the generated input files must provide everything that is needed to execute the simulation. The parser has also features to read the output files as well as collect the sample information that is needed, e.g., for the Bayesian Optimization to yield a numerical value for the optimized quantity.

Runners: Runner class implements features to execute the run for which the input files were prepared. This class relies on the compiled executable being available in the system where the framework is executed.




Fig. 3. Overview of the Enchanted-surrogates software package

3. MACHINE LEARNING ACCELERATED PEDESTAL MHD STABILITY EVALUATIONS

A typical feature of integrated pedestal performance prediction workflows, such as EPED, Europed, IPED, and IMEP, is repeated evaluations of MHD stability for a family of pedestal profiles constrained through a transport assumption [2 – 5]. When the transport assumption is encompassed by a reduced model, the repeated MHD stability simulations with models, such as HELENA and MISHKA, become the primary time-consuming part of these workflows [6, 22]. While previous research has developed surrogate models for EPED and Europed [23, 24], the key in this project is to focus on surrogating only the MHD stability part without embedding the transport assumptions. Such a surrogate model could potentially accelerate the MHD stability evaluation part in any of the integrated modelling workflows.

The proof-of-principle surrogate model for MISHKA, called KARHU, was recently published by Brunconra et al. [25]. The data generation workflow was such that Grad-Shafranov equilibria with bootstrap current included were computed with HELENA with parameterized density and temperature profiles and then MISHKA stability simulations were computed for those equilibria for toroidal mode numbers 3, 5, 7, 10, 15, 20, 30, and 50. Using Enchanted-surrogates, a database of 16000 HELENA equilibria and corresponding MISHKA stability evaluations were generated. This corresponds to more than 100 000 MISHKA simulations. The operational space corresponds to a subset of the EUROfusion pedestal database for JET [16]. This proof-of-principle model was established using modified hyperbolic tangent parameterization of the pedestal plasma profiles and applying several simplifications to reduce the dimensionality of the problem. These include, parameterized plasma shape, assuming ballooning critical pedestal profiles with the usual EPED constant of 0.076, equal ion and electron temperatures, the pedestal widths to be the same for densities and temperatures, no relative shift between the density and temperature pedestals, separatrix temperature of 100 eV, separatrix density of 25% of the pedestal top density, as well as no variation of the effective charge state of the plasma. A convolutional neural network (CNN) model was trained on this database to predict the growth rate of the most unstable mode. The mean absolute percentage error of the surrogate model was below 1% on the test set. This version of KARHU and the training database are publicly available in GitHub [26]. The present focus of the work is to expand the operational space of the surrogate model as well as relaxing many of these assumptions applied in the proof-of-principle study.

The project is also working towards extending the model with resistive features through CASTOR (Fig. 4) [27, 28, 29]. As the pedestal MHD stability is observed to be impacted by resistivity [11], this part of the project aims

to construct a multifidelity approach to surrogating pedestal MHD stability evaluations by informing the model with CASTOR. However, as CASTOR is computationally significantly more demanding than MISHKA, while a good fraction of the expected pedestal operational space is described well by ideal MHD, simply random sampling the same operational space with CASTOR would not be efficient. Instead, an active sampling scheme is being developed, where the CASTOR simulations are primarily centered around the operational space where resistivity leads to a change from the ideal MHD picture. Another challenging feature is that due to numerical reasons, the unstable eigenmodes need to be computed by using a sufficiently good initial state in the simulations. To meet this demand, the sampler that is being developed starts from sufficiently large pedestal widths to make sure that the pedestal is strongly unstable both in the ideal and resistive MHD descriptions. Then the pedestal width is gradually reduced until stable profiles are obtained in the resistive description. Hence, the sampler construction becomes somewhat more complicated than required in the original HELENA and MISHKA database generation, as the CASTOR sampler requires to conduct sequential pedestal width scan as well. A proof-of-principle implementation is expected before the end of the year 2025.




Fig. 4. An example comparison of the predicted growth rate as a function of pedestal width and toroidal mode number between MISHKA and CASTOR.

4. MACHINE LEARNING SURROGATE MODEL FOR LOCAL, LINEAR GYROKINETICS IN PEDESTAL

This branch of the project aims to develop surrogate models for GK instabilities in the pedestal based on a local, linear description with the intent to build a quasilinear transport model for pedestal plasmas [15, 16]. In the project described in this contribution, the focus is on data-efficiently generating surrogate models for local, linear GK. To achieve this objective, a GK simulation sampler is developed within Enchanted-surrogates. The data generation workflow is very similar to the data generation workflow in the MHD part of the project (Fig. 5). First equilibria are created with HELENA, including the bootstrap current, with a given plasma boundary shape, current and toroidal field, and generating pedestal profiles through mtanh-parameterization. The GENE simulation inputs are obtained from these equilibria and the associated profiles. The rotation values are estimated using the radial force balance [30]. GENE simulations are conducted using the tracer-EIFT geometry. Six radial locations are selected within the pedestal, uniformly spaced from the pedestal top to 95% of the pedestal width. Local, linear GENE simulations are conducted over a uniform grid of 10 normalized binormal wavenumbers at the scale of ions. The proof-of-principle workflow presented in this contribution is focused within the plasma parameter operational space in the vicinity of targeted JET and MAST-U plasmas [31]. To reduce the dimensionality of the space within this proof-of-principle part, data generation is focused on key pulse numbers. For JET, the Deuterium reference discharge, JET pulse number (JPN) 97781, was selected, as this was previously investigated with GENE by

Leppin, et al., providing a good base for surrogate model development [32]. For MAST-U, the selected discharge was #49108, aligning well with the overall pedestal analysis tasks within the broader EUROfusion team. For JET, the workflow is presently assuming 100 eV electron temperature at the separatrix, while for MAST-U the assumed separatrix electron temperature is 50 eV. For MAST-U, the full dataset generated by GENE consists of about 7500 local, linear simulations [31]. A fully connected neural network model was demonstrated for establishing a regression model for the linear growth rate, the diffusivity ratios between ion and electron heat and particles and electron heat. Overall promising regression performance was demonstrated, and the present focus of the work is to expand the operational space as well as to establish the corresponding model for the targeted JET plasmas.




Fig. 5. Overview of the local, linear gyrokinetic dataset generation.

5. GENERATIVE MODELS FOR PEDESTAL ANALYSIS AND PREDICTION

Even though the MHD and GK surrogates, if successful, would already be expected to provide both computational throughput and physics fidelity improvements beyond the standard EPED-like workflows, without other models for the volumetric power and particle source terms as well as SOL physics, such models would still be incomplete. Since data-driven models can also learn to directly represent distributions of experimental observations, a question emerges whether it is possible to learn to represent the gap between the known physics, encompassed by the numerical models, and observed experimental reality, represented by physics databases. To pursue this ideal, representation learning algorithms have been explored for databases of experimental observations [33 – 35]. The explored models are based on the variational autoencoder, which encodes the representation for the observed plasma state information [35, 36]. By providing conditioning information through prior- and auxiliary regression objectives, the intention is that these algorithms learn to organize the information, such that the latent representation is disentangled with respect to semantically meaningful information, such as the control parameter configuration [37]. Previous research has shown how such a model can be used for predicting the pedestal plasma state, given a control parameter configuration [33], learning machine size dependent and independent latent representations [35], as well as in applications focused on predicting dynamical evolutions of plasmas [34].

The latest focus on this part of the project has been on exploring representation learning models trained with a mixture of observations and simulation predictions (Fig. 6). In this proof-of-principle test, the model was trained with the JET EUROfusion pedestal database and the associated Europed database [5, 16]. One of the conditioning variables that the model takes is whether the observed data originates from Europed or from experiment. After training the model can generate pedestal profiles that depend on this conditioning variable, qualitatively representing the reality gap of the model. The latest activities have explored this type of a model structure in a dynamical setting with a simple 1D continuity equation for the plasmas [38].




Fig. 6. (a) Conceptual cartoon illustration of the abstract latent space manifold representing the observed experimental data (yellow-green) and the simulated data (red). (b) 2D cut from the latent space of a representation learning model with the color contour providing the associated pedestal top pressure for the given latent variable, when the model is conditioned to predict values corresponding to experimental observations. (c) Same 2D cut as in Figure (b), but the model is conditioned to predict values corresponding to European simulation results.

6. OUTLOOK

This manuscript provides an overview of a project focused on developing ML facilitated models for pedestal plasmas. Proof-of-principle implementations are emerging for linear pedestal MHD and local, linear pedestal GK surrogate models, as well as for generative models based on experimental observations. Future research activities are focused on expanding the operational space of the surrogate models, exploration of the local, linear GK surrogate model in a quasilinear transport model application, as well as investigations on how to best apply the generative models for joint distribution learning with both observed and simulated data.

ACKNOWLEDGEMENTS

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. This work was also partially supported by Research Council of Finland projects 358941 and 355460. The authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

REFERENCES

- [1] GROEBNER, R.J., SAARELMA, S., Elements of H-mode pedestal structure, Plasma Phys. Control. Fusion **65** (2023) 073001.
- [2] SNYDER, P.B., et al., A first-principled predictive model of the pedestal height and width: development, testing and ITER optimization with the EPED model, Nucl. Fusion **51** (2011) 103016.
- [3] LUDA, T., et al., Integrated modeling of ASDEX Upgrade plasmas combining core, pedestal and scrape-off layer physics, Nucl. Fusion **60** (2020) 036023.
- [4] DUNNE, M.G., et al., Global performance enhancements via pedestal optimization on ASDEX Upgrade, Plasma Phys. Control. Fusion **59** (2017) 025010.
- [5] SAARELMA, S., et al., Self-consistent pedestal prediction for JET-ILW in preparation of the DT campaign, Phys. Plasmas **26** (2019) 072501.
- [6] A. Mikhailovskii, et al, Optimization of computational MHD normal-mode analysis for tokamaks, Plasma Phys. Rep. **23** (1997) 844-857.
- [7] H.R. Wilson, et al., Numerical studies of edge localized instabilities in tokamaks, Phys. Plasmas **9** (2002) 1277-1286.
- [8] M.J. Pueschel, et al., On microinstabilities and turbulence in steep-gradient regions of fusion devices, Plasma Phys. Control. Fusion **61** (2019) 034002.

- [9] B. Chapman-Oplopoiou, et al., Composition of electron temperature gradient driven plasma turbulence in JET-ILW tokamak plasmas, *Phys. Rev. Research* **7** (2025) L012004.
- [10] D.R. Hatch, et al., Reduced models for ETG transport in the tokamak pedestal, *Phys. Plasmas* **29** (2022) 062501.
- [11] H. Nyström, et al., Effect of resistivity on the pedestal MHD stability in JET, *Nucl. Fusion* **62** (2022) 126045.
- [12] S. Saarelma, et al., Integrated modelling of H-mode pedestal and confinement in JET-ILW, *Plasma Phys. Control. Fusion* **60** (2018) 014042.
- [13] L. Horvath, et al., Pedestal particle balance studies in JET-ILW H-mode plasmas, *Plasma Phys. Control. Fusion* **65** (2023) 044003.
- [14] ANIRUDH, R., et al., 2022 Review of Data-Driven Plasma Science, *IEEE Transactions on Plasma Science* **51** (2023) 1750-1838.
- [15] JENKO, F., et al., Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** (2000) 1904-1910.
- [16] FRASSINETTI, L., et al. Pedestal structure, stability and scalings in JET-ILW: the EUROfusion JET-ILW pedestal database, *Nucl. Fusion* **61** (2021) 016001.
- [17] HE, R., CAO, J., TAN, T., Generative artificial intelligence: a historical perspective, *Nat. Sci. Rev.* **12** (2025) nwaf050.
- [18] <https://github.com/DIGIfusion/enchanted-surrogates/>
- [19] ZANISI, L., et al., Data-efficient digital twinning strategies and surrogate models of quasilinear turbulence in JET and STEP, 30th IAEA Fusion Energy Conference, Chengdu, China (2025).
- [20] <https://www.dask.org/>
- [21] JÄRVINEN, A.E., et al., Towards scalable large-scale model validation with data science, 9th Asia-Pacific Conference on Plasma Physics, Fukuoka, Japan (2025).
- [22] HUYSMANS, G.T.A., GOEDBLOED, J.P., KERNER, W., Isoparametric bicubic Hermite elements for solution of the Grad-Shafranov equation, *Proc. Int. Conf. Comp. Phys.* (1991) 371.
- [23] MENEGHINI, O., et al., Self-consistent core-pedestal transport simulations with neural network accelerated models, *Nucl. Fusion* **57** (2017) 086034.
- [24] PANERA-ALVAREZ, A., et al., EuroPED-NN: uncertainty aware surrogate model, *Plasma Phys. Control. Fusion* **66** (2024) 095012.
- [25] BRUNCRONA, A.M., et al., Machine learning surrogate model for ideal peeling-ballooning pedestal MHD stability, *Phys. Plasmas* **32** (2025) 092501.
- [26] <https://github.com/DIGIfusion/karhu>
- [27] KERNER, W., et al., CASTOR: Normal-mode Analysis of Dissipative MHD Plasmas, *J. Comp. Phys.* **142** (1998) 271.
- [28] BRUNCRONA, A.M., et al. Machine learning surrogate model for pedestal MHD stability, 51st EPS Conference on Plasma Physics, Vilnius, Lithuania (2025).
- [29] STEFANELLI, F., et al., Non-ideal MHD analysis of pedestal scans of AUG and JET, *Nucl. Fusion* **65** (2025) 086025.
- [30] LANDREMAN, M., ERNST, D.R., Local and global Fokker-Planck neoclassical calculations showing flow and bootstrap current modification in a pedestal, *Plasma Phys. Control. Fusion* **54** (2012) 115006.
- [31] NIEMELÄ, A., et al., Machine learning methods for modelling local, linear gyrokinetic simulations of MAST-U pedestal turbulence, 51st EPS Conference on Plasma Physics, Vilnius, Lithuania (2025).
- [32] LEPPIN, L.A., et al., The JET hybrid H-mode scenario from a pedestal turbulence perspective, arXiv:2405.10688 (2024).
- [33] KIT, A., et al., Developing deep learning algorithms for inferring upstream separatrix density at JET, *Nucl. Mat. Ene.* **34** (2023) 101347.
- [34] KIT, A., et al., On learning latent dynamics of the AUG plasma state, *Phys. Plasmas* **31** (2024) 032504.
- [35] JÄRVINEN, A.E., Representation learning algorithms for inferring machine independent latent features in pedestals in JET and AUG et al. *Phys. Plasmas* **31** (2024) 032508.
- [36] KINGMA, D.P., WELING, M., Auto-encoding variational Bayes, arXiv: 1312.6114 (2014).
- [37] ILSE, M., et al., DIVA: Domain invariant variational autoencoders, *Proceedings of the Third Conference on Medical Imaging with Deep Learning Proceedings of Machine Learning Research* **121** (2020) 322-348.
- [38] KIT, A, et al., On Physics-Data Generative Modeling for Core-Edge Integration in Tokamaks, 9th Asia-Pacific Conference on Plasma Physics, Fukuoka, Japan (2025).