An active learning pipeline for surrogate models of gyrokinetic turbulence.

J. Barr[1], T. Madula[1], L. Zanisi[2], A. Ho[3,4], J. Citrin[3,4], V. Gopakumar[2], and JET contributors [*]

[1] *UCL's Centre for Doctoral Training in Data Intensive Sciences, University College London;* [2] *Culham Centre for Fusion Energy, Abingdon, ;* [3] *Dutch Institute for Fundamental Energy Research, Eindhoven, Netherlands;* [4]*Science and Technology of Nuclear Fusion Group, Eindhoven University of Technology, Eindhoven, Netherlands*

Abstract

Digital twinning of a tokamak device requires fast system state inference. Physics-based computational models that predict future states are often too slow to be actionable, and thus undesirable for offline scenario planning. These tasks may be performed faster if the physics-based model is replaced by a neural network-based surrogate. Obtaining the labels to train the surrogate can be computationally expensive, additionally, some inputs may result in trivial outputs. Here we propose a two-stage active learning pipeline for digital twinning of gyrokinetic turbulence in the core of tokamak fusion plasmas. Our pipeline leverages an uncertainty-based acquisition function which greatly outperforms random acquisition and leads to a reduction of 99.6% in the amount of input-output mappings needed from the physical model without compromising on performance.

Introduction

Turbulent transport in Tokamak plasmas is a major roadblock to achieving fusion power [1]. It is thus crucial to forecast expected turbulent fluxes in the operating parameter space of a given tokamak for design, validation and optimization at the pre- and post-experiment stage. In general, a general-purpose surrogate would be expensive to obtain, as the computational effort needed to obtain the labels from all potentially needed realisations of the physical model would be prohibitively long. Instead, a surrogate can be trained directly on experimental inputs, which better capture the parameter subspace of the machine. The surrogate can then be re-evaluated and updated when new data becomes available.

Experimental data from a given machine may be redundant, and data labelling (i.e. obtaining the input-output mapping from the physical model) for similar inputs would be inefficient. Moreover, by the nature of the critical threshold characteristic of tokamak turbulence and measurement uncertainties, not all input plasma states result in unstable modes. Thus, a significant proportion of the computational budget to obtain the turbulent fluxes is spent on stable inputs, which can be wasteful. Active learning may be used to select the inputs that would be most useful to update the surrogate.

In this paper, we use active learning to build a data-efficient surrogate of plasma core turbulent transport based on the QuaLiKiz model [2, 3] and a dataset of JET shots for which QuaLikiz runs

have been obtained in [4]. The objective of Active Learning [5] is to efficiently allocate resources by only acquiring the labels that would be most beneficial to obtain. We propose a two-stage learning paradigm that exploits the binary nature of plasma core instabilities. Our framework uses a small initially labelled dataset to pre-train a stability classifier, and a regressor to predict the related turbulent fluxes. The initial dataset is then augmented by acquiring the labels of candidate inputs that are classified as likely to result in a growing instability, and that would most improve the range of validity of the surrogate. Our framework allows us to achieve the performance of a surrogate trained on the full-dataset but with only 0.4% of the data.

Data

We use an experimentally-based dataset from the JET (Joint European Torus) tokamak, which contains the turbulent transport calculation inputs and related QuaLiKiz outputs needed to make a surrogate model. This dataset was originally used in [4] to train neural network ensembles, without active learning. Here we focus on the Ion Temperature Gradient turbulence [6], for which only approximately 25% of inputs will develop turbulent transport, for a total of 4,000,000 points available. Specifically, we predict the heat flux of ions, $q_{i,ITG}$.

Active Learning

Given an unlabelled pool of experimental inputs, a classifier and a regressor are pretrained on a small random sample of data whose labels were obtained by running QuaLiKiz in [4]. Hereafter, the networks and the labelled dataset are updated following a two-stage active learning pipeline shown in Figure 1. The classifier is tasked with screening a sample of candidate points of size `CandidateSize` from the unlabelled pool. `CandidateSize` is chosen to be 10,000, which reflects the number of distinct inputs obtained from a single JET plasma discharge. An acquisition function selects which of the candidates should be labelled, and those inputs are then appended to the training data, to retrain the regressor. Specifically, we select `TopUncertain%` (set to 25%) of the inputs with the highest regressor variance. The points that were wrongly classified as unstable are stored in a buffer of size `ClassifierBuffer` (set to 200) which, when filled, is used to retrain the classifier along with the previous training data. Both the regressor and the classifiers are equipped with Dropout layers [7]. We use Monte Carlo Dropout [8] to estimate the uncertainty of the NN regressor. Although the classifier achieved an actionable performance in the pre-trained phase, more advanced acquisition functions based on both the regressor and classifier uncertainties will need to be explored.

The architecture for the classifier and regressor consists of a simple feed-forward neural network consisting of an input size of 15 and 5 hidden layers with sizes [128, 256, 512, 256, 128]. A dropout rate of 0.1 was employed in each layer along with ReLU activation functions. We used the `Adam` optimiser with a learning rate of 0.001, and a weight decay of 0.0001. The holdout set consists on 50,000 points.
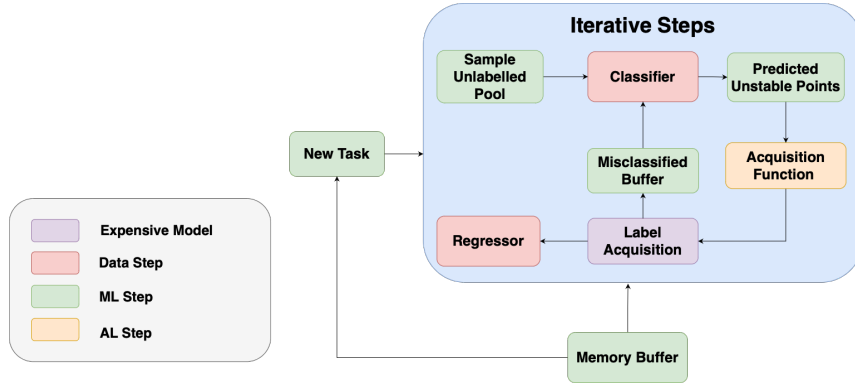
Figure 1: Schematic diagram of active learning framework. Each continual learning task consists of multiple iterations of the steps in the blue box.

Results

We conducted various experiments to investigate the effectiveness of an uncertainty-based acquisition function and the effect of the stability classifier on the pipeline performance. An initial training dataset size of 5,000 was used. Classifier performance across iterations was observed to be relatively unchanged, achieving a constant F1 score of 87%. We evaluate the performance of the Active Learning pipeline in terms of the Mean Squared Error (MSE) between the predicted and true fluxes on a holdout set,

$$MSE = \frac{1}{N} \sum_{j}^{N} (q_{j,true} - q_{j,pred})^2 \tag{1}$$

As shown in Figure 2, uncertainty-based sampling provides significantly improved performance compared to random sampling. Figure 2 shows that the combination of the classifier and the use of uncertainty-based acquisition provides a much greater decrease in MSE than using either individually. As only 25% of inputs result in a growing turbulence, the classifier stage of the pipeline provides data-efficiency.

Table 1 summarises the final test loss after 25 pipeline iterations. The combined classifier and uncertainty acquisition model achieves comparable performance to a model trained on the full dataset, but with only 0.4% of data. We conducted experiments using different initial training set sizes that are then evaluated at the same final training set size. It was found that using 5,000 points to pretrain the networks had a better final performance than when starting with 1,000 points, but further increases in the initial training set size provide increasingly marginal gains.

We note that the MSE is correlated with, but not necessarily fully reflective of, the surrogate quality in the final tokamak modelling application. Subtle features of the input-output mapping must also be captured, as shown in [9]. Future work will include more domain-specific metrics.
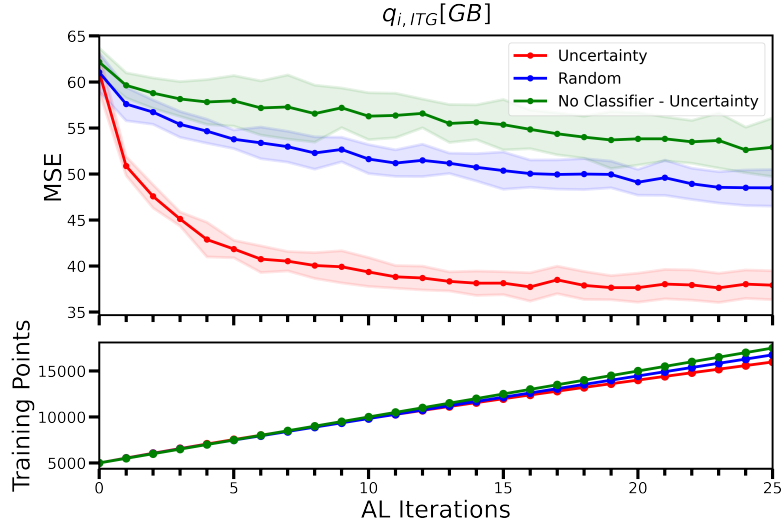
Figure 2: Test losses for $q_{i,ITG}[GB]$ using different acquisition methods. Bottom plots show the number of points sampled from the unlabelled pool at each iteration. An initial training dataset of 5,000 points was used. The lines and shaded areas are the means and one standard deviation of 10 random realisations.

Data Acquisition Method	$q_{i,ITG}[GB]$ MSE
No Classifier - Uncertainty	52.9 ± 3.2
Classifier - Random	48.5 ± 2.0
Classifier - Uncertainty	$37. \pm 1.6$
Full Dataset	38.6 ± 0.5

Table 1: The average test MSE and one standard deviation after 25 active learning pipeline iterations trained on the heat flux of ion temperature gradient turbulence in different experiments.

Conclusion

This paper investigates the use of a two-stage active learning pipeline for the modelling of turbulent transport in tokamak reactors. The inclusion of a classifier stage to identify regions of the input space that lead to growing turbulence modes improves the surrogate model performance and is more data efficient. Additionally, when used in combination with an uncertainty-based acquisition function, the performance of the pipeline matches that of full-dataset training but with only 0.4% of the data, which would have resulted in a significant reduction in computational time spent generating QuaLiKiz input-output mappings.

Acknowledgements

References
[1] EJ Doyle et al. In: *Nuclear Fusion* 47.6 (2007), S18. [2] C Bourdelle et al. In: *Plasma Physics and Controlled Fusion* 58.1 (2015), p. 014036. [3] J. Citrin et al. In: *Plasma Physics and Controlled Fusion* 59.12 (2017), p. 124005. [4] A. Ho et al. In: *Physics of Plasmas* 28 (2021), p. 032305. [5] Pengzhen Ren et al. In: *CoRR* abs/2009.00236 (2020). arXiv: 2009.00236. URL: https://arxiv.org/abs/2009.00236. [6] W Horton. In: *Reviews of Modern Physics* 71.3 (1999), p. 735. [7] Nitish Srivastava et al. In: *J. Mach. Learn. Res.* 15 (2014), pp. 1929–1958. [8] Yarin Gal and Zoubin Ghahramani. In: *international conference on machine learning*. PMLR. 2016, pp. 1050–1059. [9] Karel Lucas van de Plassche et al. In: *Physics of Plasmas* 27.2 (2020), p. 022310.