CCFE
CULHAM CENTRE FOR FUSION ENERGY

UK Atomic Energy Authority

UKAEA-CCFE-CP(25)26

L. Zanisi, A. Ho, T. Madula, J. Barr, J. Citrin, J. Buchanan, S. Pamela, V. Gopakumar

Cheap training sets for surrogate models of gyrokinetic turbulence

Cheap training sets for surrogate models of gyrokinetic turbulence

L. Zanisi, A. Ho, T. Madula, J. Barr, J. Citrin, J. Buchanan, S. Pamela, V. Gopakumar

This paper has been submitted to
29th IAEA Fusion Energy Conference (FEC 2023), London, 16-21 October 2023

CHEAP TRAINING SETS FOR SURROGATE MODELS OF GYROKINETIC TURBULENCE

L. Zanisi
Culham Centre for Fusion Energy
Abingdon, UK
Email: lorenzo.zanisi@ukaea.uk

A. Ho
Dutch Institute for Fundamental Energy Research
Eindhoven,Netherlands

T. Madula
University College London
London, UK

J. Barr
University College London
London, UK

J. Citrin
Dutch Institute for Fundamental Energy Research
Eindhoven, Netherlands

S. Pamela
Culham Centre for Fusion Energy
Abingdon, UK

V. Gopakumar
Culham Centre for Fusion Energy
Abingdon, UK

J. Buchanan
Culham Centre for Fusion Energy
Abingdon, UK

JET contributors
See author list of J. Mailloux et al 2022 Nucl. Fusion 62 042026

Abstract

Model-based plasma scenario development lies at the heart of the design and operation of future fusion powerplants. Including gyrokinetic turbulence in integrated models is essential for delivering a successful roadmap towards operation of ITER and the design of DEMO-class devices. Given the highly iterative nature of integrated models, fast machine-learning-based surrogates of gyrokinetic turbulence are fundamental to fulfil the pressing need for faster simulations opening up pulse design, optimization, and flight simulator applications. A significant bottleneck is the generation of suitably large training datasets covering a large volume in parameter space, which can be prohibitively expensive to obtain for higher fidelity codes.

In this work, we propose ADEPT (Active Deep Ensembles for Plasma Turbulence), a physics-informed, two-stage Active Learning strategy to ease this challenge. Active Learning queries a given model by means of an acquisition function that identifies regions where additional data would improve the surrogate. We provide a benchmark study using available data from the literature for the QuaLiKiz quasilinear model. We demonstrate quantitatively that the physics-informed nature of the workflow proposed reduces the need to perform simulations in stable regions of the parameter space, resulting in significantly improved data efficiency compared to non-physics informed approaches which consider a regression problem over the whole domain. We show an up to a factor of 20 reduction in training dataset size needed to achieve the same performance as random sampling. We then validate the surrogate on multichannel integrated modelling of ITG-dominated JET scenarios.

1. INTRODUCTION

Turbulent transport is the dominant transport mechanism in Tomamak plasmas and its presence 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 device. To this end, gyrokinetic models constitute a fundamental tool towards the delivery of ITER, DEMO-class reactors and beyond. However, the computational cost associated to integrating gyrokinetic models in highly iterative applications such as multichannel integrated models (e.g., JINTRAC, [2]) requires the delivery of fast and accurate surrogates. Feed-forward neural network (NN) surrogate models of the quasi-linear gyrokinetic models QuaLiKiz [3, 4, 5] and TGLF [6, 7, 8], have shown a factor 10^4 prediction speedup thus enabling real-time capable profile prediction [9], discharge optimisation studies [10] and integrated core-pedestal transport models [8, 11] at a fraction of the computational cost.

Due to the cost associated to retrieving large simulation databases to be adopted as training sets for these neural networks, previous works have focused on spanning a small volume in the input space. This restricts some of the current applications to small dimensionality and narrow range in parameter space [12, 13], or medium dimensionality [5], also sometimes based on experiments [8, 14]. Moreover, by the nature of the critical threshold characteristic of Tokamak turbulence, not all plasma states result in unstable modes. In previous work [5, 14], the consistency of the surrogate with the critical threshold behaviour was enforced by means of a physics-based loss function that encouraged a controlled extrapolation to negative values where the true output fluxes were null. Negative predictions would then be clipped to zero at inference time. Although effective, this strategy lacked in efficiency as it resulted in a large fraction of the computational budget being spent to obtain stable modes (roughly 40% in [14] across all the electrostatic modes resolved). Instead, we hypothesise that the boundary manifold between stable and unstable inputs may be learned more efficiently using a separate surrogate model.

This study proposes to build NN surrogate models of gyrokinetic turbulence by leveraging Active Learning ([15]) methods. Active Learning is a sequential sampling strategy that queries an expensive black box function (in our case a gyrokinetic model) by means of an acquisition function that identifies regions where additional data would improve the NN performance. We develop ADEPT (Active Deep Ensembles for Plasma Turbulence), a two-stage AL framework where a surrogate of the critical gradient threshold determines whether a given input will result in growing modes, and a regressor predicts the output turbulent transport fluxes. We focus on an acquisition function that queries inputs for which the output uncertainty of the NN is highest, thus maximising informativeness [16]. Deep Ensembles [17], which provide state-of-the-art uncertainty quantification capabilities for NNs, are adopted as the surrogate model.

We provide a demonstration of the ADEPT pipeline using an existing large database of QuaLiKiz simulations obtained from JET inputs [14]. For this proof-of-concept work we focus on ITG turbulence only. As the input-output mappings are already available in the dataset, we can easily test the performance of the two-stage workflow.

The paper outline is as follows. We describe the dataset in Section 2, we introduce the ADEPT methodology in Section 3.1 and we outline the integrated modelling framework in Section 4. In Section 5 we give the first main result of the paper. In Section **??** we validate the surrogate on integrated modelling of ITG-dominated JET scenarios. Finally, in Section 7 we discuss the results obtained, identify remaining issues and propose potential solutions to be explored in future work.

2. DATA

We use the existing JET-Exp-15D dataset devised in [14, 18]. The dataset is built upon 2135 JET discharges including a variety of plasma scenarios, augmented taking into account measurement uncertainties for the parameters that turbulence is most sensitive to. The output encompasses the multiple channels of transport of ITG, ETG and TEM turbulence obtained from QuaLiKiz. The raw dataset produced from all the available inputs was subjected to consistency checks to either enforce physical consistency within the data or discard abnormally large heat fluxes and abnormally small particle fluxes; see Table 6 of [14] for more details.

In the remainder of this work the focus will be on ITG turbulence for which only less than 25% of inputs in the JET-Exp-15D dataset develop turbulent transport. The transport fluxes considered (in GyroBohm units, cfr. Table 3 of [14]) are the heat flux of ions, $q_{i,ITG}$, the heat flux of electrons $q_{e,ITG}$, the momentum flux of ions, $\Pi_{i,ITG}$, the particle flux of electrons, $\Gamma_{e,ITG}$, and the particle flux of ions $\Gamma_{i,ITG}$.

3. DATA-EFFICIENT SURROGATE MODELS

3.1. Active Learning

Active Learning (AL, e.g., [15] for a review) is a sampling strategy that aims at reducing the amount of training data needed to obtain a performing surrogate. An AL system comprises three components: a learner, an oracle

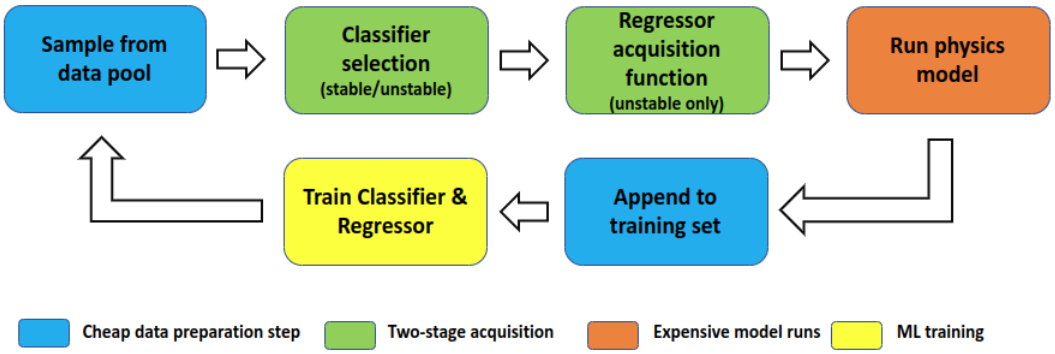

FIG. 1. Schematic diagram of the two-stage physics-informed AL workflow used in this work. Given a data pool for which only inputs are available, a classifier evaluates the likelihood of a given input in the pool resulting in unstable modes. The acquisition function is evaluated on the unstable inputs, and a batch of the most uncertain ones are selected to be run through the gyrokinetic model. The newly obtained input-output mappings are used to train both NNs. This strategy is repeated until the computational budget has exhausted or the performance of the surrogates is deemed actionable.

and a query strategy. The learner is a ML algorithm that improves its performance as more data is collected from the oracle according to the query strategy. The oracle is a costly data acquisition system that provides the training data for the learner.

The sequential strategy proposed in [16] queries the inputs for which the surrogate model's predictive uncertainty is largest,

$$x_{query} = \underset{x \in \mathcal{U}}{arg\,max}\ \sigma^2(x; D_{train,t}), \tag{1}$$

$$D_{train,t+1} = x_{query} \cup D_{train,t} \tag{2}$$

where $\sigma(x; D_{train,t})$ is the output uncertainty of a learner trained on a dataset $D_{train,t}$, t is the current iteration and $\mathcal{U}$ is a pool of inputs (e.g., the plasma states) for which the outputs (e.g., the turbulent fluxes) are not available. As indicated in the expression above, the dataset at the next iteration is enriched with data obtained from the query. Here, we adopt Batch Mode AL (e.g., [19]), which consists in performing the acquisition for the M inputs that rank highest in the model uncertainty. Batch Mode AL is more suitable for NNs, as retraining a NN with just one new sample is impractical.

3.2. Physics-informed Active Learning for gyrokinetic models with ADEPT

Below we propose ADEPT (Active Deep Ensembles for Plasma Turbulence)[1] a two-stage, physics-informed active learning strategy that delivers a significant reduction in the volume of data required to train performing surrogates. Contrary to previous work [**plascche2020**, 14], we adopt a separate NN that classifies whether an input is in a stable or unstable region. This setup preserves the physics-informed nature of the framework proposed by [5], but it splits the burden of identifying the critical gradients and regressing to turbulent fluxes between two highly specialised NNs. Given a data pool $\mathcal{U}$ of inputs, a NN classifier and a NN regressor are pretrained on a small (20,000 points) random sample of data for which the input-output mapping is available. Hereafter, for each iteration, the networks and the labelled dataset are updated following the strategy shown in Figure 1:

- The classifier is tasked with screening a sample of candidate points in the data pool $\mathcal{U}$. This is the physics-informed stage of the workflow.
- The acquisition function in eq. 1 uses the regressor's uncertainty (in our case, the epistemic uncertainty in eq. 7) to select a batch of size `acquisition_batch` candidates.
- The outputs for the input candidates selected are queried from the model of choice (QuaLiKiz in the case of this paper);

[1] We adopt the `PyTorch` library.,

- The newly available input-output mappings are appended to the training data.
- Both the regressor and the classifier NNs are trained again.

The loop above is repeated until the computational budget has exhausted, or the surrogates have reached the desired performance.

As gyrokinetic turbulence involves multichannel transport, it is necessary to maximise the information gain for all the fluxes, and therefore the acquisition function in eq. 1 becomes

$$x_{query} = \underset{x \in \mathcal{U}}{arg\,max} \sum_{k=1}^{N_{fluxes}} \sigma_k^2(x; D_{train,t}). \tag{3}$$

3.3. Surrogate uncertainty via Deep Ensembles and its uses within integrated modelling

It has long been established that NN models give overconfident predictions that are factually wrong (e.g., [20]). Equipping NNs with a notion of uncertainty in their own predictions has since become a mainstream line of research producing a rich literature [21].

Deep Ensembles [17] are the state of the art in uncertainty quantification methods for NNs. They model the target distribution as a uniformly weighted mixture model and optimise its negative log-likelihood:

$$NLL = -\frac{1}{M}\sum_{k=1}^{M} \log p_{\theta_k}(y|x) \tag{4}$$

where p_{θ_k} it the $k-$th NN parametrised by a set of parameters θ.

For classification, the objective is the usual binary cross-entropy loss. For regression problems, the Gaussian negative log-likelihood is adopted,

$$-\log p_\theta(y|x) = \frac{\log \sigma_\theta^2(x)}{2} + \frac{\big(y - \mu_\theta(x)\big)^2}{2\sigma_\theta^2(x)} + const. \tag{5}$$

A NN with two output neurons trained with the objective above will explicitly learn the mean $\mu_\theta(x)$ and variance $\sigma_\theta^2(x)$.

The mean and variance of the deep ensemble as a whole are:

$$\mu_E = \frac{1}{M}\sum_{k=1}^{M} \mu_{\theta_k} \tag{6}$$

$$\sigma^2_{Ensemble} = \underbrace{\frac{1}{M}\sum_{k=1}^{N} \sigma^2_{\theta_k}}_{\text{aleatoric}} + \underbrace{\big(\frac{1}{M}\sum_{k=1}^{M} \mu^2_{\theta_k}\big) - \mu_E^2}_{\text{epistemic}}. \tag{7}$$

The first contribution on the right hand side of the expression for the variance is the average variance between all members. The second contribution is the variance of the means of the ensemble, as shown in the last two terms on the right hand side of eq. 7. The two contributions are sometimes interpreted as the the epistemic uncertainty (i.e. the uncertainty of the model) and the aleatoric uncertainty (i.e. the irreducible noise in the data), e.g. [22]. The epistemic uncertainty is the natural choice to use in the acquisition function, as we seek to improve the inherent accuracy of the model regardless of data noise [23].

Uncertainty quantification capabilities are also a natural feature of the classifier NN. The confidence of the classifier can be defined as its output probability of a point being unstable. Probabilities close to a value of 0.5 inform downstream applications that performing a run of the original QuaLiKiz model is recommended.

3.4. Details of the training procedure

We borrow from [5] the idea of fitting NN surrogate models to the "leading flux" of a given turbulence type and the flux ratio between the leading flux and the secondary fluxes.

We train a suite of deep ensembles, each regressing to one turbulent flux, and one deep ensemble classifier for the stability boundary. We adopt 5 ensembles per model, each consisting of 8 layers with 512 parameters each and `ReLU` activation functions. Each model is trained for 200 epochs with 30 epochs of patience, a weight decay of λ=10^{-4} and a `batch_size` of 512. For the regressor we adopt the NLL loss eq. 5 and for the classifier the binary crossentropy loss. Each `acquisition_batch` consists of 512 training samples, doubling every 30 acquisitions due to the costs associated with retraining NNs on large datasets.

4. INTEGRATED MODELLING

The JINTRAC integrated modelling suite [2] was chosen for this study both due its history of integration with QLKNN [14, 5] and its relevance for ITER scenarios [24]. The JINTRAC integrated model test cases and settings used in this study were taken from those used to validate QLKNN-jetexp [14], specifically selecting:

- the H-mode carbon wall discharge (JET#73342) [4];
- and the H-mode beryllium wall discharge (JET#92436) [18].

A summary of the major JINTRAC settings used for these test cases are provided in Table 7 of [14]. In the following Sections we will compare the results of adopting either ADEPT, QLKNN-jetexp and the original QuaLiKiz within JINTRAC. Specifically, the critical gradient threshold in ADEPT will be estimated using the trained classifier, while it is estimated according to the methodology summarised in the Introduction for the QLKNN-jetexp surrogates.

ADEPT generates surrogate models that inherently include a measure of uncertainty. This characteristic can be leveraged in integrated modeling to evaluate the level of confidence that should be placed in the surrogate model's predictions, and perform a QuaLiKiz run whenever the uncertainty of surrogate is not considered acceptable. An in-depth study of the impact of the precise acceptance threshold on the integrated modelling results is outside the scope of this work.

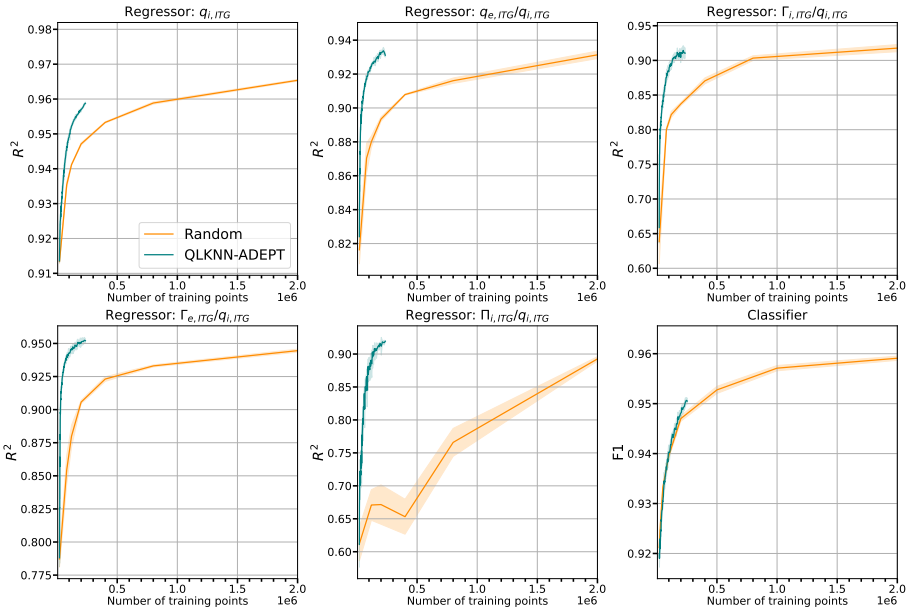

FIG. 2. Active Learning vs random sampling performance. We use the R^2 for the deep ensembles regressing to the electron heat flux and the flux ratios involving the ion heat flux and the ion momentum flux. The $F1$ score for the stability boundary classifier is also shown. The shaded areas represent the standard deviation of 5 runs with different random seed. Active Learning improves data efficiency by at least factor of 2 (and up to a factor of 20) compared random sampling.

5. RESULTS: CHEAPER TRAINING SETS WITH ACTIVE LEARNING

We evaluate the surrogates in terms of the R^2 score for the regressors and the F1 score for the classifier.

Figure 2 shows the performance of ADEPT on ITG turbulence compared to random selection as a function of number of training samples collected. It can be seen that ADEPT provides up a factor of 20 data reduction compared to random sampling. An important contribution to this success is the inclusion of the classifier stage, which allows for a more data-efficient learning of the manifold where unstable turbulent fluxes develop.

TABLE 1. THE PERFORMANCE OF ADEPT TRAINED WITH UP TO 200,000 SAMPLES COMPARED TO THE NNS PRESENTED IN [14], IN TERMS OF R^2 SCORE FOR THE FLUXES.

F1	$q_{i,ITG}$	$q_{e,ITG}/q_{i,ITG}$	$\Gamma_{i,ITG}/q_{i,ITG}$	$\Gamma_{e,ITG}/q_{i,ITG}$	$\Pi_{i,ITG}/q_{i,ITG}$
ADEPT 0.9504	0.9585	0.9366	0.9174	0.9554	0.9108
[14] 0.7791	0.951776	0.950541	0.698693	0.526750	0.914011

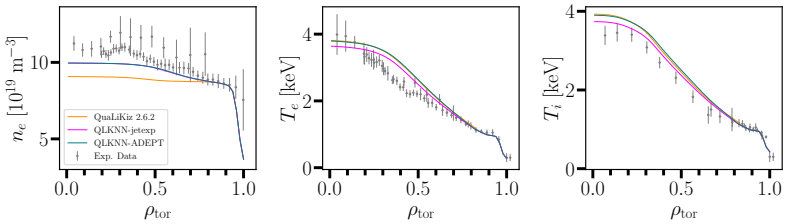

FIG. 3. Comparison of the steady-state profiles from the simulation of JET#73342.

A performance comparison on a test set between ADEPT and the surrogates presented in [14] is given in Tables 1 and **??**. Although the surrogates in Ref. [14] do not explicitly employ a separate classifier NN to model the critical gradient, they achieve a comparable effect by zeroing out all fluxes when the leading flux is predicted to be negative. Therefore, we can evaluate the F1 performance of these surrogates, as they effectively exhibit classifier-like behavior in this context. It can be seen that the performance of ADEPT in terms of the F1 score is comparable or superior to that of the NNs in Ref. [14], albeit with two orders of magnitude fewer data. In particular, a high classifier performance in the case of ADEPT is crucial in ensuring that sampling does not occur deep in the stable regions. As a proof, we have computed that the contribution of stable inputs to the training set of the classifier is 75% in the random sampling case (and, indeed, the case of Ref. [14]), while this drops to around 20% in the case of ADEPT. The surrogates of [14], instead, feature a poor Precision, showing a high number of non-zero flux predictions in the stable region. The latter surrogates, however, achieve a better Recall. ADEPT would need more data to reach the same kind of performance. The classifier is not currently included in the acquisition function explicitly, which instead will be crucial to improve its data efficiency compared to random sampling. This feature will be explored in future work.

6. VALIDATION ON ITG-DOMINATED JET DISCHARGES

Figures 3 and 4 show the steady state profiles obtained by adopting ADEPT, the original QLKNN-jetexp surrogates of Ref. [14] and the original QuaLiKiz model. The experimental data is shown here for reference, however the purpose of this test is to verify whether the surrogate models are able to reproduce the behaviour of QuaLiKiz within JINTRAC. Table 2 provides the profile-averaged relative RMS (RRMS) for these JINTRAC runs using their respective networks, with the QLKNN-jetexp reference given inside the square brackets. The RRMS is computed as:

$$RRMS = \sqrt{\frac{1}{N}\sum_i \frac{\left(Y_{NN,i} - Y_{QLK,i}\right)^2}{Y^2_{QLK,i}}} \tag{8}$$

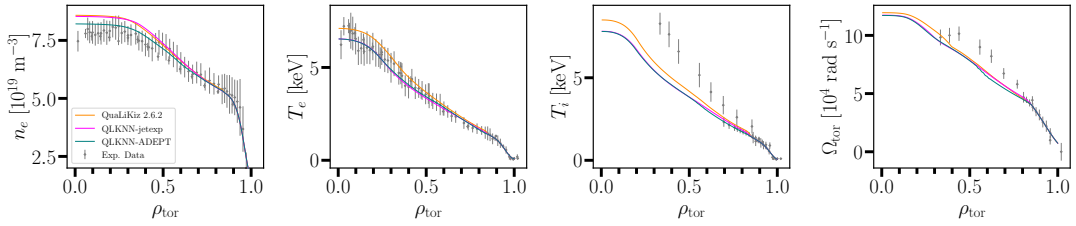

FIG. 4. Comparison of the steady-state profiles from the simulation of JET#92436.

TABLE 2. SUMMARY TABLE OF THE JINTRAC PREDICTED PROFILE RRMS WITHIN THE QUALIKIZ EVALUATION REGION. THE VALUES ARE GIVEN FOR THE QLKNN-ADEPT SIMULATION, WITH THE REFERENCE QLKNN-JETEXP SIMULATION PROVIDED WITHIN THE SQUARE BRACKETS.

	JET#73342	JET#92436
T_e RRMS	1.1% [4.1%]	8.0% [8.8%]
T_i RRMS	1.7% [3.3%]	9.1% [7.3%]
n_e RRMS	7.7% [7.7%]	2.7% [0.7%]
Ω_{tor} RRMS	– [–]	4.0% [3.0%]

where the sum is over the number of radial points, and $Y_{NN,i}$ and $Y_{QLK,i}$ indicate the profiles computed using the NN prediction and QLK respectively.

Both the surrogate models considered achieve a match with QuaLiKiz that is better than 10%. This suggests that surrogate models can effectively replace the original transport model as a drop-in replacement for obtaining steady-state profiles.

7. SUMMARY AND CONCLUSIONS

We presented ADEPT (Active Deep Ensembles for Plasma Turbulence), a two-stage physics-informed Active Learning framework for data-efficient surrogate models of gyrokinetic turbulence. ADEPT consists of a classifier NN that learns the boundary manifold between regions that are stable and unstable under linear gyrokinetic turbulence, thus simultaneously providing a model for the critical gradient and limiting the search space of a second surrogate regressing to the turbulent transport fluxes. Using an existing large dataset of QuaLiKiz simulations based on experimental JET discharges [18, 14], we have demonstrated a sizeable reduction in the size of the training dataset needed to obtain surrogates with integrated performance metrics comparable to surrogates trained by sampling at random. The reduction factor can be up to a factor of 10 or more, and it is due to both the adoption of Active Learning and the physics-informed nature of ADEPT, which alone results in increased data efficiency, as only a minority of the QuaLiKiz runs would have been performed deep into the stable regions of the parameter space, thus limiting the need to run costly and uninformative simulations. We also showed quantitative agreement with QuaLiKiz and previous surrogates in integrated modelling applications to two very different JET discharges. Finally, we identify the classifier stage of ADEPT as relevant for any other model where restricting the parameter space to a certain region with desirable properties is useful, such as the case of building surrogate models of codes modelling magnetohydronynamic instabilities [25].

REFERENCES

[1] J. D. Callen. “Transport processes in magnetically confined plasmas”. In: *Physics of Fluids B: Plasma Physics* 4.7 (1992), pp. 2142–2154. DOI: 10.1063/1.860021.

[2] M Romanelli et al. “JINTRAC: A system of codes for integrated simulation of Tokamak scenarios”. In: *Plasma and Fusion Research* 9 (Jan. 2014). DOI: 10.1585/pfr.9.3403023.

[3] C Bourdelle et al. “Core turbulent transport in tokamak plasmas: bridging theory and experiment with QuaLiKiz”. In: *Plasma Physics and Controlled Fusion* 58.1 (Dec. 2015), p. 014036. DOI: 10.1088/0741-3335/58/1/014036. URL: https://doi.org/10.1088/0741-3335/58/1/014036.

[4] J Citrin et al. “Tractable flux-driven temperature, density, and rotation profile evolution with the quasilinear gyrokinetic transport model QuaLiKiz”. In: *Plasma Physics and Controlled Fusion* 59.12 (Nov. 2017), p. 124005. DOI: 10.1088/1361-6587/aa8aeb. URL: https://doi.org/10.1088/1361-6587/aa8aeb.

[5] K. L. van de Plassche et al. “Fast modeling of turbulent transport in Fusion Plasmas using neural networks”. In: *AIP Publishing* (Feb. 2020). URL: https://aip.scitation.org/doi/10.1063/1.5134126.

[6] G. M. Staebler, J. E. Kinsey, and R. E. Waltz. “A theory-based transport model with comprehensive physics”. In: *Physics of Plasmas* 14.5 (May 2007), p. 055909. DOI: 10.1063/1.2436852. URL: https://doi.org/10.1063/1.2436852.

[7] G. M. Staebler and J. E. Kinsey. “Electron collisions in the trapped gyro-Landau fluid transport model”. In: *Physics of Plasmas* 17.12 (Dec. 2010), p. 122309. DOI: 10.1063/1.3505308. URL: https://doi.org/10.1063/1.3505308.

[8] O. Meneghini et al. “Self-consistent core-pedestal transport simulations with neural network accelerated models”. In: *Nuclear Fusion* 57.8 (July 2017), p. 086034. DOI: 10.1088/1741-4326/aa7776. URL: https://doi.org/10.1088/1741-4326/aa7776.

[9] F. Felici et al. “Real-time-capable prediction of temperature and density profiles in a tokamak using RAPTOR and a first-principle-based transport model”. In: *Nuclear Fusion* 58.9 (July 2018), p. 096006. DOI: 10.1088/1741-4326/aac8f0. URL: https://doi.org/10.1088/1741-4326/aac8f0.

[10] S. Van Mulders et al. “Rapid optimization of stationary tokamak plasmas in RAPTOR: demonstration for the ITER hybrid scenario with neural network surrogate transport model QLKNN”. In: *Nuclear Fusion* 61.8 (July 2021), p. 086019. DOI: 10.1088/1741-4326/ac0d12. URL: https://doi.org/10.1088/1741-4326/ac0d12.

[11] O. Meneghini et al. “Neural-network accelerated coupled core-pedestal simulations with self-consistent transport of impurities and compatible with ITER IMAS”. In: *Nuclear Fusion* 61.2 (Dec. 2020), p. 026006. DOI: 10.1088/1741-4326/abb918. URL: https://doi.org/10.1088/1741-4326/abb918.

[12] P. Rodriguez-Fernandez, N.T. Howard, and J. Candy. “Nonlinear gyrokinetic predictions of SPARC burning plasma profiles enabled by surrogate modeling”. In: *Nuclear Fusion* 62.7 (May 2022), p. 076036. DOI: 10.1088/1741-4326/ac64b2. URL: https://doi.org/10.1088/1741-4326/ac64b2.

[13] Ionuţ-Gabriel Farcaş, Gabriele Merlo, and Frank Jenko. “A general framework for quantifying uncertainty at scale”. In: *Communications Engineering* 1.1 (Dec. 2022). DOI: 10.1038/s44172-022-00045-0. URL: https://doi.org/10.1038/s44172-022-00045-0.

[14] A. Ho et al. “Neural network surrogate of qualikiz using jet experimental data to populate training space”. In: *AIP Publishing* (Mar. 2021). URL: https://aip.scitation.org/doi/full/10.1063/5.0038290.

[15] Charu C. Aggarwal et al. “Active learning: A survey”. English (US). In: *Data Classification*. Publisher Copyright: © 2015 by Taylor Francis Group, LLC. CRC Press, Jan. 2014, pp. 571–605. ISBN: 9781466586741. DOI: 10.1201/b17320.

[16] David J. C. MacKay. “Information-Based Objective Functions for Active Data Selection”. In: *Neural Computation* 4.4 (July 1992), pp. 590–604. ISSN: 0899-7667. DOI: 10.1162/neco.1992.4.4.590. eprint: https://direct.mit.edu/neco/article-pdf/4/4/590/812354/neco.1992.4.4.590.pdf. URL: https://doi.org/10.1162/neco.1992.4.4.590.

[17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. *Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles*. 2017. arXiv: 1612.01474 [stat.ML].

[18] A. Ho et al. “Application of Gaussian process regression to plasma turbulent transport model validation via integrated modelling”. In: *Nuclear Fusion* 59.5 (Mar. 2019), p. 056007. DOI: 10.1088/1741-4326/ab065a. URL: https://doi.org/10.1088/1741-4326/ab065a.

[19] David Holzmüller et al. “A Framework and Benchmark for Deep Batch Active Learning for Regression”. In: (2022). arXiv: 2203.09410 [stat.ML].

[20] Anh Nguyen, Jason Yosinski, and Jeff Clune. *Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images*. 2015. arXiv: 1412.1897 [cs.CV].

[21] Jakob Gawlikowski et al. *A Survey of Uncertainty in Deep Neural Networks*. 2022. arXiv: 2107.03342 [cs.LG].

[22] Fredrik K. Gustafsson, Martin Danelljan, and Thomas B. Schön. *Evaluating Scalable Bayesian Deep Learning Methods for Robust Computer Vision*. 2020. arXiv: 1906.01620 [cs.LG].

[23] Yehor Yudin et al. In: Springer Nature Switzerland, 2023, pp. 572–586. DOI: 10.1007/978-3-031-36027-5_45. URL: https://doi.org/10.1007%2F978-3-031-36027-5_45.

[24] E. Militello Asp et al. “JINTRAC integrated simulations of ITER scenarios including fuelling and divertor power flux control for H, He and DT plasmas”. In: *Nuclear Fusion* 62.12 (Oct. 2022), p. 126033. DOI: 10.1088/1741-4326/ac90d4. URL: https://dx.doi.org/10.1088/1741-4326/ac90d4.

[25] A B Mikhailovskii. “Generalized MHD for numerical stability analysis of high-performance plasmas in tokamaks”. In: *Plasma Physics and Controlled Fusion* 40.11 (Nov. 1998), pp. 1907–1920. DOI: 10.1088/0741-3335/40/11/007. URL: https://doi.org/10.1088/0741-3335/40/11/007.