CONFERENCE PRE-PRINT

STEP AND MULTIMACHINE TGLF NEURAL SURROGATES

L. Zanisi, F.J. Casson, B.S. Patel, H.G. Dudding, Z. Stancar, C. M Roach, A.P. Snodin, C. Siddle, J. Buchanan, S. Pamela, T. Brown
Culham Science Centre
Abingdon, UK
Email: lorenzo.zanisi@ukaea.uk

A. Järvinen, A. Kit
VTT Technical Research Centre of Finland, Finland

M. Marin
EPFL, Swiss Plasma Centre
SPC, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland

A. Ho, G. Snoep
MIT, Plasma Science and Fusion Center, Cambridge, Massachussets,USA

E. Vergnaud
Aix-Marseille Université, CNRS, PIIM, Marseille, France

C. Bourdelle
CEA, IRFM, Saint-Paul-lez-Durance, France

O. Meneghini, T. Neiser
General Atomics
San Diego, USA

JET contributors
See author list of C. Maggi et al. 2024 [1]

Abstract

This work summarises activities to produce fast neural network (NN) surrogate models of the TGLF quasilinear model of core turbulent transport in tokamaks. Two applications are presented. Firstly, TGLF surrogate models applicable to a host of experimental devices with aspect ratio higher than 3 are devised. These complement existing surrogates for the QuaLiKiz quasilinear model, while capturing additional poloidal shaping dependencies that are not included in QuaLiKiz, for a total of 13 dimensions spanned, and pushing towards full-radius applications in L-mode. NN surrogates for TGLF have also been generated for the equilibrium space that will be spanned in the I_p ramp up of the STEP power plant. Surrogate-enabled flux-driven integrated modelling runs are shown for an example Deuterium JET discharge and for a test STEP ramp up case, demonstrating good agreement with the simulations using TGLF directly, and with a reduction in runtime of two orders of magnitude. The achieved speed and accuracy of the surrogate models show promise towards large-scale validation of TGLF on experimental devices and high-throughput accelerated optimisation of STEP via integrated modelling. The surrogate models developed are open source and available at `https://github.com/ukaea/tglfnn-ukaea/`.

1. INTRODUCTION

Predictive modelling suitable for optimizing plasma scenarios remains impractical with first-principle based models, due to the need to evaluate turbulent transport models thousands of times within flux driven integrated modelling (such as JINTRAC and ASTRA [2, 3]) when simulating dynamics over several confinement timescales. While these limitations are already clear for current experimental machines, they are exacerbated for ITER and other power plant devices, given the longer transport timescales expected. For instance, a single transport simulation for the STEP ramp-up ($\gtrsim$ 2000s) [4] with an intermediate fidelity model such as the TGLF quasilinear model [5, 6] can take several weeks. Therefore, present optimisation strategies in integrated models for both the flat top [7] and ramp-up phases of STEP [8] are limited to low-fidelity transport models (e.g., Bohm/gyro-Bohm). The ITER ramp up was modelled with a similar level of fidelity [9]. In general, there is a pressing need for computationally feasible predict-first approaches that are necessary to design reactor-scale devices.

Surrogate models that provide a fast approximate solution to core transport given plasma states offer a solution to this computational stalemate. Quasilinear transport models such as TGLF and QuaLiKiz [10, 11] are particularly amenable for this purpose as they feature a good balance between speed and accuracy, and training sets on the scale of millions of simulations can be obtained in a matter of days on a single compute node. For instance, Neural Networks (NN) surrogate models of QuaLiKiz [12, 13] have demonstrated impressive speedups, reducing integrated transport simulation runtimes from days to mere minutes, with applications ranging from modelling and optimisation of ITER scenarios [14, 15] and flat-top predictions to challenging ramp-up modelling in JET [13]. Surrogate models of TGLF have been devised to model DIII-D discharges [16, 17].

At the same time, quasilinear models rest on an ad-hoc rule to describe turbulence saturation, and the full extent to which these assumptions affect their capability to reproduce a wide range of experimental conditions in flux-driven integrated simulations needs to be assessed to build confidence in their predictive power. Surrogate-accelerated integrated models, managed with large-scale validation tools capable of running and monitoring many thousands of simulations in parallel (such as DUQtools [18] and FUSE **??**) enable an unprecedented modelling validation across experiments and parameter spaces, allowing identification of the regions where model improvement is required. Efforts towards a large-scale database validation performance on MAST-U, NSTX and DIII-D are already under way using TGLF surrogate models [19, 20]. Validation in the parameter space of other devices would complement current efforts to help determine the range of applicability of TGLF. In particular, there is an interest to obtain predictions from flux-driven integrated transport simulations using TGLF up to the separatrix in L-mode, e.g. [21]. In this paper, a surrogate model of TGLF, TGLFNN-UKAEA-MULTIMACHINEHYPER, is provided in a space that encompasses tokamaks with aspect ratio A$\geq$3. The surrogate model is based on a hypercube that largely mimics the work that [12] carried out for QuaLiKiz, with added shaping dimensions and extending up to the separatrix. The surrogate model is demonstrated on an L-mode Deuterium JET discharge within JINTRAC.

TGLF is also the only fast quasilinear model available that includes electromagnetic and shaping effects, thus making it suitable for modelling STEP. The compact design of STEP implies limited space for the central column, which will therefore host only a small solenoid suitable for startup. A non-inductive current ramp-up will therefore be driven in a low-density regime where ECCD is most efficient, and plasma densification will be necessary to transition to the burning plasma phase where bootstrap current dominates [22]. The required auxiliary power to drive the current depends on confinement and therefore, from the modelling perspective, on transport assumptions. Current capability to model the I_p ramp-up in STEP is limited to a modified Bohm/gyro-Bohm scaling as shown in [8]. Although [8] presented an efficient optimisation scheme, this low fidelity transport assumption is considered insufficient to deploy in practice due to the very nonlinear relationship between confinement, ECCD and auxiliary power. A higher fidelity model such as TGLF would provide a better physics basis. However, a fully non-inductive ramp-up is particularly challenging to achieve, with many competing constraints and optimisation objectives at play. The complexity of the optimisation landscape will likely require querying a very large number of transport simulations to achieve the optimum. This issue is further exacerbated by the lack of data for a reactor-scale spherical tokamak, which requires performing extensive uncertainty quantification, i.e. even more simulation runs. A surrogate model of TGLF for the STEP ramp up would enable higher fidelity trajectory optimisation compared to current capabilities. Two surrogate models based on training spaces obtained with different methodologies, TGLFNN-STEP-V2 and TGLFNN-STEP-V1, are presented and applied to a test ramp-up trajectory, thus demonstrating the usefulness of surrogate models for reactor design.

This paper is structured as follows: section 2 details the strategy to build the input space for the surrogate models, Section 3 shows the standalone performance of the surrogate models, and Section 4 shows the application of the surrogate models in flux-driven transport simulations.

2. DATA

Surrogate models require a training data set of input–output pairs from which they can learn, and the scope of this data set largely determines the model's regime of applicability. Ideally, the training data would comprehensively cover all relevant input dimensions of the underlying model to ensure completeness. In practice, however, this is rarely feasible because of the 'curse of dimensionality' [23]. As dimensionality increases, learning becomes progressively more difficult: (i) sample points become increasingly sparse in high-dimensional spaces, and (ii) extending the range of even a single input dimension leads to an exponential growth in the volume of parameter space, further compounding the sampling problem. Some data–efficient Active Learning approaches [24] to alleviate this issue have been developed and demonstrated to build surrogate models of QuaLiKiz [25, 26] and GS2 [27, 28]. These strategies rely on iterative acquisition of data by running simulations only in regions of parameter space deemed most informative, and require retraining at every acquisition. The value of such strategies lies mostly in settings where medium-sized datasets need to be collected, and where simulations are very expensive. If very large databases need to be collected, and if simulators are relatively fast, as is the case for quasilinear transport models, the training time ends up dominating the computational cost, making Active Learning not worthwhile in this specific scenario. The strategies for building the input spaces of the surrogate models introduced in this work are detailed below.

2.1. 13D multimachine hypercube

A straightforward strategy to keep the number of simulation runs manageable is to simply reduce the number of dimensions spanned and adopting space-filling methods such as Latin Hypercube Sampling (LHS). A subset of important features for core transport surrogate models was identified in the work by [12] for obtaining surrogates of QuaLiKiz. The same dimensions and ranges are spanned in this work with LHS. See Table 1 in [12] for more details. Shafranov shift (Δ), triangularity (δ) and elongation (κ), which were not included in [12] due to QuaLiKiz's geometry assumptions, are also added as dimensions, with $\Delta \in [-0.3, 0]$, $\delta \in [-0.2, 0.2]$ and $\kappa \in [1.2, 1.5]$. Furthermore, the ExB shear is accounted for explicitly, with $VEXB\_SHEAR = -SIGN(I_{tor})\frac{r}{q}\frac{\partial}{\partial r}\left(\frac{V_{EXB}}{R}\right)\frac{a}{c_s} \in [-0.2, 0.2]$. The pressure gradient is computed self-consistently from the quantities in the hypercube, thus avoiding adding an extra dimension to the input space.

The structure of the TGLF input space separates major and minor radii as two independent inputs. In this work, the flux surface centroid major radius of the separatrix RMAJ_LOC= R_{maj}/a is set to 3. As the radial coordinate RMIN_LOC= $r/a \in [0, 1]$ and the gradients are normalised with respect to a, a range of aspect ratios $RMAJ_LOC/RMIN_LOC = R_{maj}/r$ is covered, specifically all aspect ratios higher than 3. The saturation rule SAT2 is chosen for this study as an example, but extension to the other saturation rules is planned. Figure 1 shows the resulting input and output space distributions. Although the original space is a Latin hypercube, the distributions shown here depart from the uniformity typically expected of Latin Hypercube Sampling, as they arise from restricting the output fluxes to values below 200 GB in absolute value — the range most likely to be queried by downstream models. As expected, this cut mostly affects regions of the space with high gradients, higher radial locations, high ion-to-electron temperature ratio, low collisionality and high safety factors. Regions of low magnetic shear are discarded, while distributions in poloidal shaping and ExB shear are unaffected.

We denominate the surrogate model trained on this space TGLFNN-UKAEA-MULTIMACHINEHYPER. Its main limitations are that it cannot handle impurity transport and electromagnetic effects, which have been shown to be important in some scenarios [29, 30]. Momentum transport is also excluded.

2.2. STEP

Table 1 shows the input TGLF dimensions included for STEP. Compared to TGLFNN-UKAEA-MULTIMACHINEHYPER, the derivatives of triangularity and elongation are also added to the parameter space. Electromagnetic effects in STEP are also important and therefore β is also included. The EXB shear is not included in this iteration but it will be in future work. The total number of dimensions for the input space considered for STEP is 15, which is higher than for TGLFNN-UKAEA-MULTIMACHINEHYPER. A strategy that allows for a higher dimensionality while keeping the number of required training data within reasonable limits consists in exploiting the full covariance of the experimental space, as done in [16] and [13], which allows to discard regions of parameter space that are unlikely to be queried in downstream applications. However, for STEP, an experimental space is not available and therefore an adaptation of this strategy is required. Specifically, the input space correlations are harvested using Pyrokinetics [31] from a suite of ramp-up transport simulations which were run with JETTO. However, running transport simulations with TGLF is expensive in the first place, and therefore obtaining an extensive input space can be costly. In this work, two different spaces are obtained from either simulations with Bohm/gyro-Bohm transport [4], which run in about two days each, or with the TGLF transport model, which take around 2 weeks each on 60 cores. From these spaces, distinct TGLF surrogate models are obtained to assess the performance in production of a cheaply produced input space (i.e. with Bohm/gyro-Bohm transport) versus an expensive to obtain input space (i.e. with TGLF transport). The surrogate models obtained from these spaces are denominated TGLFNN-STEP-V1 and TGLFNN-STEP-V2 respectively. The input dimensions considered in both cases are summarised in Table 1. As in TGLFNN-UKAEA-MULTIMACHINEHYPER, the pressure gradient is computed self-consistently from the quantities available in the input space.

For each space, a Kernel Density Estimate (KDE) was produced to obtain a smooth distribution that preserves the data covariance. The correlations were broadened artificially by 10%, partially removing the correlations and

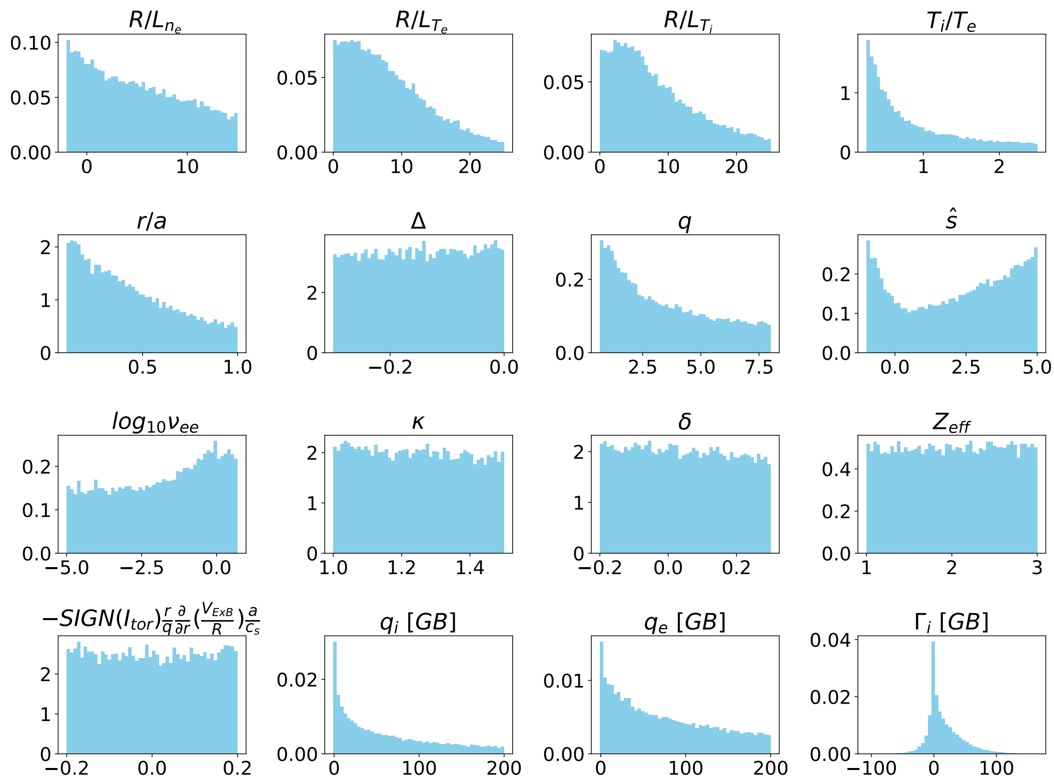

FIG. 1. The distribution of parameters in the training space of TGLFNN-UKAEA-MULTIMACHINEHYPER. Although the original space is a Latin Hypercube, the distributions shown here are not uniform as it would be expected from Latin Hypercube Sampling, and are a result of limiting the output fluxes to less than an absolute value of 200 GB.

Symbol	Parameter	STEP-V1/2	MULTIMACHINEHYPER
R/L_{n_e}	Density gradient	✓	✓
R/L_{T_e}	Electron temperature gradient	✓	✓
R/L_{T_i}	Ion temperature gradient	✓	✓
q	Safety factor	✓	✓
$\hat{s}$	Magnetic shear	✓	✓
ν^*	Collisionality	✓	✓
β	Plasma to magnetic pressure ratio	✓	×
Z_{eff}	Effective ion charge	✓	✓
T_e/T_i	Electron to ion temperature ratio	✓	✓
r/a	Flux surface centroid minor radius	✓	✓
Δ	Shafranov shift	✓	✓
δ	Plasma triangularity	✓	✓
$rd\delta/dr$	Shear in triangularity	✓	×
κ	Plasma elongation	✓	✓
$r/\kappa d\kappa/dr$	Shear in elongation	✓	×
$VEXB\_SHEAR$	ExB shear	×	✓

TABLE 1. Dimensions spanned by TGLFNN-STEP-V1/2 and TGLFNN-UKAEA-MULTIMACHINEHYPER.

Parameter	STEP-V1	STEP-V2	MULTIMACHINEHYPER
Learning rate	1.e-4	1.e-4	2.e-4
Weight decay	1.e-4	1.e-4	1.e-4
Dropout rate	0.05	0.05	0.05
Hidden size	512	512	512
Number of layers	6	6	6

TABLE 2. The NN settings from the coarse grid search that minimised the Guassian Negative Log Likelihood loss.

for the purpose of generalisation to unseen simulation trajectories. As flexibility in safety factor and shear were deemed necessary to allow for MHD stability optimisation in downstream applications, correlation between these dimensions and the remainder of the space were completely removed. 1M inputs per distribution were sampled and the respective TGLF outputs were generated. The saturation rule SAT1 was chosen as it provides a good match to MAST-U [19]; more saturation rules will be available in future work to allow capturing our epistemic uncertainty on transport in STEP. The linear solver settings of TGLF were calibrated against gradient-driven linear CGYRO [32] simulations at several times in an example ramp up case.

3. SURROGATE MODELS

3.1. Neural Network architecture and JINTRAC integration

The NN architecture adopted in this work is the same as in [25]. Briefly, the surrogates are Deep Ensembles [33] which offer a simple and scalable way of performing uncertainty quantification, a necessary feature of any surrogate model used in production. Deep Ensembles model an equaly weighted Gaussian Mixture where each ensemble member is a NN the outputs of which are the mean and standard deviation of a Gaussian distribution. The NNs are trained by minimising the gaussian negative log-likelihood as the loss function. In this work, 5 NNs per ensemble are chosen. A grid search is performed to optimise the network depth, hidden layer size, learning rate and L2 regularization for each member of the ensemble.

The TGLF fluxes considered are the ion heat flux, q_i, the electron heat flux, q_e and the ion particle flux, Γ_i which, due to ambipolarity is equal to the electron particle flux Γ_e and opposite in sign. Each TGLF flux is fitted independently by each NN. The fluxes are capped at 200 GB.

The neural networks were incorporated into the JINTRAC integrated modelling suite. Manual coding of individual networks and dependence on a specific architecture were avoided. Instead, the models were traced using TorchScript, which allows their execution as lightweight, self-contained modules independent of Python and architecture-specific definitions, thereby enabling faster, plug-and-play deployment in the integrated environment. The traced TorchScript models are called in Fortran via FTorch [34], and are accessible with a simple drop-in replacement for TGLF.

3.2. Surrogate model validation

Figure 2a shows the performance of the TGLFNN-UKAEA-MULTIMACHINEHYPER on an independent test set not observed during training, demonstrating production-ready performance.

Figures 2b and 2c shows the performance of the surrogate models for the spaces obtained from the JETTO simulations with Bohm/gyro-Bohm and TGLF transport respectively on independent test sets not observed during training. Of particular note is that the space obtained from the Bohm/gyro-Bohm transport modelling assumption produces higher fluxes than from the space obtained from using TGLF within JETTO (i.e. via more costly simulations). This is likely due to underprediction of transport in the Bohm/gyro-Bohm -based JETTO runs, and therefore higher gradients. Both the TGLFNN-STEP-V1 and TGLFNN-STEP-V2 surrogate models reproduce TGLF fluxes from a test set unseen during training, but it is expected that they will perform differently in flux-driven simulations due to the different training space spanned. Although the general trend is very good, the predictions in the low flux region tend to be more noisy. Investigation in this issue is ongoing.

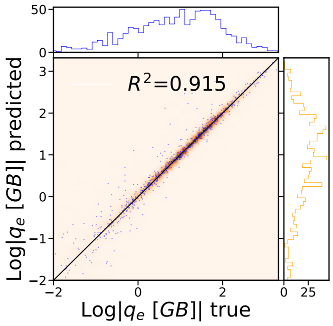

(a)

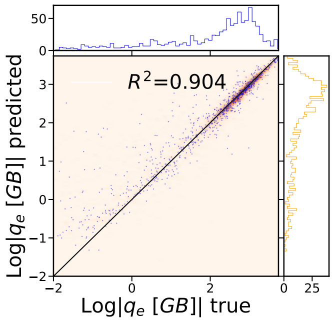

(b)

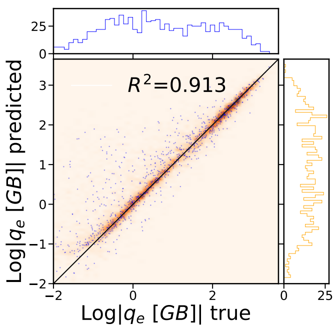

(c)

FIG. 2. (a): Comparison between the true TGLF predictions in the hypercube space applicable to several EuroFusion machines. (b): Comparison between the true TGLF predictions in the space defined by the JETTO simulations equipped with Bohm/gyro-Bohm transport and TGLFNN-STEP-V1. (c):Comparison between the true TGLF predictions in the space defined by the JETTO simulations where TGLF was used and TGLFNN-STEP-V2

Quantity	RRMS
T_e	1.8 %
T_i	3.5%
n_e	3.7%

TABLE 3. Relative Root Mean Squared Error for the profiles predicted by surrogate-enabled JINTRAC modelling of the JET discharge #89723 compared to JINTRAC with TGLF.

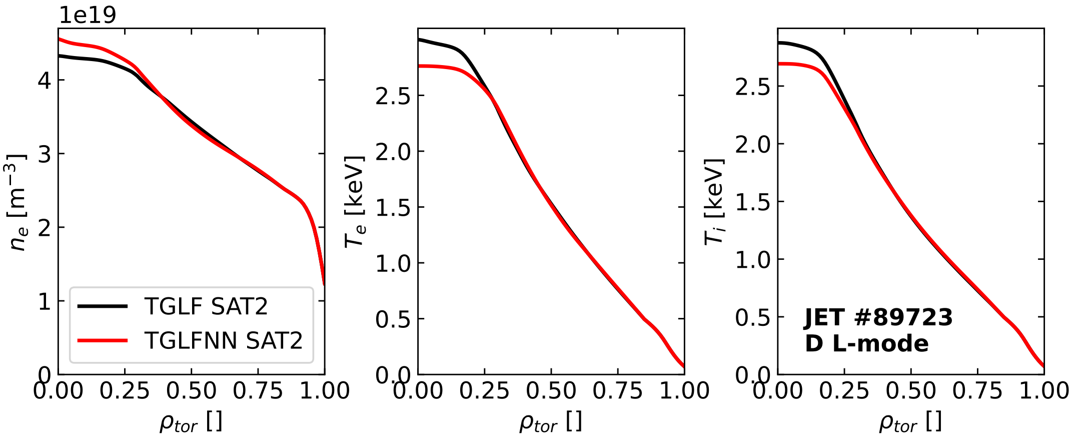

FIG. 3. Performance of the TGLFNN-UKAEA-MULTIMACHINEHYPER surrogate models in a flux-driven integrated JINTRAC run for the JET discharge #89723. Experimental data is not shown, as the main objective is to assess the quality of the surrogate model compared to runs adopting TGLF.

4. DEMONSTRATION IN INTEGRATED RUNS

4.1. Multimachine hypercube applied to JET

The TGLFNN-UKAEA-MULTIMACHINEHYPER surrogate models have been demonstrated on the JET #89723 D L-mode discharge.

Figure 3 shows the results of this exercise, and Table 3 report the RRMS error for these plots, computed as

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

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 TGLF respectively.

On the selected discharge, the surrogate-enabled JINTRAC runs achieve a sub 5% accuracy, with a runtime reduced from $\sim$4.5 hours to a mere $\sim$2 minutes. It is observed that most of the error comes from the region near the magnetic axis, likely due to the exclusion of that region in the training data. This will be addressed in a future version of the surrogate models. Nevertheless, the overall achieved accuracy indicates strong potential for success in a large validation exercise.

4.2. STEP

The TGLFNN-STEP-V1 and TGLFNN-STEP-V2 have been applied to a test STEP ramp up case [4] where TGLF was used. A comparison between TGLF and both the surrogate models is provided in Figure 4.

The runs using the TGLFNN-STEP-V2 surrogate models are generally well-behaved, although there is a loss of accuracy towards the end of the simulation, likely due to this simulation extending beyond the time provided in the training data. As expected, the TGLFNN-STEP-V1 surrogates agree less well with the ground truth compared to TGLFNN-STEP-V2, with errors compounding and driving the simulation towards a very different solution after $\sim$1500 seconds. In particular, a poorer confinement and lower temperatures are predicted, as well as a smaller bootstrap current and density. These results highlight the challenges of using a low-fidelity physics model to define the training space of a higher fidelity model.

The adoption of the surrogate models within the transport simulation reduce the runtime from 2 weeks on 60 cores to only 4 hours on 5 cores. The achieved accuracy and speed of TGLFNN-STEP-V2 are promising in view of the need for high-throughput ramp up optimisation studies in STEP.

5. CONCLUSIONS

This work presented NN surrogate models in the I_p ramp up for STEP (TGLFNN-STEP-V1 and TGLFNN-STEP-V2) as well as in a space applicable to several devices with aspect ratio greater than 3 (TGLFNN-UKAEA-MULTIMACHINEHYPER). The performance and speed achieved will allow to perform large scale validation

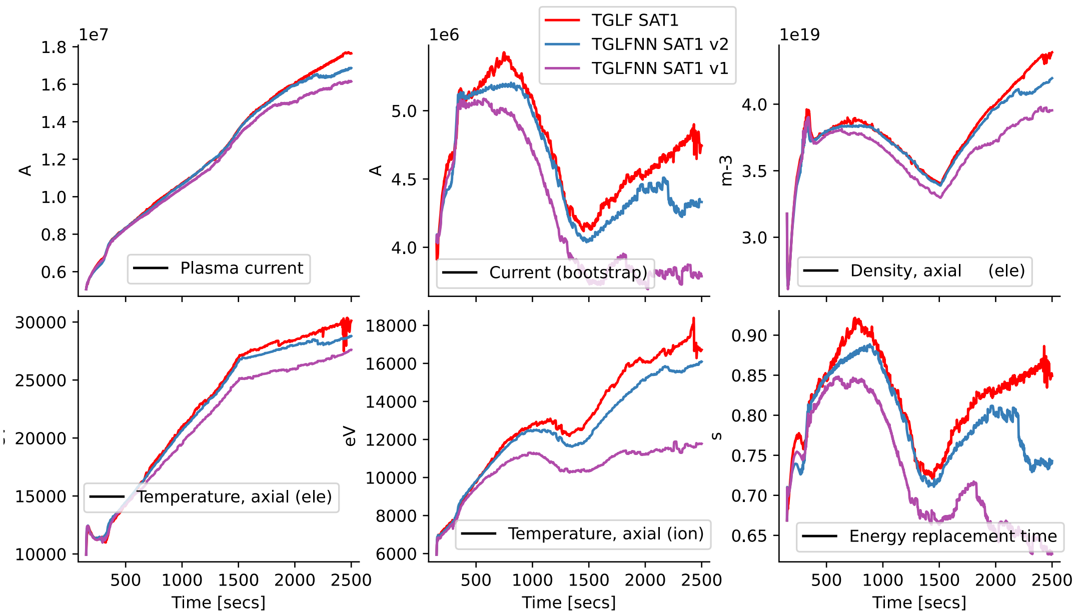

FIG. 4. STEP ramp up case obtained using TGLF (in red), TGLFNN-STEP-V1 (in purple) and TGLFNN-STEP-V2 (in blue).

of TGLF in a broad parameter space, particularly in L-mode, full-radius applications, and it will enable high-throughput optimisation for the STEP ramp up. Future work includes (i) expanding the scope of TGLFNN-UKAEA-MULTIMACHINEHYPER to increase the dimensionality of the input space which will allow multi-ion transport and momentum transport predictions; (ii) to extend the STEP TGLFNN-STEP-V2 space to higher β.

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.

REFERENCES

REFERENCES

[1] C.F. Maggi et al. "Overview of T and D–T results in JET with ITER-like wall". In: *Nuclear Fusion* 64.11 (Aug. 2024), p. 112012. ISSN: 1741-4326. DOI: 10.1088/1741-4326/ad3e16. URL: http://dx.doi.org/10.1088/1741-4326/ad3e16.

[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] G.V. Pereverzev and P.N. Yushmanov. *ASTRA - Automated System for TRansport Analysis*. 2002.

[4] F. Eriksson. "Current Ramp-up Modelling for STEP". In: *Proceedings of the 48th EPS Conference on Plasma Physics*. 2022.

[5] 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). ISSN: 1089-7674. DOI: 10.1063/1.2436852. URL: http://dx.doi.org/10.1063/1.2436852.

[6] 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). ISSN: 1089-7674. DOI: 10.1063/1.3505308. URL: http://dx.doi.org/10.1063/1.3505308.

[7] Theodore Brown et al. "Multi-Objective Bayesian Optimization for Design of Pareto-Optimal Current Drive Profiles in STEP". In: *IEEE Transactions on Plasma Science* 52.9 (Sept. 2024), pp. 3904–3909. ISSN: 1939-9375. DOI: 10.1109/tps.2024.3382775. URL: http://dx.doi.org/10.1109/TPS.2024.3382775.

[8] J. Mitchell et al. "Scenario trajectory optimization and control on STEP". In: *Fusion Engineering and Design* 192 (July 2023), p. 113777. ISSN: 0920-3796. DOI: 10.1016/j.fusengdes.2023.113777. URL: http://dx.doi.org/10.1016/j.fusengdes.2023.113777.

[9] G.M.D. Hogeweij et al. "Optimizing the current ramp-up phase for the hybrid ITER scenario". In: *Nuclear Fusion* 53.1 (Dec. 2012), p. 013008. ISSN: 1741-4326. DOI: 10.1088/0029-5515/53/1/013008. URL: http://dx.doi.org/10.1088/0029-5515/53/1/013008.

[10] 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://dx.doi.org/10.1088/1361-6587/aa8aeb.

[11] 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.

[12] 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.

[13] A. Ho et al. "Neural network surrogate of QuaLiKiz using JET experimental data to populate training space". In: *Physics of Plasmas* 28 (Mar. 2021), p. 032305. DOI: 10.1063/5.0038290.

[14] 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.

[15] J. Citrin et al. "Fast transport simulations with higher-fidelity surrogate models for ITER". In: *Physics of Plasmas* 30.6 (June 2023). DOI: 10.1063/5.0136752. URL: https://doi.org/10.1063%2F5.0136752.

[16] O. Meneghini et al. "Self-consistent core-pedestal transport simulations with neural network accelerated models". In: *Nuclear Fusion* 57 (June 2017). DOI: 10.1088/1741-4326/aa7776.

[17] 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.

[18] Victor Azizi et al. *Duqtools: Dynamic uncertainty quantification for Tokamak reactor simulations modelling*. 2024. DOI: 10.48550/ARXIV.2409.13529. URL: https://arxiv.org/abs/2409.13529.

[19] Tom F. Neiser et al. "Large database validation of TGLF on DIII-D and MAST-U plasmas". In: *Proceedings of the 66th Annual Meeting of the APS Division of Plasma Physics*. Monday–Friday, October 7–11, 2024. American Physical Society. Atlanta, Georgia, USA, Oct. 2024.

[20] Tom Neiser et al. "Database generation for validation of TGLF and retraining of neural network accelerated TGLF-NN". In: *APS Division of Plasma Physics Meeting Abstracts*. Vol. 2022. APS Meeting Abstracts. Jan. 2022, GP11.010, GP11.010.

[21] N. Bonanomi et al. "Time-dependent full-radius integrated modeling of the DTT tokamak main plasma scenarios". In: *Nuclear Fusion* 65.1 (Nov. 2024), p. 016005. ISSN: 1741-4326. DOI: 10.1088/1741-4326/ad8edb. URL: http://dx.doi.org/10.1088/1741-4326/ad8edb.

[22] F. J. Casson et al. "A Non-Inductive Ramp-Up Scenario for STEP". In: *Proceedings of the 22nd International Spherical Torus Workshop (ISTW 2024)*. 21–24 October 2024. Oxford, UK, Oct. 2024.

[23] Richard Bellman. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957. ISBN: 978-0-691-07951-6.

[24] 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.

[25] L. Zanisi et al. *Efficient training sets for surrogate models of tokamak turbulence with Active Deep Ensembles*. 2023. arXiv: 2310.09024 [physics.plasm-ph]. URL: https://arxiv.org/abs/2310.09024.

[26] A. Ho et al. "Efficient dataset construction using active learning and uncertainty-aware neural networks for plasma turbulent transport surrogate models". In: *Physics of Plasmas* 32.9 (Sept. 2025). ISSN: 1089-7674. DOI: 10.1063/5.0276455. URL: http://dx.doi.org/10.1063/5.0276455.

[27] Michael Barnes et al. *GS2 v8.1.2*. 2022. DOI: 10.5281/ZENODO.6882296. URL: https://zenodo.org/record/6882296.

[28] W. A Hornsby et al. "Gaussian process regression models for the properties of micro-tearing modes in spherical tokamaks". In: *Physics of Plasmas* 31.1 (Jan. 2024). ISSN: 1089-7674. DOI: 10.1063/5.0174478. URL: http://dx.doi.org/10.1063/5.0174478.

[29] J Citrin et al. "Electromagnetic stabilization of tokamak microturbulence in a high-regime". In: *Plasma Physics and Controlled Fusion* 57.1 (Nov. 2014), p. 014032. ISSN: 1361-6587. DOI: 10.1088/0741-3335/57/1/014032. URL: http://dx.doi.org/10.1088/0741-3335/57/1/014032.

[30] N. Bonanomi et al. "Light impurity transport in JET ILW L-mode plasmas". In: *Nuclear Fusion* 58.3 (Jan. 2018), p. 036009. ISSN: 1741-4326. DOI: 10.1088/1741-4326/aaa4d3. URL: http://dx.doi.org/10.1088/1741-4326/aaa4d3.

[31] Bhavin S. Patel et al. "Pyrokinetics - A Python library to standardise gyrokinetic analysis". In: *Journal of Open Source Software* 9.95 (2024), p. 5866. DOI: 10.21105/joss.05866. URL: https://doi.org/10.21105/joss.05866.

[32] J. Candy, E.A. Belli, and R.V. Bravenec. "A high-accuracy Eulerian gyrokinetic solver for collisional plasmas". In: *Journal of Computational Physics* 324 (Nov. 2016), pp. 73–93. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2016.07.039. URL: http://dx.doi.org/10.1016/j.jcp.2016.07.039.

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

[34] Jack Atkinson et al. "FTorch: a library for coupling PyTorch models to Fortran". In: *Journal of Open Source Software* 10.107 (Mar. 2025), p. 7602. ISSN: 2475-9066. DOI: 10.21105/joss.07602. URL: http://dx.doi.org/10.21105/joss.07602.