AI-AUGMENTED SCENARIO DESIGN AND CLASSICAL CONTROL OF TOKAMAK PLASMAS

[1]A. AGNELLO, [2]N. AMORISCO, [1]A. GARROD, [1]A. ROSS, [1]G.K. HOLT, [1]A. KEATS,
[1]P. CAVESTANY-OLIVARES ,[1]M. MARSHALL,
[2]K. PENTLAND, [2]G. MCARDLE, [2]C. VINCENT, [2] J. BUCHANAN, [2]S. PAMELA

[1] STFC Hartree Centre, Sci-Tech Daresbury, Warrington (United Kingdom)
[2] United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon (UK)

Email: adriano.agnello@stfc.ac.uk

1. VIRTUAL CIRCUITS FOR REAL-TIME CONTROL

In real-time control of a tokamak experiment, the plasma is shaped by magnetic fields from surrounding poloidal field (PF) coils, which must also incorporate the stray field from Ohmic coils that drive a plasma current [1]. Feed-forward coil sweeps are pre-computed, and control loops are designed to correct for any departures of the plasma control targets, on 0.1ms timescales. This is achieved through Jacobian *shape matrices* of derivatives of the control targets *vs* PF coil currents, and in particular through their pseudo-inverse *Virtual Circuits (VCs)* [2]. These quantify how the requested PF coil currents should differ to correct the measured observables towards the desired targets. VCs are pre-computed with finite-difference derivatives of shape targets *vs* PF currents, for a small number of reference stages in a shot. Different configurations require different VCs, so the accuracy of control is limited by the availability of VCs closest to the ones for a configuration observed in real time. With a large library of Grad-Shafranov (GS) equilibria, one can train fast and accurate emulators to the values of control targets as functions of the PF coil currents and plasma parameters (plasma current, toroidal current density profiles) [3]. To solve the forward GS problem, we use the purpose-built FreeGSNKE codebase [4,5,6]. FreeGSNKE has robust solvers of the GS equation that are robust against known nonlinear instabilities from Picard iterations. Being written in pure Python, it is seamlessly interfaceable with other libraries for parameter space exploration and machine learning.

2. GENERATING A LARGE TRAINING SET, TRAINING AND TUNING EMULATORS

As detailed by [7], entirely synthetic libraries can be built: starting from some reasonable parameter choices, a MCMC random walk sequentially samples the input parameters, privileging equilibrium configurations with desirable properties, avoiding parameter combinations resulting in problematic profiles or numerical issues in the integration of the GS equation. The random walk enables faster convergence of the GS solver, and the step sizes are adaptively adjusted. For work tailored on the MAST-U machine description [8], we start from snapshots of experiments as initial seeds for each MCMC chain. We obtain 5 million synthetic equilibria, including configurations encountered in MAST-U campaigns so far (including limited, diverted and super-X), and extend to yet-unseen configurations. To make the MCMC sampling even more efficient, we have extended the sequential sampler [9] to an asynchronous and scalable multi-walker algorithm, with surrogates of the score function to encourage steps towards promising parameter choices [10]. This approach works well for any problem that requires the traversal of a non-trivial simulation space. We train fully-connected, feed-forward neural networks (FNNs) to reproduce shape-control targets *vs* coil currents, plasma current, internal inductance and poloidal beta. We use early-stopping and learning-rate shrinkage by monitoring the FNN loss on a 20% validation set. We develop a Thompson sampler and annealing algorithm to optimise the architecture and learning hyperparameters [9,11]. Optimal models have 5-10 layers with 100-200 nodes per layer. The overall residual scatter, relative to the intrinsic scatter in the data, is below 1% and the relative bias is 0.03%.

3. VIRTUAL CIRCUIT VALIDATION, SHAPING CURRENTS

As a general testbed for VCs, we examine how equilibria change under PF current shifts from VCs computed via finite differences with FreeGSNKE. In general, linearisation and VCs perform relatively well on many equilibria, but some targets are harder to decouple (e.g. divertor gap and strike-point), as they would require higher-order changes around the lower X-point. The Jacobians of the FNNs, with respect to input currents, are also a good approximation of the finite-difference Jacobians. This is not a trivial result, as the universal approximation property is proved in the L^p norm but not on derivatives. The robustness in Jacobians likely stems from the very limited overfitting of the FNNs. Encouraged by these findings, FNN Jacobians are being implemented for real-time control on MAST-U.

In real-time control, incomplete information is available on currents in all metals, including passive structures, that can affect the shape of the equilibrium and hence the choice of appropriate VCs. At design stage, those are typically incorporated via an inverse GS solution to infer *shaping currents* that the active coils should have, with zero currents on passive structures, to produce the same equilibrium configuration as from a real shot. We have developed lean models to infer shaping currents in real time throughout a shot, based on the histories of measured coil and plasma currents, without the need to perform real-time GS equilibrium reconstruction.

4. REINFORCEMENT LEARNING

Reinforcement-learning (RL) policies naturally account for transients in vessel currents, which are neglected in the construction of (pre-computed or emulated) VCs [12,13,14]. Exploiting the stateful representation of FreeGSNKE, we train RL agents to control a set of shape targets, developing the FreeGSNKE-RL package [15]. An initial exploration shows that the best results are achieved when the RL agents must only learn corrections on top of pre-calculated feed-forward "coil sweeps". Previous result in the literature have trained RL agents to learn the entire control policy, however the wide span in timescales (0.1ms for control over a 0.5s shot) hinders the efficient learning of zeroth-order sweeps [12,16]. Our preliminary results suggest that RL agents may improve with feed-forward information.

ACKNOWLEDGEMENTS

This work was funded under the *Fusion Computing Lab* collaboration between the STFC Hartree Centre and the UK Atomic Energy Authority.

REFERENCES

[1] Ariola, M., "Integrated Plasma Shape and Boundary Flux Control on JET Tokamak", *Fusion Science and Technology*, vol. 53, no. 3, pp. 789–805, 2008. doi:10.13182/FST08-A1735

[2] McArdle, G., Pangione, L., and Kochan, M., "The MAST Upgrade plasma control system", *Fusion Engineering and Design*, vol. 159, Art. no. 111764, 2020. doi:10.1016/j.fusengdes.2020.111764.

[3] Wai, J. T., Boyer, M. D., and Kolemen, E., "Neural net modeling of equilibria in NSTX-U", *Nuclear Fusion*, vol. 62, no. 8, Art. no. 086042, IOP, 2022. doi:10.1088/1741-4326/ac77e6

[4] AMORISCO, N.C., AGNELLO, A., HOLT, G.K. ET AL., FreeGSNKE: A Python-based dynamic free-boundary toroidal plasma equilibrium solver, *Physics of Plasmas*, Volume 31, Issue 4, id.042517, 2024, DOI:10.1063/5.0188467

[5] AMORISCO, N.C., gitlab.stfc.ac.uk/farscape-ws3/freegsnke

[6] Pentland, K., et al., "Validation of the static forward Grad-Shafranov equilibrium solvers in FreeGSNKE and Fiesta using EFIT++ reconstructions from MAST-U", arXiv:2407.12432, 2024. doi:10.48550/arXiv.2407.12432.

[7] AGNELLO, A., AMORISCO, N.C., KEATS, A., ET AL., Emulation techniques for scenario and classical control design of tokamak plasmas, *Physics of Plasmas*, Volume 31, Issue 4, id.043901, 2024, DOI: 10.1063/5.0187822

[8] Fishpool, G., "MAST-upgrade divertor facility and assessing performance of long-legged divertors", *Journal of Nuclear Materials*, vol. 438, pp. S356–S359, 2013. doi:10.1016/j.jnucmat.2013.01.067

[9] AGNELLO, A., github.com/farscape-project/freegs-libgen

[10] AGNELLO, A., GARROD, A., MARSHALL, M. gitlab.stfc.ac.uk/optimising-expensive-optimisation/HiveMind

[11] AGNELLO, A., GARROD, A., github.com/farscape-project/TRANSP-emu

[12] Degrave, J., "Magnetic control of tokamak plasmas through deep reinforcement learning", *Nature*, vol. 602, no. 7897, pp. 414–419, 2022. doi:10.1038/s41586-021-04301-9

[13] Seo, J., "Feedforward beta control in the KSTAR tokamak by deep reinforcement learning", *Nuclear Fusion*, vol. 61, no. 10, Art. no. 106010, IOP, 2021. doi:10.1088/1741-4326/ac121b

[14] Zhang, Y. C., Wang, S., Yuan, Q. P., Xiao, B. J., and Huang, Y., "Real-time feedback control of β_{p} based on deep reinforcement learning on EAST", *Plasma Physics and Controlled Fusion*, vol. 66, no. 5, Art. no. 055014, IOP, 2024. doi:10.1088/1361-6587/ad3749

[15] HOLT, G.K. AGNELLO, A., NGUYEN. N., ET AL., gitlab.stfc.ac.uk/farscape-ws3/freegsnkerl

[16] Tracey, B. D., "Towards practical reinforcement learning for tokamak magnetic control", *Fusion Engineering and Design*, vol. 200, Art. no. 114161, 2024. doi:10.1016/j.fusengdes.2024.114161