
Progress towards high fidelity collisional-radiative model surrogates for rapid in-situ evaluation

Nathan A. Garland
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87544
ngarland@lanl.gov

Romit Maulik
Argonne Leadership Computing Facility
Argonne National Laboratory
Lemont, IL 60439

Qi Tang
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87544

Xian-Zhu Tang
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87544

Prasanna Balaprakash
Argonne Leadership Computing Facility and
Mathematics and Computer Science Division
Argonne National Laboratory
Lemont, IL 60439

Abstract

The advancement of future fusion plasma transport modeling platforms relies on having an accurate understanding of the composition and radiation properties of the plasma. These quantities can be obtained from solving a collisional-radiative (CR) model at each time step or iteration of a transport code. However, even compact, approximate CR models can be computationally onerous to evaluate and in-situ evaluation of these models within a larger plasma transport code leads to a rigid bottleneck. As a way to bypass this bottleneck, we propose deploying artificial neural network surrogates to allow rapid evaluation of the necessary plasma quantities. In this work we explore the computational benefits of the proposed method, and outline a data-driven adaptive sampling routine to allow autonomous refined sampling of the parameter space in order to ensure a broad and meaningful set of training data is assembled to train the network.

1 Introduction

Research and construction towards ITER is being performed across the world. ITER will be the world's largest tokamak, which is a device to magnetically confine a hot plasma in the shape of a torus with the goal of demonstrating a path to viable fusion energy. Presently, the biggest threat to ITER's safe operation are plasma disruptions. A disruption is a sudden termination of the plasma, and can lead to extreme heat deposition and relativistic electron beam impact on reactor walls. In either case, the end result is predicted to be unacceptable hardware damage, costing exorbitant time and financial resources. Presently, impurity injection is the reference ITER mitigation strategy, with neon being a likely candidate for injection into the hydrogenic plasma [1].

Accurate modeling of fusion plasmas with added impurities is an indispensable resource to understanding the physics and mitigation plan of tokamak disruptions [2]. Given some impurity of atomic

number, Z , and density, n_I , added to a plasma of a fixed deuterium density, n_D , and described by an electron temperature, T_e , the associated collision and radiative events described by atomic physics sets the populations of the neutral atom and Z possible ion stages with density n_j , where $0 \leq j \leq Z$. Further, knowing the radiation potential of the plasma due to the population of various excited states within each ion stage is crucial understanding the ability of a hot plasma to radiatively cool.

To access the desired quantities, n_j , and radiation potential, P_{Rad} , a collisional-radiative (CR) model is used to solve for the populations of ions in various excited energy states. A CR model is a coupled set of rate ordinary differential equations (ODEs), solved as a rate matrix problem, that describe transitions in to, or out of, ionic excited state populations, $n_{j,\alpha}$, where the index j denotes ion stage, and α denotes the excited state level. In the steady-state, which is a reasonable assumption for certain fusion plasma regimes, the rate matrix problem is $R(\mathbf{n})\mathbf{n} = 0$, where R is the rate matrix composed of elements R_{ij} describing gain and loss from element i to j , and \mathbf{n} is the state vector of excited states of ion populations, $n_{j,\alpha}$. Solution of this system can be computationally challenging when \mathbf{n} is large or R is near singular. Once \mathbf{n} is known, the sum of radiation cooling rates due to each excited ion state can then be computed by the CR model, and then passed back to the plasma transport model [3, 4, 5].

Ideally, one would solve a CR model at each time-step or iteration of a plasma transport code, however the computational bottleneck created by this requirement is undesirable. Even relatively compact, $\mathcal{O}(10^1 - 10^2)$, approximate CR models, such as that used in this present work [6, 7], can take up to dozens of seconds to solve. While state-of-the-art, fine structure $\mathcal{O}(10^6 - 10^9)$ CR models can take minutes or hours to obtain one solution [8, 9]. Presently, in order to obtain these quantities within plasma transport codes the fusion community has resorted to greatly simplified models that limit the fidelity of physics retained in the model [10]. In this study we seek an alternative approach to allow rapid evaluation of necessary quantities in the form of a generalizable artificial neural network (ANN) surrogate. In doing so, we hope to push forward the limits of fusion plasma modeling, aided by machine learning, in order to facilitate greater understanding of the critical tokamak disruption mitigation issue for ITER.

2 Surrogate modeling strategy

In the following section, we evaluate the construction of surrogate models for our CR dataset so that dramatically inexpensive forward models may be obtained using an ANN. Notably, our goal is to not just obtain a cheap surrogate but to also balance the offline cost of representative data set generation for training this surrogate. This is obtained via an adaptive sampling explained below. We use two surrogate models in this research: a low fidelity (LF) and high fidelity (HF) surrogate. The two models differ in the ability to accurately represent the CR input-output relationship and are used for two distinct purposes in our surrogate development campaign. These models are given by a low-fidelity random forest regressor (RFR) [11], and a high fidelity ANN that may act as a universal approximator [12] trained by backpropagation [13].

Firstly, the surrogate we seek to learn will provide a map of the input field of: electron temperature, impurity density, and deuterium density, T_e, n_I, n_D , to an output field of the total radiative power loss and ion stage populations for deuterium and the added impurity, $P_{Rad}, n_0^D, n_{+1}^D, n_0^I, n_{+1}^I, \dots, n_{+(Z-1)}^I, n_{+Z}^I$ where Z is the atomic number of the impurity of concern. In this study we assume $Z = 10$ for injected neon impurity, which is the current strategy outlined for ITER. Expansion to variable Z will be explored in future studies, but as the atomic physics of each atom, even those adjacent to neon on the periodic table, can be significantly different we limit ourselves to one atomic target initially. The input field quantities are limited to the domain of $1\text{eV} \leq T_e \leq 1000\text{ eV}$, $10^{13}\text{ cm}^{-3} \leq n_I \leq 10^{15}\text{ cm}^{-3}$, and $10^{13}\text{ cm}^{-3} \leq n_D \leq 10^{15}\text{ cm}^{-3}$.

The underlying mechanism of our adaptive sampling strategy uses the ensemble property of the RFR wherein each tree in the forest may be queried for a prediction of a scalar metric, to estimate prediction variances in the multidimensional parameter space. If the different trees in an RFR show a wide range of predictions (i.e., there is widespread disagreement in the predictions of trees in a forest) in some region of the finely sampled space, we run full-order CR evaluations in that region, append this to our master training data set and retrain the RFR as well as the multi-layer perceptron (MLP). This allows us to adaptively sample higher dimensional spaces which may be challenging to cover due to the curse of dimensionality. To initialize this adaptive algorithm, a first (relatively coarse) sampling of the input space is constructed to train our RFR following which the adaptive sampling is performed

iteratively. The end of an adaptive sampling iteration is accompanied by a re-training of the MLP surrogate on the expanded training data set. The validation performance of this MLP is recorded for the purpose of a posteriori model selection. Note that the sampling at each iteration (and at the start of the adaptive training data augmentation) utilizes the Latin Hypercube sampling (LHS) method for experimental design. We note that a choice of scalar metric for the LF surrogate must be made in order to guide resampling. In this research we employ two physically important quantities in this role, one being the average charge state of the impurity ions $\bar{Z} = \sum_{j=0}^Z j n_j / \sum_{j=0}^Z n_j$ where Z is the atomic number of the impurity, and n_j are ion populations of a charge j . The second scalar is the total radiative power loss, P_{rad} , which is a function computed by the forward CR model dependent on the range of excited state populations.

The hyperparameters that need to be selected a-priori for this framework include the number of samples required for the first training iteration, the number of new samples for each data set augmentation, the total computational budget for training the surrogate and the traditional hyperparameters of the MLP (e.g., the learning rate, early stopping criterion, architecture) and the RFR (the number of trees, maximum depth of the forest). Our sampling strategy is outlined in Algorithm 1 where the hyperparameters of the associated components are as follows: $N_{init} = 50$ initial samples, $N_{resample} = 0.1N_{init}$ re-samples per data set augmentation iteration, and a budget of $N_{budget} = 1950$ samples to make 2000 samples the maximum possible. The MLP was chosen to have two hidden layers, 50 neurons wide each, between the input and output layers of dimension three and 14 respectively. The ReLU activation function was employed in this network. The learning rate was fixed as 10^{-3} with the Adam algorithm [14] used for optimization. An early stopping criteria was set to halt learning after 100 epoch of failing to reduce validation loss. Hyperparameters of the RFR were chosen as 50 trees in the forest, with a maximum depth of eight. Stopping criteria for the overall adaptive sampling framework was set to be when a validation R^2 result of 0.95 or greater is obtained.

Algorithm 1: Adaptive CR surrogate training sampling

```

Generate  $N_{init}$  initial samples of forward CR model evaluation;
Train HF surrogate;
while  $N_{samples} < N_{budget}$  do
    Train LF surrogate to map input field to scalar metric;
    Evaluate fine sample space grid using LF surrogate. Sort by variance;
    for  $i = 1 : N_{samples}/10$  do
        | Take input field of highest variance LF prediction. Generate new sample;
    end
    Train HF surrogate on expanded sample set. Evaluate surrogate metrics.;
    if  $R^2 > R_{goal}^2$  then
        | Stop.
    end
end

```

3 Experiments

As a point of reference for the performance of the adaptive sampling framework, an ANN surrogate was simply trained on a fixed data set of size 3375, spanning a space of $15 \times 15 \times 15$ samples determined via LHS. The hyperparameters outlined in the previous Section were employed in training this surrogate, and a test R^2 of 0.999 was obtained against an unseen test set. To demonstrate the ability of this ANN, from this unseen test set, a kernel density estimate plot showing the predicted relationship between two output field quantities is shown in Figure 2. All assessments in this study were conducted with Tensorflow 2.0.0 under Python 3.7.9 on a 2.9GHz Intel Core i9 CPU with 32GB of RAM.

3.1 Adaptive sampling performance

Employing the network architectures and hyperparameter choices outlined in the previous section, we performed analyses to examine the ability of the adaptive sampling framework to produce an ANN

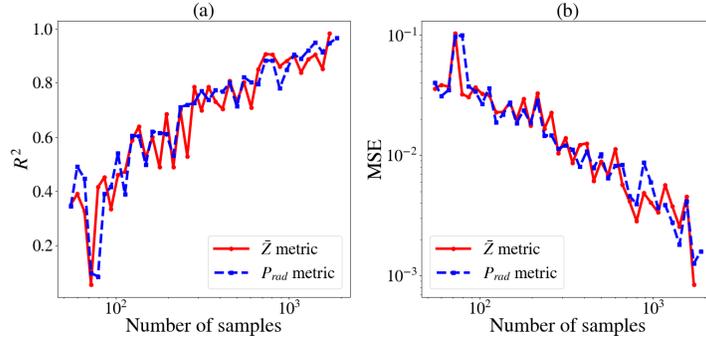


Figure 1: Validation metrics for (a) R^2 and (b) mean-squared error (MSE) of adaptively trained MLP surrogate when trained on a given master training data set size. Clear, but relatively minor, differences are seen between using either \bar{Z} or P_{rad} as the decision metric for the RFR low fidelity surrogate.

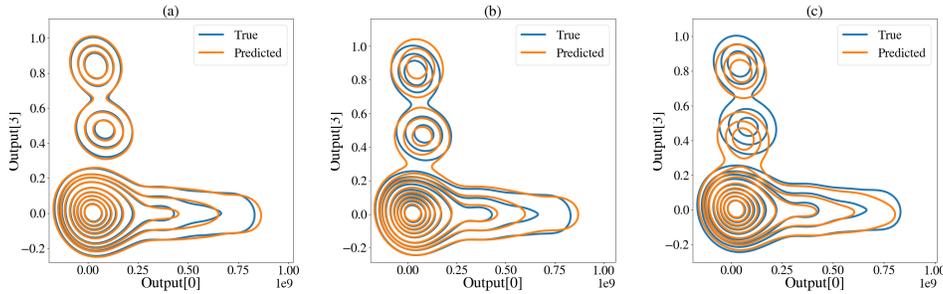


Figure 2: Kernel density estimate plots demonstrating the true and predicted relationship between output field quantities of total radiation P_{rad} and the population density of the Ne^{+10} ion for (a) MFP conventionally trained on fixed set of 3375 samples, (b) MFP adaptively trained on 1719 samples by using \bar{Z} as a guiding metric, and (c) MFP adaptively trained on 1890 samples by using P_{rad} as a guiding metric.

surrogate trained on largely autonomously gathered training data. Employing the two scalar metrics of \bar{Z} and P_{rad} in RFR decision making, we found both cases produced a steady gain in validation R^2 performance and drop in validation MSE as more samples were acquired, as shown in Figure 1. Ultimately the adaptively trained MFP surrogates passed the R^2 threshold of 0.95 and final surrogate models were obtained with a validation R^2 of 0.983 and MSE of 8.422×10^{-4} for the \bar{Z} metric case, after 1719 samples were acquired. Similarly the case of the P_{rad} metric obtained a validation R^2 of 0.965 and MSE of 1.587×10^{-3} after 1890 samples were acquired.

Testing the adaptively trained MLP surrogates on an unseen data set, compared to that from training or validation, we present kernel density estimate (KDE) plots in Figure 2 showing the relationship between the output field quantities of total radiation P_{rad} and the population density of the Ne^{+10} ion as a way to demonstrate the performance of predicted quantities versus the true values. Here we see, compared to the ANN trained on a larger training set, the \bar{Z} guided surrogate performs reasonably well, with the exception of some translation of the prediction over the two upper islands in the plot. Similarly, the surrogate guided by the P_{rad} metric does not accurately reproduce the upper two islands of the KDE. We note that testing R^2 values of 0.982 and 0.967 were obtained for the adaptively trained networks informed by \bar{Z} and P_{rad} respectively. Given the primary motivator of constructing a framework for surrogate training was for rapid in situ evaluation, a timing comparison for 1000 forward pass evaluations of the ANN surrogate and original CR model FORTRAN executable was performed to highlight the compute time savings. The mean surrogate execution time was approximately 4.2 ms, while the mean CR model forward pass execution was 2.7 s.

4 Conclusion and Future Work

In this work we have proposed employing strategically trained, fit-for-purpose surrogate models for evaluation of quantities normally accessed by collisional-radiative modeling in a thermal fusion plasma. We have demonstrated the vast savings in execution time of such a surrogate, which would reduce computation bottlenecks in flagship plasma simulation codes compared to if a conventional CR model solve was performed. We have demonstrated the capability of an adaptive sampling framework to automatically generate training data in statistically meaningful areas of our parameter space. This training technique then allows the possibility of (i) faster acquisition of a sufficiently trained ANN surrogate, and (ii) assurance that data assembled for training is representative and meaningful towards training the surrogate - rather than simply being random. While the present CR model employed in this study to develop the training method is comparatively fast compared to more accurate and expensive models, and generation of many thousands of random training data points for a conventionally trained ANN is not too burdensome, we note that great savings will be found when this method is applied to more accurate, but time consuming, forward models - where a single CR model solve can take on the order of an hour. It is thus intended in future to utilize the adaptive sampling training method towards creation of ANN based surrogates for accurate CR models that would allow rapid evaluation of accurate ion and radiation properties for multiple atom types within a fusion plasma following the injection of an impurity species, such as neon or argon. Future work avenues will explore transfer learning of the ANN (rather than retraining), and parallelized sampling, for example employing DeepHyper [15].

Broader Impact

The ITER project will explore a range of tokamak plasma operating conditions, with the hope of eventually producing a burning, self-sustaining plasma to demonstrate nuclear fusion. The biggest threat to the ability for ITER to safely operate is that of plasma disruptions, which can produce extreme heat being deposited to the reactor walls, or the generation of relativistic electron beams which can impact the reactor walls. Either scenario could be catastrophic to the containment vessel surrounding the plasma. To mitigate this disruption scenario, the current ITER strategy is to inject impurity particles into the plasma. To plan for this scenario, and to maximize the chances of success of this strategy, large-scale plasma physics models require information on the ion population in the plasma, as well as the radiation produced by the excited states of these ions. In order to retain a high fidelity of physics, while allowing rapid evaluation of these quantities, we propose employing data-driven surrogates to allow forward pass evaluation on the order of milliseconds, instead of seconds or longer, depending on the sophistication level of the original physics-based forward model. To ensure efficient training from the forward physics model, we propose the accumulation of a statistically meaningful training data set in an adaptive manner. This is achieved by focusing on greater data generation in regions that display greater uncertainty from ensemble predictions of low fidelity models. Therefore, the final data set represents the parameter space of interest in an efficient manner. In doing so, we hope to promote a step-forward in plasma modeling capabilities by allowing sophisticated atomic and molecular physics to be included in-situ with flagship fusion plasma simulation codes.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, through the SciDAC-RAPIDS Institute under Argonne National Laboratory contract DE-AC02-06CH11357. This work was also jointly supported by the U.S. Department of Energy through the Magnetic Fusion Theory Program and the Tokamak Disruption Simulation (TDS) SciDAC project at Los Alamos National Laboratory (Contract No. 89233218CNA000001). Additional support is provided by Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20200356ER. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. DOE or the United States Government.

References

- [1] M Lehnen and S Maruyama. Executive report of the Disruption Mitigation Workshop. Technical Report ITR-18-002, ITER Organization, St-Paul-lez-Durance, March 2018.
- [2] M. Lehnen, K. Aleynikova, P. B. Aleynikov, D. J. Campbell, P. Drewelow, N. W. Eidietis, Yu. Gasparyan, R. S. Granetz, Y. Gribov, N. Hartmann, E. M. Hollmann, V. A. Izzo, S. Jachmich, S. H. Kim, M. Kočan, H. R. Koslowski, D. Kovalenko, U. Kruezi, A. Loarte, S. Maruyama, G. F. Matthews, P. B. Parks, G. Pautasso, R. A. Pitts, C. Reux, V. Riccardo, R. Roccella, J. A. Snipes, A. J. Thornton, and P. C. de Vries. Disruptions in ITER and strategies for their control and mitigation. *Journal of Nuclear Materials*, 463:39–48, August 2015.
- [3] E. M. Hollmann, I. Bykov, N. W. Eidietis, J. L. Herfindal, A. Lvovskiy, R. A. Moyer, P. B. Parks, C. Paz-Soldan, A. Yu. Pigarov, D. L. Rudakov, D. Shiraki, and J. Watkins. Study of argon expulsion from the post-disruption runaway electron plateau following low-Z massive gas injection in DIII-D. *Physics of Plasmas*, 27(4):042515, April 2020.
- [4] L. Hesslow, O. Embréus, G. J. Wilkie, G. Papp, and T. Fülöp. Effect of partially ionized impurities and radiation on the effective critical electric field for runaway generation. *Plasma Physics and Controlled Fusion*, 60(7):074010, June 2018.
- [5] Christopher J. McDevitt, Zehua Guo, and Xian-Zhu Tang. Avalanche mechanism for runaway electron amplification in a tokamak plasma. *Plasma Physics and Controlled Fusion*, 61(5):054008, April 2019.
- [6] H. K. Chung, M. H. Chen, W. L. Morgan, Y. Ralchenko, and R. W. Lee. FLYCHK: Generalized population kinetics and spectral model for rapid spectroscopic analysis for all elements. *High Energy Density Physics*, 1(1):3–12, December 2005.
- [7] Nathan A. Garland, Hyun-Kyung Chung, Christopher J. Fontes, Mark C. Zammit, James Colgan, Todd Elder, Christopher J. McDevitt, Timothy M. Wildey, and Xian-Zhu Tang. Impact of a minority relativistic electron tail interacting with a thermal plasma containing high-atomic-number impurities. *Physics of Plasmas*, 27(4):040702, April 2020.
- [8] Y. Ralchenko. *Modern Methods in Collisional-Radiative Modeling of Plasmas*. Springer Series on Atomic, Optical, and Plasma Physics. Springer International Publishing, 2016.
- [9] C. J. Fontes, H. L. Zhang, J. Abdallah Jr, R. E. H. Clark, D. P. Kilcrease, J. Colgan, R. T. Cunningham, P. Hakel, N. H. Magee, and M. E. Sherrill. The Los Alamos suite of relativistic atomic physics codes. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 48(14):144014, 2015.
- [10] D. G. Whyte, T.E. Evans, and A.G.E. Mellman. In *Proceedings of the 24th European Conference on Controlled Fusion and Plasma Physics, Berchtesgaden*, volume 21A, page 1137. European Physical Society, Petit-Lancy, 1997.
- [11] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [12] Kurt Hornik, Maxwell Stinchcombe, Halbert White, et al. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [13] Robert Hecht-Nielsen. Theory of the backpropagation neural network. In *Neural networks for perception*, pages 65–93. Elsevier, 1992.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2017.
- [15] Romit Maulik, Romain Egele, Bethany Lusch, and Prasanna Balaprakash. Recurrent neural network architecture search for geophysical emulation. *arXiv preprint arXiv:2004.10928*, 2020.