Uncertainty based Online Ensemble on Non-Stationary Data for Fusion Science

Kishansingh Rajput[1, 2], Malachi Schram[1], Brian Sammuli[3], Sen Lin[2]
[1]Thomas Jefferson National Accelerator Facility, Newport News, VA, USA 23606
[2]Department of Computer Science, University of Houston, Houston, TX, USA 77024
[3]General Atomics, PO Box 85608, San Diego, California 92186-5608 USA

Abstract

Machine Learning (ML) is poised to play a pivotal role in the development and operation of next-generation fusion devices. Fusion data shows non-stationary behavior with distribution drifts in the data, resulted by both experimental evolution and machine wear-and-tear. ML models assume stationary distribution and fail to maintain performance when encountered with non-stationary data streams. Online learning can be used to continuously adapt the models with new data as it is acquired. However, traditional online learning can suffer from short-term performance degradation, as ground truth is not available before making the prediction. To address this challenge, we propose an uncertainty aware ensemble approach for online learning, where a Deep Gaussian Process Approximation (DGPA) technique is leveraged for calibrated uncertainty estimation and the uncertainty values are then used to guide a meta-algorithm that produces predictions based on ensemble of learners. Moreover, DGPA also provides uncertainty estimation along with the predictions for decision makers. This paper demonstrates that the proposed method outperforms traditional online learning approach, and a naive ensemble without uncertainty guidance by about 7% and 6%, respectively, on B-coil deflection prediction at DIII-D Fusion Facility.

1 Introduction

The pursuit of controlled nuclear fusion as a viable energy source has led to the development of increasingly complex fusion facilities, such as tokamaks [1], which demand precise control, real-time monitoring, and predictive capabilities to optimize performance and ensure operational safety. Traditional solutions in fusion research have relied heavily on physics-based modeling, often requiring high-fidelity simulations that are computationally expensive and not easily adaptable in real time. As such, data driven especially Machine Learning (ML) based solutions are poised to play a pivotal role in the development and operation of next-generation fusion devices. ML has been emerging as transformative tool for fusion research and operations [2]. By integrating sensor data with advanced computational models, ML enables modeling complex processes, supporting tasks such as diagnostics, forecasting, anomaly detection, and decision support.

One particularly compelling application is the modeling and monitoring of the toroidal field (TF) coils in tokamaks. These coils are responsible for producing the strong magnetic fields required for plasma confinement, and are subject to extreme electromagnetic forces during high-field operation. At the DIII-D tokamak [3], calibration experiments have revealed that under certain full-field configurations, the outer legs of the TF coils (B-coils) experience lateral deflections large enough to cause mechanical interference and transient electrical shorts [4]. These deflections are sensitive to subtle mechanical slip between unbonded coil turns, thermal expansion, and the evolving electromagnetic loading profile—all of which are difficult to model analytically in real time. Traditional finite element

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

simulations have helped establish upper bounds on allowable lateral loads, but these models are not well suited for fast inference, adaptive prediction, or incorporation of real-time data. This opens an opportunity for ML methods—particularly those emphasizing uncertainty quantification, and online learning to augment or replace traditional models with fast, adaptive ML models that support control-relevant decision-making.

Fusion data demonstrates non-stationary behavior due to shot-to-shot drifts arising from both experimental evolution and aging of the equipment. On the other hand, traditional ML assumes stationary distribution and directly deploying an ML model trained offline in fusion facilities would not work well under online non-stationary data streams. To address this, online learning [5] has evolved as a promising technique to continuously adapt ML models with new data. However, one of the major challenges in using supervised models in online learning framework is that the ground truth information is only available after making the prediction. This limits the benefit of online learning in highly non-stationary settings such as fusion experiments.

In this paper, we propose an uncertainty-aware ensemble-based meta-algorithm to improve online learning on non-stationary data streams and provide a reliable uncertainty estimation with predictions. We use the data collected at DIII-D fusion facility at General Atomics to predict TF coils deflection based on plasma input parameters. Moreover, we demonstrate that the proposed uncertainty-aware meta-algorithm outperforms standard online learning approach and naive ensemble-based online learning by about 7% and 6%, respectively.

2 Methods

For modeling, we employ Deep Neural Networks (DNNs) composed of convolutional, maxpool, dense, and dropout layers. The model is trained to produce expected coil deflection based on plasma input parameters in a supervised manner with Mean Absolute Error (MAE) loss and Adam optimizer with lerning rate of 10^{-4}. We leverage Deep Gaussian Process Approximation (DGPA) [6] that has been previously shown to provide reliable out-of-distribution (OOD) uncertainty quantification in particle accelerator applications [7]. The reliability of uncertainty quantification in DGPA comes from the assumption that it follows the similar behavior as traditional Gaussian Process (GP) which is designed to provide distance based uncertainties, farther a data sample from training distribution, higher is the uncertainty prediction. This is one of the reason, we chose DGPA as uncertainty quantification method in our use case. It leads to two fold benefits: ***a)*** provide confidence boundaries along with the predictions for the decision makers, and ***b)*** use it as an estimation of error on future data (before ground truth becomes available) to guide the meta-algorithm in online learning.

2.1 Online Learning

Online learning frameworks update ML models sequentially as new data become available, rather than retraining on the entire dataset, thereby enabling dynamic adaptation to evolving data distributions. This paradigm is particularly well-suited for our application, where the data stream exhibits non-stationary behavior. To accommodate the sequential nature of the data, we adopt a sliding-window batched training strategy in which the model is fine-tuned from its previously learned weights using the most recent sample together with a small buffer of recent historical data. For abrupt drifts, shorter buffer size is favored, whereas for slow and gradual drifts longer buffer size is more appropriate. In fusion experiments, the drifts can attain different behaviors depending on maintenance schedules, experimental evolution and equipment wear-and-tear. As such, a single fixed buffer size for the online batched learning is not ideal. In this study we use an ensemble of learners with different buffer sizes belonging to different time scales at DIII-D facility, starting with 1 shot (immediate scale), 5 shots (roughly an hour), 20 shots (about half a day), 40 shots (about a day worth of data), and 200 shots (a week full of data).

2.2 Deep Gaussian Process Approximation (DGPA)

Gaussian process (GP) [8] models provide principled uncertainty quantification by leveraging kernel functions to assess similarity between samples, making them intrinsically distance-aware and effective at detecting OOD data. However, traditional GPs scale poorly with large datasets due to their high computational complexity, limiting their applicability to high-dimensional problems. To overcome these limitations, the DGPA model has been proposed [9], which integrates the expressive capacity of

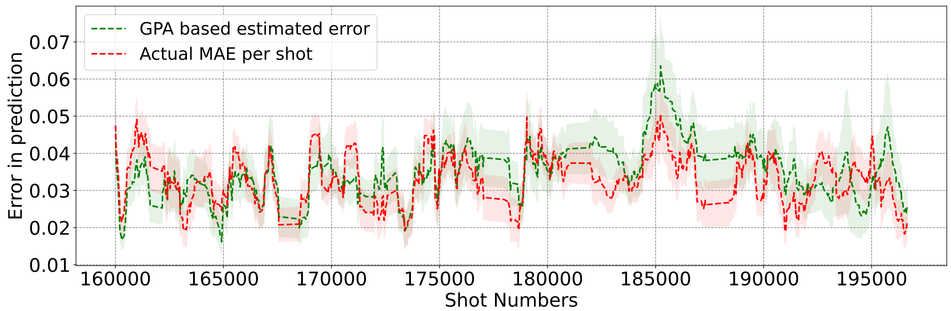

Figure 1: A comparison of average error estimated based on uncertainty quantification and actual average error per shot after ground truth information becomes available. The bands show statistical variation with 10 trials. The error estimations are highly correlated with the actual MAE errors, captured by Pearson correlation coefficient of 0.59.

DNNs with a fixed-size GP approximation. Specifically, it uses a Radial Basis Function (RBF) kernel approximation using Random Fourier Features (RFFs) at the output layer of the neural network. This hybrid structure enables the DGPA to combine the scalability of DNNs with the uncertainty-aware characteristics of GPs, resulting in reliable in-distribution predictions and improved OOD uncertainty quantification. Moreover, calibrated uncertainties from DGPA models can also provide an estimate of error on predictions, where the predictions with lower uncertainties are expected to be more accurate than those with higher uncertainty. We leverage an uncertainty toolbox [10] along with callbacks in our models to calibrate uncertainty estimation after every training session. Figure 1 demonstrates a strong correlation between average model error and estimated average error per shot based on uncertainty estimation from the DGPA model.

2.3 DGPA guided Online Learning

To adapt to various drift types, we employ an ensemble of ML models trained on different buffer sizes of historical data in an incremental manner. However, relying solely on naive ensembles can lead to misleading results due to the lack of knowledge about expected error from individual models at the time of combining predictions. To address this challenge, we propose leveraging uncertainty quantification from DGPA models as an estimate of the expected error per prediction. This enables guidance for meta-algorithms to generate more accurate ensemble predictions. Our proposed meta-algorithm is outlined in Algorithm 1, which demonstrates how DGPA-based uncertainty quantification can be leveraged to perform an informed weighted average of predictions from ensemble.

Algorithm 1 Uncertainty-Aware Online Ensemble Learning for Non-Stationary Data Streams

Input : Ensemble size n; a base model (θ) trained on data up to time T, initial buffers for each model $\mathcal{B} = \{\mathcal{B}_1, \mathcal{B}_2 \ldots, \mathcal{B}_n\}$, learning rate η
Output : Ensemble predictions $\hat{y}^{(t)}$ over time

Initialization: Create an ensemble of models $\mathcal{E} = \{\theta_1, \theta_2 \ldots, \theta_n\}$, initialized as $\theta_i = \theta$
for *each time step* $t = T+1, T+2, \ldots$ **do**
  Obtain prediction and uncertainty estimation $\hat{y}_i^t, \sigma_i^{(t)} \leftarrow \mathrm{f}_{\theta_i}(x^{(t)})\ i = 1, \ldots, n$
  Compute ensemble prediction: $\hat{y}^{(t)} = \sum_{i=1}^{n} w_i^{(t)} \hat{y}_i^{(t)}$ with $w_i \propto \frac{1}{\sigma_i^2}$, and $\sum_{i=1}^{n} w_i^{(t)} = 1$
  Compute combined standard deviation $\sigma_{\bar{x}} = \left(\sum_{i=1}^{n} \sigma_i^{-2}\right)^{-\frac{1}{2}}$
  Receive new ground truth $\mathbf{y}^t$ from the experiment and update each rolling buffer $\mathcal{B}_i$ with $\mathbf{y}^t$;
  for *each model* θ_i *and respective buffer* $\mathcal{B}_i$ **do**
    Compute loss: $L = \frac{1}{|\mathcal{B}_i|} \sum_{(x^{(j)}, y^{(j)}) \in \mathcal{B}_i} \ell(f_{\theta_i}(x^{(j)}), y^{(j)})$
    Backward pass: compute gradients $\nabla_{\theta_i} L$
    Update parameters: $\theta_i \leftarrow \theta_i - \eta \nabla_{\theta_i} L$

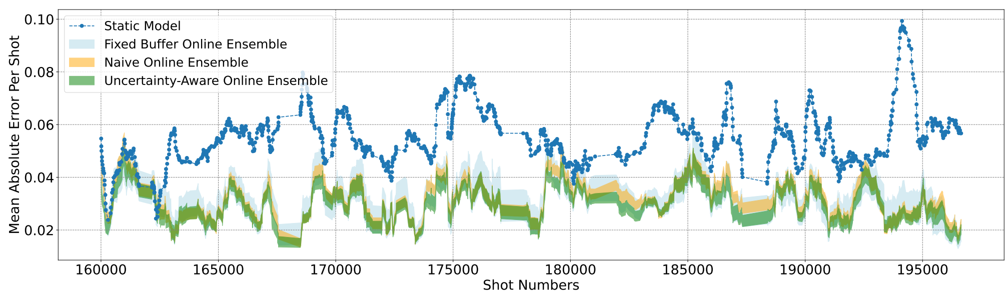

Figure 2: Comparison of Mean Absolute Error (MAE) per shot among different methods with 10 trials to produce statistically robust analysis. The figure shows moving average over 20 points to clearly show the trend. Naive ensemble with different buffer size models is slightly more stable (narrow bands) than fixed buffer size ensemble but mean performance is similar between the two (with overall MAE of 0.0323 and 0.0320). On the other hand, the proposed uncertainty aware online ensemble approach (with overall MAE of 0.0300) outperforms fixed buffer size model, and naive ensemble by 7% and 6% respectively.

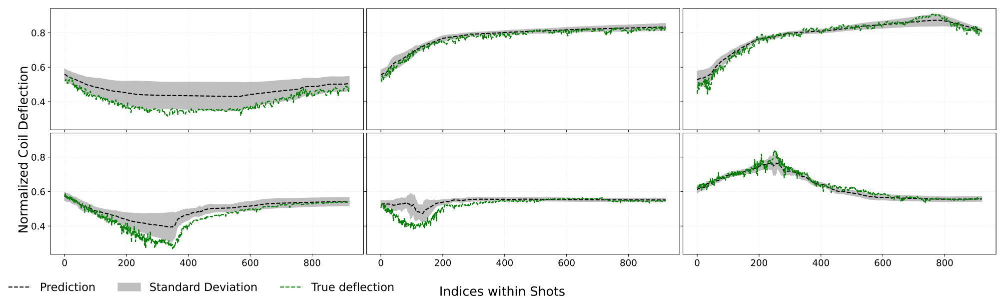

Figure 3: Visualization of uncertainty estimation quality from uncertainty aware online ensemble method. It preserves the quality of the original DGPA uncertainty estimation in the combined predictions. These examples show that the uncertainty estimation is highly correlated with error in predictions.

3 Results

In our experiments, we compare four different approaches as follows: ***a)*** Static model that is trained on historical data before shot number 160000 (taken in 2014) and is not updated on any new data; ***b)*** Naive Online learning with a single model (we choose buffer size of 20 shots for this comparison as it appears outperform other buffer sizes in single model approach); ***c)*** Naive online ensemble that uses average of predictions from n different models incrementally trained with different buffer sizes, approximately representing daily, weekly, monthly, and yearly timescales; ***d)*** Uncertainty aware online ensemble that uses uncertainty guided weighted average of the predictions from n models that are same as *c)* above.

To provide a statistically robust comparison, we perform 10 trials of each method. Figure 2 shows the comparison between these methods. The bands in the figure represents statistical variations over multiple trials. It demonstrates that the uncertainty aware ensemble method with Mean Absolute Error (MAE) of 0.0300 outperforms naive ensemble (MAE 0.0320), and a single model approach (MAE 0.0323) by about 6% and 7% respectively. Figure 3 shows examples where the error from meta algorithm is slightly high, however, the combined uncertainty estimation preserves the original property of the DGPA model and it provides a very accurate estimation of the expected error.

4 Conclusion

We have presented a novel online ensemble method based on reliable uncertainty estimation that outperforms naive ensembles, and standard online learning by 6% and 7%, respectively. We leverage DGPA based uncertainty estimation that has been shown to provide reliable OOD uncertainties along with predictions to guide our meta algorithm that combines the predictions from ensemble of models that use rolling buffers with different sizes for incremental training data. Our method with different buffer sizes for ensemble is particularly suitable for non-stationary online learning scenarios where drift is expected at different time scales. However, it is important to note that the algorithm can be generalized to other applications that may benefit from different types of learners. In future, we would like to further improve UQ based error estimation and extend our approach to include other types of models in the ensemble including different architectures, and potentially use attention mechanism with different scales to address drifts. In addition, we would like to deploy this approach at DIII-D facility to evaluate it's performance in real time.

5 Acknowledgement

Work supported by Jefferson Lab and General Atomics internal research and development. This manuscript has been authored by Jefferson Science Associates (JSA) operating the Thomas Jefferson National Accelerator Facility for the U.S.Department of Energy under Contract No. DE-AC05-06OR23177. Additionally, this material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-FC02-04ER54698, along with Office of Fusion Energy Sciences Award No. DE-SC0024426 (Fusion Data Platform). The US government retains, and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes.

References

[1] C. Gliss, C. Bachmann, S. Ciattaglia, B. Drumm, M. Gracia Camacho, I. Moscato, T. Mull, and I. Palermo. Integrated design of tokamak building concepts including ex-vessel maintenance. *Fusion Engineering and Design*, 177:113068, 2022.

[2] D. P. Schissel, R. M. Nazikian, and T. Gibbs. Summary report from the mini-conference on digital twins for fusion research. *Physics of Plasmas*, 32(5):050601, 05 2025.

[3] M. E. Fenstermacher, J. Abbate, S. Abe, T. Abrams, M. Adams, B. Adamson, N. Aiba, T. Akiyama, P. Aleynikov, E. Allen, et al. Diii-d research advancing the physics basis for optimizing the tokamak approach to fusion energy. *Nuclear Fusion*, 62(4), 04 2022.

[4] E.E. Reis and E. Chin. Structural response of the diii-d toroidal field coil to increased lateral loads. In *20th IEEE/NPSS Symposium onFusion Engineering, 2003.*, pages 577–580, 2003.

[5] Steven C.H. Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021.

[6] Malachi Schram, Kishansingh Rajput, Karthik Somayaji NS, Peng Li, Jason St. John, and Himanshu Sharma. Uncertainty aware machine-learning-based surrogate models for particle accelerators: Study at the fermilab booster accelerator complex. *Phys. Rev. Accel. Beams*, 26:044602, Apr 2023.

[7] Kishansingh Rajput, Malachi Schram, and Karthik Somayaji. Uncertainty aware deep learning for particle accelerators, 2023.

[8] Christopher Williams and Carl Rasmussen. Gaussian processes for regression. In D. Touretzky, M.C. Mozer, and M. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1995.

[9] Malachi Schram, Kishansingh Rajput, Karthik Somayaji NS, Peng Li, Jason St. John, and Himanshu Sharma. Uncertainty aware machine-learning-based surrogate models for particle accelerators: Study at the fermilab booster accelerator complex. *Phys. Rev. Accel. Beams*, 26:044602, Apr 2023.

[10] Youngseog Chung, Ian Char, Han Guo, Jeff Schneider, and Willie Neiswanger. Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification, 2021.