Meta-Learned Bayesian Optimization for Energy Yield in Inertial Confinement Fusion

Vineet Gundecha[1†], Ricardo Luna Gutierrez[1†], Sahand Ghorbanpour[1†], Rahman Ejaz[2†], Varchas Gopalaswamy[2†], Riccardo Betti[2†], Desik Rengarajan[1], Soumyendu Sarkar[1†*]

[1]Hewlett Packard Enterprise, [2]University of Rochester

{vineet.gundecha, rluna, sahand.ghorbanpour, desik.rengarajan
soumyendu.sarkar}@hpe.com
{reja, vgopalas, betti}@lle.rochester.edu

Abstract

With the growing demand for clean energy, fusion presents a promising path to sustainable power generation. Inertial confinement fusion (ICF) experiments trigger nuclear reactions by firing lasers at a fuel target, typically composed of deuterium and tritium. These experiments are costly and require complex optimization of the laser pulse shape across multiple shots to maximize energy yield. Even though Bayesian Optimization (BO) has been commonly used to optimize such expensive scientific experiments, vanilla BO methods do not leverage prior knowledge of the function from simulations or past experiments and fail to achieve high sample efficiency. In this work, we adapted and explored BO meta-learning techniques for ICF that either meta-learn the BO surrogate model, the acquisition function, or both from simulations. Our results demonstrate that the three meta-learning techniques we investigated, Meta-Learning Acquisition Functions for BO (MetaBO), Rank-Weighted Gaussian Process Ensemble (RGPE), and Neural Acquisition Processes (NAP), drastically reduce the number of experiments needed to achieve a satisfactory yield in ICF simulations.

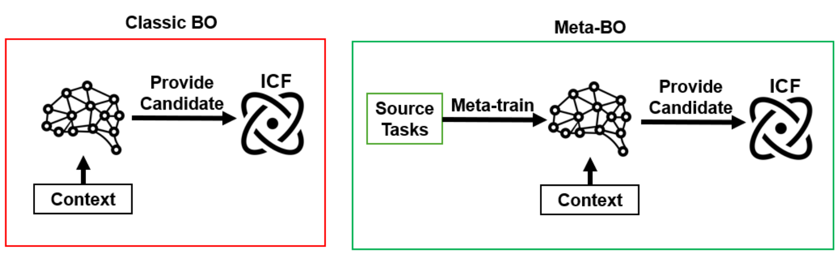

Figure 1: Difference between Meta-BO and Classic BO approaches. Meta-BO methods leverage knowledge collected from previous optimizations or simulations to increase optimization effectiveness.

1 Introduction

The global energy demand is rapidly raising, the International Energy Agency (IEA) projects that this demand will double by 2050 (Eia, 2015). Currently, more than 80% of the world's energy supply comes from fossil fuels like petroleum, coal, and natural gas (Bahrami et al., 2019). Nuclear fusion

*Corresponding author. †These authors contributed equally.

holds the promise for limitless, clean energy. Inertial confinement fusion (ICF) initiates nuclear fusion by rapidly compressing fuel targets, typically composed of deuterium and tritium, using high-energy lasers. ICF experiments are complex and exceptionally costly because of the extreme conditions needed to achieve nuclear fusion. The high expenses and the limited number of ICF facilities lead to very few opportunities for experimentation, often just a few times a year. So the development and study of ICF optimization techniques that can utilize data in a sample-efficient manner is critical. To our knowledge, this work is the first to study the application of meta-learning BO techniques in ICF, augmenting BO from previous experiments or simulations. We specifically study three classes of BO meta-learning techniques:

- MetaBO (Volpp et al., 2020) which meta-learns the acquisition function using Reinforcement Learning (RL).
- Rank-Weighted Gaussian Process Ensemble (RGPE) (Feurer et al., 2022) which uses an ensemble of Gaussian Processes as the surrogate model.
- Neural Acquisition Process (NAP) (Maraval et al., 2023) which simultaneously meta-learns the surrogate model and the acquisition process in an end-to-end manner.

We show that meta-learning improves the performance of BO by increasing the effectiveness of the optimization process in ICF.

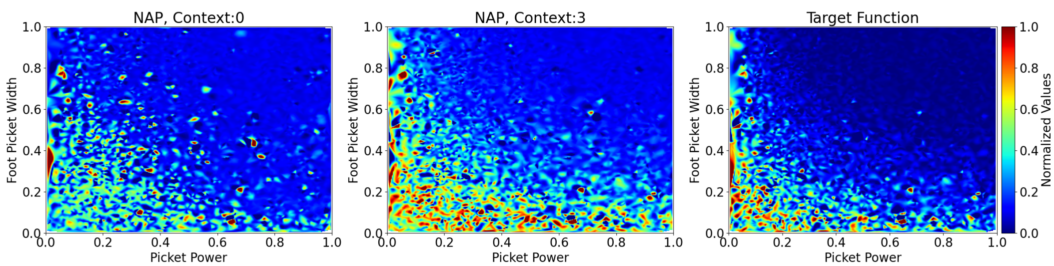

Figure 2: NAP's meta-learned surrogate predictions: (left) without any context points; (middle) after three context points; (right) the optimization target function. NAP achieves rapid adaptation with high sample efficiency, which is ideal for the limited ICF experiments possible on a shot day, given only a few shot days occur annually. For clarity, two of the five dimensions are plotted here.

2 Background

2.1 ICF

ICF is a technique used to achieve nuclear fusion by subjecting a tiny fuel pellet, usually made from a blend of deuterium and tritium (hydrogen isotopes), to extremely high temperatures and pressures. This is typically done using high-powered lasers. The aim is to create the right environment for the nuclei to overcome their natural repulsion and fuse, releasing substantial energy (Betti and Hurricane, 2016; Lees et al., 2021; Gopalaswamy et al., 2019, 2024). The energy output for a given payload is influenced by the laser pulse, which delivers a massive burst of energy within a short time (3 ns). The shape of this pulse is controlled by 5 parameters, which are vital for ICF optimization. An example of a pulse is shown in Figure 3. Other related optimization studies include (Gutierrez et al., 2024; Shmakov et al., 2023).

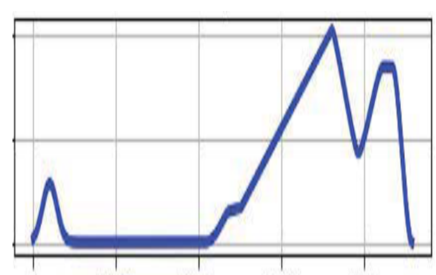

Figure 3: An example laser pulse shape used in ICF experiments. The shape is controlled by 5 parameters. The parameters are adjusted during an optimization campaign to maximize energy yield.

ICF experiments are extremely expensive. They rely on advanced, powerful laser systems due to the demanding conditions necessary for achieving nuclear fusion. The high costs of the experiments and the limited number of facilities equipped to perform ICF experiments lead to very infrequent experiment opportunities, often limited to ten to twenty shots annually (five to ten per experimental campaign). To accelerate ICF development, the exploration and study of sample-efficient techniques for ICF optimization is crucial.

2.2 Bayesian Optmization

In BO, we sequentially maximize an expensive to evaluate black-box function $f(x)$ for variables $x \in X$, where X is the input domain. BO techniques operate in two steps. First, based on previously collected evaluation data, we fit a probabilistic surrogate model to emulate $f(x)$ allowing us to make probabilistic predictions of the function's behavior on unobserved input points. Given the first step's probabilistic predictions, the second step in BO optimizes an AF that trades off exploration and exploitation to find a new input query, $x_{new} \in X$, which is chosen for evaluation.

The goal in BO is to find the global optimum x^* as:

$$x^* = \underset{x \in X}{\operatorname{argmax}}\, f(x) \tag{1}$$

2.3 Meta Bayesian Optimization

Meta-BO approaches aim to improve the optimization of new unseen target black-box functions by leveraging knowledge from a set of N related source tasks (functions) $\mathcal{F}$ (Maraval et al., 2023).

In Meta-BO, we assume that the model has access to this knowledge in the form of N datasets $\mathcal{D}_1, ..., \mathcal{D}_N$ collected from evaluations in the set of N source tasks. Each dataset $\mathcal{D}_n$ consist of e_n evaluations of $f_n(x) \in \mathcal{F}$ for all $n \in [1 : N]$, such as $\mathcal{D}_n = \{(x_n^i, y_n^i)\}_{i=1}^{e_n}$, where $y_n^i = f_n(x_n^i)$. These datasets are leveraged by the Meta-BO approaches to meta-learn a surrogate model, an AF, or both. In real-world applications such as ICF, the source task used for meta-learning can be built from simulations.

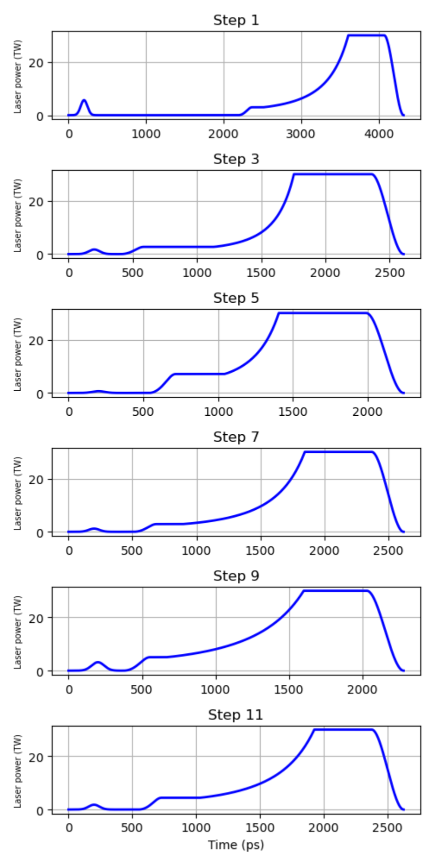

Figure 4: The progression of the pulse shape over the course of the optimization trajectory as recommended by NAP.

2.3.1 MetaBO

MetaBO (Volpp et al., 2020) uses RL to meta-learn an AF on source tasks that are drawn from the same function class as the target task. It replaces the AF with a neural network that is able to identify and exploit structural properties of a class of functions. Specifically it uses the Proximal Policy Optimization (PPO) method as proposed in (Schulman et al., 2017) for training. A major drawback of this technique is the use of a discrete grid when finding the maximum of the AF. This is done in order to save on computational cost during training. We replace this with a continuous optimization algorithm during the evaluation phase only.

2.3.2 RGPE

RGPE (Feurer et al., 2022) proposes a new method for Bayesian optimization by combining knowledge from past optimization runs. RGPE ranks past optimization tasks according to their similarity to the current task and creates a weighted combination of Gaussian process models, each trained on these ranked tasks. This approach allows selective transfer of knowledge, focusing on the most relevant past experiences and ultimately improving the accuracy and efficiency of the surrogate model.

2.3.3 Neural Acquisition Process (NAP)

While most Meta-BO techniques either meta-learn the surrogate or the AF, NAP provides state-of-the-art performance by jointly learning both in an end-to-end manner using Transformer Neural Processes (TNP) (Nguyen and Grover, 2023; Maraval et al., 2023), a class of models that combine the flexibility and performance of transformers with the properties of stochastic processes. NAP is trained using PPO in conjunction with an auxiliary supervised loss to build its AF.

3 Experiments

We evaluate the performance of different Meta-BO approaches, and compare against vanilla BO approaches, in the task of energy yield optimization for ICF. For dataset creation, we utilized the LOTUS library (Ejaz et al., 2024) to generate various laser pulse profiles based on a custom parametrization. These parameters determine the laser power and timing, which were then used as inputs for LILAC (Delettrez et al., 1987), a simulator for laser-driven fusion physics. The chosen laser pulse shape significantly affects the experimental results (energy yield), influencing both the compression of the fusion fuel and the development of hydrodynamic instabilities (Williams et al., 2021). To develop a response surface reflecting different entropy shapes, we varied 5 parameters related to the laser pulse. Using Latin hypercube sampling, we generated 50,000 samples within the design constraints of the laser system. These pulse shapes were tested with the LILAC simulator on a fixed fusion fuel target, and the resulting neutron yields were analyzed to construct a response surface based on these 5 parameters. Additionally, to create a diverse evaluation, we generated two source tasks and one test task by modifying the simulator's physics models, specifically, the equation of state, which affects shock behavior in fusion fuel, thus changing the response surfaces across different simulation versions. This divergence allows us to approximate the potential differences encountered in sim-to-real scenarios for ICF experiments.

3.1 Meta-learning Performance

We compare the Meta-BO approaches with the standard BO approach, which utilizes a GP surrogate model combined with classical acquisition functions like UCB and EI (Garnett, 2023). Additionally, we include Random Search as a baseline. The performance of the approaches is evaluated in terms of simple regret Volpp et al. (2020); Maraval et al. (2023).To emulate a real-world setting, we restrict the number of samples in our experiments to 15. We run all experiments using 5 different seeds. The results of this evaluation, shown in Figure 5, demonstrate the comparative performance of the methods. The meta-learning approaches such as RGPE and NAP demonstrate superior performance compared to the classic BO baselines, highlighting the benefits of meta-learning in the ICF domain. Notably, NAP exhibits exceptional performance by reaching optima with only a few samples, indicating its potential to accelerate progress in ICF development.

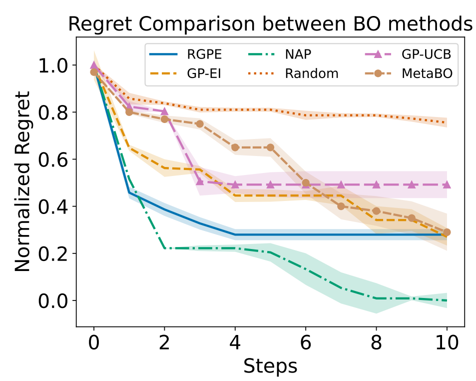

Figure 5: Comparison between classic BO methods and Meta-BO methods in ICF. NAP is able to achieve optimal performance in few-samples. The shaded areas represent a ± 1 standard error. Regrets are normalized.

3.2 NAP's Surrogate Predictions

To understand the superior performance achieved by NAP, we examine the predictions made by its surrogate model and assess how well it adapts. To conduct this evaluation, we compare the real target function (test task) with NAP's predictions in a 2D space. For NAP, we first evaluate its initial predictions with zero context, so the surrogate information only comes from the meta-learning pretraining (Context:0). Moreover, we assess how NAP adapts by querying it after three samples have been collected from the target function (Context:3). The evaluation results in Figure 2 show that NAP effectively uses information from source tasks to identify regions with a higher likelihood of finding optimal solutions. Furthermore, after incorporating just a few samples, NAP accurately adjusts its predictions to match the target function closely. Figure 4 shows the progression of the pulse shape over the course of an optimization trajectory, as recommended by NAP. NAP proposes larger adjustments at the early stages, while at later stages it does some minor finetunning.

3.3 Experimental Settings

For the NAP experiments, we used the default hyper-parameters provided in the original implementation (Maraval et al., 2023). Table 1 shows the hyper-parameters for PPO, and Table 2 shows those for the TNP.

Parameter	**Value**
Learning rate for gradient descent	$3 \cdot 10^{-5}$
Learning rate decay	Linear decay to 0 over 2000 iterations
Number of training PPO iterations	2000
Horizon of episodes used in training	24
Trajectories collected per iteration	60
Total numbers of transformer updates	90,000
Mini-batch size	32
Weight of auxiliary loss in total loss (λ)	1.0
Weight of the value function loss in total loss	1.0
Generalized Advantage Estimator-λ	0.98
Discount factor γ	0.98
Clip of importance sampling ratio ϵ	0.15
L2 gradient clipping	0.5

Table 1: Hyperparameters used for NAP's PPO

Parameter	**Value**
Number of buckets in the output histogram	1000
Point-wise feed-forward dimension of Transformer	1024
Embedding dimension of Transformer	512
Number of self-attention layers of Transformer	6
Number of self-attention heads of Transformer	4
Dropout rate of Transformer	0.0
Softmax temperature to compute π from α	0.1 for training, argmax otherwise
Value function network	Linear(2, 512), Tanh, Linear(512, 1)

Table 2: Architecture for the Transformer Neural Processes in NAP

4 Conclusion and Future Work

We show that Meta-BO techniques can significantly improve the effectiveness of ICF simulations. High sample efficiency can reduce the cost of experiments, accelerating the path to fusion energy and contributing to a more sustainable planet. We plan to extend these techniques to multi-fidelity Bayesian Optimization, incorporating data from higher-fidelity simulations and past experiments to further improve optimization.

5 Acknowledgement

This work was financially supported by the U.S. Department of Energy under Award No. DE-SC0024408, Hewlett Packard Enterprise, and the University of Rochester.

References

U. Eia, Annual energy outlook 2015 with projections to 2040, United States Energy Information Administration, Washington (2015).

A. Bahrami, A. Teimourian, C. O. Okoye, H. Shiri, Technical and economic analysis of wind energy potential in uzbekistan, Journal of cleaner production 223 (2019) 801–814.

M. Volpp, L. P. Fröhlich, K. Fischer, A. Doerr, S. Falkner, F. Hutter, C. Daniel, Meta-learning acquisition functions for transfer learning in bayesian optimization, 2020. URL: `https://arxiv.org/abs/1904.02642`. `arXiv:1904.02642`.

M. Feurer, B. Letham, F. Hutter, E. Bakshy, Practical transfer learning for bayesian optimization, 2022. URL: `https://arxiv.org/abs/1802.02219`. `arXiv:1802.02219`.

A. Maraval, M. Zimmer, A. Grosnit, H. B. Ammar, End-to-end meta-bayesian optimisation with transformer neural processes, 2023. `arXiv:2305.15930`.

R. Betti, O. A. Hurricane, Inertial-confinement fusion with lasers, Nature Physics 12 (2016) 435–448. URL: `https://doi.org/10.1038/nphys3736`. doi:`10.1038/nphys3736`.

A. Lees, R. Betti, J. P. Knauer, V. Gopalaswamy, D. Patel, K. M. Woo, K. S. Anderson, E. M. Campbell, D. Cao, J. Carroll-Nellenback, R. Epstein, C. Forrest, V. N. Goncharov, D. R. Harding, S. X. Hu, I. V. Igumenshchev, R. T. Janezic, O. M. Mannion, P. B. Radha, S. P. Regan, A. Shvydky, R. C. Shah, W. T. Shmayda, C. Stoeckl, W. Theobald, C. Thomas, Experimentally inferred fusion yield dependencies of omega inertial confinement fusion implosions, Phys. Rev. Lett. 127 (2021) 105001. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.127.105001`. doi:`10.1103/PhysRevLett.127.105001`.

V. Gopalaswamy, R. Betti, J. P. Knauer, N. Luciani, D. Patel, K. M. Woo, A. Bose, I. V. Igumenshchev, E. M. Campbell, K. S. Anderson, K. A. Bauer, M. J. Bonino, D. Cao, A. R. Christopherson, G. W. Collins, T. J. B. Collins, J. R. Davies, J. A. Delettrez, D. H. Edgell, R. Epstein, C. J. Forrest, D. H. Froula, V. Y. Glebov, V. N. Goncharov, D. R. Harding, S. X. Hu, D. W. Jacobs-Perkins, R. T. Janezic, J. H. Kelly, O. M. Mannion, A. Maximov, F. J. Marshall, D. T. Michel, S. Miller, S. F. B. Morse, J. Palastro, J. Peebles, P. B. Radha, S. P. Regan, S. Sampat, T. C. Sangster, A. B. Sefkow, W. Seka, R. C. Shah, W. T. Shmyada, A. Shvydky, C. Stoeckl, A. A. Solodov, W. Theobald, J. D. Zuegel, M. G. Johnson, R. D. Petrasso, C. K. Li, J. A. Frenje, Tripled yield in direct-drive laser fusion through statistical modelling, Nature 565 (2019) 581–586.

V. Gopalaswamy, C. A. Williams, R. Betti, D. Patel, J. P. Knauer, A. Lees, D. Cao, E. M. Campbell, P. Farmakis, R. Ejaz, K. S. Anderson, R. Epstein, J. Carroll-Nellenbeck, I. V. Igumenshchev, J. A. Marozas, P. B. Radha, A. A. Solodov, C. A. Thomas, K. M. Woo, T. J. B. Collins, S. X. Hu, W. Scullin, D. Turnbull, V. N. Goncharov, K. Churnetski, C. J. Forrest, V. Y. Glebov, P. V. Heuer, H. McClow, R. C. Shah, C. Stoeckl, W. Theobald, D. H. Edgell, S. Ivancic, M. J. Rosenberg, S. P. Regan, D. Bredesen, C. Fella, M. Koch, R. T. Janezic, M. J. Bonino, D. R. Harding, K. A. Bauer, S. Sampat, L. J. Waxer, M. Labuzeta, S. F. B. Morse, M. Gatu-Johnson, R. D. Petrasso, J. A. Frenje, J. Murray, B. Serrato, D. Guzman, C. Shuldberg, M. Farrell, C. Deeney, Demonstration of a hydrodynamically equivalent burning plasma in direct-drive inertial confinement fusion, Nature Physics 20 (2024) 751–757.

R. L. Gutierrez, S. Ghorbanpour, V. Gundecha, R. Ejaz, V. Gopalaswamy, R. Betti, A. Naug, D. Rengarajan, A. R. Babu, P. Faraboschi, S. Sarkar, Explainable meta bayesian optimization with human feedback for scientific applications like fusion energy, in: NeurIPS 2024 Workshop on Tackling Climate Change with Machine Learning, 2024.

A. Shmakov, A. Naug, V. Gundecha, S. Ghorbanpour, R. L. Gutierrez, A. R. Babu, A. Guillen, S. Sarkar, Rtdk-bo: High dimensional bayesian optimization with reinforced transformer deep kernels, in: 2023 IEEE 19th International Conference on Automation Science and Engineering (CASE), IEEE, 2023, pp. 1–8.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, 2017. URL: `https://arxiv.org/abs/1707.06347`. `arXiv:1707.06347`.

T. Nguyen, A. Grover, Transformer neural processes: Uncertainty-aware meta learning via sequence modeling, 2023. arXiv:2207.04179.

R. Ejaz, V. Gopalaswamy, A. Lees, C. Kanan, D. Cao, R. Betti, Deep learning-based predictive models for laser direct drive at the omega laser facility, Physics of Plasmas 31 (2024). doi:10.1063/5.0195675.

J. Delettrez, R. Epstein, M. C. Richardson, P. A. Jaanimagi, B. L. Henke, Effect of laser illumination nonuniformity on the analysis of time-resolved x-ray measurements in uv spherical transport experiments, Phys. Rev. A 36 (1987) 3926–3934. URL: https://link.aps.org/doi/10.1103/PhysRevA.36.3926. doi:10.1103/PhysRevA.36.3926.

C. A. Williams, R. Betti, V. Gopalaswamy, A. Lees, High yields in direct-drive inertial confinement fusion using thin-ice DT liner targets, Physics of Plasmas 28 (2021) 122708. URL: https://doi.org/10.1063/5.0069372. doi:10.1063/5.0069372.

R. Garnett, Bayesian Optimization, Cambridge University Press, 2023.