Fewer Shots, Better Implosions: Sample-Efficient Optimization for Inertial Confinement Fusion

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

[1]Hewlett Packard Enterprise, [2]University of Rochester
{vineet.gundecha, rluna, sahand.ghorbanpour, soumyendu.sarkar}@hpe.com
{reja, vgopalas, betti, alee}@lle.rochester.edu

Abstract

The global demand for clean energy has brought Inertial Confinement Fusion (ICF) to the forefront of sustainable power research. Due to the high cost and limited availability of ICF experiments, optimization methods must achieve exceptional sample efficiency. Bayesian Optimization (BO) is a standard tool for expensive black-box function optimization, yet it suffers from a "cold-start" problem, neglecting prior knowledge from simulations and past experiments. We propose a Meta-Bayesian Optimization (Meta-BO) approach that integrates prior tasks into the BO loop. Our method introduces boundary-box constraints, dual acquisition strategies, and interpretability candidates. Our method (MBO) achieves substantial performance gains over existing BO approaches in ICF energy-yield optimization, demonstrating sample efficiency and accelerating the path to a sustainable fusion power source.

1 Introduction

The urgent need for sustainable energy has placed Inertial Confinement Fusion (ICF) at the center of modern scientific research. By using powerful lasers to compress and heat fuel pellets to fusion conditions, ICF offers a path toward a nearly limitless, carbon-free power source. Despite this promise, the path to ignition is challenging. ICF experiments are exceptionally complex, costly, and limited to only a handful of opportunities annually. This reality makes sample efficiency, the ability to find an optimal solution with the fewest experiments possible, the single most critical factor in a successful research campaign.

Bayesian Optimization (BO) has been the go-to method for optimizing expensive, black-box environments (Feurer et al., 2015; Snoek et al., 2012; Wang et al., 2023; Shmakov et al., 2023; Gundecha et al., 2024; Gutierrez et al., 2024; Ghorbanpour et al., 2024). However, traditional BO begins each new optimization campaign with no prior knowledge, essentially starting from scratch every time. This approach is inefficient and wastes valuable experimental opportunities.

To solve this, we study Meta-Bayesian Optimization (Meta-BO) as a paradigm for ICF optimization (Bai et al., 2023). Instead of starting from a blank slate, Meta-BO learns from a library of related source tasks, such as simulations. It then uses this pre-trained knowledge to rapidly adapt its surrogate model and acquisition function to a new target experiment. This approach allows the optimization to begin in a much more informed state, leading to faster convergence and a higher probability of finding a superior solution.

*These authors contributed equally.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2025.

In this work, we adapt a Transformer Learning Processes (TNPs) (Nguyen and Grover, 2023) and propose a set of improvements to the standard BO loop for the ICF domain. Our findings provide a tangible and impactful solution for ICF practitioners. By empowering researchers with a tool that makes their experimental campaigns dramatically more effective, our work directly accelerates progress toward the ultimate goal of fusion energy.

2 Background

2.1 ICF

Inertial Confinement Fusion (ICF) harnesses nuclear fusion by compressing and heating a small fuel pellet, typically a deuterium-tritium mixture, to extreme temperatures and pressures. This is primarily achieved using high-power lasers that deliver a massive, precisely shaped energy burst within a few nanoseconds. The goal is to create conditions where atomic nuclei can overcome their natural repulsion and fuse, releasing substantial energy. The efficiency of this energy release is critically dependent on the laser pulse shape, which during an ICF experiment is optimized by controlling two essential parameters, the foot power and the picket power (Betti and Hurricane, 2016; Lees et al., 2021; Gopalaswamy et al., 2019, 2024). A representation of the pulse shape is shown in Figure 3.

2.2 Bayesian Optimization

Bayesian Optimization (BO) is a sequential decision-making framework for optimizing black-box objective functions f whose evaluations may be expensive or time-consuming (Garnett, 2023). It aims to find the global optimum $x^* = \text{argmax}_{x \in \mathcal{X}} f(x)$, where $\mathcal{X} \subseteq \mathbb{R}^d$ is the input domain and x^* denotes the maximizer.

BO operates by constructing a surrogate model of f from past evaluations, then using this model to guide the selection of the next query point $x \in \mathcal{X}$ via an acquisition function (AF). Gaussian Processes (GPs) are a common choice for surrogates due to their flexibility and well-calibrated uncertainty estimates. Popular AFs include the Upper Confidence Bound (UCB) (Jones, 2001; Srinivas et al., 2012) and Expected Improvement (EI) (Močkus et al., 1978; Jones et al., 1998), which balance exploration of uncertain regions with exploitation of promising areas (Garnett, 2023).

2.3 Meta-Bayesian Optimization

Meta-Bayesian Optimization (Meta-BO) extends BO by leveraging knowledge from N related source tasks $\mathcal{F}$ to accelerate the optimization of a new, unseen target black-box function (Maraval et al., 2023). The prior knowledge is provided in the form of datasets $\mathcal{D}_1, \ldots, \mathcal{D}_N$, where each $\mathcal{D}_n = \{(x_n^i, y_n^i)\}_{i=1}^{e_n}, \quad y_n^i = f_n(x_n^i), \quad f_n \in \mathcal{F}$, contains e_n evaluations from the n-th source function.

While GPs have been the dominant surrogate model in single-task BO, Meta-BO often benefits from uncertainty-aware meta-learning methods such as Neural Processes (NPs) (Garnelo et al., 2018). NPs combine the expressive power of neural networks with the probabilistic nature of GPs, and are trained in a meta-learning framework to enable rapid adaptation to new functions (Jha et al., 2023).

Given a context set $\mathcal{D}_c$ of observed input–output pairs and an unlabeled set of target inputs $\mathcal{X}_T = \{x^i\}_{i=1}^{l}$, NPs produce a predictive distribution $p(\cdot \mid \mathcal{X}_T, \mathcal{D}_c)$ that approximates the true posterior over the target outputs $\mathcal{Y}_T$.

Building on recent state-of-the-art advances (Nguyen and Grover, 2023; Maraval et al., 2023; Feng et al., 2023; Müller et al., 2023), this work adopts Transformer Neural Processes (TNPs), a variant of NPs that employs transformer architectures for enhanced representation learning and scalability.

3 Meta-Bayesian Optimization for ICF

Transformer Neural Processes (TNPs) have demonstrated impressive performance as surrogate models in meta-learning settings (Nguyen and Grover, 2023; Gutierrez et al., 2024; Gundecha et al., 2024). We build on prior TNP-based Meta-BO approaches and propose three refinements tailored to ICF:

- We introduce AF **boundary boxes** to enforce diversity in candidate experiments.
- We propose a **dual acquisition function** strategy that offers practitioners diverse candidate solutions.
- We offer interpretable **response surface representations** of the surrogate's predictions and AF values to provide practitioners with deeper insights about the optimization.

These enhancements are specifically designed to address the unique constraints of real-world, high-stakes experiments.

4 Boundary Boxes

For ICF, the high cost and limited availability of experimental time make it essential to avoid redundant trials, which would otherwise be a waste of valuable resources. To enforce this practical constraint, we introduce the concept of **boundary boxes** to a standard AFs, such as EI or UCB. Instead of allowing the AF to propose points anywhere in the search space, we define and mask regions around previously sampled points, making them "invisible" to the AF and thus preventing re-sampling. The sizes of these regions can be dynamically changed and are defined by the ICF experts during an online experiment.

Let the search space be $\mathcal{X} \subset \mathbb{R}^D$, where D is the number of experimental parameters and the set of previously sampled points be $\mathcal{X}_{\text{prev}} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{k-1}\}$. We define a hyperrectangular boundary box B_i around each point $\mathbf{x}_i$ to be excluded from the search. This box is defined by a vector of radii $\mathbf{r} = (r_1, \ldots, r_D)$, where each r_j specifies the half-width of the box in dimension j.

The boundary box B_i around a point $\mathbf{x}_i$ is formally given by:

$$B_i = \{\mathbf{x} \in \mathcal{X} \mid |\mathbf{x}_j - \mathbf{x}_{i,j}| \leq r_j, \quad \forall j \in \{1, \ldots, D\}\}.$$

We then define the total excluded region as the union of all such boxes, $\mathcal{B} = \bigcup_{i=1}^{k-1} B_i$. The standard acquisition function $AF(\mathbf{x})$ is then modified to create a masked acquisition function, $AF_{\text{masked}}(\mathbf{x})$, which effectively discourages sampling within the excluded regions:

$$AF_{\text{masked}}(\mathbf{x}) = \begin{cases} AF(\mathbf{x}) & \text{if } \mathbf{x} \notin \mathcal{B} \\ -\infty & \text{if } \mathbf{x} \in \mathcal{B} \end{cases}$$

The next experimental point is then chosen by maximizing this masked function over the entire search space, ensuring it is a predefined distance away from all previously tested points: $\mathbf{x}_k = \arg\max_{\mathbf{x} \in \mathcal{X}} AF_{\text{masked}}(\mathbf{x})$. This mechanism guarantees the exploration of novel parameter spaces while adhering to practical experimental requirements.

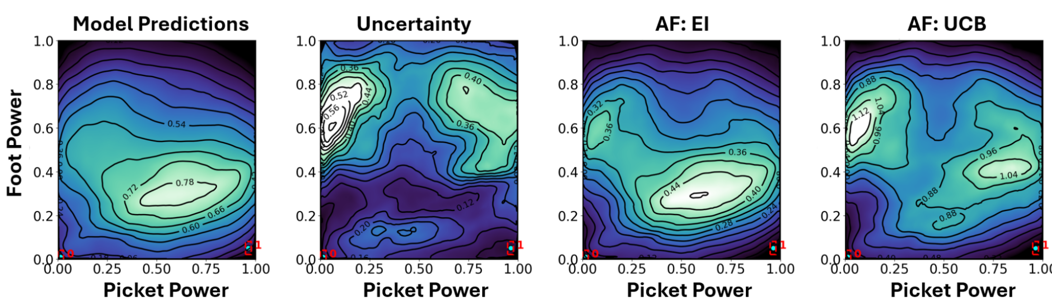

Figure 1: Illustration of an ICF trajectory in the early stages of optimization. EI primarily proposes to exploit regions with high surrogate predictions, while UCB prioritizes exploration of uncertain regions. Such visualizations provide ICF practitioners with insights to select the candidate most likely to succeed. The numbering represents the order in which samples were taken. The red boxes represent the boundary boxes defined by the ICF expert.

4.1 Dual Acquisition Function

In high-stakes scientific optimization, a single acquisition function recommendation may not always provide the most useful guidance to a practitioner. To offer a more comprehensive view of the search space, we propose a **dual acquisition function** approach. Instead of a single candidate point, we provide two distinct points, each offering a different perspective on where to conduct the next experiment. In our approach, the **first candidate** is proposed using EI as AF, while the **second candidate** is chosen by UCB.

To ensure these recommendations are meaningfully different from one another, we apply our boundary box mechanism between the two proposed candidates. This ensures the candidates are not in the immediate vicinity of each other, guaranteeing that both options represent a meaningful step in the optimization process.

4.2 ICF Response Surface Representations

For ICF, it is crucial to offer users transparent and interpretable representations of the predictions and decisions generated by complex black-box models, such as the one proposed in this paper. To enhance user understanding and trust, we provide intuitive visual representations that convey both the predicted outputs of the TNPs and the associated uncertainties in those predictions. Additionally, to facilitate a deeper insight into the optimization process, we present visualizations of the AF across the response or search space, enabling users to comprehend how the model navigates and prioritizes different regions during decision-making. Figure 1 illustrates an example of these informative representations, highlighting the interplay between model predictions, uncertainty quantification, and optimization guidance.

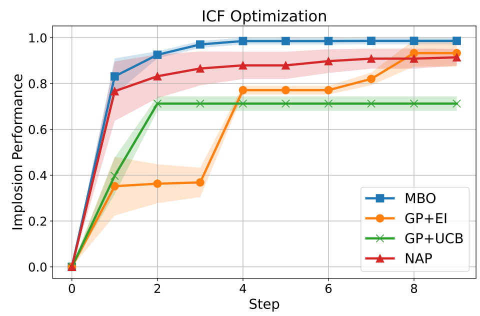

Figure 2: Comparison between BO methods and MBO in ICF. The shaded areas represent a ± 1 standard error. Performance is normalized.

5 Experimental Evaluation

We evaluate the performance of different BO approaches, in the task of energy yield optimization for ICF. We compare our approach (MBO) with classical BO baselines, which utilizes a GP surrogate model combined with classical acquisition functions like UCB and EI (Garnett, 2023; Chung et al., 2020; Vazirani et al., 2024; Hatfield et al., 2019; Li et al., 2023; Vazirani et al., 2023). Additionally, we include NAP (Maraval et al., 2023), the state-of-the-art method for Meta-BO, as a baseline (Gundecha et al., 2024). For MBO, a human expert selected the candidate to be sampled, choosing from the proposals generated by the two acquisition functions. To simulate a real-world constraint, we limited the optimization to 10 total samples (shots). Our evaluation involved running all approaches on 6 distinct test tasks, using 10 different initialization seeds for robustness. The performance of these different methods is presented in Figure 2, quantified by the mean normalized energy yield of the proposed candidates. This yield is reported as the best candidate point found up to that point during each optimization run. Our results demonstrate MBO's superior performance over all baselines, highlighting its benefits for the ICF domain. Experimental details can be found in Section C (Appendix).

6 Conclusion

Our results demonstrate that our approach significantly improves the effectiveness of inertial confinement fusion (ICF) experimental campaigns by achieving high sample efficiency. By providing optimized experimental solutions with minimal required shots, MBO reduces operational costs, accelerates experimental cycles, and maximizes the scientific return of limited shot opportunities. Consequently, MBO offers a practical and scalable path toward faster progress in fusion energy research and a more efficient utilization of high-value experimental facilities.

References

M. Feurer, A. Klein, K. Eggensperger, J. Springenberg, M. Blum, F. Hutter, Efficient and robust automated machine learning, in: C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 28, Curran Associates, Inc., 2015. URL: `https://proceedings.neurips.cc/paper_files/paper/2015/file/11d0e6287202fced83f79975ec59a3a6-Paper.pdf`.

J. Snoek, H. Larochelle, R. P. Adams, Practical bayesian optimization of machine learning algorithms, in: F. Pereira, C. Burges, L. Bottou, K. Weinberger (Eds.), Advances in Neural Information Processing Systems, volume 25, Curran Associates, Inc., 2012. URL: `https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf`.

X. Wang, Y. Jin, S. Schmitt, M. Olhofer, Recent advances in bayesian optimization, ACM Comput. Surv. 55 (2023). URL: `https://doi.org/10.1145/3582078`. doi:`10.1145/3582078`.

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.

V. Gundecha, R. L. Gutierrez, S. Ghorbanpour, R. Ejaz, V. Gopalaswamy, R. Betti, D. Rengarajan, S. Sarkar, Meta-learned bayesian optimization for energy yield in inertial confinement fusion, in: NeurIPS 2024 Workshop on Physical Sciences, 2024. URL: `https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_4.pdf`.

R. L. Gutierrez, S. Ghorbanpour, V. Gundecha, R. Ejaz, V. Gopalaswamy, R. Betti, A. Naug, D. Rengarajan, A. R. Babu, P. Faraboschi, et al., 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.

S. Ghorbanpour, R. L. Gutierrez, V. Gundecha, D. Rengarajan, A. R. Babu, S. Sarkar, Llm enhanced bayesian optimization for scientific applications like fusion, in: NeurIPS 2024 Workshop on Physical Sciences, 2024.

T. Bai, Y. Li, Y. Shen, X. Zhang, W. Zhang, B. Cui, Transfer learning for bayesian optimization: A survey, 2023. `arXiv:2302.05927`.

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

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

D. R. Jones, A taxonomy of global optimization methods based on response surfaces, Journal of Global Optimization 21 (2001) 345–383. doi:10.1023/A:1012771025575.

N. Srinivas, A. Krause, S. M. Kakade, M. W. Seeger, Information-theoretic regret bounds for gaussian process optimization in the bandit setting, IEEE Transactions on Information Theory 58 (2012) 3250–3265. URL: http://dx.doi.org/10.1109/TIT.2011.2182033. doi:10.1109/tit.2011.2182033.

J. Močkus, J. Tiesis, A. Žilinskas, The application of bayesian methods for seeking the extremum, in: L. Dixon, G. Szegö (Eds.), Towards Global Optimization, volume 2, North-Holland, 1978, pp. 117–129.

D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box functions, Journal of Global Optimization 13 (1998) 455–492. doi:10.1023/A:1008306431147.

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

M. Garnelo, D. Rosenbaum, C. J. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. W. Teh, D. J. Rezende, S. M. A. Eslami, Conditional neural processes, 2018. URL: https://arxiv.org/abs/1807.01613. arXiv:1807.01613.

S. Jha, D. Gong, X. Wang, R. E. Turner, L. Yao, The neural process family: Survey, applications and perspectives, 2023. arXiv:2209.00517.

L. Feng, H. Hajimirsadeghi, Y. Bengio, M. O. Ahmed, Latent bottlenecked attentive neural processes, 2023. arXiv:2211.08458.

S. Müller, N. Hollmann, S. P. Arango, J. Grabocka, F. Hutter, Transformers can do bayesian inference, 2023. arXiv:2112.10510.

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, 2024.

V. Gundecha, R. Luna Gutierrez, S. Ghorbanpour, R. Ejaz, V. Gopalaswamy, R. Betti, A. Naug, P. Faraboschi, S. Sarkar, Meta-learned bayesian optimization for energy yield in inertial confinement fusion, 2024. URL: https://www.climatechange.ai/papers/neurips2024/35.

Y. Chung, I. Char, W. Neiswanger, K. Kandasamy, A. O. Nelson, M. D. Boyer, E. Kolemen, J. Schneider, Offline contextual bayesian optimization for nuclear fusion, 2020. URL: https://arxiv.org/abs/2001.01793. arXiv:2001.01793.

N. N. Vazirani, R. Sacks, B. M. Haines, M. J. Grosskopf, D. J. Stark, P. A. Bradley, Bayesian batch optimization for molybdenum versus tungsten inertial confinement fusion double shell target design, Statistical Analysis and Data Mining: An ASA Data Science Journal 17 (2024) e11698. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/sam.11698. doi:https://doi.org/10.1002/sam.11698. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/sam.11698.

P. W. Hatfield, S. J. Rose, R. H. H. Scott, The blind implosion-maker: Automated inertial confinement fusion experiment design, Physics of Plasmas 26 (2019). URL: http://dx.doi.org/10.1063/1.5091985. doi:10.1063/1.5091985.

Z. Li, Z. Q. Zhao, X. H. Yang, G. B. Zhang, Y. Y. Ma, H. Xu, F. Y. Wu, F. Q. Shao, J. Zhang, Hybrid optimization of laser-driven fusion targets and laser profiles, 2023. URL: https://arxiv.org/abs/2305.15434. arXiv:2305.15434.

N. N. Vazirani, M. J. Grosskopf, D. J. Stark, P. A. Bradley, B. M. Haines, E. N. Loomis, S. L. England, W. A. Scales, Coupling multi-fidelity xrage with machine learning for graded inner shell design optimization in double shell capsules, Physics of Plasmas 30 (2023) 062704. URL: https://doi.org/10.1063/5.0129565. doi:10.1063/5.0129565.

D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017. URL: https://arxiv.org/abs/1412.6980. arXiv:1412.6980.

G. De Ath, R. M. Everson, A. A. M. Rahat, J. E. Fieldsend, Greed is good: Exploration and exploitation trade-offs in bayesian optimisation, ACM Transactions on Evolutionary Learning and Optimization 1 (2021) 1–22. URL: http://dx.doi.org/10.1145/3425501. doi:10.1145/3425501.

A. S. Raihan, H. Khosravi, S. Das, I. Ahmed, Accelerating material discovery with a threshold-driven hybrid acquisition policy-based bayesian optimization, 2023. URL: https://arxiv.org/abs/2311.09591. arXiv:2311.09591.

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. Ejaz, V. Gopalaswamy, R. Luna, V. Gundecha, A. Lees, R. Betti, S. Ghorbanpour, S. Sarkar, C. Kanan, Optimizing the performance of direct-drive implosion experiments using meta-bayesian optimization, Bulletin of the American Physical Society (2024).

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, et al., Understanding the fusion yield dependencies in omega dt-layered implosion experiments using a physics-based statistical mapping model, Physics of Plasmas 30 (2023).

A MBO's Surrogate Predictions

To understand the superior performance achieved by our method, 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 against MBO predictions. For MBO, first we evaluate its initial predictions when using one context point (1 sample), so the information provided to the surrogate comes only from the first evaluation sample. Moreover, we assess how MBO adapts by querying it after three samples have been collected from the target function (3 Samples). Figure 4 shows the results of the evaluation. From these results, we can see that MBO effectively uses information from source tasks to identify regions with a higher likelihood of finding optimal solutions. Furthermore, after incorporating just a few samples, MBO accurately adjusts its predictions to match the target function closely.

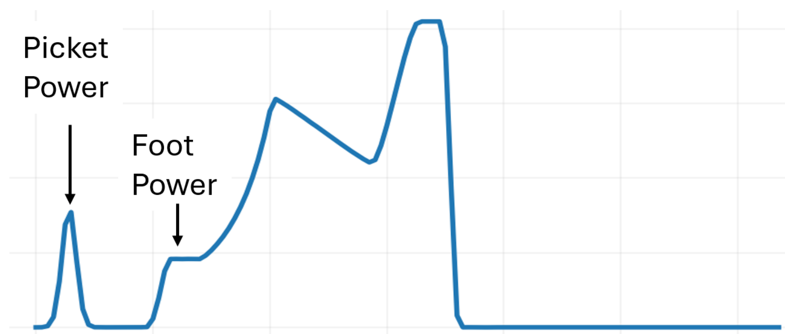

Figure 3: Representation of a laser pulse shape used in ICF experiments. During online optimization, the foot and picket powers are tuned to improve implosion performance

B MBO vs NAP Adaptation

To better understand the superior performance of MBO against NAP, the state-of-art-approach of Meta-BO which has shown good results in optimizing ICF objectives (Gundecha et al., 2024), we compare the predictions of both model's surrogate. Figure 5 shows the results of this comparison. For each trajectory we follow the candidates proposed by each method. We can observe that the predictions of MBO are more accurate and smoother, Moreover, MBO's surrogate predictions are more interpretable.

C Experimental Details

For all experiments involving MBO, we employed the autoregressive variant of Transformer Neural Processes (TNPs-A) (Nguyen and Grover, 2023). The model was configured using the hyperparameters recommended in the original work (Table 1). Optimization was performed using Adam (Kingma and Ba, 2017) with a cosine annealing learning rate schedule. Random seeds were set using Python's built-in `random` library.

All training and evaluation were carried out on a server equipped with an AMD EPYC 7542 32-core CPU and four NVIDIA H100 GPUs. These experiments ran for approximately 180 hours at $\sim 80\%$ GPU utilization.

D EI and UCB Candidate Trajectory

Previous work has compared acquisition functions such as EI and UCB, showing that UCB tends to favor exploration while EI prioritizes exploitation (De Ath et al., 2021; Raihan et al., 2023). We observe the same behavior in the ICF domain, as illustrated in Figure 6. In the early stages of optimization, EI is inherently more *exploitative*, focusing on regions predicted to outperform the current best observation. In contrast, UCB is more *explorative*, as it explicitly incorporates

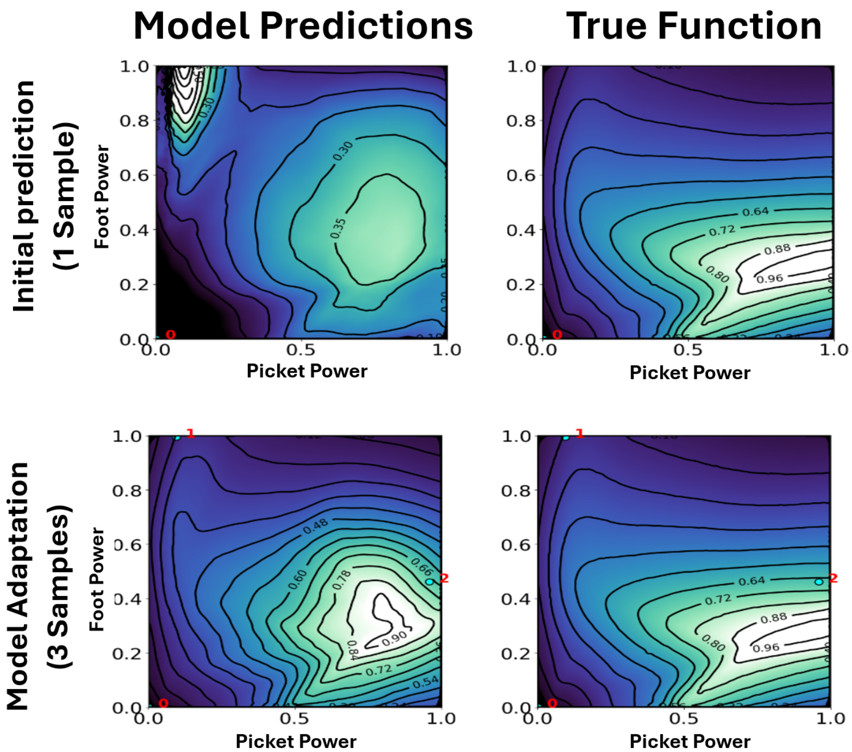

Figure 4: MBO's meta-learned surrogate predictions: (top) without only one context point; (bottom) after three context points; (right) the optimization target function. We can observe that our approach achieves quick adaptation with high sample efficiency, closely approximating the true function in just three samples. This property is ideal for the limited ICF experiments possible on a shot day.

TNPs' Hyperparameters	
Model dimension	64
Embedding layers	4
Feed forward dimension	128
Attention heads	8
Transformer layers	6
Dropout	0.0
Learning rate	$5e-5$

Table 1: Hyperparameters used by the TNPs for all experiments.

predictive uncertainty and encourages sampling in less-visited areas of the search space. By uniting these two perspectives in our dual acquisition strategy, we provide scientists with complementary candidate proposals, one guided by exploitation and the other by exploration. This not only makes the optimization process more transparent and interpretable, but also empowers practitioners to exercise informed scientific judgment when selecting the next experiment, a critical advantage in high-stakes settings where each trial is costly and limited.

D.1 ICF Dataset

To construct the ICF dataset (the collection of source tasks used to train our meta-learning models) we employed the LOTUS library to generate laser pulse shapes using a custom parameterization (Ejaz et al., 2024). The resulting laser power and timing profiles (laser pulse shape) served as inputs to LILAC (Delettrez et al., 1987), a physics-based simulator for laser-driven fusion. For a fixed fusion fuel target, the laser pulse shape determines key outcomes such as neutron yield, influencing both target compression and the growth of hydrodynamic instabilities (Williams et al., 2021).

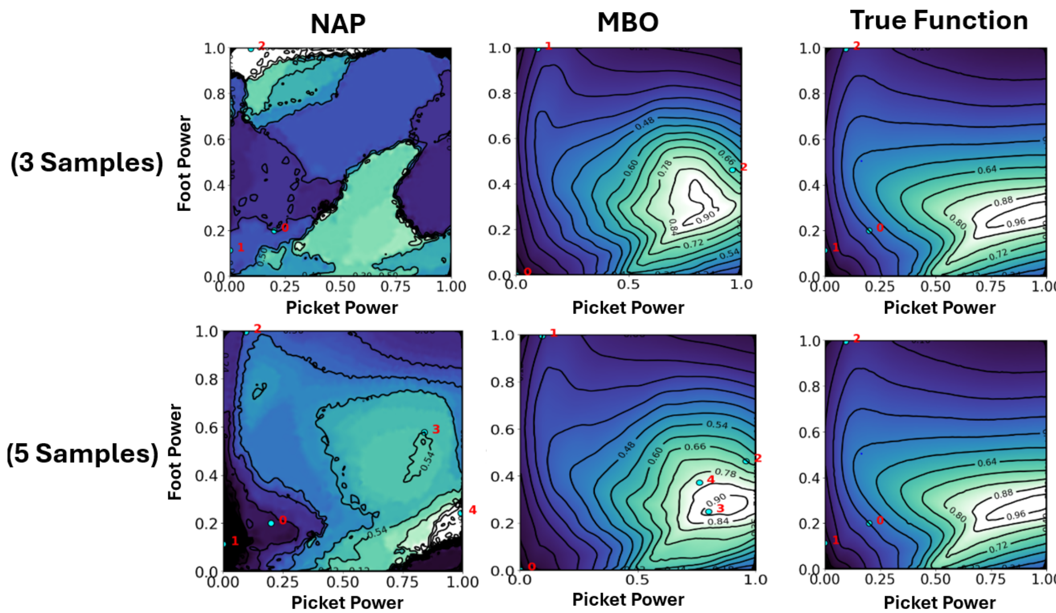

Figure 5: MBO's meta-learned surrogate prediction against NAP's surrogate prediction. MBO's predictions are more accurate, smoother and interpretable.

To create the response surface, we varied two front-end parameters of the pulse shape representation, generating 1000 samples via Latin hypercube sampling while respecting the laser system's design limits. Each pulse shape was simulated in LILAC with the same target configuration, and the resulting optimization objectives (neutron yield and ρR), which are the ICF performance metrics under consideration (Ejaz et al., 2024), were recorded to define the response surface over the two parameters.

We generated 16 source tasks, 6 validation tasks, and 6 test tasks by perturbing the simulator's physics models, such as the equation of state, thermal transport properties, and inclusion of 3D degradation effects(Lees et al., 2023), these alter the response surfaces and generates a variety of source tasks to learn from and evaluate on. These physics variations mimic the mismatch between simulations and real experiments, enabling us to evaluate our method under realistic domain-shift conditions.

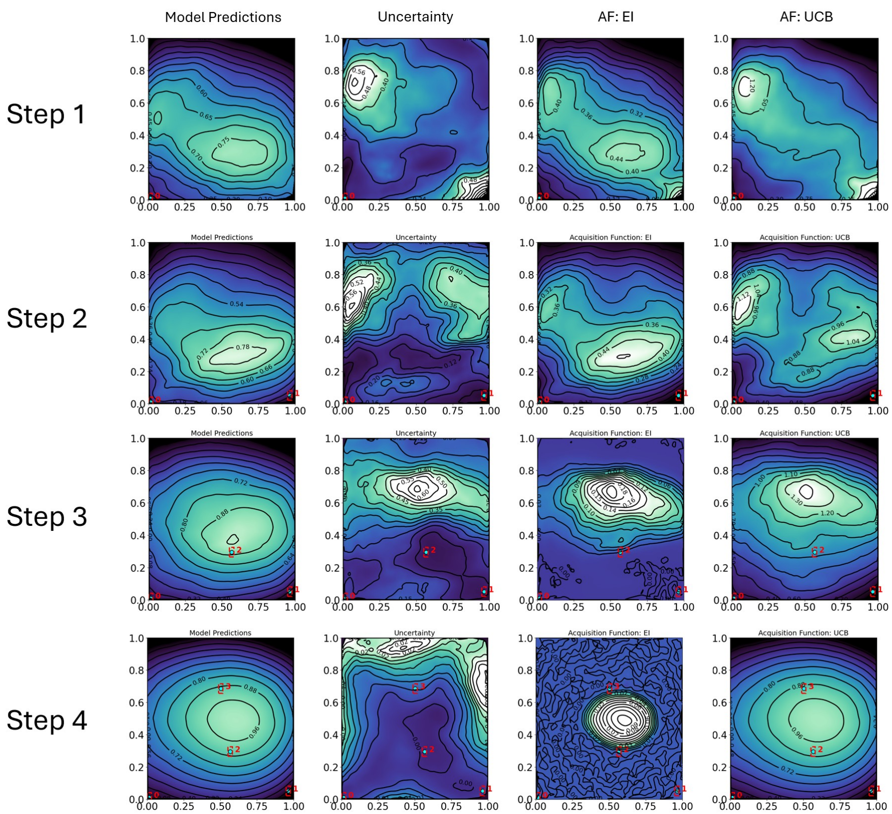

Figure 6: Comparison of acquisition strategies. In early stages, EI prioritizes regions predicted to yield improvements over the current best (exploitative), while UCB balances the mean prediction with uncertainty, sampling in less-explored areas (explorative). On later stages, both AF converge to a similar point. Together, they provide complementary candidate proposals for high-stakes optimization.