
OptICF: Sample-Efficient Optimization of Implosion Outcomes in Inertial Confinement Fusion

Ricardo Luna Gutierrez^{1*}, Vineet Gundecha^{1*}, Rahman Ejaz^{2*},
Varchas Gopaldaswamy^{2*}, Riccardo Betti^{2*}, Sahand Ghorbanpour^{1*}, Aarne Lees^{1*},
Soumyendu Sarkar^{1*†}

¹Hewlett Packard Enterprise ²University of Rochester
{vineet.gundecha, rluna, sahand.ghorbanpour, soumyendu.sarkar}@hpe.com
{reja, vgopaldas, 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 (OptICF) 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 [1–4]. 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 [5]. 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.

In this work, we build on uncertainty-aware Meta-Bayesian Optimization [6] and extend it to the unique demands of Inertial Confinement Fusion (ICF). Specifically, we propose three innovations:

*These authors contributed equally.

†Corresponding author.

(i) **acquisition function boundary boxes**, which explicitly enforce diversity and prevent redundant sampling under scarce experimental budgets, (ii) a **dual acquisition strategy** that proposes complementary exploitative and explorative candidates, giving practitioners robust and actionable choices, and (iii) **response surface visualizations** that reveal both surrogate predictions and acquisition landscapes, enhancing interpretability and practitioner trust. Together, these improvements transform Meta-BO from a purely algorithmic tool into a practical, transparent, and sample-efficient framework for high-stakes scientific optimization.

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 critical, y dependent on the laser pulse’s shape, which during an ICF experiment is optimized by controlling two essential parameters, the foot-power and the picket-power [7–15].

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 [16]. It aims to find the global optimum $x^* = \operatorname{argmax}_{x \in \mathcal{X}} f(x)$; where $\mathcal{X} \subset \mathbb{R}^d$ 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) [17, 18] and Expected Improvement (EI) [19, 20], which balance exploration of uncertain regions with exploitation of promising areas [16].

2.3 Meta-Bayesian Optimization

Meta-Bayesian Optimization (Meta-BO) extends BO by leveraging knowledge from N related source tasks F to accelerate the optimization of a new, unseen target black-box function [21]. The prior knowledge is provided in the form of datasets $D_1; \dots; D_N$, where each $D_n = \{(x_n^i, y_n^i)_{i=1}^{e_n}\}$; $y_n^i = f_n(x_n^i)$; $f_n \in F$; 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) [22]. 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 [23].

Given a context set D_c of observed input–output pairs and an unlabeled set of target inputs $X_T = \{x^i\}_{i=1}^j$, NPs produce a predictive distribution $p(y | X_T; D_c)$ that approximates the true posterior over the target outputs Y_T .

Building on recent state-of-the-art advances [6, 21, 24, 25], 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 recently shown strong performance as surrogate models in meta-learning settings [6, 26, 27]. However, directly applying prior TNP-based Meta-BO approaches to Inertial Confinement Fusion (ICF) is insufficient: ICF experiments are costly, scarce, and demand

both diversity in exploration and transparency for domain experts. To address these challenges, we introduce a Meta-Bayesian Optimization framework with three key innovations:

- **Acquisition function boundary boxes** that explicitly prevent redundant sampling and enforce diversity, ensuring that limited experimental opportunities are maximally informative.
- **Dual acquisition function optimization**, which provides complementary candidate proposals (exploitative and explorative), empowering practitioners with choice and robustness in high-stakes decision-making.
- **Interpretable response surface visualizations** of both surrogate predictions and acquisition landscapes, bridging the gap between black-box optimization and scientific insight.

Together, these contributions extend Meta-BO beyond “faster optimization” to a setting where sample efficiency, practitioner trust, and actionable decision support are critical. While we develop and validate our approach in ICF, the proposed framework is broadly applicable to other domains where experiments are scarce, expensive, and high-impact.

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 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 $X \subseteq \mathbb{R}^D$, where D is the number of experimental parameters and the set of previously sampled points be $X_{\text{prev}} = \{\mathbf{x}_1, \dots, \mathbf{x}_{k-1}\}$. 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 $\mathbf{r} = (r_1, \dots, r_D)$, 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:

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

We then define the total excluded region as the union of all such boxes, $B = \bigcup_{i=1}^{k-1} B_i$. 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:

$$AF_{\text{masked}}(\mathbf{x}) = \begin{cases} AF(\mathbf{x}) & \text{if } \mathbf{x} \notin B \\ 1 & \text{if } \mathbf{x} \in B \end{cases}$$

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: $\mathbf{x}_k = \arg \max_{\mathbf{x} \in X} AF_{\text{masked}}(\mathbf{x})$. This mechanism guarantees the exploration of novel parameter spaces while adhering to practical experimental requirements.

4.1 Dual Acquisition Function

In high-stakes scientific optimization, relying on a single acquisition function can be limiting: practitioners must commit to either an exploitative or explorative strategy, even when uncertainty is high and experimental shots are scarce. To overcome this, we introduce a **dual acquisition function** approach that simultaneously presents two complementary candidate points. Specifically, the **first candidate** is proposed using EI, which emphasizes exploitation of promising regions, while the **second candidate** is generated via UCB, which prioritizes exploration of uncertain areas.

To ensure that these recommendations are not redundant, we apply our boundary-box mechanism between them, forcing meaningful separation in the search space. This guarantees that both options represent distinct scientific hypotheses, empowering practitioners with interpretable alternatives rather than a single opaque recommendation.

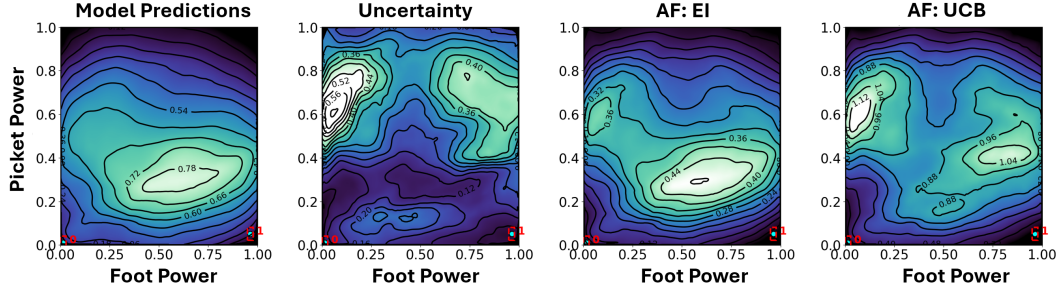


Figure 1: Illustration of an ICF trajectory in the early stages of optimization. EI primarily proposes to exploit regions with high surrogate predictions, while UCB prioritizes exploration of uncertain regions. Such visualizations provide ICF practitioners with insights to select the candidate most likely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

4.2 ICF Response Surface Representations

In high-stakes applications, 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.

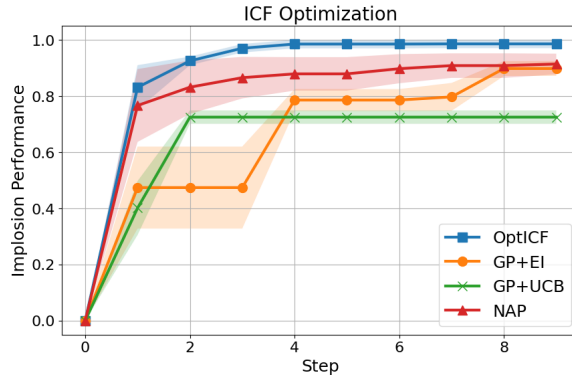


Figure 2: Comparison between BO methods and OptICF in ICF. The shaded areas represent a 1 standard error. Performance is normalized.

5 Experimental Evaluation

We evaluate our method on the task of energy yield optimization for ICF, comparing it against several BO baselines. Specifically, we benchmark our approach (OptICF) against classical BO methods that employ GPs with acquisition functions such as UCB and EI [16, 28–32], as well as against NAP [21], the current state-of-the-art in Meta-BO [27]. For OptICF, the candidate to sample between both AFs was selected by a human based on the shown representations. Performance is measured as the maximum energy yield observed up to each step of the optimization. To reflect real-world constraints, we restrict the number of evaluations to only 10 samples per run, and average results over 10 random seeds. Figure 2 summarizes the outcomes: OptICF consistently outperforms all baselines, underscoring its ability to deliver greater sample efficiency in the ICF domain. Further experimental details are provided in Appendix D.

6 Conclusion

Our results demonstrate that our approach improves the effectiveness of ICF experimental campaigns by achieving high sample efficiency. By providing sample efficient solutions, OptICF reduces experimental cost and maximizes the value of limited shot opportunities, offering a practical path toward faster progress in fusion energy research.

References

- [1] M. Feurer, A. Klein, K. Eggenberger, 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/11d0e6287202fcd83f79975ec59a3a6-Paper.pdf .
- [2] 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 .
- [3] 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 .
- [4] A. Shmakov, A. Naug, V. Gundecka, 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.
- [5] T. Bai, Y. Li, Y. Shen, X. Zhang, W. Zhang, B. Cui, Transfer learning for bayesian optimization: A survey, 2023arXiv:2302.05927 .
- [6] T. Nguyen, A. Grover, Transformer neural processes: Uncertainty-aware meta learning via sequence modeling, 2023arXiv:2207.04179 .
- [7] R. Betti, O. A. Hurricane, Inertial-con nement fusion with lasers, *Nature Physics* 12 (2016) 435–448. URL <https://doi.org/10.1038/nphys3736> . doi:10.1038/nphys3736 .
- [8] A. Lees, R. Betti, J. P. Knauer, V. Gopaldaswamy, 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 con nement 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 .
- [9] V. Gopaldaswamy, 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.
- [10] V. Gopaldaswamy, 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. Shulderg, M. Farrell, C. Deeney, Demonstration of a hydrodynamically equivalent burning plasma in direct-drive inertial con nement fusion, *Nature Physics* 20 (2024) 751–757.

- [11] V. Gundecha, R. L. Gutierrez, S. Ghorbanpour, R. Ejaz, V. Gopaldaswamy, 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.
- [12] R. L. Gutierrez, S. Ghorbanpour, V. Gundecha, R. Ejaz, V. Gopaldaswamy, 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.
- [13] 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.
- [14] R. L. Gutierrez, V. Gundecha, R. Ejaz, V. Gopaldaswamy, R. Betti, S. Ghorbanpour, A. Lees, S. Sarkar, Im-lpg: Inverse modeling approach to laser pulse shape generation in inertial confinement fusion, in: NeurIPS 2025 AI for Science Workshop, 2025.
- [15] 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), 2023, pp. 1–8. doi:10.1109/CASE56687.2023.10260520
- [16] R. Garnett, Bayesian Optimization, Cambridge University Press, 2023.
- [17] 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.
- [18] 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
- [19] J. Mockus, J. Tiesis, A. ilinskas, 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.
- [20] 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
- [21] A. Maraval, M. Zimmer, A. Grosnit, H. B. Ammar, End-to-end meta-bayesian optimisation with transformer neural processes, 2023. arXiv:2305.15930
- [22] 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
- [23] S. Jha, D. Gong, X. Wang, R. E. Turner, L. Yao, The neural process family: Survey, applications and perspectives, 2023. arXiv:2209.00517
- [24] L. Feng, H. Hajimirsadeghi, Y. Bengio, M. O. Ahmed, Latent bottlenecked attentive neural processes, 2023. arXiv:2211.08458
- [25] S. Müller, N. Hollmann, S. P. Arango, J. Grabocka, F. Hutter, Transformers can do bayesian inference, 2023. arXiv:2112.10510
- [26] R. L. Gutierrez, S. Ghorbanpour, V. Gundecha, R. Ejaz, V. Gopaldaswamy, 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.
- [27] V. Gundecha, R. Luna Gutierrez, S. Ghorbanpour, R. Ejaz, V. Gopaldaswamy, 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>

- [28] Y. Chung, I. Char, W. Neiswanger, K. Kandasamy, A. O. Nelson, M. D. Boyer, E. Kolemen, J. Schneider, Of ine contextual bayesian optimization for nuclear fusion, 2020. [URLs: //arxiv.org/abs/2001.01793](https://arxiv.org/abs/2001.01793) . arXiv:2001.01793 .
- [29] 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 con nement 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](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 .
- [30] P. W. Hat eld, S. J. Rose, R. H. H. Scott, The blind implosion-maker: Automated inertial con nement fusion experiment design, Physics of Plasmas 26 (2019). [URL: https://dx.doi.org/10.1063/1.5091985](https://dx.doi.org/10.1063/1.5091985) . doi:10.1063/1.5091985 .
- [31] 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 pro les, 2023. [URLs: //arxiv.org/abs/2305.15434](https://arxiv.org/abs/2305.15434) . arXiv:2305.15434 .
- [32] 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- delity 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](https://doi.org/10.1063/5.0129565) . doi:10.1063/5.0129565 .
- [33] 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: https://dx.doi.org/10.1145/3425501](https://dx.doi.org/10.1145/3425501) . doi:10.1145/3425501 .
- [34] A. S. Raihan, H. Khosravi, S. Das, I. Ahmed, Accelerating material discovery with a threshold-driven hybrid acquisition policy-based bayesian optimization, 2023. [URLs: //arxiv.org/abs/2311.09591](https://arxiv.org/abs/2311.09591) . arXiv:2311.09591 .
- [35] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017. [URLs: //arxiv.org/abs/1412.6980](https://arxiv.org/abs/1412.6980) . arXiv:1412.6980 .
- [36] 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: 1063/5.0195675 .
- [37] 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. [URLs: //link.aps.org/doi/10.1103/PhysRevA.36.3926](https://link.aps.org/doi/10.1103/PhysRevA.36.3926) . doi:10.1103/PhysRevA.36.3926 .
- [38] C. A. Williams, R. Betti, V. Gopalaswamy, A. Lees, High yields in direct-drive inertial con nement fusion using thin-ice DT liner targets, Physics of Plasmas 28 (2021) 122708. [URL: https://doi.org/10.1063/5.0069372](https://doi.org/10.1063/5.0069372) . doi:10.1063/5.0069372 .

A OptICF’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 OptICF predictions. For OptICF, 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 OptICF adapts by querying it after three samples have been collected from the target function (3 Samples). Figure 3 shows the results of the evaluation. From these results, we can see that OptICF effectively uses information from source tasks to identify regions with a higher likelihood of finding optimal solutions. Furthermore, after incorporating just a few samples, OptICF accurately adjusts its predictions to match the target function closely.

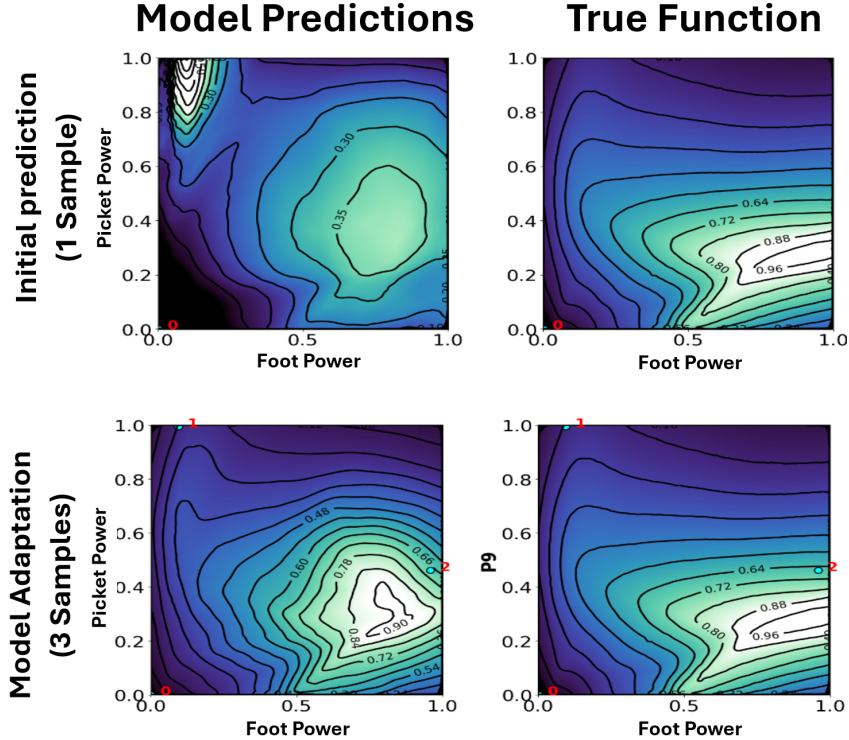


Figure 3: OptICF’s meta-learned surrogate predictions: (top) without only one context point; (bottom) after three context points; (right) the optimization target function. We can observe that our approach achieves quick adaptation with high sample efficiency, closely approximating the true function in just three samples. This property is ideal for the limited ICF experiments possible on a shot day.

B OptICF vs NAP Adaptation

To gain deeper insight into the superior performance of OptICF over NAP, we compare their surrogate predictions against the state-of-the-art Meta-BO approach, which has demonstrated strong results in optimizing ICF objectives [27]. Figure 4 presents the outcomes of this comparison. For each trajectory, we trace the candidates proposed by the respective methods. The results show that MBO produces predictions that are not only more accurate and smoother but also yield surrogate outputs that are easier to interpret.

C 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 [33, 34]. We observe the same behavior in the ICF domain, as illustrated in Figure 5. In the early stages of optimization, EI is inherently more

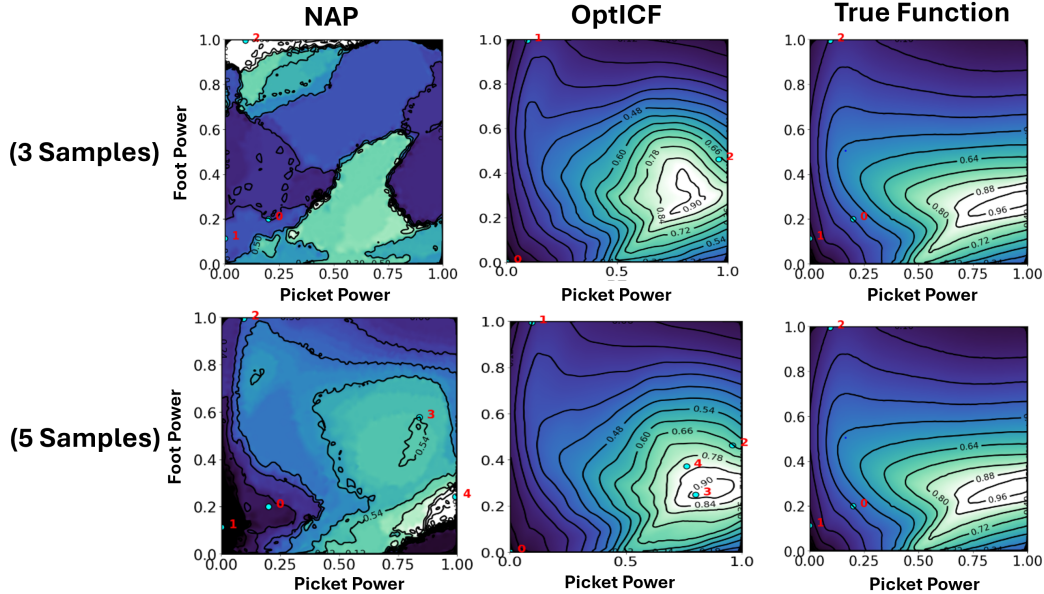


Figure 4: MBO’s meta-learned surrogate prediction against NAP’s surrogate prediction. MBO’s predictions are more accurate, smoother and interpretable.

exploitative, focusing on regions predicted to outperform the current best observation. In contrast, UCB is more *explorative*, as it explicitly incorporates 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 Experimental Details

For all experiments involving MBO, we employed the autoregressive variant of Transformer Neural Processes (TNPs-A) [6]. The model was configured using the hyperparameters recommended in the original work (Table 1). Optimization was performed using Adam [35] 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 80% GPU utilization.

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.

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

Figure 5: 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.

[36]. The resulting laser power and timing profiles served as inputs to LILAC [37], 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 [38].

To create the response surface, we varied two front-end parameters of the pulse shape representation, generating 50,000 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 neutron yields 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, specifically, the equation of state, which governs shock propagation in the fuel material and alters the response surface. These physics variations mimic the mismatch between simulations and real experiments, enabling us to evaluate our method under realistic domain-shift conditions.