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

Ricardo Luna Gutierrez, Vineet Gundecha, Rahman Ejaz, Varchas Gopalaswamy, Riccardo Betti, Sahand Ghorbanpour, Aarne Lees, Soumyendu Sarkar

Hewlett Packard Enterprise, University of Rochester

Motivation

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.

The efficiency of energy release in ICF 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.

❑Challenges:

- ICF experiments are extremely expensive, limiting researchers to very few experimental shots (low sample budget).
- The fusion gain function is a complex, non-linear black box that must be optimized under strict budgetary constraints.
- Standard Bayesian Optimization (BO) begins with no prior knowledge, wasting previous historical or simulation data.

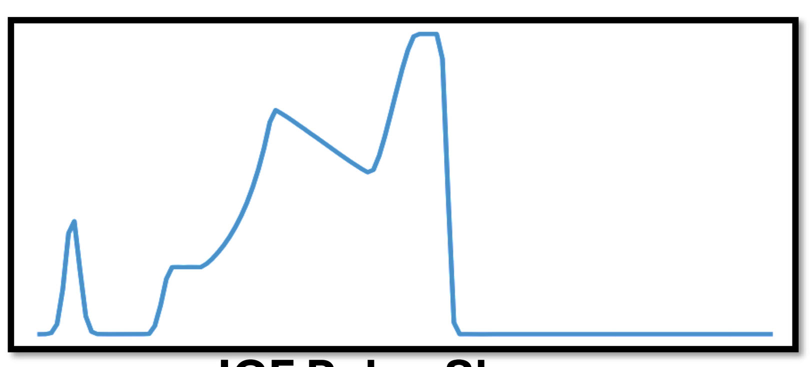

ICF Pulse Shape

Surrogate Quick Adaptation

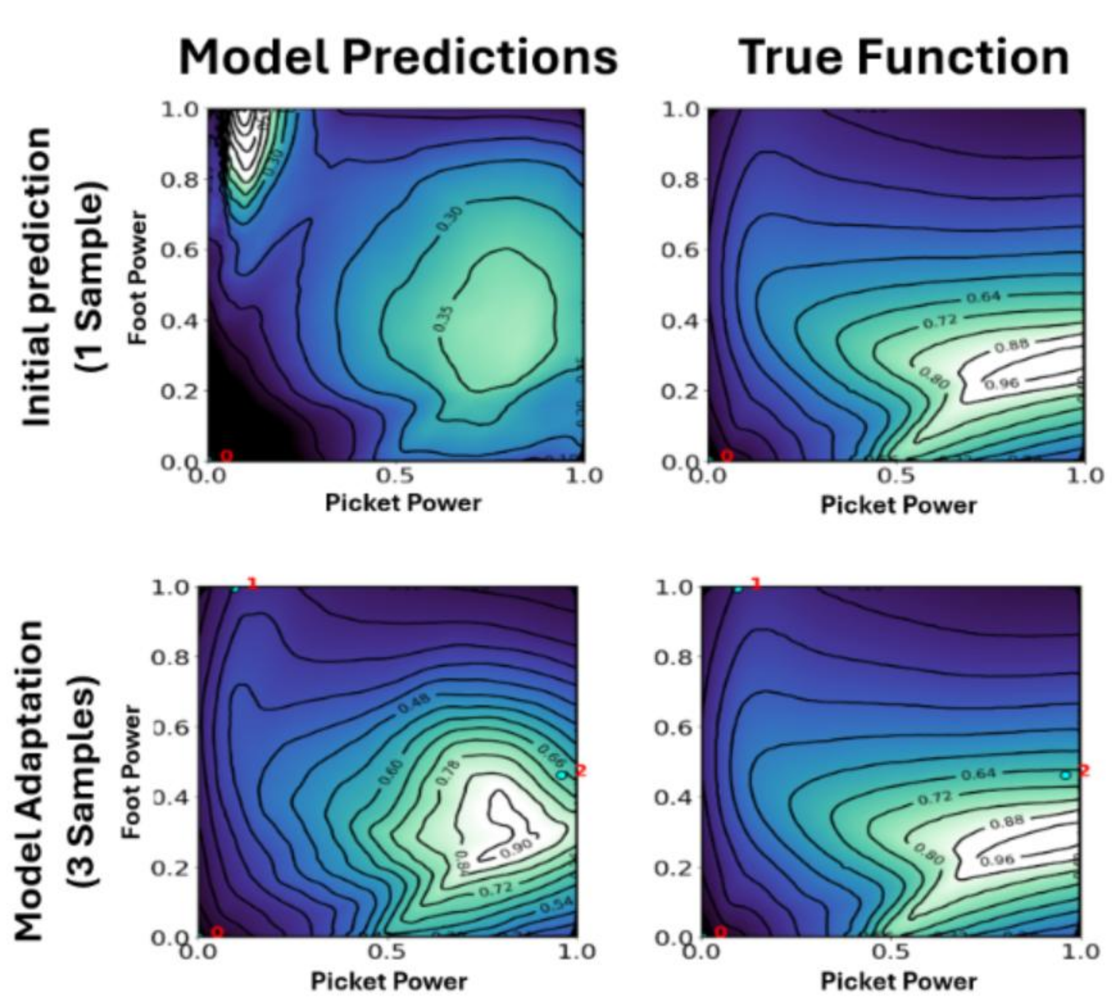

MBO's meta-learned surrogate predictions: (top) without only one context point; (bottom) after three context points; (right) the optimization target function. We can observe that our approach achieves quick adaptation with high sample efficiency, closely approximating the true function in just three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

Results

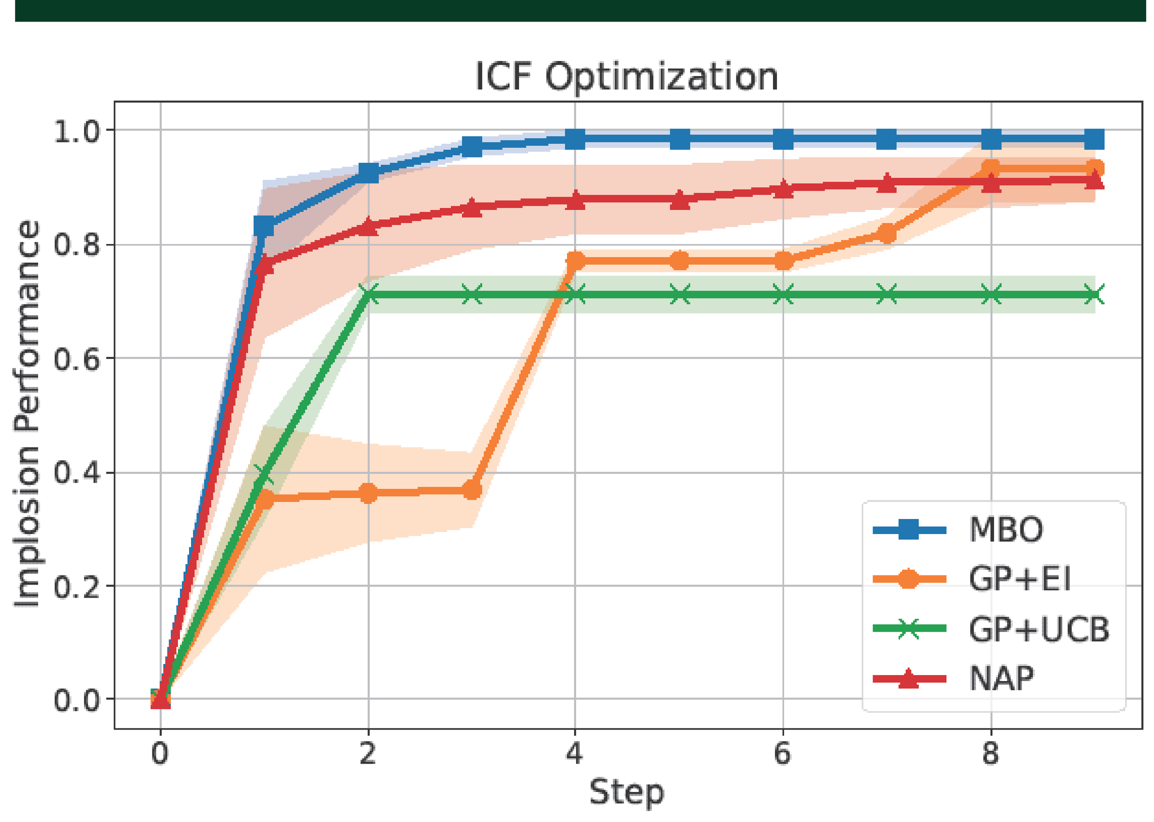

Comparison between BO methods and MBO in ICF. The shaded areas represent a ±1standard error. Performance is normalized.

Our Solution: Meta-Bayesian Optimization

We propose to leverage Meta-Bayesian Optimization (Meta-BO) to overcome the cold-start problem. The key is to leverage vast datasets from low-fidelity simulation tasks or previous experiments (source tasks) to construct a highly informed prior for the optimization of the expensive high-fidelity experimental task (target task).

Meta-Learning Knowledge Transfer

Source Task Generation

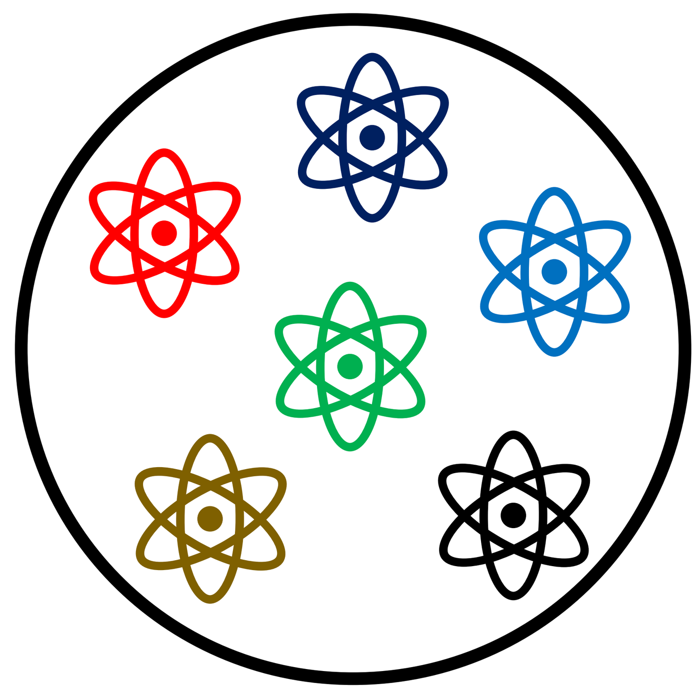

We generate a multitude of Source Tasks by systematically varying parameters of the physics-based simulator.

Meta-Training

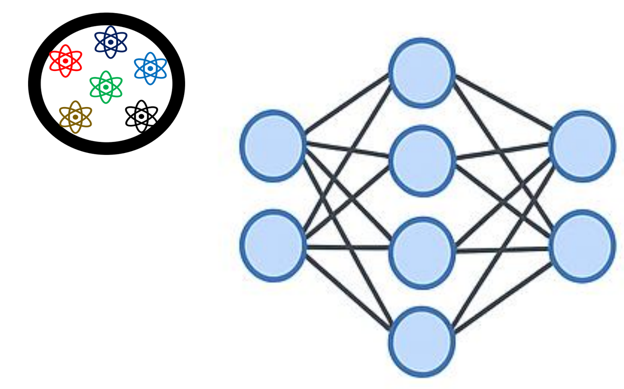

Transformer Neural Processes (TNPs) are meta-trained across this distribution of source tasks. The model learns the meta-knowledge of what an "ICF gain curve" generally looks like across different physics regimes.

Transfer Informed Prior

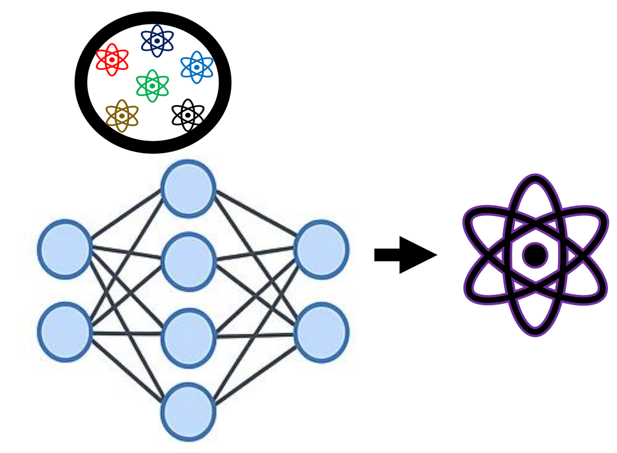

The learned TNPs are deployed for the Target Task. This pre-trained state acts as a highly informed functional prior, drastically reducing the initial uncertainty in the search space and enabling intelligent exploration.

Enhanced Optimization

Our framework integrates advanced optimization strategies using the TNP's predictive distribution:

- **Interpretable Candidates:** We provide intuitive visual representations that convey both the predicted outputs of the surrogate model and the associated uncertain-ties in those predictions. Additionally, to facilitate 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.

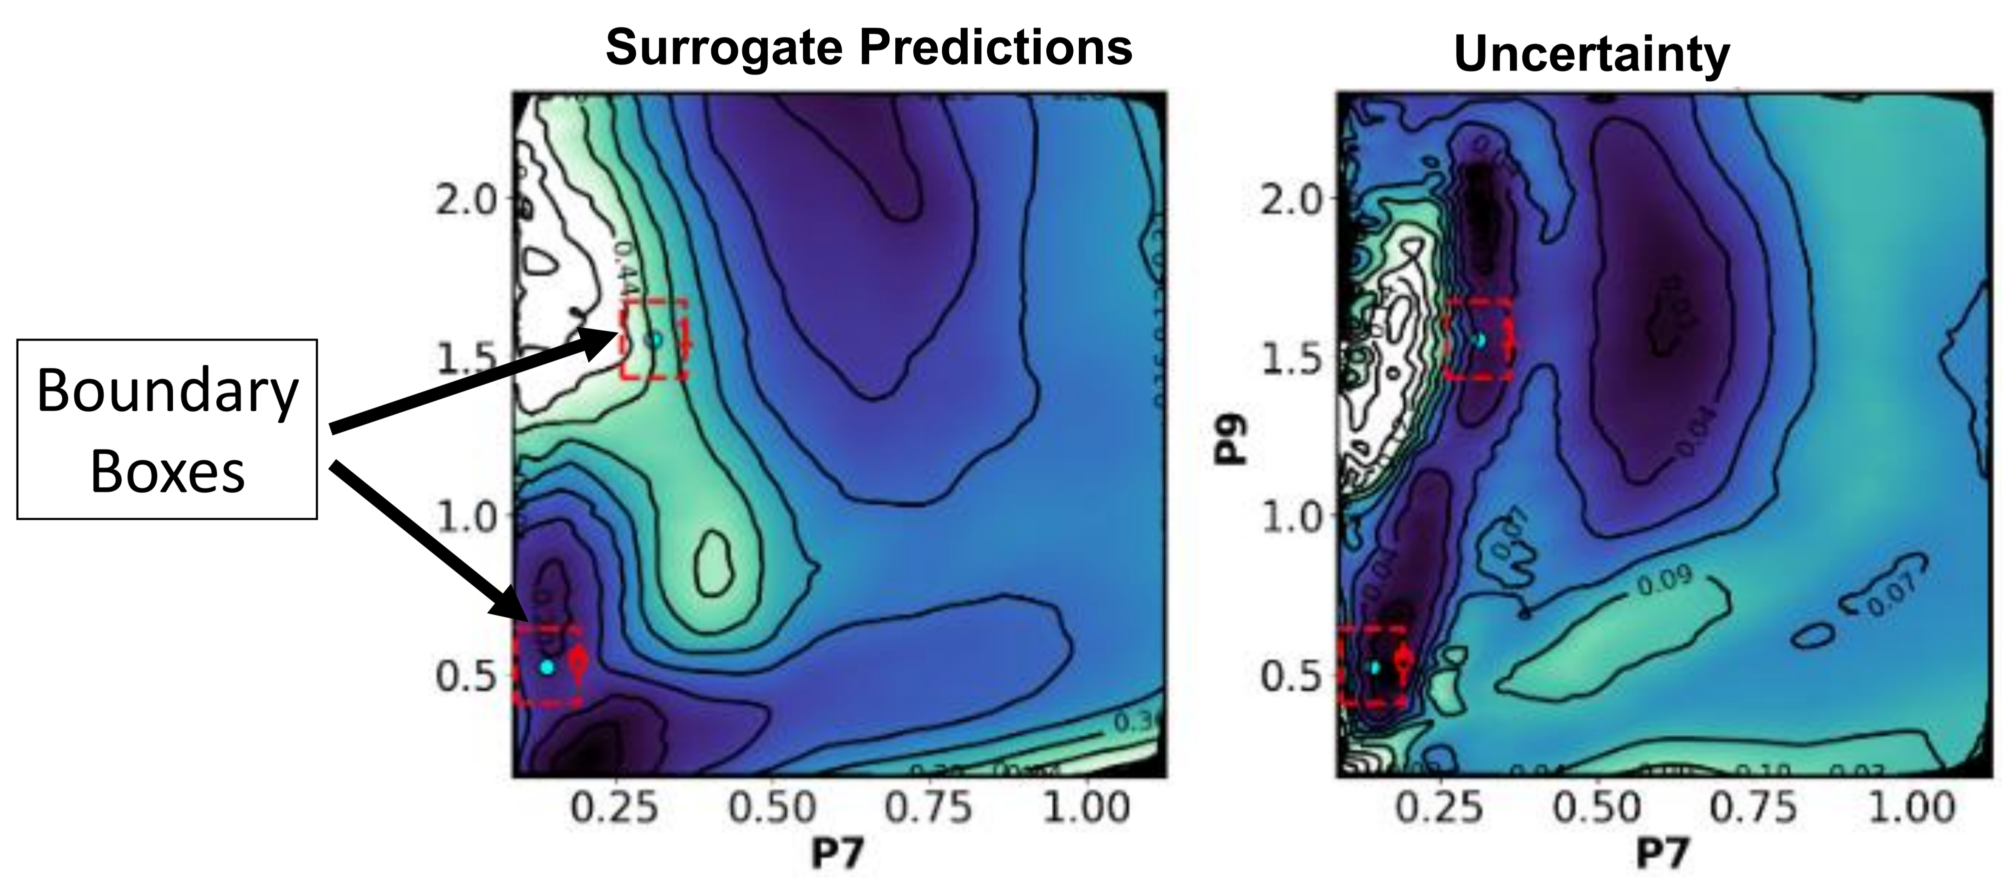

- **Boundary-Box Constraints:** To avoid redundant trials, which would otherwise be a waste of valuable resources, we introduce boundary boxes to a standard AFs. Instead of allowing the AF to propose points anywhere in the search space, we define and mask regions around previously sampled points, making them "invisible" to the AF and thus preventing re-sampling.
- **Dual Acquisition Strategy:** 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 we use EI and UCB. To ensure these recommendations are meaningfully different from one another, we apply our boundary box mechanism between the two proposed candidates.