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], Avisek Naug[1], Paolo Faraboschi[1], Soumyendu Sarkar[1]

[1]Hewlett Packard Labs @ HPE
[2]University of Rochester

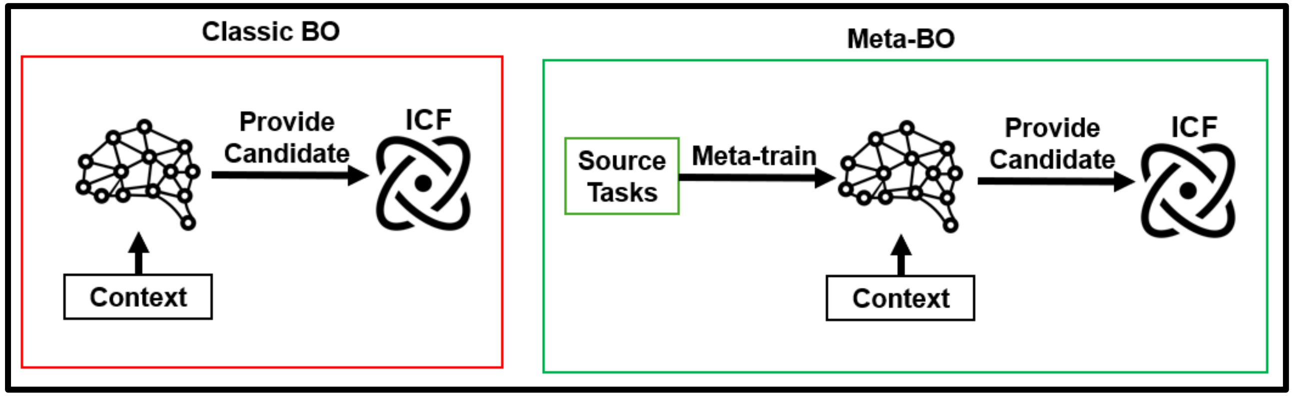

This material is based upon work supported by the Department of Energy National under Award No. DE-SC0024408, Hewlett Packard Enterprise and the University of Rochester.

Motivation

The Rising Global Energy Demand

- The global energy demand is rapidly rising.
- More than 80% of the world's energy supply comes from fossil fuels.

The Promise of Nuclear Fusion

- Nuclear Fusion: Offers potential for limitless, clean energy.
- **Inertial Confinement Fusion (ICF):**
  - ICF achieves nuclear fusion by subjecting a tiny fuel pellet to immense temperatures and pressures, typically using lasers.
  - The goal is to maximize the energy produced in a shot, by optimizing the parameters which form a pulse shape.
  - **Challenges**:
    - High costs and extreme conditions for experiments.
    - Limited number of facilities equipped to conduct ICF experiments.
    - Experiments are very sparse.

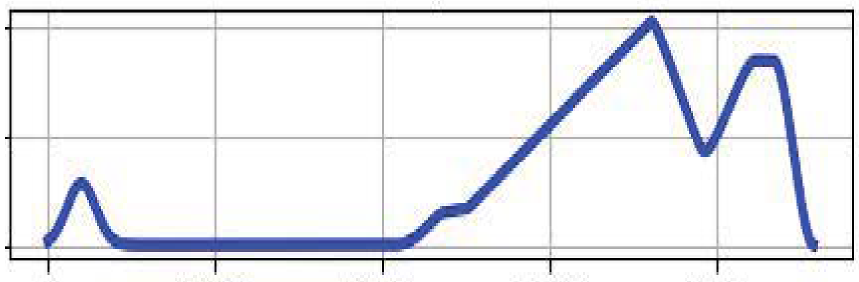

ICF Pulse Shape

Bayesian Optimization

Goal: Find the global optimum of a costly black-box function over a given input domain.

BO operates in two main steps:

Surrogate Modeling:

- Use previous evaluation data to fit a probabilistic surrogate model of the true black-box function.
- The surrogate model allows for probabilistic predictions for unobserved inputs.

Acquisition Function Optimization:

- Optimize an Acquisition Function (AF) that balances:
  - Exploration (uncertain regions in the search space)
  - Exploitation (promising areas based on previous evaluations)
- Selects next query point for evaluation.

Classical BO techniques start the optimization from scratch for new functions, without utilizing knowledge from related black-box functions.

Meta Bayesian Optimization

Improves optimization on new, unseen target black-box functions by using knowledge from related source functions (tasks).

Source Knowledge:

- Assumes access to related source tasks or to datasets collected from evaluations on them.

Meta-Learning Targets:

- Meta Bayesian Optimization uses this prior knowledge to meta-learn a surrogate model, an acquisition function, or both.

Meta Bayesian Optimization

- **MetaBO** (Volpp et al., 2020).
- **Rank-Weighted Gaussian Process Ensemble (RGPE)** (Feurer et al., 2022).
- **Neural Acquisition Process (NAP)** (Maraval et al., 2023).

MetaBO

MetaBO uses Reinforcement Learning (RL) to meta-learn an adaptive acquisition function that can generalize across related tasks.

The acquisition function is replaced with a **neural network**, which can identify and exploit **structural properties** shared across a class of functions.

- **Meta-Training Phase**:
  - Source tasks are used to evaluate the acquisition function, using GP as surrogate model.
  - Using **Proximal Policy Optimization (PPO)**, MetaBO trains the neural network-based acquisition function.
- **Target Task Optimization**:
  - The learned acquisition function is applied to the target task just like classical BO.

MetaBO Training

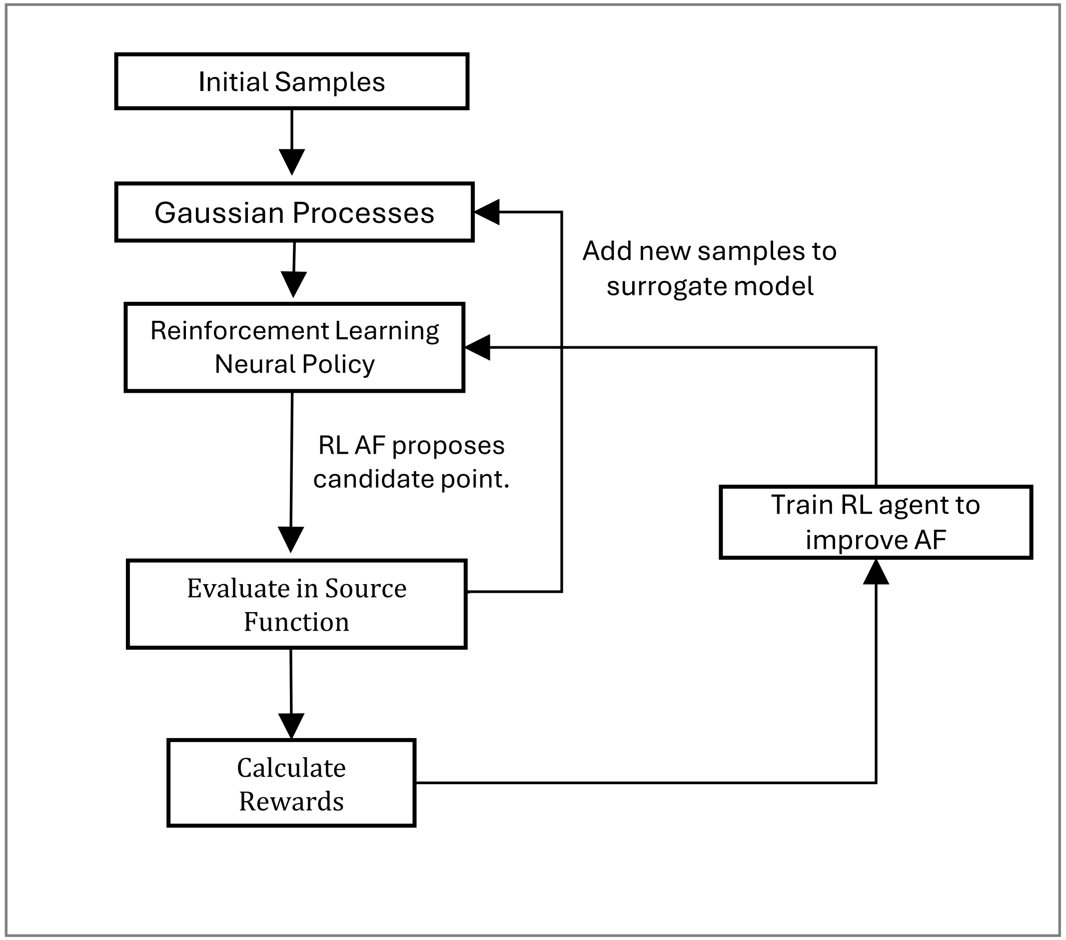

Rank-Weighted Gaussian Process Ensemble (RGPE)

RGPE is designed to optimize a target task by utilizing previously solved source tasks that share similarities with it.

Rank-Weighted Ensemble of Source Tasks:

- RGPE takes samples from source tasks and builds a separate **Gaussian Process (GP) model** for each task.
- For the target task, RGPE then **weights each GP model based on how relevant the task is** to the new target task.
- RGPE combines the weighted GPs from source tasks with a GP fitted directly on the target task data.
- This ensemble GP represents a mixture of both **source and target data**, prioritizing information from the source tasks that most closely align with the target.

RGPE Process

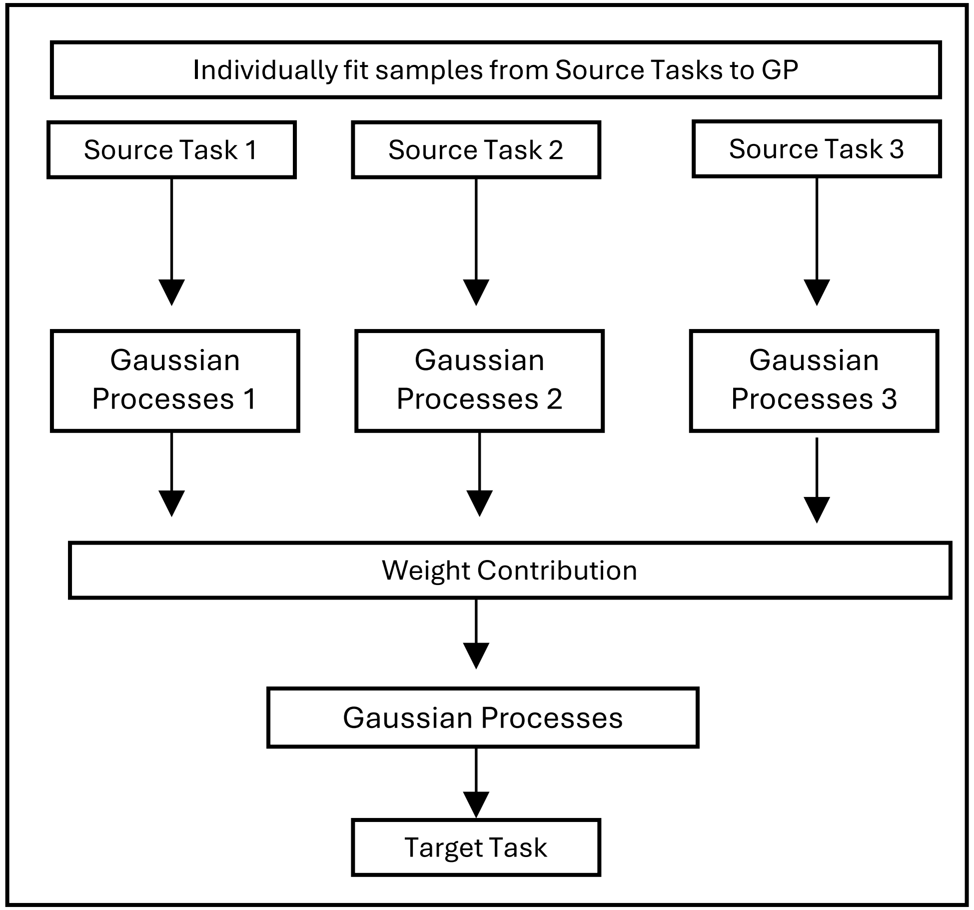

Neural Acquisition Processes (NAP)

NAP jointly meta-learns both the surrogate model and the acquisition function.

- Leverages **Transformer Neural Processes**:
    - Combines the flexibility and performance of transformers with properties of stochastic processes.

- **Training Process**:
    - Proximal Policy Optimization (PPO) is used for training.
    - An auxiliary Supervised Loss is used to introduce inductive bias, enhancing the optimization performance

NAP Training

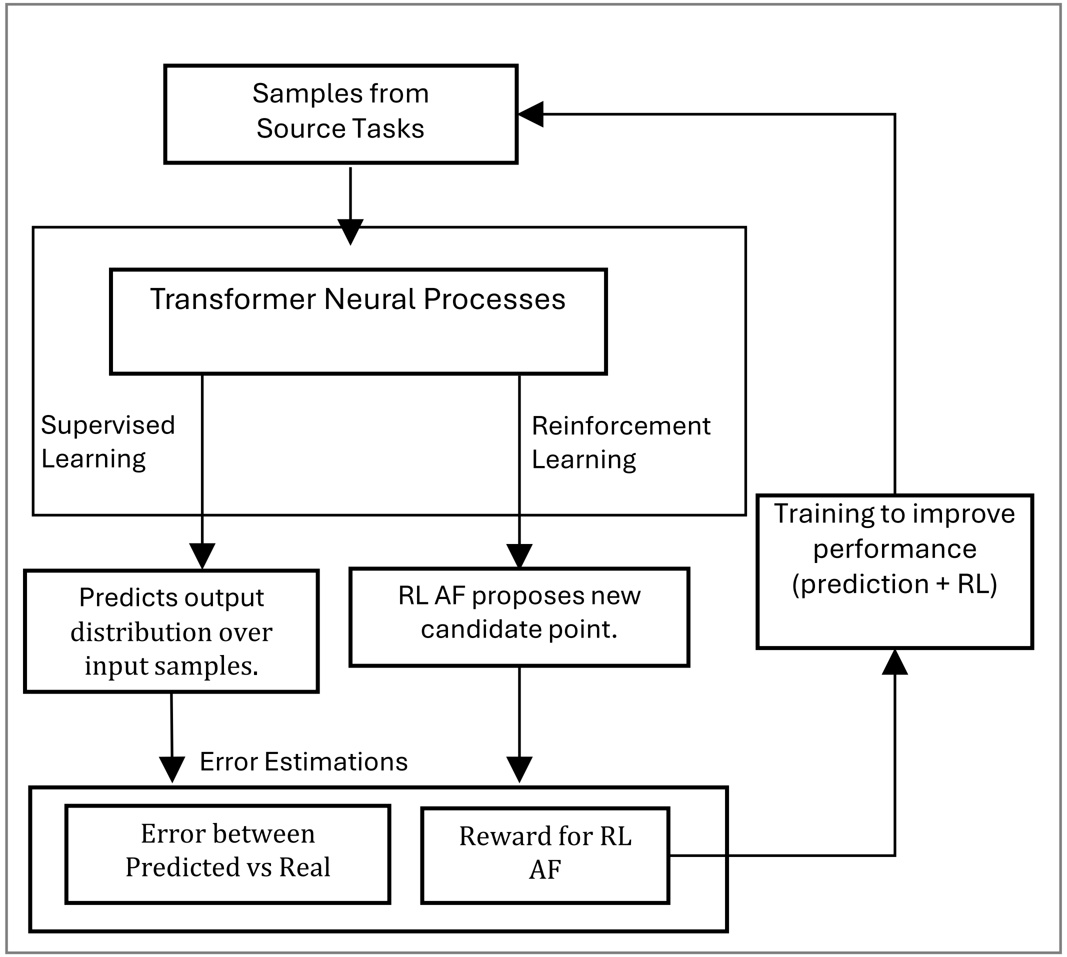

Meta Bayesian Optimization for ICF

Climate Change AI

NEURAL INFORMATION PROCESSING SYSTEMS

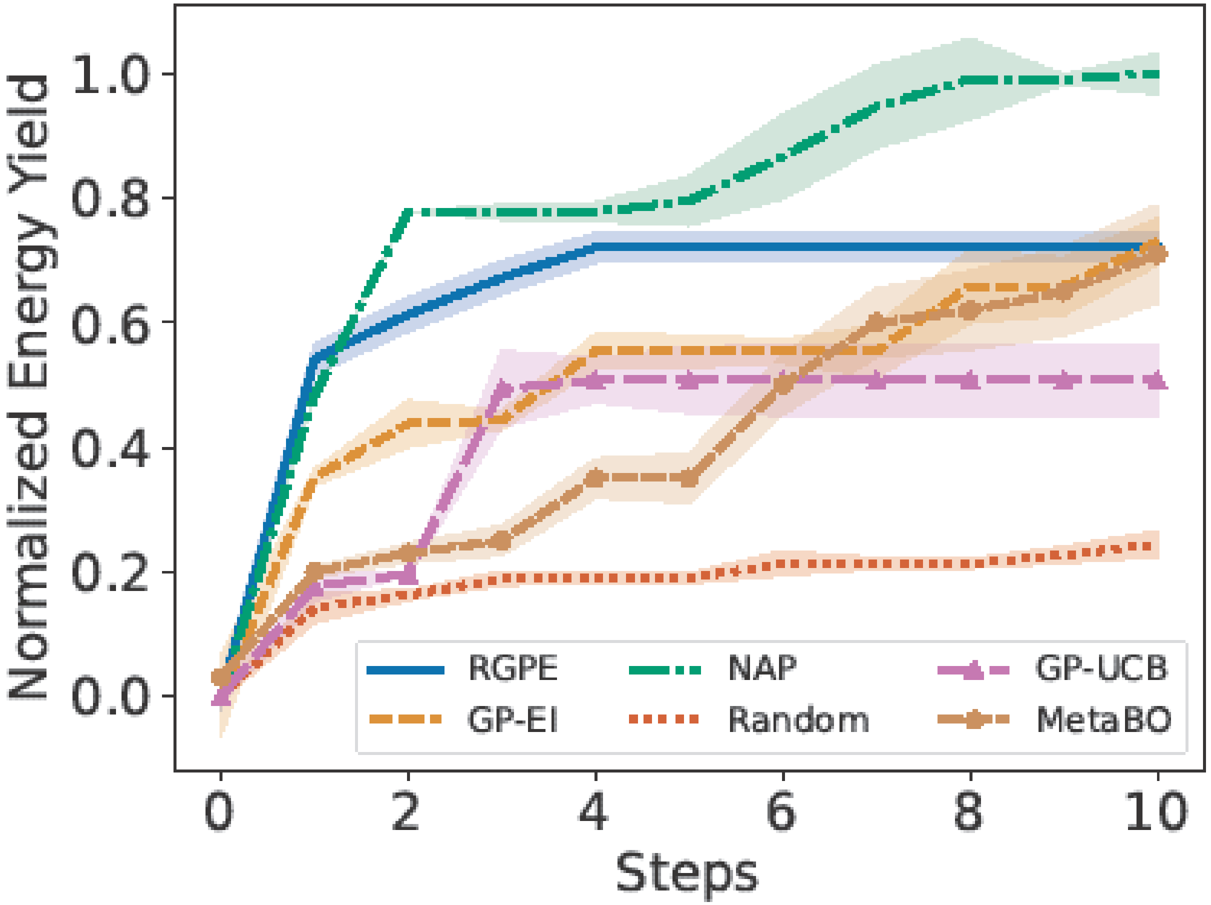

Comparison between classic BO methods (EI and UCB) and Meta-BO methods in ICF. NAP can achieve optimal performance in few samples.

Meta Bayesian Optimization Adaptation

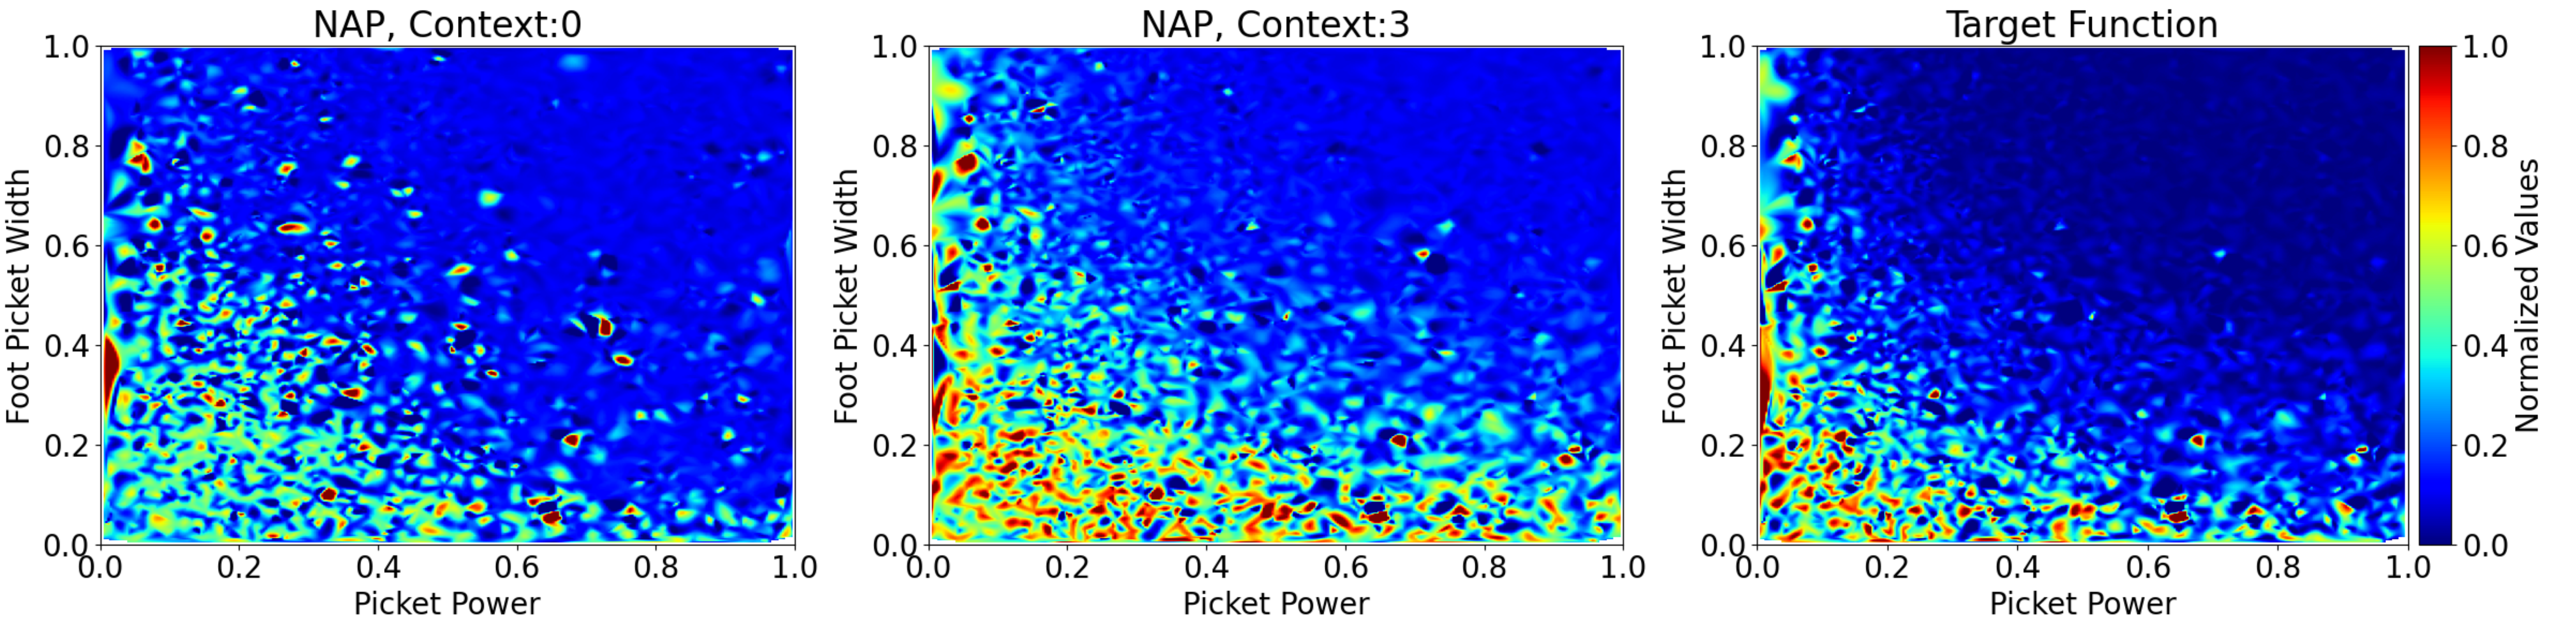

NAP's meta-learned surrogate predictions: (**left**) without any context points; (**middle**) after three context points; (**right**) target function. NAP achieves quick adaptation with high sample efficiency.

Thank you!

rluna@hpe.com