Multi-task Bayesian optimization of laser-plasma accelerators

Remi Lehe

Lawrence Berkeley National Laboratory, USA

March 6, 2024

Many thanks to the team!

DESY

A. Ferran Pousa

S. Jalas

M. Kirchen

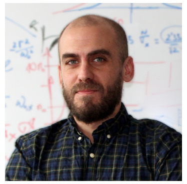

S. Martinez de la Ossa

M. Thévenet

BERKELEY LAB

J. L. Vay

A. Huebl

Argonne NATIONAL LABORATORY

S. Hudson

J. Larson

A. Ferran-Pousa et al., ***Bayesian optimization of laser-plasma accelerators assisted by reduced physical models****, PRAB (2023)*

Plasma-based acceleration: an emerging acceleration technology

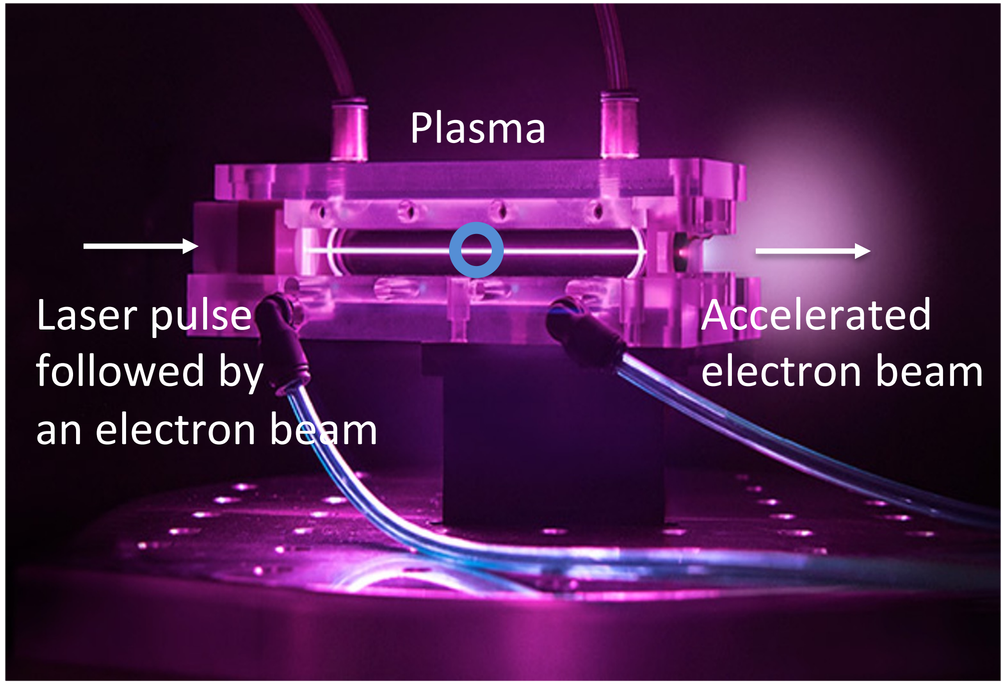

Main advantage:
~ 50 GeV/m accelerating gradient

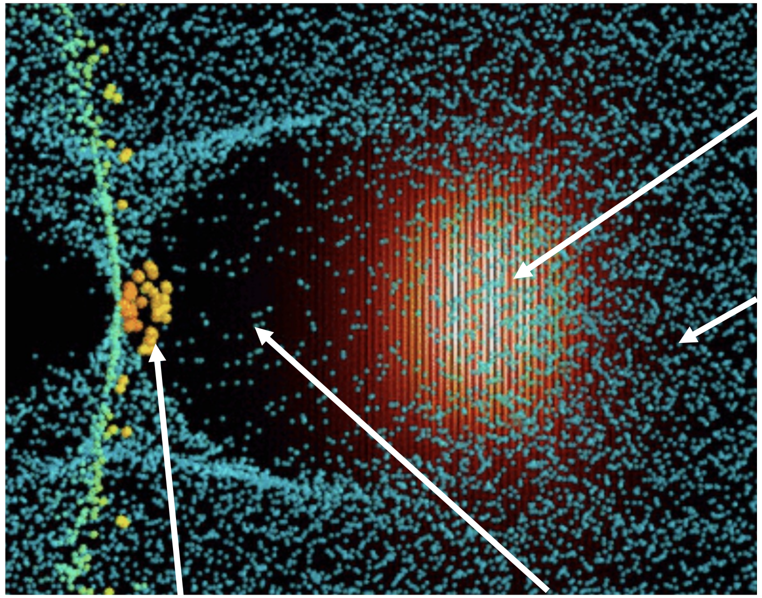

Can be modeled with Particle-In-Cell (PIC) simulations

Example of design optimization: beamloading in a plasma accelerator

- **Simulation-based design optimization** is a common workflow for plasma-based acceleartors

In:
current profile of the beam

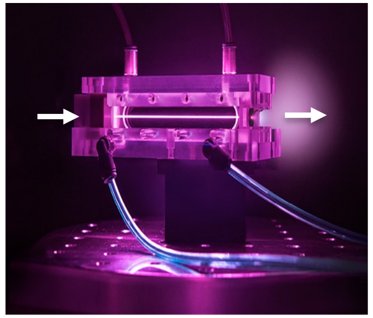

Out:
beam quality

- Example in this presentation: **beamloading optimization**
  Tune the **current profile** to maximize **beam quality** after acceleration
  - 4 input parameters that parametrize the initial **delay** and **current profile** of the beam
  - Single objective function: (quantifies **beam quality**)

$$f \propto \frac{Q \times E}{\sigma_E}$$

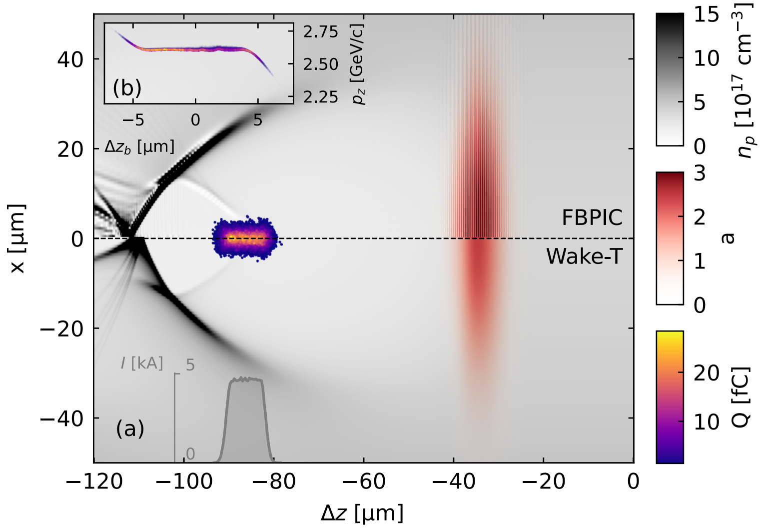

- **Bayesian optimization** is an efficient optimization method.
  Here we use **multi-fidelity Bayesian optimization.**

Multi-fidelity Bayesian optimization

Intuitive idea:

- inexpensive, **low-fidelity** simulations for **broad parameter exploration**
- expensive, **high-fidelity** simulations for few, well-targeted simulations

Different fidelities may mean:

- Different resolutions
- Different geometries (e.g. cylindrical vs full 3D)
- Different approximations (e.g. quasi-static vs full PIC)

Quantified by a fidelity parameter *s*, which is passed to the (modified) GP, along with ***x*** (tuning parameters), for each point

FBPIC ("high-fidelity", s=1)

- Full-PIC
- On the example setup: ~45 min per simulation, on 1 GPU

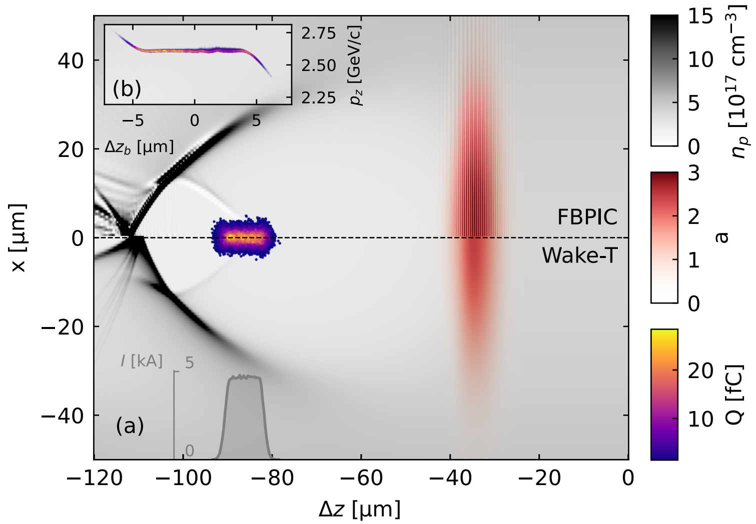

Wake-T ("low-fidelity", s=0)

- Quasistatic + laser envelope
- On the example setup: ~5 min per simulation, on 1 CPU core

See **Axel Huebl's presentation** on Thursday 5:40 PM regarding models with different fidelities

Multi-fidelity optimization requires a fidelity-aware Gaussian process

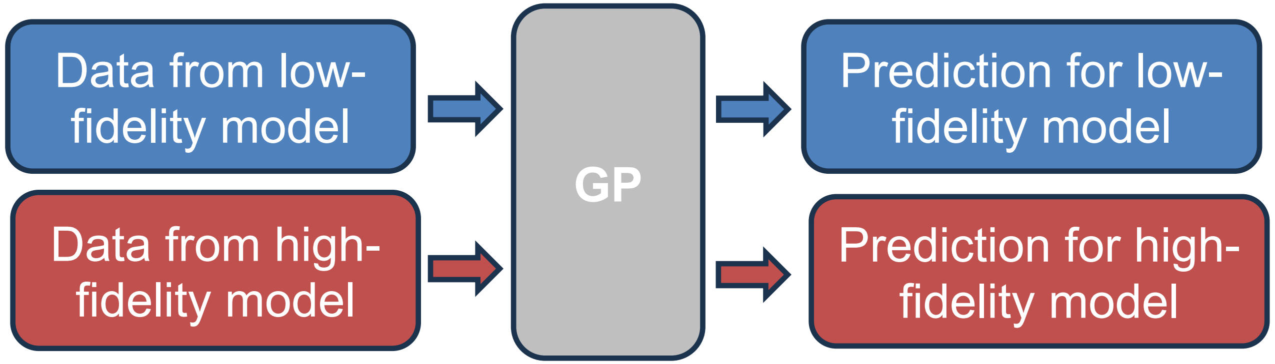

- Automatically evaluates **the level of correlation** between low-fidelity and high-fidelity data
- When strongly-correlated: can use low-fidelity data to **inform predictions** on high-fidelity data

Strongly-correlated case:

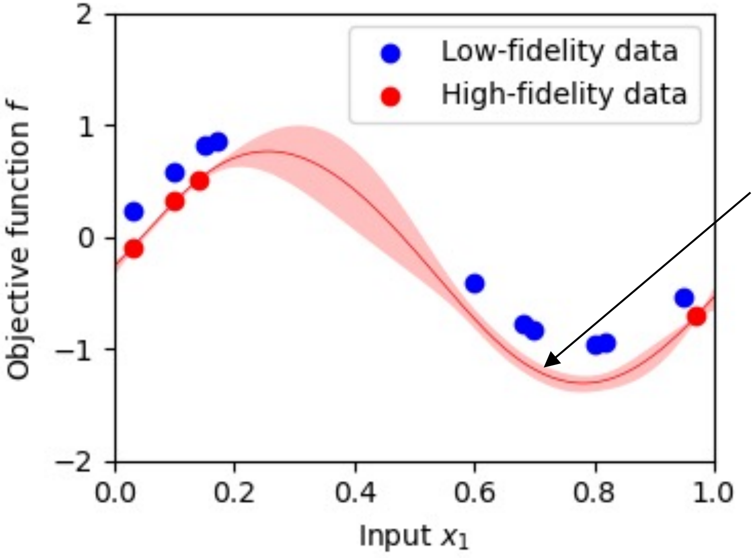

Low uncertainty, despite the absence of high-fidelity data

Un-correlated case:

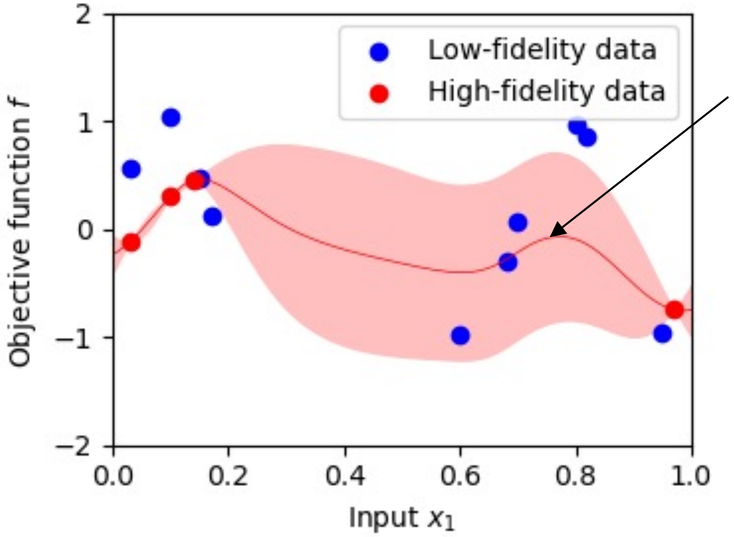

High uncertainty ; low-fidelity data is ignored

Multi-fidelity optimization requires a fidelity-aware Gaussian process

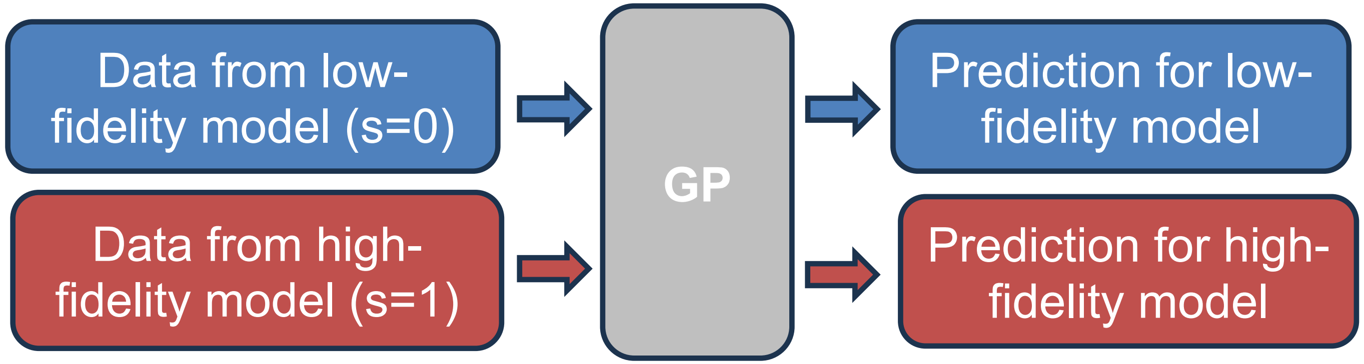

Strongly-correlated case:

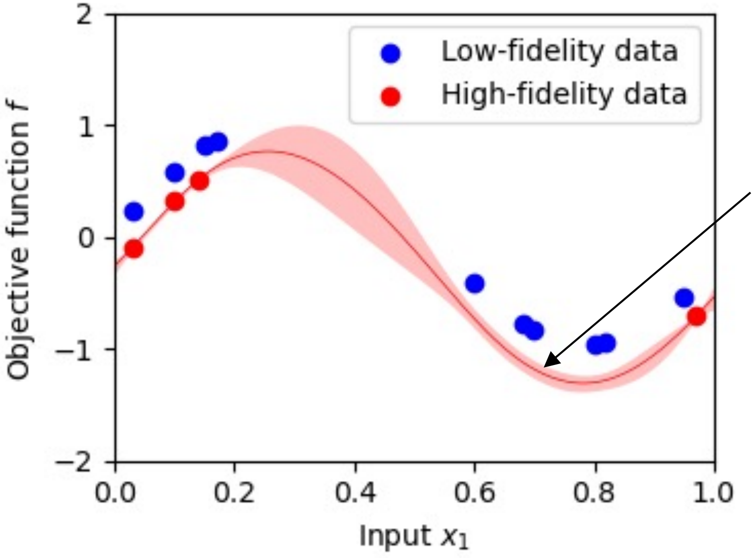

Low uncertainty, despite the absence of high-fidelity data

- The input space for the GP contains both **x** (tuning parameters) and *s* (fidelity)
- The kernel is usually assumed to be separable

$$k((s,\boldsymbol{x}),(s',\boldsymbol{x}')) = \tilde{\kappa}(||s-s'||)\kappa(\boldsymbol{x},\boldsymbol{x}')$$

- The **lengthscale hyperparameter** of $\tilde{\kappa}$ quantifies how correlated the different fidelities are, and is automatically tuned during **hyperparameter optimization**

e.g. *Bonilla, **Multi-task Gaussian Process Prediction**, NeurIPS (2007)*

The components of a multi-fidelity Bayesian optimization

- **Fidelity-aware Gaussian Process**
  e.g. *Bonilla, **Multi-task Gaussian Process Prediction**, NeurIPS (2007)*

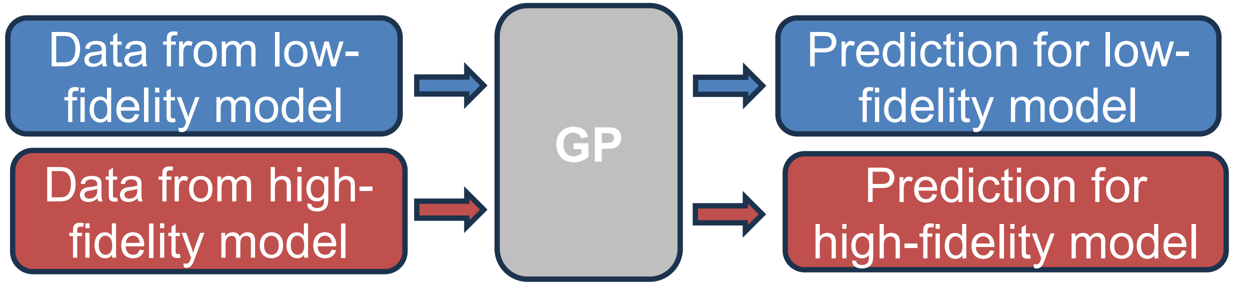

- **Procedure to suggest which points to evaluate (i.e. simulate) next, and at which fidelity**
  Recent work using multi-fidelity Bayesian optimization in plasma-based acceleration:
  - **Dynamic selection** of the fidelity for each simulation:
    *F. Irshad et al., **Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration**, PRR (2023)*
  - Batches of **fixed numbers** of low-fidelity and high-fidelity simulation (our work):
    *A. Ferran-Pousa et al., **Bayesian optimization of laser-plasma accelerators assisted by reduced physical models**, PRAB (2023)*

See also: *R. Roussel et al., **Bayesian Optimization Algorithms for Accelerator Physics**, arXiv:2312.05667 (2023)*

The multi-task Bayesian optimization algorithm alternatively runs high and low-fidelity simulations.

Algorithm:
Swersky et al., NeurIPS (2013)
Letham et al., arxiv 1904.01049 (2019)
(implemented in the Ax package)

- Repeatedly runs **fixed-size batches** of:
  - n_1 low-fidelity simulations
  - n_2 high-fidelity simulations
  (here n_1=90, n_2=3)

- Simulations are evaluated at the **most promising points** in parameter space, accordingly to the predictions of the GP prediction **for the high-fidelity**

- The model itself is updated after each batch

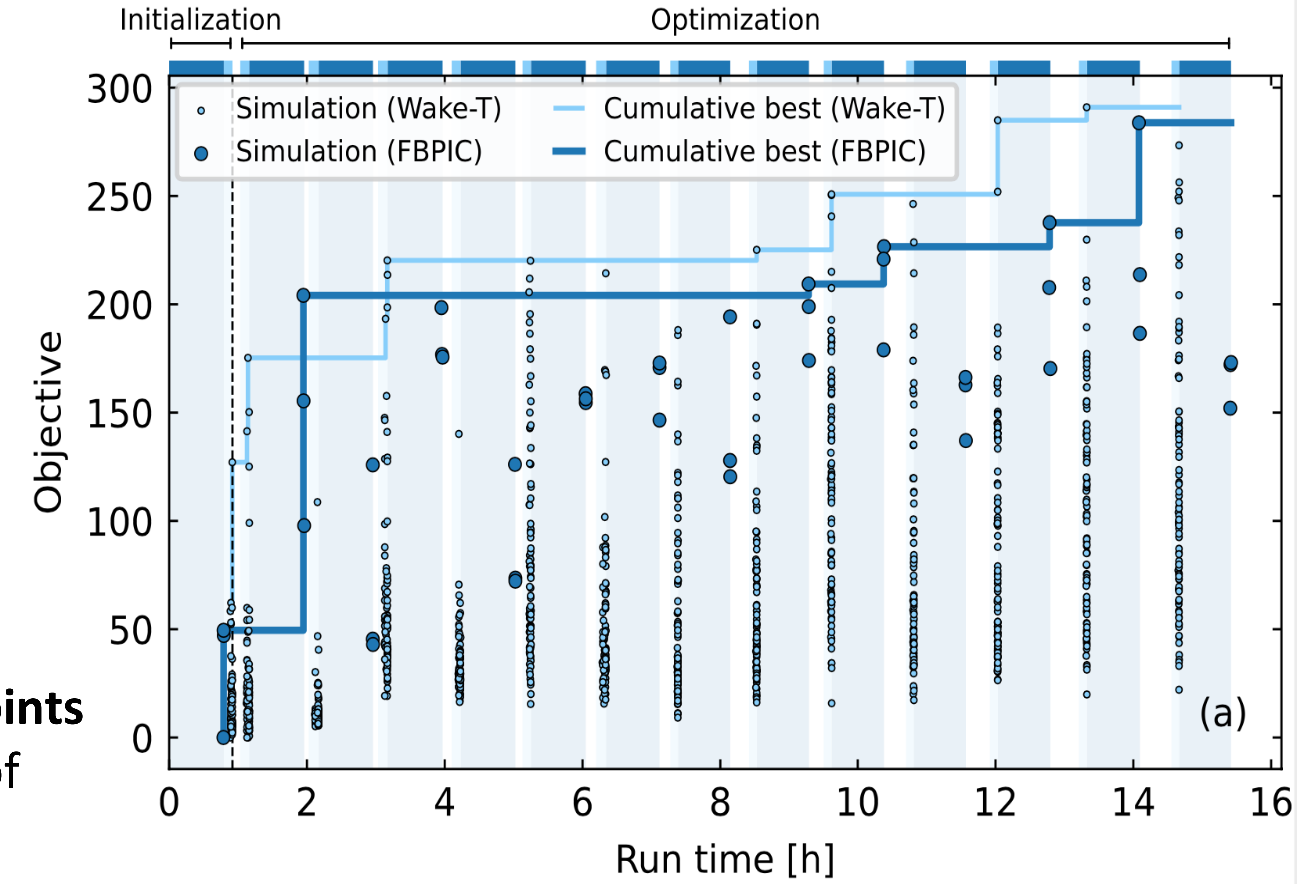

Orchestration of the simulations is done via `optimas`

`github.com/optimas-org/optimas`

Multi-fidelity Bayesian optimization can be significantly faster.

- 4 input parameters that parametrize the longitudinal beam profile, in order to optimize beamloading
- Objective function: $f \propto \frac{Q \times E}{\sigma_E}$

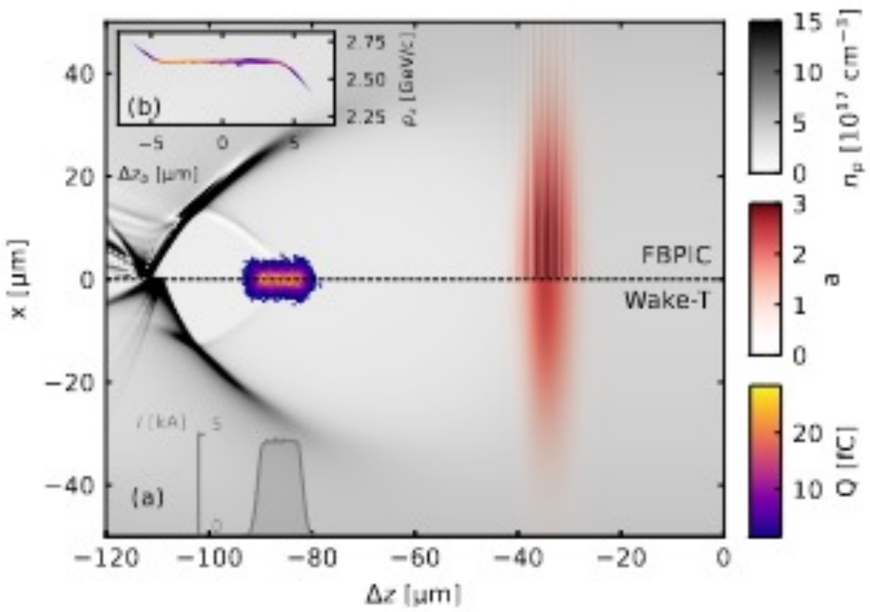

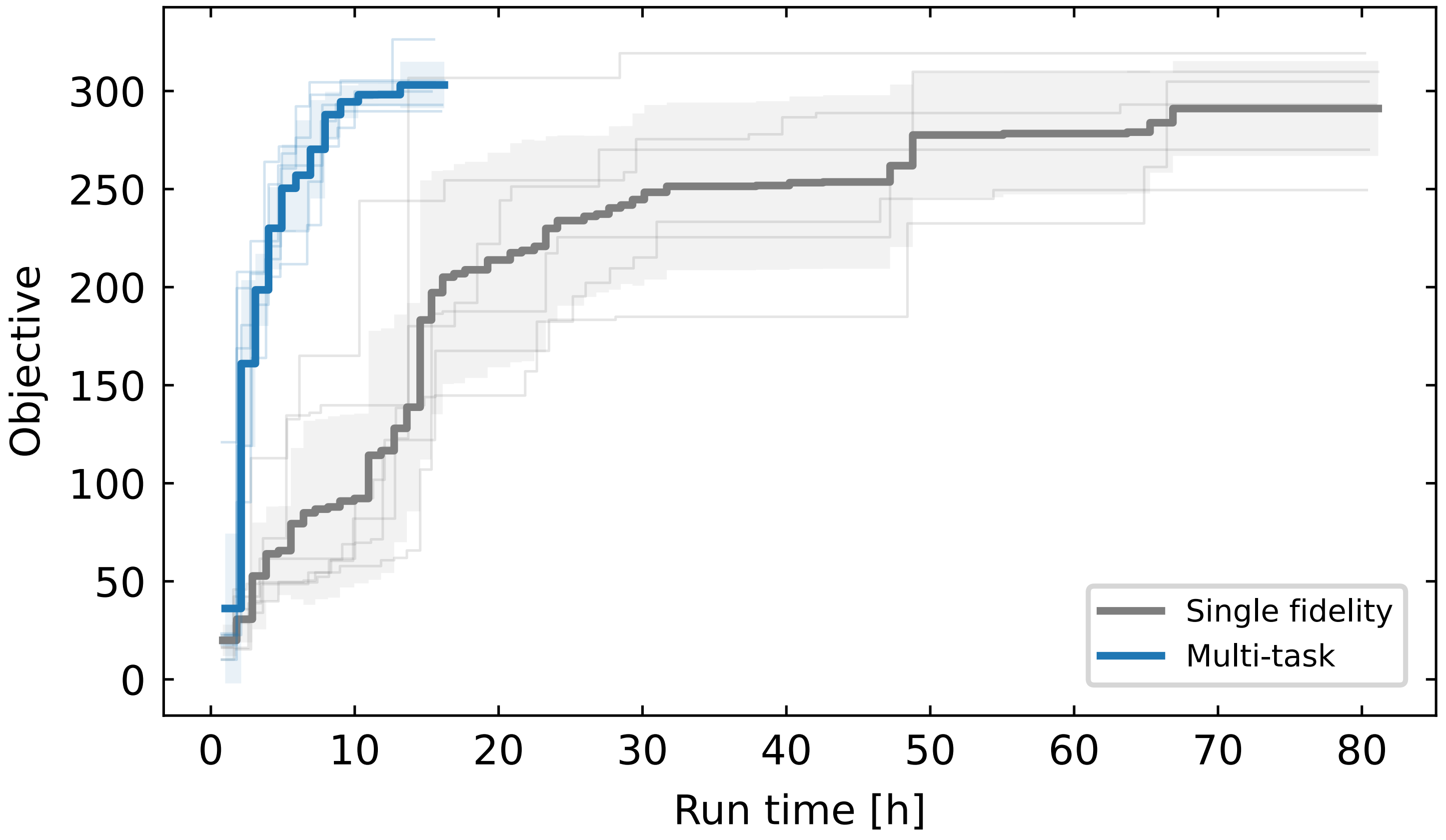

(6 independent run in each case, with different random seeds)

Optimization is robust when low-fidelity does not match high-fidelity.

High-fidelity:
FBPIC, fixed resolution

Low-fidelity:
Wake-T, varying resolution

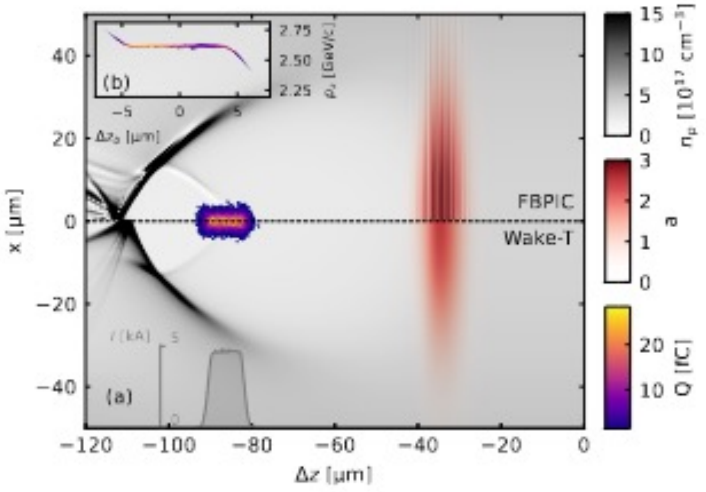

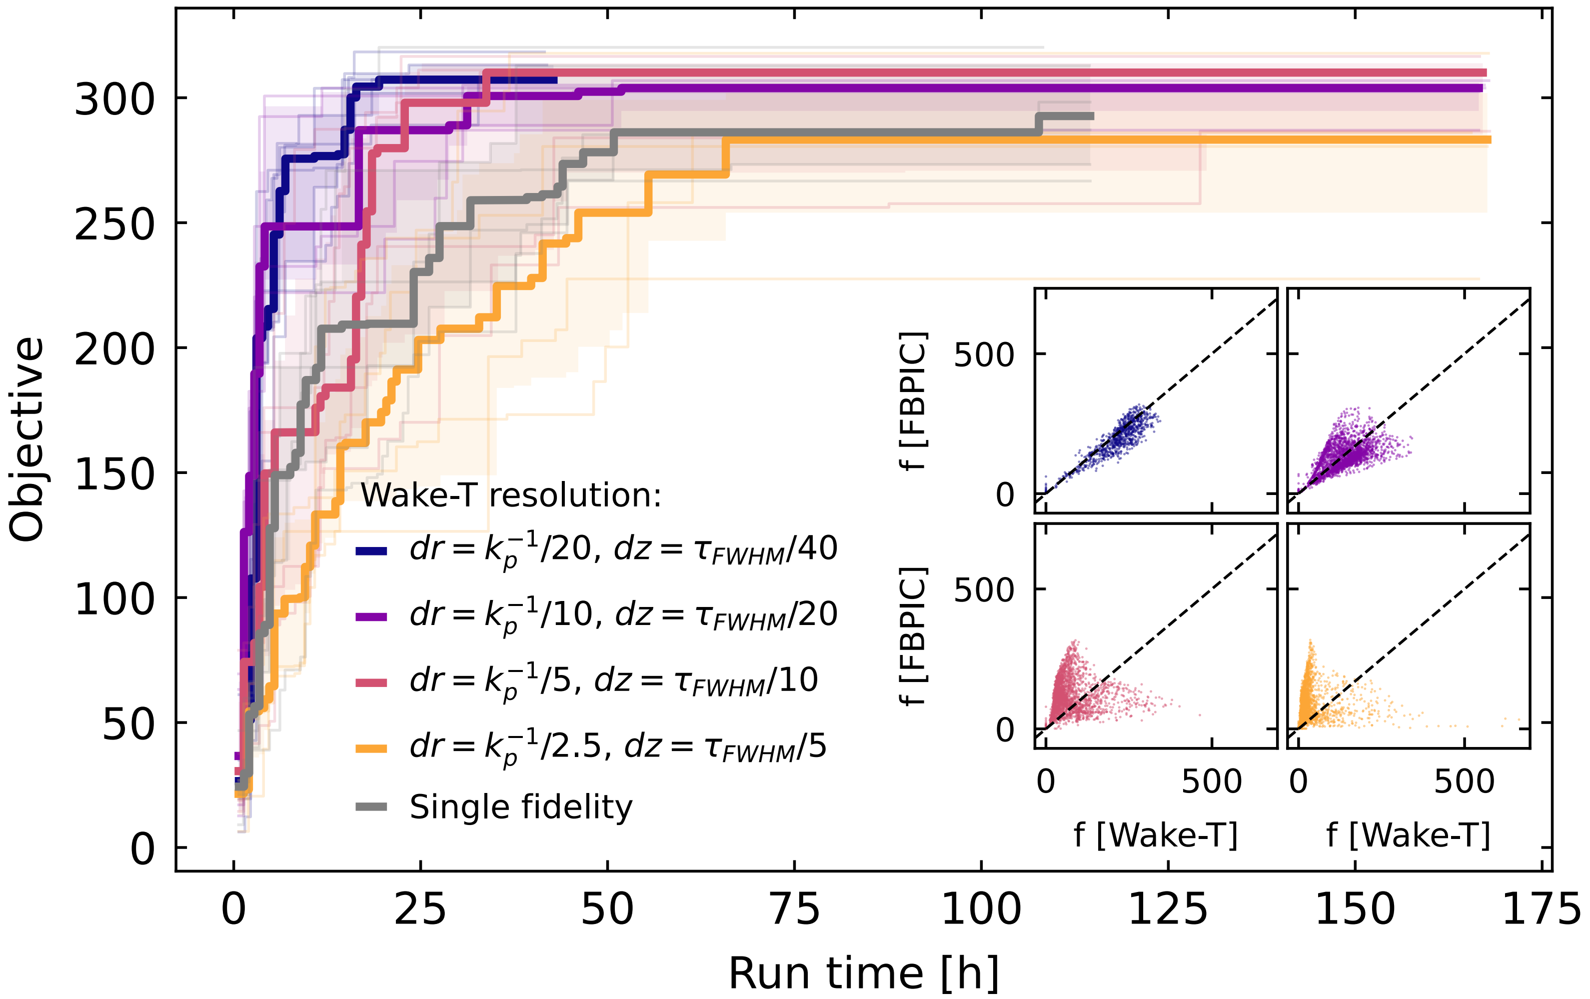

Caveat: generating next points can take a significant time

- Here: GP and acquisition function evaluated on a GPU (Ax implementation)
- Theoretical scaling for training the GP $\propto N^3$ where N is the number of datapoints collected so far

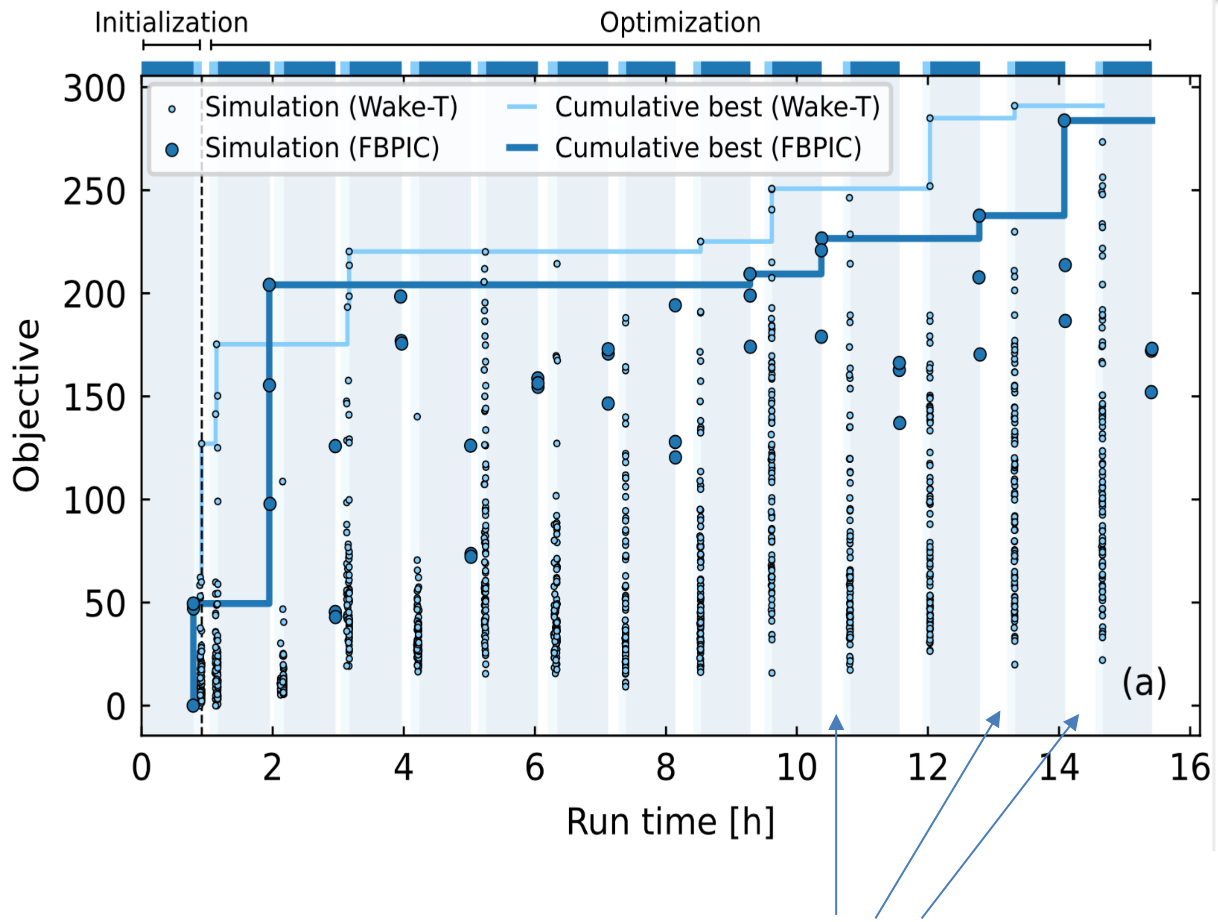

White gaps: time it takes for the algorithm to decide which points to evaluate in the next batch

Conclusion

- Multi-fidelity Bayesian optimization can make use of **inexpensive, low-fidelity** simulations to **accelerate optimization**
- Anecdotal evidence seems to indicate that it is **relatively robust** when low-fidelity and high-fidelity don't match
- Cost of fitting the GP and optimizing the acquisition function can be become a **bottleneck** when using a large amount of low-fidelity data

Next step: use multi-fidelity Bayesian optimization (and other ML models) for simulation-aided optimization of experiments

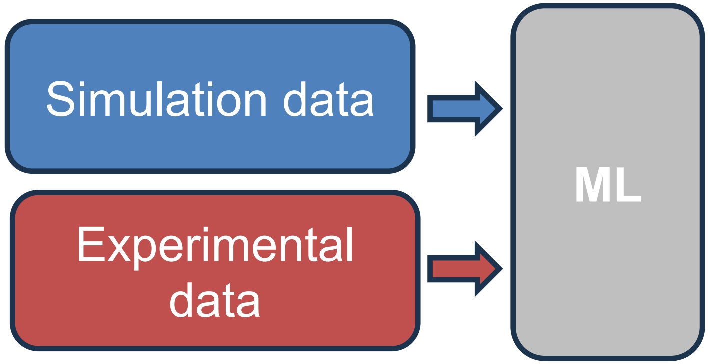

Recently received funding as part of **2024 LDRD at LBNL**

We welcome ideas and collaboration on **efficient combination** of experimental and simulation data!

Berkeley Lab Laser Accelerator (BELLA) Center
Experiments with Terawatt and Petawatt lasers

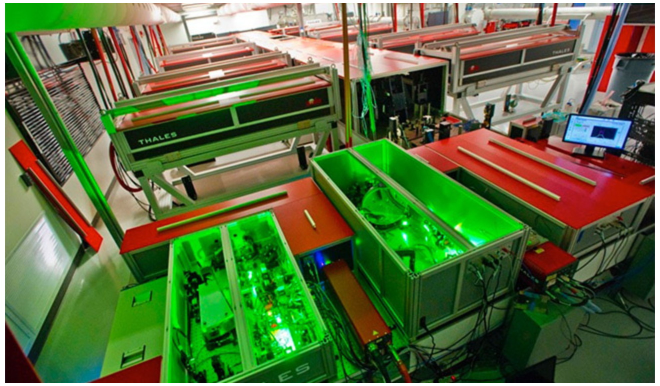

Accelerator Modeling Program
Modeling with large-scale Particle-In-Cell simulations

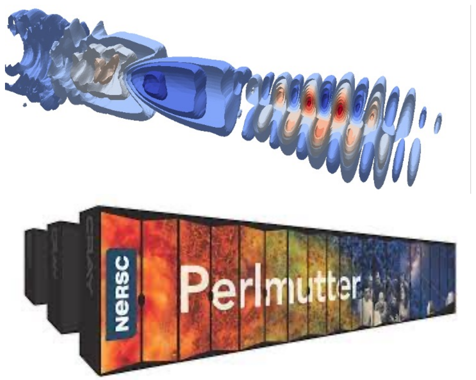

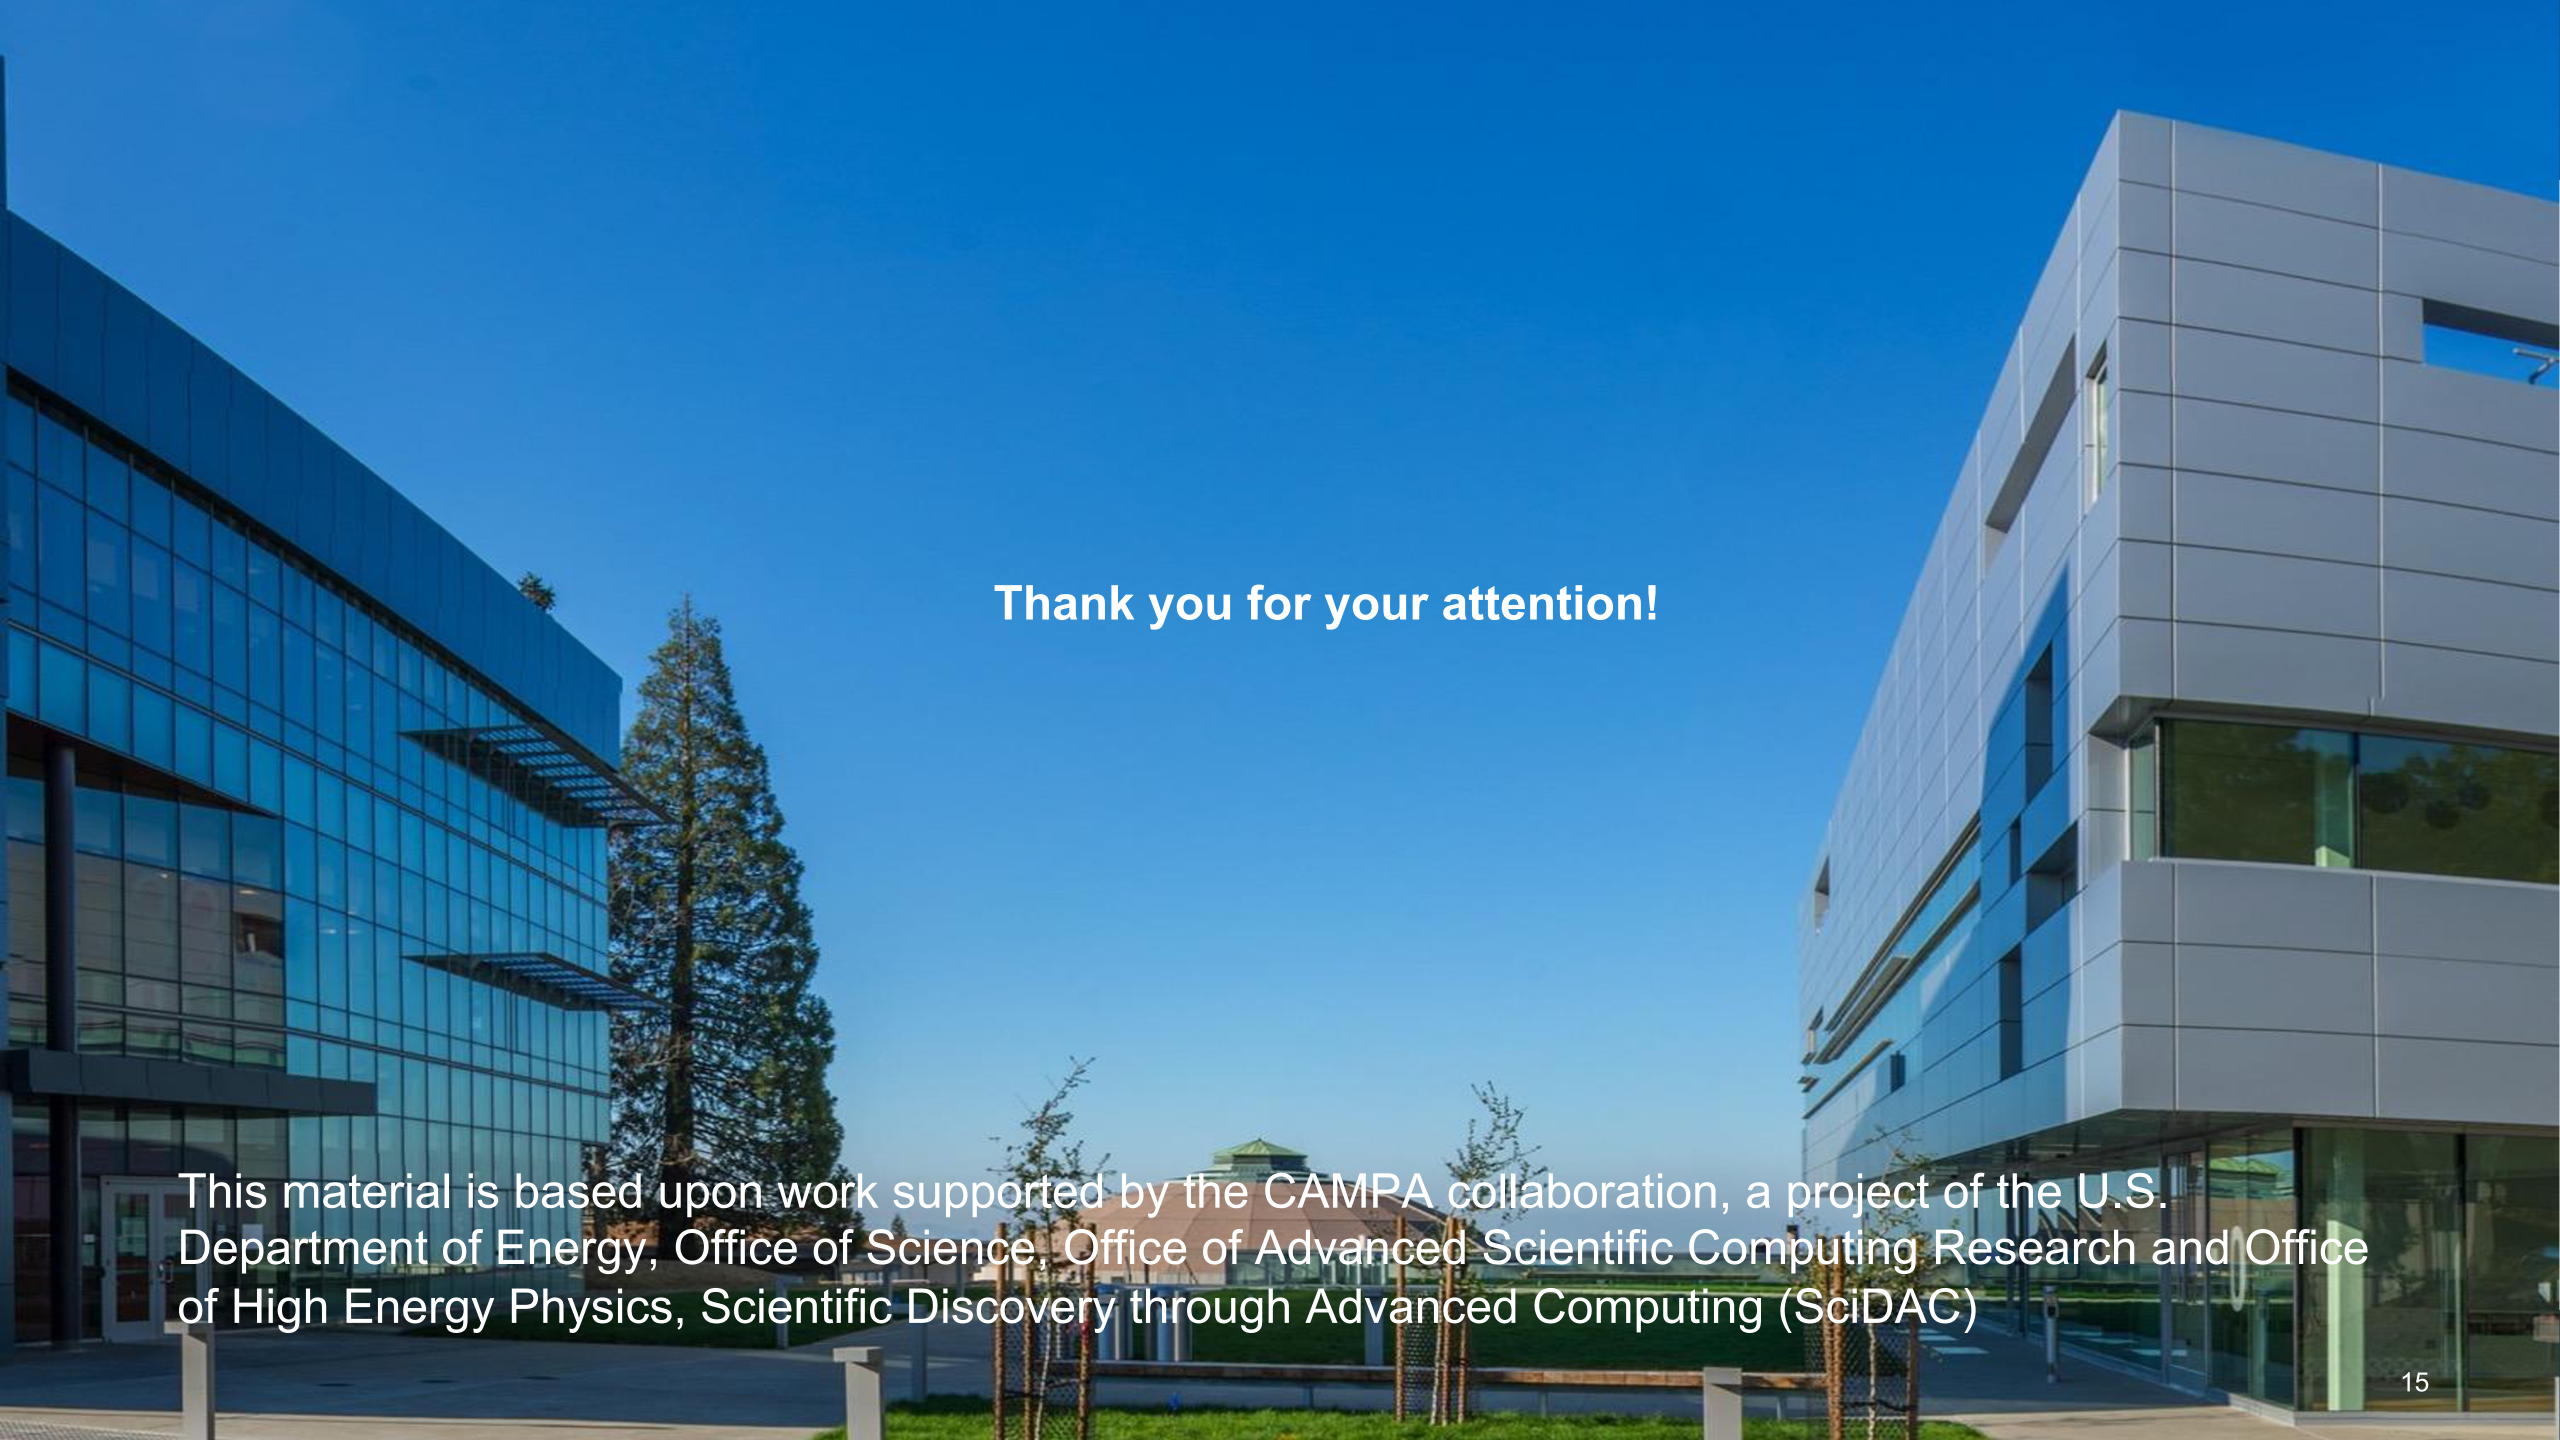

Thank you for your attention!

This material is based upon work supported by the CAMPA collaboration, a project of the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of High Energy Physics, Scientific Discovery through Advanced Computing (SciDAC)