Contribution ID: **194**

Type: **Oral**

Hybrid Bayesian Optimisation / Evolution Strategy applied to the design of uni-axially driven ICF targets

Wednesday, 29 November 2023 09:15 (35 minutes)

Inertial confinement fusion (ICF) relies on the implosion of precision engineered capsules containing DT fuel. The implosion is initiated by a driver, usually a laser, and the target may feature one or more outer shells to enable driver coupling. First Light Fusion's (FLF) novel approach separates the design of the target into a fuel capsule and a shock amplifier, which is uni-axially driven by a hyper-velocity projectile. The target design process seeks to maximise the neutron yield of the implosion. This is typically accomplished through optimising the design using radiation-hydrodynamics (rad-hydro) simulations and comparison with experiments. The optimisation process can require thousands of high-fidelity simulations, taking a large amount of compute resource.

For these reasons, numerical optimisation techniques are of key interest to the ICF community and several algorithms have previously been proposed. Bayesian Optimisation (BO) is a promising approach to these problems, as it excels at finding the global optimum of expensive black box objectives. Within the Bayesian optimisation loop, a machine learned model (or emulator) is constructed from an initial set of simulation runs and optimised in place of the actual function. This suggests a best set of design parameters; and evaluations at those parameters (made by running more rad-hydro simulations) are used to update the emulator. The process is repeated to rapidly locate the optimum. While extremely powerful, Bayesian Optimisation has extensive configuration options and effective use requires detailed knowledge of the problem being optimised.

In this work, we describe a comprehensive Bayesian Optimisation (BO) capability tailored for use in ICF target design problems. The BO routines are implemented using BoTorch, an open-source Python framework. Gaussian process (GP) models are used as emulators, which fit a normally distributed family of response surfaces to the simulation samples, the initial set of such being generated by an optimised Latin Hypercube space-filling design. GP models can struggle with noisy and discontinuous response surfaces common in ICF simulation. This is solved by using a data-learned transform on the objective. Asynchronous batch execution (using the "Kriging Believer"heuristic for conditioning the GP on pending inputs) is implemented, allowing full use of HPC compute resources. Black box constraints, commonly found in ICF problems, are handled through multiple output tasks (the objective and the constrained outputs) learned simultaneously and used to constrain the space of input parameters that are suggested by the BO algorithm.

The approach is benchmarked against a robust and commonly used optimisation algorithm, CMA-ES (Covariance Matrix Adaptation Evolution Strategy) and the EGO (Efficient Global Optimisation) Bayesian Optimisation algorithm available in the DAKOTA toolkit (developed by Sandia National Laboratory). Our approach outperforms both CMA-ES and EGO for optimising the performance of our amplifier design (a 9D problem) and on other synthetic benchmarks. Finally, as BO is well suited for locating global optima but can be slow to refine to the exact location, we develop an algorithm that combines Bayesian Optimisation with CMA-ES, called BOCMA, which outperforms both CMA-ES and BO in terms of the number of iterations required.

Speaker's Affiliation

First Light Fusion Ltd.

Member State or IGO/NGO

United Kingdom

Primary author: Dr EDWARDS, Thomas (First Light Fusion Ltd.)

Co-authors: Dr ALLISON, James (First Light Fusion Ltd.); Dr BORDAS, Rafel (First Light Fusion Ltd.)

Presenter: Dr EDWARDS, Thomas (First Light Fusion Ltd.)

Track Classification: AI