An Efficient Classifier-Based Surrogate Assisted Evolutionary Algorithm

Christopher M. Pierce and Ivan Bazarov

Abstract

Multi-objective genetic optimization is an important tool in the design and operation of modern particle accelerators. Unfortunately, the objective functions of interest to accelerator physicists are often both expensive to compute and don't benefit from large scale parallelization at the level of the individual. Here, we present our progress on developing a new highly efficient surrogate assisted evolutionary algorithm that, in a wide variety of test problems, is 10x faster than conventional genetic optimizers (NSGA-II). In contrast to some Bayesian approaches, this optimizer continues to offer parallelization at the population level. To achieve this speed-up, we show that a neural-network-based binary classifier can approximately model the Pareto-domination relationship among individuals in the optimization problem. We then investigate how this information can then be used to discriminate proposed individuals which are unlikely to improve the Pareto front and prevent the expensive objective functions from being evaluated on it. We present our progress in including constraints into the problem and show initial work on applying this algorithm to accelerator specific problems.

Beamline Optimization is Expensive

PHYSICAL REVIEW ACCELERATORS AND BEAMS **23**, 070101 (2020)

Low intrinsic emittance in modern photoinjector brightness

Christopher M. Pierce,* Matthew B. Andorf, Edmond Lu, Colwyn Gulliford, Ivan V. Bazarov, and Jared M. Maxson

Cornell Laboratory for Accelerator-based Sciences and Education, Cornell University, Ithaca, New York 14853, USA

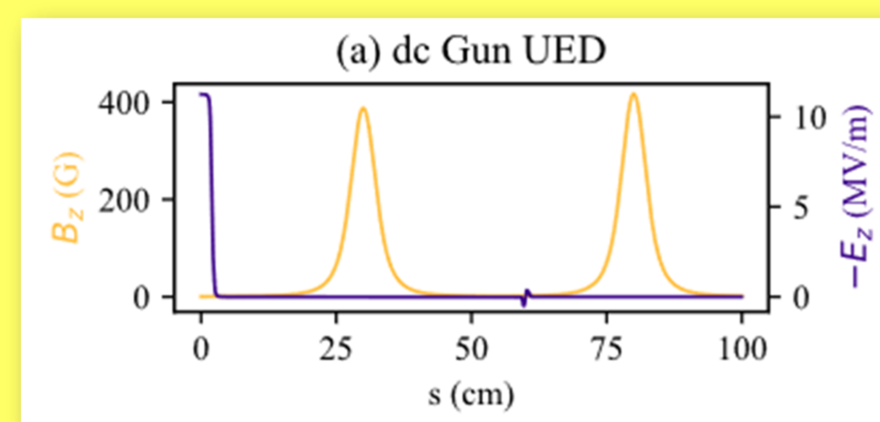

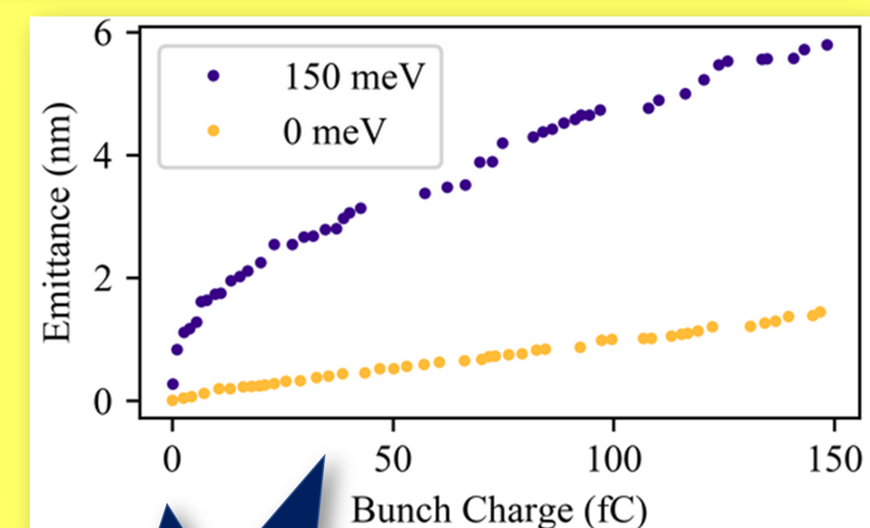

These optimizations involved a detailed space charge model that took 100k macroparticles to ensure accurate results

5 CPU Years of Compute Time!

Is there some way to reduce the workload?

A Domination Sorting Model

Domination: The solution A "dominates" solution B if it is at least as good as B in all objectives and is better than the solution for at least one objective.

Idea: Train a neural network based classifier to predict the comparison operator Objective(A) > Objective(B) for all objectives in the optimization. Use information from the network's predictions to avoid running your expensive simulation on individuals that aren't expected to improve the Pareto front.

Symmetric Classifier: The operator being modeled has symmetry such that if a > b is true then b > a is false (ignoring equality). We enforce it by using the same feature extraction layers on each input. Then the outputs of those layers are concatenated and fed into special symmetric dense layers that enforce the symmetry and predict the final probabilities.

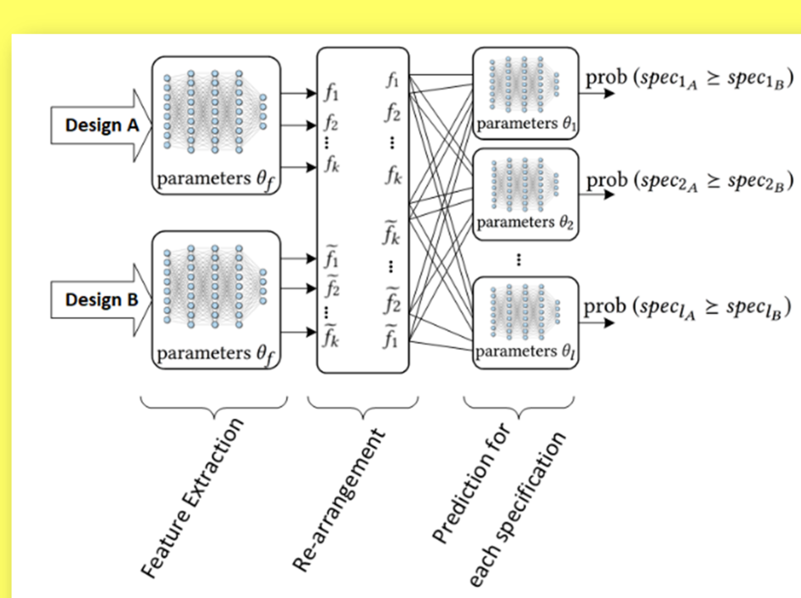

Adding Model Info to Optimizers

The algorithm NSGA-II is naturally extended with information about previous solutions from the model. After the proposed children are generated, perform a test domination sort using the surrogate model. Only accept children that will make it to the next population. This step can be repeated as many times as necessary to achieve a full population of children that the model says will improve the Pareto front.

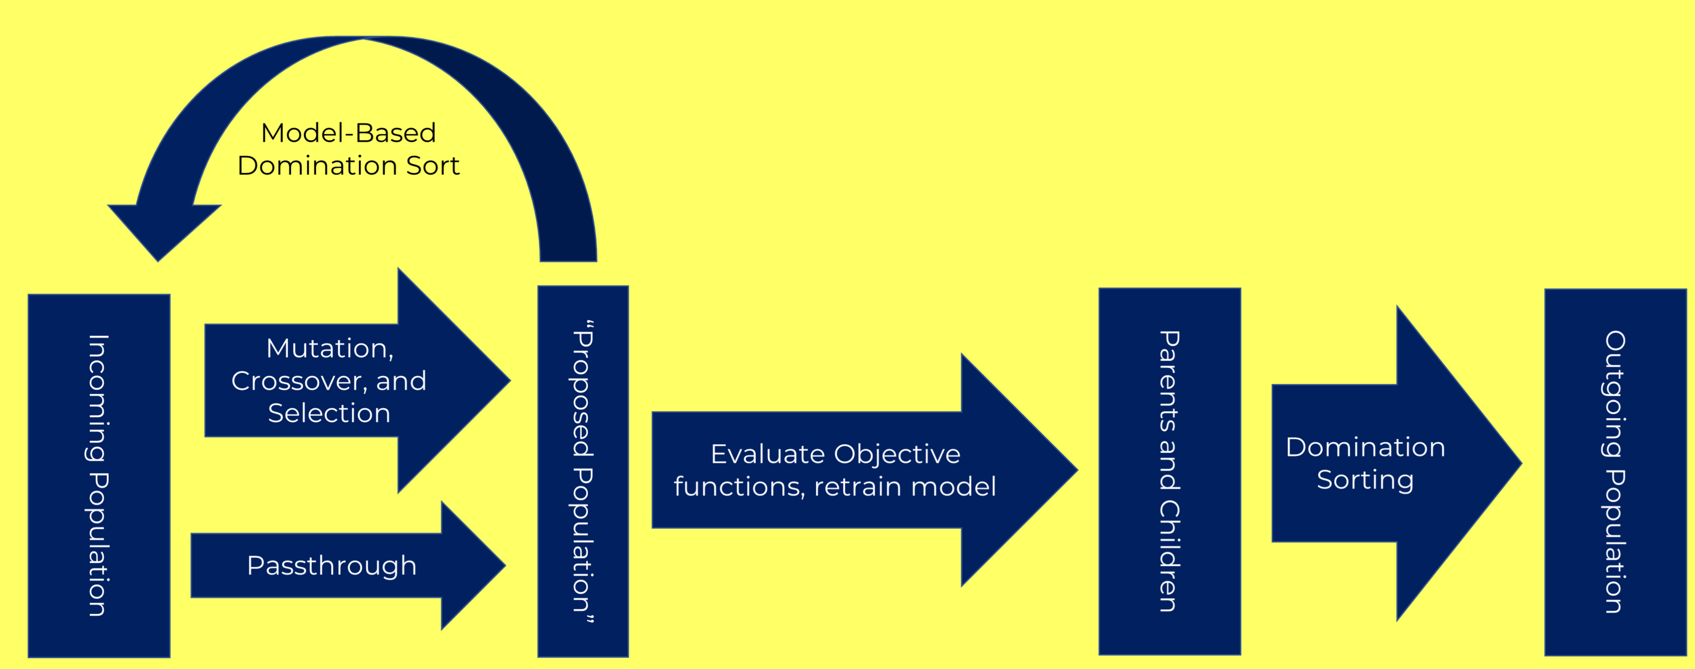

10x Speedup over Conventional Algorithm! (Orange == this work)

Comparison is against officially endorsed benchmark implementation of NSGA-II

Error in Solution (Smaller is Better)

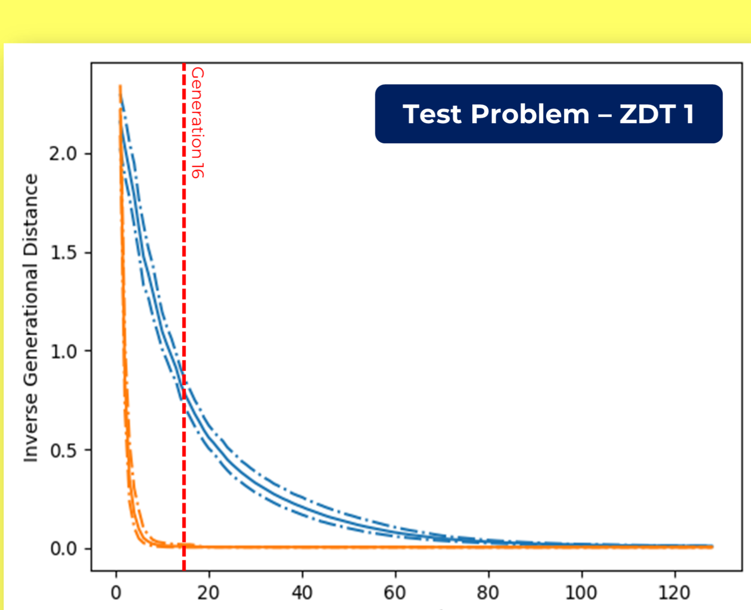

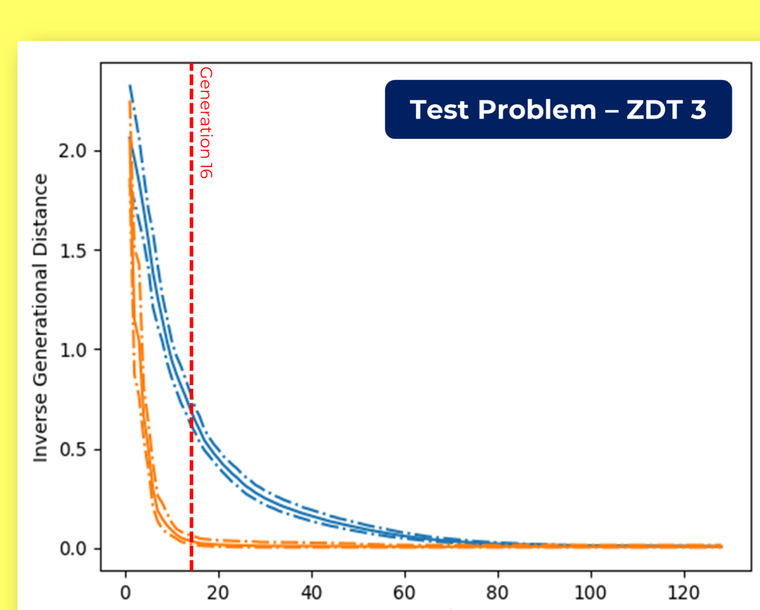

Number of Function Evaluation Batches (Smaller is Better)

Pareto Fronts Compared at Generation 16