HIGH-PRECISION REGRESSORS FOR PARTICLE PHYSICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Monte Carlo Simulations of physics processes at particle colliders like the Large Hadron Collider at CERN take up a major fraction of the computational budget. For some simulations, a single data point takes seconds, minutes, or even hours to compute from first principles. Since the necessary number of data points per simulation is on the order of 10^9 – 10^{12}, machine learning regressors can be used in place of physics simulators to significantly reduce this computational burden. However, this task requires high precision regressors that can deliver data with relative errors less than 1% or even 0.1% over the entire domain of the function. In this paper, we develop optimal training strategies and tune various machine learning regressors to satisfy the high-precision requirement. We leverage symmetry arguments from particle physics to optimize the performance of the regressors. We find that at lower dimensions, Boosted Decision Trees far outperform Neural Networks while at higher dimensions Neural Networks perform better. Inspired by ResNets, we design a Deep Neural Network with skip connections that outperform fully connected Deep Neural Networks. We show that these regressors can speed up simulations by a factor of 10^3 – 10^6 over the first-principles computations currently used in Monte Carlo simulations. Additionally, using symmetry arguments derived from particle physics, we reduce the number of regressors necessary for each simulation by an order of magnitude. Our work can significantly reduce the training and storage burden of Monte Carlo simulations at current and future collider experiments.