



Contribution ID: 215

Type: Poster

New RooFit Developments on Performance Optimization

Thursday 27 October 2022 11:00 (30 minutes)

RooFit is a toolkit for statistical modeling and fitting, and together with RooStats it is used for measurements and statistical tests by most experiments in particle physics, particularly the LHC experiments. As the LHC program progresses, physics analyses become more computationally demanding. Therefore, recent RooFit developments were focused on performance optimization, in particular to speed up the minimization of the negative log likelihood when fitting a model to a dataset.

Two such improvements will be discussed in this session: gradient-based CPU parallelization and batched computations. The former strategy parallelizes the calculation of the gradient in the line search approach (MIGRAD) used for minimum likelihood estimation in RooFit. Here, the parallelization approach and computational tools used will be discussed. The second strategy comprises a restructuring of the computational graph associated with a model and dataset in order to allow for batched computations. With batched computations RooFit can evaluate batches of events simultaneously per computational graph node, rather than event by event. This simultaneous computation can be either supported by vectorization or GPU parallelization.

Throughout this session, there will be an emphasis on detailed benchmarking and how it was used to optimize various parts of the developed performance improvements, including load balancing and the reduction of communication overhead. Benchmarks are primarily shown for cutting-edge Higgs combination fits, where the developed improvements were intended to achieve order-of-magnitude improvements in execution wall time.

Significance

RooFit is a library that is used in a large number of physics analyses, especially those from the LHC experiments. Physicists that use RooFit in any of their analyses can benefit from these new performance improvements, especially in those that are limited by computational runtime. From a computational perspective, the benchmarks and optimizations tailored to cutting-edge Higgs combination fits provide interesting insights into the computational challenges provided by such problems and how they could be tackled.

References

Experiment context, if any

Authors: Dr BURGARD, Carsten (Technische Universitaet Dortmund (DE)); REMBSER, Jonas (CERN); MONETA, Lorenzo (CERN); Dr BOS, Patrick (Nikhef National institute for subatomic physics (NL)); VERKERKE, Wouter

(Nikhef National institute for subatomic physics (NL)); WOLFFS, Zef (Nikhef National institute for subatomic physics (NL))

Presenter: WOLFFS, Zef (Nikhef National institute for subatomic physics (NL))

Session Classification: Poster session with coffee break

Track Classification: Track 2: Data Analysis - Algorithms and Tools