

## The effort to support automatic differentiation (AD) in RooFit

In this talk, we report on the effort to support automatic differentiation (AD) in RooFit, a toolkit for statistical modeling and fitting used by many HEP/NP experiments that is part of ROOT. The new AD backend improves both the performance and numeric stability of likelihood minimizations, for which we will provide several examples in this contribution. Our approach is to extend RooFit with a tool that generates overhead-free C++ code for a full likelihood function built from RooFit functional models. Gradients are then generated using Clad, a compiler-based source-code-transformation AD tool, using this C++ code. After presenting promising results from a proof-of-concept with this pipeline applied to a HistFactory model at the ACAT 2022 conference, we reported on the integration inside ROOT and showcased more general benchmarks at CHEP 2023. Following this last milestone, work focused on evolving the Minuit 2 minimizer backed to make better use of the automatic gradient and extended the code generation with support for more RooFit models. In this workshop, we will present updated benchmarks where all numeric differentiation is avoided on the Minuit 2 side, as well as new results with the RooFit AD backed applied to cutting-edge ATLAS Higgs analysis benchmarks for the first time.

These results show that the RooFit AD backend is the prime choice for combined binned likelihoods with many parameters, yielding minimization times one order of magnitude below RooFit's other backends and improving the fit convergence rate.