

# A Faster, More Intuitive RooFit

Stephan Hageböck<sup>1,\*</sup>

<sup>1</sup>CERN, Esplanade des Particules 1, 1211 Geneva 23, Switzerland

**Abstract.** RooFit and RooStats, the toolkits for statistical modelling in ROOT, are used in most searches and measurements at the Large Hadron Collider as well as at  $B$  factories. Larger datasets to be collected at e.g. the High-Luminosity LHC will enable measurements with higher precision, but will require faster data processing to keep fitting times stable. In this work, a simplification of RooFit's interfaces and a redesign of its internal dataflow is presented. Interfaces are being extended to look and feel more STL-like to be more accessible both from C++ and Python to improve interoperability and ease of use, while maintaining compatibility with old code. The redesign of the dataflow improves cache locality and data loading, and can be used to process batches of data with vectorised SIMD computations. This reduces the time for computing unbinned likelihoods by a factor four to 16. This will allow to fit larger datasets of the future in the same time or faster than today's fits.

## 1 Introduction

RooFit [1] is a C++ package for statistical modelling distributed with ROOT [2]. RooFit allows to define computation graphs to connect observables, parameters, functions and PDFs<sup>1</sup> to likelihood models, which can be fit to data and be used for statistical tests. RooFit is shipped with RooStats, a toolkit for performing statistical tests with RooFit models. It provides tools such as Toy Monte Carlo tests, setting limits and computing significances. Further, HistFactory provides tools to create RooFit models from a collection of ROOT histograms.

RooFit was originally developed for the BaBar collaboration, and is now used for statistical inference across many experiments in High-Energy Physics, e.g., at  $B$  factories and at the Large Hadron Collider. It is crucial for the final steps of most analyses. It was designed for single-core processors, and was neither optimised for large caches nor SIMD computations. This work stands at the beginning of efforts to modernise RooFit to speed up fits, and make it easier to use from both C++ and Python. This will enable researchers to analyse larger datasets and devise more elaborate statistical models.

## 2 Modernising RooFit's Interfaces

In RooFit, any collection of mathematical entities such as parameters, observables, functions or PDFs are saved or passed to functions using the classes RooArgSet and RooArgList, sets and lists of RooFit objects.

---

\*e-mail: stephan.hageboeck@cern.ch

<sup>1</sup>Probability Density Functions

The most common operation is iterating through these collections, both during fitting and when e.g. inspecting the values of parameters on the user side. This favours array-like data structures like e.g. `std::vector`, but internally, RooFit's collections were using a linked list with optional hash lookup. Iterating through a RooFit collection requires the following C++ code in ROOT 6.16:

```

1 TIterator* it = pdf.getParameters(obs)->createIterator();
2 RooAbsArg* p;
3 while ((p=(RooAbsArg*)it->Next())) {
4     p->Print();
5 }
6 delete it;

```

To speed up iterating and to provide an STL-like interface for RooFit's collections, the linked list in RooFit's collections was replaced by a `std::vector`. Functions such as `begin()`, `end()`, `size()`, `operator[]` were implemented to allow for STL-like handling of the collections. In ROOT 6.18, iterating through the list of observables of a PDF as above can therefore be achieved as follows:

```

1 for (auto p : *pdf.getParameters(obs))
2     p->Print();

```

The STL-like interface allows to reduce heap allocations, and replaces `while (p = Next())` loops by range-based for loops. This reduces code clutter, memory leaks, dangling pointers and variable shadowing, and speeds up iterating through collections by 20% to 25%. Random access, which was slow with large linked lists, now completes in constant time. Depending on how often collections are iterated, typical workflows in RooFit are sped up from 5% to 21% [3]. Fits with a binned ATLAS likelihood model [4] completed 19% faster while yielding identical results.

Modernising the C++ interfaces is also beneficial for using RooFit from Python. Since ROOT has a C++ interpreter, it can dynamically generate Python bindings for C++ objects [5]. In ROOT 6.16, a loop over the parameters of a PDF would have to imitate the C++ code in the first listing. However, with an STL-like interface, Python iterators are generated automatically. The equivalent Python loop using ROOT 6.18 therefore reads:

```

1 for p in pdf.getParameters(obs):
2     p.Print()

```

In ROOT 6.22, the interfaces of category classes will be modernised in a similar way. Defining and printing category numbers and names compares as follows:

ROOT 6.22	ROOT 6.18
<pre> 1 RooCategory cat("cat", "Lep. mult."); 2 cat.defineTypes( 3     {"0Lep", "1Lep", "2Lep", "3Lep"}, 4     { 0,     1,     2,     3 }); </pre>	<pre> 1 RooCategory cat("cat", "Lep. mult."); 2 cat.defineType("0Lep", 0); 3 cat.defineType("1Lep", 1); 4 cat.defineType("2Lep", 2); 5 cat.defineType("3Lep", 3); </pre>
<pre> 5 for (const auto&amp; name_idx : cat) { 6     std::cout &lt;&lt; name_idx.first &lt;&lt; ", " 7     &lt;&lt; name_idx.second &lt;&lt; std::endl; 8 } </pre>	<pre> 6 TIterator* typeIt = cat.typeIterator(); 7 RooCatType* catType; 8 while ( (catType = 9     dynamic_cast&lt;RooCatType*&gt;(typeIt-&gt;Next()) ) 10     != nullptr) { 11     std::cout &lt;&lt; catType.getVal() &lt;&lt; ", " 12     &lt;&lt; catType.GetName() &lt;&lt; std::endl; 13 } 14 delete typeIt; </pre>

## 2.1 Old Interfaces Remain Supported

Given that the old interfaces of classes such as `RooAbsCollection` and `RooAbsCategory` (+ derived classes) are in use, they cannot be deprecated without forcing users to update existing code. Therefore, the new interfaces are provided as an addition, while old interfaces remain supported. Updating to the new interfaces leads to faster, shorter and more type-safe code, but it is not required.

To enable such kind of backward compatibility, legacy iterators/functions were implemented, which mimic the functionality of the original objects. Occasionally, this requires extra virtual calls or heap allocations, leading to slow downs of 1 % to 3 %, but the new interfaces allow for 20 % speed up. The most critical iterators in `Roofit` have been modernised, and more iterators will be replaced as the modernisation of `Roofit` continues. To detect uses of inefficient interfaces in user code, users can add `#define R__SUGGEST_NEW_INTERFACE` before including ROOT headers, which will trigger deprecation warnings with the `clang++`, `g++` and `MSVC` compilers. Further, old interfaces will be documented in special sections of `Roofit`'s reference guide [6] to aid users in modernising existing code.

More updates of interfaces are planned, such as easier importing of data from e.g. ROOT's `RDataFrame`, STL containers or numpy arrays. The release of the new PyROOT [5] will further enable designing more pythonic interfaces that don't have to closely imitate the C++ syntax.

## 3 More Stable and Faster Built-in PDFs

In modernising `Roofit`, the Hypatia2 [7] (ROOT 6.20) and Johnson distributions [8] (ROOT 6.18, see figure) were added. Although any C++ function can be used as a PDF in `Roofit` (see section 4), providing optimised implementations for these frequently used PDFs is beneficial. Apart from providing a citeable default, built-in PDFs can be extended with extra checks, and optimised computations can be applied. The advantages of built-in PDFs are:




Figure 1. Johnson distribution

**Analytic Integrals** In order to normalise functions, these must be integrated over the definition range of their observables. `Roofit` integrates all functions numerically, unless a function for analytic integration is overridden. This allows for faster and more accurate normalisation.

**Generating Samples** `Roofit` generates data samples using the accept/reject method, unless a generator function is implemented. For many distributions, more efficient sampling strategies are known, which can only be employed for built-in PDFs.

**Checking parameters** Just-in-time compiled PDFs as in section 4 will accept parameters and observables with arbitrary definition range. When trying to evaluate a function outside of its definition range, computations might yield negative probabilities, infinity or NaN. This can slow down or entirely prohibit fitting distributions to data, since the minimiser has no means of determining whether a parameter is in the allowed range. Starting with ROOT 6.18, the definition ranges for parameters of built-in PDFs can be checked by the PDF implementation.

**SIMD Computations** Starting with ROOT 6.20, computations can be sped up significantly by evaluating large batches of data using SIMD computations. This is discussed in section 5.

For the built-in Johnson distribution, for example, previously unstable fits were found to converge six times faster than with the commonly used interpreted Johnson formula in LHCb.

## 4 Just-in-time Compiled PDFs

Although implementing a PDF as a RooFit class is the fastest and most accurate way of building likelihood models, RooFit was supporting interpreted PDFs using the class RooGenericPdf. It takes strings of function expressions, and interprets these using an instance of ROOT 5's TFormula. With the arrival of cling, the just-in-time C++ compiler and interpreter, TFormula was updated to use just-in-time compiled expressions. In ROOT 6.20, the new TFormula was integrated in RooFit, allowing for faster computations (compiled with optimisations instead of interpreted). This also enables using previously-defined functions or functions loaded from a library as in the following example:

```
1 // In a library or in included code
2 double func(double x, double a) { return a*x*x + 1.;}
3 [...]
4 // When building fit model
5 RooRealVar x("x", "Observable", 2.); // Define observable
6 RooRealVar a("a", "Parameter", 3.); // Define parameter
7 RooGenericPdf pdf("pdf", "func(x, a)", {x, a}); //evaluate func and normalise
```

## 5 Faster PDF Computations

### 5.1 Why Unbinned Fits Are Slow in RooFit

When RooFit fits PDFs to data, an expression is evaluated for each entry in the dataset to compute the likelihood of observing an event. RooFit achieves this by writing values from the rows of a dataset into the leaves of a computation graph, and the probability of the top node is evaluated by calling evaluate and normalisation functions of daughter nodes. Each node caches its last value, and therefore constant branches of the computation graph (e.g. branches that only depend on parameters) are only computed once. This cycle, however, repeats for every entry in the dataset such that all branches that depend on observables have to be recomputed every time. The total number of function calls is therefore proportional to the number of entries in the dataset and to the number of (non-constant) nodes in the graph,  $N_{\text{Data}} \cdot N_{\text{Nodes}}$ . For a small mathematical expression with 10 nodes and one million events, this already amounts to considerably more than 10 million function calls because additional calls for the normalisation of PDFs and for invalidating the node-local caches are necessary.

Furthermore, loading only single values into the nodes of the computation graph is hostile to CPU caches. It is possible that when a node is being revisited to load the next entry, data have been evicted from the cache(s). This means that RooFit runs inefficient on modern CPUs because of poor data locality and inefficient memory access patterns.


### 5.2 Three Times Faster Computations with Higher Data Locality

To improve the data locality and reduce the number of function calls, a RooFit-internal interface for batched likelihood computations was implemented. Data are passed between nodes using a `std::span`. Inputs for computations are directly read from arrays, while outputs are written to contiguous memory that is owned by the node that is running a computation.

Instead of computing only one probability per node, all probabilities for all entries in the dataset can be computed in a few function calls.<sup>2</sup> Since such computations operate on contiguous floating point numbers, data locality, caching and data prefetching improve. This

---

<sup>2</sup>RooFit also supports a multi-process mode, in which case smaller batches of contiguous data are processed, which are divided among multiple workers.



**Figure 2.** Speed up for computing the likelihoods of datasets of 100 000 to 300 000 events for various likelihood models. Using ROOT 6.20, the fast RooFit batch interface is timed against the normal single-value computations for different CPUs, compilers and instruction sets. On an intermediate-level CPU that supports AVX2 instructions (■), a speed up of 4x to 9x can be expected. On CPUs supporting AVX512 instruction sets, the speed up ranges from 4x to 16x. For the “ChiSquarePdf” (top), almost no SIMD functions were available. The speed up of 3x is the result of batched computations with faster data loading.

speeds up computations three to four times without loss of precision. Figure 2 shows the combined speed up of batched computations and additionally vectorisation with SIMD instructions (section 5.3) against classic single-value RooFit computations. For the topmost entry (“ChiSquarePdf”), almost no SIMD instructions could be used. This speed up is mostly due to faster data loading and reducing function calls.

When PDFs that support the fast interface are used together with legacy PDFs, compatibility is ensured by a generic batch computation function that runs RooFit’s classic single-value computations for each entry in the dataset, writes results into an array, and passes the results on to the next node of the graph using `std::span`.

### 5.3 4x to 16x Faster Computations with SIMD

By converting RooFit’s data access patterns to reading from array-like structures, RooFit could be extended with single-instruction-multiple-data computations (SIMD). Optimisation reports of the `clang++-8` compiler were used to design computation functions that allow for automatic vectorisation with AVX2 instructions. This allows to increase the throughput, since four double-precision numbers can be processed simultaneously (8 with AVX512). This requires using auto-vectorisable, inlinable mathematical functions, and entries in arrays cannot depend on other entries. Auto-vectorisable functions are provided by the VDT [9] package, of which logarithm and exponential function were used most frequently. VDT approximates standard library functions using Padé polynomials, which can be inlined into other computations, and optimised by the compiler. ROOT can be compiled without VDT, though, in which case the standard (non-inlinable) implementations are used. This usually prevents automatic vectorisation.

To further aid the compiler optimiser, computations were rewritten such that data dependencies between elements are eliminated, and that as many intermediate results as possible can be held in CPU registers. Excessive branching in complicated functions was replaced by branch-less computations such as  $1 \cdot A + 0 \cdot B$  to select  $A$  and its counterpart to select  $B$ . Although more CPU cycles are needed because both branches are computed, higher throughput is achieved because SIMD instructions process four or eight entries simultaneously. Finally, calls to non-inlinable functions and complicated reductions were removed where possible.

The result of this work is shown in fig. 2. For various RooFit PDFs and composite models that use multiple PDFs, the run time of the optimised SIMD computation with VDT functions is compared to classic RooFit single-value computations. Depending on the complexity of the computations and on the level of automatic compiler optimisations, computations speed up 3x to 16x. The theoretical speed up of 12x (3x from data loading, 4x for AVX2) cannot be reached because of Amdahl's law. Not all parts of RooFit's code can be replaced with SIMD code, and were therefore not sped up. Moreover, in order to use SIMD instructions, more instructions may have to be issued (e.g. compute two Padé polynomials to approximate the `exp` function instead of calling the standard implementation).

The level of automatic compiler optimisations has a strong effect on the speed up. Figure 2 demonstrates that both clang 8 and 9 have powerful optimisers for AVX2 instructions, but gcc significantly outperforms clang when ROOT is compiled with AVX512 instructions. It can therefore be expected that as compiler optimisers improve, the speed up for RooFit increases, and that differences between the compilers will shrink.

Batch computations with vectorisation have been released in ROOT 6.20. These can be switched on by adding the “BatchMode” argument to the fitting instruction:

```
pdf.fitTo(data, BatchMode(true)); // Evaluate likelihood using fast batch mode
```

Unit tests ensure that the optimised computation functions yield the same results as the classic RooFit functions. When the fast VDT approximations are used, the relative difference for probabilities is usually below  $1.0 \times 10^{-14}$ , for log-likelihoods below  $2.0 \times 10^{-14}$ , and for fit parameters below  $1.0 \times 10^{-5}$ , which is except for corner cases orders of magnitude smaller than the statistical error of the fit parameters.

A drawback of the current implementation is that for maximal performance, ROOT has to be compiled targeting the instruction set that the CPU supports. The possibility of shipping a small library with computation functions for a few different architectures (e.g. SSE4 + AVX2) is being investigated.

## 6 Summary

RooFit's interfaces are being modernised, and computations are being sped up with fast data loading and SIMD computations. The single-thread performance of unbinned fits in RooFit has been increased by several factors without requiring code changes on the user side. In conjunction with work on parallelising computations [10], a speed up by more than an order of magnitude can be expected.

The work on modernising RooFit's interfaces and improving its performance will continue, especially for binned fits such as HistFactory models.

## References

- [1] W. Verkerke, D.P. Kirkby, econf **C0303241**, MOLT007 (2003), physics/0306116
- [2] R. Brun, F. Rademakers, Nucl. Instrum. Methods Phys. Res. **A389**, 81 (1997)
- [3] S. Hageboeck, L. Moneta, *Making RooFit Ready for Run 3*, in *19th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT2019)* (2020), unpublished
- [4] ATLAS Collaboration, JHEP **01**, 069 (2015), 1409.6212
- [5] E. Tejedor Saavedra, S. Wunsch, M. Galli, *A new PyROOT: Modern, Interoperable and more Pythonic*, EPJ Web Conf. (this volume) (2020)
- [6] The ROOT Team, *Root reference guide* (2020), <https://root.cern/doc/master/classRooAbsCategory.html>
- [7] D.M. Santos, F. Dupertuis, Nucl. Instrum. Methods Phys. Res. A **764**, 150 (2014)
- [8] N.L. Johnson, Biometrika **36**, 149 (1949)
- [9] D. Piparo, V. Innocente, T. Hauth, J. Phys. Conf. Ser. **513**, 052027 (2014)
- [10] P. Bos et al., *Faster roofitting: Automated parallel computation of collaborative statistical models*, EPJ Web Conf. (this volume) (2020)