



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
Main Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2020

zfit: Scalable pythonic fitting

Eschle, Jonas ; Puig Navarro, Albert ; Silva Coutinho, Rafael ; Serra, Nicola

Abstract: Statistical modeling is a key element in many scientific fields and especially in High-Energy Physics (HEP) analysis. The standard framework to perform this task in HEP is the C++ ROOT/RooFit toolkit; with Python bindings that are only loosely integrated into the scientific Python ecosystem. In this paper, zfit, a new alternative to RooFit written in pure Python, is presented. Most of all, zfit provides a well defined high-level API and workflow for advanced model building and fitting, together with an implementation on top of TensorFlow, allowing a transparent usage of CPUs and GPUs. It is designed to be extendable in a very simple fashion, allowing the usage of cutting-edge developments from the scientific Python ecosystem in a transparent way. The main features of zfit are introduced, and its extension to data analysis, especially in the context of HEP experiments, is discussed.

DOI: <https://doi.org/10.1016/j.softx.2020.100508>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-189044>

Journal Article

Published Version



The following work is licensed under a Creative Commons: Attribution 4.0 International (CC BY 4.0) License.

Originally published at:

Eschle, Jonas; Puig Navarro, Albert; Silva Coutinho, Rafael; Serra, Nicola (2020). zfit: Scalable pythonic fitting. SoftwareX, 11:100508.

DOI: <https://doi.org/10.1016/j.softx.2020.100508>



Original software publication

zfit: Scalable pythonic fitting

Jonas Eschle*, Albert Puig Navarro, Rafael Silva Coutinho, Nicola Serra

Physik-Institut, Universität Zürich, Zürich, Switzerland



ARTICLE INFO

Article history:

Received 20 December 2019
 Received in revised form 24 April 2020
 Accepted 12 May 2020

Keywords:

Model fitting
 Data analysis
 Statistical inference
 Python

ABSTRACT

Statistical modeling is a key element in many scientific fields and especially in High-Energy Physics (HEP) analysis. The standard framework to perform this task in HEP is the C++ ROOT/RooFit toolkit; with Python bindings that are only loosely integrated into the scientific Python ecosystem. In this paper, zfit, a new alternative to RooFit written in pure Python, is presented. Most of all, zfit provides a well defined high-level API and workflow for advanced model building and fitting, together with an implementation on top of TensorFlow, allowing a transparent usage of CPUs and GPUs. It is designed to be extendable in a very simple fashion, allowing the usage of cutting-edge developments from the scientific Python ecosystem in a transparent way. The main features of zfit are introduced, and its extension to data analysis, especially in the context of HEP experiments, is discussed.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Code metadata

Current code version	0.3.6
Permanent link to code/repository used for this code version	https://github.com/ElsevierSoftwareX/SOFTX_2019_376
Legal Code License	BSD-3
Code versioning system used	Git
Software code languages, tools, and services used	Python, TensorFlow
Compilation requirements, operating environments & dependencies	Linux and MacOS supported (Windows functionality unknown); Python 3.6+; dependencies (also specified in <i>requirements.txt</i>): <code>tensorflow >= 1.14.0, < 2,</code> <code>tensorflow_probability >= 0.6.0, < 0.8,</code> <code>scipy >= 1.2</code> <code>uproot, pandas, numpy, iminuit,</code> <code>typing, colorlog, texttable, ordered-set</code> https://zfit.readthedocs.io/en/0.3.6/ zfit@physik.uzh.ch
If available Link to developer documentation/manual	
Support email for questions	

Software metadata

Current software version	0.3.6
Permanent link to executables of this version	https://github.com/zfit/zfit/releases/tag/0.3.6
Legal Software License	BSD-3
Computing platforms/Operating Systems	Unix-like (Windows untested)
Installation requirements & dependencies	dependencies (also specified in <i>requirements.txt</i>): <code>tensorflow >= 1.14.0, < 2,</code> <code>tensorflow_probability >= 0.6.0, < 0.8,</code> <code>scipy >= 1.2</code> <code>uproot, pandas, numpy, iminuit,</code> <code>typing, colorlog, texttable, ordered-set</code> https://zfit.readthedocs.io/en/0.3.6/
If available, link to user manual—if formally published include a reference to the publication in the reference list	
Support email for questions	zfit@physik.uzh.ch

* Corresponding author.

E-mail addresses: Jonas.Eschle@cern.ch (J. Eschle), albert.puig.navarro@gmail.com (A. Puig Navarro), rafael.silva.coutinho@cern.ch (R. Silva Coutinho), nicola.serra@cern.ch (N. Serra).

<https://doi.org/10.1016/j.softx.2020.100508>

2352-7110/© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Data collected by experiments such as the Large Hadron Collider (LHC) at CERN or ultra-energetic cosmic ray detectors are analyzed in terms of physics-motivated mathematical models. Typically, these models have a parametric form, and the values of their parameters are the quantities of interest in a given analysis. Once a model has been defined, a *fitting* procedure that minimizes the disagreement between the data and the model is applied to find the optimal values of these parameters, a crucial step in almost every analysis.

Model fitting as done in High Energy Physics (HEP) has some peculiarities and differs from other fields. Models are often multidimensional, built by composing complicated, custom shapes, and have a non-trivial normalization. In addition, two further requirements need to be fulfilled: on one hand, reasonable scaling with data size and model complexity, motivated by the large data samples being collected at the LHC; and, on the other hand, the capability to not only find optimal values for the parameters, but to be able to calculate their uncertainties with enough flexibility taking all correlations into account to allow for advanced statistical treatment, as uncertainties of measurements are as important as their central values.

Since recent years, HEP analysis is moving more and more towards the usage of Python over the more traditional C++. The moving force behind this paradigm shift is the existence of an extremely rich scientific computing Python ecosystem [1,2], which provides a collection of efficient, easy-to-use tools for data analysis that are shared across a very large community, including scientists from various fields and industry, especially in data science, and which one can build new applications on.

The *de facto* standard for model fitting in HEP is RooFit [3], which is integrated in the ROOT toolkit [4] and provides a very powerful object-oriented architecture, as well as plotting capabilities and an advanced statistics module. It is written in C++ and can be accessed in Python through the ROOT Python bindings. These bindings come with its own set of non-negligible problems: complex memory management that easily leads to memory leaks, difficult integration with the scientific computing Python stack, lack of extendability from Python, and a complex installation procedure due to the need to install the full ROOT toolkit.

General-purpose Python-based libraries such as SciPy [5], Imfit [6] or TensorFlow Probability [7] do not fulfill all HEP fitting requirements, and therefore cannot be used *as-is*. Python-based HEP-specific fitting packages, such as probfit [8], TensorProb [9] or TensorFlowAnalysis [10], have tried to address these issues and can be considered as proof-of-concepts that HEP-centric model fitting is possible in Python; none of them, however, offers the same fitting functionalities and ease of use as RooFit does while still integrating well with the scientific computing Python ecosystem.

The goal of zfit is to fill this important gap in the HEP software ecosystem and provide a well defined API and workflow for model fitting together with a Python-based package that is fast, scalable and extensible.

2. Software description

The structure of a generic model fitting workflow is illustrated in Fig. 1. Its basic goal is to maximize the agreement of a set of data points to a probability distribution described by a set of parameters. The measurement of the metric of the goodness – or, more precisely, the disagreement – of the predicted *model* with the experimental *data* is referred to as *loss* function. One of the most used loss functions in HEP is based on the likelihood function, in particular the negative log-likelihood, which

is minimized in order to estimate the parameters of the model. Once the optimal values of the model parameters have been determined, several statistical techniques can be used to calculate their *uncertainties*.

The structure defined by *model-data-loss-minimize* can also be used to describe a typical Deep Learning workflow—even more clearly by replacing *model* by “Neural Network” and *minimize* by “Training”. This similarity has profound implications in the design of the zfit implementation, as it allows to use an industry standard Deep Learning tool such as TensorFlow [11] (TF) as its computational backend by using the low-level API that consists of mathematical functions. The immediate benefits for zfit in terms of performance arise from key features of TF such as the capacity to be run on heterogeneous architectures, mathematical optimizations such as automatic differentiation (required by the most used minimization algorithms in HEP), and operation caching. This puts the performance of zfit in the same order of magnitude as the existing standard as shown in Appendix.

zfit does not intend to be a large, monolithic package and has therefore a clear, restricted scope: model fitting and sampling with a well defined and consistent API which is mostly independent of the implementation. Tasks such as data processing, plotting and statistical analysis are left for other packages to tackle.

The abstraction of the fitting process in the steps shown in Fig. 1 is one of the main features of the architecture of zfit: by developing each step as a separate “brick” of the library, it becomes possible to tightly integrate with the scientific computing Python ecosystem, and use the available tools to improve performance, reduce code duplication and add extra functionality. Each of these bricks (*model*, *data*, *loss*, *minimizer* and *fit result*) provides override hooks for all its core functionality, allowing to completely customize or replace any aspect of their behavior. As hinted at above, these objects deal internally with TF. Since using TF in the correct way adds another layer of difficulty, most of these additional complications are hidden from the user and the library offers a similar user experience as the one found in other model fitting libraries, independent of the actual backend.

Model fitting handles multiple dimensions. There are observables, the dependent variables, over which a model is usually normalized and parameters, which will be adjusted during the minimization of the loss. In contrast to other frameworks such as RooFit, zfit makes a strong distinction between observables and parameters; this removes any ambiguity in terms of normalization, allows to be closer to Deep Learning and the optimized way of how TF works. The observables and corresponding limits in which the model fit is performed are handled by the *Space* object. Space objects can be single- or multi-dimensional, and, thanks to the use of observables, can be combined using the Python arithmetic operators to tackle arbitrarily complex use cases.

Models are defined inside a given Space and implement a shape function that returns a value given some input data. Custom models are easily implemented using their functional form and by default receive standard attributes such as numerical integration and sampling. These generic functionalities can also be extended to analytical integration and specific sampling if provided by the user. Models can be combined to build more complex models (including increasing the number of dimensions) while keeping automatic integration, observable handling and event sampling. The normalization is done over a range specified in the *obs* or further customized by overriding the *_normalization* method manually.

Models are parametrized by *Parameter* objects, which can be used in the implementation of the shape function like any other scalar. *Independent parameters* are the only type of objects that are left free in the fit. *Composed parameters* can be built by combining

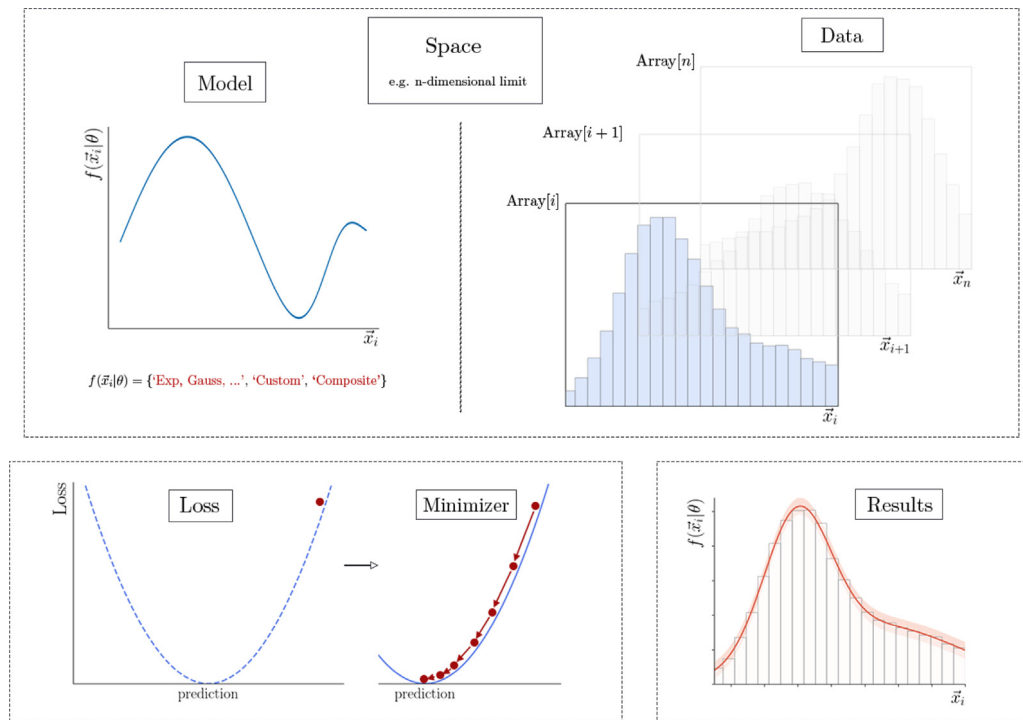


Fig. 1. Schematic zfit workflow view of the individual structural blocks.

parameters through mathematical operators and can be used like independent parameters.

The loading, ordering and processing of data in zfit is handled uniformly, no matter their source, by the *Data* object. While only the most important formats are currently implemented – ROOT trees, Numpy arrays and Pandas DataFrames, as well as TF tensors – the simple API of the *Data* object makes it easy to add additional data handling capabilities.

Data and Models are combined to build a *Loss* function, which can either be pre-defined such as the negative log-likelihood or completely arbitrary. The computation of the value of the loss for a given configuration of the model is efficiently implemented in zfit thanks to TF. It is also possible to add constraints to the loss in order to modify its behavior. Constraints are additional terms to the likelihood that often represent some external information on the parameter. While the most used constraints in HEP are directly implemented in zfit, such as Gaussian constraints, arbitrary custom functions can be also used.

Minimization is performed through stateless *Minimizer* objects. Since the loss calculation is usually the most intensive part of the minimization problem, in practice it is possible to wrap any existing Python minimizer to be used by zfit. Several minimizers are included in zfit, such as the SciPy optimizers, the most-used minimizer in HEP, called Minuit [12], or the TF implementation of the Adam optimizer [13].

The results of the minimization process are returned in a *FitResult* object, which stores all the information about the fit parameters, the minimization process, as well as the used instances of the *Minimizer* and the *Loss*. This includes the function minimum and the covariance matrix.

While the approximate uncertainties related to the fit parameters can be calculated directly inside zfit using the Hessian matrix or a simple profiling method, the calculation of precise parameter uncertainties and confidence intervals are left to specialized packages, such as those in Refs. [14,15], which do not need any object other than the *FitResult*.

3. Illustrative examples

The architecture and features of zfit are better visualized with a typical HEP example: a fit to the invariant mass of multiple particles including both signal and background events, extended to multiple dimensions, and introducing a control channel.

3.1. Basic usage

Let us assume the domain of interest of the mass observable, in this case the B^0 meson mass, is in the range [4.5, 6.0] GeV. In zfit this is expressed with a *Space* defining our domain:

```
obs = zfit.Space(obs="mass", limits=(4.5, 6.0))
```

The best interpretation of the observable at this stage is that it defines the name and range of the observable axis.

Data can be loaded for example from a numpy array, where the *obs* of the *Space* are matched by order with the columns of the array:

```
data = zfit.Data.from_numpy(array=numpy_array_data, obs=obs)
```

with *numpy_array_data* being a numpy array holding our data. The signal component of the fit will be modeled using a Gaussian (normal), which has two free parameters—its mean and its width:

```
mu = zfit.Parameter("mu", 5.2, 4.5, 6.0)
sigma = zfit.Parameter("sigma", 20, 1, 40)
```

where the arguments for the parameters are: a unique name, an initial value, a lower and an upper bound. The bounds are not mandatory, the parameter will be floating in both cases. These can be used to instantiate the *Model* as:

```
signal = zfit.pdf.Gauss(obs=obs, mu=mu, sigma=sigma)
```

While most commonly-used models are provided within the pdf module for convenience, the philosophy of zfit is to provide a clear API with powerful mechanisms to facilitate that the users and the community can implement their own. Making use of another built-in model, the background that is parametrized as an exponential:

```
bkg = zfit.pdf.Exponential(-0.001, obs=obs)
```

In order to build the sum of these two models, an additional parameter `frac` is used to describe the fraction of each species¹:

```
frac_sig = zfit.Parameter("signal fraction", 0.5, 0, 1)
model = zfit.pdf.SumPDF(pdf=[signal, bkg], fracs=frac_sig)
```

With the model and the data built, a loss function can be constructed. In this case, the goal is to perform an unbinned maximum likelihood fit, so the corresponding Loss object is initialized as

```
nll = zfit.loss.UnbinnedNLL(model=model, data=data)
```

Finally, the minimization process is executed, first instantiating a Minimizer object – in this case using the Minuit minimizer – and then running the minimization algorithm:

```
minimizer = zfit.minimize.MinuitMinimizer()
result = minimizer.minimize(nll)
```

The outcome of this minimization is stored in a `FitResult` object, which provides access to all details that occurred during the convergence process and to the best values estimated for the free parameters together with a nice printout of the parameters and uncertainties. For example, one could obtain the fitted value of the signal mean as²

```
if result.converged:
    mu_value = result.params[mu]["value"]
```

The `params` object stores additional information such as estimated uncertainties if they have been calculated.

3.2. Advanced usage

The next step is to extend the model to take into account extra dimensions of the problem, in this case the angular distributions of the decay products of the B^0 meson. To do so, the signal distribution is modeled with a custom model, which is created by inheriting from one of the base classes provided by zfit:

```
from zfit import z

class AngularSignalPDF(zfit.pdf.ZPDF):
    _PARAMS = ["FH", "AFB"]
    _N_OBS = 1

    def _unnormalized_pdf(self, x):
        # retrieve data
        costheta = z.unstack_x(x)
        # retrieve parameters
        FH = self.params["FH"]
        AFB = self.params["AFB"]
        # build the shape
        pdf = (3 / 4 * (1 - FH) * (1 - z.square(costheta))
              + 0.5 * FH + AFB * costheta)
        return pdf
```

¹ Following a similar approach to the reference implementation of RooFit, the fraction of the last species corresponds to $1 - \sum \text{frac}_i$.

² Notice that `mu`, the parameter object, and not `"mu"`, the name of it, is used as the key.

Instead of `BasePDF`, we use here an even simpler class, `ZPDF`, which only requires to specify the naming of the parametrization via the `_PARAMS` attribute and allows, optionally, to define the dimensionality through `_N_OBS`. The latter also allow to create n-dimensional Models as `x` will be a list of length `_N_OBS` containing the data. Internally, the class handles the management of the free parameters, which can be accessed by their parametrization name through the `params` attribute, which holds the parameters. It also provides a simple way to access the input data with the `z.unstack_x`. The module `z` provides some TF functions wrapped for better handling of dtypes as well as additional methods and can be used mixed with pure TF functions.

```
# create the parameters
fh = zfit.Parameter("FH", 1.)
afb = zfit.Parameter("AFB", 0.5)

obs_costheta = zfit.Space("theta", limits=(-1, 1))

angular = AngularSignalPDF(obs=obs_costheta, FH=fh, AFB=afb)
```

To create a two-dimensional model, one way is to multiply the models describing the different dimensions. Given that their observables are different, the new model will automatically be defined in the combined space of the two:

```
combined_model = angular * model
```

This example can be extended to simultaneously fit a control dataset from which a more precise determination of the parameter `mu` can be obtained. To do so, a second Gaussian is created to model this control mode:

```
obs_control = zfit.Space("control", limits=(5.2, 5.6))
sigma_control = zfit.Parameter("sigma control channel", 2, 0, 5)
gauss_control = zfit.pdf.Gauss(obs=obs_control,
                               mu=mu # common mu
                               sigma=sigma_control)
```

A simultaneous Loss can be built by either two independent Loss functions or by providing both models and datasets in a list as

```
nll = zfit.loss.UnbinnedNLL([combined_model, gauss_control],
                           [data, data_control])
```

where `data_control` has been created analogously to `data`. This simultaneous Loss can be minimized as shown before.

4. Impact and conclusions

The zfit package is a crucial piece required to complete the transition of HEP analysis to a stack fully based on the scientific computing Python ecosystem. The importance of this paradigm change should not be underestimated: given the limited availability of resources in HEP, especially in areas such as software development, the possibility of integrating physics analysis within a well-developed and well-supported software ecosystem – allowing to focus mainly on the HEP-specific parts of the analysis software – will have long term repercussions in the quality and usability of HEP analysis software.

The package itself is not only an implementation of a versatile and powerful model fitting library, but also specifies a carefully designed and well discussed API of the most crucial components and the workflow. With the aim to be a standard in its kind, other packages such as high-level statistics and plotting tools, can be built against the standardized API instead of a concrete implementation, strongly reducing package dependencies. In addition, other packages can focus on the specific implementation of a certain task, such as building a specialized Loss, Model or a Minimizer, without the need to also implement the other parts of the workflow themselves.

Finally, fits in HEP analysis often consist of large datasets and complicated models, and require far more computing power than a single core. Scalability in terms of size and complexity of the data and models and the possibility to run on large computing infrastructures, including accelerators such as GPUs, is therefore a necessity. TensorFlow, the graph based computing backend in zfit, was built for the sole purpose of efficient large scale computing and takes care of the non-trivial task of optimal workload parallelization. This does not only yield state-of-the-art performance, but also enables the HEP community to keep up with the latest algorithm implementations and low-level computing instructions at virtually no cost.

In summary, a new high-level statistical modeling and fitting library purely designed within, and for, the Python ecosystem is presented. Future plans for zfit include full support to binned models, further extension on baseline HEP features as well as standardized, human-readable serialization of objects.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We are grateful to Anton Poluektov, Chris Burr and Igor Babuschkin for demonstrating the potential of unbinned model fitting within the context of TensorFlow, which inspired this work. We also thank the Zurich LHCb Group, Matthieu Marinangeli, Josh Bendavid, Lukas Heinrich and the HSF community, especially Scikit-HEP project members, for useful discussions. A. Puig, R. Silva Coutinho, J. Eschle and N. Serra gratefully acknowledge the support by the Swiss National Science Foundation (SNF) under contracts 168169, 174182 and 182622.

Appendix. Performance

A series of studies are performed to evaluate the execution time of zfit in comparison to the conventional RooFit library. The performance relies on two aspects, *i.e.* the complexity of the fit (given by the number of free parameters) and the data sample size. In this study, the sum of 9 Gaussian functions with a common mean and width parameter, but with different central values, is used as baseline model.

Ensembles of Monte Carlo pseudoexperiments are generated and fitted with different setup configurations.³

- i. Nominal zfit implementation on CPU, using an analytic gradient provided by TF. In this case, an initial run is done to remove the graph compiler time. While not significant, this provides a more realistic estimation.
- ii. Alternatively the analytic gradient provided by TF is disabled, denoted by the addition “nograd”. Instead, the Minuit minimizer calculates a numerical approximation of the gradient internally.
- iii. The default CPU zfit settings is replaced by the GPU implementation.
- iv. A standard RooFit implementation using the Python bindings with PyROOT (version 6.16.00). The parallelization is done when invoking the `fitTo` method and equals the number of cores available.

³ The fitting performance is examined using a 12 core Intel i7 8850H with 2.60 GHz and 32 GB RAM, and a Nvidia P1000 with 4 GB RAM.

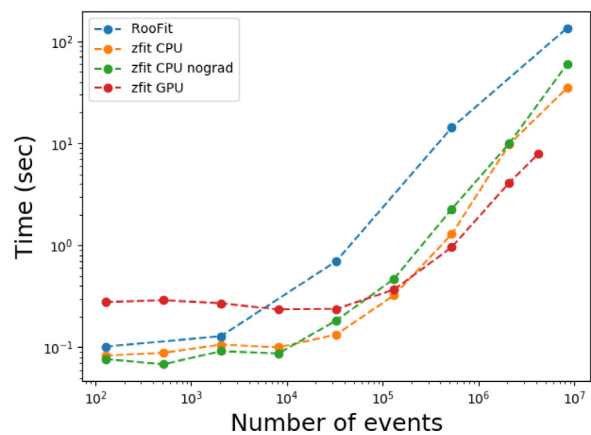


Fig. A.2. Averaged time per fit for a sum of Gaussian with shared mean and width as a function of the number of generated events.

For each setup, twenty toys are generated with sample sizes from 128 to 8 million events (except for the GPU, where it goes only up to 4 million⁴) and the average fitting time is used as reference point. Note that for simplicity the Minuit algorithm is used in all cases.

Fig. A.2 compares the performance of different fit scenarios and can be summarized as follows.

- As the number of events increase, the execution time of RooFit monotonically increases.⁵
- There is no advantage using parallelization for very few calculations, since the overhead of splitting and collecting the results is dominant. With increasing number of events this gets negligible and thus the execution time of zfit increases way slower than that for RooFit. This also comes from the fact that more events mean a more stable loss shape, so the minimizer used in zfit performs better.
- The GPU is highly efficient in computing thousands of events in parallel. For only a few data points though, the overhead of moving data back and forth dominates, resulting in a worse performance in this regime.

As a conclusion, the speed for a multicore CPU system is comparable between zfit and RooFit. For large number of events, zfit typically outperforms with a similar response for both GPU and a multicore system.

References

- [1] Oliphant TE. Python for scientific computing. *Comput Sci Eng* 2007;9(3):10–20. <http://dx.doi.org/10.1109/MCSE.2007.58>.
- [2] Millman KJ, Aivazis M. Python for scientists and engineers. *Comput Sci Eng* 2011;13(2):9–12. <http://dx.doi.org/10.1109/MCSE.2011.36>.
- [3] Verkerke W, Kirkby DP. The RooFit toolkit for data modeling. *eConf* 2003;C0303241:MOLT007. [arXiv:physics/0306116](https://arxiv.org/abs/physics/0306116).
- [4] Brun R, Rademakers F. ROOT: An object oriented data analysis framework. *Nucl Instrum Meth* 1997;A389:81–6. [http://dx.doi.org/10.1016/S0168-9002\(97\)00048-X](http://dx.doi.org/10.1016/S0168-9002(97)00048-X).
- [5] Jones E, Oliphant T, Peterson P, et al. SciPy: Open source scientific tools for python. 2001–. [Online; accessed 11/07/2019] URL <http://www.scipy.org/>.

⁴ The GPU used in these tests has a relative small memory. Performing larger-than-memory computations and multi-GPU is still work in progress.

⁵ It is worth mentioning that the RooFit parallelization strategy is rather simple and with very few overheads; several improvements in this direction are ongoing. It however hints an important point of using TF in zfit, namely that this kind of work is factored out and a superior performance can come for free.

- [6] Newville M, Otten R, Nelson A, Ingargiola A, Stensitzki T, Allan D, Fox A, Carter F, Michał, Pustakhod D, Ram Y, Glenn, Deil C, Stuermer, Beelen A, Frost O, Zobrist N, Pasquevich G, Hansen ALR, Spillane T, Caldwell S, Polloreno A, andrewhannum, Zimmermann J, Borreguero J, Fraine J, deep 42-thought, Maier BF, Gamari B, Almarza A. Lmfit/lmfit-py. 2019, <http://dx.doi.org/10.5281/zenodo.598352>.
- [7] Dillon JV, Langmore I, Tran D, Brevdo E, Vasudevan S, Moore D, Patton B, et al. TensorFlow distributions. 2017, [arXiv:1711.10604](https://arxiv.org/abs/1711.10604).
- [8] Ongmongkolkul P, Deil C, hsiang Cheng C, Pearce A, Rodrigues E, Schreiner H, Marinangeli M, Geiger L, Dembinski H. Scikit-hep/probfit. 2018, <http://dx.doi.org/10.5281/zenodo.1477852>.
- [9] Babuschkin I, Burr C, Neubert S, Brügge K. TensorProb. 2019, URL <https://tensorprob.github.io/tensorprob/>.
- [10] Poluektov A, Mauri A, Merli A, Mathad A, Martinelli M, Morris A. Using tensorflow for amplitude fits - the TensorFlowAnalysis package. 2018, <http://dx.doi.org/10.5281/zenodo.1415412>.
- [11] Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, et al. TensorFlow: Large-scale machine learning on heterogeneous systems. 2015, Software available from tensorflow.org URL <https://www.tensorflow.org/>.
- [12] James F, Roos M. Minuit: A system for function minimization and analysis of the parameter errors and correlations. *Comput Phys Comm* 1975;10:343–67. [http://dx.doi.org/10.1016/0010-4655\(75\)90039-9](http://dx.doi.org/10.1016/0010-4655(75)90039-9).
- [13] Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. In: 3rd International Conference on Learning Representations, San Diego: 2015, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [14] Marinangeli M. Lauztat. 2019, <http://dx.doi.org/10.5281/zenodo.3249383>.
- [15] Marinangeli M, Pollack B, Rodrigues E, Eschle J. Scikit-hep/hepstats. 2019, <http://dx.doi.org/10.5281/zenodo.3519200>.