

# 1 UnROOT: an I/O library for the CERN ROOT file 2 format written in Julia

3 **Tamás Gál<sup>1, 2</sup>, Jerry (Jiahong) Ling<sup>3</sup>, and Nick Amin<sup>4</sup>**

4 **1** Erlangen Centre for Astroparticle Physics **2** Friedrich-Alexander-Universität Erlangen-Nürnberg **3**  
5 Harvard University **4** University of California, Santa Barbara

DOI: [10.21105/joss.0XXXX](https://doi.org/10.21105/joss.0XXXX)

## Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

Editor: [Editor Name](#) ↗

Submitted: 01 January XXXX

Published: 01 January XXXX

## License

Authors of papers retain  
copyright and release the work  
under a Creative Commons  
Attribution 4.0 International  
License ([CC BY 4.0](#)).

## 6 Summary

7 `UnROOT.jl` is a pure Julia implementation of CERN ROOT (Brun & Rademakers, 1997) files  
8 I/O (`.root`) that is fast, memory-efficient, and composes well with Julia's high-performance  
9 iteration, array, and multi-threading interfaces.

## 10 Statement of need

11 The High-Energy Physics (HEP) community has been troubled by the two-language problem  
12 for a long time. Often, physicists would start prototyping with a `Python` front-end which  
13 glues to a `C/C++/Fortran` back-end. Soon they will hit a task which is extremely hard to  
14 express in columnar (i.e. “vectorized”) style, a type of problems which are normally tackled  
15 with libraries like `numpy` (Harris et al., 2020) or `pandas` (The pandas development team,  
16 2020). This usually leads to either writing `C++` kernels and interface them with `Python`, or  
17 porting the prototype to `C++` and start to maintain two code bases including the wrapper code.  
18 Both options are engineering challenges for physicists who usually have no or little background  
19 in software engineering. Many steps of this process are critical, like identifying bottlenecks,  
20 creating an architecture which is both performant and maintainable at the same time while  
21 still being user-friendly and logically structured. Using a `Python` front-end and dancing across  
22 language barriers also hinders the ability to parallelize tasks that are conceptually trivial most  
23 of the time.

24 `UnROOT.jl` attempts to solve all of the above by choosing Julia, a high-performance language  
25 with simple and expressive syntax (Bezanson et al., 2017). Julia is designed to solve the two-  
26 language problem in general. This has been studied for HEP specifically as well (Stanitzki &  
27 Strube, 2021). Analysis software written in Julia can freely escape to a `for`-loop whenever  
28 vectorized-style processing is not flexible enough, without any performance degradation. At  
29 the same time, `UnROOT.jl` transparently supports multi-threading and multi-processing by  
30 simply providing data structures which are a subtype of `AbstractArray`, the built-in abstract  
31 type for array-like objects, which allows to interface with array-routines from other packages  
32 easily, thanks to multiple dispatch, one of the main features of Julia.

## 33 Features and Functionality

34 The ROOT dataformat is flexible and mostly self-descriptive. Users can define their own data  
35 structures (`C++` classes) which derive from ROOT classes and serialise them into directories,  
36 trees and branches. The information about the deserialisation is written to the output file

37 (therefore: self-descriptive) but there are some basic structures and constants needed to  
 38 bootstrap the parsing process. One of the biggest advantages of the ROOT data format is the  
 39 ability to store jagged structures like nested arrays of structs with different sub-array lengths.  
 40 In high-energy physics, such structures are preferred to resemble e.g. particle interactions and  
 41 detector responses as signals from different hardware components, combined into a tree of  
 42 events.

43 UnROOT.jl understands the core structure of ROOT files, and is able to decompress and  
 44 deserialize instances of the commonly used TH1, TH2, TDirectory, TTree etc. ROOT  
 45 classes. All basic C++ types for TTree branches are supported as well, including their nested  
 46 variants. Additionally, UnROOT.jl provides a way to hook into the deserialisation process of  
 47 custom types where the automatic parsing fails. By the time of writing, UnROOT is already  
 48 used successfully in the data analysis of the KM3NeT neutrino telescope (Adriá n-Martínez  
 49 et al., 2016) and the CMS detector (Ehatäht, 2020).

50 Opening and loading a TTree lazily – i.e. without reading the whole data into memory – is  
 51 simple:

```

julia> using UnROOT

julia> f = ROOTFile("test/samples/NanoAODv5_sample.root")
ROOTFile with 2 entries and 21 streamers.
test/samples/NanoAODv5_sample.root
  Events
    "run"
    "luminosityBlock"
    "event"
    "HTXS_Higgs_pt"
    "HTXS_Higgs_y"
    ...

julia> mytree = LazyTree(f, "Events", ["Electron_dxy", "nMuon", r"Muon_(pt|eta)$"])
  Row  Electron_dxy  nMuon  Muon_eta  Muon_pt
      Vector{Float32}  UInt32  Vector{Float32}  Vector{Float32}

  1  [0.000371]      0      []         []
  2  [-0.00982]     2      [0.53, 0.229] [19.9, 15.3]
  3  []              0      []         []
  4  [-0.00157]     0      []         []
  ...
  
```

52 As seen in the above example, the entries in the columns are multi-dimensional and jagged.  
 53 The LazyTree object acts as a table which supports sequential or parallel iteration, selections  
 54 and filtering based on ranges or masks, and operations on whole columns:

```

for event in mytree
  # ... Operate on event
end

Threads.@threads for event in mytree # multi-threading
  # ... Operate on event
end

mytree.Muon_pt # a column as a lazy vector of vectors
  
```

55 The LazyTree is designed as `<: AbstractArray` which makes it compose well with the rest of  
56 the Julia ecosystem. For example, syntactic loop fusion<sup>1</sup> or Query-style tabular manipulations  
57 provided by packages like `Query.jl`<sup>2</sup> without any additional code support just work out-of-  
58 the-box.

## 59 Comparison with existing software

60 This section focusses on the comparison with other existing ROOT I/O solutions in the Julia  
61 universe, however, one honorable mention is `uproot` (Pivarski et al., 2021), which is a purely  
62 Python-based ROOT I/O library and played (still plays) an important role for the development  
63 of `UnROOT.jl` as it is by the time of writing the most complete and best documented ROOT  
64 I/O implementation.

- 65 ▪ `UpROOT.jl` is a wrapper for the aforementioned `uproot` Python package and uses  
66 `PyCall.jl`<sup>3</sup> as a bridge, which means that it relies on Python as a glue language.  
67 In addition to that, `uproot` itself utilises the C++ library `AwkwardArray` (Pivarski  
68 et al., 2018) to efficiently deal with jagged data in ROOT files. Most of the features  
69 of `uproot` are available in the Julia context, but there are intrinsic performance and  
70 usability drawbacks due to the three language architecture.
- 71 ▪ `ROOT.jl`<sup>4</sup> is one of the oldest Julia ROOT packages. It uses C++ bindings to directly  
72 wrap the ROOT framework and therefore is not limited to I/O. Unfortunately, the `Cxx.jl`<sup>5</sup>  
73 package which is used to generate the C++ glue code does not support Julia 1.4 or  
74 later. The multi-threaded features are also limited.

## 75 Conclusion

76 `UnROOT.jl` is an important package in high-energy physics and related scientific fields where  
77 the ROOT dataformat is established, since the ability to read and parse scientific data is  
78 certainly the first mandatory step to open the window to a programming language and its  
79 package ecosystem. `UnROOT.jl` has demonstrated tree processing speeds at the same level  
80 as the C++ ROOT framework in per-event iteration as well as the Python-based `uproot` library  
81 in chunked iteration.

## 82 References

- 83 Adriá n-Martínez, S., Ageron, M., Aharonian, F., Aiello, S., Albert, A., Ameli, F., Anassontzis,  
84 E., Andre, M., Androulakis, G., Anghinolfi, M., Anton, G., Ardid, M., Avgitas, T., Bar-  
85 barino, G., Barbarito, E., Baret, B., Barrios-Martí, J., Belhorma, B., Belias, A., ... Zúñiga,  
86 J. (2016). Letter of intent for KM3NeT 2.0. *Journal of Physics G: Nuclear and Particle*  
87 *Physics*, 43(8), 084001. <https://doi.org/10.1088/0954-3899/43/8/084001>
- 88 Bezanson, Jeff., Edelman, Alan., Karpinski, Stefan., & Shah, V. B. (2017). Julia: A fresh  
89 approach to numerical computing. *SIAM Review*, 59(1), 65–98. [https://doi.org/10.1137/](https://doi.org/10.1137/141000671)  
90 [141000671](https://doi.org/10.1137/141000671)

<sup>1</sup><https://julialang.org/blog/2017/01/moredots/>

<sup>2</sup><https://github.com/queryverse/Query.jl>

<sup>3</sup><https://github.com/JuliaPy/PyCall.jl>

<sup>4</sup><https://github.com/JuliaHEP/ROOT.jl>

<sup>5</sup><https://github.com/JuliaInterop/Cxx.jl>

- 91 Brun, R., & Rademakers, F. (1997). ROOT: An object oriented data analysis framework.  
92 *Nucl. Instrum. Meth. A*, 389, 81–86. [https://doi.org/10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X)
- 93 Ehatäht, K. (2020). NANOAOB: a new compact event data format in CMS. *EPJ Web Conf.*,  
94 245, 06002. <https://doi.org/10.1051/epjconf/202024506002>
- 95 Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau,  
96 D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Kerkwijk,  
97 M. H. van, Brett, M., Haldane, A., Río, J. F. del, Wiebe, M., Peterson, P., ... Oliphant,  
98 T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- 100 Pivarski, J., Osborne, I., Ifrim, I., Schreiner, H., Hollands, A., Biswas, A., Das, P., Roy  
101 Choudhury, S., & Smith, N. (2018). *Awkward array* (Version 1.9.0rc4) [Computer soft-  
102 ware]. Zenodo. <https://doi.org/10.5281/zenodo.6522027>
- 103 Pivarski, J., Schreiner, H., Smith, N., Burr, C., Kalinkin, D., Stark, G., Hartmann, N., Davis,  
104 D., O’Neil, R., Novak, A., Greiner, B., Stanislaus, B., ChristopheRappold, Deaconu, C.,  
105 Cervenkov, D., Rübenach, J., Bendavid, J., Lieret, K., Peresano, M., ... Held, A. (2021).  
106 *Scikit-hep/uproot4: 4.1.3* (Version 4.1.3) [Computer software]. Zenodo. <https://doi.org/10.5281/zenodo.5539722>
- 108 Stanitzki, M., & Strube, J. (2021). Performance of julia for high energy physics analyses.  
109 *Computing and Software for Big Science*, 5(1), 1–11.
- 110 The pandas development team. (2020). *Pandas* (latest) [Computer software]. Zenodo.  
111 <https://doi.org/10.5281/zenodo.3509134>

DRAFT