

## Streamlined data analysis in Python

---

**L. Altenkort,<sup>a</sup> D. A. Clarke,<sup>b,\*</sup> J. Goswami<sup>c</sup> and H. Sandmeyer<sup>a,1</sup>**

<sup>a</sup>*Universität Bielefeld, Fakultät für Physik,  
D-33615 Bielefeld, Germany*

<sup>b</sup>*Department of Physics and Astronomy, University of Utah,  
Salt Lake City, Utah 84112, United States*

<sup>c</sup>*RIKEN Center for Computational Science,  
Kobe 650-0047, Japan*

*E-mail:* [clarke.davida@gmail.com](mailto:clarke.davida@gmail.com)

Python is a particularly appealing language to carry out data analysis, owing in part to its user-friendly character as well as its access to well maintained and powerful libraries like NumPy and SciPy. Still, for the purpose of analyzing data in a lattice QCD context, some desirable functionality is missing from these libraries. Moreover, scripting languages tend to be slower than compiled ones. To help address these points we present the AnalysisToolbox, a collection of Python modules to facilitate lattice QCD data analysis. Some highlighted features include general-purpose jackknife and bootstrap routines; modules for reading in and storing gauge configurations; a module to carry out hadron resonance gas model calculations; and convenience wrappers for SciPy integration, curve fitting, and splines. These features are sped up behind the scenes using parallelization and just-in-time compilation.

*The 40th International Symposium on Lattice Field Theory, LATTICE2023 31st July - 4th Aug, 2023  
Batavia, Illinois, USA*

---

<sup>1</sup>Work conducted while still at Bielefeld.

\*Speaker

## 1. Motivation

Lattice QCD calculations are broadly separated into three distinct steps: one generates configurations, makes measurements on those configurations, then carries out data analysis on the measurements. These steps are roughly ranked in decreasing order of computational demand. Hence the first two steps usually have to be carried out on powerful hardware using high-performance code. The last step can often be done on one's own laptop. There are many general techniques that apply to a large class of statistical problems, so it's useful to have code where your implementation of that technique can be applied in as many situations as possible. This strategy saves you time; moreover if you've implemented the technique correctly, it prevents you from making any error having to again implement that technique any further time.

At this stage of analysis it's therefore valuable to use Python. Python has many characteristics that make it particularly appealing to lattice practitioners doing data analysis. For instance it lets you implement classes to keep code well organized; doesn't require almost any advanced knowledge of software development; has many already existing, well tested libraries<sup>1</sup> that can be easily imported or exist already in common virtual environments such as Anaconda; and is very flexible and user friendly. On the other hand, it's slow relative to compiled languages such as C++ and even the up-and-coming scripting language, Julia. At the level of data analysis this often doesn't matter, but it can, for instance when implementing systematic corrections depending on double or triple integrals or solving slowly-converging optimization problems.

In these proceedings, we present a partial answer to some of the aforementioned challenges, the *AnalysisToolbox*, which is publicly available on GitHub [1]. Statements in these proceedings apply to release `v1.1.0`. Hauke Sandmeyer wrote the basis for this package in the context of HotQCD collaboration. From that basis, we tried to keep only the most general routines, like those that can jackknife an arbitrary function of a time series. In addition, much effort was made to speed up such routines, for example through just-in-time compilation or parallelization.

## 2. An invitation

We begin with some examples how the *AnalysisToolbox* can be used to carry out physics calculations. The hadron resonance gas (HRG) model is a low-temperature statistical physics model that imagines that hadrons are the only degrees of freedom. Comparing HRG with lattice data is often utilized to, for example, determine the temperature at which hadronic states dissociate into quarks [2]. In Listing 1 we show how the `HRG` class can be used to determine the baryon-number fluctuation  $\chi_2^B$  in an HRG. As one can see from the `gen_chi` call, arbitrary conserved-charge cumulants are supported.

To complete a lattice calculation, one must extrapolate to the continuum limit; hence straightforward continuum-limit extrapolations are something one will do fairly often. In the next example, we show how the *AnalysisToolbox* can be used to perform general continuum-limit extrapolations.

---

<sup>1</sup>Nowadays machine learning is taking a major foothold in lattice field theory, with lattice practitioners developing algorithms for e.g. gauge field generation and phase transition identification. Python has access to many machine-learning libraries, and it may be useful in a lattice field theory context to leverage these.

```

1 import numpy as np
2 import latqcdtools.base.logger as logger
3 from latqcdtools.physics.HRG import HRG
4 from latqcdtools.base.readWrite import readTable, writeTable
5 from latqcdtools.base.initialize import initialize, finalize
6
7 # Write terminal output to log file. Includes git commit hash.
8 initialize('HRG.log')
9
10 # Pick a temperature range in MeV
11 T = np.arange(100, 166, 1)
12
13 # Read in hadron names, masses, charges, baryon number, strangeness,
14 # charm, and degeneracy factor. This table is provided with AnalysisToolbox.
15 QMHRG_table = '../latqcdtools/physics/HRGtables/QM_hadron_list_ext_strange_2020.txt'
16 hadrons, M, Q, B, S, C, g = readTable(QMHRG_table, usecols=(0,1,2,3,4,5,6),
17                                     dtype="U11,f8,i8,i8,i8,i8,i8")
18 w = np.array([1 if ba==0 else -1 for ba in B])
19
20 # Instantiate HRG object.
21 QMhrhg = HRG(M,g,w,B,S,Q,C)
22
23 # This computation is vectorized since T is a numpy array.
24 logger.info('Computing chi2B.')
25 chi = QMhrhg.gen_chi(T, B_order=2, Q_order=0, S_order=0, C_order=0,
26                    muB_div_T=0.3, muQ_div_T=0, muS_div_T=0, muC_div_T=0)
27
28 # Output T and chi2B as columns in this table.
29 writeTable("chi2B.txt", T, chi, header=['T [MeV]', 'chi2B (QMHRG)'])
30
31 finalize()

```

**Listing 1:** An example HRG calculation. In this case we work at fixed  $\mu_B/T = 0.3$  and compute  $\chi_2^B$  as function of temperature. What is needed is as much relevant input knowledge about hadron bound states as is known. This is read from a table in lines 16-17. The table of information is provided with the AnalysisToolbox. The list is also available as an ancillary file in the arXiv version of [3]. Line 18 determines for each species whether the gas is bosonic or fermionic. Finally the `HRG` class is instantiated in line 21, which besides cumulants like the one computed on line 25, contains many methods for various thermodynamic observables such as the pressure and entropy.

For instance suppose we want to perform a continuum-limit extrapolation to determine the deconfinement transition temperature  $T_d$  in pure SU(3). The order parameter for this phase transition is given by the Polyakov loop,  $P$ . The transition is first-order in the thermodynamic limit, where  $\langle |P| \rangle$  as function of temperature would jump discontinuously at  $T_d$ . At finite volume, this abrupt jump becomes smooth, and  $T_d$  is estimated by the inflection point of the curve.

In Listing 2 we show how such an extrapolation is achieved with the AnalysisToolbox, along with error estimation, plotting the results, and carrying out a statistical comparison with the known literature value. We assume you already<sup>2</sup> have results for  $\langle |P| \rangle$  at various  $N_\tau$ , which we read in from tables of the form `Nt6.txt`. For each  $N_\tau$ , this code estimates the inflection point of  $\langle |P| \rangle$  as a

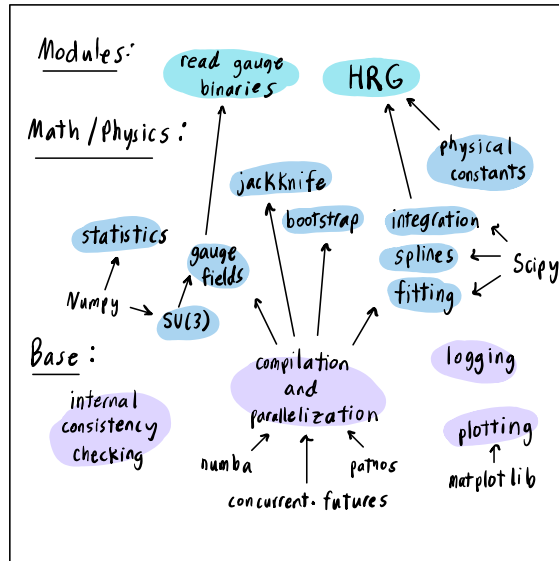
<sup>2</sup>Incidentally, the AnalysisToolbox already has code that will assist with computing some Polyakov-loop observables, such as  $|P|$  or the susceptibility  $\chi_{|P|}$ , if you started with measurements at the configuration level.

```

1 import numpy as np
2 import latqcdtools.base.logger as logger
3 from latqcdtools.base.readWrite import readTable
4 from latqcdtools.base.printErrorBars import get_err_str
5 from latqcdtools.base.initialize import initialize, finalize
6 from latqcdtools.math.num_deriv import diff_deriv
7 from latqcdtools.math.spline import getSpline
8 from latqcdtools.statistics.statistics import gaudif
9 from latqcdtools.statistics.bootstr import bootstr_from_gauss
10 from latqcdtools.physics.continuumExtrap import continuumExtrapolate
11 from latqcdtools.physics.constants import r0_phys
12 from latqcdtools.physics.lattice_params import latticeParams
13
14 initialize('cont.log')
15
16 Nts = [6,8,10,12,14,16,18,20]
17 Tds, Tderrs = [], []
18
19 for Nt in Nts:
20
21     T = []
22     Ns = Nt*3
23
24     # Read in Polyakov loop measurements,
25     beta, PM, PE = readTable('ploop/Nt'+str(Nt)+'.txt',usecols=(0,1,2))
26
27     # Create array of temperatures in physical units
28     for b in beta:
29         lp = latticeParams(Ns, Nt, b, scaleType='r0')
30         T.append( lp.getT() )
31     t = np.linspace(T[0],T[-1],1001)
32
33     # Extract Td from inflection point of <|P|> vs T using natural spline
34     def getTd(pm):
35         spl = getSpline(T, pm, natural=True)
36         dPdT = diff_deriv(t, spl)
37         maxIndex = np.argmax(dPdT)
38         return t[maxIndex]
39
40     # Error in Td estimate comes from 1000 Gaussian bootstrap samples
41     Td, Tde = bootstr_from_gauss(getTd, PM, PE, 1000)
42     Tds.append(Td)
43     Tderrs.append(Tde)
44
45     # Perform O(a^4) continuum-limit extrapolation
46     result, result_err, chidof = continuumExtrapolate( Nts, Tds, Tderrs, order=2,
47                                                       xtype="Nt", show_results=True,
48                                                       plot_results=True )
49
50     # Do a Z-test against literature result,
51     r0 = r0_phys(year=2014, units="MeVinv")
52     Tdr0, Tdr0e = r0 * result[0], r0 * result_err[0]
53     Tdr0_lit, Tdr0_lite = 0.7457, 0.0045
54     logger.info('q(ours vs. lit) =', gaudif(Tdr0,Tdr0e,Tdr0_lit,Tdr0_lite))
55
56     finalize()

```

**Listing 2:** Example continuum-limit extrapolation. As one sees in line 22, this calculation was performed on ensembles with aspect ratio 3. Excluding the imports, we were able to read in  $\langle |P| \rangle$  vs.  $T$  tables, fit those data with splines, estimate inflection points, bootstrap those estimates, continuum extrapolate those results, plot them, and compare the final estimate using a Z-test to the literature value with about 25 commands. Taken from a pedagogical project for first- and second-year undergraduate students [4].



**Figure 1:** Figure 1: Rough sketch of organizational hierarchy. Base modules encapsulate combinations of well maintained Python modules, such as SciPy and NumPy. These are used to construct and enhance math and physics objects, which in turn build up high-level modules. Here the configuration reader and hadron resonance gas model are given as two examples.

function of  $T$  to get  $T_d(N_\tau)$ . Temperatures are calculated in MeV using a 2017 parameterization [5] of the Sommer scale  $r_0/a$  [6]. This procedure is wrapped in a user-defined function<sup>3</sup> `getTc`, so that errors in the  $\langle|P|\rangle$  data can be conveniently propagated into the error in  $T_d(N_\tau)$  using a Gaussian bootstrap<sup>4</sup>, which is done on line 41. We stress that this scheme allows the bootstrap to be completely agnostic to the form and contents of the to-be-bootstrapped function.

Having the `Nts`, `Tds`, and `Tderrs`, we are ready to perform a continuum-limit extrapolation on line 46. This will perform<sup>5</sup> an extrapolation to second order in  $a^2$ , i.e.  $O(a^4)$ , print the fit results to screen, and create a plot of the extrapolation for you. The arrays `result` and `result_err` contain the best fit parameters along with their errors, with `result[0]` being the continuum-limit  $T_d$ . In the last block we compare our result with  $T_d r_0$  determined in Francis *et al.* [7]. The temperatures calculated in this code implicitly had units of MeV, hence we use the 2014 result [8] for  $r_0$  in physical units converted to 1/MeV. Finally we call `gaudif` to carry out a Z-test.

### 3. Code design

As mentioned in Sec. 1, there exist many well tested and actively maintained Python packages for various kinds of data analysis. In general we try to take the philosophy that we cannot write

<sup>3</sup>Within this function is a `getSpline` method that wraps various SciPy spline methods and `diff_deriv`, which is our implementation of a central-difference numerical derivative.

<sup>4</sup>There is also an ordinary bootstrap, in case you would rather work with data on the configuration level.

<sup>5</sup>Under the hood, this tries several different `scipy.optimize` algorithms and chooses the one with the lowest reduced  $\chi^2$ . If you would like to try higher-order fits with few data, we also support Bayesian priors.

something better than such libraries, especially given our limited personnel. Therefore, many of the methods and classes in the AnalysisToolbox are wrappers for or inherit from existing methods and classes in SciPy [9] and NumPy [10].

In Fig. 1 we outline the code’s organizational hierarchy using some example modules. At the lowest level we collect methods and classes that we think all other modules might use. This includes internal methods that streamline<sup>6</sup> and check for consistency<sup>7</sup>. Error messages are output with our custom logger, which uses parts of Python’s `logging` library. Our plotting library consists mainly of convenience wrappers for Matplotlib [11].

Most importantly at this level is a module for speeding up Python calculations. Many kinds of computations easily profit from parallelization; for example a jackknife can be easily parallelized by having each jackknife sample independently evaluated in parallel. Python has many parallelization libraries available to it. In our code, we have abstracted this away with a `parallel_function_eval` method. This takes as arguments a function and an array, then evaluates the function on the array in parallel. Hence it can be used to replace for loops. While parallelization in Python is generally less challenging than other languages, it can still be tricky to pass arguments besides the to-be-parallelized array to the function, and the abstraction helps with this. Moreover this abstraction lets one switch out different parallelizers without modifying any higher-level code.

Sections of Python code using nested C-like for loops are often tremendously slow. In some instances, they may not be easy to vectorize, and may be slow even after some straightforward parallelization. For some such cases, Python has available to it a just-in-time compilation library called Numba [12]. On the other hand, compilation also takes some time, and it could be that for small enough problem sizes, the just-in-time compilation time is similar to the run time. Hence we introduce a `@compile` decorator along with `numbaON()` and `numbaOFF()` methods to conveniently toggle compilation for large swathes of code on and off.

At the next level up in the hierarchy we have physics and math modules. When possible, we try to use already existing methods in NumPy and SciPy, for instance for curve fitting, integration, matrix multiplication, and very basic statistics. At the time of writing, numerical differentiation is slated to no longer be supported by SciPy, and we have not been able to make other automatic differentiators work well in the code, so we do our own numerical differentiation using central differences. More advanced statistical tools not already in NumPy or SciPy include the bootstrap and jackknife. As discussed in Sec. 2, the jackknife and bootstrap are agnostic to the problem at hand, which saves the user from constantly having to reimplement or tailor a jackknife approach to each new problem. Moreover they call `parallel_function_eval`, so they are sped up under the hood. Our class for SU(3) matrices inherits from NumPy matrices to leverage that already existing functionality, and we implement a few methods special to SU(3) like reunitarization. Parallelization, compilation, and the SU(3) class are synthesized in a gauge field class, so that one can store gauge fields and somewhat quickly compute average plaquettes.

<sup>6</sup>For example some methods may expect NumPy arrays, and hence return an array, even if you wanted to pass a simple scalar. The `unvector` method checks if something is a single-element, array-like object and, if it is, removes the outermost brackets.

<sup>7</sup>For instance checking whether arrays have the same length or objects are of the correct type. In particular we try to make sure that, when the code gets something it didn’t expect, it tells you exactly what values it received.

At the highest level are modules requiring a synthesis of many lower-lying modules in the code. The gauge field reader, inspired by `qcdutils` [13], is a good example of this. It reads in binary and stores the result in a gauge field object. Mirroring gauge field readers in typical lattice codes, we also check things like the average plaquette. At the moment we can only read in NERSC format binaries and lack checksum calculation.

Finally we mention our unit tests, which are intended to make sure results remain stable under changes to the code. These tests are usually very simple and include comparisons against trusted results and analytic results or comparisons between multiple computation strategies.

## 4. Existing features

### 4.1 Math and statistics

The math and statistics methods of the code are independent of any physics context and hence are applicable well beyond the lattice QCD. These methods include numerical differentiation and integration; spline fitting; curve fitting with [14] and without Bayesian priors; ordinary and Gaussian bootstrap; jackknife; autocorrelation calculation [15]; and the statistical Z-test and T-test.

### 4.2 Physics

We feature physics modules that might be of interest both to lattice practitioners and those studying finite-temperature QCD phenomenology like HotQCD parameterizations of, e.g.,  $af_K(\beta)$  and  $r_1m_s(\beta)$  [3, 8]; physical parameters and their errors from the PDG, e.g.  $m_\pi$  and  $m_\rho$ ; the hadron resonance gas model [16, 17]; the QCD equation of state [18–20]; the QCD beta function; the static quark potential and Polyakov loop observables; and critical exponents for various universality classes.

### 4.3 Interfacing

Having evolved in the context of HotQCD and MILC projects, the code interfaces with some software and conventions of these groups, for instance parsing HotQCD, MILC, and SIMULATE-QCD [21, 22] ensemble names to extract metadata and jackknifing C. Schmidt’s `DenseCode` output [23]. We also try to make the code flexible to conventions in the broader lattice community, like reading<sup>8</sup> in NERSC-format gauge configurations and reading `.gpl` files [24].

## 5. Future work and outlook

We have implemented features in this code as we encountered them in various lattice projects. Many of these methods are sufficiently general that we believe they can especially be of use to other lattice practitioners. Some methods, such as the QCD EoS, should also be useful to phenomenologists, and still others like the jackknife, bootstrap, and fitting modules are usable for general data analysis, even outside of a physics context.

As the ILDG progresses [25], we intend to dovetail some of our efforts with theirs, for instance supporting `QCdml` markup and reading in ILDG configurations.

<sup>8</sup>We have found a gauge field reader to be of use in Python, since some projects explore applying machine learning at the configuration level, and many machine learning tools already exist for Python.

## Acknowledgements

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Proj. No. 315477589-TRR 211 and Proj. No. 460248186 (PUNCH4NFDI).

We have had several other contributors to the AnalysisToolbox code over the years. In particular we thank O. Kaczmarek, L. Mazur, M. Sarkar, C. Schmidt, H.-T. Shu, and T. Ueding.

DAC would like to thank his students Kai Ebira, Daeton McClure, Coleman Rohde, and Ed Wright for producing the data showcased by the example script `main_continuumExtrapolate.py` and for their hard work and enthusiasm learning about lattice QCD. He would also like to thank University of Utah's Science Research Initiative for providing a framework to give undergraduates some opportunity to try research right away.

## References

- [1] "AnalysisToolbox." <https://github.com/LatticeQCD/AnalysisToolbox>. 10.5281/zenodo.8368581.
- [2] S. Sharma, *Charm fluctuations in (2+1)-flavor QCD at high temperature*, *PoS LATTICE2022* (2023) 191 [2212.11148].
- [3] HotQCD collaboration, *Second order cumulants of conserved charge fluctuations revisited: Vanishing chemical potentials*, *Phys. Rev. D* **104** (2021) [2107.10011].
- [4] D. A. Clarke et al., "Undergraduate project: Determination of  $T_d$  in pure SU(3) LGT." <https://github.com/clarkedavida/Computational-Methods-for-Hot-Nuclear-Matter>, 2023.
- [5] Y. Burnier, H. T. Ding, O. Kaczmarek, A. L. Kruse, M. Laine, H. Ohno et al., *Thermal quarkonium physics in the pseudoscalar channel*, *JHEP* **11** (2017) 206 [1709.07612].
- [6] R. Sommer, *A New way to set the energy scale in lattice gauge theories and its applications to the static force and  $\alpha_s$  in SU(2) Yang-Mills theory*, *Nucl. Phys. B* **411** (1994) 839 [hep-lat/9310022].
- [7] A. Francis, O. Kaczmarek, M. Laine, T. Neuhaus and H. Ohno, *Critical point and scale setting in SU(3) plasma: An update*, *Phys. Rev. D* **91** (2015) 096002 [1503.05652].
- [8] HotQCD collaboration, *Equation of state in (2+1)-flavor QCD*, *Phys. Rev. D* **90** (2014) 094503 [1407.6387].
- [9] P. Virtanen et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Methods* **17** (2020) 261.
- [10] C. R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357.
- [11] J. D. Hunter, *Matplotlib: A 2d graphics environment*, *Computing in Science & Engineering* **9** (2007) 90.



- [12] S. K. Lam, A. Pitrou and S. Seibert, *Numba: A llvm-based python jit compiler*, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6, 2015.
- [13] M. Di Pierro, *QCDUtils*, [1202.4813](#).
- [14] HPQCD collaboration, *Constrained curve fitting*, *Nucl. Phys. B Proc. Suppl.* **106** (2002) 12 [[hep-lat/0110175](#)].
- [15] B. A. Berg, *Markov Chain Monte Carlo Simulations and Their Statistical Analysis*. World Scientific, Singapore, 2004.
- [16] J. Goswami, F. Karsch, C. Schmidt, S. Mukherjee and P. Petreczky, *Comparing conserved charge fluctuations from lattice QCD to HRG model calculations*, *Acta Phys. Polon. Supp.* **14** (2021) 251 [[2011.02812](#)].
- [17] HotQCD collaboration, *Fluctuations and Correlations of net baryon number, electric charge, and strangeness: A comparison of lattice QCD results with the hadron resonance gas model*, *Phys. Rev. D* **86** (2012) 034509 [[1203.0784](#)].
- [18] HotQCD collaboration, *Equation of state and speed of sound of (2+1)-flavor QCD in strangeness-neutral matter at nonvanishing net baryon-number density*, *Phys. Rev. D* **108** (2023) 014510 [[2212.09043](#)].
- [19] HotQCD collaboration, *The isentropic equation of state of (2+1)-flavor QCD: An update based on high precision Taylor expansion and Pade-resummed expansion at finite chemical potentials*, *PoS LATTICE2022* (2023) 149 [[2212.10016](#)].
- [20] A. Bazavov et al., *The QCD Equation of State to  $O(\mu_B^6)$  from Lattice QCD*, *Phys. Rev. D* **95** (2017) 054504 [[1701.04325](#)].
- [21] L. Altenkort, D. Bollweg, D. A. Clarke, O. Kaczmarek, L. Mazur, C. Schmidt et al., *HotQCD on multi-GPU Systems*, *PoS LATTICE2021* (2022) 196 [[2111.10354](#)].
- [22] HotQCD collaboration, *SIMULATEQCD: A simple multi-GPU lattice code for QCD calculations*, [2306.01098](#).
- [23] C. R. Allton, M. Doring, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch et al., *Thermodynamics of two flavor QCD to sixth order in quark chemical potential*, *Phys. Rev. D* **71** (2005) 054508 [[hep-lat/0501030](#)].
- [24] P. Lepage and C. Gohlke, “gplepage/lqfit: lqfit version 13.0.1.” <https://doi.org/10.5281/zenodo.7931361>, 2023.
- [25] F. Karsch, H. Simma and T. Yoshie, *The International Lattice Data Grid – towards FAIR data*, *PoS LATTICE2022* (2023) 244 [[2212.08392](#)].