

Uniform SAMPLING with BOLTZmann

Matthieu Dien Normandie Université, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France
 Email: matthieu.dien@unicaen.fr MartinPépin Normandie Université, UNICAEN, ENSICAEN, CNRS, GREYC,
 14000 Caen, France
 Email: martin.pepin@unicaen.fr

Abstract—USAIN BOLTZ is a fast Python library for the uniform random generation of tree-like structures. It allows the user to specify both (1) the data structure they wish to sample, using simple combinators similar to those of context-free grammars, and (2) their memory representation. The underlying algorithms are optimised Boltzmann samplers allowing to get approximate-size uniform random generation in linear time. Experimental results show that USAIN BOLTZ matches the performance of the experimental Arbogen package for OCaml, and out-performs the Boltzmann brain Haskell library, while being easier to integrate into existing scientific tools such as Sagemath.

Index Terms—Combinatorics, Trees, Random generation

I. INTRODUCTION

Discrete structures are often used in sciences to model a concept or an object belonging to the original domain of study. For example, tree structures represent lineage relations between people (history), words, languages (linguistics), or genes (biology); and they also structure information (tree diagrams in probabilities or data structures in computer science). In order to understand the underlying patterns governing the shape of the structures, random generation has proven to be a useful experimentation tool. It is also a basic block of other algorithms. For example, in the context of software engineering, it allows to test the scalability of applications [1], to perform *unit tests* [2], [3], or to study the asymptotic complexity of algorithms. In bioinformatics, the random generation of RNA secondary structures is also used to design RNA molecules for nano-materials or therapeutics [4].

In these examples, the goal is to simulate the uniform distribution: each outcome (of the same size) has the same probability. Various generic approaches have been developed to devise uniform samplers. Notable examples include Monte Carlo techniques, based on Markov chains [5], [6]; the famous recursive method from Nijenhuis and Wilf [7], [8], [9], which applies to any data structure admitting a combinatorial decomposition; and finally the *Boltzmann method* [10], which can generically sample any data structure admitting a combinatorial *specification* in the sense of [11].

This last method generally out-performs the others by achieving *linear* time sampling if one can tolerate a small imprecision in the size of the generated object. This is why we opted for Boltzmann sampling in our library USAIN BOLTZ.

This research was partially supported by the ANR PPS project ANR-19-CE48-0014 and the “DYNNET” project, co-funded by the Normandy County Council and the European Union in the framework of the ERDF-ESF operational program 2014-2020.

Our implementation¹ has fast C++ routines in its core and is exposed as a Python [12] library so that it integrates well into the rich software environment that exists in Python for scientific computing, notably thanks to Sagemath [13]. Our implementation is open-source (GPLv3 license) to support the open science dynamics and the free software movement.

The rest of this paper is organised as follows. Section II gives some key facts on the theory behind USAIN BOLTZ. Section III describes our running example. Section IV shows how to use the library and explain some relevant features. Then Section V presents a performance comparison with other existing tools. And finally Section VI concludes with possible improvements.

II. THE BOLTZMANN METHOD IN A NUTSHELL

The Boltzmann method works on *combinatorial specifications*, which are at the core of analytic combinatorics [11]. They offer a language with basic constructions and combinators allowing to describe data structures by a system of equations. The most important such constructions are the disjoint union, the Cartesian product, and some base cases, which we illustrate here in an example. A more complete list can be found in [11, Theorem I.1 p.27]. Here is a specification of the set of binary trees in this formalism:

$$\mathcal{B} = \mathcal{E} + \mathcal{Z} \times \mathcal{B} \times \mathcal{B}. \quad (1)$$

This reads as follows: a binary tree \mathcal{B} is either a leaf \mathcal{E} or an internal node with two children $\mathcal{Z} \times \mathcal{B} \times \mathcal{B}$. The \times symbol represents a Cartesian product (you need two \mathcal{B} s to form a pair of children) and the $+$ symbol represents a *disjoint* union.

A key component of such descriptions is the notion of *size*. Here the symbol \mathcal{Z} encodes the fact that an internal node “accounts for one” in the size of a tree, hereby defining the size of a binary tree as its number of internal nodes. By opposition, the symbol \mathcal{E} , represents a single object of size zero. It is essential for the theory of Boltzmann samplers that the number of objects of a given size described by such a specification remains finite. This gives a well-founded notion of “uniform object of size n ”. Given a set \mathcal{A} described by such a specification, the Boltzmann method provides an automatic way to build algorithms for the random sampling of elements of \mathcal{A} according to a specific probability distribution: the *Boltzmann model* [10, p. 581]. It offers two guarantees. First, two objects of the same size have the same probability to be drawn (uniformity by size). And second, when coupled

¹available at <https://gitlab.com/ParComb/usain-boltz>

with (early) rejection as described in [10, p. 601], it allows to draw objects of size close to a target size n (that is in $[(1 - \epsilon)n; (1 + \epsilon)n]$ for some $\epsilon > 0$) in *linear time*.

III. EXAMPLE OF APPLICATION: RNA SECONDARY STRUCTURES

As an illustrative example that we will follow along this paper, consider the RNA (RiboNucleic Acid) secondary structures as defined in [4]. These are a simplified view of the structures resulting from the natural folding of RNA strands in their environment. According to [4], RNA strands (chains of nucleotides A , U , C , and G) fold in a hierarchical fashion where matching pairs ($A-U$ and $C-G$) create non-crossing links (chemical bounds), as modelled below:

$$\begin{aligned} S &= \mathcal{B} \times \mathcal{Z} \times (S + \mathcal{E}) + \mathcal{B} \times \mathcal{Z} \times S \times \mathcal{Z} \times (S + \mathcal{E}) \\ \mathcal{B} &= \mathcal{U}_A + \mathcal{U}_U + \mathcal{U}_C + \mathcal{U}_G \end{aligned} \quad (2)$$

Here, S represents a secondary structure, \mathcal{Z} counts nucleotides, \mathcal{E} represents an empty strand, and \mathcal{B} distinguishes between the four kinds of nucleotides using *marker* variables. The marker construction \mathcal{U}_x does not contribute to the size (like \mathcal{E}), but it “marks” and keeps a count of some features of the data structure. This allows to get statistics from the sampler and to bias the generation (see below). The second term $\mathcal{B} \times \mathcal{Z} \times S \times \mathcal{Z} \times (S + \mathcal{E})$, encodes a link between the first nucleotide $\mathcal{B} \times \mathcal{Z}$ and a matching nucleotide \mathcal{Z} , within a structure.

IV. USAIN BOLTZ

The first step to get a Boltzmann sampler using USAIN BOLTZ is to give it a combinatorial specification. They are represented in Python using a hierarchy of classes, one for each possible construction and one to represent the whole system of equations. The most notable constructions are illustrated below on the example of RNA secondary structures and the complete list can be found online in our package’s documentation².

```
B, S, z, empty = RuleName("B"),
↳ RuleName("S"), Atom(), Epsilon()
A, U, C, G = Marker("A"), Marker("U"),
↳ Marker("C"), Marker("G")
grammar = Grammar({S: B * z * z * (S + empty) +
↳ B * z * S * z * (S + empty),
B: A + U + C + G})
```

Note that to be able to bind \mathcal{B} and S in the Grammar, we must declare S and B as RuleNames beforehand. Also, since specifications are regular Python expressions it is possible to define expressions outside of the grammar as short-hands, as we did here with z and $empty$ for instance.

Then, one needs to create a Generator instance, which initialises the C++ sampler that runs under the hood, and provides an interface to it. The simplest way to call this sampler is then via the `sample` method which takes a size window as an argument and outputs a `result` object storing a generated structure (of size in the window) in `result.obj` and some statistics in `result.sizes`.

```
generator = Generator(grammar)
result = generator.sample((10000, 15000))
# result.sizes = {A: 2080, U: 2063, G: 2087,
↳ C: 2110, z: 10558}
```

It is also possible to *tune* the generator (thanks to `pa-ganini` [14]) towards producing more As and to apply size-based rejection to A , to bias its output. For instance you can tell that you want 3000 Adenine nucleotides in expectation and reject any quantity lower than 2800 as so:

```
biased_gen = Generator(grammar,
↳ expectations={A: 3000, z: 10000})
res = biased_gen.sample({z: (10000, 15000),
↳ A: (2800, 15000)})
# res.sizes = {U: 2788, C: 2760, G: 2700, A:
↳ 2803, z: 13986}
```

A. Generator optimisations

1) *Simulation*: Although the generation of objects of size in $[(1 - \epsilon)n; (1 + \epsilon)n]$ is linear in average, it has a non-negligible cost in practice because of the useless allocations incurred by the rejections. To circumvent this issue, we propose an alternative rejection procedure that “simulates” the generation of an object by only computing its size rather than actually constructing it. When a size in the targeted size window is found, we reset the pseudo random number generator to the state in which it was just before the simulation and the actual generator is run instead of the simulator. This technique gives a significant speed-up as shown in Figure 1.

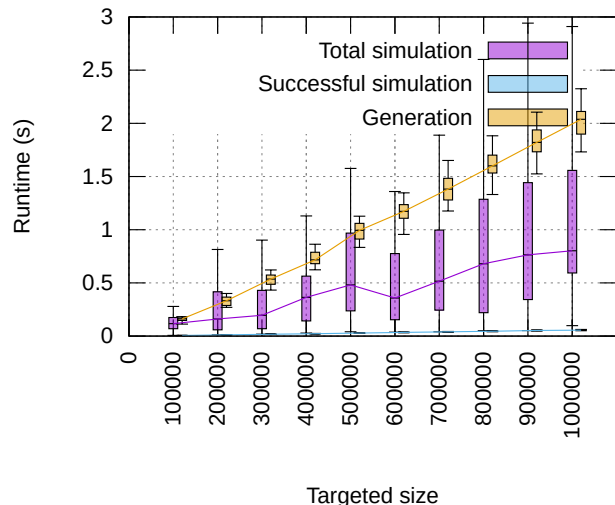


Fig. 1. Comparison of the runtime spent in the simulation and generation algorithms for RNA secondary structures (2) where $\epsilon = 0.1$ and n varies. In pink is the total simulation time, in blue is the portion of that runtime spent in the last simulation (the successful one), and in yellow is the runtime spent in the final generation. For selected values of n , we CPU-timed the algorithms 51 times with different seeds and show here their median and interquartile range (IQR) in a box plot. The whiskers extend to the most distant measure that lies within 1.5 times the IQR.

2) *Non-recursive implementation*: Another aspect of our implementation is that we implemented both the simulation and the generation routines in a non-recursive fashion. The original description of the algorithm is not tail-recursive and uses a stack space that is of the same order as the *height* of the

²usain-boltz.readthedocs.io/en/latest/generated/usainboltz.grammar.html

generated tree, that is \sqrt{n} in most cases (see [15]). This cannot be neglected since stack space is often a limited resource. Our implementation takes the form of a stack machine described in the PhD thesis of one of the authors [16, p. 120].

B. Builders

A last feature of our implementation, which we believe is unique to USAIN BOLTZ and makes it easy to integrate inside another system, is the “builder” mechanism. The idea is to let to the user of the library the choice of how the generated objects are constructed rather than enforcing one data representation. This is achieved by letting the user pass to the generator a “builder” function for each symbol of the specification. These functions take a partially built structure as an argument (a nested tuple in which all recursive sub-structures are already fully built) and complete the construction.

For instance, here is an RNA secondary structure of size 15 generated with the default builders using tuples to represents the elements of a Cartesian product:

```
(U, 'z', (C, 'z', 'epsilon'), 'z', (U, 'z',
  ↪ (A, 'z', (G, 'z', (G, 'z',
(C, 'z', 'epsilon')))), 'z', (C, 'z', (G,
  ↪ 'z', (C, 'z', (A, 'z', (U, 'z',
'epsilon'), 'z', 'epsilon'))))))
```

One possible, more suitable, data representation is to use list of markers to store the list of bases, and to store an offset together with the first base of each link. To use this representation, we must pass a builder for S to the generator. This builder will receive either a tuple of type $B * z * (S + \text{empty})$ or a tuple of type $B * z * S * z * (S + \text{empty})$ where the S objects have *already* been built. Note that we keep the default builder for B as it already does what we want.

```
dual = {A: U, U: A, C: G, G: C}
build_S_or_E = union_builder(lambda x: x,
  ↪ lambda eps: [])
def build_prefix(t): # case: B * z * (S +
  ↪ Epsilon())
    base, z, se = t
    return [base] + build_S_or_E(se)
def build_matching(t): # case: B * S * z *
  ↪ (S + Epsilon())
    base, z, s, z, se = t
    return [(base, len(s) + 1)] + s +
  ↪ [dual[base]] + build_S_or_E(se)
generator.set_builder(S,
  ↪ union_builder(build_prefix,
  ↪ build_matching))
```

Note the use of the `union_builder` function here. In the case of the union of two classes (e.g. $(S + \text{empty})$), in order to avoid potential ambiguities, the generator tags the generated object with an integer indicating from which component of the union it comes. The `union_builder` function automates the process of calling the right builder depending on this tag, so that the user only has to specify a sub-builder for each component of the union. For instance, to build an element from $(S + \text{empty})$, we must provide two builders. In case we receive an S , the object has already been built so we pass the identity function, and in case we receive an `Epsilon()`, we return the empty secondary structure `[]`. Finally we pass the builder to the generator by

calling the `set_builder` method. Using the same seed as in the previous example the generator now yields:

```
[(U, 2), C, A, (U, 5), A, G, G, C, A, C, G,
  ↪ C, (A, 2), U, U]
```

V. PERFORMANCE COMPARISON

At the moment, few tools currently implement the Boltzmann method. An OCaml prototype called Arbogen [17] has been developed by various people including the authors. It has few features but implements the simulation mechanism. Another tool called Boltzmann brain has been written in Haskell [18], [14], [19], in the form of a compiler that takes a Haskell data type as an input and produces a Haskell sampler, and also as a template Haskell library. Unfortunately, due to some (probably minor) bugs in Boltzmann brain, we could not compare its performance with that of USAIN BOLTZ on the RNA example. So we compared the three tools on the simpler example of binary trees (1) which works across all tools. The results are presented in Figure 2 and the sources of the experiments are available at <https://gitlab.com/ParComb/usain-boltz/-/tree/master/benchmarks>.

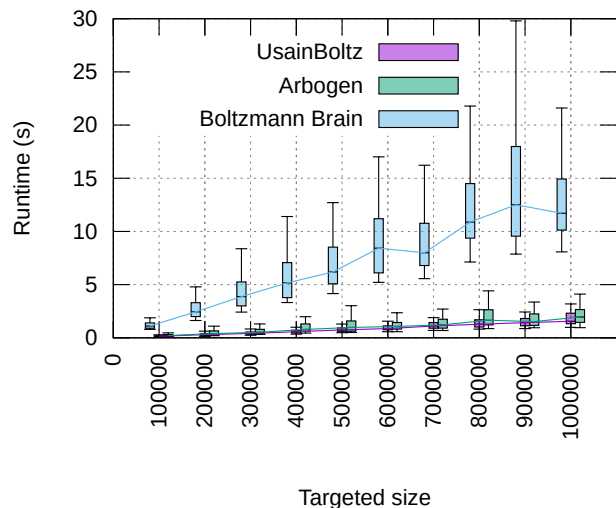


Fig. 2. Comparison of the runtimes of USAIN BOLTZ, Arbogen, and Boltzmann brain for generating binary trees. The setup is the same as for the simulation benchmark of Figure 1: for each size, run the generator 51 times and draw some box plots to describe the runtime discrepancy between the different runs. We can see a clear difference here between Boltzmann brain and the other two tools, which we explain by the simulation trick (see Section IV-A1).

VI. PERSPECTIVES

USAIN BOLTZ has more features than we could present in this paper, but the literature on Boltzmann sampling is rich of extensions we want to implement. For instance, the so-called Pólya operators, allowing to deal with symmetries, have been added to the Boltzmann framework in [20], [21]. We already implemented the most common one (MSET) but some work remains to be done in this direction. Moreover, USAIN BOLTZ also supports labelled specifications [10, Sec.4, p. 593], but no form of constrained labellings [22], [23] yet, which is also planned for the future. Finally, some internal tuning of the

algorithms is delegated to the open source `paganini` [14], [24] library. But in rare occasions we fall out of its scope and need a more general tool such as `NewtonGF` [25]. However it is bound to the proprietary Maple computer algebra system which makes the interface difficult. We would like to provide our own Python implementation of this tool.

REFERENCES

- [1] A. Mougnot, A. Darrasse, X. Blanc, and M. Soria, "Uniform random generation of huge metamodel instances," in *Model Driven Architecture - Foundations and Applications*, R. F. Paige, A. Hartman, and A. Rensink, Eds. Berlin, Heidelberg: Springer, 2009, pp. 130–145.
- [2] K. Claessen and J. Hughes, "Quickcheck: a lightweight tool for random testing of haskell programs," *ACM SIGPLAN notices*, vol. 46, no. 4, pp. 53–64, 2011.
- [3] B. Canou and A. Darrasse, "Fast and sound random generation for automated testing and benchmarking in objective caml," in *Proceedings of the 2009 ACM SIGPLAN workshop on ML*, 2009, pp. 61–70.
- [4] Y. Ponty, "Ensemble algorithms and analytic combinatorics in rna bioinformatics and beyond," 2020.
- [5] W. Jockusch, J. Propp, and P. Shor, "Random domino tilings and the arctic circle theorem," 1998.
- [6] J. G. Propp and D. B. Wilson, "Exact sampling with coupled markov chains and applications to statistical mechanics," *Random Structures & Algorithms*, vol. 9, no. 1-2, pp. 223–252, 1996.
- [7] A. Nijenhuis and H. Wilf, *Combinatorial Algorithms: For Computers and Hand Calculators*, 2nd ed. USA: Academic Press, Inc., 1978.
- [8] P. Flajolet, P. Zimmermann, and B. Van Cutsem, "A calculus for the random generation of labelled combinatorial structures," *Theoretical Computer Science*, vol. 132, no. 2, pp. 1–35, 1994.
- [9] C. Martínez and X. Molinero, "A generic approach for the unranking of labeled combinatorial classes," *Random Structures & Algorithms*, vol. 19, no. 3-4, pp. 472–497, 2001.
- [10] P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer, "Boltzmann samplers for the random generation of combinatorial structures," *Combinatorics, Probability and Computing*, vol. 13, no. 4-5, pp. 577–625, Jul. 2004.
- [11] P. Flajolet and R. Sedgewick, *Analytic Combinatorics*. Cambridge University Press, 2009.
- [12] Python Core Team, "Python." [Online]. Available: <https://www.python.org>
- [13] The Sage Developers, *SageMath, the Sage Mathematics Software System (Version 8.5.0)*, 2018, <https://www.sagemath.org>.
- [14] M. Bendkowski, O. Bodini, and S. Dovgal, "Polynomial tuning of multiparametric combinatorial samplers," in *Proceedings of the Fifteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*. SIAM, 2018, pp. 92–106.
- [15] P. Flajolet and A. M. Odlyzko, "The average height of binary trees and other simple trees," *Journal of Computer and System Sciences*, vol. 25, pp. 171–213, Jan. 1981.
- [16] M. Pépin, "Quantitative and algorithmic analysis of concurrent programs," Ph.D. dissertation, Sorbonne Université, 9 2021.
- [17] M. Dien, A. Genitrini, M. Ghanem, M. Pépin, F. Peschanski, and X. Zhan, "arbogen: a fast uniform random tree generator." [Online]. Available: <https://github.com/fredokun/arbogen>
- [18] M. Bendkowski, "boltzmann-brain: Analytic sampler compiler for combinatorial systems." [Online]. Available: <https://github.com/maciej-bendkowski/boltzmann-brain>
- [19] —, "Automatic compile-time synthesis of entropy-optimal boltzmann samplers," 2022.
- [20] P. Flajolet, É. Fusy, and C. Pivoteau, "Boltzmann sampling of unlabelled structures," in *2007 Proceedings of the Fourth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*. SIAM, 2007, pp. 201–211.
- [21] M. Bodirsky, É. Fusy, M. Kang, and S. Vigerske, "Boltzmann samplers, Pólya theory, and cycle pointing," *SIAM Journal on Computing*, vol. 40, no. 3, pp. 721–769, 2011.
- [22] O. Bodini, O. Roussel, and M. Soria, "Boltzmann samplers for first-order differential specifications," *Discrete Applied Mathematics*, vol. 160, no. 18, pp. 2563–2572, Dec. 2012.
- [23] M. Dien, "Processus concurrents et combinatoire des structures croissantes : analyse quantitative et algorithmes de génération aléatoire," Ph.D. dissertation, Université Pierre et Marie Curie - Paris VI, 9 2017. [Online]. Available: <http://www.theses.fr/2017PA066210>
- [24] M. Bendkowski, "paganini: Multiparametric tuner for combinatorial specifications." [Online]. Available: <https://github.com/maciej-bendkowski/paganini>
- [25] C. Pivoteau, B. Salvy, and M. Soria, "Algorithms for combinatorial structures: Well-founded systems and newton iterations," *Journal of Combinatorial Theory, Series A*, vol. 119, no. 8, pp. 1711–1773, Nov. 2012.