

ARTICLE

Tuning as convex optimisation: a polynomial tuner for multi-parametric combinatorial samplers

Maciej Bendkowski¹, Olivier Bodini² and Sergey Dovgal^{2*}

¹Theoretical Computer Science Department, Faculty of Mathematics and Computer Science, Jagiellonian University, Łojasiewicza 6, 30-348 Kraków, Poland and ²Institut Galilée, Université Paris 13, 99 Avenue Jean Baptiste Clément, 93430 Villetaneuse, France.

*Corresponding author. E-mail: vit.north@gmail.com

(Received 22 April 2020; revised 21 April 2021; accepted 31 May 2021; first published online 15 December 2021)

Abstract

Combinatorial samplers are algorithmic schemes devised for the approximate- and exact-size generation of large random combinatorial structures, such as context-free words, various tree-like data structures, maps, tilings, RNA molecules. They can be adapted to combinatorial specifications with additional parameters, allowing for a more flexible control over the output profile of parametrised combinatorial patterns. One can control, for instance, the number of leaves, profile of node degrees in trees or the number of certain sub-patterns in generated strings. However, such a flexible control requires an additional and nontrivial tuning procedure. Using techniques of convex optimisation, we present an efficient tuning algorithm for multi-parametric combinatorial specifications. Our algorithm works in polynomial time in the system description length, the number of tuning parameters, the number of combinatorial classes in the specification, and the logarithm of the total target size. We demonstrate the effectiveness of our method on a series of practical examples, including rational, algebraic, and so-called Pólya specifications. We show how our method can be adapted to a broad range of less typical combinatorial constructions, including symmetric polynomials, labelled sets and cycles with cardinality lower bounds, simple increasing trees or substitutions. Finally, we discuss some practical aspects of our prototype tuner implementation and provide its benchmark results.

Keywords: boltzmann samplers; multiparametric tuning; convex optimisation; combinatorial specifications; analytic combinatorics; context-free grammars; random sampling; self-concordant barriers

2020 MSC Codes: Primary: 68Q87; Secondary: 05A99, 68R05

1. Introduction

Random generation of combinatorial structures forms a prominent research area of theoretical computer science. Its wide applications include such topics as simulation of large physical statistical models [LB14], automated software testing [Pał+11, Cla+00] and counterexample construction for interactive theorem provers [Par+15], statistical analysis of queueing networks [BBR14], RNA design [Ham+19], or network theory, where one of the major challenges is to devise a realistic model of random graphs reflecting the properties of real-world networks [Bar16].

Given a formal specification defining a set of combinatorial structures, such as graphs, proteins, or tree-like data structures, we are interested in designing an efficient algorithmic sampling

Maciej Bendkowski was partially supported by the Polish National Science Center grant 2016/21/N/ST6/01032 and a French Government Scholarship awarded by the French Embassy in Poland. Olivier Bodini and Sergey Dovgal were supported by the French project ANR project MetACOnC, ANR-15-CE40-0014. The current paper is an extended version of [doi:10.1137/1.9781611975062.9](https://doi.org/10.1137/1.9781611975062.9) presented at ANALCO'18.

© The Author(s), 2021. Published by Cambridge University Press.

scheme, generating such structures according to some prescribed and rigorously controlled distribution. For instance, being interested in sampling certain plane trees following a *uniform* outcome distribution, where plane trees with an equal number of nodes share the same probability of being constructed, we want to obtain a *combinatorial sampler* satisfying these input requirements.

There exists a number of different sampling techniques in the literature. Depending on the considered class of structures, there exist different *ad-hoc* methods, such as the prominent sampler for plane binary trees due to Rémy [Ré85]. If no direct sampling technique is applicable, a common technique is to use *rejection sampling* in which one generates objects from a larger, yet simpler class and rejects unwanted samples. Although usually straightforward to implement, rejection sampling may quickly become infeasible, especially if the rejection rate grows exponentially fast. Another method, quite popular in physical modelling, is various Monte Carlo Markov Chain algorithms. If each of the chain states has an equal number of transitions, then the stationary distribution is in fact uniform. Let us remark that this technique and its modifications were successfully applied in sampling random walks and dimer models [PW96].

One of the earliest examples of a universal sampling template is Nijenhuis and Wilf's recursive method [NW78], later systematised by Flajolet, Zimmermann and Van Cutsem [FZC94]. In this approach, the input specification, as well as the combinatorial structures it defines, are recursively decomposed into primitive building blocks. Accordingly, sampling such objects follows closely the recursive structure of their specification. The generation scheme is split into two stages – an initial preprocessing phase, and the proper sampling itself. During the former, a set of decision probabilities based on a *target size* n is computed and stored for later use. The probabilities are chosen so to guarantee a uniform distribution among structures of size n constructed in the latter phase. Consequently, the sampling process reduces to a series of random decisions following precomputed distributions, dictating how to compose the output structure.

Although quite general, the recursive method poses considerable practical limitations. In both phases, the designed algorithm can manipulate integers of size exponential in the target size n , turning its effective bit complexity to $O(n^{3+o(1)})$, compared to $\Theta(n^2)$ arithmetic operations required. Denise and Zimmermann reduced later the average-case bit complexity of the recursive method to $O(n \log n)$ in time and $O(n)$ in space using a certified floating-point arithmetic optimisation [DZ99]. Regardless, worst-case space bit complexity remained $O(n^2)$ as well as bit complexity for specifications defining non-algebraic languages. Remarkably, for rational languages Bernardi and Giménez [BG12] recently linked the floating-point optimisation of Denise and Zimmermann with a specialised divide-and-conquer scheme reducing further the worst-case space bit complexity and the average-case time bit complexity to $O(n)$.

For many years, the exact-size sampling paradigm was the *de facto* standard in combinatorial generation. In many applications, however, such a precision is not necessary, and the outcome size may fluctuate around some target value n . Such a relaxed paradigm was made possible with the seminal paper of Duchon, Flajolet, Louchard and Schaeffer who proposed a universal sampler construction framework of so-called *Boltzmann samplers* [Duc+04]. The key idea of their approach is to relax the previous exact-size setting and allow for *approximate-size* samplers, generating structures within a target size window $[(1 - \varepsilon)n, (1 + \varepsilon)n]$ centred around some input size n . Like in the recursive method, Boltzmann samplers closely follow the recursive structure of the input specification. However now, instead of directly manipulating large integers or floating-point numbers in order to compute respective decision probabilities, the preprocessing phase uses numerical *oracles* to evaluate systems of generating functions corresponding to the specified combinatorial structures.

Throughout the years, a series of important extensions and improvements of Boltzmann samplers was proposed. Let us mention, for instance, linear approximate-size (and quadratic exact-size) Boltzmann samplers for planar graphs [Fus05], general-purpose samplers for unlabelled structures [FFP07], efficient samplers for plane partitions [BFP10] or the cycle pointing operator for Pólya structures [Bod+11]. The framework of Boltzmann samplers was moreover generalised onto differential specifications [BRS12, Bod+16]. Finally, let us mention linear

exact-size samplers for Catalan and Motzkin trees exploiting the shape of their holonomic specifications [BBJ13].

What was left open since the initial work of Duchon, Flajolet, Loucard, and Schaeffer was the development of (i) efficient, general-purpose Boltzmann oracles providing effective means of evaluating combinatorial systems within their disks of convergence, and (ii) an automated tuning procedure controlling the expected parameter value sizes of generated objects. The former problem was finally addressed by Pivoteau, Salvy and Soria [PSS12] who defined a rapidly converging combinatorial variant of the Newton oracle by lifting the combinatorial version of Newton's iteration of Bergeron, Labelle and Leroux [BLL98] to a new numerical level. In principle, using their Newton iteration and an appropriate use of binary search, it became possible to approximate the singularity of a given algebraic combinatorial system with arbitrarily high precision. However, even if the singularity ρ is estimated with high precision, say 10^{-10} , its approximation quality does not correspond to an equally accurate approximation of the generating function values at ρ , often not better than 10^{-2} . Precise evaluation at z close to ρ requires an extremely accurate precision of z . Fortunately, it is possible to trade-off the evaluation precision for an additional rejection phase using the idea of analytic samplers [BLR15] retaining the uniformity even with rough evaluation estimates.

Nonetheless, frequently in practical applications such as for instance software testing, uniform distribution of outcome structures might not be the most effective choice [Art+15]. In fact, it can be argued that most software bugs are minuscule, neglected corner cases, which will not be caught using large, typical instances of random data, see [Pał+11, RNL08]. In such cases, additional control over the internal structure of generated objects is required, cf. [Pał12]. Non-uniform generation schemes are also required in genomics [DRT00]. Patterns observed in real genomic sequences are tested against randomness and therefore sequences with given nucleotide frequencies need to be sampled. Random generation becomes more involved when the properties do not relate to simple motifs, but, for example, relates to secondary protein structure [Ham+19] or to evolution histories [CPW20].

In [DRT00], a multi-parametric random generation framework for context-free languages was suggested, while the question of efficient numerical tuning of the required generating function arguments was left open. In [BP10] Bodini and Ponty proposed a multidimensional Boltzmann sampler model, developing a tuning algorithm meant for the random generation of words from context-free languages with a given target letter frequency vector. This was a major improvement over [DRT00]; however, their algorithm converges only in an *a priori* unknown vicinity of the target tuning variable vector. In practice, it is therefore possible to control no more than a few tuning parameters at the same time.

In the present paper, we propose a novel polynomial-time tuning algorithm based on convex optimisation techniques, overcoming the previous convergence difficulties. We demonstrate the effectiveness of our approach with several examples of rational, algebraic, Pólya structures, and labelled structures. Remarkably, with our new method, we are able to easily handle large combinatorial systems with thousands of combinatorial classes and tuning parameters.

In order to illustrate the effectiveness of our approach, we implemented a prototype Python library called Paganini meant to provide suitable tuning vectors, and a standalone sampler generator Boltzmann Brain. Our software is freely available as open source^{1,2}.

The paper is structured as follows. In Section 2 we outline basic concepts of combinatorial random sampling, including the principles of the recursive method and Boltzmann sampling. In Section 2.2 we introduce a new (to our best knowledge) admissible *combination* operation of combinatorial classes, and explain its usefulness in certain auxiliary transformations for combinatorial specifications used during the construction of the tuning problem. In the following Section 3, we show how to express the tuning problem as a convex optimisation problem and

¹<https://github.com/maciej-bendkowski/paganini>

²<https://github.com/maciej-bendkowski/boltzmann-brain>

provide important simplifications for context-free grammars, labelled structures, increasing trees and other types of structures. In Section 4, we provide a detailed complexity analysis of the convex optimisation problems obtained from various combinatorial specifications using the notion of self-concordant barriers. Then, in Section 5 we address the problem of finding an optimal target expectation for Boltzmann samplers when the target size is constrained in a finite interval $[n_1, n_2]$. Next, in Section 6 we describe our prototype implementation. Finally, in Section 7 we illustrate the effectiveness of our approach providing several exemplary applications of our tuning algorithm.

2. Combinatorial random sampling

2.1 Admissible combinatorial classes

Let us consider the neutral class \mathcal{E} , commonly denoted as 1, consisting of a single object of size zero, and its atomic counterpart \mathcal{Z} , which is a class consisting of a single object of size one. Both are equipped with a finite set of admissible operators, such as the disjoint union $+$, Cartesian product \times , and sequence SEQ, see [FS09, Section I.2]. Depending on whether we consider labelled or unlabelled structures, we allow for more expressive admissible operators including for instance the multiset MSET, set SET, or cycle CYC constructions. In such a setting, combinatorial specifications we consider in the current paper are finite systems of equations (possibly mutually recursive) built from elementary classes \mathcal{E} , \mathcal{Z} , and admissible operators.

Example 1. Consider the following joint specification for \mathcal{T} and \mathcal{Q} . In the combinatorial class \mathcal{T} of trees, nodes at even level (the root starts at level one) have either no or exactly two children, whereas each node at odd level has an arbitrary number of non-planarily ordered children:

$$\begin{cases} \mathcal{T} = \mathcal{Z} \text{ MSET}(\mathcal{Q}) \\ \mathcal{Q} = \mathcal{Z} + \mathcal{Z}\mathcal{T}^2 \end{cases} \quad (1)$$

In order to distinguish, in other words *mark*, some additional combinatorial parameters we consider the following natural multivariate extension of specifiable classes.

Definition 2 (Specifiable k -parametric combinatorial classes). A specifiable k -parametric combinatorial class is a combinatorial specification built, in a possibly recursive manner, from k distinct atomic classes $\mathcal{Z}_1, \dots, \mathcal{Z}_k$, the neutral class \mathcal{E} , and admissible operators. In particular, a vector $\mathcal{C} = (\mathcal{C}_1, \dots, \mathcal{C}_m)$ forms a specifiable k -parametric combinatorial class if its specification can be written down as

$$\begin{cases} \mathcal{C}_1 = \Phi_1(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k) \\ \vdots \\ \mathcal{C}_m = \Phi_m(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k) \end{cases}, \quad (2)$$

where the right-hand side expressions $\Phi_i(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k)$ are composed from $\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k$, admissible operators, and the neutral class \mathcal{E} .

Example 3. Let us continue our running example, see (1). Note that we can introduce two additional marking classes \mathcal{U} and \mathcal{V} , into (2), each of weight zero, turning our example into a k -specifiable combinatorial class:

$$\begin{cases} \mathcal{T} = \mathcal{U}\mathcal{Z} \text{ MSET}(\mathcal{Q}), \\ \mathcal{Q} = \mathcal{V}\mathcal{Z} + \mathcal{Z}\mathcal{T}^2. \end{cases} \quad (3)$$

Here, \mathcal{U} is meant to mark the occurrences of nodes at odd levels, whereas \mathcal{V} is meant to mark leaves at even levels. In effect, we *decorate* the univariate specification with explicit information regarding the internal structural patterns of our interest.

Much like in their univariate variants, k -parametric combinatorial specifications are naturally linked to ordinary multivariate generating functions, see [FS09, Chapter III].

Definition 4 (Multivariate generating functions). The multivariate ordinary generating function $C(z_1, \dots, z_k)$ in variables z_1, \dots, z_k associated to a specifiable k -parametric combinatorial class \mathcal{C} is defined as

$$C(z_1, \dots, z_k) = \sum_{p_1 \geq 0, \dots, p_k \geq 0} c_p \mathbf{z}^p, \tag{4}$$

where $c_p = c_{p_1, \dots, p_k}$ denotes the number of structures with p_i atoms of type \mathcal{Z}_i , and \mathbf{z}^p denotes the product $z_1^{p_1} \cdots z_k^{p_k}$. In the sequel, we call p the (*composition*) size of the structure.

In this setting, we can easily lift the usual univariate generating function construction rules associated with admissible constructions to the realm of multivariate generating functions associated to specifiable k -parametric combinatorial classes.

2.2 Combination operator

Many examples of combinatorial structures are naturally expressed as collections of disjoint unions and do not require to explicitly *exclude* certain undesired configurations. Consequently, right-hand side expressions $\Phi_i(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k)$ of their combinatorial specifications tend to be composed of summands with positive coefficients. One notable exception, however, is specifications constructed using the *inclusion-exclusion principle* which, due to the induced class subtraction, may cause considerable difficulties for both tuning and sampling.

From the sampling perspective, exclusion in the combinatorial specification requires additional rejection, whose cost may undesirably dominate over the cost of sampling. Even worse, later in Section 3 we will argue that, in general, negative coefficients do not allow to express the tuning problem in convex optimisation form.

Remarkably, in some cases it is possible to rewrite the initial system with negative coefficients in such a way that the resulting system contains only positive terms. A typical example of such a situation is the construction of a non-empty product containing at most one object from each of the classes $\mathcal{C}_1, \dots, \mathcal{C}_d$. Symbolically:

$$(1 + \mathcal{C}_1)(1 + \mathcal{C}_2) \cdots (1 + \mathcal{C}_d) - 1. \tag{5}$$

Note that (5) is a combinatorial class consisting of non-empty tuples of length d , such that their i -th coordinate is either \mathcal{E} or an object from \mathcal{C}_i . The above concise specification explicitly excludes the empty tuple, however introduces subtraction in the specification. We can eliminate the subtraction by expanding all brackets in (5), however, such a naive transformation produces $2^d - 1$ summands in the resulting specification. Instead, we propose another transformation.

Definition 5 (Combination operator). Let $\mathcal{C}_1, \dots, \mathcal{C}_d$ be combinatorial classes. The admissible k -combination or k -selection $\text{SELECT}_k(\mathcal{C}_1, \dots, \mathcal{C}_d)$ of $\mathcal{C}_1, \dots, \mathcal{C}_d$ is a combinatorial class consisting of all tuples (c_1, \dots, c_d) such that

- c_i is either empty, that is \mathcal{E} , or an element of \mathcal{C}_i , and
- exactly $\binom{d}{k}$ elements of (c_1, \dots, c_d) are non-empty.

Example 6. Consider the classes $\mathcal{C}_1, \dots, \mathcal{C}_d$. Note that

$$\text{SELECT}_k(\mathcal{C}_1, \dots, \mathcal{C}_d) \cong \sum_{1 \leq i_1 < i_2 < \dots < i_k \leq d} \mathcal{C}_{i_1} \times \mathcal{C}_{i_2} \times \dots \times \mathcal{C}_{i_k}. \tag{6}$$

In words, the class $\text{SELECT}_k(\mathcal{C}_1, \dots, \mathcal{C}_d)$ is isomorphic with the disjoint union of all ordered k -products of classes in $\mathcal{C}_1, \dots, \mathcal{C}_d$. In particular, $\text{SELECT}_0(\mathcal{C}_1, \dots, \mathcal{C}_d) \cong 1$. Furthermore, it holds

$\text{SELECT}_d(C_1, \dots, C_d) \cong \prod_i C_i$. It should be noted that the isomorphism in (6) cannot be replaced with strict equality, as $\text{SELECT}_k(C_1, \dots, C_d)$ consists of tuples of length d whereas the right-hand side sum consists of tuples of length k .

Let us denote $\text{SELECT}_k(C_1, \dots, C_d)$ as \mathcal{S}_k . Note that

$$(1 + C_1)(1 + C_2) \cdots (1 + C_d) - 1 = \mathcal{S}_1 + \cdots + \mathcal{S}_d. \tag{7}$$

The resulting classes \mathcal{S}_k are, in fact, *elementary symmetric polynomials* consisting of all $\binom{d}{k}$ summands of the expanded $\prod_{i=1}^d (1 + C_i)$ of length k . Since a direct, verbose representation of $\text{SELECT}_k(C_1, \dots, C_d)$ is exponential in d , we suggest an indirect dynamic programming approach, allowing us to obtain a much more succinct representation of each of the selection operations using a total of $O(d^2)$ of auxiliary combinatorial classes.

Let $\mathcal{P}_{k,m}$ be a subset of \mathcal{S}_k in which products are restricted to involve classes from C_1, \dots, C_m . Note that $\mathcal{P}_{k,m}$ is empty if and only if $m < k$. Moreover, $\mathcal{P}_{k,d} = \mathcal{S}_k$. We start with

$$\begin{aligned} \mathcal{P}_{1,1} &= C_1 \\ \mathcal{P}_{1,m+1} &= \mathcal{P}_m + C_{m+1}. \end{aligned} \tag{8}$$

Now, suppose that we have computed $\mathcal{P}_{0,m}, \mathcal{P}_{1,m}, \dots, \mathcal{P}_{k,m}$ for all values of m , and wish to compute the next row of classes corresponding to $k + 1$. We start to iterate m in increasing order. For all m such that $m < k + 1$, we set $\mathcal{P}_{k+1,m} = \emptyset$. Otherwise, if $m \geq k + 1$ we note that

$$\mathcal{P}_{k+1,m} = \mathcal{P}_{k+1,m-1} + C_m \times \mathcal{P}_{k,m-1}. \tag{9}$$

The correctness of the above scheme can be proven by induction. With its help, we can compute a lower-triangular matrix of the symbolic representations for $\mathcal{P}_{k,m}$ and so also the elementary symmetric polynomials \mathcal{S}_k . As a by-product, we can therefore efficiently rewrite the initial system (5) into an equivalent one with positive terms.

This transformation can be used as an auxiliary construction in other combinatorial specifications. For example, we apply the described technique in Section 7.4 in order to sample from a multi-parametric class

$$\text{MSET}(\text{SEQ}(\mathcal{Z}_1)\text{SEQ}(\mathcal{Z}_2) \cdots \text{SEQ}(\mathcal{Z}_d) - 1). \tag{10}$$

2.3 Boltzmann samplers and the recursive method

There exists a number of different sampling techniques in the literature. In the current paper, we focus on two most prominent, general purpose frameworks — the recursive method [NW78] and Boltzmann samplers [Duc+04]. In what follows, we focus specifically on their distinguishing features.

2.3.1 The recursive method

The three basic building bricks of the *recursive method* are, similarly to the symbolic method and admissible classes, disjoint union, Cartesian product, and pointing. In this framework, the counting sequences of the combinatorial classes need to be readily available, as they determine the branching probabilities of the random generation process. The three operations are processed as follows.

Disjoint union. If the counting sequence $(C_n)_{n \geq 0}$ of the target class \mathcal{C} is given by $C_n = A_n + B_n$ where A_n and B_n are the counting sequences of the classes \mathcal{A} and \mathcal{B} , respectively, then with probability $\frac{A_n}{A_n + B_n}$ an object from \mathcal{A} is constructed, otherwise an object from class \mathcal{B} is constructed.

Cartesian product. If the counting sequence of the target class \mathcal{C} satisfies the equation $C_n = \sum_{k=0}^n A_k \cdot B_{n-k}$, then a tuple of sizes $(k, n - k)$ is chosen with probability $\frac{A_k B_{n-k}}{C_n}$ and objects from \mathcal{A} and \mathcal{B} of respective sizes k and $n - k$ are generated.

Pointing. If \mathcal{A} is a combinatorial class, then an object in a pointed class $\Theta\mathcal{A}$ is isomorphic to an object from \mathcal{A} which has a distinguished atom. Now, if the counting sequence of \mathcal{A} is A_n , then the counting sequence of $\Theta\mathcal{A}$ is given by $n \cdot A_n$. Having one of the samplers for \mathcal{A} or $\Theta\mathcal{A}$, it is easy to obtain the other one by either distinguishing a label uniformly at random or, the other way round, by forgetting which label is distinguished.

Note that the pointing operator substantially enriches the set of admissible specifications and plays a central rôle in the design of recursive samplers. For example, the two labelled operations, SET and CYC are bijectively transformed using the pointing operation as follows:

$$\Theta_{\text{SET}}(\mathcal{A}) = \text{SET}(\mathcal{A}) \times \Theta\mathcal{A} \quad \text{and} \quad \Theta_{\text{CYC}}(\mathcal{A}) = \text{SEQ}(\mathcal{A}) \times \Theta\mathcal{A}. \tag{11}$$

We are not aware of a direct and efficient sampling algorithm based on the recursive method for unlabelled structures involving the MSET and CYC operators.

2.3.2 Boltzmann samplers

While the recursive method is applicable to specifications irrespective of the analyticity of their generating functions, Boltzmann samplers work only with classes whose generating functions are analytic. However, unlike the recursive method, Boltzmann samplers do not require the underlying counting series. Instead, they rely on the values of the generating functions. As a consequence, the size of the outcome structure is no longer fixed, but follows a *Boltzmann distribution* $\mathbb{P}(\text{size} = k) = \frac{A_k z^k}{\sum_{n \geq 0} A_n z^n}$ with z being its parameter. Nevertheless, conditioned on size, such samplers generate a uniformly chosen object from the target class.

Boltzmann samplers support three basic operations, that is disjoint union, Cartesian product, and substitution. For instance, assuming that $\varphi(x)$ is analytic, we can consider a family of simply generated trees satisfying

$$T(z) = z\varphi(T(z)). \tag{12}$$

Let us remark however, that while substitutions of type $\varphi(T(z))$ are usually easier to handle, substitutions in form of $T(\varphi(z))$ are much more involved. Note that the recursive method does not easily support substitutions. Now, let us focus on how these three basic operations are processed.

Disjoint union. Consider a target class \mathcal{C} with a generating function $C(x)$. Let \mathcal{A} and \mathcal{B} be two combinatorial classes with generating functions $A(x)$ and $B(x)$, respectively, such that $\mathcal{C} = \mathcal{A} + \mathcal{B}$. Then, with probability $\frac{A(x)}{A(x)+B(x)}$ an object from \mathcal{A} is drawn, otherwise an object from class \mathcal{B} is generated.

Cartesian product. Consider a target class \mathcal{C} with a generating function $C(x)$. Let \mathcal{A} and \mathcal{B} be two combinatorial classes with generating functions $A(x)$ and $B(x)$, respectively, such that $\mathcal{C} = \mathcal{A} \times \mathcal{B}$. Then, an independent pair of recursively generated objects from \mathcal{A} and \mathcal{B} is drawn.

Substitution. Let $C(x) = \varphi(B(x))$ where $\varphi(t) = \sum_i \phi_i t^i$. Then, a Boltzmann sampler for \mathcal{C} is obtained as follows. Fix $t = B(x)$ and sample a random integer k from the distribution $\mathbb{P}(k) = \phi_k t^k / \varphi(t)$. Finally, draw k independent copies of recursively sampled objects from \mathcal{B} .

The substitution rule, in particular, provides Boltzmann samplers for the labelled SET, SEQ and CYC constructions. It turns out that such samplers also cover a wide interesting family of combinatorial classes and constructions, including unlabelled structures, first-order differential specifications, Hadamard product and Dirichlet generating series, see [Bod10] for further details.

We summarise the most common rules in Tables 1, 2, and 3. We write $X \implies \Gamma$ to denote a procedure generating random objects from Γ based on the random discrete distribution X – we draw an integer r from X and then, repeatedly and independently, invoke the respective sampler Γ r times. As a result, we obtain an r -tuple of sampled objects. For more details we refer the reader to [Duc+04] and [FFP07].

Table 1. Multivariate generating functions and corresponding Boltzmann samplers $\Gamma\mathcal{C}(\mathbf{z})$.

Class	Description	$C(\mathbf{z})$	$\Gamma\mathcal{C}(\mathbf{z})$
Neutral	$\mathcal{C} = \{\varepsilon\}$	$C(\mathbf{z}) = 1$	ε
Atom	$\mathcal{C} = \{\square_i\}$	$C(\mathbf{z}) = z_i$	\square_i
Union	$\mathcal{C} = \mathcal{A} + \mathcal{B}$	$A(\mathbf{z}) + B(\mathbf{z})$	$\text{Bern}\left(\frac{A(\mathbf{z})}{C(\mathbf{z})}, \frac{B(\mathbf{z})}{C(\mathbf{z})}\right) \longrightarrow \Gamma\mathcal{A}(\mathbf{z}) \mid \Gamma\mathcal{B}(\mathbf{z})$
Product	$\mathcal{C} = \mathcal{A} \times \mathcal{B}$	$A(\mathbf{z}) \times B(\mathbf{z})$	$(\Gamma\mathcal{A}(\mathbf{z}); \Gamma\mathcal{B}(\mathbf{z}))$
Sequence	$\mathcal{C} = \text{SEQ}(\mathcal{A})$	$(1 - A(\mathbf{z}))^{-1}$	$\text{Geom}(1 - A(\mathbf{z})) \implies (\Gamma\mathcal{A}(\mathbf{z}))$

Table 2. Multivariate generating functions and corresponding Boltzmann samplers $\Gamma\mathcal{C}(\mathbf{z})$ for unlabelled constructions.

Class	Description	$C(\mathbf{z})$	$\Gamma\mathcal{C}(\mathbf{z})$
MultiSet	$\text{MSET}(\mathcal{A})$	$\exp\left(\sum_{m=1}^{\infty} \frac{1}{m} A(\mathbf{z}^m)\right)$	see [Bod10, Algorithm 1]
Cycle	$\text{CYC}(\mathcal{A})$	$\sum_{m=1}^{\infty} \frac{\varphi(m)}{m} \ln \frac{1}{1 - A(\mathbf{z}^m)}$	see [Bod10, Algorithm 2]

Table 3. Multivariate generating functions and corresponding Boltzmann samplers $\Gamma\widehat{\mathcal{C}}(\mathbf{z})$ for labelled constructions.

Class	Description	$\widehat{C}(\mathbf{z})$	$\Gamma\widehat{\mathcal{C}}(\mathbf{z})$
Set	$\text{SET}(\mathcal{A})$	$\exp(\widehat{A}(\mathbf{z}))$	$\text{Pois}(\widehat{A}(\mathbf{z})) \implies (\Gamma\widehat{A}(\mathbf{z}))$
Cycle	$\text{CYC}(\mathcal{A})$	$\log(1 - \widehat{A}(\mathbf{z}))^{-1}$	$\text{Loga}(\widehat{A}(\mathbf{z})) \implies (\Gamma\widehat{A}(\mathbf{z}))$

2.3.3 Multi-parametric Boltzmann samplers

Consider a multi-parametric combinatorial class \mathcal{C} with a multivariate generating function $C(\mathbf{z})$. Let $\omega \in \mathcal{C}$ be a combinatorial object with composition size \mathbf{p} . Then, a multi-parametric Boltzmann sampler $\Gamma\mathcal{C}(\mathbf{z})$ outputs ω with probability

$$\mathbb{P}_{\mathbf{z}}(\omega) = \frac{\mathbf{z}^{\mathbf{p}}}{C(\mathbf{z})}. \tag{13}$$

Such samplers can be constructed from multi-parametric combinatorial specifications in the same way as ordinary Boltzmann samplers are constructed. When the expressions of generating functions involve different values of the tuning variables, then these variables, after substitution, yield new branching probabilities.

Proposition 7 (Log-exp transform of the tuning problem). *Let $\mathbf{N} = (N_1, \dots, N_k)$ be the random vector where N_i equals the number of atoms of type \mathcal{Z}_1 in a random combinatorial structure returned by the k -parametric Boltzmann sampler $\Gamma\mathcal{C}(\mathbf{z})$. Then, the expectation vector $\mathbb{E}_{\mathbf{z}}(\mathbf{N})$ and the covariance matrix $\text{Cov}_{\mathbf{z}}(\mathbf{N})$ are given by*

$$\mathbb{E}_{\mathbf{z}}(N_i) = \left. \frac{\partial}{\partial \xi_i} \log C(e^{\xi}) \right|_{\xi = \log \mathbf{z}} \quad \text{and} \quad \text{Cov}_{\mathbf{z}}(\mathbf{N}) = \left[\left. \frac{\partial^2}{\partial \xi_i \partial \xi_j} \log C(e^{\xi}) \right]_{i,j=1}^k \right|_{\xi = \log \mathbf{z}}. \tag{14}$$

Hereafter, we use $e^{\mathbf{z}}$ to denote coordinate-wise exponentiation.

Proof. Let the following *nabla* notation denote the vector of derivatives (so-called gradient vector) with respect to the variable vector $\mathbf{z} = (z_1, \dots, z_k)$:

$$\nabla_{\mathbf{z}}f(\mathbf{z}) = \left(\frac{\partial}{\partial z_1}f(\mathbf{z}), \dots, \frac{\partial}{\partial z_k}f(\mathbf{z}) \right)^\top. \tag{15}$$

Let $(c_n)_{n \geq 0}$ be the counting sequence of the combinatorial class \mathcal{C} . The probability generating function $p(\mathbf{u}|\mathbf{z})$ for N with \mathbf{u} as an argument and \mathbf{z} as a parameter takes the form

$$p(\mathbf{u}|\mathbf{z}) = \sum_{n \geq 0} \frac{c_n \mathbf{z}^n \mathbf{u}^n}{C(\mathbf{z})} = \frac{C(\mathbf{u} \bullet \mathbf{z})}{C(\mathbf{z})}, \tag{16}$$

where \bullet denotes component-wise vector multiplication. Hence, the expected value and the covariance of N can be immediately expressed through its probability generating function as

$$\begin{aligned} \mathbb{E}_{\mathbf{z}}N &= \nabla_{\mathbf{u}}p(\mathbf{u}|\mathbf{z})|_{\mathbf{u}=\mathbf{1}} \\ \text{Cov}_{\mathbf{z}}(N) &= \left[\nabla_{\mathbf{u}}^2 p(\mathbf{u}|\mathbf{z}) + \text{diag}(\nabla_{\mathbf{u}}p(\mathbf{u}|\mathbf{z})) - \nabla_{\mathbf{u}}p(\mathbf{u}|\mathbf{z})\nabla_{\mathbf{u}}^\top p(\mathbf{u}|\mathbf{z}) \right]_{\mathbf{u}=\mathbf{1}}. \end{aligned} \tag{17}$$

The proof is finished by expanding the log-exp transform in (14) and comparing the result to the expressions obtained from the probability generating functions. \square

Remark 8. The function $\gamma(\mathbf{z}) := \log C(e^{\mathbf{z}})$ is convex because its matrix of second derivatives, as a covariance matrix, is positive semi-definite inside the set of convergence. This crucial assertion will later prove central to the design of our tuning algorithm. The expressions for the expectation and the covariance matrix are similar to those obtained in [BR83] for central limit theorem for multivariate generating functions.

Remark 9. Uni-parametric recursive samplers of Nijenhuis and Wilf take, as well as Boltzmann samplers, a system of generating functions as their input. This system can be modified by putting fixed values of tuning variables, in effect altering the corresponding branching probabilities. The resulting distribution of the random variable corresponding to a weighted recursive sampler coincides with the distribution of the Boltzmann-generated variable conditioned on the structure size. As a corollary, the tuning procedure that we discuss in the following section is also valid for the exact-size approximate-frequency recursive sampling.

2.4 Complexity of exact parameter sampling

While the current paper is devoted to tuning of the parameters *in expectation*, let us pause for a moment and ask the following, natural question – what is the complexity of *exact-parameter* sampling for multi-parametric combinatorial specifications?

In the current section, we show that, unless both the classes of decision problems solvable in randomised RP and nondeterministic NP polynomial time are equal, then already for unambiguous context-free languages there exists no fully polynomial-time algorithm for almost-uniform exact-parameter sampling problem. Since it is widely conjectured that $\text{RP} \neq \text{NP}$, cf. [WG01], this infeasibility result justifies parameter tuning *in expectation*, which can be regarded as a continuous relaxation of the exact-parameter problem variant.

Consider a context-free grammar G with derivation rules in form of

$$A_i \rightarrow T_{i,j}, \tag{18}$$

where A_i is a non-terminal symbol, and the right-hand side expression $T_{i,j}$ is a (possibly empty) word consisting of both terminal and non-terminal symbols. Recall that the context-free grammar G is said to be *unambiguous* if each word its generates has a unique derivation. Let a_1, \dots, a_d be distinct terminal symbols. The exact multi-parametric sampling problem for unambiguous

context-free grammars can be stated as follows – given natural numbers n_1, n_2, \dots, n_d , sample uniformly at random a word of length $n = n_1 + \dots + n_d$ from the language $L(G)$ generated by G , such that the number of occurrences of each terminal symbol a_j is equal to n_j .

Example 10. As a simple illustrating example of the discussed problem, consider the following grammar B generating (unambiguously) all binary words over the alphabet $\Sigma = \{0, 1\}$:

$$B \rightarrow 0B \mid 1B \mid \varepsilon. \tag{19}$$

Recall that ε denotes the empty word. In this example, given numbers n_0 and n_1 , the multi-parametric sampling problem asks to generate a uniformly random binary word over Σ which has exactly n_0 0s, and n_1 1s.

In what follows we prove that, in general, the problem of exact-size multi-parametric sampling (even if the specification does not involve loops) can be reduced to the $\#P$ -complete $\#2$ -SAT problem, which asks to count the number of satisfiable variable assignments of a given 2-CNF formula. As suggested to us by Sergey Tarasov in personal communication, such a complexity result might be folklore knowledge; however, we did not manage to find it in the literature. We therefore take the liberty to fill this gap. Detailed definitions from complexity theory can be found in the papers referenced during the proof of the following theorem.

Theorem 11 (Infeasibility of exact parameter sampling). *Unless $NP = RP$, there is no fully polynomial-time algorithm for almost-uniform multi-parametric sampling from unambiguous context-free grammars.*

Proof. We start by showing that extracting the coefficients $[z_1^{k_1} \dots z_m^{k_m}]F_i(z_1, \dots, z_m)$ of a multivariate generating $F_i(z_1, \dots, z_m)$ satisfying a system of polynomial equations in form of

$$\begin{cases} F_1(z_1, \dots, z_m) = \Phi_1(z_1, \dots, z_m, F_1, \dots, F_n) \\ \vdots \\ F_n(z_1, \dots, z_m) = \Phi_n(z_1, \dots, z_m, F_1, \dots, F_n) \end{cases}, \tag{20}$$

is $\#P$ -hard. We proceed by reduction from the $\#P$ -complete $\#2$ -SAT problem.

Consider a 2-SAT formula

$$F = \bigwedge_{j=1}^m (\alpha_j \vee \beta_j), \tag{21}$$

with n Boolean variables x_1, \dots, x_n and m clauses. For each clause $(\alpha_j \vee \beta_j)$ we create a distinct (complex) variable c_j . Next, for each literal $x \in \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$ we introduce a corresponding multivariate generating function $X(c_1, \dots, c_m)$ (where $X \in \{X_1, \dots, X_n, \bar{X}_1, \dots, \bar{X}_n\}$) defined as

$$X(c_1, \dots, c_m) = \prod_{j=1}^m c_j^{\mathbf{1}_{c_j}(x)}, \tag{22}$$

where $x \mapsto \mathbf{1}_A(x)$ denotes the indicator function with respect to a set A (i.e. $\mathbf{1}_A(x) = 1$ if $x \in A$, and $\mathbf{1}_A(x) = 0$ if $x \notin A$). Finally, we create a multivariate generating function

$$H(c_1, \dots, c_m) = \prod_{i=1}^n (X_i(c_1, \dots, c_m) + \bar{X}_i(c_1, \dots, c_m)). \tag{23}$$

Clearly, all of these generating functions form a polynomial system of equations. Intuitively, generating functions $X_i(c_1, \dots, c_m)$ are products of variables c_j corresponding to clauses which

become satisfied once the respective literal x_i is true. Furthermore, $H(c_1, \dots, c_m)$ encodes all possible variable assignments. Indeed, if we expand all brackets in (23) we obtain 2^n summands, each corresponding to a distinct variable assignment – occurrences of $X_i(c_1, \dots, c_m)$ in each summand encode setting the respective variable x_i to true, whereas $\bar{X}_i(c_1, \dots, c_m)$ encode setting the respective variable x_i to false.

Let us consider an arbitrary summand H_φ in the expanded (23) corresponding to some variable assignment $\varphi: \{x_1, \dots, x_n\} \rightarrow \{\text{True}, \text{False}\}$. Note that once we unfold the respective definitions of $X_i(c_1, \dots, c_m)$ and $\bar{X}_i(c_1, \dots, c_m)$, H_φ becomes a monomial consisting of c_j 's satisfied by φ . Each clause $(\alpha_j \vee \beta_j)$ in F consists of two literals, hence its corresponding variable c_j can occur at most twice in H_φ . Consequently, the number of satisfiable assignments satisfying F is equal to

$$[c_1^{\geq 1} \dots c_n^{\geq 1}]H(c_1, \dots, c_n), \tag{24}$$

that is to the number of monomials H_φ in which each clause variable occurs at least once. In order to obtain the number (24) of satisfiable assignments using a single, exact coefficient extraction, we note that

$$[c_1^2 \dots c_n^2]H(c_1, \dots, c_n) \prod_{i=1}^n (1 + c_i) = [c_1^{\geq 1} \dots c_n^{\geq 1}]H(c_1, \dots, c_n). \tag{25}$$

And so, it is as hard to extract the coefficients of a generating function as solving a #2-SAT instance. Nevertheless, it should be noticed that not every hard enumeration problem automatically corresponds to a hard uniform random sampling problem.

In order to complete the proof, we use the celebrated [JVV86, Theorem 6.4] which proves that if there exists a fully polynomial almost-uniform random sampling algorithm (with an exponentially small error), then there exists a fully polynomial randomised approximation scheme for the counting problem as well (within a polynomially small error). Moreover, unless $\text{RP} = \text{NP}$, there exists not fully polynomial randomised approximation scheme for #2-SAT [WG01]. Hence, indeed the theorem statement must hold. \square

3. Multi-parametric tuning

A combinatorial specification typically involves several classes $\mathcal{C}_1, \dots, \mathcal{C}_m$ which are defined in a mutually recursive manner. Let us denote $(\mathcal{C}_1, \dots, \mathcal{C}_m)$ as \mathcal{C} . For tuning and sampling purposes, one particular class in \mathcal{C} is chosen. For example, consider a system

$$\begin{cases} \mathcal{C}_1 = \Phi_1(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k) \\ \vdots \\ \mathcal{C}_m = \Phi_m(\mathcal{C}, \mathcal{Z}_1, \dots, \mathcal{Z}_k). \end{cases} \tag{26}$$

Suppose that we want to sample the objects from the class \mathcal{C}_1 and we fix the expected values of parameters $\mathcal{Z}_1, \dots, \mathcal{Z}_k$ to be N_1, \dots, N_k , respectively. Let $C_1(z_1, \dots, z_k), \dots, C_m(z_1, \dots, z_k)$ be the related generating functions. Then, the system of polynomial equations corresponding to the tuning problem (see e.g. [Duc+04, Proposition 2.1]) is given by

$$N_i = z_i \frac{\partial_{z_i} C_1(z_1, \dots, z_k)}{C_1(z_1, \dots, z_k)} \quad \text{for } i = 1, \dots, k. \tag{27}$$

Note that, in general, the values of the tuning parameters cannot be obtained by independently solving each of the equations (27). Each of the functions is depending on all the arguments at the same time (see also Figure 1). Hence, we propose an alternative procedure to achieve this goal.

From a technical point of view, the easiest case for theoretical analysis is combinatorial specifications corresponding to unambiguous context-free grammars. Here, the system defining the

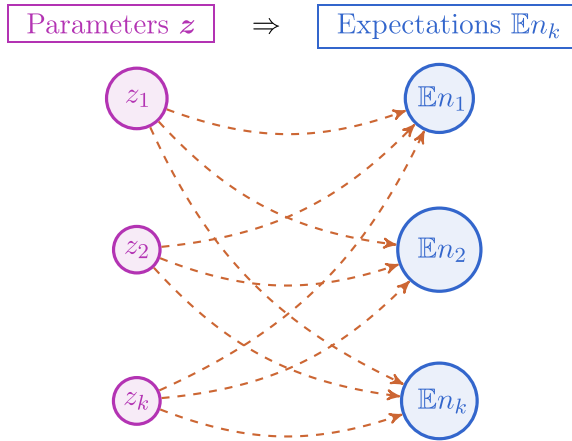


Figure 1. Dependency of parameters and expectations for multi-parametric tuning. Circle radii represent the possible values of the parameters.

generating functions becomes a *well-founded system* of polynomial equations, see [PSS12]. The most general framework of Boltzmann sampling comprises much more cases, including labelled objects, Pólya structures, or first- and second-order differential specifications. There also exist specifications whose equations include subtractions (related to the inclusion-exclusion principle), substitutions, and in these cases, different sampling strategies should be applied, for example recursive sampling or sampling with rejections, see Section 2.3. In this section, we are only concerned with tuning, setting thus all of these sampling issues aside.

3.1 Tuning as a convex optimisation problem

It turns out that instead of solving a system of polynomial equations (27) involving the derivatives of the generating function, a simpler convex optimisation problem, involving only the values of the generating functions, can be considered. Having the derivatives of this function comes as an advantage, because it allows to use first-order subroutines to solve the optimisation problem. The following theorem contains the most general form of our tuning approach. For convenience, we will write $f(\cdot) \rightarrow \min_z, f(\cdot) \rightarrow \max_z$ to denote the minimisation (maximisation, respectively) problem of the target function $f(\cdot)$ with respect to the vector variable z .

Theorem 12 (Tuning as convex optimisation). *Let $F(z_1, \dots, z_d)$ be a formal power series with non-negative coefficients analytic in an open d -dimensional set $\Omega \subseteq \{z_1 > 0, \dots, z_d > 0\}$. Assume that the solution of the multi-parametric tuning problem*

$$N_i = z_i \frac{\partial_{z_i} F(z_1, \dots, z_k)}{F(z_1, \dots, z_k)} \quad \text{for } i = 1, \dots, k, \tag{28}$$

belongs to Ω .

Let \mathbf{N} denote the vector (N_1, \dots, N_d) . Then, the tuning problem (28) is equivalent to the following convex optimisation problem over real variables φ and $\xi = (\xi_1, \dots, \xi_d)$:

$$\begin{cases} \varphi - \mathbf{N}^\top \xi \rightarrow \min_{\xi, \varphi}, \\ \varphi \geq \log F(e^\xi) \end{cases} \tag{29}$$

provided that the arguments of F meet its domain, and the logarithm is well-defined. The respective tuning parameters z^* satisfy then $z^* = e^\xi$.

Proof. Using a log-exp transformation, we note that the tuning problem (28) is equivalent to $\nabla_{\zeta} \log F(e^{\zeta}) = N$, cf. Proposition 7. Since the right-hand side vector N is equal to $\nabla_{\zeta}(N^{\top} \zeta)$, the tuning problem is further equivalent to

$$\nabla_{\zeta} \left(\log F(e^{\zeta}) - N^{\top} \zeta \right) = 0. \tag{30}$$

Note that the function under the gradient is a sum of a convex and linear function, and so necessarily convex itself. We can therefore equivalently express (30) as a convex minimisation problem in form of

$$\log F(e^{\zeta}) - N^{\top} \zeta \rightarrow \min_{\zeta}. \tag{31}$$

This problem can be reduced to a standard form (29) by adding an auxiliary variable φ . □

Remark 13. We do not require the exponents of the generating function to be non-negative integers. Depending on the specific application, they might, for instance, be positive real or rational numbers [Ham+19]. Even the non-negativity requirement could be omitted as well, as illustrated by a following univariate example. Let the exponents of the generating function $F(z)$ belong to the set A which may potentially include negative elements. Then, the log-exp transform of the function $F(z)$ is a composition of a convex function and a set of linear functions, which is again, convex:

$$\log F(e^z) = \log \left(\sum_{s \in A} a_s e^{zs} \right).$$

In contrast, it is crucial that the *coefficients* of the generating functions remain non-negative, as otherwise, the logarithm of the sum of exponents with potentially negative weights a_s ceases to be convex.

In subsequent sections we show that having a combinatorial specification is a for the generating function $F(z_1, \dots, z_d)$ is a great advantage – instead of requiring additional oracles providing the values of the generating function and its derivatives, a more direct approach is available.

3.2 Unambiguous context-free grammars

In the current section, we refine our general tuning procedure for *regular* and *unambiguous context-free* specifications, avoiding any use of external oracles. Recall that such systems are often used in connection with the Drmota–Lalley–Woods framework (see [FS09, Section VII.6]). However, contrary to [FS09, p. VII.6], we do not distinguish linear and non-linear cases. Instead, we develop a general result allowing to cover both cases. Also note that we do not require the system to be strongly connected, replacing this condition by a weaker requirement that every state is reachable from the initial one.

Theorem 14 (Tuning with finite parameter expectation). *Let $\mathcal{C} = \Phi(\mathcal{C}, \mathcal{Z})$ be a multi-parametric system with $\mathcal{C} = (C_1, \dots, C_m)$ and $\Phi = (\Phi_1, \dots, \Phi_m)$, where all the functions Φ_1, \dots, Φ_m are positive polynomials. Suppose that in the dependency graph corresponding to Φ all the states are reachable from the initial state C_1 . Let $N = (N_1, \dots, N_k)$ be the vector of target atom occurrences of each type. Fix the expectations N_i of the parameters of objects sampled from C_1 to $\mathbb{E}_z N = \mathbf{v}$. Then, the tuning vector z is equal to e^{ξ} where ξ comes from the convex problem:*

$$\begin{cases} c_1 - \mathbf{v}^{\top} \xi \rightarrow \min_{\xi, c} \\ c \geq \log \Phi(e^c, e^{\xi}). \end{cases} \tag{32}$$

Hereafter, e^{ξ} and $\log \Phi$ denote coordinate-wise exponentiation and logarithm, respectively.

Proof. Consider the vector \mathbf{z}^* such that $\mathbb{E}_{\mathbf{z}^*}(N) = \mathbf{v}$. Let \mathbf{c} denote the logarithms of the values of generating functions at point $\mathbf{z}^* = e^{\xi^*}$. Clearly, for such a choice of the vectors \mathbf{c} and $\xi = \xi^*$ all inequalities in (32) become equalities.

Let us show that if the point (\mathbf{c}, ξ) is optimal, then all the inequalities in (32) become equalities. Firstly, consider the case when the inequality

$$c_1 \geq \log \Phi_1(e^{\mathbf{c}}, e^{\xi}), \tag{33}$$

does not turn to an equality. Certainly, there is a *gap* and the value c_1 can be decreased without affecting the validity of other inequalities. In doing so, the target function value is decreased as well. Hence, the point (\mathbf{c}, ξ) cannot be optimal.

Now, suppose that the initial inequality does turn to equality, however $c_k > \log \Phi_k(e^{\mathbf{c}}, e^{\xi})$ for some $k \neq 1$. Since each of the states is reachable from the initial state \mathcal{C}_1 , it means that there exists a path $P = c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_k$ (indices are chosen without loss of generality) in the corresponding dependency graph. Note that for pairs of consecutive variables (c_i, c_{i+1}) in P , the function $\log \Phi_i(e^{\mathbf{c}}, e^{\xi})$ is strictly monotonic in c_{i+1} (as it is obtained as a log-exp transform of a positive polynomial and it references c_{i+1}). In such a case we can decrease c_{i+1} so to assure that $c_i > \log \Phi_i(e^{\mathbf{c}}, e^{\xi})$ while the point (\mathbf{c}, ξ) remains feasible. Decreasing c_{i+1}, c_i, \dots, c_1 in order, we finally arrive at a feasible point with a decreased target function value. In consequence, (\mathbf{c}, ξ) could not have been optimal to begin with.

So, eventually, the optimisation problem reduces to minimising the expression, subject to the system of equations $\mathbf{c} = \log \Phi(e^{\mathbf{c}}, e^{\xi})$ or, equivalently, $C(\mathbf{z}) = \Phi(C(\mathbf{z}), \mathbf{z})$ and can be therefore further reduced to Theorem 12. □

Remark 15. Let us note that the above theorem extends to the case of labelled structures with SET and CYC operators. For unlabelled Pólya operators like MSET or CYC, we have to truncate the specification to bound the number of substitutions. In consequence, it becomes possible to sample corresponding unlabelled structures, including partitions, functional graphs, series-parallel circuits, etc. For more details, see Section 6.1 .

Singular Boltzmann samplers (also defined in [Duc+04]) are the limit variant of ordinary Boltzmann samplers with an infinite expected size of generated structures. In their multivariate version, samplers are considered *singular* if their corresponding variable vectors belong to the boundary of the respective convergence sets. We present an alternative option to tune such singular samplers corresponding to Drmota–Lalley–Woods framework [FS09, Section VII.6], provided that their corresponding dependency graphs are strongly connected.

Theorem 16 (Tuning with infinite parameter expectation). *Let $\mathbf{C} = \Phi(\mathbf{C}, \mathcal{Z}, \mathbf{U})$ be a strongly connected multi-parametric system of positive polynomial equations with $\mathbf{C} = (C_1, \dots, C_m)$, the atomic class \mathcal{Z} marking the corresponding structure size, and $\mathbf{U} = (U_1, \dots, U_k)$ being a vector (possibly empty) of distinguished atoms. Assume that the target expected frequencies of the atoms U_i are given by the vector α . Then, the variables (z, \mathbf{u}) that deliver the tuning of the corresponding singular Boltzmann sampler are the result of the following convex optimisation problem, where $z = e^{\xi}$, $\mathbf{u} = e^{\eta}$:*

$$\begin{cases} \xi + \alpha^\top \eta \rightarrow \max_{\xi, \eta, \mathbf{c}} \\ \mathbf{c} \geq \log \Phi(e^{\mathbf{c}}, e^{\xi}, e^{\eta}). \end{cases} \tag{34}$$

Proof. By similar reasoning as in the previous proof, we can show that the maximum is attained when all the inequalities turn to equalities. Indeed, suppose that at least one inequality is strict, say $c_j > \log \Phi_j(e^{\mathbf{c}}, e^{\xi}, e^{\eta})$.

Because all right-hand sides of each of the inequality are monotonic with respect to c_j , we note that the value c_j can be slightly decreased by choosing a sufficiently small distortion $\varepsilon > 0$, turning

all the equalities containing c_j in the right-hand side $\log \Phi_i(e^c, e^\xi, e^\eta)$ into strict inequalities. Clearly, we can repeat this process until all equalities turn into inequalities.

Finally, let us focus on the target function. Again, because all right-hand sides of each inequality are monotonic with respect to ξ , we can slightly increase its value and increase the target function so to remain inside the feasible set.

Let us fix $\mathbf{u} = e^\eta$. For rational and algebraic grammars, within the Drmota–Lalley–Woods framework, see for instance [Drm97], the corresponding generating function singular approximation takes the form

$$C(z, \mathbf{u}) \sim a_0(\mathbf{u}) - b_0(\mathbf{u}) \left(1 - \frac{z}{\rho(\mathbf{u})}\right)^t. \tag{35}$$

If $t < 0$, then the asymptotically dominant term becomes $-b_0 \left(1 - \frac{z}{\rho(\mathbf{u})}\right)^t$. In this case, tuning the target expected frequencies corresponds to solving the following equation as $z \rightarrow \rho(\mathbf{u})$:

$$\text{diag}(\mathbf{u}) \frac{[z^n] \nabla_{\mathbf{u}} C(z, \mathbf{u})}{[z^n] C(z, \mathbf{u})} = n\alpha. \tag{36}$$

Let us substitute the asymptotic expansion (35) into (36) to track how \mathbf{u} depends on α :

$$\text{diag}(\mathbf{u}) \frac{[z^n] t b_0(\mathbf{u}) \left(1 - \frac{z}{\rho(\mathbf{u})}\right)^{t-1} z \frac{\nabla_{\mathbf{u}} \rho(\mathbf{u})}{\rho^2(\mathbf{u})}}{[z^n] b_0(\mathbf{u}) \left(1 - \frac{z}{\rho(\mathbf{u})}\right)^t} = -n\alpha. \tag{37}$$

Only dominant terms are accounted for. Then, by the binomial theorem

$$\text{diag}(\mathbf{u}) b_0(\mathbf{u}) \frac{t}{n} \binom{t-1}{n} \frac{z \nabla_{\mathbf{u}} \rho(\mathbf{u})}{\rho^2(\mathbf{u})} b_0(\mathbf{u})^{-1} \binom{t}{n}^{-1} = -\alpha. \tag{38}$$

With $z = \rho(\mathbf{u})$, as $n \rightarrow \infty$, we obtain after cancellations

$$\text{diag}(\mathbf{u}) \frac{\nabla_{\mathbf{u}} \rho(\mathbf{u})}{\rho(\mathbf{u})} = -\alpha, \tag{39}$$

which can be rewritten as

$$\nabla_{\eta} \log \rho(e^\eta) = -\alpha. \tag{40}$$

Passing to exponential variables (40) becomes

$$\nabla_{\eta} (\xi(\eta) + \alpha^\top \eta) = 0. \tag{41}$$

As we already discovered, the dependence $\xi(\eta)$ is given by the system of equations because the maximum is achieved only when all inequalities turn to equations. That is, tuning the singular sampler is equivalent to maximising $\xi + \alpha^\top \eta$ over the set of feasible points. \square

Remark 17. For ordinary and singular samplers, the corresponding feasible set remains the same; what differs is the optimised target function. Singular samplers correspond to imposing an infinite target size. In practice, however, the required singularity is almost never known *exactly* but rather calculated up to some feasible finite precision. The tuned structure size is therefore enormously large, but still, nevertheless, finite. In this context, singular samplers provide a natural *limiting* understanding of the tuning phenomenon and as such, there are several possible ways of proving Theorem 16.

Figure 2 illustrates the feasible set for the class of binary trees and its transition after applying the log-exp transform, turning the set into a convex collection of feasible points. In both figures,




Figure 2. Binary trees $B \geq z + zB^2$ and log-exp transform of the feasible set. The black curve denotes the principal branch of the generating function $B(z)$ corresponding to the class of binary trees.

the singular point is the rightmost point on the plot. Ordinary sampler tuning corresponds to finding the tangent line which touches the set, given the angle between the line and the abscissa axis.

Remark 18. As an interesting by-product, our tuning algorithm provides a way of obtaining the singularities of a system and the values of the generating functions at the point of this singularity.

Remark 19. It is also possible to consider singular tuning the case of non-strongly connected specifications. However, practically speaking, it should be noted that a notion of a *singular sampler* for non-strongly connected specifications such as

$$F(z) = \frac{1}{1 - 2T(z)} \quad \text{and} \quad T(z) = ze^T(z), \tag{42}$$

is ambiguous – both singular samplers for T and F admit different values of the variable z . For T its $z = e^{-1}$ whereas for F we have $z = e^{-1/2}/2$.

If the substitutions in the dependency graph of the specification include only *subcritical compositions* (see [FS09, Section VI.9]) it is possible to incrementally tune its parts in topological order. Further theoretical analysis of singular samplers involving supercritical compositions is somewhat more complicated, but can be developed as well. Nevertheless, the easiest practical way to derive a singular tuner is to tune the target class with a large, yet finite object size.

3.3 Labelled and unlabelled structures

Systems originating from labelled specifications, that is whose generating functions are typically of exponential type, feature such admissible operators as SET or CYC, both in their unrestricted and cardinality restricted variants. Consider a labelled multi-parametric combinatorial class \mathcal{A} with a generating function $A(z)$. Then, the resulting exponential generating functions obtained using these operators take form

$$\begin{aligned} \text{SET}_k(\mathcal{A})(z) &= \frac{A(z)^k}{k!} & \text{CYC}_k(\mathcal{A})(z) &= \frac{A(z)^k}{k} \\ \text{SET}(\mathcal{A})(z) &= e^{A(z)} & \text{CYC}(\mathcal{A})(z) &= \log \frac{1}{1 - A(z)}. \end{aligned} \tag{43}$$

Classes whose definition involves one of the above operators can be incorporated into the convex optimisation problem using a log-exp transformation $F(z) \mapsto e^\varphi$ similarly to the case context-free grammars, see Section 3.2 . Broadly speaking, the application of such operators results in a composition with one of the basic functions

$$\frac{x^k}{k!}, \quad \frac{x^k}{k}, \quad e^x \quad \text{and} \quad \log \frac{1}{1 - x}, \tag{44}$$

expressing, respectively, the exponential generating functions for the class of labelled sets with k elements, labelled cycles with k elements, and both unrestricted labelled sets and cycles.

On the other hand, ordinary generating functions, used for enumeration of unlabelled structures, feature such operators as MSET, PSET, and CYC, standing for the multiset, set, and a cycle constructions, respectively. These, in contrast, are evaluated differently than their labelled counterparts. Specifically, if applied to a class with an ordinary generating function $A(z) = \sum_{n>0} a_n z^n$, we obtain, respectively,

$$\exp\left(\sum_{k \geq 1} \frac{A(z^k)}{k}\right), \quad \exp\left(\sum_{k \geq 1} \frac{(-1)^{k-1} A(z^k)}{k}\right) \quad \text{and} \quad \sum_{k \geq 1} \frac{\varphi(k)}{k} \log \frac{1}{1 - A(z^k)}. \quad (45)$$

Here $\varphi(k)$ denotes Euler’s totient function.

Let us note that both the first and the third expressions, after log-exp transformations, become convex. The resulting infinite series can be then truncated at a finite threshold, given the fact that if the corresponding singular value $\rho < 1$, the sequence $A(z^k)$ converges at geometrical speed to $A(0)$. Such a truncation is common practice and has been applied a number of time in the context of sampling unlabelled structures, see for example [BLR15, FFP07]. A more detailed discussion regarding these transformations is given in Section 6.1 .

Unfortunately, the remaining PSET operator, in its aforementioned series form with negative coefficients (45), does not easily fit the convex programming framework. In principle, an alternative form could be used

$$\text{PSET}(\mathcal{A})(z) = \prod_{n>0} ((1 + z_1^{n_1}) \cdots (1 + z_d^{n_d}))^{a_n}, \quad (46)$$

though in this case, it is not clear how to convert the problem into a polynomially tractable form.

3.4 Increasing trees

Boltzmann samplers for first- and second-order differential specifications have been developed in [BRS12, Bod+16]. In particular, in [BRS12], the authors solve the problem of Boltzmann sampling from specifications of type $\mathcal{T}' = \mathcal{F}(\mathcal{Z}, \mathcal{T})$. There exist several particular cases of such a differential specification which admit an explicit solution of the corresponding differential equation. For example, a differential equation is said to be *autonomous* if it does not depend explicitly on the independent variable, that is $\mathcal{T}' = \mathcal{F}(\mathcal{T})$. In this case, the underlying differential equation can be solved by separation of variables:

$$\begin{cases} T'(z) = F(T(z)) \\ T(0) = t_0 \end{cases} \quad \Rightarrow \quad \int_{t_0}^T \frac{dt}{F(t)} = z(T). \quad (47)$$

The final solution $T(z)$ is obtained by inverting $z(T)$, which, in the case of a differential equation with combinatorial origin, can be obtained using binary search.

Several different strategies for evaluating $\int \frac{1}{F(t)}$ may be also suggested. For instance, if $F(t)$ is a polynomial in t , all its complex roots (together with multiplicity structure) can be efficiently found through a numerical procedure³, see [Zen05]. In the context of multi-parametric tuning, we

³Efficient polynomial root-finding is directly related to Smale’s 17th problem about finding complex solutions of complex polynomial systems, recently positively solved by Lairez, see [Lai17].

are mostly interested in the case when $F(T)$ also depends on auxiliary variables $\mathbf{u} = (u_1, \dots, u_d)$, which need to be tuned. In this case, the roots of the polynomial $F(T)$ become dependent on \mathbf{u} .

Theorem 20 (Multi-parametric increasing trees). *Suppose that the generating function $T(z, \mathbf{u})$ corresponding to a family of increasing trees is described by a functional equation*

$$T'_z(z, \mathbf{u}) = F(T(z, \mathbf{u}), \mathbf{u}) \quad \text{with} \quad T(0) = t_0, \tag{48}$$

and that T is a formal power series with non-negative coefficients, $\mathbf{u} = (u_1, \dots, u_d)$.

Let N_0, N_1, \dots, N_d be the expected size, and expected parameter value of u_1, \dots, u_d , respectively. Then, the tuning problem is equivalent to a convex optimisation problem over real variables $\varphi, \zeta, \eta_0, \dots, \eta_d$:

$$\begin{cases} \varphi - N_0\zeta - N_1\eta_1 - \dots - N_d\eta_d \rightarrow \min, \\ \zeta \leq \log \int_{t_0}^{e^\varphi} \frac{dt}{F(t, e^{\eta_1}, \dots, e^{\eta_d})} \end{cases}, \tag{49}$$

where the solution of the tuning problem is obtained by assigning

$$z = e^\zeta, \quad u_1 = e^{\eta_1}, \quad \dots, \quad u_d = e^{\eta_d}, \quad \text{and} \quad T(z, u_1, \dots, u_d) = e^\varphi. \tag{50}$$

Proof. Following Theorem 12, multi-parametric tuning is equivalent to the following convex optimisation problem over real variables $\varphi, \zeta, \eta_1, \dots, \eta_d$:

$$\begin{cases} \varphi - N_0\zeta - N_1\eta_1 - \dots - N_d\eta_d \rightarrow \min, \\ \varphi \geq \log T(e^\zeta, e^{\eta_1}, \dots, e^{\eta_d}) \end{cases} \tag{51}$$

where $T(z, u_1, \dots, u_d)$ is the solution of differential equation $T'_z(z, \mathbf{u}) = F(T(z, \mathbf{u}), \mathbf{u})$ with initial conditions $T(0, \mathbf{u}) = t_0$. The target solution is given by $(z, u_1, \dots, u_d) = (e^\zeta, e^{\eta_1}, \dots, e^{\eta_d})$ and $F(z, u_1, \dots, u_d) = e^\varphi$. Let us denote by $Z(\tau, \eta_1, \dots, \eta_d)$ the inverse function of $\log T(e^\zeta, e^\eta)$ with respect to ζ , so that

$$\log T(e^{Z(\tau, \eta_1, \dots, \eta_d)}, e^{\eta_1}, \dots, e^{\eta_d}) = \tau. \tag{52}$$

Since T is a formal power series with non-negative coefficients, the function $\log T(e^\zeta, e^{\eta_1}, \dots, e^{\eta_d})$ is convex and increasing in both ζ and η . Therefore, its inverse function with respect to ζ is a concave increasing function. Moreover, the function $Z(\tau, \eta_1, \dots, \eta_d)$ is jointly concave in all of its arguments. That is because a function $f(z)$ is concave if and only if its hypograph $\text{hyp}f = \{(y, z) : y \leq f(z)\}$ is a convex set.

The hypograph of the inverse function $Z(\tau, \eta)$ directly corresponds to the epigraph of the initial function $\log T(e^z, e^\eta)$. Consequently, a convex constraint $\varphi \geq \log T(e^\zeta, e^{\eta_1}, \dots, e^{\eta_d})$ can be replaced by a different convex constraint obtained by taking the inverse with respect to ζ , resulting in a new convex optimisation problem

$$\begin{cases} \varphi - N_0\zeta - N_1\eta_1 - \dots - N_d\eta_d \rightarrow \min, \\ \zeta \leq Z(\varphi, \eta_1, \dots, \eta_d). \end{cases} \tag{53}$$

In order to construct $Z(\varphi, \eta_1, \dots, \eta_d)$, we use the solution of the autonomous differential equation, see (47). Denote by $z(t, \mathbf{u})$ the inverse of the generating function $T(z, \mathbf{u})$, that is

$$T(z(t, \mathbf{u}), \mathbf{u}) = t, \tag{54}$$

and recall that $z(t, \mathbf{u})$ is given by

$$z(t, \mathbf{u}) = \int_{t_0}^t \frac{d\tau}{F(\tau, \mathbf{u})}. \tag{55}$$

Finally, a direct calculation shows that $Z(\varphi, \eta_1, \dots, \eta_d) = \log z(e^\varphi, e^{\eta_1}, \dots, e^{\eta_d})$. □

Remark 21. With an external oracle available, systems of differential equations involving more than one variable can still be tuned using Theorem 12. However, there is little hope to generalise Theorem 20 onto multivariate differential equations, even if we assume that all its differential equations are autonomous. In fact, by introducing an additional dimension to the problem, any first-order system of differential equations

$$\frac{d}{dz} \mathbf{y}(z) = \mathbf{F}(z, \mathbf{y}(z)), \tag{56}$$

can be transformed into a system of autonomous differential equations. Moreover, starting from dimension four, systems of functional equations can admit solutions which exhibit chaotic behaviour, see [FS09, Remark VII.51].

On the other hand, any system of differential equations of arbitrarily high order

$$\mathbf{y}^{(n)}(z) = \mathbf{F}(z, \mathbf{y}(z), \mathbf{y}^{(1)}(z), \dots, \mathbf{y}^{(n-1)}(z)), \tag{57}$$

can be reduced to a system of first-order differential equations by expanding the dimension space, and therefore, Boltzmann samplers for such specifications may apply. As discussed in [BRS12], for such systems, only the first sampling step is the most expensive one, because it requires to generate a random variable defined on the interval $[0, \rho)$, where ρ is the radius of convergence of the formal power series. For all the consecutive operations, the support of the required random variables is separated from ρ , so a variety of approximation methods, including Runge–Kutta can be applied.

3.5 Other types of combinatorial structures

The general technique described in Section 3.1 can be applied to any analytic multivariate generating function provided that the solution to the multiparametric tuning problem lies inside the respective domain of analyticity. It can be applied even to those coming from somewhat exotic systems, including partial differential equations, systems with negative coefficients, catalytic equations, systems with non-trivial substitutions, etc. As long as the oracle providing values and the derivatives of the generating functions is given, the source of the equations is irrelevant.

Still, for most of these systems no efficient oracle is known. Typically, the following methods, which could be called the *oracles* for simplicity, can be used for the mentioned types of functional equations

- Runge–Kutta method for ordinary differential equations,
- Grid approximation methods for partial differential equations,
- Gröbner bases for systems of polynomial equations, or
- Plain coefficient-wise evaluation for arbitrary systems.

All of these methods, however, do not guarantee polynomial complexity in terms of the bit length of the target size and the number of equations. When additional parameters are introduced, then the approach that we propose in Section 3.1 reduces the number of equations for an oracle to solve, but adds the optimisation procedure on the top. Specifically, with this approach, the tuning system of equations (28) is not anymore included as an input to the oracle, and the tuning

parameters are chosen by the optimisation procedure, and so are considered fixed real numbers from the viewpoint of the oracle.

Remark 22. It is worth noting that for certain combinatorial classes beyond the scope of the current paper, the tuning problem does not have a solution in the analyticity domain. The simplest non-trivial example (with an infinite number of objects in a class) is the case of unrooted trees $U(z) = \sum_{n \geq 0} n^{n-2} \frac{z^n}{n!}$: near its dominant singular point $z = e^{-1}$, the Puiseux expansion takes form

$$U(z) = \frac{1}{2} - (1 - ez) + \frac{2^{3/2}}{3}(1 - ez)^{3/2} + O((1 - ez)^2).$$

This means that the maximal possible expected value of an unrooted tree generated from Boltzmann distribution is equal to

$$\lim_{z \rightarrow e^{-1}} \frac{z \partial_z U(z)}{U(z)} = 2.$$

However, for most examples of this kind, the Boltzmann sampler itself is not directly available, not speaking of the fact that there would be no practical motivation to apply it for expectations higher than the maximal available value.

Still, there exist techniques different from the Boltzmann sampling which can be used to draw the objects from such classes in a controlled manner. The reason why the solution does not exist in the above example is because locally around the singularity ρ , the singular term of the generating function takes form $(1 - \frac{z}{\rho})^\alpha$, where $\alpha > 1$. By repeatedly applying pointed derivative to such a class, the singular term can be modified, and the exponent α shall be consequently replaced by a value strictly smaller than 1. From the combinatorial viewpoint, pointed derivative serves to distinguish atoms in the underlying object, which can be un-distinguished later after the sampling is performed. The possible drawbacks of this approach are that firstly, successive pointing can increase the size of the grammar, and secondly, the resulting distribution is not anymore Boltzmann, which is not surprising, because the tuning problem does not have a solution in the Boltzmann case.

It is still important to discover and understand controlled sampling techniques in the cases not covered by admissible tuning in the sense of (28): when inclusion-exclusion is used, or even in more sophisticated cases involving stretched exponents or any kinds of essential singularities.

4. Complexity of convex optimisation

In the previous section, we have reduced the problem of multiparametric tuning to a problem of convex optimisation. In this section, we provide the complexity results, case by case. Let us introduce the key elements of the interior-point optimisation and discuss how to evaluate its complexity.

A convex optimisation problem can be expressed as minimisation of a linear multivariate function (or, in an equivalent formulation, a convex one) subject to a set of constraints, expressed as inequalities of the form $f_i(z) \leq 0$ involving multivariate convex functions $f_i(z)$. In addition, affine inequalities and linear equations are often treated separately. Each of these constraints describes a convex set, and therefore, an intersection of these sets is also convex. A classical convex optimisation program can be formulated as

$$\begin{cases} c^\top z \rightarrow \min_z \\ f_i(z) \leq 0 \text{ for } i = 1, \dots, m \end{cases}, \tag{58}$$

where f_1, \dots, f_m are convex functions of the vector argument \mathbf{z} . Such programs are widely believed to be *solvable* in polynomial time, due to the popularisation of interior-point methods. Nevertheless, a detailed answer regarding the algorithmic complexity of convex optimisation needs a careful investigation.

For instance, if $f_i(\mathbf{z})$ are linear functions with integer coefficients, then the solution to the optimisation problem is a rational number. In this case the problem is known to be solvable *exactly* using $O(n^{3.5}L)$ arithmetic operations, where n denotes the number of variables, and L is the bit length of the coefficient representation [Kar84]. Remarkably, the existence of a polynomial algorithm whose arithmetic complexity does not depend on L (so-called *strongly-polynomial algorithms*) is listed among 18 problems in Smale's celebrated list [Sma98]. If the functions $f_i(\mathbf{z})$ are not necessarily linear, the optimal solution does not have to be a rational number and so we usually ask for an ε -approximation, instead.

Nesterov and Nemirovskii [NN94] developed a theoretically and practically efficient framework covering a large subset of convex programs. Their method relies on the notion of a *self-concordant barrier* whose construction depends on the problem formulation and needs to exploit the global structure of the optimisation problem. The number of Newton iterations required by their method to solve the optimisation problem (58) can be bounded by

$$O\left(\sqrt{\nu} \log\left(\frac{\nu\mu_0}{\varepsilon}\right)\right) \quad (59)$$

where ε is the required precision, ν is the complexity of the constructed barrier (typically proportional to the problem dimension), and μ_0 is related to the choice of the starting point. More specifically, the value of the target function is guaranteed to lie within an ε -neighbourhood of the optimal value. Since μ_0 does not depend on the target precision, it is often omitted in the final analysis [PW00]. Further on, we compute ν for various types of combinatorial specifications.

The cost of each Newton iteration is a sum of the two parts: the cost of the composition of the matrix of second derivatives and the cost of the inversion of this matrix. In Section 4.2, we specify in more detail what is a Newton iteration in the interior-point method. While the cost of the inversion part is assumed to be cubic in the number of variables, the cost of the composition of the matrix of second derivatives depends on the problem structure.

One of the most significant achievements of interior-point programming was a proof of existence of a *universal family of barriers*, which admits a $O(n)$ self-concordance complexity for any convex body in \mathbb{R}^n . Unfortunately, such an existence result seems to be currently only of theoretical interest – a construction of such a barrier requires computing the convex body's volume, which in itself is known to be an NP-hard problem. In practice, for each concrete convex optimisation problem the barriers have to be constructed separately. There is no known constructive, general-purpose polynomial-time barrier construction algorithm for convex optimisation problems.

To summarise, the practical complexity of the convex optimisation is determined by the complexity parameter ν of the associated self-concordant barrier, which requires the problem structure to be exploited in order to be constructed. More details on the algorithmic complexity of convex optimisation methods can be found in [PW00].

4.1 Disciplined Convex Programming

Disciplined Convex Programming (DCP, sometimes also referred to as CVX) is another framework meant for modelling convex optimisation problems [GBY06]. Both the objective and constraints are built using certain *atomic expressions* such as $\log(\cdot)$, $\|\cdot\|_2$, or $\log \sum e^{\cdot}$, and a restricted set of *composition rules* following basic principles of convex analysis. The inductive definition of expressions allows the framework to track their curvature and monotonicity, and, in consequence, automatically determine whether the given program is indeed a valid convex program. Essentially, DCP can be viewed as both a domain-specific language and a compiler which transforms the given

convex program to a mixture of linear, quadratic, and exponential conic optimisation problems. Such problems are passed to corresponding solvers in the necessary standard form. DCP covers therefore a strict subset of convex programs of Nesterov and Nemirovskii’s framework, but serves well as a prototyping interface for provably convex programs. In our implementation, see Section 6, we rely on DCP with two particular conic solvers, a second-order Embedded Conic Solver [DCB13] and recently developed first-order Splitting Conic Solver [ODo+16].

Some of the classical combinatorial constructions, for example MSET, CYC or $SET_{>0}$ can be expressed in DCP, however involve using (at least in principle) an infinite amount of summands. Some of the operators, for example, MSET can be efficiently represented in a truncated form, cf. [BLR15], but for others it is intrinsically impossible to provide a reasonable truncation level, since a large number of summands take non-negligible values. Therefore, for such constructions DCP is not the preferable, and the classical interior-point methods should be selected instead. In what follows, we provide some additional theoretical background on self-concordant barriers and propose our constructions which can be further used to cover these operations.

4.2 Self-concordant barriers

While the precise descriptions of the interior-point optimisation schemes and the definitions of the self-concordant barriers can be found in [Nem04], we only need to borrow a couple of statements. We intentionally do not give the definition of a self-concordant barrier, but instead use a sufficient construction condition. More specifically, we rely on the fact that self-concordant barriers are functions assigned to various convex sets in \mathbb{R}^n and we will only need to construct such sets for *epigraphs* of convex functions, defined as

$$\text{epi } f = \{(t, x) \in \mathbb{R}^2 \mid f(x) \leq t\}. \tag{60}$$

Lemma 23 ([Nem04, Proposition 9.2.2]). *Let $f(x)$ be a three times continuously differentiable real-valued convex function on the ray $\{x > 0\}$ such that*

$$|f'''(x)| \leq 3\beta x^{-1}f''(x), \quad x > 0. \tag{61}$$

for some parameter $\beta > 0$.
Then, the function

$$F(t, x) = -\log(t - f(x)) - \max[1, \beta^2] \log x, \tag{62}$$

is a self-concordant barrier with complexity parameter $\nu = 1 + \max[1, \beta^2]$ for the two-dimensional convex domain $\{(t, x) \in \mathbb{R}^2 \mid x > 0, f(x) \leq t\}$.

The two important operations for composing convex problems are addition and composition (the composition of a convex function with an *increasing* convex function is also convex). Each composition is treated by introducing an additional variable; for example, if $F(x)$ is convex and increasing, and $G(x)$ is convex, then the epigraph of the composition $\{(t, x) \mid F(G(x)) \leq t\}$ can be expressed as a projection of a three-dimensional set

$$\{(x, t) \mid \exists y \mid F(y) \leq t, G(x) \leq y\}. \tag{63}$$

The behaviour of self-concordant barriers with respect to taking linear combinations or with respect to combining multiple dimensions is described by the following proposition.

Lemma 24 ([Nem04, Proposition 3.1.1]). *The following propositions hold:*

- (i) *Let F_i be self-concordant barriers with complexity parameters ν_i for closed convex domains $Q_i \subset \mathbb{R}^{\nu_i}$, $i = 1, \dots, m$. Then the function*

$$F(x_1, \dots, x_m) = F_1(x_1) + \dots + F_m(x_m), \tag{64}$$

is a $(\sum_i \nu_i)$ -self-concordant barrier for $Q = Q_1 \times \dots \times Q_m$.

(ii) Let F_i be self-concordant barriers with complexity parameters v_i for closed convex domains $Q_i \subset \mathbb{R}^n$ and $\alpha_i \geq 1$ be reals, $i = 1, \dots, m$. Assume that $Q = \cap_{i=1}^m Q_i$ has a non-empty interior. Then the function

$$F(x) = \alpha_1 F_1(x) + \dots + \alpha_m F_m(x), \tag{65}$$

is a $(\sum_i \alpha_i v_i)$ -self-concordant barrier for Q .

Self-concordant barrier design is one of the central issues in modern convex optimisation – knowing these barriers allows one to reduce any convex program to a standard form [Nes98, Section 4.2.6]. The reduction proceeds as follows.

$$\begin{cases} f_0(\mathbf{z}) \rightarrow \min \\ f_i(\mathbf{z}) \leq 0 \text{ for } i = 1, \dots, m \end{cases} \Leftrightarrow \begin{cases} \tau \rightarrow \min \\ f_0(\mathbf{z}) \leq \tau \\ f_i(\mathbf{z}) \leq \kappa \text{ for } i = 1, \dots, m \\ \tau \leq C, \kappa \leq 0. \end{cases} \tag{66}$$

Here, C is a technical constant, bounding the value of the target function from above. Suppose that $F_0(\mathbf{z}, \tau), F_1(\mathbf{z}, \kappa), \dots, F_m(\mathbf{z}, \kappa)$ are the v_0 -self-concordant and v_i -self-concordant barrier functions for the epigraphs of the functions $f_0(\mathbf{z}), f_1(\mathbf{z}), \dots, f_m(\mathbf{z})$, respectively. Then, according to Lemmas 23 and 24, the resulting self-concordant barrier

$$\widehat{F}(\mathbf{z}, \tau, \kappa) = F_0(\mathbf{z}, \tau) + \sum_{i=1}^m F_i(\mathbf{z}, \kappa) - \log(C - \tau) - \log(-\kappa), \tag{67}$$

has a complexity parameter

$$\widehat{v} = v_0 + \sum_{i=1}^m v_i + 2. \tag{68}$$

As specified in [NN94, Section 2], each Newton step is a step of the form

$$x_{i+1} = x_i - \mu_i \cdot (\widehat{F}''(x_i))^{-1} \widehat{F}'(x_i),$$

where μ_i is a dampening parameter defined by the procedure, and $x_i \in \mathbb{R}^m$, where m is the dimension of the parameter space. The computation of the dampening parameter, as well as the Newton step, involves computing the inverse of the matrix of second derivatives. It costs $O(m^2 L)$ arithmetic operations to compose this matrix, which may potentially vary if some exotic constructions are introduced, and $O(m^3)$ steps to compute its numerical inverse.

Finally, let us note that there is room for a variety of additional fine-tuned optimisation techniques. For instance, the Newton iterations can use faster than $O(n^3)$ matrix inversion algorithms, use $O(n^\omega)$ matrix multiplication algorithms, or even use algorithms dedicated to sparse matrices.

Remark 25. The DCP software, that we use in our prototype tuner implementation, uses a completely different principle and expresses the problem in a standard form using so-called *exponential cones*. The approach with barriers gives here more theoretical guarantees as it allows to cover a broader range of constructions, but was not chosen for prototyping reasons.

4.3 Barriers for context-free specifications

In this section, we give rigorous complexity bounds for the case of context-free grammars, see Section 3.2. Each of the constraints in the auxiliary convex optimisation problem that we construct takes form

$$x_i \geq \log \left(\sum_j \exp(\mathbf{a}_{ij}^\top \mathbf{x}) \right) \quad \text{for } i = 1, \dots, m, \tag{69}$$

where \mathbf{x} is the vector of unknowns, and \mathbf{a}_{ij} are vectors of coefficients. Such a convex constraint can be rewritten into an equivalent form

$$e^{x_i} \geq \sum_j \exp(\mathbf{a}_{ij}^\top \mathbf{x}) \iff \sum_j \exp((\mathbf{a}_{ij} - \mathbf{e}_i)^\top \mathbf{x}) \leq 0, \tag{70}$$

where \mathbf{e}_i is i th basis vector. Each of the functions $\exp((\mathbf{a}_{ij} - \mathbf{e}_i)^\top \mathbf{x})$ is convex, as $(\mathbf{a}_{ij} - \mathbf{e}_i)^\top \mathbf{x}$ is linear and e^x is both increasing and convex. Clearly, their sum is also convex. Using the sum-substitution technique from Lemma 24, we conclude that in order to provide a fast convex optimisation interior-point procedure, it is sufficient to construct a self-concordant barrier for the epigraph of e^x . Note that $\mathbf{epi} e^x = \{(t, x) \in \mathbb{R}^2 \mid e^x \leq t\}$ can be equivalently rewritten as $\{(t, x) \in \mathbb{R}^2 \mid t > 0, -\log t \leq -x\}$. Furthermore, a linear variable change $x \mapsto -x$, and variable swapping $(x, t) \mapsto (t, x)$ converts it into an epigraph of the convex function $f(x) = -\log x$:

$$\mathbf{epi}[-\log(x)] = \{(t, x) \in \mathbb{R}^2 \mid x > 0, -\log x \leq t\}. \tag{71}$$

Finally, we note that the condition from Lemma 23 is fulfilled with $\beta = \frac{2}{3}$. Therefore, a logarithmic barrier

$$F(t, x) = -\log(t + \log(x)) - \log x, \tag{72}$$

is a 2-self-concordant barrier for an epigraph of $-\log(x)$. This also gives a 2-self-concordant barrier for an epigraph of e^x . Consequently, each unambiguous context-free specification can be converted into a standard form for an interior-point method with total barrier complexity parameter $O(L)$ where L is the total number of terms in the specification.

4.4 Cycle and positive set constructions

Among the basic labelled operators, the CYC operator cannot be easily fitted into the more restricted DCP framework that we use in our prototype implementation. The same issue concerns all of the restricted versions of labelled operators including $\text{SET}_{>0}$, $\text{SET}_{>k}$, and $\text{CYC}_{>k}$. We provide further, more heuristic analysis of the logarithmic barriers for these operators and hint as to why they can be efficiently supported. We do not provide practical implementations of these modified versions, however, in principle, they can be handled as well.

Let us start with two particular functions obtained as a log-exp transform of the CYC and $\text{SET}_{>0}$ constructors, see Section 6.1. These two functions are

$$L(x) = \log \log \frac{1}{1 - e^x} \quad \text{and} \quad E(x) = \log(e^{e^x} - 1). \tag{73}$$

Note that the domain of $L(x)$ is equal to $\{x \in \mathbb{R} \mid x < 0\}$. The convexity of each of $E(x)$ and $L(x)$ is ensured by noticing that they are log-sum-exp expressions with an infinite number of positive summands:

$$L(x) = \log \sum_{n=1}^{\infty} \frac{e^{kx}}{k} \quad \text{and} \quad E(x) = \log \sum_{n=1}^{\infty} \frac{e^{kx}}{k!}. \tag{74}$$

Unfortunately, we cannot afford an infinite number of summands, as adding each summand contributes to the final barrier complexity. Instead, we are going to explicitly construct self-concordant barriers for the epigraphs of $L(x)$ and $E(x)$ with the help of Lemma 23.

Table 4. $|f'''(\hat{x})| = \frac{3\beta}{x} f''(\hat{x})$ for different β and corresponding \hat{x} .

β	6	9	12	15	30
$ \hat{x} - 3\beta $	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-11}

The case of $f(x) = L(x)$. Direct numerical evaluation shows that when $0 < x \lesssim 3\beta$ it holds $|f'''(x)| \leq \frac{3\beta}{x} f''(x)$. Relative errors of \hat{x} satisfying $|f'''(\hat{x})| = \frac{3\beta}{x} f''(\hat{x})$ for varying values of β is listed in Table 4.

Let us notice the practical consequences of this fact. For most typical combinatorial systems, tuning the expected size to n gives the following behaviour of the tuning parameter:

- $x_n = \rho (1 - O(n^{-1}))$ for algebraic and rational systems, and
- $x_n = O(\text{Poly}(n))$ for some specifications with singularity ρ at infinity, such as entire functions (e.g. labelled sets or similar constructs).

In fact, we are not aware of any natural example combinatorial system where the tuning parameter approaches the singular value at an exponential speed $x_n = \rho(1 - e^{-n})$.

Therefore if we choose, say, $\beta = 10$, then according to Lemma 23, the resulting logarithmic barrier will be β^2 -self-concordant for $x \lesssim 30$. In practical terms, $\beta = 10$ covers expected values up to $N \sim e^{30} \approx 10^{13}$. More generally, for $N \rightarrow \infty$, the complexity of the whole optimisation scheme is proportional to a square root of the sum of all the barrier complexities, which yields an additional $\log N$ factor for the resulting barrier complexity, and for the complexity of the final tuning algorithm.

The case of $f(x) = E(x)$. If we try to apply Lemma 23 directly to $E(x)$, we discover that the conditions of the lemma cannot be satisfied. Therefore, we pre-process the epigraph of $E(x)$ by taking its inverse, similarly as we did for the 2-self-concordant barrier corresponding the exponent function e^x . Specifically, we transform **epi** E into

$$\mathbf{epi} E = \{(x, t) \in \mathbb{R}^2 \mid E(x) \leq t\} = \{(x, t) \in \mathbb{R}^2 \mid t > 0, x \leq \log \log (1 + e^t)\}. \tag{75}$$

It can be proven that $f(x) = -\log \log (1 + e^x)$ is convex for $x > 0$. Moreover, we can numerically verify that the condition of Lemma 23 is satisfied for all $x \leq 10^{100}$. We conjecture that it is valid for all $x \in \mathbb{R}$, but, as we discussed above, in practice we only need it to be valid for relatively small values of x .

4.4.1 Restricted versions of cycles and sets.

It is also possible to consider cycles and sets containing more than k elements, so that we need to construct self-concordant barriers for the functions

$$L_{>k}(x) = \log \left(\log \frac{1}{1 - e^x} - \sum_{n=1}^k \frac{e^{kx}}{k} \right) \quad \text{and} \quad E_{>k}(x) = \log \left(e^{e^x} - \sum_{n=0}^k \frac{e^{kx}}{k!} \right). \tag{76}$$

For the cycle construction, we again empirically discover that for $x \lesssim 3\beta$ the logarithmic barrier is self-concordant. We have numerically tested this assertion for $k \in \{0, 1, 2, 3\}$ and $\beta \leq 10$.

For the restricted construction $\text{SET}_{>k}$, we proceed in the same way as we did for $\text{SET}_{>0}$ by constructing the epigraph of the negative inverse function. In this case, however, the inverse function cannot be expressed explicitly, but can only be computed numerically. All the numerical computations can be done efficiently using the binary search for the inverse function, and implicit derivative theorem for its derivatives. Again, this gives an algorithmic definition of a self-concordant barrier for restricted sets without the necessity to operate with infinite sums.

4.4.2 Unlabelled restricted cycles and multi-sets.

Note that the unlabelled version of the CYC operator expressed by

$$\text{CYC}(\mathcal{A})(z) = \sum_{k \geq 1} \frac{\varphi(k)}{k} \log \frac{1}{1 - A(z^k)} \tag{77}$$

is in fact a weighted sum of already familiar labelled operators, for which we have already provided efficient self-concordant barriers. Again, restricted variations of the corresponding operators $\text{MSET}_{>0}$, $\text{MSET}_{>k}$ and $\text{CYC}_{>k}$ cannot be efficiently treated by DCP, but the usage of self-concordant barriers can provide some insight. For example, $\text{MSET}_{>0}$ defined as

$$\text{MSET}_{>0}(\mathcal{A})(z) = \exp \left(\sum_{k \geq 0} \frac{A(z^k)}{k} \right) - 1, \tag{78}$$

can be represented as a composition with $e^x - 1$, whose log-exp transform is $E(x) = \log(e^{e^x} - 1)$.

Next, an unlabelled cycle construction with a lower bound on the number of components $\text{CYC}_{\geq m}$ can be expressed as

$$\text{CYC}_{\geq m}(\mathcal{A})(z) = \sum_{k \geq 1} \frac{\varphi(k)}{k} \log_{\geq \lceil m/k \rceil} \frac{1}{1 - A(z^k)} \quad \text{where} \quad \log_{\geq m} \frac{1}{1 - x} = \sum_{k \geq m} \frac{x^k}{k}. \tag{79}$$

Again, as in the case of unlabelled compositions, the sequence of values of the generating functions $(A(z^k))_{k=1}^{\infty}$ tends to $A(0) = 0$ at the speed of a geometric progression. Note that this gives an efficient composition scheme for unlabelled cycles as well, using

$$L_{>k}(x) = \log \left(\log \frac{1}{1 - e^x} - \sum_{n=1}^k \frac{e^{kx}}{k} \right). \tag{80}$$

Finally, let us sketch how to express restricted unlabelled sets using the self-concordant barriers for restricted labelled sets. We start by considering two examples, $\text{MSET}_{\geq 2}$ and $\text{MSET}_{\geq 3}$.

For sets with at least two elements we have

$$\text{MSET}_{\geq 2}(\mathcal{A})(z) = \text{MSET}(\mathcal{A})(z) - 1 - A(z) = \exp \left(\sum_{k \geq 1} \frac{A(z^k)}{k} \right) - 1 - A(z). \tag{81}$$

If we set $h_1 = A(z)$ and $h_2 = \sum_{k \geq 2} \frac{A(z^k)}{k}$, then the resulting value can be rewritten as

$$\exp(h_1 + h_2) - 1 - h_1 = \exp(h_1 + h_2) - 1 - (h_1 + h_2) + h_2. \tag{82}$$

The required composition scheme is obtained by composing with the log-exp transform of the function $e^x - x - 1$ which is $E_{\geq 2}(x) := E_{>1}(x) = \log(e^{e^x} - e^x - 1)$.

For sets with at least three elements it holds

$$\text{MSET}_{\geq 3}(\mathcal{A})(z) = \exp \left(\sum_{k \geq 1} \frac{A(z^k)}{k} \right) - 1 - A(z) - \frac{A(z)^2}{2} - \frac{A(z^2)}{2}. \tag{83}$$

Again, by introducing auxiliary variables $h_1 = A(z)$, $h_2 = \frac{A(z^2)}{2}$, $h_3 = \sum_{k \geq 3} \frac{A(z^k)}{k}$, and $H = h_1 + h_2 + h_3$, we can rewrite the operator as

$$\exp(h_1 + h_2 + h_3) - 1 - h_1 - \frac{h_1^2}{2} - h_2 = \exp(H) - 1 - H - \frac{H^2}{2} + h_3 + R \tag{84}$$




Figure 3. Optimal biased expectation.

where R is a positive quantity obtained by subtracting $\frac{h_1^2}{2}$ from $\frac{(h_1+h_2+h_3)^2}{2}$. Again, we obtain a composition scheme by composing with $E_{\geq 3}(x) = E_{>2}(x)$.

Finally, for sets with general restrictions the above scheme can be repeated for multi-sets with forbidden cardinalities by subtracting $MSET_{=m}$ obtained by

$$MSET_{=m}(\mathcal{A})(z) = [u^m] \exp \left(\sum_{k \geq 1} \frac{A(z^k)u^k}{k} \right). \tag{85}$$

Afterwards, the result can be represented by adding missing parts to $MSET_{>m}$. We do not provide a detailed analysis of this case, as for large m this procedure becomes less and less efficient due to the potential super-polynomial increase in m of the number of the required terms.

5. Optimal biased expectation and tuning precision

In the current section, we address the choice of the optimal expected value for the Boltzmann sampler with prescribed anticipated rejection domain. We also investigate the precision of the control vector \mathbf{x} required to guarantee a linear time, approximate-size anticipated rejection sampling scheme.

In what follows, we analyse the rejection cost independently for each parameter. That is to say, we evaluate the cost of the anticipated rejection until the value of the investigated parameter N falls into its prescribed tolerance window $[n_1, n_2]$. The sum of these costs taken over all the parameters is an upper bound for costs of a global rejection scheme (which, of course, could be better because of some overlapping rejections).

We distinguish two types of behaviours. The first concerns parameters admitting a singular behaviour of type $\left(1 - \frac{z}{\rho}\right)^{-\alpha}$ which corresponds, for instance, to the size parameter of *peak* and *flat* distributions described in [BGR15]. In this case, we will prove that for each α the required precision is of order $O\left(\frac{1}{N}\right)$ where N is the approximate size of the expected outcome.

The second case concerns so-called *bumpy* distributions, appearing in numerous situations, such as parameter analysis in strongly connected algebraic or rational grammars. Since the distribution is concentrated around its mean value, the needed precision is related to the standard deviation and is, usually, of order $O\left(\frac{1}{\sqrt{N}}\right)$. In contrast, the size of the target window is not linear, as for the non-concentrated case, but instead is of order \sqrt{N} , or of a different order n^γ .

In the first case, we show that the standard target tuning expectation $\mathbb{E}(N) = n := \frac{n_1+n_2}{2}$ involving the admissible window $[n_1, n_2]$ does *not* minimise the overall cost of rejections, if anticipated rejection is taken into account. Instead, $\mathbb{E}(N) = \beta n$, for some suitable parameter $\beta(n_1, n_2)$ should be used, cf. Figure 3. In what follows, we describe how to find the optimal value of β and how it improves the rejection cost, compared to the usual $n := \frac{n_1+n_2}{2}$ parameter.

5.1 Non-concentrated case

For simplicity, we assume that the analysis concerns the size parameter z . Let $C_{\mathbf{u}}(z)$ be the generating function of a class \mathcal{C} for which the parameters \mathbf{u} are fixed (i.e. introduce some additional

weights). Assume that $C_u(z)$ is Δ -analytic and there exist analytic functions $\alpha(z)$, $\beta(z)$, and a corresponding singularity ρ (all depending on u) such that as $z \rightarrow \rho$ it holds

$$C(z) \sim a(z) - b(z) \left(1 - \frac{z}{\rho}\right)^{-\alpha}. \tag{86}$$

Let $C^{<n_1}(z)$, $C^{>n_2}(z)$ and $C^{[n_1, n_2]}(z)$ denote the generating functions corresponding to the subclasses of objects of size strictly smaller than n_1 , strictly greater than n_2 , and in between (inclusively) n_1 and n_2 , respectively. Let T_n stand for the cumulative size of the generated and rejected objects produced by a sampler calibrated with parameter x , and an admissible size window $[(1 - \varepsilon)n, (1 + \varepsilon)n]$. From [Duc+04, Theorem 7.3], the corresponding probability generating function

$$F(u, x) = \sum_k \mathbb{P}(T_n = k) u^k, \tag{87}$$

satisfies

$$F(u, x) = \left(1 - \frac{1}{C(x)} (C^{<(1-\varepsilon)n}(ux) + C^{>(1+\varepsilon)n}(x) u^{\lceil(1+\varepsilon)n\rceil})\right)^{-1} \frac{C^{\lceil(1-\varepsilon)n, (1+\varepsilon)n\rceil}(x)}{C(x)}, \tag{88}$$

and so

$$\mathbb{E}(T_n) = \frac{x \frac{d}{dx} C^{<(1-\varepsilon)n}(x) + \lceil(1 + \varepsilon)n\rceil C^{>(1+\varepsilon)n}(x)}{C^{\lceil(1-\varepsilon)n, (1+\varepsilon)n\rceil}(x)}. \tag{89}$$

Remark 26. Notice that even if the studied parameter does not denote the object size, the expectation $\mathbb{E}(T_n)$ stays essentially similar to (89). Involved generating functions correspond then to the subclasses of objects whose *size contribution of the studied parameter* is smaller, greater, or in between the respective window parameters. It ensues that the following analysis for the size parameter can be readily translated to others parameters.

Proposition 27. *Let us consider an approximate-size, anticipated rejection Boltzmann sampler $\Gamma(C)$ for a class \mathcal{C} of the above kind. Assume that $\Gamma(C)$ is calibrated using $x_n = \rho \left(1 - \frac{\delta}{n}\right)$ for some parameter δ . Let T_n denote the cumulative size of objects rejected by $\Gamma(C)$. Then, as n tends to infinity*

$$\mathbb{E}(T_n) \sim n \kappa_c(\varepsilon, \alpha, \delta) \quad \text{where}$$

$$\kappa_c(\varepsilon, \alpha, \delta) = \frac{\int_{w=0}^{1-\varepsilon} w^\alpha e^{-\delta w} dw + \int_{w=1+\varepsilon}^{\infty} (1 + \varepsilon) w^{\alpha-1} e^{-\delta w} dw}{\int_{w=1-\varepsilon}^{1+\varepsilon} w^{\alpha-1} e^{-\delta w} dw}. \tag{90}$$

Proof. This theorem is very similar to [BGR15, Theorem 3.2]. It has an analogous proof using Euler–Maclaurin summation, with the difference that we do not anymore suppose that $\alpha = \delta$. Let us emphasise again that keeping these two parameters distinct is the key to obtaining faster samplers with lower anticipated rejection cost. \square

Note that $\kappa_c(\varepsilon, \alpha, \delta)$ is independent of n . Moreover, $\kappa_c(\varepsilon, \alpha, \delta)$ is bounded for each $\delta > 0$. Consequently, $\Gamma(C)$ calibrated with parameter x_n of precision of order $O\left(\frac{1}{n}\right)$ has expected linear time complexity.

Remark 28. The value of $\kappa_c(\varepsilon, \alpha, \delta)$ is also independent of the weight u , depending only on the type of the singularity α and the prescribed tolerance ε . Note that it provides strong stability

Table 5. Pairs of values $(\delta_{\min}, \kappa_c(\varepsilon, \alpha, \delta_{\min}))$ for respective parameters ε and α . Numerical values are rounded up to the second decimal point.

	$\varepsilon = 0.2$	$\varepsilon = 0.1$	$\varepsilon = 0.05$	$\varepsilon = 0.01$
$-\alpha = -2$	(2.77, 2.37)	(2.66, 5.58)	(2.61, 11.95)	(2.58, 62.78)
$-\alpha = -3/2$	(2.25, 2.66)	(2.16, 6.17)	(2.12, 13.12)	(2.09, 68.58)
$-\alpha = -1$	(1.74, 3.06)	(1.66, 6.97)	(1.62, 14.72)	(1.60, 76.51)
$-\alpha = -1/2$	(1.23, 3.66)	(1.16, 8.18)	(1.13, 17.18)	(1.11, 88.38)
$-\alpha = 1/2$	(0.25, 7.37)	(0.22, 15.68)	(0.21, 32.10)	(0.20, 162.91)

guarantees regarding the sampler’s performance under the change of involved parameters; most notably, in the vicinity of the values of u where the type of the singularity changes discontinuously, cf. [BP10, Ban12]. Even in this particular cases, our result guarantees the linearity of the anticipated rejection sampling scheme.

With an explicit formula (90) for $\kappa_c(\varepsilon, \alpha, \delta)$ it is possible to optimise its value for specific values of ε and α . In other words, improve the multiplicative constant in the expected number T_n . Table 5 provides some exemplary values for δ_{\min} and the corresponding $\kappa_c(\varepsilon, \alpha, \delta_{\min})$.

Remark 29. The careful reader might be surprised by the fact that δ_{\min} is not equal to α , as it is suggested in the seminal paper [Duc+04]. The reason behind this is the fact that anticipated rejection creates a small bias in the distribution, initially not taken into account. For instance, for $\alpha = 1$ and tolerance $\varepsilon = 0.1$, the best choice for β is not 1, but 1.6572067. Notably, this decreases the rejection complexity from $8.05n$ to $6.97n$.

In consequence of the introduced bias, it is no longer optimal to find x by solving $\mathbb{E}_x(N) = n$. Instead, we have to make a small correction accounting for the effect of anticipated rejection. Luckily, this is easy to provide. Recall that, asymptotically, $\mathbb{E}_x(N) = n$ is attained for $x_n = \rho(1 - \frac{\alpha}{n})$. After computing the expectation, it follows that in order to obtain $x_n = \rho(1 - \frac{\delta_{\min}}{n})$ we have to solve the corrected $\mathbb{E}_x(N) = \frac{\alpha}{\delta_{\min}}n$, instead.

5.2 Concentrated cases

In following part we consider cases where the investigated size parameter tends asymptotically to a Gaussian law, provided that the target expectation is fixed and large. Let $C(z)$ be, as usual, the generating function corresponding to the class \mathcal{C} , and N be the random variable representing the size of an object generated according to the associated Boltzmann distribution. As previously, we want to evaluate the rejection cost and find an appropriate bias parameter which minimises its value.

Before we begin, let us consider a few examples of concentrated distributions analysed using the following generalised quasi-powers theorem.

Lemma 30 (Generalised quasi-powers, [FS09, Theorem IX.13]). *Assume that, for u in a fixed complex neighbourhood Ω of 1, the probability generating functions $p_n(u)$ of non-negative discrete random variables X_n admit representations of the form*

$$p_n(u) = \exp(h_n(u)) (1 + o(1)), \tag{91}$$

uniformly with respect to u , where each $h_n(u)$ is analytic in Ω . Assume also the conditions

$$h'_n(1) + h''_n(1) \rightarrow \infty \quad \text{and} \quad \frac{h'''(u)}{(h'_n(1) + h''_n(1))^{3/2}} \rightarrow 0, \tag{92}$$

as $n \rightarrow \infty$, uniformly for $u \in \Omega$.

Then, the random variable

$$X_n^* = \frac{X_n - \mathbb{E}X_n}{\sqrt{\text{Var } X_n}} = \frac{X_n - h'_n(1)}{(h'_n(1) + h''_n(1))^{1/2}}, \tag{93}$$

converges in distribution to $\mathcal{N}(0, 1)$.

In what follows, we apply the generalised quasi-powers theorem to a series of probability generating functions capturing the outcome size distributions of example Boltzmann samplers. Therefore, the discrete random variables of interest X_n denote these outcome sizes, where n stands for the target expected size, while $n \rightarrow \infty$. We use the standard formulas for the mean and variance corresponding to the size N of objects sampled according to the Boltzmann distribution parametrised with x :

$$\mathbb{E}_x N = \mu(x) = \frac{x C'(x)}{C(x)} \quad \text{Var}_x N = \sigma(x)^2 = x \mu(x)'. \tag{94}$$

For convenience, we also use x_n to denote the solution of the tuning equation $\mu(x_n) = n$.

5.2.1 Bicoloured sets

Consider the combinatorial class \mathcal{C} consisting of finite sets with atoms of two different types (colours). The corresponding generating function satisfies $C(x) = e^{2x}$. Accordingly, the mean value $\mu(x)$ and standard deviation $\sigma(x)$ satisfy

$$\mu(x) = 2x \quad \text{and} \quad \sigma(x) = \sqrt{2x}, \tag{95}$$

whereas $x_n = n/2$.

The probability generating function $p_n(u)$ capturing the size distribution calibrated with parameter x_n satisfies therefore

$$p_n(u) = \frac{C(x_n u)}{C(x_n)} = e^{n(u-1)}. \tag{96}$$

A direct calculation reveals that all the conditions of Lemma 30 hold, and the size distribution tends to a Gaussian law as $x_n \rightarrow \infty$. In particular, we have

$$\mathbb{E}_{x_n} N \sim n \quad \text{and} \quad \text{Var}_{x_n} N \sim n. \tag{97}$$

5.2.2 Involutions

Consider the combinatorial class of involutions, that is n -element permutations π satisfying $\pi \circ \pi = \text{id}$. The respective generating function satisfies $C(x) = \exp\left(x + \frac{x^2}{2}\right)$. Consequently, the mean value $\mu(x)$, variance $\sigma(x)^2$, and tuning parameter x_n satisfy

$$\mu(x) = x(1+x), \quad \sigma(x)^2 = x(1+2x), \quad \text{and} \quad x_n = \frac{\sqrt{1+4n}-1}{2}. \tag{98}$$

And so

$$p_n(u) = \exp\left(-\frac{1}{4}(u-1)\left(\left(\sqrt{1+4n}-1\right)(u-1)-2n(u+1)\right)\right). \tag{99}$$

Again, the premises of Lemma 30 can be easily verified. In the end, we obtain a limit Gaussian distribution where

$$\mathbb{E}_{x_n} N \sim n \quad \text{and} \quad \text{Var}_{x_n} N \sim \frac{1}{2}\left(1+4n-\sqrt{1+4n}\right). \tag{100}$$

5.2.3 Set partitions

Consider the combinatorial class of set partitions for which we have $C(x) = e^{e^x - 1}$. A direct computation provides the identities

$$\mu(x) = xe^x, \quad \sigma(x)^2 = xe^x(1+x), \quad \text{and} \quad x_n = W(n), \tag{101}$$

where $W(n)$ is the Lambert function defined as the positive solution of $W(n)e^{W(n)} = n$.

In this case, $p_n(u) = \exp(h_n(u)) (1 + o(1))$ where $h_n(u) = e^{uW(n)} - e^{W(n)}$. We can easily check that

$$h'_n(1) = n, \quad h''_n(1) = e^{W(n)}W(n)^2, \quad \text{and} \quad h'''_n(u) = W(n)^3 e^{uW(n)}. \tag{102}$$

Since $W(n) = \ln n - \ln \ln n + o(1)$ we note that $h'_n(1) + h''_n(1) \rightarrow \infty$. Moreover

$$\frac{h'''_n(u)}{(h'_n(1) + h''_n(1))^{3/2}} = \frac{W(n)^3 e^{uW(n)}}{(n(W(n) + 1))^{3/2}} \sim (\ln n)^{3/2} n^{u-3/2}, \tag{103}$$

which for u fixed near one tends to 0 as $n \rightarrow \infty$.

Hence, by Lemma 30 the distribution tends, again, to a Gaussian law. In the limit we obtain

$$\mathbb{E}_{x_n}(N) \sim n \quad \text{and} \quad \text{Var}_{x_n}(N) \sim n(1 + W(n)). \tag{104}$$

5.2.4 Fragmented permutations

Next, consider the class of fragmented permutations, that is sets of non-empty labelled sequences, see [FS09, Example VIII.7, p. 562]. The corresponding generating function satisfies $C(x) = \exp\left(\frac{x}{1-x}\right)$. Consequently, we find that

$$\mu(x) = \frac{x}{(1-x)^2}, \quad \sigma(x)^2 = x \left(\frac{2x}{(1-x)^3} + \frac{1}{(1-x)^2} \right) \quad \text{and} \quad x_n = \frac{1 + 2n - \sqrt{1 + 4n}}{2n}. \tag{105}$$

Note that as $n \rightarrow \infty$, the tuning parameter $x_n \rightarrow 1$. Again, we verify that, as $x_n \rightarrow \infty$, the function $h_n(u) = \frac{ux_n}{1-ux_n} - \frac{x_n}{1-x_n}$ satisfies both

$$h'(x_n) + h''(x_n) \rightarrow \infty \quad \text{and} \quad \frac{h'''(x_n u)}{(h'_n(x_n) + h''_n(x_n))^{3/2}} \rightarrow 0. \tag{106}$$

However, now $h_n(u)$ is not analytic at $u = 1$. We cannot therefore apply Lemma 30. Nonetheless, it is still possible to prove that the limiting distribution is Gaussian, using the explicit formula for the number of fragmented permutations from [FS09].

From the above examples, we see that there is a variety of different behaviours regarding the variance of the limit Gaussian distribution. In what follows we will show how the mean value is connected to the anticipated rejection cost (89) in the Gaussian case.

Theorem 31 (*Tuning precision for concentrated distributions*). *Let the size parameter asymptotically follow a Gaussian law with expectation $\mu(x)$ and standard deviation $\sigma(x)$, and let $X(n)$ be the inverse function of $\mu(x)$, that is $\mu(X(n)) = n$. Denote by $x_{n,\delta}$ the biased tuning value*

$$x_{n,\delta} := x_n + \delta \sigma(x_n) \frac{dX(n)}{dn}. \tag{107}$$

Then the anticipated rejection cost T_n with the target tolerance window $[n - \varepsilon \sigma(x_n), n + \varepsilon \sigma(x_n)]$ satisfies, when n tends to infinity:

$$\mathbb{E}_{x_{n,\delta}}(T_n) \sim n\kappa(\varepsilon, \delta), \quad \kappa(\varepsilon, \delta) = \frac{1 + \Phi(-\varepsilon - \delta) - \Phi(\varepsilon - \delta)}{\Phi(\varepsilon - \delta) - \Phi(-\varepsilon - \delta)}, \tag{108}$$

where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-w^2/2} dw$ is the Gaussian distribution function. The minimal cost is achieved when $\delta = \delta_{\min} = 0$.

Proof. Using the Euler–MacLaurin estimate similarly to [BGR15], we obtain the following estimates for $C^{\lfloor n_1, n_2 \rfloor}(x)$, $C^{>n_2}(x)$ and $\frac{d}{dx} C^{<n_1}(x)$, when $x = x_n$ as $n \rightarrow \infty$:

$$\begin{aligned} \frac{C^{\lfloor n_1, n_2 \rfloor}(x)}{C(x)} &= \sum_{n=n_1}^{n_2} \frac{[z^n]C(z)x^n}{C(x)} \sim \frac{1}{\sqrt{2\pi}\sigma(x)} \int_{n_1}^{n_2} e^{-\frac{1}{2}\left(\frac{n-\mu(x)}{\sigma(x)}\right)^2} dn \\ &= \Phi\left(\frac{n_2 - \mu(x)}{\sigma(x)}\right) - \Phi\left(\frac{n_1 - \mu(x)}{\sigma(x)}\right), \\ \frac{C^{>n_2}(x)}{C(x)} &\sim 1 - \Phi\left(\frac{n_2 - \mu(x)}{\sigma(x)}\right), \text{ and} \\ \frac{x \frac{d}{dx} C^{<n_1}(x)}{C(x)} &= \sum_{n=0}^{n_1} n \frac{[z^n]C(z)x^n}{C(x)} \sim \frac{1}{\sqrt{2\pi}\sigma(x)} \int_0^{n_1} [(n - \mu(x)) + \mu(x)] e^{-\frac{1}{2}\left(\frac{n-\mu(x)}{\sigma(x)}\right)^2} dn \\ &\sim \frac{2\sigma(x)}{\sqrt{2\pi}} \left[1 - e^{-\frac{1}{2}\left(\frac{n_1-\mu(x)}{\sigma(x)}\right)^2} \right] + \mu(x)\Phi\left(\frac{n_1 - \mu(x)}{\sigma(x)}\right). \end{aligned} \tag{109}$$

Since $\mu(x) \gg \sigma(x)$ for $x = x_n$ as $n \rightarrow \infty$, the second summand dominates the first summand in the last expression.

Combining these quantities, and by substituting into $\mathbb{E}(T_n) = \frac{x C'^{<n_1}(x) + n_2 C^{>n_2}(x)}{C^{\lfloor n_1, n_2 \rfloor}(x)}$, we get the estimated cost of anticipated rejection:

$$\mathbb{E}_x(T_n) = \frac{\mu(x)\Phi\left(\frac{n_1 - \mu(x)}{\sigma(x)}\right) + n_2 \left[1 - \Phi\left(\frac{n_2 - \mu(x)}{\sigma(x)}\right) \right]}{\Phi\left(\frac{n_2 - \mu(x)}{\sigma(x)}\right) - \Phi\left(\frac{n_1 - \mu(x)}{\sigma(x)}\right)}. \tag{110}$$

Let us choose $x = x_{n,\delta}$ in such a way that $\mu(x_{n,\delta}) = n + \delta\sigma(x_n)$. Using the first two terms of the Taylor expansion of $\mu(x_{n,\delta})$ around x_n , we can show that this value is asymptotically $x_{n,\delta} \sim x_n + \delta\sigma(x_n) \frac{dX(n)}{dn}$. This finally allows to obtain $\kappa(\varepsilon, \delta) = \lim_{n \rightarrow \infty} \frac{\mathbb{E}T_n}{n}$:

$$\kappa(\varepsilon, \delta) = \frac{1 + \Phi(-\varepsilon - \delta) - \Phi(\varepsilon - \delta)}{\Phi(\varepsilon - \delta) - \Phi(-\varepsilon - \delta)}. \tag{111}$$

This expression achieves its minimum when $\Phi(\varepsilon - \delta) - \Phi(-\varepsilon - \delta)$ achieves its maximum. Its derivative has only one root $\delta = 0$ which corresponds to the global maximum. Therefore, $\delta_{\min} = 0$. □

Remark 32. In the case when the standard deviation $\sigma(x_n)$ is negligible compared to the mean value $\mu(x_n)$ as $n \rightarrow \infty$, the rejection cost can be shown to be asymptotically equal

$$\mathbb{E}_{x_{n,\delta}} T_n \sim n\kappa(\varepsilon, \delta) = n \frac{\int_{(-\infty, n_1) \cup (n_2, +\infty)} p(x) dx}{\int_{(n_1, n_2)} p(x) dx}, \tag{112}$$

where $p(x)$ is the probability density function of the limiting distribution. In the case when the limiting distribution is not symmetric, and the probability density function is known, the corresponding bias can be computed as $\delta_{\min} = \arg \min \kappa(\varepsilon, \delta)$.

6. Paganini: a multi-parametric tuner prototype

To illustrate the effectiveness of our tuning procedure, we developed Paganini⁴ — a lightweight Python library implementing the *tuning-as-convex-optimisation* idea. Our software relies on cvxpy, a Python-embedded modelling language for Disciplined Convex Programming (DCP) [GBY06]. With its help, Paganini is able to automatically compose, and solve adequate optimisation problems so to compute the parameter vector corresponding to the user-defined expectations.

6.1 Implementation details

Due to the imposed restrictions of Disciplined Convex Programming, Paganini supports a strict, though substantial subset of admissible constructions. In both the labelled and unlabelled case, Paganini provides the basic empty class ε , disjoint sum $+$, and Cartesian product \times operations. More involved constructions are briefly discussed below.

6.1.1 Sequence operator

Consider $\text{SEQ}(\mathcal{A})$ for some class \mathcal{A} (either labelled or unlabelled). By definition, $\text{SEQ}(\mathcal{A}) = \varepsilon + \mathcal{A} \times \text{SEQ}(\mathcal{A})$ and so we can treat $\text{SEQ}(\mathcal{A})$ as a new, auxiliary variable with a corresponding definition of the above shape. The log-exp transform of $\text{SEQ}(\mathcal{A})$ takes then the form of an elementary DCP log-sum-exp function:

$$\text{SEQ}(\mathcal{A})(z) = \frac{1}{1 - \mathcal{A}(z)} \xrightarrow{\text{log-exp}} \sigma \geq \log(1 + e^{\alpha + \sigma}), \tag{113}$$

where $e^\alpha = \mathcal{A}(z)$ and $e^\sigma = \text{SEQ}(\mathcal{A})(z)$.

Likewise, since

$$\begin{aligned} \text{SEQ}_{=k}(\mathcal{A})(z) &= \mathcal{A}(z)^k, & \text{SEQ}_{\geq k}(\mathcal{A})(z) &= \frac{\mathcal{A}(z)^k}{1 - \mathcal{A}(z)}, \quad \text{and} \\ \text{SEQ}_{\leq k}(\mathcal{A})(z) &= \sum_{i=0}^k \mathcal{A}(z)^i, \end{aligned} \tag{114}$$

it is readily possible to translate SEQ with its restricted variants into valid DCP constraints.

6.1.2 Pólya structures

Consider $\text{MSET}(\mathcal{A})$ and its log-exp variant:

$$\text{MSET}(\mathcal{A})(z) = \exp\left(\sum_{k \geq 1} \frac{\mathcal{A}(z^k)}{k}\right) \xrightarrow{\text{log-exp}} \mu \geq \sum_{k \geq 1} \frac{e^{\alpha k}}{k}, \tag{115}$$

where $e^{\alpha k} = \mathcal{A}(z^k)$ and $e^\mu = \text{MSET}(\mathcal{A})(z)$.

The right-hand side of the MSET log-exp transform is an infinite sum of exponents with positive weights. For practical purposes, we can therefore truncate the series at a finite threshold and notice that it conforms with DCP rules. What remains is to construct constraints for the respective *diagonals* $\mathcal{A}(z^k)$ based on the definition of $\mathcal{A}(z)$.

It is also possible to handle the restricted $\text{MSET}_{=k}$. Following [FFP07, Section 2.4] we notice that $\text{MSET}_{=k}(\mathcal{A})(z)$ can be expressed as

⁴see <https://github.com/maciej-bendkowski/paganini> and <http://paganini.readthedocs.io>.

$$\text{MSET}_{=k}(\mathcal{A})(z) = \sum_{P \in \mathcal{P}_k} \prod_{i=1}^k \frac{A(z^i)}{i^{n_i} n_i!}, \tag{116}$$

where \mathcal{P}_k consists of so-called *partition sequences* of size k , that is sequences (n_i) of natural numbers satisfying the condition $\sum_{i=1}^k i n_i = k$. In this form, (116) unfolds to a polynomial in $A(z), A(z^2), \dots, A(z^k)$ with positive coefficients. Consequently, $\text{MSET}_{=k}$ can be converted to a DCP constraint just like a regular algebraic equation. Following the same idea, we can handle $\text{MSET}_{\leq k}$ as

$$\text{MSET}_{\leq k}(\mathcal{A})(z) = 1 + \sum_{i=1}^k \text{MSET}_{=i}(\mathcal{A})(z). \tag{117}$$

The $\text{MSET}_{\geq k}$ variant is much more involved. Since the difference of convex functions is not necessarily a convex function itself, we cannot directly translate the defining

$$\text{MSET}_{\geq k}(\mathcal{A}) = \text{MSET}(\mathcal{A}) - \text{MSET}_{<k}(\mathcal{A}), \tag{118}$$

into a valid DCP constraint. Let us recall, however, that for $k = 1$ in the case when $\mathcal{A} = z_1 + z_2 + \dots + z_d$ it is possible to rewrite $\text{MSET}_{\geq k}$ so to avoid subtraction altogether and compose a corresponding DCP constraint (cf. Section 7.4 and more generally, Section 2.2).

For general k , a more heuristic approach using *Disciplined Convex-Concave Programming* (DCCP) might be preferred, see [She+16]. Consider the following *exp* transform:

$$\text{MSET}_{\geq 1}(\mathcal{A})(z) \xrightarrow{\text{exp}} e^\mu + 1 \geq \exp \left(\sum_{k \geq 1} \frac{e^{\alpha_k}}{k} \right), \tag{119}$$

where $e^{\alpha_k} = A(z^k)$ and $e^\mu = \text{MSET}(\mathcal{A})_{\geq 1}(z)$.

Here we have two convex expressions on both sides of the inequality. Although such constraints do not conform with DCP rules, they are allowed in the DCCP framework, and can be therefore included in the problem statement.

Now, let us focus the cycle construction $\text{CYC}(\mathcal{A})$. Note that

$$\text{CYC}(\mathcal{A})(z) = \sum_{k \geq 1} \frac{\varphi(k)}{k} \log \frac{1}{1 - A(z^k)} \xrightarrow{\text{log-exp}} \gamma \geq \log \left(\sum_{k \geq 1} \frac{\varphi(k)}{k} \log \frac{1}{1 - e^{\alpha_k}} \right), \tag{120}$$

where $e^{\alpha_k} = A(z^k)$, $e^\gamma = \text{CYC}(\mathcal{A})(z)$, and $\varphi(k)$ is the Euler totient function.

Unfortunately, such a constraint does not meet the requirements of DCP, even though its right-hand side is convex. On the other hand, the restricted $\text{CYC}_{=k}(\mathcal{A})$ satisfies

$$\text{CYC}_{=k}(\mathcal{A})(z) = \frac{1}{k} \sum_{i|k} \varphi(i) A(z^i)^{\frac{k}{i}}, \tag{121}$$

and so it is possible to express its log-exp transform as a standard DCP log-sum-exp function. In order to emulate the unrestricted $\text{CYC}(\mathcal{A})$ operator, one can either use DCCP, reformulating (120) as a DCCP constraint, or use the relation

$$\text{CYC}(\mathcal{A})(z) = \sum_{k \geq 0} \text{CYC}_{=k}(\mathcal{A})(z), \tag{122}$$

with a (heuristically chosen) truncation threshold. See further comments about the cycle construction and the function $\log \log \frac{1}{1 - e^x}$ in Section 4.4.

Finally, consider the power set construction $\text{PSET}(\mathcal{A})$:

$$\text{PSET}(\mathcal{A})(z) = \exp \left(\sum_{k \geq 1} \frac{(-1)^{k-1}}{k} A(z^k) \right) \xrightarrow{\text{log-exp}} \pi \geq \sum_{k \geq 1} \frac{(-1)^{k-1} e^{\alpha k}}{k}, \quad (123)$$

where $e^{\alpha k} = A(z^k)$ and $e^\pi = \text{PSET}(\mathcal{A})(z)$.

Due to the alternating summation and subtraction in (123), the right-hand side of the constraint cannot be expressed as an elementary, convex, DCP expression. Consequently, it is not supported in our prototype implementation.

6.1.3 Labelled constructions

Consider the labelled set operator $\text{SET}(\mathcal{A})$. Recall that for both the restricted and unrestricted variants we have

$$\begin{aligned} \text{SET}(\mathcal{A})(z) = e^{A(z)} &\xrightarrow{\text{log-exp}} \sigma \geq e^\alpha \\ \text{SET}_{=k}(\mathcal{A})(z) = \frac{1}{k!} A(z)^k &\xrightarrow{\text{log-exp}} \sigma_k \geq \log \frac{e^{\alpha k}}{k!}, \end{aligned} \quad (124)$$

where $e^\alpha = A(z)$, $e^\sigma = \text{SET}(\mathcal{A})(z)$, and $e^{\sigma k} = \text{SET}_{=k}(\mathcal{A})(z)$.

In this form, it is clear that both log-exp transformations can be expressed in terms of elementary DCP functions. Hence, both $\text{SET}(\mathcal{A})$ and $\text{SET}_{=k}(\mathcal{A})$ can be effectively handled.

Let us now focus on the final cycle operator $\text{CYC}(\mathcal{A})$. Note that

$$\begin{aligned} \text{CYC}(\mathcal{A})(z) = \log \frac{1}{1 - A(z)} &\xrightarrow{\text{log-exp}} \gamma \geq \log \log \frac{1}{1 - e^\alpha} \\ \text{CYC}_{=k}(\mathcal{A})(z) = \frac{1}{k} A(z)^k &\xrightarrow{\text{log-exp}} \gamma_k \geq \log \frac{e^{\alpha k}}{k}, \end{aligned} \quad (125)$$

where $e^\alpha = A(z)$, $e^\gamma = \text{CYC}(\mathcal{A})(z)$, and $e^{\gamma k} = \text{CYC}_{=k}(\mathcal{A})(z)$.

The log-exp transform of $\text{CYC}_{=k}$ is an elementary DCP constraint. Unfortunately, the same does not hold for CYC and, in general, expressions in form of

$$\log \log \frac{1}{1 - e^x}. \quad (126)$$

Although (126) is convex, it cannot be modelled in terms of basic DCP functions. Consequently, heuristic approaches (as discussed above) or alternative convex programming techniques like the interior point method should be applied, instead. More labelled constructions, and the ways to bypass this limitation are discussed in Section 4.4.

6.2 Sampler construction

Given a combinatorial specification enriched with user-specified parameters and their target expectations, it is possible to *mechanically* compute respective tuning values, and compile dedicated samplers. To illustrate this point, we implemented a *sampler compiler* called Boltzmann Brain⁵ supporting algebraic, and in particular rational, specifications.

Constructed samplers for general algebraic specifications use the principle of *anticipated rejection* whereas samplers for rational specifications implement the idea of *interruptible sampling*, see [BGR15]. In addition, both sampler types are endowed with near-optimal decision trees based on established branching probabilities [KY76]. Combined, these optimisations lead to remarkably efficient samplers, supporting all of the algebraic specifications included in the current paper, see Section 7.

⁵see <https://github.com/maciej-bendkowski/boltzmann-brain>.



Figure 4. Examples of admissible tiles.




Figure 5. Eight random $n \times 7$ tilings of areas in the interval $[500;520]$ using in total 95 different tiles.

7. Applications

In this section, we present several examples illustrating the wide range of applications of our tuning techniques.

7.1 Polyomino tilings.

We start with a benchmark example of a rational specification defining $n \times 7$ rectangular tilings using up to 126 different tile variants (a toy example of so-called transfer matrix models, cf. [FS09, Chapter V.6, Transfer matrix models]).

We begin the construction with defining the set T of admissible tiles. Each tile $t \in T$ consists of two horizontal layers. The base layer is a single connected block of width $w_t \leq 6$. The second layer, placed on top of the base one, is a subset (possibly empty) of w_t blocks, see Figure 4. For presentation purposes, each tile is given a unique, distinguishable colour.

Next, we construct the asserted rational specification following the general construction method of defining a deterministic automaton with one state per each possible partial tiling configuration using the set T of available tiles. Tracking the evolution of attainable configurations while new tiles arrive, we connect relevant configurations by suitable transition rules in the automaton. Finally, we (partially) minimise the constructed automaton removing states unreachable from the initial empty configuration. Once the automaton is created, we tune the tiling sampler such that the target colour frequencies are uniform, that is each colour occupies, on average, approximately $\frac{1}{126} \approx 0.7936\%$ of the outcome tiling area. Figure 5 depicts an exemplary tiling generated by our sampler.

Remark 33. The automaton corresponding to our $n \times 7$ tiling sampler consists of more than 2 000 states and 28, 000 transitions. Pushing this toy construction to its extreme, we were able to

Table 6. Empirical frequencies of the node degree distribution.

Node degree	0	1	2	3	4	5	6	7	8	9
Tuned frequency	---	---	1.00%	1.00%	1.00%	1.00%	1.00%	1.00%	1.00%	1.00%
Observed frequency	35.925%	56.168%	0.928%	0.898%	1.098%	0.818%	1.247%	0.938%	1.058%	0.918%
Default frequency	50.004%	24.952%	12.356%	6.322%	2.882%	1.984%	0.877%	0.378%	0.169%	0.069%




Figure 6. Two random plane trees with degrees in the set $D = \{0, \dots, 9\}$. On the left, a tree of size in between 500 and 550; on the right, a tree of size in the interval $[10, 000; 10, 050]$.

develop a sampler for $n \times 9$ tilings, using 1022 different tiles. The corresponding automaton has more than 19, 000 states and 357, 000 edges.

We remark that both examples are a notable improvement over the work of Bodini and Ponty [BP10] who were able to sample $n \times 6$ tilings using 7 different tiles with a corresponding automaton consisting of roughly 1 500 states and 3 200 transitions.

7.2 Simply-generated trees with node degree constraints.

Next, we give an example of simple varieties of plane trees with fixed sets of admissible node degrees, satisfying the general equation

$$y(z) = z\phi(y(z)) \quad \text{for some polynomial } \phi: \mathbb{C} \rightarrow \mathbb{C}. \tag{127}$$

Let us consider the case of plane trees where nodes have degrees in the set $D = \{0, \dots, 9\}$, that is $\phi(y(z)) = a_0 + a_1y(z) + a_2y(z)^2 + \dots + a_9y(z)^9$. Here, the numbers $a_0, a_1, a_2, \dots, a_9$ are non-negative real coefficients. We tune the corresponding algebraic specification so to achieve a target frequency of 1% for all nodes of degrees $d \geq 2$. Frequencies of nodes with degrees $d \leq 1$ are left undistorted. For presentation purposes, all nodes with equal degree are given the same unique, distinguishable colour. Figure 6 depicts two exemplary trees generated in this manner.

Empirical frequencies for the right tree of Figure 6 and a simply-generated tree of size in between 10,000 and 10,050 with default node degree frequencies are included in Table 6.

We briefly remark that for this particular problem, Bodini, David and Marchal proposed a different, bit-optimal sampling procedure for random trees with given partition of node degrees [BDM16].



Figure 7. On the left, a random λ -term of size in the interval [500;550]; on the right, a larger example of a random λ -term of size between 10,000 and 10,050.

7.3 Variable distribution in plain λ -terms.

To exhibit the benefits of distorting the intrinsic distribution of various structural patterns in algebraic data types, we present an example specification defining so-called plain λ -terms with explicit control over the distribution of de Bruijn indices.

In their nameless representation due to de Bruijn [Bru72] λ -terms are defined by the formal grammar $L ::= \lambda L | (LL) | D$ where $D = \{0, \underline{1}, \underline{2}, \dots\}$ is an infinite denumerable set of so-called indices (cf. [Ben+17, GG16]). Assuming that we encode de Bruijn indices as a sequence of successors of zero (i.e. use a unary base representation), the class \mathcal{L} of plain λ -terms can be specified as $\mathcal{L} = \mathcal{Z}\mathcal{L} + \mathcal{Z}\mathcal{L}^2 + \mathcal{D}$ where $\mathcal{D} = \mathcal{Z}\text{SEQ}(\mathcal{Z})$. In order to control the distribution of de Bruijn indices, we need a more explicit specification for de Bruijn indices. For instance:

$$\mathcal{D} = \mathcal{U}_0\mathcal{Z} + \mathcal{U}_1\mathcal{Z}^2 + \dots + \mathcal{U}_k\mathcal{Z}^{k+1} + \mathcal{Z}^{k+2}\text{SEQ}(\mathcal{Z}).$$

Here, we roll out the $k + 1$ initial indices and assign distinct marking variables to each one of them, leaving the remainder sequence intact. In doing so, we are in a position to construct a sampler tuned to enforce a uniform distribution of 8% among all marked indices, that is indices $0, \underline{1}, \dots, \underline{8}$, distorting in effect their intrinsic geometric distribution.

Figure 7 illustrates two random λ -terms with such a new distribution of indices. For presentation purposes, each index in the left picture is given a distinct colour.

Empirical frequencies for the right term of Figure 7 and a plain λ -term of size in between 10,000 and 10,050 with default de Bruijn index frequencies are included in table 7.

Let us note that algebraic data types, an essential conceptual ingredient of various functional programming languages such as Haskell or OCaml, and the random generation of their inhabitants satisfying additional structural or semantic properties is one of the central problems present in the field of property-based software testing (see, e.g. [CH00, Pal12]). In such an approach to software quality assurance, programmer-declared function invariants (so-called properties) are checked using random inputs, generated accordingly to some predetermined, though usually not rigorously controlled, distribution. In this context, our techniques provide a novel and effective

Table 7. Empirical frequencies (with respect to the term size) of index distribution.

Index	<u>0</u>	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	<u>6</u>	<u>7</u>	<u>8</u>
Tuned frequency	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%	8.00%
Observed frequency	7.50%	7.77%	8.00%	8.23%	8.04%	7.61%	8.53%	7.43%	9.08%
Default frequency	21.91%	12.51%	5.68%	2.31%	0.74%	0.17%	0.20%	0.07%	- - -

approach to generating random algebraic data types with fixed average frequencies of type constructors. In particular, using our methods it is possible to *boost* the intrinsic frequencies of certain desired subpatterns or *diminish* those which are unwanted.

7.4 Weighted partitions.

Integer partitions are one of the most intensively studied objects in number theory, algebraic combinatorics and statistical physics. Hardy and Ramanujan obtained the famous asymptotics which has later been refined by Rademacher [FS09, Chapter VIII]. In his article [Ver96], Vershik considers several combinatorial examples related to statistical mechanics and obtains the limit shape for a random integer partition of size n with $\alpha\sqrt{n}$ parts and summands bounded by $\theta\sqrt{n}$. Let us remark that Bernstein, Fahrbach, and Randall [BFR18] have recently analysed the complexity of exact-size Boltzmann sampler for weighted partitions. In the model of ideal gas, there are several particles (bosons) which form a so-called assembly of particles. The overall energy of the system is the sum of the energies $\Lambda = \sum_{i=1}^N \lambda_i$ where λ_i denotes the energy of i -th particle. We assume that energies are positive integers. Depending on the energy level λ , there are $j(\lambda)$ possible available states for each particle; the function $j(\lambda)$ depends on the physical model. Since all the particles are indistinguishable, the generating function $P(z)$ for the number of assemblies $p(\Lambda)$ with energy Λ takes the form

$$P(z) = \sum_{\Lambda=0}^{\infty} p(\Lambda)z^{\Lambda} = \prod_{\lambda>0} \frac{1}{(1 - z^{\lambda})^{j(\lambda)}} . \tag{128}$$

In the model of d -dimensional harmonic trap (also known as the Bose-Einstein condensation) according to [CMZ99, HHA97, LR08] the number of states for a particle with energy λ is $d + \lambda - \binom{1}{\lambda}$ so that each state can be represented as a multiset with λ elements having d different colours. Accordingly, an assembly is a multiset of particles (since they are bosons and hence indistinguishable) therefore the generating function for the number of assemblies takes the form

$$P(z) = \text{MSET}_{\geq 1}(\text{MSET}_{\geq 1}(\mathcal{Z}_1 + \dots + \mathcal{Z}_d)) . \tag{129}$$

It is possible to control the expected frequencies of colours using our tuning procedure and sample resulting assemblies as Young tableaux. Each row corresponds to a particle whereas the colouring of the row displays the multiset of included colours, see Figure 8. We also generated weighted partitions of expected size 1000 (which are too large to display) with tuned frequencies of 5 colours, see Table 8.

Let us briefly explain our generation procedure. Boltzmann sampling for the outer MSET operator is described in [FFP07]. The sampling of inner $\text{MSET}_{\geq 1}(\mathcal{Z}_1 + \dots + \mathcal{Z}_d)$ is more delicate. The generating function for this multiset can be written as

Table 8. Empirical frequencies of colours observed in random partition.

Colour index	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>5</u>	size
Tuned frequency	0.03	0.07	0.1	0.3	0.5	1000
	0.03	0.08	0.07	0.33	0.49	957
	0.03	0.06	0.09	0.28	0.54	1099
Observed frequency	0.03	0.08	0.09	0.34	0.46	992
	0.04	0.07	0.1	0.31	0.49	932
	0.04	0.09	0.1	0.25	0.52	1067




Figure 8. Young tableaux corresponding to Bose–Einstein condensates with expected numbers of different colours. Notation $[c_1, c_2, \dots, c_k]$ provides the expected number c_j of the j -th colour, c_k^m is a shortcut for m occurrences of c_k .

$$\text{MSET}_{\geq 1}(z_1 + \dots + z_d) = \prod_{i=1}^d \frac{1}{1 - z_i} - 1 . \tag{130}$$

In order to correctly calculate the branching probabilities, we introduce slack variables s_1, \dots, s_d satisfying $(1 + s_i) = (1 - z_i)^{-1}$. Boltzmann samplers for the newly determined combinatorial classes $\Gamma \mathcal{S}_i$ are essentially Boltzmann samplers for $\text{SEQ}_{\geq 1}(Z_i)$. Let us note that after expanding brackets the expression becomes

$$\text{MSET}_{\geq 1}(z_1 + \dots + z_d) = (s_1 + \dots + s_d) + (s_1 s_2 + \dots + s_{d-1} s_d) + \dots + s_1 s_2 \dots s_d . \tag{131}$$

The total number of summands is $2^d - 1$ where each summand corresponds to choosing some subset of colours. Finally, we pre-compute all the symmetric polynomials and efficiently handle the branching process in quadratic time using a dynamic programming approach discussed in Section 2.2.

7.5 Multi-partite rooted labelled trees

Consider a family of rooted labelled trees, such that the children of each node are not ordered. The exponential generating function of such trees $T(z)$ satisfies the equation

$$T(z) = ze^{T(z)} . \tag{132}$$



Figure 9. Random coloured Cayley tree of size 1665 drawn with Kamada–Kawai algorithm [KK89].

In this example, we suggest an alteration of this model, where the nodes on each level have distinct colours. We consider a periodic system of colouring, where the levels $1, 2, \dots, d$ have distinct colours, and then the colours repeat, that is the levels $d + 1, \dots, 2d$ have the same colours as the levels $1, 2, \dots, d$. Let u_1, \dots, u_d be the marking variables for the respective colours, and let T_1, \dots, T_d denote multivariate exponential generating functions for trees whose root is coloured respectively, with the colour $1, 2, \dots, d$. Then, these functions satisfy the system of functional equations

$$\begin{aligned}
 T_1(z, u_1, \dots, u_d) &= zu_1 e^{T_2(z, u_1, \dots, u_d)}, \\
 T_2(z, u_1, \dots, u_d) &= zu_2 e^{T_3(z, u_1, \dots, u_d)}, \\
 &\vdots \\
 T_d(z, u_1, \dots, u_d) &= zu_d e^{T_1(z, u_1, \dots, u_d)}.
 \end{aligned}
 \tag{133}$$

Using our software (see Section 6), we implement the multi-parametric tuning and sampling when $d = 10$ and the proportions of vertices of the respective colours are sorted in an arithmetic progression $0.01, 0.03, 0.05, 0.07, 0.09, 0.11, 0.13, 0.15, 0.17, 0.19$.

An example of a resulting tree of size 1665 is shown in Figure 9, and the empirical frequencies of the colours inside this tree are given in Figure 10.

Note that even in such a simple example where the dependency graph forms a cycle, and the structure of the equations is symmetrical, no simpler tuning procedure is available — no variable can be tuned separately from the others, and the resulting variable values do not form the same arithmetic pattern as the target weights. Also, let us point to a curious pattern for the numerical values of the tuning parameters can be observed in Table 9.

7.6 Otter trees

Starting from the seminal paper of Otter [Ott48], unlabelled tree-like structures play an important role in chemistry, phylogenetics [Pen82] and synthetic biology [FVS17]. Their study also helps to discover new methods for numerical solution of partial differential equations, involving automatic differentiation and construction of expression trees [MCA18].

In this section, we consider a relatively simple example of rooted unlabelled binary trees, such that the children of each node are not ordered. An additional assumption that the leaves are coloured in d distinct colours gives the following specification:

$$T(z, u_1, \dots, u_d) = z \sum_{i=1}^d u_i + \text{MSET}_2(T(z, u_1, \dots, u_d)),
 \tag{134}$$

Table 9. Numerical values for multi-partite rooted labelled trees parameters.

	z	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}
value	0.3	0.009	1.88	1.37	1.29	1.26	1.25	1.24	1.23	1.23	3.52




Figure 10. Histogram of colour frequencies for Cayley tree.




Figure 11. Random coloured Otter tree of size 1434 drawn with Kamada-Kawai algorithm [KK89].

where the $MSET_2$ operator is defined as

$$MSET_2(T(z, u_1, \dots, u_d)) = \frac{T(z, u_1, \dots, u_d)^2 + T(z^2, u_1^2, \dots, u_d^2)}{2}. \tag{135}$$

Recall that the corresponding univariate model was considered in [BLR15] where the authors constructed a system of quadratic equations which could be solved in reverse. As in the previous example (see Section 7.5), we solve the multi-parametric tuning problem when $d = 10$, and the expected colour frequencies form an arithmetic progression $0.01, 0.03, \dots, 0.19$.

After setting an appropriate truncation level (the probability of having nodes on a level h is exponentially decreasing in h), we solve the tuning problem and generate the corresponding trees, see Figure 11 for the generated tree of size 1434, and Figure 12 for empirical frequency histogram. The numerical values of the tuning parameters are given in Table 10.

Table 10. Numerical values of tuning parameters for multicoloured Otter trees.

	u_{1Z}	u_{2Z}	u_{3Z}	u_{4Z}	u_{5Z}	u_{6Z}	u_{7Z}	u_{8Z}	u_{9Z}	u_{10Z}
value	0.005	0.015	0.025	0.035	0.044	0.054	0.063	0.072	0.081	0.09

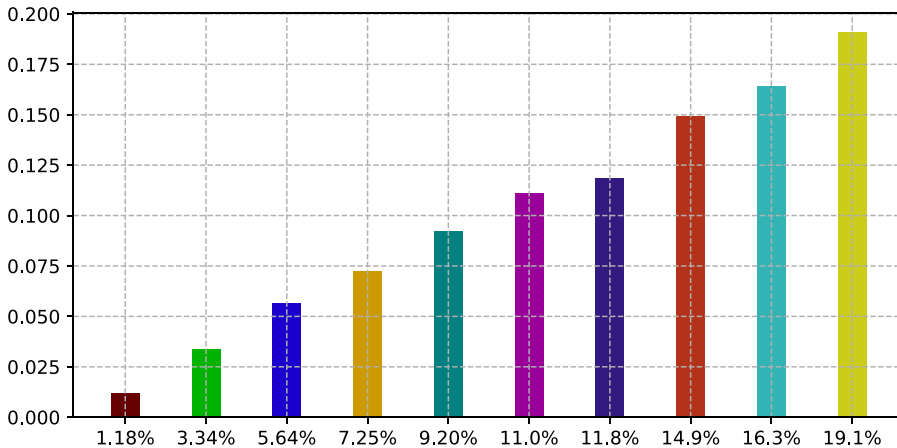


Figure 12. Histogram of colour frequencies for Otter tree.

7.7 Substitution-closed permutation classes

Permutation patterns stem from a growing body of research which originated from Knuth’s study of sorting algorithms and also later from the study of sorting networks by Tarjan and Pratt. Recently, the authors of [Bas17] have obtained a method which allows to automatically construct a combinatorial specification for permutation classes avoiding given set of permutations.

Let us start by giving a few basic definitions. A permutation $\sigma = (\sigma(1), \sigma(2), \dots, \sigma(n))$ is said to *contain a pattern* $\pi = (\pi(1), \dots, \pi(k))$ if σ has a subsequence whose terms have the same relative ordering as π . A permutation of length n is said to be *simple* if it does not contain *intervals* of length strictly in between 1 and n , where an *interval* is a contiguous sequence of indices $\{i \mid a \leq i \leq b\}$ such that the set of values $\{\sigma(i) \mid a \leq i \leq b\}$ is also contiguous. And so, for instance, the permutation from Figure 13 is not simple, since it contains an interval $\{1, 2, 3\}$ whose image is $\{5, 6, 7\}$.

Many interesting permutation classes can be described in the augmented language of context-free specifications. Given a permutation $\sigma \in S_m$ and non-empty permutations τ_1, \dots, τ_m , the *inflation* of σ by τ_1, \dots, τ_m is the permutation obtained by replacing each entry $\sigma(i)$ with an interval having the same relative ordering as τ_i . If τ_1, \dots, τ_m belong, respectively, to the classes $\mathcal{T}_1, \dots, \mathcal{T}_m$, such an inflation is denoted as $\sigma[\mathcal{T}_1, \dots, \mathcal{T}_m]$. While from the counting point of view, $\sigma[\mathcal{T}_1, \dots, \mathcal{T}_m]$ is isomorphic to a Cartesian product $\mathcal{T}_1 \times \dots \times \mathcal{T}_m$, it is useful to explicitly keep the external permutation σ for sampling and construction purposes.

Interestingly, Albert and Atkinson describe a specification for any substitution-closed class of permutations [AA05].

Proposition 34 ([AA05, Lemma 11]). *Suppose that a class \mathcal{C} is substitution-closed and contains 12 and 21. Let \mathcal{S} be the class of all simple permutations contained in \mathcal{C} . Then, \mathcal{C} satisfies the following system of equations*




Figure 13. Visualisation of a random permutation 76512834.

$$\begin{aligned}
 \mathcal{C} &= \{\bullet\} + 12[\mathcal{C}^+, \mathcal{C}] + 21[\mathcal{C}^-, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}] \\
 \mathcal{C}^+ &= \{\bullet\} + 21[\mathcal{C}^-, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}] \\
 \mathcal{C}^- &= \{\bullet\} + 12[\mathcal{C}^+, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}].
 \end{aligned} \tag{136}$$

It is, therefore, possible to endow each of the simple permutations $\pi \in \mathcal{S}$ by a distinguished marking variable u_π and insert them into the specification:

$$\begin{aligned}
 \mathcal{C} &= \{\bullet\} + 12[\mathcal{C}^+, \mathcal{C}] + 21[\mathcal{C}^-, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} u_\pi \cdot \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}] \\
 \mathcal{C}^+ &= \{\bullet\} + 21[\mathcal{C}^-, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} u_\pi \cdot \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}] \\
 \mathcal{C}^- &= \{\bullet\} + 12[\mathcal{C}^+, \mathcal{C}] + \sum_{\pi \in \mathcal{S}} u_\pi \cdot \pi[\mathcal{C}, \mathcal{C}, \dots, \mathcal{C}].
 \end{aligned} \tag{137}$$

Finally, note that by tuning the expectations attached to u_π we alter the expected frequencies of inflations used during the construction of a permutation.

Acknowledgements

We are grateful to Anne Bouillard, Matthieu Dien, Michael Grant, Hsien-Kuei Hwang, Élie de Panafieu, Martin Pépin, Bruno Salvy for their valuable feedback and fruitful discussions. We are especially grateful to Cedric Chauve, Yann Ponty, and Sebastian Will for showing us numerous applications of Boltzmann sampling in biology, and to Sergey Tarasov for asking a question about computational complexity of random sampling. We appreciate the anonymous referee’s comments that helped to improve the paper and in particular a comment about the existence of the tuning problem solution which resulted in Remark 22.

References

- [AA05] Albert, M. H. and Atkinson, M. D. (2005) Simple permutations and pattern restricted permutations. *Disc. Math.* 300(1–3) 1–15.
- [Art+15] Arts, T., Hughes, J., Norell, U. and Svensson, H. (2015) Testing AUTOSAR software with QuickCheck. In *2015 IEEE Eighth International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*, pp. 1–4.
- [Ban+12] Banderier, C., Bodini, O., Ponty, Y. and Bouzid, H.T. (2012) On the diversity of pattern distributions in rational language. In *Proceedings of the Ninth Workshop on Analytic Alg. and Combinatorics*, pp. 107–115.
- [Bar16] Barabási, A.-L. (2016) *Network Science*. Cambridge University Press.
- [Bas+17] Bassino, F., Bouvel, M., Pierrot, A., Pivoteau, C. and Rossin, D. (2017) An algorithm computing combinatorial specifications of permutation classes. *Disc. Appl. Math.* 224 16–44.
- [BBD18] Bendkowski, M., Bodini, O. and Dovgal, S. (2018) Polynomial tuning of multiparametric combinatorial samplers. In *2018 Proceedings of the Fifteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pp. 92–106. eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9781611975062.9>.
- [BBJ13] Bacher, A., Bodini, O. and Jacquot, A. (2013) Exact-size sampling for Motzkin trees in linear time via Boltzmann samplers and holonomic specification. In *Proceedings of the Meeting on Analytic Algorithmics and Combinatorics*, pp. 52–61.
- [BBR14] Bouillard, A., Bušić, A. and Rovetta, C. (2014) Perfect sampling for closed queueing networks. *Perform. Eval.* 79 146–159.
- [BDM16] Bodini, O., David, J. and Marchal, P. (2016) Random-bit optimal uniform sampling for rooted planar trees with given sequence of degrees and applications. In *Conference on Algorithms and Discrete Applied Mathematics*, pp. 97–114.
- [Ben+17] Bendkowski, M., Grygiel, K., Lescanne, P. and Zaionc, M. (2017) Combinatorics of λ -terms: a natural approach. *J. Logic Comput.*
- [BFP10] Bodini, O., Fusy, É. and Pivoteau, C. (2010) Random sampling of plane partitions. *Comb. Prob. Comput.* 19(2) 201–226.
- [BFR18] Bernstein, M., Fahrback, M. and Randall, D. (2018) Analyzing Boltzmann samplers for Bose–Einstein condensates with Dirichlet generating functions. In *2018 Proceedings of the Fifteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pp. 107–117. eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9781611975062.10>.
- [BG12] Bernardi, O. and Giménez, O. (2012). A linear algorithm for the random sampling from regular languages. *Algorithmica* 62(1) 130–145.
- [BGR15] Bodini, O., Genitrini, A. and Rolin, N. (2015) Pointed versus singular Boltzmann samplers: a comparative analysis. *Pure Math. Appl.* 25(2) 115–131.
- [BLL98] Bergeron, F., Labelle, G. and Leroux, P. (1998). *Combinatorial Species and Tree-Like Structures*. Vol. 67. Cambridge University Press.
- [BLR15] Bodini, O., Lumbroso, J. and Rolin, N. (2015) Analytic samplers and the combinatorial rejection method. In *Proceedings of the Meeting on Analytic Algorithmics and Combinatorics*, pp. 40–50.
- [Bod+11] Bodirsky, M., Fusy, É., Kang, M. and Vigerske, S. (2011) Boltzmann samplers, Pólya theory, and cycle pointing. *SIAM J. Comp.* 40(3) 721–769.
- [Bod+16] Bodini, O., Dien, M., Fontaine, X., Genitrini, A. and Hwang, H.-K. (2016) Increasing diamonds. In *Latin American Symposium on Theoretical Informatics*, pp. 207–219.
- [Bod10] Bodini, O. (2010) Autour de la génération aléatoire sous modèle de Boltzmann [On random generation under Boltzmann models]. Habilitation. Université Pierre et Marie Curie.
- [BP10] Bodini, O. and Ponty, Y. (2010) Multi-dimensional Boltzmann sampling of context-free languages. In *21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA 2010)*. Vol. AM.
- [BR83] Bender, E. A. and Bruce Richmond, L. (1983) Central and local limit theorems applied to asymptotic enumeration II: Multivariate generating functions. *J. Comb. Theory, A* 34(3) 255–265.
- [BRS12] Bodini, O., Roussel, O. and Soria, M. (2012) Boltzmann samplers for first-order differential specifications. *Disc. App. Math.* 160(18) 2563–2572.
- [Bru72] de Bruijn, N. G. (1972) Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* 75(5) 381–392.
- [CH00] Claessen, K. and Hughes, J. (2000) QuickCheck: A lightweight tool for random testing of Haskell programs. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming*, pp. 268–279.
- [Cla+00] Clarke, E., Grumberg, O., Jha, S., Lu, Y. and Veith, H. (2000) Counterexample-guided abstraction refinement. *Computer Aided Verification* (E. A. Emerson and A. P. Sistla, eds), Springer Berlin Heidelberg, pp. 154–169. ISBN: 978-3-540-45047-4.
- [CMZ99] Chase, K. C., Mekjian, A. Z. and Zamick, L. (1999) ‘Canonical and microcanonical ensemble approaches to Bose-Einstein condensation: The thermodynamics of particles in harmonic traps. *Euro. Phys. J. B-Condens. Matt. Comp. Syst.* 8(2) 281–285.

- [CPW20] Chauve, C., Ponty, Y. and Wallner, M. (2020) Counting and sampling gene family evolutionary histories in the duplication-loss and duplication-loss-transfer models. *J. Math. Biol.*, 1–36.
- [DCB13] Domahidi, A., Chu, E. and Boyd, S. (2013) ECOS: An SOCP solver for embedded systems. In *2013 European Control Conference (ECC)*. IEEE, pp. 3071–3076.
- [Drm97] Drmota, M. (1997) Systems of functional equations. *Rand. Struct. & Alg.* **10**(1–2) 103–124.
- [DRT00] Denise, A., Roques, O. and Termier, M. (2000) Random generation of words of contextfree languages according to the frequencies of letters. In *Mathematics and Computer Science*. Springer, pp. 113–125.
- [Duc+04] Duchon, P., Flajolet, P., Louchard, G. and Schaeffer, G. (2004) Boltzmann samplers for the random generation of combinatorial structures. *Comb. Prob. Comput.* **13**(4–5) 577–625.
- [DZ99] Denise, A. and Zimmermann, P. (1999) Uniform random generation of decomposable structures using floating-point arithmetic. *Theor. Comput. Sci.* **218**(2) 233–248.
- [FFP07] Flajolet, P., Fusy, É. and Pivoteau, C. (2007) Boltzmann sampling of unlabelled structures. In *Proceedings of the Meeting on Analytic Algorithmics and Combinatorics*, pp. 201–211.
- [FS09] Flajolet, P. and Sedgewick, R. (2009) *Analytic Combinatorics*. Cambridge University Press. ISBN: 978-0-521-89806-5.
- [Fus05] Fusy, É. (2005) Quadratic exact size and linear approximate size random generation of planar graphs. *Discr. Math. & Theor. Comp. Sci.*, 125–138.
- [FVS17] Fichtner, M., Voigt, K. and Schuster, S. (2017) The tip and hidden part of the iceberg: Proteinogenic and non-proteinogenic aliphatic amino acids. *Biochimica et Biophysica Acta (BBA)-Gen. Sub.* **1861**(1) 3258–3269.
- [FZC94] Flajolet, P., Zimmermann, P. and Van Cutsem, B. (1994) A calculus for the random generation of labelled combinatorial structures. *Theor. Comput. Sci.* **132**(1) 1–35.
- [GBY06] Grant, M., Boyd, S. and Ye, Y. (2006) Disciplined convex programming. In *Global Optimization*, pp. 155–210.
- [GG16] Gittenberger, B. and GoŁebiewski, Z. (2016) On the number of lambda terms with prescribed size of their de Bruijn representation. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS*, pp. 40:1–40:13.
- [Ham+19] Hammer, S., Wang, W., Will, S. and Ponty, Y. (2019) Fixed-parameter tractable sampling for RNA design with multiple target structures. *BMC Bioinform.* **20**(1) 209.
- [HHA97] Haugset, T., Haugerud, H. and Andersen, J. O. (1997) Bose-Einstein condensation in anisotropic harmonic traps. *Phys. Rev. A* **55**(4) 2922.
- [JVV86] Jerrum, M. R., Valiant, L. G. and Vazirani, V. V. (1986) Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.* **43** 169–188.
- [Kar84] Karmarkar, N. A new polynomial-time algorithm for linear programming. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pp. 302–311.
- [KK89] Kamada, T. and Kawai, S. (1989) An algorithm for drawing general undirected graphs. *Inform. Process. Lett.* **31**(1) 7–15.
- [KY76] Knuth, D. and Yao, A. (1976) *Algorithms and Complexity: New Directions and Recent Results*. Academic Press. Chap. The complexity of nonuniform random number generation.
- [Lai17] Lairez, P. (2017) A deterministic algorithm to compute approximate roots of polynomial systems in polynomial average time. *Found. Comput. Math.* **17**(5) 1265–1292.
- [LB14] Landau, D. P. and Binder, K. (2014). *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge University Press.
- [LR08] Lucietti, J. and Rangamani, M. (2008) Asymptotic counting of BPS operators in superconformal field theories. *J. Math. Phys.* **49**(8) 082301.
- [MCA18] Mullier, O., Chapoutot, A. and dit Sandretto, J. A. (2018) Validated computation of the local truncation error of Runge–Kutta methods with automatic differentiation. *Optim. Methods Softw.* **33**(4–6) 718–728.
- [Nem04] Nemirovski, A. (2004) *Interior Point Polynomial Time Methods in Convex Programming*. Lecture Notes.
- [Nes98] Nesterov, Y. (1998) Introductory lectures on convex programming volume I: Basic course. *Lecture Notes* **3**(4) 5.
- [NN94] Nesterov, Y. and Nemirovskii, A. (1994) *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM.
- [NW78] Nijenhuis, A. and Wilf, H. S. (1978). *Combinatorial Algorithms*. 2nd ed. Academic Press.
- [ODo+16] O’Donoghue, B., Chu, E., Parikh, N. and Boyd, S. (2016) Conic optimization via operator splitting and homogeneous self-dual embedding. *J. Opt. Th. & App.* **169**(3) 1042–1068.
- [Ott48] Otter, R. (1948) The number of trees. *Ann. Math.*, 583–599.
- [Pa+11] PaŁka, M. H., Claessen, K., Russo, A. and Hughes, J. (2011) Testing an optimising compiler by generating random lambda terms. In *Proceedings of the 6th International Workshop on Automation of Software Test, AST 2011*. ACM, pp. 91–97. ISBN: 978-1-4503-0592-1.
- [PaŁ+11] PaŁka, M. H. (2012) Random structured test data generation for black-box testing. PhD thesis. Chalmers University of Technology.
- [Par+15] Paraskevopoulou, Z., Hrttcu, C., Dénès, M., Lampropoulos, L. and Pierce, B. C. (2015) Foundational property-based testing. In *Interactive Theorem Proving*, C. Urban and X. Zhang, eds. Springer International Publishing, pp. 325–343. ISBN: 978-3-319-22102-1.
- [Pen82] Penny, D. (1982) Testing the theory of evolution by comparing phylogenetic trees constructed from five different protein sequences. *Nature* **297** 197–200.
- [PSS12] Pivoteau, C., Salvy, B. and Soria, M. (2012) Algorithms for combinatorial structures: Well-founded systems and Newton iterations. *J. Comb. Th. A* **119**(8) 1711–1773.

- [PW00] Potra, F. A. and Wright, S. J. (2000) Interior-point methods. *J. Comput. Appl. Math.* 124(1–2) 281–302.
- [PW96] Propp, J. G. and Wilson, D. B. (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algor.* 9(1–2) 223–252.
- [Rà85] Rémy, J.-L. (1985) Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire. *ITA* 19(2) 179–195.
- [RNL08] Runciman, C., Naylor, M. and Lindblad, F. (2008) Smallcheck and lazy smallcheck: Automatic exhaustive testing for small values. In *Proceedings of the First ACM SIGPLAN Symposium on Haskell, Haskell '08*. ACM, pp. 37–48. ISBN: 978-1-60558-064-7.
- [She+16] Shen, X., Diamond, S., Gu, Y. and Boyd, S. (2016) Disciplined convex-concave programming. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 1009–1014.
- [Sma98] Smale, S. (1998) Mathematical problems for the next century. *Math. Intell.* 20(2) 7–15.
- [Ver96] Vershik, A. M. (1996) Statistical mechanics of combinatorial partitions, and their limit shapes. *Funct. Anal. Appl.* 30(2) 90–105.
- [WG01] Welsh, D. and Gale, A. (2001) The complexity of counting problems. *Asp. Compl. de Gruyter Ser. Log. Appl.*, pp. 115–154.
- [Zen05] Zeng, Z. (2005) Computing multiple roots of inexact polynomials. *Math. Comput.* 74(250) 869–903.