Practical volume estimation by a new annealing schedule for cooling convex bodies

Apostolos Chalkis[1], Ioannis Z. Emiris[2], and Vissarion Fisikopoulos[3]

1 **Department of Informatics & Telecommunications**
National & Kapodistrian University of Athens, Greece
achalkis@di.uoa.gr
2 **Department of Informatics & Telecommunications**
National & Kapodistrian University of Athens, Greece, and
ATHENA Research & Innovation Center, Greece
emiris@di.uoa.gr
3 **Department of Informatics & Telecommunications**
National & Kapodistrian University of Athens, Greece
vfisikop@di.uoa.gr

Abstract

We experimentally study the problem of estimating the volume of convex bodies, focusing on H- and V-polytopes, as well as zonotopes. Although a lot of effort is devoted to practical algorithms for H-polytopes there is no such method for the latter two representations. We propose a new, practical method for all representations which also improves upon the performance of existing methods on H-polytopes.

1 Introduction

Volume computation of a convex body in general dimension is a fundamental problem in discrete geometry. In the past 28 years randomized algorithms, for this problem, have made great progress. The two existing [5], [2] practical methods and the corresponding implementations are based on theoretical results, but they make some practical adjustments and show experimentally that they estimate volumes with small errors and high probability. Our new practical method can be used for general convex bodies but in this paper we focus on convex polytopes. A convex polytope P can be given as (a) an intersection of q halfspaces (H-polytope), (b) a convex hull of a set of points (V-polytope) and (c) a Minkowski sum of k segments (zonotope). We assume that an H-polytope is given by a matrix $A \in \mathbb{R}^{q\times d}$ and a vector $b \in \mathbb{R}^q$, s.t. $P = \{x \mid Ax \leq b\}$ and a zonotope by a matrix $G \in \mathbb{R}^{d\times k}$ which contains the k segments column-wise.

Exact volume computation is #P-hard for H- and V-polytopes, including zonotopes [6]. There are several implementations in packages such as `VINCI` or `qHull` but, as expected, they do not scale beyond, say, $d \geq 15$ dimensions. The first approximation algorithm, is given in [4] with complexity $O^*(d^{23})$.

The main approach relies on a Multiphase Monte Carlo (MMC) sequence of convex bodies $P_0 \subseteq \cdots \subseteq P_m = P$ such that rejection sampling would efficiently estimate $\mathrm{vol}(P_i)/\mathrm{vol}(P_{i-1})$, i.e. sample uniform points from P_i and reject/accept points in P_{i-1}. Assuming P is well-rounded and also that the unit ball B_d is the largest inscribed ball in P, defining P_1. Then each convex body P_i is defined by the intersection of a scaled copy of B_d with P while the largest ball, which defines $P_m = P$, is an enclosing ball of P. Then we estimate $\mathrm{vol}(P)$ through the telescopic product (2). The critical complexity issue is to minimize the length of the sequence in MMC, called m, while each ratio remains large enough to use rejection sampling. In [7] the sequence of balls in MMC is defined deterministically for each instance,

35th European Workshop on Computational Geometry, Utrecht, The Netherlands, March 18–20, 2019.
This is an extended abstract of a presentation given at EuroCG'19. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

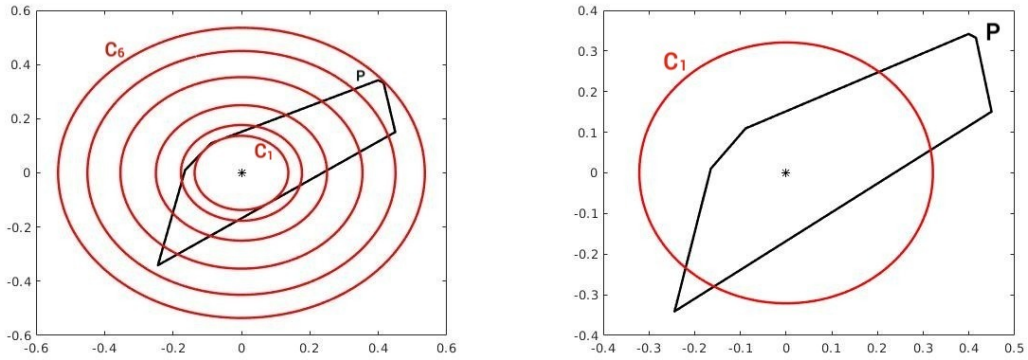

Figure 1 Balls in MMC in [7] left and from the annealing schedule right (r=0.25, δ=0.05)

while $m = O(d \lg d)$. Our method uses a new annealing schedule to define a sparser sequence of balls, (Fig. 1), and a new practical convergence criterion for each ratio in the MMC in order to minimize the number of sampled points to estimate each ratio. Moreover for zonotopes of order ≤ 4 we do not use balls in MMC but we define a centrally symmetric convex polytope which reduces the number of bodies (or phases) significantly.

Current state-of-the-art software is based on the above paradigms and, for H-polytopes, typically uses Hit-and-Run (HnR). `VolEsti` [5], which scales up to hundreds of dimensions, uses Coordinate-Direction HnR. We shall also juxtapose the software of [2] (for H-polytopes), which implements [1] with an annealing schedule.

These implementations can not handle efficiently zonotopes or V-polytopes as they require an inscribed ball (ideally the largest one). Additionally the software of [2] requires the number of facets which is typically exponential in the dimension for both zonotopes and V-polytopes. Our software outperforms for $d \leq 100$ software for H-polytopes by [2] and [5]. Moreover we provide the first practical method for V-polytopes and zonotopes that scales to high dimensions (currently 100 for V-polytopes and low-order zonotopes).

We introduce some notions from statistics and refer to [3] for details. Given a random sample of size ν from a random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ with unknown variance σ^2, the (one tailed) t-test checks the null hypothesis that the population mean exceeds a specified value μ_0 using the statistic $t = \frac{\bar{x}-\mu_0}{s/\sqrt{\nu}} \sim t_{\nu-1}$, where $\bar{x}$ is the sample mean, s the sample standard deviation and $t_{\nu-1}$ is the t-student distribution with $\nu - 1$ degrees of freedom. Given a significance level $\alpha > 0$ we test the null hypothesis for the mean value of the population, $H_0 : \mu \leq \mu_0$. We reject H_0 if,

$$t \geq t_{\nu-1,\alpha} \Rightarrow \bar{x} \geq \mu_0 + t_{\nu-1,\alpha} s/\sqrt{n}, \tag{1}$$

where $t_{\nu-1,\alpha}$ is the critical value of the t-student distribution. Inequality 1 implies $\Pr(H_0 \text{ true} \mid \text{reject } H_0) = \alpha$. Otherwise we fail to reject H_0.

2 Volume algorithm

Our method introduces some new algorithmic features. The MMC of the algorithm constructs a sequence of convex bodies $C_1 \supseteq \cdots \supseteq C_m$ intersecting the given polytope P; we introduce a new annealing schedule in order to minimize m. A typical choice for the C_i's is a sequence of co-centric balls but any set of convex bodies can be used in our method. We re-write the

telescopic product in [7] as follows:

$$\mathrm{vol}(P) = \frac{\frac{\mathrm{vol}(P_m)}{\mathrm{vol}(C_m)}}{\frac{\mathrm{vol}(P_1)}{\mathrm{vol}(P_0)}\frac{\mathrm{vol}(P_2)}{\mathrm{vol}(P_1)}\cdots\frac{\mathrm{vol}(P_m)}{\mathrm{vol}(P_{m-1})}}\mathrm{vol}(C_m), \quad \text{where } P_0 = P, P_i = C_i \cap P. \tag{2}$$

The behavior of our method is parametrized by: the error of approximation ϵ, cooling parameters $0 < r < 1$, $\delta > 0$, s.t. $0 < r + \delta < 1$ which are used in the schedule, significance level $\alpha > 0$ of the statistical tests, ν the degrees of freedom for the t-student used in t-tests, parameter N that controls the number of points νN generated in P_i, and n the length of the sliding window. From the telescopic product (2) it is clear that in practical estimations C_m has to be a convex body whose volume is computed much faster than $\mathrm{vol}(P)$ (ideally by a closed formula) and which can be sampled efficiently.

The annealing schedule specifies $C_1 \supseteq \cdots \supseteq C_m$ using the following two statistical tests:

testL$(P_1, P_2, r, \delta, \alpha, \nu, N)$:	**testR**$(P_1, P_2, r, \alpha, \nu, N)$:
$H_0:\ \mathrm{vol}(P_2)/\mathrm{vol}(P_1) \geq r + \delta$	$H_0:\ \mathrm{vol}(P_2)/\mathrm{vol}(P_1) \leq r$
Successful if H_0 is **rejected**	**Successful** if H_0 is **rejected**

These tests are being used by annealing schedule to restrict each ratio $r_i = \mathrm{vol}(P_{i+1})/\mathrm{vol}(P_i)$ in the interval $[r, r + \delta]$ with high probability in order to avoid unnecessarily big ratios in MMC. Then we can use rejection sampling to estimate efficiently each ratio. Given P_i, `testL` is used to define $P_{i+1} \subseteq P_i$ s.t. ratio $\mathrm{vol}(P_{i+1})/\mathrm{vol}(P_i)$ is not too large, while `testR` is used so that the ratio is not too small, with high probability. In general, if we sample N uniform points from a body P_i then random variable X that counts points in P_{i+1}, follows $X \sim b(N, r_i)$, the binomial distribution, and random variable $Y = X/N \sim \mathcal{N}(r_i, r_i(1-r_i)/N)$ is Gaussian. If we sample νN points from P_i and split the sample into ν sublists of length N, the corresponding ν ratios are experimental values that follow $\mathcal{N}(r_i, r_i(1 - r_i)/N)$ and can be used to check both null hypotheses for r_i in `testL` and `testR`. Using the mean μ_0 of the ν ratios, r_i is restricted to $[r, r + \delta]$ with high probability when the following holds:

$$r + \delta - t_{\nu-1,\alpha}\frac{s}{\sqrt{\nu}} > \mu_0 > r + t_{\nu-1,\alpha}\frac{s}{\sqrt{\nu}}, \tag{3}$$

Initialization of the annealing schedule is to compute the body C' s.t. the volume of $C' \cap P$ could be efficiently estimated using rejection sampling, i.e. sampling from C' and accepting points in $C' \cap P$. Body C' is also used for the stopping criterion: the annealing schedule stops in the i-th step if `testR`$(P_i, C' \cap P)$ succeeds, which means that the $\mathrm{vol}(P_i)$ is close enough to $\mathrm{vol}(C' \cap P)$, so that rejection sampling can be used. Then set $m = i + 1$ and $C_m = C'$, $P_m = C_m \cap P$. When balls are used in the MMC, the smallest ball C_m is not an inscribed ball and the largest one, C_1, is not an enclosing ball as in [7]. Hence in practice the number of phases in [5] is an upper bound, with high probability, on the number of phases of our method, when $0 < r + \delta < 1/2$. Fig. 1 shows the sequence of balls for a given polytope P with our method (m=1) and in [7] (m=6). The ratios our method estimates are: $\mathrm{vol}(P_1)/\mathrm{vol}(P_0)$ and $\mathrm{vol}(P_1)/\mathrm{vol}(C_1)$, where $P_0 = P$, $P_1 = C_1 \cap P$.

The annealing schedule returns m bodies and we estimate $m + 1$ volume ratios. For fixed step i and each sample point generated in P_i, we keep the value of the i-th ratio. We store the last n such values in a queue called sliding window denoted by W whose length is n. We update W for each new point by inserting the new ratio and by popping out the oldest ratio in W. For each ratio r_i, we bound error by ϵ_i s.t. $\sum_{i=0}^{m} \epsilon_i^2 = \epsilon^2$ then, from standard error propagation analysis, (2) estimates $\mathrm{vol}(P)$ with error at most ϵ.

At step i, let $\hat{\mu}$ be the mean, s the st.d. of W and $Pr = 3/4$. Using $p = (1 + Pr)/2$, where $z_p = \sqrt{2} \cdot \text{erf}^{-1}(2p - 1)$, we consider the interval $[\hat{\mu} - z_p s, \hat{\mu} + z_p s]$, where erf is the Gauss error function. The values that the sliding window contains are not independent, so we can not define a confidence interval using t-student distribution, but we experimentally show that the following practical criterion is very efficient:

$$\text{if } \frac{2z_p s}{\hat{\mu} - z_p s} \leq \epsilon_i/2, \text{ then declare convergence.} \tag{4}$$

In practice we set $n = O(d^2)$ so that our method estimates volumes with error $\leq \epsilon$ with high probability.

We use Hit-and-Run (HnR) with uniform target distribution for sampling from P_i at step i of the annealing schedule. One step of HnR is described below. For a value t we return a point after t iterations.

Hit-and-Run(P, p): Convex polytope P, current point $p \in P$
Pick a uniformly distributed line ℓ through p
Return a uniform point on the chord $\ell \cap P$

For zonotopes each step in both Coordinate-Directions HnR and Random-Directions HnR solves the following LP to compute one extreme point on $\ell \cap P$: minimize α, $s.t. \, p + \alpha v = \sum_{i=1}^{k} \lambda_i g_i$, $-1 \leq \lambda_i \leq 1$. For the second extreme point, keep the same constraints and minimize $-\alpha$. This LP uses the basic feasible solution of the first one.

Moreover, for zonotopes we study different types of convex bodies than ball for the MMC sequence. $G^T G$ has $k - d$ zero eigenvalues; the corresponding eigenvectors form matrix $Q \in \mathbb{R}^{k \times (k-d)}$. The intersection of the hypercube $[-1, 1]^k$ with the d-dimensional affine subspace defined by $Q^T = 0$ equals a d-dimensional polytope C in $\mathbb{R}^k$. SVD yields an orthonormal basis for the linear constraints, and its orthogonal complement $W_\perp$:

$$Q = USV^T = \begin{bmatrix} W \\ W_\perp \end{bmatrix}^T \begin{bmatrix} S_1 & 0 \\ 0 & 0 \end{bmatrix} V^T.$$

Let $Ay \leq b$, $A \in \mathbb{R}^{2k \times k}$ be an H-representation of $[-1, 1]^k$, then $Mx \leq b$, $M = AW_\perp^T (GW_\perp^T)^{-1} \in \mathbb{R}^{d \times d}$ is an H-representation of the full-dimensional, centrally symmetric polytope $C \subseteq P$ with $\leq 2k$ facets to be used in MMC. Each C_i arises from parallel shifting of the facets of C. This C improves the schedule when order is low, i.e. ≤ 4.

3 Implementation and experiments

We perform extended experiments analyzing various aspects of our method such as practical complexity and how is affected by the bodies used in MMC and we compare our implementation with the `matlab` code of [2] and `C++` package `VolEsti` [5]. Our `C++` software is open source[1]. When we use balls in MMC we call our implementation `CoolingBall` and when we use the H-polytope for zonotopes we call it `CoolingHpoly`. We call the implementation of [2] `CoolingGaussian` and that of [5] `SeqOfBalls`.

Set $r = 0.1$ and $\delta = 0.05$ in order for the next convex body in MMC to have about 10% of the volume of the previous one; let s.l. be $\alpha = 0.10$. We set the number of points sampled in P_i per step to be $\nu N = 1200 + 2d^2$ and $\nu = 10$. Set the length of the sliding window $n = 2d^2 + 250$ and the step of HnR $t = 1$.

[1] `https://github.com/TolisChal/volume_approximation/tree/v_poly`

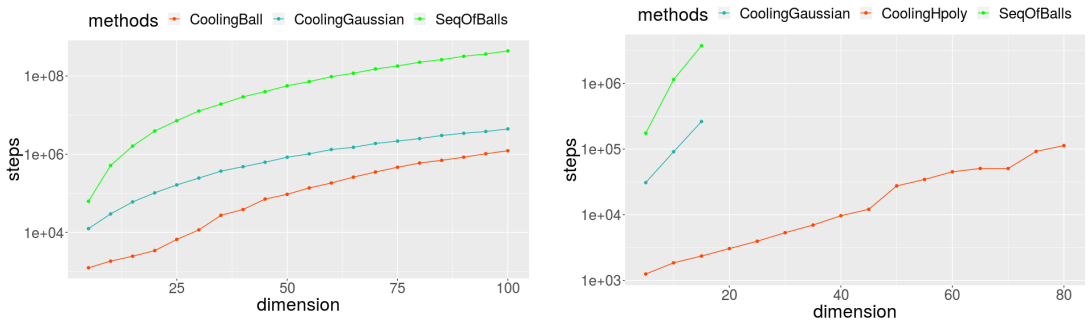

Figure 2 Left: the number of steps for for unit cubes in H-representation, $d = 5, 10, \ldots, 100$. Right: the number of steps for random zonotopes of order 2, $d = 5, 10, \ldots, 80$. In both plots we use log_{10} scale for the y-axis

To study the practical complexity of our method we experimentally correlate the total number of HnR steps with the dimension. In the right plot in Fig. 2 we compare the number of steps for random zonotopes. `CoolingGaussian` fails to estimate volumes for $d > 15$ as the upper bound for the number of facets is the bottleneck for this implementation while `SeqOfBalls` takes $> 1hr$ for $d > 15$. In the left plot we notice that our method is faster than both `CoolingGaussian` and `SeqOfBalls` for $d \leq 100$. In Table 1 we estimate the volumes of random zonotopes. The number of phases for high-order zonotopes is $m = 1$ as our methods defines just an enclosing ball and applies rejection sampling, whereas for low-order zonotopes the H-polytope we defined reduces significantly the number of phases and run-time. The maximum number of phases for zonotopes (up to what our software computes in $<$ 10hr) is $m \leq 3$. None of the volumes of the zonotopes in Table 1 can be computed using exact computations in practice. To define a random zonotope z-d-k we choose a random direction for each segment $s \in S$, where $\sum_{s \in S} s$, and pick a random length in the interval $[0, \sqrt{d}]$.

z-d-k	*Body*	*order*	Vol	m	e	*steps*	*time*
z-10-1000	Ball	100	2.62e+29	1	0.1	0.1400e+04	130.1
z-15-1500	Ball	100	5.00e+45	1	0.1	0.1650e+04	506.1
z-20-2000	Ball	100	2.79e+62	1	0.1	0.2000e+04	1428
z-60-90	Hpoly	1.5	5.81e+82	2	0.1	5.355e+04	943.9
z-80-120	Hpoly	1.5	8.48e+114	3	0.1	12.35e+04	4180
z-100-150	Hpoly	1.5	2.32+149	3	0.1	15.43e+04	10060
z-80-160	Hpoly	2	2.01e+131	3	0.2	11.31e+04	5356
z-100-200	Hpoly	2	5.27e+167	3	0.2	15.25e+04	34110

Table 1 *Body* the type of body in MMC; *order* is k/d, Vol the estimated volume; m the number of phases in MMC; ϵ the requested error; *time* the time in seconds; e the input value for error.

Acknowledgements. AC and IE acknowledge support by the EU Horizon 2020 research and innovation programme under grant agreement No 734242 (Project LAMBDA). They are both members of team AROMATH joint between NKUA and INRIA Sophia-Antipolis (France).

References

1 B. Cousins and S. Vempala. Bypassing KLS: Gaussian cooling and an $o^*(n^3)$ volume algorithm. In *Proc. ACM STOC*, pages 539–548, 2015.

2 B. Cousins and S. Vempala. A practical volume algorithm. *Mathematical Programming Computation*, 8, June 2016.

3 H. Cramer. *Mathematical methods of statistics*. Princeton University Press, 1946.

4 M. Dyer, A. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. URL: `http://doi.acm.org/10.1145/102782.102783`.

5 I.Z. Emiris and V. Fisikopoulos. Practical polytope volume approximation. *ACM Trans. Math. Soft.*, 44(4):38:1–38:21, 2018. Prelim. version: Proc. Symp. Comp. Geometry, 2014. URL: `http://doi.acm.org/10.1145/3194656`, `doi:10.1145/3194656`.

6 E. Gover and N. Krikorian. Determinants and the volumes of parallelotopes and zonotopes. *Linear Algebra and its Applications*, 413:28–40, 2010.

7 L. Lovász, R. Kannan, and M. Simonovits. Random walks and an $o^*(n^5)$ volume algorithm for convex bodies. *Random Structures and Algorithms*, 11:1–50, 1997.