

Efficient Learning of Simplices

Joseph Anderson

*Computer Science and Engineering
Ohio State University*

ANDEJOSE@CSE.OHIO-STATE.EDU

Navin Goyal

Microsoft Research India

NAVINGO@MICROSOFT.COM

Luis Rademacher

*Computer Science and Engineering
Ohio State University*

LRADEMAC@CSE.OHIO-STATE.EDU

Abstract

We show an efficient algorithm for the following problem: Given uniformly random points from an arbitrary n -dimensional simplex, estimate the simplex. The size of the sample and the number of arithmetic operations of our algorithm are polynomial in n . This answers a question of Frieze, Jerrum and Kannan [Frieze et al. \(1996\)](#). Our result can also be interpreted as efficiently learning the intersection of $n + 1$ half-spaces in \mathbb{R}^n in the model where the intersection is bounded and we are given polynomially many uniform samples from it. Our proof uses the local search technique from Independent Component Analysis (ICA), also used by [Frieze et al. \(1996\)](#). Unlike these previous algorithms, which were based on analyzing the fourth moment, ours is based on the third moment.

We also show a direct connection between the problem of learning a simplex and ICA: a simple randomized reduction to ICA from the problem of learning a simplex. The connection is based on a known representation of the uniform measure on a simplex. Similar representations lead to a reduction from the problem of learning an affine transformation of an n -dimensional ℓ_p ball to ICA.

Keywords: independent component analysis, randomized reductions, learning convex bodies, method of moments

1. Introduction

We are given uniformly random samples from an unknown convex body in \mathbb{R}^n , how many samples are needed to approximately reconstruct the body? It seems intuitively clear, at least for $n = 2, 3$, that if we are given sufficiently many such samples then we can reconstruct (or learn) the body with very little error. For general n , it is known to require $2^{\Omega(\sqrt{n})}$ samples [Goyal and Rademacher \(2009\)](#) (see also [Klivans et al. \(2008\)](#) for a similar lower bound in a different but related model of learning). This is an information-theoretic lower bound and no computational considerations are involved. As mentioned in [Goyal and Rademacher \(2009\)](#), it turns out that if the body has few facets (e.g. polynomial in n), then polynomial in n samples are sufficient for approximate reconstruction. This is an information-theoretic upper bound and no efficient algorithms (i.e., with running time $\text{poly}(n)$) are known. (We remark that to our knowledge the same situation holds for polytopes with $\text{poly}(n)$ vertices.) In this paper we study the reconstruction problem for

the special case when the input bodies are restricted to be (full-dimensional) simplices. We show that in this case one can in fact learn the body efficiently. More precisely, the algorithm knows that the input body is a simplex but only up to an affine transformation, and the problem is to recover this affine transformation. This answers a question of (Frieze et al., 1996, Section 6).

The problem of learning a simplex is also closely related to the well-studied problem of learning intersections of half-spaces. Suppose that the intersection of $n + 1$ half-spaces in \mathbb{R}^n is bounded, and we are given $\text{poly}(n)$ uniformly random samples from it. Then our learning simplices result directly implies that we can learn the $n + 1$ half-spaces. This also has the advantage of being a proper learning algorithm, meaning that the output of the algorithm is a set of $n + 1$ half-spaces, unlike many of the previous algorithms.

Previous work. Perhaps the first approach to learning simplices that comes to mind is to find a minimum volume simplex containing the samples. This can be shown to be a good approximation to the original simplex. (Such minimum volume estimators have been studied in machine learning literature, see e.g. Schölkopf et al. (2001) for the problem of estimating the support of a probability distribution. We are not aware of any technique that applies to our situation and provides theoretical guarantees.) However, the problem of finding a minimum volume simplex is in general NP-hard Packer (2002). This hardness is not directly applicable for our problem because our input is a random sample and not a general point set. Nevertheless, we do not have an algorithm for directly finding a minimum volume simplex; instead we use ideas similar to those used in Independent Component Analysis (ICA). ICA studies the following problem: Given a sample from an affine transformation of a random vector with independently distributed coordinates, recover the affine transformation (up to some unavoidable ambiguities). Frieze et al. (1996) gave an efficient algorithm for this problem (with some restrictions on the allowed distributions, but also with some weaker requirements than full independence) along with most of the details of a rigorous analysis (a complete analysis of a special case can be found in Arora et al. (2012); see also Vempala and Xiao (2011) for a generalization of ICA to subspaces along with a rigorous analysis). The problem of learning parallelepipeds from uniformly random samples is a special case of this problem. Frieze et al. (1996) asked if one could learn other convex bodies, and in particular simplices, efficiently from uniformly random samples. Nguyen and Regev (2009) gave a simpler and rigorous algorithm and analysis for the case of learning parallelepipeds with similarities to the popular FastICA algorithm of Hyvärinen (1999). The algorithm in Nguyen and Regev (2009) is a first order algorithm unlike Frieze et al.’s second order algorithm.

The algorithms in both Frieze et al. (1996); Nguyen and Regev (2009) make use of the fourth moment function of the probability distribution. Briefly, the fourth moment in direction $u \in \mathbb{R}^n$ is $\mathbb{E}(u \cdot X)^4$, where $X \in \mathbb{R}^n$ is the random variable distributed according to the input distribution. The moment function can be estimated from the samples. The independent components of the distribution correspond to local maxima or minima of the moment function, and can be approximately found by finding the local maxima/minima of the moment function estimated from the sample.

More information on ICA including historical remarks can be found in Hyvärinen et al. (2001); Comon and Jutten (2010). Ideas similar to ICA have been used in statistics in

the context of projection pursuit since the mid-seventies. It is not clear how to apply ICA to the simplex learning problem directly as there is no clear independence among the components. Let us note that [Frieze et al. \(1996\)](#) allow certain kinds of dependencies among the components, however this does not appear to be useful for learning simplices.

Learning intersections of half-spaces is a well-studied problem in learning theory. The problem of PAC-learning intersections of even two half-spaces is open, and there is evidence that it is hard at least for sufficiently large number of half-spaces: E.g., [Klivans and Sherstov \(2009\)](#) prove that learning intersections of n^ϵ half-spaces in \mathbb{R}^n (for constant $\epsilon > 0$) is hard under standard cryptographic assumptions (PAC-learning is possible, however, if one also has access to a membership oracle in addition to random samples [Kwek and Pitt \(1998\)](#)). Because of this, much effort has been expended on learning when the distribution of random samples is some simple distribution, see e.g. [Klivans and Sherstov \(2007\)](#); [Vempala \(2010a,b\)](#) and references therein. This line of work makes substantial progress towards the goal of learning intersections of k half-spaces efficiently, however it falls short of being able to do this in time polynomial in k and n ; in particular, these algorithms do not seem to be able to learn simplices. The distribution of samples in these works is either the Gaussian distribution or the uniform distribution over a ball. [Frieze et al. \(1996\)](#) and [Goyal and Rademacher \(2009\)](#) consider the uniform distribution over the intersection. Note that this requires that the intersection be bounded. Note also that one only gets positive samples in this case unlike other work on learning intersections of half-spaces. The problem of learning convex bodies can also be thought of as learning a distribution or density estimation problem for a special class of distributions.

[Gravin et al. \(2012\)](#) show how to reconstruct a polytope with N vertices in \mathbb{R}^n , given its first $O(nN)$ moments in $(n+1)$ random directions. In our setting, where we have access to only a polynomial number of random samples, it's not clear how to compute moments of such high orders to the accuracy required for the algorithm of [Gravin et al. \(2012\)](#) even for simplices.

As was pointed out by a COLT referee, efficient learnability of a simplex follows from the recent work on learning latent Dirichlet allocations [Anandkumar et al. \(2012b\)](#). That as well as the current paper use low-degree moments although the way the structure is extracted from the moments is different: [Anandkumar et al. \(2012b\)](#) use eigenvalue decompositions. The more recent work of [Anandkumar et al. \(2012a\)](#) uses symmetric orthogonal tensor decompositions to learn various latent variable models and in particular Dirichlet allocations. [Anandkumar et al. \(2012a\)](#) use a fixed-point iteration similar to the one we use in our first method for learning simplices. [Anandkumar et al. \(2012a\)](#) and our work are independent.

Our results For clarity of the presentation, we use the following machine model for the running time: a random access machine that allows the following exact arithmetic operations over real numbers in constant time: addition, subtraction, multiplication, division and square root.

The estimation error is measured using total variation distance, denoted d_{TV} (see [Section 2](#)).

Theorem 1 *There is an algorithm ([Algorithm 1](#) below) such that given access to random samples from a simplex $S_{INPUT} \subseteq \mathbb{R}^n$, with probability at least $1 - \delta$ over the sample and*

the randomness of the algorithm, it outputs $n + 1$ vectors that are the vertices of a simplex S so that $d_{TV}(S, S_{INPUT}) \leq \epsilon$. The algorithm runs in time polynomial in n , $1/\epsilon$ and $1/\delta$.

As mentioned earlier, our algorithm uses ideas from ICA. Our algorithm uses the third moment instead of the fourth moment used in certain versions of ICA. The third moment is not useful for learning symmetric bodies such as the cube as it is identically 0. It is however useful for learning a simplex where it provides useful information, and is easier to handle than the fourth moment. One of the main contributions of our work is the understanding of the third moment of a simplex and the structure of local maxima. This is more involved than in previous work as the simplex has no obvious independence structure, and the moment polynomial one gets has no obvious structure unlike for ICA.

The probability of success of the algorithm can be “boosted” so that the dependence of the running time on δ is only linear in $\log(1/\delta)$, see Appendix C.

While our algorithm for learning simplices uses techniques for ICA, we have to do substantial work to make those techniques work for the simplex problem. We also show a more direct connection between the problem of learning a simplex and ICA: a randomized reduction from the problem of learning a simplex to ICA. The connection is based on a known representation of the uniform measure on a simplex as a normalization of a vector having independent coordinates. Similar representations are known for the uniform measure in an n -dimensional ℓ_p ball (denoted ℓ_p^n) Barthe et al. (2005) and the *cone measure* on the boundary of an ℓ_p^n ball Schechtman and Zinn (1990); Rachev and Ruschendorf (1991); Song and Gupta (1997) (see Appendix A for the definition of the cone measure). These representations lead to a reduction from the problem of learning an affine transformation of an ℓ_p^n ball to ICA. These reductions show connections between estimation problems with no obvious independence structure and ICA. They also make possible the use of any off-the-shelf implementation of ICA. However, the results here do not supersede our result for learning simplices because to our knowledge no rigorous analysis is available for the ICA problem when the distributions are the ones in the above reductions.

Idea of the algorithm. The new idea for the algorithm is that after putting the samples in a suitable position (see below), the third moment of the sample can be used to recover the simplex using a simple FastICA-like algorithm. We outline our algorithm next.

As any full-dimensional simplex can be mapped to any other full-dimensional simplex by an invertible affine transformation, it is enough to determine the translation and linear transformation that would take the given simplex to some canonical simplex. As is well-known for ICA-like problems (see, e.g., Frieze et al. (1996)), this transformation can be determined *up to a rotation* from the mean and the covariance matrix of the uniform distribution on the given simplex. The mean and the covariance matrix can be estimated efficiently from a sample. A convenient choice of an n -dimensional simplex is the convex hull of the canonical vectors in \mathbb{R}^{n+1} . We denote this simplex Δ_n and call it the *standard simplex*. So, the algorithm begins by picking an arbitrary invertible affine transformation T that maps \mathbb{R}^n onto the hyperplane $\{x \in \mathbb{R}^{n+1} : \mathbb{1} \cdot x = 1\}$. We use a T so that $T^{-1}(\Delta_n)$ is an isotropic¹ simplex. In this case, the algorithm brings the sample set into isotropic position and embeds it in \mathbb{R}^{n+1} using T . After applying these transformations we may assume (at

1. See Section 2.

the cost of small errors in the final result) that our sample set is obtained by sampling from an unknown rotation of the standard simplex that leaves the all-ones vector (denoted $\mathbb{1}$ from now on) invariant (thus this rotation keeps the center of mass of the standard simplex fixed), and the problem is to recover this rotation.

To find the rotation, the algorithm will find the vertices of the rotated simplex approximately. This can be done efficiently because of the following characterization of the vertices: Project the vertices of the simplex onto the hyperplane through the origin orthogonal to $\mathbb{1}$ and normalize the resulting vectors. Let V denote this set of $n + 1$ points. Consider the problem of maximizing the third moment of the uniform distribution in the simplex along unit vectors orthogonal to $\mathbb{1}$. Then V is the complete set of local maxima and the complete set of global maxima (Theorem 11). A fixed point-like iteration (inspired by the analysis of FastICA Hyvärinen (1999) and of gradient descent in Nguyen and Regev (2009)) starting from a random point in the unit sphere finds a local maximum efficiently with high probability. By the analysis of the coupon collector’s problem, $O(n \log n)$ repetitions are highly likely to find all local maxima.

Idea of the analysis. In the analysis, we first argue that after putting the sample in isotropic position and mapping it through T , it is enough to analyze the algorithm in the case where the sample comes from a simplex S that is close to a simplex S' that is the result of applying a rotation leaving $\mathbb{1}$ invariant to the standard simplex. The closeness here depends on the accuracy of the sample covariance and mean as an estimate of the input simplex’s covariance matrix and mean. A sample of size $O(n)$ guarantees ((Adamczak et al., 2010, Theorem 4.1), (Srivastava and Vershynin, 2011, Corollary 1.2)) that the covariance and mean are close enough so that the uniform distributions on S and S' are close in total variation. We show that the subroutine that finds the vertices (Subroutine 1), succeeds with some probability when given a sample from S' . By definition of total variation distance, Subroutine 1 succeeds with almost as large probability when given a sample from S (an argument already used in Nguyen and Regev (2009)). As an additional simplifying assumption, it is enough to analyze the algorithm (Algorithm 1) in the case where the input is isotropic, as the output distribution of the algorithm is equivariant with respect to affine invertible transformations as a function of the input distribution.

Organization of the paper. We start with some preliminaries in Sec. 2. In Sec. 3 we give an algorithm that estimates individual vertices of simplices in a special position; using this algorithm as a subroutine in Sec. 4 we give the algorithm for the general case. Sec. 5 gives the probabilistic results underlying the reductions from learning simplices and ℓ_p^n balls to ICA. Sec. 6 explains those reductions.

2. Preliminaries

An n -simplex is the convex hull of $n + 1$ points in \mathbb{R}^n that do not lie on an $(n - 1)$ -dimensional affine hyperplane. It will be convenient to work with the standard n -simplex Δ^n living in \mathbb{R}^{n+1} defined as the convex hull of the $n + 1$ canonical unit vectors e_1, \dots, e_{n+1} ; that is

$$\Delta^n = \{(x_0, \dots, x_n) \in \mathbb{R}^{n+1} : x_0 + \dots + x_n = 1 \text{ and } x_i \geq 0 \text{ for all } i\}.$$

The canonical simplex Ω^n living in \mathbb{R}^n is given by

$$\{(x_0, \dots, x_{n-1}) \in \mathbb{R}^n : x_0 + \dots + x_{n-1} \leq 1 \text{ and } x_i \geq 0 \text{ for all } i\}.$$

Note that Δ^n is the facet of Ω^{n+1} opposite to the origin.

The total variation distance between two probability measures is $d_{TV}(\mu, \nu) = \sup_A |\mu(A) - \nu(A)|$ for measurable A . For two compact sets $K, L \subseteq \mathbb{R}^n$, we define the total variation distance $d_{TV}(K, L)$ as the total variation distance between the corresponding uniform distributions on each set. It can be expressed as

$$d_{TV}(K, L) = \begin{cases} \frac{\text{vol } K \setminus L}{\text{vol } K} & \text{if } \text{vol } K \geq \text{vol } L, \\ \frac{\text{vol } L \setminus K}{\text{vol } L} & \text{if } \text{vol } L > \text{vol } K. \end{cases}$$

Define the d -th power sum as

$$p_d(u_0, \dots, u_n) = u_0^d + \dots + u_n^d.$$

For a vector $u = (u_0, u_1, \dots, u_n)$, we define

$$u^{(2)} = (u_0^2, u_1^2, \dots, u_n^2).$$

Vector $\mathbb{1}$ denotes the all ones vector (the dimension of the vector will be clear from the context).

A random vector $X \in \mathbb{R}^n$ is *isotropic* if $\mathbb{E}(X) = 0$ and $\mathbb{E}(XX^T) = I$. A compact set in \mathbb{R}^n is isotropic if a uniformly distributed random vector in it is isotropic. The inradius of an isotropic n -simplex is $\sqrt{(n+2)/n}$, the circumradius is $\sqrt{n(n+2)}$.

For a point $x \in \mathbb{R}^n$, $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$ is the standard ℓ_p norm. The unit ℓ_p^n ball is defined by

$$B_p^n = \{x \in \mathbb{R}^n : \|x\|_p \leq 1\}.$$

The Gamma distribution is denoted as $\text{Gamma}(\alpha, \beta)$ and has density $f(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} 1_{x \geq 0}$, with shape parameters $\alpha, \beta > 0$. $\text{Gamma}(1, \lambda)$ is the exponential distribution, denoted $\text{Exp}(\lambda)$. The Gamma distribution also satisfies the following additivity property: If X is distributed as $\text{Gamma}(\alpha, \beta)$ and Y is distributed as $\text{Gamma}(\alpha', \beta)$, then $X + Y$ is distributed as $\text{Gamma}(\alpha + \alpha', \beta)$.

3. Subroutine for finding the vertices of a rotated standard simplex

In this section we solve the following simpler problem: Suppose we have $\text{poly}(n)$ samples from a rotated copy S of the standard simplex, where the rotation is such that it leaves $\mathbb{1}$ invariant. The problem is to approximately estimate the vertices of the rotated simplex from the samples.

We will analyze our algorithm in the coordinate system in which the input simplex is the standard simplex. This is only for convenience in the analysis and the algorithm itself does not know this coordinate system.

As we noted in the introduction, our algorithm is inspired by the algorithm of [Nguyen and Regev \(2009\)](#) for the related problem of learning hypercubes and also by the FastICA

algorithm in [Hyvärinen \(1999\)](#). New ideas are needed for our algorithm for learning simplices; in particular, our update rule is different. With the right update rule in hand the analysis turns out to be quite similar to the one in [Nguyen and Regev \(2009\)](#).

We want to find local maxima of the sample third moment. A natural approach to do this would be to use gradient descent or Newton’s method (this was done in [Frieze et al. \(1996\)](#)). Our algorithm, which only uses first order information, can be thought of as a fixed point algorithm leading to a particularly simple analysis and fast convergence. Before stating our algorithm we describe the update rule we use.

The k -th moment $m_k(u)$ over Δ^n is the function $u \mapsto \mathbb{E}_{X \in \Delta_n}((u \cdot X)^k)$. In [Appendix F](#) we show:

$$m_3(u) = \frac{(p_1(u))^3 + 3p_1(u)p_2(u) + 2p_3(u)}{(n+1)(n+2)(n+3)}.$$

We will use the abbreviation $C_n = (n+1)(n+2)(n+3)/6$. Then, from the expression for $m_3(u)$ we get

$$\nabla m_3(u) = \frac{1}{6C_n} \left(3p_1(u)^2 \mathbb{1} + 3p_2(u) \mathbb{1} + 6p_1(u)u + 6u^{(2)} \right).$$

Solving for $u^{(2)}$ we get

$$\begin{aligned} u^{(2)} &= C_n \nabla m_3(u) - \frac{1}{2} p_1(u)^2 \mathbb{1} - \frac{1}{2} p_2(u) \mathbb{1} - p_1(u)u \\ &= C_n \nabla m_3(u) - \frac{1}{2} (u \cdot \mathbb{1})^2 \mathbb{1} - \frac{1}{2} (u \cdot u)^2 \mathbb{1} - (u \cdot \mathbb{1})u. \end{aligned} \tag{1}$$

While the above expressions are in the coordinate system where the input simplex is the canonical simplex, the important point is that all terms in the last expression can be computed in any coordinate system that is obtained by a rotation leaving $\mathbb{1}$ invariant. Thus, we can compute $u^{(2)}$ as well independently of what coordinate system we are working in. This immediately gives us the algorithm below. We denote by $\hat{m}_3(u)$ the sample third moment, i.e., $\hat{m}_3(u) = \frac{1}{t} \sum_{i=1}^t (u \cdot r_i)^3$ for t samples. This is a polynomial in u , and the gradient is computed in the obvious way. Moreover, the gradient of the sample moment is clearly an unbiased estimator of the gradient of the moment; a bound on the deviation is given in the analysis ([Lemma 2](#)). For each evaluation of the gradient of the sample moment, we use a fresh sample.

It may seem a bit alarming that the fixed point-like iteration is squaring the coordinates of u , leading to an extremely fast growth (see [Equation 1](#) and [Subroutine 1](#)). But, as in other algorithms having quadratic convergence like certain versions of Newton’s method, the convergence is very fast and the number of iterations is small. We show below that it is $O(\log(n/\delta))$, leading to a growth of u that is polynomial in n and $1/\delta$. The boosting argument described in the introduction makes the final overall dependence in δ to be only linear in $\log(1/\delta)$.

We state the following subroutine for \mathbb{R}^n instead of \mathbb{R}^{n+1} (thus it is learning a rotated copy of Δ^{n-1} instead of Δ^n). This is for notational convenience so that we work with n instead of $n+1$.

Subroutine 1 Find one vertex of a rotation of the standard simplex Δ^{n-1} via a fixed point iteration-like algorithm

Input: Samples from a rotated copy of the n -dimensional standard simplex (for a rotation that leaves $\mathbb{1}$ invariant).

Output: An approximation to a uniformly random vertex of the input simplex.

Pick $u(1) \in S^{n-1}$, uniformly at random.

for $i = 1$ to r **do**

$$u(i+1) := C_{n-1} \nabla \hat{m}_3(u(i)) - \frac{1}{2}(u(i) \cdot \mathbb{1})^2 \mathbb{1} - \frac{1}{2}(u(i) \cdot u(i))^2 \mathbb{1} - (u(i) \cdot \mathbb{1})u(i).$$

Normalize $u(i+1)$ by dividing by $\|u(i+1)\|_2$.

end for

Output $u(r+1)$.

Lemma 2 *Let $c > 0$ be a constant, $n > 20$, and $0 < \delta < 1$. Suppose that Subroutine 1 uses a sample of size $t = 2^{17} n^{2c+22} (\frac{1}{\delta})^2 \ln \frac{2n^5 r}{\delta}$ for each evaluation of the gradient and runs for $r = \log \frac{4(c+3)n^2 \ln n}{\delta}$ iterations. Then with probability at least $1 - \delta$ Subroutine 1 outputs a vector within distance $1/n^c$ from a vertex of the input simplex. With respect of the process of picking a sample and running the algorithm, each vertex is equally likely to be the nearest.*

See Appendix D for the proof.

4. Learning simplices

In this section we give our algorithm for learning general simplices, which uses the subroutine from the previous section. The learning algorithm uses an affine map $T : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1}$ that maps some isotropic simplex to the standard simplex. We describe now a way of constructing such a map: Let A be a matrix having as columns an orthonormal basis of $\mathbb{1}^\perp$ in \mathbb{R}^{n+1} . To compute one such A , one can start with the $(n+1)$ -by- $(n+1)$ matrix B that has ones in the diagonal and first column, everything else is zero. Let $QR = B$ be a QR-decomposition of B . By definition we have that the first column of Q is parallel to $\mathbb{1}$ and the rest of the columns span $\mathbb{1}^\perp$. Given this, let A be the matrix formed by all columns of Q except the first. We have that the set $\{A^T e_i\}$ is the set of vertices of a regular n -simplex. Each vertex is at distance

$$\sqrt{\left(1 - \frac{1}{n+1}\right)^2 + \frac{n}{(n+1)^2}} = \sqrt{\frac{n}{n+1}}$$

from the origin, while an isotropic simplex has vertices at distance $\sqrt{n(n+2)}$ from the origin. So an affine transformation that maps an isotropic simplex in \mathbb{R}^n to the standard simplex in \mathbb{R}^{n+1} is $T(x) = \frac{1}{\sqrt{(n+1)(n+2)}} Ax + \frac{1}{n+1} \mathbb{1}_{n+1}$.

To simplify the analysis, we pick a new sample $r(1), \dots, r(t_3)$ to find every vertex, as this makes every vertex equally likely to be found when given a sample from an isotropic simplex. (The core of the analysis is done for an isotropic simplex; this is enough as the algorithm's first step is to find an affine transformation that puts the input simplex

Algorithm 1 Learning a simplex.

Input: Error parameter $\epsilon > 0$. Probability of failure parameter $\delta > 0$. Oracle access to random points from some n -dimensional simplex S_{INPUT} .

Output: $V = \{v(1), \dots, v(n+1)\} \subseteq \mathbb{R}^n$ (approximations to the vertices of the simplex).

Estimate the mean and covariance using $t_1 = \text{poly}(n, 1/\epsilon, 1/\delta)$ samples $p(1), \dots, p(t_1)$:

$$\mu = \frac{1}{t_1} \sum_i p(i), \quad \Sigma = \frac{1}{t_1} \sum_i (p(i) - \mu)(p(i) - \mu)^T.$$

Compute a matrix B so that $\Sigma = BB^T$ (say, Cholesky decomposition).

Let $U = \emptyset$.

for $i = 1$ to m (with $m = \text{poly}(n, \log 1/\delta)$) **do**

 Get $t_3 = \text{poly}(n, 1/\epsilon, \log 1/\delta)$ samples $r(1), \dots, r(t_3)$ and use μ, B to map them to samples $s(i)$ from a nearly-isotropic simplex: $s(i) = B^{-1}(r(i) - \mu)$.

 Embed the resulting samples in \mathbb{R}^{n+1} as a sample from an approximately rotated standard simplex: Let $l(i) = T(s(i))$.

 Invoke Subroutine 1 with sample $l(1), \dots, l(t_3)$ to get $u \in \mathbb{R}^{n+1}$.

 Let \tilde{u} be the nearest point to u in the affine hyperplane $\{x : x \cdot \mathbb{1} = 1\}$. If \tilde{u} is not within $1/\sqrt{2}$ of a point in U , add \tilde{u} to U . (Here $1/\sqrt{2}$ is half of the edge length of the standard simplex.)

end for

Let

$$V = BT^{-1}(U) + \mu = \sqrt{(n+1)(n+2)}BA^T \left(U - \frac{1}{n+1}\mathbb{1} \right) + \mu.$$

in approximately isotropic position. The fact that this approximation is close in total variation distance implies that it is enough to analyze the algorithm for the case of exact isotropic position, the analysis carries over to the approximate case with a small loss in the probability of success. See the proof below for the details.) A practical implementation may prefer to select one such sample outside of the for loop, and find all the vertices with just that sample—an analysis of this version would involve bounding the probability that each vertex is found (given the sample, over the choice of the starting point of gradient descent) and a variation of the coupon collector’s problem with coupons that are not equally likely.

See Appendix E for the proof of Theorem 1.

5. Probabilistic Results

In this section we show the probabilistic results underlying the reductions from learning simplices and ℓ_p^n balls to ICA. The results are Theorems 4 and 6. They each show a simple non-linear rescaling of the respective uniform distributions that gives a distribution with independent components (Definition 3).

Theorem 4 below gives us, in a sense, a “reversal” of the representation of the cone measure on ∂B_p^n , see in Theorem 9. Given any random point in the standard simplex, we can apply a simple non-linear scaling and recover a distribution with independent components.

Definition 3 *We say that a random vector X has independent components if it is an affine transformation of a random vector having independent coordinates.*

Theorem 4 *Let X be a uniformly random vector in the $(n - 1)$ -dimensional standard simplex Δ_{n-1} . Let T be a random scalar distributed as $\text{Gamma}(n, 1)$. Then the coordinates of TX are iid as $\text{Exp}(1)$.*

Moreover, if $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an invertible linear transformation, then the random vector $TA(X)$ has independent components.

The proof of this theorem and others in this Section can be found in Appendix G.

The next lemma complements the main result in Barthe et al. (2005), Theorem 1 (Theorem 10 in Appendix A). They show a representation of the uniform distribution in B_p^n , but they do not state the independence that we need for our reduction to ICA.

Lemma 5 *Let $p \in [1, \infty)$. Let $G = (G_1, \dots, G_n)$ be iid random variables with density proportional to $\exp(-|t|^p)$. Let W be a nonnegative random variable with distribution $\text{Exp}(1)$ and independent of G . Then the random vector*

$$\frac{G}{(\|G\|_p^p + W)^{1/p}}$$

is independent of $(\|G\|_p^p + W)^{1/p}$.

Proof idea. We aim to compute the joint density, showing that it is a product of individual densities. To avoid complication, we raise everything to the p th power, which eliminates extensive use of the chain rule involved in the change of variables. \square

With this in mind, we show now our analog of Theorem 4 for B_p^n .

Theorem 6 *Let X be a uniformly random vector in B_p^n . Let T be a random scalar distributed as $\text{Gamma}((n/p) + 1, 1)$. Then the coordinates of $T^{1/p}X$ are iid, each with density proportional to $\exp(-|t|^p)$. Moreover, if $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an invertible linear transformation, then the random vector given by $T^{1/p}A(X)$ has independent components.*

This result shows that one can obtain a vector with independent components from a sample in a linearly transformed ℓ_p ball. In Section 6 we show that they are related in such a way that one can recover the linear transformation from the independent components via ICA.

6. Learning problems that reduce to ICA

Independent component analysis is a certain computational problem and an associated family of algorithms. Suppose that X is a random n -dimensional vector whose coordinates are independently distributed. The coordinates’ distributions are unknown and not necessarily

identical. The ICA problem can be stated as follows: given samples from an affine transformation $Y = AX + b$ of X , estimate A and b (up to a certain intrinsic indeterminacy: permutation and scaling of the columns of A). We will state more precisely below what is expected of an ICA algorithm.

We show randomized reductions from the following two natural statistical estimation problems to ICA:

Problem 1 (simplex) *Given uniformly random points from an n -dimensional simplex, estimate the simplex.*

This is the same problem of learning a simplex as in the rest of the paper, we just restate it here for clarity.

To simplify the presentation for the second problem, we ignore the estimation of the mean of an affinely transformed distribution. That is, we assume that the ℓ_p^n ball to be learned has only been *linearly* transformed.

Problem 2 (linearly transformed ℓ_p^n balls) *Given uniformly random points from a linear transformation of the ℓ_p^n -ball, estimate the linear transformation.*

These problems do not have an obvious independence structure. Nevertheless, known representations of the uniform measure in an ℓ_p^n ball and the *cone measure* (defined in Appendix A) on the surface of an ℓ_p^n ball can be slightly extended to map a sample from those distributions into a sample with independent components by a non-linear scaling step. The use of a non-linear scaling step to turn a distribution into one having independent components has been done before [Sinz and Bethge \(2010, 2008\)](#), but there it is applied *after* finding a transformation that makes the distribution axis-aligned. This alignment is attempted using ICA (or variations of PCA) on the original distribution [Sinz and Bethge \(2010, 2008\)](#), without independent components, and therefore the use of ICA is somewhat heuristic. One of the contributions of our reduction is that the rescaling we apply is “blind”, namely, it can be applied to the original distribution. In fact, the distribution does not even need to be isotropic (“whitened”). The distribution resulting from the reduction has independent components and therefore the use of ICA on it is well justified.

The reductions are given in Algorithms 2 and 3. To state the reductions, we denote by $ICA(s(1), s(2), \dots)$ the invocation of an ICA routine. It takes samples $s(1), s(2), \dots$ of a random vector $Y = AX + \mu$, where the coordinates of X are independent, and returns an approximation to a square matrix M such that $M(Y - \mathbb{E}(Y))$ is isotropic and has independent coordinates. The theory of ICA ([Comon, 1994](#), Theorem 11) implies that if X is isotropic and at most one coordinate is distributed as a Gaussian, then such an M exists and it satisfies $MA = DP$, where P is a permutation matrix and D is a diagonal matrix with diagonal entries in $\{-1, 1\}$. We thus need the following definition to state our reduction: Let $c_{p,n} = (\mathbb{E}_{X \in B_p^n}(X_1^2))^{1/2}$. That is, the uniform distribution in $B_p^n/c_{p,n}$ is isotropic.

As we do not state a full analysis of any particular ICA routine, we do not state explicit approximation guarantees.

Algorithm 2 works as follows: Let X be an $(n + 1)$ -dimensional random vector with iid coordinates distributed as $\text{Exp}(1)$. Let V be the matrix having columns $(v(i), 1)$ for

Algorithm 2 Reduction from Problem 1 to ICA

Input: A uniformly random sample $p(1), \dots, p(t)$ from an n -dimensional simplex S .

Output: Vectors $\tilde{v}(1), \dots, \tilde{v}(n+1)$ such that their convex hull is close to S .

Embed the sample in \mathbb{R}^{n+1} : Let $p'(i) = (p(i), 1)$.

For every $i = 1, \dots, t$, generate a random scalar $T(i)$ distributed as $\text{Gamma}(n+1, 1)$.

Let $q(i) = p'(i)T(i)$.

Invoke $ICA(q(1), \dots, q(t))$ to obtain an approximately separating matrix \tilde{M} .

Compute the inverse of \tilde{M} and multiply every column by the sign of its last entry to get a matrix \tilde{A} .

Remove the last row of \tilde{A} and return the columns of the resulting matrix as $\tilde{v}(1), \dots, \tilde{v}(n+1)$.

$i = 1, \dots, n+1$. Let Y be random according to the distribution that results from scaling in the algorithm. Theorem 4 implies that Y and VX have the same distribution. Also, $X - \mathbb{1}$ is isotropic and Y and $V(X - \mathbb{1}) + V\mathbb{1}$ have the same distribution. Thus, the discussion about ICA earlier in this section gives that the only separating matrices M are such that $MV = DP$ where P is a permutation matrix and D is a diagonal matrix with diagonal entries in $\{-1, 1\}$. That is, $VP^T = M^{-1}D$. As the last row of V is all ones, the sign change step in Algorithm 2 undoes the effect of D and recovers the correct orientation.

Algorithm 3 Reduction from Problem 2 to ICA

Input: A uniformly random sample $p(1), \dots, p(t)$ from $A(B_p^n)$ for a known parameter $p \in [1, \infty)$, where $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an unknown invertible linear transformation.

Output: A matrix \tilde{A} such that $\tilde{A}B_p^n$ is close to $A(B_p^n)$.

For every $i = 1, \dots, t$, generate a random scalar $T(i)$ distributed as $\text{Gamma}((n/p) + 1, 1)$.

Let $q(i) = p(i)T(i)^{1/p}$.

Invoke $ICA(q(1), \dots, q(t))$ to obtain an approximately separating matrix \tilde{M} .

Output $\tilde{A} = c_{p,n}^{-1} \tilde{M}^{-1}$.

Similarly, Algorithm 3 works as follows: Let X be a random vector with iid coordinates, each with density proportional to $\exp(-|t|^p)$. Let Y be random according to the distribution that results from scaling in the algorithm. Theorem 6 implies that Y and AX have the same distribution. Also, $X/c_{p,n}$ is isotropic and we have Y and $Ac_{p,n}(X/c_{p,n})$ have the same distribution. Thus, the discussion about ICA earlier in this section gives that the only separating matrices M are such that $MAc_{p,n} = DP$ where P is a permutation matrix and D is a diagonal matrix with diagonal entries in $\{-1, 1\}$. That is, $AP^T D^{-1} = c_{p,n}^{-1} M^{-1}$. The fact that B_p^n is symmetric with respect coordinate permutations and sign changes implies that $AP^T D^{-1} B_p^n = AB_p^n$ and is the same as $c_{p,n}^{-1} M^{-1}$. When $p \neq 2$, the assumptions in the discussion above about ICA are satisfied and Algorithm 3 is correct. When $p = 2$, the distribution of the scaled sample is Gaussian and this introduces ambiguity with respect to rotations in the definition of M , but this ambiguity is no problem as it is counteracted by the fact that the l_2 ball is symmetric with respect to rotations.

7. Conclusion

We showed, in two different ways, that the problem of learning simplices can be solved efficiently using techniques for ICA. We also showed that when the sample is one that may not satisfy the requirement of independent components, we can efficiently obtain from it a sample that guarantees this property and from which the original distribution can be estimated. Many questions remain: Can we do this for other polytopes? Can we do this when the points come from the Gaussian distribution with labels instead of the uniform distribution in the polytope? In particular, does any one of the two techniques that we used in this paper for learning simplices extend to learning polytopes or to latent variable models?

It is not clear how to extend our method to learn intersections of more than $n + 1$ half-spaces. The main hurdle seems to be that the nice symmetric form of the moment polynomial that we get for the simplex is no longer there, and it's not clear how the polynomial is related to the features of the polytope. We don't even know whether or not there is a one-to-one correspondence between polytopes and the moment polynomials (of low constant degree). Interestingly, it is also not clear how to learn intersections of n (or fewer, but not much fewer) half-spaces, for similar reasons: First note that we cannot consider uniform distributions over such sets because they are unbounded. As mentioned before, there is considerable literature when one considers other distributions, such as the Gaussian distribution. However, our moment-based method runs into the same difficulty as before, namely the relation between the moment polynomial and the half-spaces defining the polyhedron is not clear.

As pointed out by a COLT referee, our second method for learning simplices, namely direct reduction to ICA, can possibly be used to learn the more general Dirichlet distribution on the simplex when combined with the results of Barthe et al. (2010). This is an interesting direction for future work.

Acknowledgments

We thank Santosh Vempala for telling us about the polytope learning problem, the approach of using higher-order moments and for helpful discussions. We also thank Keith Ball, Alexander Barvinok, Franck Barthe, Mikhail Belkin, Adam Kalai, Assaf Naor, Aaditya Ramdas, Roman Vershynin and James Voss for helpful discussions. We thank anonymous COLT referees for useful remarks and pointers to the literature.

References

- Radosław Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *J. Amer. Math. Soc.*, 23(2):535–561, 2010. ISSN 0894-0347. doi: 10.1090/S0894-0347-09-00650-X. URL <http://dx.doi.org/10.1090/S0894-0347-09-00650-X>.
- Anima Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *CoRR*, abs/1210.7559, 2012a.

- Animashree Anandkumar, Daniel Hsu, and Sham M. Kakade. A method of moments for mixture models and hidden markov models. In *COLT*, 2012b. <http://arxiv.org/abs/1203.0683>.
- Sanjeev Arora, Rong Ge, Ankur Moitra, and Sushant Sachdeva. Provable ICA with unknown Gaussian noise, and implications for Gaussian mixtures and autoencoders. In *NIPS*, 2012. arXiv:1206.5349.
- F. Barthe, O. Guédon, S. Mendelson, and A. Naor. A probabilistic approach to the geometry of the ℓ_p^n -ball. *The Annals of Probability*, 33(2):480–513, 2005.
- F. Barthe, F. Gamboa, L. Lozada-Chang, and A. Rouault. Generalized Dirichlet distributions on the ball and moments. *ALEA Lat. Am. J. Probab. Math. Stat.*, 7:319–340, 2010. ISSN 1980-0436.
- Patrick Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, third edition, 1995. ISBN 0-471-00710-2. A Wiley-Interscience Publication.
- P. Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.
- Pierre Comon and Christian Jutten, editors. *Handbook of Blind Source Separation*. Academic Press, 2010.
- Alan M. Frieze, Mark Jerrum, and Ravi Kannan. Learning linear transformations. In *FOCS*, pages 359–368, 1996.
- Navin Goyal and Luis Rademacher. Learning convex bodies is hard. In *COLT*, 2009.
- Nick Gravin, Jean Lasserre, Dmitrii V. Pasechnik, and Sinai Robins. The inverse moment problem for convex polytopes. *Discrete & Computational Geometry*, 48(3):596–621, 2012.
- A. Grundmann and H. M. Moeller. Invariant integration formulas for the n -simplex by combinatorial methods. *SIAM J. Numer. Anal.*, 15:282–290, 1978.
- A. K. Gupta and D. Song. L_p -norm spherical distribution. *J. Statist. Plann. Inference*, 60(2):241–260, 1997. ISSN 0378-3758. doi: 10.1016/S0378-3758(96)00129-2. URL [http://dx.doi.org/10.1016/S0378-3758\(96\)00129-2](http://dx.doi.org/10.1016/S0378-3758(96)00129-2).
- Aapo Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, 1999.
- Aapo Hyvärinen, Juha Karhunen, and Erkki Oja. *Independent Component Analysis*. Wiley, 2001.
- Adam R. Klivans and Alexander A. Sherstov. A lower bound for agnostically learning disjunctions. In *COLT*, pages 409–423, 2007.
- Adam R. Klivans and Alexander A. Sherstov. Cryptographic hardness for learning intersections of halfspaces. *J. Comput. Syst. Sci.*, 75(1):2–12, 2009.

- Adam R. Klivans, Ryan O’Donnell, and Rocco A. Servedio. Learning geometric concepts via Gaussian surface area. In *FOCS*, pages 541–550, 2008.
- Stephen Kwek and Leonard Pitt. PAC learning intersections of halfspaces with membership queries. *Algorithmica*, 22(1/2):53–75, 1998.
- Jean B. Lasserre and Konstantin E. Avrachenkov. The multi-dimensional version of $\int a^b x^p dx$. *American Math. Month.*, 108(2):151–154, 2001.
- David G. Luenberger and Yinyu Ye. *Linear and nonlinear programming*. International Series in Operations Research & Management Science, 116. Springer, New York, third edition, 2008. ISBN 978-0-387-74502-2.
- Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *J. Cryptology*, 22(2):139–160, 2009.
- Asa Packer. NP-hardness of largest contained and smallest containing simplices for V- and H-polytopes. *Discrete & Computational Geometry*, 28(3):349–377, 2002.
- S.T. Rachev and L. Ruschendorf. Approximate independence of distributions on spheres and their stability properties. *The Annals of Probability*, 19(3):1311–1337, 1991.
- G. Schechtman and J. Zinn. On the volume of the intersection of two l_p^n balls. In *Proc. Amer. Math. Soc.*, volume 110, pages 217–224, 1990.
- Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.
- Fabian Sinz and Matthias Bethge. L_p -nested symmetric distributions. *J. Mach. Learn. Res.*, 11:3409–3451, 2010. ISSN 1532-4435.
- Fabian H. Sinz and Matthias Bethge. The conjoint effect of divisive normalization and orientation selectivity on redundancy reduction. In *NIPS*, pages 1521–1528, 2008.
- D. Song and A. K. Gupta. L_p -norm uniform distribution. *Proc. Amer. Math. Soc.*, 125(2):595–601, 1997. ISSN 0002-9939. doi: 10.1090/S0002-9939-97-03900-2. URL <http://dx.doi.org/10.1090/S0002-9939-97-03900-2>.
- Nikhil Srivastava and Roman Vershynin. Covariance estimation for distributions with $2 + \epsilon$ moments, 2011.
- Santosh Vempala. A random-sampling-based algorithm for learning intersections of halfspaces. *J. ACM*, 57(6):32, 2010a.
- Santosh Vempala. Learning convex concepts from Gaussian distributions with PCA. In *FOCS*, pages 124–130, 2010b.
- Santosh S. Vempala and Ying Xiao. Structure from local optima: Learning subspace juntas via higher order PCA. *CoRR*, abs/1108.3329, 2011.

Appendix A. Additional preliminaries

Let B_n denote the n -dimensional Euclidean ball.

The complete homogeneous symmetric polynomial of degree d in variables u_0, \dots, u_n , denoted $h_d(u_0, \dots, u_n)$, is the sum of all monomials of degree d in the variables:

$$h_d(u_0, \dots, u_n) = \sum_{k_0 + \dots + k_n = d} u_0^{k_0} \dots u_n^{k_n} = \sum_{0 \leq i_0 \leq i_1 \leq \dots \leq i_d \leq n} u_{i_0} u_{i_1} \dots u_{i_d}.$$

Lemma 7 *Let K, L be two compact sets in \mathbb{R}^n . Let $0 < \alpha \leq 1 \leq \beta$ such that $\alpha K \subseteq L \subseteq \beta K$. Then $d_{TV}(K, L) \leq 2(1 - (\alpha/\beta)^n)$.*

Proof We have $d_{TV}(\alpha K, \beta K) = 1 - (\alpha/\beta)^n$. Triangle inequality implies the desired inequality. \blacksquare

Lemma 8 *Consider the coupon collector's problem with n coupons where every coupon occurs with probability at least α . Let $\delta > 0$. Then with probability at least $1 - \delta$ all coupons are collected after $\alpha^{-1}(\log n + \log 1/\delta)$ trials.*

Proof The probability that a particular coupon is not collected after that many trials is at most

$$(1 - \alpha)^{\alpha^{-1}(\log n + \log 1/\delta)} \leq e^{-\log n - \log 1/\delta} = \delta/n.$$

The union bound over all coupons implies the claim. \blacksquare

The cone measure on the surface ∂K of centrally symmetric convex body K in \mathbb{R}^n [Barthe et al. \(2005\)](#); [Schechtman and Zinn \(1990\)](#); [Rachev and Ruschendorf \(1991\)](#); [Song and Gupta \(1997\)](#) is defined by

$$\mu_K(A) = \frac{\text{vol}(ta; a \in A, 0 \leq t \leq 1)}{\text{vol}(K)}.$$

It is easy to see that $\mu_{B_p^n}$ is uniform on ∂B_p^n for $p \in \{1, 2, \infty\}$.

From [Schechtman and Zinn \(1990\)](#) and [Rachev and Ruschendorf \(1991\)](#) we have the following representation of the cone measure on ∂B_p^n :

Theorem 9 *Let G_1, G_2, \dots, G_n be iid random variables with density proportional to $\exp(-|t|^p)$. Then the random vector $X = G/\|G\|_p$ is independent of $\|G\|_p$. Moreover, X is distributed according to $\mu_{B_p^n}$.*

From [Barthe et al. \(2005\)](#), we also have the following variation, a representation of the uniform distribution in B_p^n :

Theorem 10 *Let $G = (G_1, \dots, G_n)$ be iid random variables with density proportional to $\exp(-|t|^p)$. Let Z be a random variable distributed as $\text{Exp}(1)$, independent of G . Then the random vector*

$$V = \frac{G}{(\sum_{i=1}^n |G_i|^p + Z)^{1/p}}$$

is uniformly distributed in B_p^n .

See e.g. [\(Billingsley, 1995, Section 20\)](#) for the change of variable formula in probability.

Appendix B. The local and global maxima of the 3rd moment of the standard simplex and the isotropic simplex

In this section we study the structure of the set of local maxima of the third moment as a function of the direction (which happens to be essentially $u \mapsto \sum u_i^3$ as discussed in Section F). This is not necessary for our algorithmic result, however it gives insight into the geometry of the third moment (the location of local maxima/minima and stationary points) and suggests that more direct optimization algorithms like gradient descent and Newton's method will also work, although we will not prove that.

Theorem 11 *Let $K \subseteq \mathbb{R}^n$ be an isotropic simplex. Let X be random in K . Let $V = \{x_i\}_{i=1}^{n+1} \subseteq \mathbb{R}^n$ be the set of normalized vertices of K . Then V is a complete set of local maxima and a complete set of global maxima of $F : S^{n-1} \rightarrow \mathbb{R}$ given by $F(u) = \mathbb{E}((u \cdot X)^3)$.*

Proof idea: Embed the simplex in \mathbb{R}^{n+1} . Show that the third moment is proportional to the complete homogeneous symmetric polynomial of degree 3, which for the relevant directions is proportional to the sum of cubes. To conclude, use first and second order optimality conditions to characterize the set of local maxima.

Proof Consider the standard simplex

$$\Delta^n = \text{conv}\{e_1, \dots, e_{n+1}\} \subseteq \mathbb{R}^{n+1}$$

and identify it with V via a linear map $A : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ so that $A(\Delta^n) = V$. Let Y be random in Δ^n . Consider $G : S^n \rightarrow \mathbb{R}$ given by $G(v) = m_3(v) = \mathbb{E}((v \cdot Y)^3)$. Let $U = \{v \in \mathbb{R}^{n+1} : v \cdot \mathbb{1} = 0, \|v\| = 1\}$ be the equivalent feasible set for the embedded problem. We have $G(v) = cF(Av)$ for any $v \in U$ and some constant $c > 0$ independent of v . To get the theorem, it is enough to show that the local maxima of G in U are precisely the normalized versions of the projections of the canonical vectors onto the hyperplane orthogonal to $\mathbb{1} = (1, \dots, 1)$. According to Section F, for $v \in U$ we have

$$G(v) \propto p_3(v).$$

Using a more convenient but equivalent constant, we want to enumerate the local maxima of the problem

$$\begin{aligned} \max \quad & \frac{1}{3} p_3(v) \\ \text{s.t.} \quad & v \cdot v = 1 \\ & v \cdot \mathbb{1} = 0 \\ & v \in \mathbb{R}^{n+1}. \end{aligned} \tag{2}$$

The Lagrangian function is

$$L(v, \lambda_1, \lambda_2) = \frac{1}{3} \sum_i v_i^3 - \lambda_1 \sum_i v_i - \lambda_2 \frac{1}{2} \left(\left(\sum_i v_i^2 \right) - 1 \right).$$

The first order condition is $\nabla_v L = 0$, that is,

$$v_i^2 = \lambda_1 + \lambda_2 v_i \quad \text{for } i = 1, \dots, n+1. \tag{3}$$

Consider this system of equations on v for any fixed λ_1, λ_2 . Let $f(x) = x^2, g(x) = \lambda_1 + \lambda_2 x$. The first order condition says $f(v_i) = g(v_i)$, where f is convex and g is affine. That is, the v_i s can take at most two different values. As our optimization problem (2) is symmetric under permutation of the coordinates, we conclude that, after putting the coordinates of a point v in non-increasing order, if v is a local maximum of (2), then v must be of the form

$$v = (a, \dots, a, b, \dots, b),$$

where $a > 0 > b$ and there are exactly α a s and β b s, for $\alpha, \beta \in \{1, \dots, n\}$.

We will now study the second order necessary condition (SONC) to eliminate from the list of candidates all vectors with $\alpha > 1$. It is easy to see that the surviving vectors are exactly the promised scaled projections of the canonical vectors. These vectors must all be local and global maxima: At least one of them must be a global maximum as we are maximizing a continuous function over a compact set and all of them have the same objective value so all of them are local and global maxima.

The SONC at v asks for the Hessian of the Lagrangian to be negative semidefinite when restricted to the tangent space to the constraint set at v (Luenberger and Ye, 2008, Section 11.5). We compute the Hessian (recall that $v^{(2)}$ is the vector of the squared coordinates of v):

$$\begin{aligned} \nabla_v L &= v^{(2)} - \lambda_1 \mathbb{1} - \lambda_2 v \\ \nabla_v^2 L &= 2 \operatorname{diag}(v) - \lambda_2 I \end{aligned}$$

where $\operatorname{diag}(v)$ is the $(n+1)$ -by- $(n+1)$ matrix having the entries of v in the diagonal and 0 elsewhere.

A vector in the tangent space is any $z \in \mathbb{R}^{n+1}$ such that $z \cdot \mathbb{1} = 0, v \cdot z = 0$, and definiteness of the Hessian is determined by the sign of $z^T \nabla_v^2 L z$ for any such z , where

$$z^T \nabla_v^2 L z = \sum_{i=1}^{n+1} z_i^2 (2v_i - \lambda_2).$$

Suppose v is a critical point with $\alpha \geq 2$. To see that such a v cannot be a local maximum, it is enough to show $2a > \lambda_2$, as in that case we can take $z = (1, -1, 0, \dots, 0)$ to make the second derivative of L positive in the direction z .

In terms of α, β, a, b , the constraints of (2) are $\alpha a + \beta b = 0, \alpha a^2 + \beta b^2 = 1$, and this implies $a = \sqrt{\frac{\beta}{\alpha(n+1)}}, b = -\sqrt{\frac{\alpha}{\beta(n+1)}}$. The inner product between the first order condition (3) and v implies $\lambda_2 = \sum v_i^3 = \alpha a^3 + \beta b^3$. It is convenient to consider the change of variable $\gamma = \alpha/(n+1)$, as now candidate critical points are parameterized by certain discrete values

of γ in $(0, 1)$. This gives $\beta = (1 - \gamma)(n + 1)$, $a = \sqrt{(1 - \gamma)/(\gamma(n + 1))}$ and

$$\begin{aligned} \lambda_2 &= (n + 1) \left[\gamma \left(\frac{1 - \gamma}{\gamma(n + 1)} \right)^{3/2} \right. \\ &\quad \left. - (1 - \gamma) \left(\frac{\gamma}{(1 - \gamma)(n + 1)} \right)^{3/2} \right] \\ &= \frac{1}{\sqrt{(n + 1)\gamma(1 - \gamma)}} [(1 - \gamma)^2 - \gamma^2] \\ &= \frac{1}{\sqrt{(n + 1)\gamma(1 - \gamma)}} [1 - 2\gamma]. \end{aligned}$$

This implies

$$\begin{aligned} 2a - \lambda_2 &= \frac{1}{\sqrt{(n + 1)\gamma(1 - \gamma)}} [2(1 - \gamma) - 1 + 2\gamma] \\ &= \frac{1}{\sqrt{(n + 1)\gamma(1 - \gamma)}}. \end{aligned}$$

In $(0, 1)$, the function given by $\gamma \mapsto 2a - \lambda_2 = \frac{1}{\sqrt{(n+1)\gamma(1-\gamma)}}$ is convex and symmetric around $1/2$, where it attains its global minimum value, $2/\sqrt{n + 1}$, which is positive. \blacksquare

Appendix C. Improving the dependence on the probability of success

The probability of success of the algorithm can be “boosted” so that the dependence of the running time on δ is only linear in $\log(1/\delta)$ as follows: The following discussion uses the space of simplices with total variation distance as the underlying metric space. Let ϵ be the target distance. Take an algorithm that succeeds with probability $5/6$ and error parameter ϵ' to be fixed later (such as Algorithm 1 with $\delta = 1/6$). Run the algorithm $t = O(\log 1/\delta)$ times to get t simplices. By a Chernoff-type argument, at least $2t/3$ simplices are within ϵ' of the input simplex with probability at least $1 - \delta/2$.

By sampling, we can estimate the distances between all pairs of simplices with additive error less than $\epsilon'/10$ in time polynomial in $t, 1/\epsilon'$ and $\log 1/\delta$ so that all estimates are correct with probability at least $1 - \delta/2$. For every output simplex, compute the number of output simplices within estimated distance $(2 + 1/10)\epsilon'$. With probability at least $1 - \delta$ both of the desirable events happen, and then necessarily there is at least one output simplex, call it S , that has $2t/3$ output simplices within estimated distance $(2 + 1/10)\epsilon'$. Any such S must be within $(3 + 2/10)\epsilon'$ of the input simplex. Thus, set $\epsilon' = \epsilon/(3 + 2/10)$.

Appendix D. Proof of Lemma 2

Note that if we condition on the sample, different vertices are not equally likely over the randomness of the algorithm. That is, if we try to find all vertices running the algorithm multiple times on a fixed sample, different vertices will be found with different likelihoods.

Proof Our analysis has the same outline as that of [Nguyen and Regev \(2009\)](#). This is because the iteration that we get is the same as that of [Nguyen and Regev \(2009\)](#) except that cubing is replaced by squaring (see below); however some details in our proof are different. In the proof below, several of the inequalities are quite loose and are so chosen to make the computations simpler.

We first prove the lemma assuming that the gradient computations are exact and then show how to handle samples. We will carry out the analysis in the coordinate system where the given simplex is the standard simplex. This is only for the purpose of the analysis, and this coordinate system is not known to the algorithm. Clearly, $u(i+1) = (u(i)_1^2, \dots, u(i)_n^2)$. It follows that,

$$u(i+1) = (u(1)_1^{2^i}, \dots, u(1)_n^{2^i}).$$

Now, since we choose $u(1)$ randomly, with probability at least $(1 - (n^2 - n)\delta')$ one of the coordinates of $u(1)$ is greater than all the other coordinates in absolute value by a factor of at least $(1 + \delta')$, where $0 < \delta' < 1$. (A similar argument is made in [Nguyen and Regev \(2009\)](#) with different parameters. We briefly indicate the proof for our case: The probability that the event in question does not happen is less than the probability that there are two coordinates $u(1)_a$ and $u(1)_b$ such that their absolute values are within factor $1 + \delta'$, i.e. $1/(1 + \delta') \leq |u(1)_a|/|u(1)_b| < 1 + \delta'$. The probability that for given a, b this event happens can be seen as the Gaussian area of the four sectors (corresponding to the four choices of signs of $u(1)_a, u(1)_b$) in the plane each with angle less than $2\delta'$. By symmetry, the Gaussian volume of these sectors is $2\delta'/(\pi/2) < 2\delta'$. The probability that such a pair (a, b) exists is less than $2\binom{n}{2}\delta'$.) Assuming this happens, then after r iterations, the ratio between the largest coordinate (in absolute value) and the absolute value of any other coordinate is at least $(1 + \delta')^{2^r}$. Thus, one of the coordinates is very close to 1 and others are very close to 0, and so $u(r+1)$ is very close to a vertex of the input simplex.

Now we drop the assumption that the gradient is known exactly. For each evaluation of the gradient we use a fresh subset of samples of t points. Here t is chosen so that each evaluation of the gradient is within ℓ_2 -distance $1/n^{c_1}$ from its true value with probability at least $1 - \delta''$, where c_1 will be set at the end of the proof. An application of the Chernoff bound yields that we can take $t = 200n^{2c_1+4} \ln \frac{2n^3}{\delta''}$; we omit the details. Thus all the r evaluations of the gradient are within distance $1/n^{c_1}$ from their true values with probability at least $1 - r\delta''$.

We assumed that our starting vector $u(1)$ has a coordinate greater than every other coordinate by a factor of $(1 + \delta')$ in absolute value; let us assume without loss of generality that this is the first coordinate. Hence $|u(1)_1| \geq 1/\sqrt{n}$. When expressing $u^{(2)}$ in terms of the gradient, the gradient gets multiplied by $C_{n-1} < n^3$ (we are assuming $n > 20$), keeping this in mind and letting $c_2 = c_1 - 3$ we get for $j \neq 1$

$$\frac{|u(i+1)_1|}{|u(i+1)_j|} \geq \frac{u(i)_1^2 - 1/n^{c_2}}{u(i)_j^2 + 1/n^{c_2}} \geq \frac{u(i)_1^2(1 - n^{-(c_2-1)})}{u(i)_j^2 + 1/n^{c_2}}.$$

If $u(i)_j^2 > 1/n^{c_2-c_3}$, where $1 \leq c_3 \leq c_2 - 2$ will be determined later, then we get

$$\begin{aligned} |u(i+1)_1|/|u(i+1)_j| &> \frac{1 - 1/n^{c_2-1}}{1 + 1/n^{c_3}} \cdot \left(\frac{u(i)_1}{u(i)_j}\right)^2 \\ &> (1 - 1/n^{c_3})^2 \left(\frac{u(i)_1}{u(i)_j}\right)^2. \end{aligned} \quad (4)$$

Else,

$$\begin{aligned} |u(i+1)_1|/|u(i+1)_j| &> \frac{1/n - 1/n^{c_2}}{1/n^{c_2-c_3} + 1/n^{c_2}} \\ &> \left(1 - \frac{1}{n^{c_3}}\right)^2 \cdot n^{c_2-c_3-1} \\ &> \frac{1}{2}n^{c_2-c_3-1}, \end{aligned}$$

where we used $c_3 \geq 1$ and $n > 20$ in the last inequality.

We choose c_3 so that

$$\left(1 - \frac{1}{n^{c_3}}\right)^2 (1 + \delta') > (1 + \delta'/2). \quad (5)$$

For this, $\delta' \geq 32/n^{c_3}$ or equivalently $c_3 \geq (\ln(32/\delta'))/\ln n$ suffices.

For c_3 satisfying (5) we have $(1 - \frac{1}{n^{c_3}})^2(1 + \delta')^2 > (1 + \delta')$. It then follows from (4) that the first coordinate continues to remain the largest in absolute value by a factor of at least $(1 + \delta')$ after each iteration. Also, once we have $|u(i)_1|/|u(i)_j| > \frac{1}{2}n^{c_2-c_3-1}$, we have $|u(i')_1|/|u(i')_j| > \frac{1}{2}n^{c_2-c_3-1}$ for all $i' > i$.

(4) gives that after r iterations we have

$$\begin{aligned} \frac{|u(r+1)_1|}{|u(r+1)_j|} &> (1 - 1/n^{c_3})^{2+2^2+\dots+2^r} \left(\frac{u(1)_1}{u(1)_j}\right)^{2^r} \\ &\geq (1 - 1/n^{c_3})^{2^{r+1}-2} (1 + \delta')^{2^r}. \end{aligned}$$

Now if r is such that $(1 - 1/n^{c_3})^{2^{r+1}-2} (1 + \delta')^{2^r} > \frac{1}{2}n^{c_2-c_3-1}$, we will be guaranteed that $|u(r+1)_1|/|u(r+1)_j| > \frac{1}{2}n^{c_2-c_3-1}$. This condition is satisfied if we have $(1 - 1/n^{c_3})^{2^{r+1}} (1 + \delta')^{2^r} > \frac{1}{2}n^{c_2-c_3-1}$, or equivalently $((1 - 1/n^{c_3})^2 (1 + \delta'))^{2^r} \geq \frac{1}{2}n^{c_2-c_3-1}$. Now using (5) it suffices to choose r so that $(1 + \delta'/2)^{2^r} \geq \frac{1}{2}n^{c_2-c_3-1}$. Thus we can take $r = \log(4(c_2 - c_3)(\ln n)/\delta')$.

Hence we get $|u(r+1)_1|/|u(r+1)_j| > \frac{1}{2}n^{c_2-c_3-1}$. It follows that for $u(r+1)$, the ℓ_2 -distance from the vertex $(1, 0, \dots, 0)$ is at most $8/n^{c_2-c_3-2} < 1/n^{c_2-c_3-3}$ for $n > 20$; we omit easy details.

Now we set our parameters: $c_3 = 1 + (\ln(32/\delta'))/\ln n$ and $c_2 - c_3 - 3 = c$ and $c_1 = c_2 + 3 = 7 + c + \ln(32/\delta')/\ln n$ satisfies all the constraints we imposed on c_1, c_2, c_3 . Choosing $\delta'' = \delta'/r$, we get that the procedure succeeds with probability at least $1 - (n^2 - n)\delta' - r\delta'' > 1 - n^2\delta'$. Now setting $\delta' = \delta/n^2$ gives the overall probability of error δ , and the number of samples and iterations as claimed in the lemma. \blacksquare

Appendix E. Proof of Theorem 1

Proof [Proof of Theorem 1] As a function of the input simplex, the distribution of the output of the algorithm is equivariant under invertible affine transformations. Namely, if we apply an affine transformation to the input simplex, the distribution of the output is equally transformed.² The notion of error, total variation distance, is also invariant under invertible affine transformations. Therefore, it is enough to analyze the algorithm when the input simplex is in isotropic position. In this case $\|p(i)\| \leq n+1$ (see Section 2) and we can set $t_1 \leq \text{poly}(n, 1/\epsilon', \log(1/\delta))$ so that $\|\mu\| \leq \epsilon'$ with probability at least $1 - \delta/10$ (by an easy application of Chernoff's bound), for some ϵ' to be fixed later. Similarly, using results from (Adamczak et al., 2010, Theorem 4.1), a choice of $t_1 \leq n\epsilon'^{-2} \text{polylog}(1/\epsilon') \text{polylog}(1/\delta)$ implies that the empirical second moment matrix

$$\bar{\Sigma} = \frac{1}{t_1} \sum_i p(i)p(i)^T$$

satisfies $\|\bar{\Sigma} - I\| \leq \epsilon'$ with probability at least $1 - \delta/10$. We have $\Sigma = \bar{\Sigma} - \mu\mu^T$ and this implies $\|\Sigma - I\| \leq \|\bar{\Sigma} - I\| + \|\mu\mu^T\| \leq 2\epsilon'$. Now, $s(1), \dots, s(t_3)$ is an iid sample from a simplex $S' = B^{-1}(S_{INPUT} - \mu)$. Simplex S' is close in total variation distance to some isotropic simplex³ S_{ISO} . More precisely, Lemma 12 below shows that

$$d_{TV}(S', S_{ISO}) \leq 12n\epsilon', \tag{6}$$

with probability at least $1 - \delta/5$.

Assume for a moment that $s(1), \dots, s(t_3)$ are from S_{ISO} . The analysis of Subroutine 1 (fixed point-like iteration) given in Lemma 2 would guarantee the following: Successive invocations to Subroutine 1 find approximations to vertices of $T(S_{ISO})$ within Euclidean distance ϵ'' for some ϵ'' to be determined later and $t_3 = \text{poly}(n, 1/\epsilon'', \log 1/\delta)$. We ask for each invocation to succeed with probability at least $1 - \delta/(20m)$ with $m = n(\log n + \log 20/\delta)$. Note that each vertex is equally likely to be found. The choice of m is so that, if all m invocations succeed (which happens with probability at least $1 - \delta/20$), then the analysis of the coupon collector's problem, Lemma 8, implies that we fail to find a vertex with probability at most $\delta/20$. Overall, we find all vertices with probability at least $1 - \delta/10$.

But in reality samples $s(1), \dots, s(t_3)$ are from S' , which is only *close* to S_{ISO} . The estimate from (6) with appropriate $\epsilon' = \text{poly}(1/n, \epsilon'', \delta)$ gives

$$d_{TV}(S', S_{ISO}) \leq \frac{\delta}{10} \frac{1}{t_3 m},$$

which implies that the total variation distance between the joint distribution of all $t_3 m$ samples used in the loop and the joint distribution of actual samples from the isotropic

2. To see this: the equivariance of the algorithm as a map between distributions is implied by the equivariance of the algorithm on any given input sample. Now, given the input sample, if we apply an affine transformation to it, this transformation is undone except possibly for a rotation by the step $s(i) = B^{-1}(r(i) - \mu)$. A rotation may remain because of the ambiguity in the characterization of B . But the steps of the algorithm that follow the definition of $s(i)$ are equivariant under rotation, and the ambiguous rotation will be removed at the end when B is applied again in the last step.

3. The isotropic simplex S_{ISO} will typically be far from the (isotropic) input simplex, because of the ambiguity up to orthogonal transformations in the characterization of B .

simplex S_{ISO} is at most $\delta/10$, and this implies that the loop finds approximations to all vertices of $T(S_{ISO})$ when given samples from S' with probability at least $1 - \delta/5$. The points in U are still within Euclidean distance ϵ'' of corresponding vertices of $T(S_{ISO})$.

To conclude, we turn our estimate of distances between estimated and true vertices into a total variation estimate, and map it back to the input simplex. Let $S'' = \text{conv} T^{-1}U$. As T maps an isotropic simplex to a standard simplex, we have that $\sqrt{(n+1)(n+2)}T$ is an isometry, and therefore the vertices of S'' are within distance $\epsilon''/\sqrt{(n+1)(n+2)}$ of the corresponding vertices of S_{ISO} . Thus, the corresponding support functions are uniformly within

$$\epsilon''' = \epsilon''/\sqrt{(n+1)(n+2)}$$

of each other on the unit sphere. This and the fact that $S_{ISO} \supseteq B_n$ imply

$$(1 - \epsilon''')S_{ISO} \subseteq S'' \subseteq (1 + \epsilon''')S_{ISO}.$$

Thus, by Lemma 7, $d_{TV}(S'', S_{ISO}) \leq 1 - \frac{1-\epsilon'''}{1+\epsilon'''} \leq 1 - (1 - \epsilon''')^{2n} \leq 2n\epsilon''' \leq 2\epsilon''$ and this implies that the total variation distance between the uniform distributions on $\text{conv} V$ and the input simplex is at most $2\epsilon''$. Over all random choices, this happens with probability at least $1 - 2\delta/5$. We set $\epsilon'' = \epsilon/2$. \blacksquare

Lemma 12 *Let S_{INPUT} be an n -dimensional isotropic simplex. Let Σ be an n -by- n positive definite matrix such that $\|\Sigma - I\| \leq \epsilon < 1/2$. Let μ be an n -dimensional vector such that $\|\mu\| \leq \epsilon$. Let B be an n -by- n matrix such that $\Sigma = BB^T$. Let S be the simplex $B^{-1}(S_{INPUT} - \mu)$. Then there exists an isotropic simplex S_{ISO} such that $d_{TV}(S, S_{ISO}) \leq 6n\epsilon$.*

Proof We use an argument along the lines of the orthogonal Procrustes problem (nearest orthogonal matrix to B^{-1} , already in (Nguyen and Regev, 2009, Proof of Theorem 4)): Let UDV^T be the singular value decomposition of B^{-1} . Let $R = UV^T$ be an orthogonal matrix (that approximates B^{-1}). Let $S_{ISO} = RS_{INPUT}$.

We have $S = UDV^T(S_{INPUT} - \mu)$. Let $\sigma_{min}, \sigma_{max}$ be the minimum and maximum singular values of D , respectively. This implies:

$$\begin{aligned} \sigma_{min}UV^T(S_{INPUT} - \mu) &\subseteq S \subseteq \sigma_{max}UV^T(S_{INPUT} - \mu), \\ \sigma_{min}(S_{ISO} - R\mu) &\subseteq S \subseteq \sigma_{max}(S_{ISO} - R\mu). \end{aligned} \tag{7}$$

As $S_{ISO} \supseteq B_n$, $\|\mu\| \leq 1$, R is orthogonal and S_{ISO} is convex, we have

$$S_{ISO} - R\mu \supseteq (1 - \|\mu\|)S_{ISO}.$$

Also,

$$\begin{aligned} S_{ISO} - R\mu &\subseteq S_{ISO} + \|\mu\|B_n \\ &\subseteq S_{ISO}(1 + \|\mu\|). \end{aligned}$$

This in (7) gives

$$\sigma_{min}(1 - \|\mu\|)S_{ISO} \subseteq S \subseteq \sigma_{max}(1 + \|\mu\|)S_{ISO}.$$

This and Lemma 7 imply

$$d_{TV}(S, S_{ISO}) \leq 2 \left(1 - \left(\frac{\sigma_{min}(1 - \|\mu\|)}{\sigma_{max}(1 + \|\mu\|)} \right)^n \right).$$

The estimate on Σ gives $\sigma_{min} \geq \sqrt{1 - \epsilon}$, $\sigma_{max} \leq \sqrt{1 + \epsilon}$. Thus

$$\begin{aligned} d_{TV}(S, S_{ISO}) &\leq 2 \left(1 - \left(\frac{1 - \epsilon}{1 + \epsilon} \right)^{3n/2} \right) \\ &\leq 2 \left(1 - (1 - \epsilon)^{3n} \right) \\ &\leq 6n\epsilon. \end{aligned}$$

■

Appendix F. Computing the moments of a simplex

In this section we present a formula for the moment over Δ^n . Similar more general formulas appear in [Lasserre and Avrachenkov \(2001\)](#). We will use the following result from [Grundmann and Moeller \(1978\)](#) for $\alpha_i \geq 0$:

$$\int_{\Omega^{n+1}} x_0^{\alpha_0} \cdots x_n^{\alpha_n} dx = \frac{\alpha_0! \cdots \alpha_n!}{(n + 1 + \sum_i \alpha_i)!}.$$

From the above we can easily derive a formula for integration over Δ^n :

$$\int_{\Delta^n} x_0^{\alpha_0} \cdots x_n^{\alpha_n} dx = \sqrt{n + 1} \cdot \frac{\alpha_0! \cdots \alpha_n!}{(n + \sum_i \alpha_i)!}.$$

Now

$$\begin{aligned} &\int_{\Delta^n} (x_0 u_0 + \cdots + x_n u_n)^k dx \\ &= \sum_{k_0 + \cdots + k_n = k} \binom{k}{k_0!, \dots, k_n!} u_0^{k_0} \cdots u_n^{k_n} \int_{\Delta^n} x_0^{k_0} \cdots x_n^{k_n} dx \\ &= \sum_{k_0 + \cdots + k_n = k} \binom{k}{k_0!, \dots, k_n!} u_0^{k_0} u_0^{k_0} \cdots u_n^{k_n} \frac{\sqrt{n + 1} \cdot k_0! \cdots k_n!}{(n + \sum_i k_i)!} \\ &= \frac{k! \sqrt{n + 1}}{(n + k)!} \sum_{k_0 + \cdots + k_n = k} u_0^{k_0} \cdots u_n^{k_n} \\ &= \frac{k! \sqrt{n + 1}}{(n + k)!} h_k(u). \end{aligned}$$

The variant of Newton's identities for the complete homogeneous symmetric polynomial gives the following relations which can also be verified easily by direct computation:

$$3h_3(u) = h_2(u)p_1(u) + h_1(u)p_2(u) + p_3(u),$$

$$2h_2(u) = h_1(u)p_1(u) + p_2(u) = p_1(u)^2 + p_2(u).$$

Divide the above integral by the volume of the standard simplex $|\Delta_n| = \sqrt{n+1}/n!$ to get the moment:

$$\begin{aligned} m_3(u) &= \frac{3!\sqrt{n+1}}{(n+3)!} h_3(u)/|\Delta_n| \\ &= \frac{2(h_2(u)p_1(u) + h_1(u)p_2(u) + p_3(u))}{(n+1)(n+2)(n+3)} \\ &= \frac{(p_1(u)^3 + 3p_1(u)p_2(u) + 2p_3(u))}{(n+1)(n+2)(n+3)}. \end{aligned}$$

Appendix G. Proofs of probabilistic results

Proof In the case where $p = 1$, Theorem 9 restricted to the positive orthant implies that for random vector $G = (G_1, \dots, G_n)$, if each G_i is an iid exponential random variable $\text{Exp}(1)$, then $(G/\|G\|_1, \|G\|_1)$ has the same (joint) distribution as (X, T) . Given the measurable function $f(x, t) = xt$, $f(X, T)$ has the same distribution as $f(G/\|G\|_1, \|G\|_1)$. That is, XT and G have the same distribution⁴.

For the second part, we know $TA(X) = A(TX)$ by linearity. By the previous argument the coordinates of TX are independent. This implies that $A(TX)$ has independent components. ■

Proof [of Lemma 5] It is enough to show the claim conditioning on the orthant in which G falls, and by symmetry it is enough to prove it for the positive orthant. Let random variable $H = (G_1^p, G_2^p, \dots, G_n^p)$. Since raising (strictly) positive numbers to the p th power is injective, it suffices to show that the random vector

$$X = \frac{H}{\sum_{i=1}^n H_i + W}$$

is independent of the random vector $Y = \sum_{i=1}^n H_i + W$.

First, let U be the interior of the support of (X, Y) , that is $U = \{x \in \mathbb{R}^n : x_i > 0, \sum_i x_i < 1\} \times \{y \in \mathbb{R} : y > 0\}$ and consider $h : U \rightarrow \mathbb{R}^n$ and $w : U \rightarrow \mathbb{R}$ where

$$h(x, y) = xy$$

and

$$w(x, y) = y - \sum_{i=1}^n h(x, y)_i = y - \sum_{i=1}^n x_i \cdot y = y \left(1 - \sum_{i=1}^n x_i\right).$$

The random vector (H, W) has a density $f_{H,W}$ supported on $V = \text{int } \mathbb{R}_+^{n+1}$ and

$$(x, y) \mapsto (h(x, y), w(x, y))$$

4. See (Gupta and Song, 1997, Theorem 1.1) for a similar argument in this context.

is one-to-one from U onto V . Let $J(x, y)$ be the determinant of its Jacobian. This Jacobian is

$$\begin{pmatrix} y & 0 & \cdots & 0 & x_1 \\ 0 & y & \cdots & 0 & x_2 \\ \vdots & & & & \vdots \\ 0 & 0 & \cdots & y & x_n \\ -y & -y & \cdots & -y & 1 - \sum_{i=1}^n x_i \end{pmatrix}$$

which, by adding each of the first n rows to the last row, reduces to

$$\begin{pmatrix} y & 0 & \cdots & 0 & x_1 \\ 0 & y & \cdots & 0 & x_2 \\ \vdots & & & & \vdots \\ 0 & 0 & \cdots & y & x_n \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix},$$

the determinant of which is trivially $J(x, y) = y^n$.

We have that $J(x, y)$ is nonzero in U . Thus, (X, Y) has density $f_{X,Y}$ supported on U given by

$$f_{X,Y}(x, y) = f_{H,W}(h(x, y), w(x, y)) \cdot |J(x)|.$$

It is easy to see⁵ that each $H_i = G_i^p$ has density $\text{Gamma}(1/p, 1)$ and thus $\sum_{i=1}^n H_i$ has density $\text{Gamma}(n/p, 1)$ by the additivity of the Gamma distribution. We then compute the joint density

$$\begin{aligned} f_{X,Y}(x, y) &= f_{H,W}(h(x, y), w(x, y)) \cdot y^n \\ &= f_{H,W}\left(xy, y\left(1 - \sum_{i=1}^n x_i\right)\right) \cdot y^n. \end{aligned}$$

Since W is independent of H ,

$$f_{X,Y}(x, y) = f_W\left(y\left(1 - \sum_{i=1}^n x_i\right)\right) \cdot y^n \prod_{i=1}^n f_{H_i}(x_i y)$$

where

$$\begin{aligned} \prod_{i=1}^n \left(f_{H_i}(x_i y)\right) \cdot f_W\left(y\left(1 - \sum_{i=1}^n x_i\right)\right) \cdot y^n &\propto \prod_{i=1}^n \left[e^{-x_i y} (x_i y)^{\frac{1}{p}-1}\right] \exp\left(-y\left(1 - \sum_{i=1}^n x_i\right)\right) y^n \\ &\propto \left(\prod_{i=1}^n x_i^{\frac{1}{p}-1}\right) y^{n/p}. \end{aligned}$$

The result follows. ■

5. See for example (Barthe et al., 2005, proof of Theorem 3).