

On the Computability of Continuous Maximum Entropy Distributions with Applications

Jonathan Leake
KTH

Nisheeth K. Vishnoi
Yale University

April 17, 2020

Abstract

We initiate a study of the following problem: Given a continuous domain Ω along with its convex hull \mathcal{K} , a point $A \in \mathcal{K}$ and a prior measure μ on Ω , find the probability density over Ω whose marginal is A and that minimizes the KL-divergence to μ . This framework gives rise to several extremal distributions that arise in mathematics, quantum mechanics, statistics, and theoretical computer science. Our technical contributions include a polynomial bound on the norm of the optimizer of the dual problem that holds in a very general setting and relies on a “balance” property of the measure μ on Ω , and exact algorithms for evaluating the dual and its gradient for several interesting settings of Ω and μ . Together, along with the ellipsoid method, these results imply polynomial-time algorithms to compute such KL-divergence minimizing distributions in several cases. Applications of our results include: 1) an optimization characterization of the Goemans-Williamson measure [15] that is used to round a positive semidefinite matrix to a vector, 2) the computability of the entropic barrier for polytopes studied by [7], and 3) a polynomial-time algorithm to compute the barycentric quantum entropy of a density matrix that was proposed as an alternative to von Neumann entropy in the 1970s [3, 32, 37]: this corresponds to the case when Ω is the set of rank one projections matrices and μ corresponds to the Haar measure on the unit sphere. Our techniques generalize to the setting of Hermitian rank k projections using the Harish-Chandra-Itzykson-Zuber formula [21, 24], and are applicable even beyond, to adjoint orbits of compact Lie groups.

Contents

1	Introduction	3
1.1	Our contributions	4
1.2	Applications	5
2	Preliminaries	6
3	The maximum entropy framework	9
4	Formal statement of our results	10
4.1	Mathematical and computational results	10
4.2	Applications	13
5	Technical overview	14
5.1	Proof overview: bounding box	15
5.2	Proof overview: counting oracle for \mathcal{P}_1 and \mathcal{V}_1	16
5.3	Proof overview: extending the counting oracle for \mathcal{P}_1 to \mathcal{P}_k	20
6	Bounding box	20
6.1	General bounding box	20
6.2	Rank- k projections	22
6.3	Convex bodies	22
7	Counting oracle for \mathcal{P}_k	23
7.1	Algorithm for $k = 1$	23
7.1.1	Evaluating the dual integral	24
7.1.2	Computing the gradient	28
7.2	Algorithm for $k > 1$	29
7.3	Sampling from \mathcal{P}_1	31
8	Computing maximum entropy measures	35
8.1	The ellipsoid framework	35
8.2	Rank- k Projections	36
9	The Goemans-Williamson measure	37
9.1	Goemans-Williamson measure on \mathcal{V}_1 maximizes entropy	37
9.2	Goemans-Williamson measure projected to the unit sphere does not maximize entropy	39
10	Generalization of the maximum entropy framework to Lie groups	40
A	The dual formulation and strong duality	44
A.1	The dual formulation	44
A.2	Strong duality under Slater's condition	45
A.3	Slater's condition holds for general Ω and μ	46
A.4	Slater's condition for \mathcal{P}_k and \mathcal{V}_1	47
B	The Schur-Horn theorem	48
C	Closeness of the approximate distribution	49

1 Introduction

Entropy maximizing distributions. Let Ω be a subset of \mathbb{R}^d and let $\mathcal{K} = \text{hull}(\Omega)$ denote the convex hull of Ω . Suppose one is given an $A \in \mathcal{K}$. A natural question arises: *Is there a canonical way to choose a probability measure supported on Ω that can be used to express A as a convex combination of points on Ω ?* When Ω is a discrete and finite set, this problem has been extensively studied and a canonical probability distribution was proposed by Jaynes [25, 26]: among all probability distributions that can be used to express A as a convex combination of points in Ω , pick the one that maximizes the Shannon entropy. These distributions are referred to as maximum entropy (max-entropy) distributions and arise in machine learning, statistics, mathematics, and theoretical computer science (TCS). In TCS, these distributions have found many uses due to duality, connections to polynomials, and algorithms to compute them [20, 36, 2, 14, 11, 1]; see [38].

In this paper we initiate a study of the computability when Ω is a continuous (and often nonconvex) manifold. Examples of interest include

$$\mathcal{V}_1 := \{vv^\top : v \in \mathbb{R}^n\},$$

$$\mathcal{P}_1 := \{vv^* : v \in \mathbb{C}^n, \|v\|_2 = 1\},$$

the set of rank k Hermitian projection matrices

$$\mathcal{P}_k := \{Y : Y \in \mathbb{C}^{n \times n}, \text{Tr}(Y) = k, Y = Y^*, Y^2 = Y\}$$

(related to the Grassmannian), or a convex body (in which case $\mathcal{K} = \Omega$).

Unlike the discrete setting, in the continuous setting the notion of finding a max-entropy distribution is not well-defined since a canonical notion of entropy does not necessarily exist. We instead consider relative entropy, Kullback-Leibler (KL) divergence with respect to a prior measure μ on Ω that corresponds to the density function $f(X) \equiv 1$ for all $X \in \Omega$. For all of the manifolds mentioned above, there is a canonical measure that has this property and is called the uniform measure; see Section 2. This leads us to the following infinite dimensional convex optimization problem which gives a canonical way to write A as a convex combination of points in Ω : Find a measure ν on Ω that is continuous with respect to μ and, subject to the constraint that the expected point in \mathcal{K} with respect to ν is A , ν minimizes the KL divergence to μ . Note that, by choice, ν is as close to the distribution μ as possible; hence we call it a maximum entropy distribution.

The class of extremal entropy maximizing distributions that arise in this manner have several properties that have led to their appearance, implicitly or explicitly, in several different areas:

- the work of Klartag (inspired by a work of Gromov) on the isotropic constant [28, 16],
- the work of Khatri and Mardia on the Matrix Bingham distribution in statistics with applications to various scientific and engineering problems [6, 27, 22],
- as shown here, the work of Goemans and Williamson on rounding semidefinite programs [15],
- the works of Güler, Bubeck and Eldan on barrier functions for interior point methods [18, 19, 7],
- the works of Band, Park, and Slater that defined the barycentric quantum entropy and proposed it as an alternative to the von Neumann entropy in the 1970s [3, 32, 37].

Computability of entropy maximizing distributions. One of the reasons why the entropy maximizing problem defined earlier is interesting (and unifies the above problems) is *duality*: the dual optimization problem roughly has the form:

$$\inf_Y \langle Y, A \rangle + \log \int_{X \in \Omega} e^{-\langle Y, X \rangle} d\mu(X),$$

where $\langle \cdot, \cdot \rangle$ is an inner product and μ is the given measure. If strong duality holds, it can be shown that the optimal distribution ν^* to the entropy maximizing problem can be described by the optimizer Y^* to the dual above: $\nu^*(X) \propto e^{-\langle Y^*, X \rangle}$ for $X \in \Omega$. As for computability of ν^* , Y^* lives in a small, convex, and finite dimensional (same dimension as \mathcal{K}) domain. Hence in principle, one could hope to represent ν^* efficiently. However, bounding the running time of a optimization method to find Y^* reduces to 1) a bounding some norm of Y^* and, 2) coming up with efficient algorithms to compute $\int_{X \in \Omega} e^{-\langle Y^*, X \rangle} d\mu(X)$ for matrices Y with that norm. These are the main problems studied in this paper.

1.1 Our contributions

The main contributions of this paper are to initiate a formal study of the computability of entropy maximizing distributions on continuous domains, to present an ellipsoid method-based framework to compute them, to derive polynomial time algorithms for computing maximum entropy distributions for specific manifolds mentioned earlier, and to present implications to some of the applications listed above.

The continuous maximum entropy framework and duality. Our general framework is presented Section 3. The focus is on the setting when the manifold Ω and the base measure μ is fixed to either the set of all rank one matrices over reals (\mathcal{V}_1) with the measure induced by Lebesgue measure on \mathbb{R}^n , or the set of all rank k projections over complexes (\mathcal{P}_k) for $k \geq 1$ with the appropriate Haar measure. The input consists of an element A (which is a matrix in the cases of interest) and the goal is to compute a representation for ν^* that is the KL-divergence minimizing distribution to μ with marginal A . We start by writing down the dual of this optimization problem (Section A.1) and showing that strong duality holds under Slater’s condition – that there is a density function that is strictly positive (and bounded) on Ω and has marginal A (Section A.2). This is implied by the condition that A is in the relative interior of the convex hull \mathcal{K} of Ω , which we then show is true quite generally in Sections A.3 and A.4. Strong duality then implies that the optimal measure ν^* is determined by the optimal dual solution Y^* as $\nu^*(X) \propto e^{-\langle Y^*, X \rangle}$; see Theorem 4.1.

Norm of the optimal dual solution. However, to solve the dual convex program one needs, at the bare minimum, that the norm of Y^* is reasonably bounded. It is not difficult to see that as A tends to the boundary of \mathcal{K} , the optimal measure is concentrated on a face of \mathcal{K} implying that the norm of Y^* must tend to infinity. Thus, one needs some assumption on the “interiority” of A to ensure polynomial time computability. The situation is exacerbated by the fact that the Y^* appears in the exponent and, hence, to have any hope of computability of the entropy maximizing distribution, the bound on Y^* should be *polynomial* in the bit complexity of A . Unlike the case when Ω is discrete (studied in [36]), the fact that the base measure μ is continuous makes it harder. Our main contribution towards the problem of bounding the norm of Y^* involves identifying a certain “balance” property of the measure μ on the manifold Ω (Definition 4.1) and showing that, roughly, $\|Y^*\| \leq \text{poly}(d, 1/\eta)$ where η is the distance of A from the boundary of \mathcal{K} ; see Theorem 4.2.

We show that this balance property holds for a wide class of manifolds and obtain as corollaries a bound of $\text{poly}(n, 1/\eta)$ for both $\Omega = \mathcal{P}_k$ (Corollary 6.4) and when Ω is an n -dimensional convex body (Corollary 6.6). This bounding box result is quite general and expected to find further applications.

Computing the integral in the dual for matrix manifolds. A bound on the norm of Y^* allows us to show that we can use the ellipsoid method to solve the dual convex program, provided the measure μ is balanced on Ω , and we can evaluate the dual and its gradient at a specified Y of norm up to that of Y^* . The tasks of evaluating the dual and its gradient essentially reduce to the computation of the integral $\int_{X \in \Omega} e^{-\langle Y, X \rangle} d\mu(X)$. In the case when $\Omega = \mathcal{V}_1$ with μ being the measure induced by the Lebesgue measure, we observe that the dual optimization problem is finite only when $Y \succ 0$, and thus we need to evaluate the integral only for such a Y . The integral above then turns out to have a simple formula: roughly, $\log \det Y$ (Proposition 9.3).

In other interesting cases, computing such an integral turns out to be a nontrivial task. In the case when $\Omega = \mathcal{P}_1$ and μ_1 is the uniform measure induced by the Haar measure on the complex unit sphere, we first note that the entropy maximizing measure cannot be obtained by solving the problem first for \mathcal{V}_1 and then “projecting” it on the sphere; see Section 9.2. Then, we note that the integral does not reduce to a product of n integrals as in the Lebesgue case, and there is no easy way around this. We need an algorithm to integrate the density $e^{-v^* Y v}$ over the complex unit sphere where the only thing we know about Y is that it is Hermitian. Neither the density is log-concave, nor the support (unit sphere) is convex. Our main contribution here is to give an *exact* algorithm to compute this integral whose running time depends *single exponentially* on the bit complexity of the input Y to it (Theorem 4.5). As remarked earlier, because Y is being exponentiated, this is the best one can hope for and also turns out to be sufficient to obtain polynomial time algorithms for computing maximum entropy distributions on \mathcal{P}_1 .

Interestingly, the algorithm to compute this integral and its proof relies on a connection between the manifold \mathcal{P}_1 and the probability simplex in n dimensions. Specifically, one can naturally push forward the entropy maximizing measure from \mathcal{P}_1 to a log-linear measure on the corresponding simplex. There are then algorithms to sample from such a density function on the simplex to estimate such an integral; however, to obtain an $1 + \delta$ approximation to it, the running time of these methods depends polynomially on $1/\delta$ instead $\log 1/\delta$. We give an exact algorithm to compute this integral. Our method relies on Laplace transforms, is elementary, and a significant effort is needed to deal with the case when Y has repeated eigenvalues. Importantly, this viewpoint also leads us to an exact algorithm for computing such an integral for \mathcal{P}_k for $k > 1$ using the Harish-Chandra-Itzykson-Zuber formula [21, 24, 13, 39]; see Theorem 7.7.

Efficient algorithm via the ellipsoid method. Our general ellipsoid method-based algorithm requires 1) a full dimensional embedding of $\text{hull}(\Omega)$ in a d -dimensional real Hilbert space, 2) μ is a balanced measure on Ω , 3) Ω is contained in a ball of radius r , 4) the point A is in the η -interior of $\text{hull}(\Omega)$ and, 5) that we have an exact counting/integrating oracle. It runs in time polynomial in $d, 1/\eta, \log r$ and $\log 1/\varepsilon$, to solve the dual problem to an additive ε ; see Theorem 4.4. Our bound on the norm of Y^* and exact algorithms to compute the dual objective/gradient for the case of \mathcal{P}_k imply a polynomial time algorithm to compute the entropy maximizing measure in this case when A is in the polynomial interior of $\text{hull}(\mathcal{P}_k)$; see Corollary 4.9.

1.2 Applications

SDP rounding. One approach to semi-definite programming (SDP) based approximation algorithms, starting with the work of Goemans-Williamson [15] for the maximum cut problem, is SDP

rounding. Here, typically, A is a positive semi-definite (PSD) matrix, that is computed using a SDP relaxation to some non-convex problem, and one of the goals is to *round* A to a vector. This involves choosing a distribution on the set \mathcal{V}_1 defined above, and typical choices have been somewhat magical and lack an explanation. In the Goemans-Williamson setting, A is an $n \times n$ PSD matrix, and the density ν on Ω they choose to express A as a convex combination is as follows: pick a vector $v \in \mathbb{R}^n$ from the normal distribution with covariance matrix A . We show that this distribution is the maximum entropy distribution ν^* (corresponding to A) on \mathcal{V}_1 with base measure induced by the Lebesgue measure on \mathbb{R}^n , thus giving an optimization characterization of this measure; see Corollary 4.12. The proof relies on strong duality and a closed form expression for the dual objective integral on \mathcal{V}_1 ; see Theorem 4.1.

Quantum entropy. In quantum mechanics, a density matrix ρ is a trace one complex $n \times n$ PSD matrix and describes the statistical state of a system. The extreme points in the set of density matrices are the pure states or \mathcal{P}_1 . von Neumann defined a notion of entropy [40] of ρ that is computed by first writing ρ as a convex combination $\sum_{i=1}^n \lambda_i u_i u_i^*$, where $\{u_i\}_{i \in [n]}$ is an orthonormal basis for \mathbb{C}^n , and then computing the negative Shannon entropy of the λ_i 's. While the von Neumann entropy is a mathematically elegant notion, it was vigorously argued in the 1970s that it does not capture the uncertainty in ρ [3, 32, 37]. In fact, von Neumann's way to write ρ as a convex combination of pure states can be viewed as “the most terse”, or entropy minimizing one. In the same papers, an alternative way to define entropy of a density matrix was suggested – as the entropy of the entropy maximizing distribution with marginal ρ – and referred to as the barycentric quantum entropy. Unlike the von Neumann entropy, that has a simple formula ($-\text{Tr } \rho \log \rho$), the barycentric entropy did not have an efficient algorithm that could compute it. Our algorithm to compute entropy maximizing distributions for \mathcal{P}_1 mentioned above directly implies a polynomial time algorithm to compute the barycentric entropy of a density matrix (that is sufficiently in the interior) along with the probability density that achieves it; see Corollary 4.11.

Entropic barrier function. Bubeck and Eldan in [7] proved that the entropic barrier of a convex body $K \subseteq \mathbb{R}^d$ is a $(1+o(1))n$ -self-concordant barrier on K . Roughly speaking, this barrier function, for a point in K is defined to be the optimal value of a dual maximum entropy optimization problem when $\Omega = K$ and the measure is the Lebesgue measure on K . The computability of this barrier function for a point K is not known in general. One obstacle is to get a reasonable bound on the norm of the optimal dual solution. An almost direct consequence of Theorem 4.2 implies such a bound for points that are sufficiently in the interior of K ; see Corollary 6.6.

2 Preliminaries

Notation. Let $\mathbb{C}, \mathbb{R}, \mathbb{R}_+, \mathbb{N}$ denote the complex, real, nonnegative real, and natural numbers respectively. For $k, n \in \mathbb{N}$, let $\mathbb{C}^{k \times n}$ and $\mathbb{R}^{k \times n}$ denote the sets of $k \times n$ complex and real matrices respectively. A matrix $M \in \mathbb{C}^{n \times n}$ is said to be Hermitian if $A = A^*$ where $*$ denotes the conjugate transpose. A Hermitian matrix M is said to be PD (positive definite) and PSD (positive semidefinite) if its eigenvalues are positive and nonnegative respectively. For an $n \times n$ matrix X , we define $\text{diag}(X)$ to be the length- n vector of the diagonal entries of X . If x is a vector, then we define $\text{diag}(x)$ to be the diagonal matrix with entries the entries of x . For any $k, n \in \mathbb{N}$, we equip the vector space $\mathbb{C}^{k \times n}$ with the Frobenius inner product $\langle Y, Z \rangle := \text{Tr}(YZ^*)$. We also denote $\|Y\| := \sqrt{\langle Y, Y \rangle}$. Note that $\langle Y, Z \rangle \in \mathbb{R}$ whenever Y, Z are Hermitian, so that the set of $n \times n$ Hermitian matrices is a real Hilbert space of dimension n^2 . Also $\langle Y, Z \rangle \geq 0$ whenever Y, Z are

PSD. We further let $B_\varepsilon(Y)$ denote the open ε -ball centered at Y in the space in which Y lives (e.g., the $n \times n$ Hermitian matrices). Finally, we let $\text{hull}(S)$ denote the convex hull of a set S in some ambient vector space.

Manifolds. In general, we let Ω be any smooth manifold that is embedded in a d -dimensional real Hilbert space V with inner product $\langle \cdot, \cdot \rangle$. Let $\mathcal{L}(X) = B$ denote the affine space in which $\text{hull}(\Omega)$ is full dimensional, i.e., every element $X \in \text{hull}(\Omega)$ satisfies the equation $\mathcal{L}(X) = B$. The concrete manifolds we consider are collections of matrices with some structure. In particular, for fixed $n \in \mathbb{N}$, consider the following manifold within $\mathbb{C}^{n \times n}$. An $n \times n$ *rank- k PSD projection* is a PSD matrix with k eigenvalues equal to 1 and the rest equal to 0.

$$\mathcal{P}_k = \mathcal{P}_k(n) := \{n \times n \text{ rank-}k \text{ PSD projections}\}.$$

Note that \mathcal{P}_k is also a manifold within the space of $n \times n$ Hermitian matrices.¹ Other manifolds we consider are the complex unit sphere $S_{\mathbb{C}}^n \subset \mathbb{C}^n$ (which is related to \mathcal{P}_1), the manifold of all rank one matrices (not necessarily trace one): $\mathcal{V}_1 := \{vv^\top : v \in \mathbb{R}^n\}$, and a convex body $K \subset \mathbb{R}^n$.

We would also like to consider the convex hull of a given manifold Ω . To make sense of such a notion, we need to consider the manifold as being embedded in some ambient vector space. This ambient space often the space of $n \times n$ Hermitian matrices in our examples. In general, we refer to the elements of $\text{hull}(\Omega)$ as *marginals* or *marginals matrices*.

Group actions. It is useful to understand the symmetries of some of the manifolds mentioned above in terms of groups that act on them. Recall that an $n \times n$ unitary matrix is an invertible matrix U for which $U^{-1} = U^*$, and an $n \times n$ orthogonal matrix is an invertible matrix O for which $O^{-1} = O^\top$. The unitary and orthogonal groups ($U(n)$ and $O(n)$) act on the manifolds discussed above as follows:

- $U(n)$ acts on column vectors in $S_{\mathbb{C}}^n$ and on $\text{hull}(S_{\mathbb{C}}^n)$ by left multiplication.
- $U(n)$ acts on \mathcal{P}_k and on $\text{hull}(\mathcal{P}_k)$ by conjugation.
- $O(n)$ acts on \mathcal{V}_1 and on $\text{hull}(\mathcal{V}_1)$ by conjugation.

Note that the actions of $U(n)$ on $S_{\mathbb{C}}^n$ and on \mathcal{P}_1 are compatible in the sense that for $x \in S_{\mathbb{C}}^n$ and $U \in U(n)$, we have $(Ux)(Ux)^* = U(xx^*)U^*$ where $xx^* \in \mathcal{P}_1$.

Relative interior. The convex set $\text{hull}(\Omega)$ is not necessarily full dimensional in the ambient Hilbert space. To define a notion of interior for $\text{hull}(\Omega)$, we restrict to the minimal affine subspace in which Ω lives (this is given explicitly by $\mathcal{L}(X) = B$ discussed above). More generally, we make the following definition.

Definition 2.1 (Relative interior) Fix a convex subset S in a vector space V , and let $V_{\mathcal{L},B}$ be the minimal affine subspace in which S lives. We say that $Y \in V$ is in the η -interior of S (for $\eta > 0$) if

$$B_\eta(Y) \cap V_{\mathcal{L},B} \subseteq S.$$

We say that Y is in the interior of S if there exists $\eta > 0$ such that Y is in the η -interior of S .

Here we usually consider $S = \mathcal{P}_k(n)$, and we will be interested in the case where $\eta \geq \frac{1}{\text{poly}(n)}$.

¹Note that \mathcal{P}_k is homeomorphic to a Grassmannian, i.e., the manifold of k -dimensional subspaces within an n -dimensional space. The homeomorphism is explicitly given as the map which sends a rank- k PSD projection to the k -dimensional subspace given by its image.

Measures and densities. Often, the manifolds Ω we consider have some geometric structure (e.g., it is a manifold with a group action), and we want to consider measures which interact nicely with this structure. To make sure this happens, we restrict to the class of measures which are given by continuous density functions on Ω . To make sense of this, we need a natural base measure μ on Ω which corresponds to the density function $f(X) \equiv 1$. (E.g., in the case of $\Omega = \mathbb{C}^n$ or $\Omega = \mathbb{R}^n$, the Lebesgue measure often plays this role.) In particular, the support of μ should be equal to Ω .

In the case of $\Omega = \mathcal{P}_k$, there is a canonical measure which is appropriately called the uniform measure: we define μ_k be the unique unitarily invariant measure on \mathcal{P}_k , where $U(n)$ acts by conjugation (as discussed above). Hence, equivalently (and more formally), we restrict to the class of measures on \mathcal{P}_k which are absolutely continuous with respect to μ_k . We prove here the existence of μ_k , a classical result.

Proposition 2.1 (Existence of μ_k) *There exists a distribution μ_k on \mathcal{P}_k (which we call the uniform distribution). If X is a random variable distributed according to μ_k , then X and UXU^* have the same distribution for any unitary U .*

Proof: Pick random complex unit vectors $v_i \in \mathbb{C}^{n+1-i}$ for $i \in [k]$. Note that $v_1 \in \mathbb{C}^n$. Now, map v_2 into $v_1^\perp \cong \mathbb{C}^{n-1}$, map v_3 into $\{v_1, v_2\}^\perp \cong \mathbb{C}^{n-2}$, etc. to obtain a collection of k orthogonal vectors in \mathbb{C}^n . Form an $k \times n$ matrix P by letting the v_i be the rows of P . Defining $X := P^*P \in \mathcal{P}_k$ gives a distribution μ_k on \mathcal{P}_k .

For unitary invariance, note that this property holds for the choice of v_1 by construction. This can then be inductively applied to v_2, \dots, v_k by composing the given unitary with the appropriate projection. ■

We also consider the standard Lebesgue measure on \mathbb{R}^n for convex bodies and its pushforward measure μ through the map $v \mapsto vv^\top$ on \mathcal{V}_1 . Note that $S_{\mathbb{C}}^n$ also has a canonical unitarily invariant measure, usually called the Haar measure. The pushforward of this measure through the map $v \mapsto vv^*$ yields the unitarily invariant measure μ_1 on \mathcal{P}_1 .

Integration/Counting oracle. We are interested in computing the following exponential integral for a given Y in our Hilbert space V .

Definition 2.2 (Exponential integrals) *Fix $n \in \mathbb{N}$ and let μ be a measure with support Ω , a manifold embedded in the real Hilbert space V . We define the following function on an input $Y \in V$:*

$$\mathcal{E}(Y) = \mathcal{E}_\mu(Y) := \log \int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X).$$

Whenever $\mu = \mu_k$ and $\Omega = \mathcal{P}_k$, we use the following shorthand notation $\mathcal{E}_k(Y)$. We sometimes also refer to these integrals as exponential integrals.

A strong integration/counting oracle for Ω and μ outputs two quantities, given an element Y from the ambient Hilbert space V of Ω :

1. $\mathcal{E}_\mu(Y)$
2. the matrix $\nabla \mathcal{E}_\mu(Y)$, defined so that the following holds for any $Z \in V$:

$$\langle \nabla \mathcal{E}_\mu(Y), Z \rangle = \frac{d}{dt} \mathcal{E}_\mu(Y + tZ) \Big|_{t=0}.$$

Primal	Dual
$\sup_{\nu} \left[- \int_{\Omega} \nu(X) \log(\nu(X)) d\mu(X) \right]$	$\inf_Y F_A(Y) = \inf_Y \left[\langle Y, A \rangle + \log \int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X) \right]$
subject to:	subject to:
$\nu : \Omega \rightarrow \mathbb{R}_{\geq 0}, \mu\text{-measurable}$	$\mathcal{L}(Y) = 0$
$\int_{\Omega} X \nu(X) d\mu(X) = A$	
$\int_{\Omega} \nu(X) d\mu(X) = 1$	

Figure 1: Primal and dual maximum entropy convex programs for A in the interior of $\text{hull}(\Omega)$.

In the case of $\Omega = \mathcal{P}_k$, Y and Z are Hermitian. Further, since the measure μ_k is unitarily invariant, we can assume that Y is diagonal and expect the running time of the counting oracle should depend polynomially on n and the number of bits needed to represent e^{-y_i} for any i , where y_1, \dots, y_n are the eigenvalues (diagonal elements) of Y .

As we will show, in the special case when $\Omega = \mathcal{V}_1$ and μ is the pushforward of the Lebesgue measure, we can compute the integral $\mathcal{E}_\mu(Y)$ exactly in time polynomial in the bit complexity of Y due to a direct formula. This happens because the measure μ is a product measure, which is not the case for μ_k .

3 The maximum entropy framework

In this section we present our maximum entropy convex program. Fix a manifold Ω in a d -dimensional real Hilbert space with inner product $\langle \cdot, \cdot \rangle$, and let $\mathcal{L}(X) = B$ denote the corresponding affine space containing Ω . Let μ be the base measure on Ω and A in $\mathcal{K} := \text{hull}(\Omega)$. Our goal is to find a density function ν with marginal A that minimizes the KL-divergence with respect to μ .

We use the shorthand $\text{Prim}_\mu(A)$ (or $\text{Prim}_k(A)$ if $\mu = \mu_k$) to refer to this primal optimization program. We mainly consider the case of $\mu = \mu_k$ and $\Omega = \mathcal{P}_k$ or $\Omega = \mathcal{V}_1$ with μ the pushforward of Lebesgue measure. In these cases Y will come from some subspace of the $n \times n$ Hermitian matrices. Drawing from the intuition that these base measures are uniform over the manifold, and hence in some sense maximize entropy, we say the KL-divergence minimizing measure is entropy maximizing. However, we note that this framework is also applicable to other base measures, in particular to the case when Ω is a convex body in \mathbb{R}^d and μ is the Lebesgue measure. The fact that the entropy integral (without the minus sign) is convex as a function of the density ν follows from the fact that this integral is precisely the KL divergence between the probability distribution corresponding to ν and the distribution μ . Convexity of the KL divergence for probability distributions is then a well-known fact.

Efficiently solving this convex program directly is a priori impossible as the support of ν is infinite. To find a succinct representation for the optimal ν^* , we turn to the dual program (see Section A.1 for a derivation), which gives us a nice representation of the max-entropy density function ν^* . We often use the shorthand $\text{Dual}_\mu(A)$ (or $\text{Dual}_k(A)$ if $\mu = \mu_k$) to refer to this program.

In the case of \mathcal{P}_k with uniform measure μ_k , the optimal solution to $\text{Dual}_k(A)$ is given by a

Hermitian matrix Y^* . By strong duality (see Theorem 4.1), this in turn shows that the max-entropy density function ν^* takes on a nice form:

$$\nu^*(X) \propto e^{-\langle Y^*, X \rangle}.$$

As a note, in the case of $\Omega = \mathcal{P}_k$ this matrix Y^* is only unique up to a shift by a multiple of the identity matrix. Issues arising from non-uniqueness can be handled by restricting to the minimal affine subspace in which $\text{hull}(\mathcal{P}_k)$ lives, as referred to in the discussion surrounding Definition 2.1. However, as A tends to the boundary of $\text{hull}(\Omega)$, Y^* can be seen to tend to infinity as the support of the measure ν^* tends to lower dimensions.

4 Formal statement of our results

4.1 Mathematical and computational results

Our first result shows that strong duality holds.

Theorem 4.1 (Strong duality) *Let Ω be a manifold that is embedded in a d -dimensional real Hilbert space with an inner product $\langle \cdot, \cdot \rangle$, and let μ be a measure supported on Ω . For any A in the relative interior of the convex hull of Ω , the optimal values of the primal and dual objective functions coincide, and the corresponding max-entropy distribution has density function of the following form for some Y^* :*

$$\nu^*(X) \propto e^{-\langle Y^*, X \rangle}.$$

The proof of this result uses standard techniques and appears in the appendix (Sections A.2 and A.3). This result applied to \mathcal{P}_k and μ_k shows that optimizing $\text{Dual}_k(A)$ is in fact equivalent to optimizing $\text{Prim}_k(A)$, and therefore the max-entropy measure has the exponential form described above.

With strong duality in hand, we focus on the computability of the optimal matrix Y^* for the dual program. To do this we use a version of the ellipsoid algorithm (see Theorem 8.1 and the algorithm that follows), for which we need two things.

First, we need an upper bound on some norm of the dual optimal solution. If Y^* is the optimal solution, then the number of iterations of the ellipsoid algorithm depends on $\log \|Y^*\|$. That said, it may seem that a bound depending on $e^{1/\eta}$, where η is such that $B_\eta(A) \subset \text{hull}(\Omega)$, is enough to achieve polynomial dependence on $\frac{1}{\eta}$. However, this is not enough, since the integral appearing in the dual is polynomially dependent on the number of bits needs to represent e^{-y_i} , where the y_i 's are the entries or eigenvalues of a given input Y . Hence, we actually need polynomial dependence on $\frac{1}{\eta}$, which is achieved in our bounding box result below. Note that this issue is not surprising, as it crops up in exactly the same way in the discrete maximum entropy case (see [36]).

We give here a bounding box result which is more general than we need for the rank- k projections case ($\Omega = \mathcal{P}_k$ and $\mu = \mu_k$). It relies on a key “balance” property of the measures. This notion extends important properties of the discrete uniform measure to continuous measures on manifolds and is one of the key notions we introduce.

Definition 4.1 (Balanced measure) *A measure μ is said to be balanced if for any $\delta > 0$ and $X \in \Omega$, we have that at least $\exp(-\text{poly}(\delta^{-1}, d))$ of the mass of μ is contained in the δ -ball about X (where d is the dimension of the ambient space in which $\text{hull}(\Omega)$ lives).*

We see in Definition 6.2 how this notion can be used to give a more refined notion of interior (beyond the η parameter discussed above). Conceptually, it allows us to give an measure-theoretic relaxation of the notion of a separating hyperplane.

Theorem 4.2 (Bounding box) *Let μ be a measure supported on a manifold Ω embedded in a d -dimensional real Hilbert space. Suppose that μ is balanced, in the sense of Definition 4.1. Further, let A be an element of the η -interior of the convex hull of Ω . Then there is an optimal solution Y^* to the dual program such that:*

$$\|Y^*\| \leq \text{poly}(\eta^{-1}, d).$$

Corollary 6.4 and Corollary 6.6 give bounds for rank- k projections and convex bodies as corollaries.

Remark 4.3 *Our bounding box result significantly generalizes the discrete case (Theorem 2.7 in [36]). Uniform distribution in the discrete case has atoms of uniformly strictly positive (at worst singly-exponentially small) mass at all points, and this implies a bound on optimal dual solutions. In the continuous case this is no longer true, the notion of balance then fills the gap.*

Second, at each step of the ellipsoid algorithm, we need to be able to evaluate the dual objective function and its gradient at given input Y . The hardest part of such a computation comes in evaluating \mathcal{E}_μ , the exponential integral portion of the objective function. We show that if we have access to such an evaluation oracle, then under very general conditions, we can compute the maximum entropy distribution.

Theorem 4.4 (Ellipsoid method-based general algorithm) *Let μ be a balanced measure with support on a manifold Ω embedded in a d -dimensional real Hilbert space. Let the affine space in which Ω lies, $\mathcal{L}(X) = B$, be given as input (\mathcal{L}, B) . Assume that Ω is contained in a ball of radius r . There exists an algorithm that, given A in the η -interior of $\text{hull}(\Omega)$, any $\varepsilon > 0$, and a strong counting/integration oracle for the exponential integral $\mathcal{E}_\mu(Y)$, returns Y° such that*

$$F_A(Y^\circ) \leq F_A(Y^*) + \varepsilon$$

where F_A is the objective function for the dual program $\text{Dual}_\mu(A)$, and Y^* is an optimum of the dual program. The running time of the algorithm is polynomial in d , η^{-1} , $\log(\varepsilon^{-1})$, $\log(r)$, and the number of bits needed to represent A , \mathcal{L} , and B .

Our next result says that in fact we have an efficient strong counting oracle for \mathcal{E}_k on the domain \mathcal{P}_k with measure μ_k .

Theorem 4.5 (Counting oracle) *There is an algorithm that, given $n \in \mathbb{N}$, $k \in [n]$, an $n \times n$ real diagonal matrix $Y = \text{diag}(y)$, and a $\delta > 0$, returns numbers \bar{E}, \bar{G} such that*

1. $|\bar{E} - \mathcal{E}_k(Y)| \leq \delta$
2. $|\bar{G} - \nabla \mathcal{E}_k(Y)| \leq \delta$,

where \mathcal{E}_k is the exponential integral defined above (and in Definition 2.2). The running time of the algorithm is polynomial in n , $\log(\frac{1}{\delta})$, and the number of bits needed to represent e^{-y_i} for any $i \in [n]$.

The proof of this theorem for $k = 1$ is elementary but relies on the interesting connection between the complex unit sphere and the probability simplex. This connection also yields an exact sampling algorithm; see Proposition 7.10. For $k > 1$, the proof of the theorem above relies on the Harish-Chandra-Itzykson-Zuber formula [21], [24]; see Theorem 7.7.

Remark 4.6 In the case of \mathcal{V}_1 with the pushforward of Lebesgue measure, there is an exact formula to compute the corresponding dual optimum for positive definite marginals A : $Y^* = \frac{1}{2}A^{-1}$; see Corollary 9.4. Positive-definiteness of the input Y is in fact required for the dual objective to be finite, which is in stark contrast with the \mathcal{P}_k case where any Hermitian matrix is allowed. These points suggest a conceptual divide between the Lebesgue measure case and the rank- k projections case. We do not expect such a formula for Y^* in the case of \mathcal{P}_1 and, indeed, the lack of one has been one of the obstacles for efficient algorithms for quantum barycentric entropy and computing the normalizing constant of the matrix Bingham distribution.

Remark 4.7 In this paper we primarily consider the best possible setting where the running time of the counting oracle depends logarithmically on the accuracy. We refer to such counting oracles as exact. We note that our framework does allow for counting oracles where the dependence is polynomially in $1/\delta$.

Remark 4.8 Guler in [17] studies the characteristic function of a convex cone. In our language, the characteristic function of a cone is the exponential integral $\mathcal{E}_K(y)$ with respect to the Lebesgue measure on the dual cone K :

$$\mathcal{E}_K(y) = \log \int_K e^{-\langle y, x \rangle} dx.$$

For the case of homogeneous convex cones, Guler gives a nice way to construct explicit formulas for the characteristic function. (A homogeneous cone is a cone K such that for all $u, v \in K$ there is a linear isomorphism of K which maps u to v . Orthants, Lorentz cones, and semidefinite cones are all homogeneous. See Sections 3 and 7 of [17] for more details.) Given a fixed vector e in the interior of K , any other vector y in the interior of K , and an automorphism A_y of K mapping e to y , the dual objective for K can be written up to additive constant as:

$$F_\theta(y) = \langle y, \theta \rangle - \log \int_K e^{-\langle y, x \rangle} dx = \langle y, \theta \rangle - \frac{1}{2} \log(\det(A_y A_y^\top)).$$

Such an explicit formula gives a route to efficiently computing the dual objective function in this case.

The bounding box and counting oracle for μ_k and \mathcal{P}_k then imply that the ellipsoid method-based algorithm from Theorem 4.4 gives a polynomial time algorithm for approximately computing Y^* , the optimum of the program $\text{Dual}_k(A)$.

Corollary 4.9 (Ellipsoid method-based efficient algorithm for \mathcal{P}_k) There exists an algorithm that, given $n \in \mathbb{N}$, $k \in [n]$, a trace- k PD matrix A in the η -interior of the convex hull of the set of $n \times n$ rank- k PSD projection matrices (i.e., $\text{hull}(\mathcal{P}_k)$), and an $\varepsilon > 0$, returns a Hermitian matrix Y° such that

$$F_A(Y^\circ) \leq F_A(Y^*) + \varepsilon,$$

where F_A is the dual objective function and Y^* is an optimal solution to the dual program $\text{Dual}_k(A)$. The running time of the algorithm is polynomial in n , $\frac{1}{\eta}$, $\log(\frac{1}{\varepsilon})$, and the number of bits needed to represent A .

We further discuss the closeness of the distributions associated to Y° and Y^* from the previous Corollary in Appendix C.

Remark 4.10 Notice that the dependence on $\frac{1}{\eta}$ means that we do not achieve a polynomial time algorithm for A near the boundary of $\text{hull}(\mathcal{P}_k)$. This dependence comes from the fact that the bounding box (Theorem 4.2) is dependent on $\frac{1}{\eta}$. One may then naturally ask whether this bounding box dependence can be improved. It turns out that it cannot in this case, see Remark 6.5. Note that this differs from the discrete case, where in [38] the authors are able to remove this $\frac{1}{\eta}$ dependence under certain assumptions on the polytope.

4.2 Applications

Barycentric quantum entropy. In [37], Slater discusses the notion of barycentric quantum entropy of a density matrix, and compares it to that of von Neumann entropy. His investigation of this notion was prompted by the work of Band and Park [3, 32], who critiqued the use of von Neumann entropy as a good indicator of the uncertainty of the given density matrix. In particular, they argue that a better notion of entropy would relate to distributions on all possible pure states, whereas the von Neumann entropy is derived from the discrete distribution on the pure states corresponding to eigenvectors of the matrix. In response to this, Slater defines a notion of quantum entropy in terms of a max-entropy program on the set of all pure states. He then goes on to show how one might determine the quantum entropy in a few specific cases.

Definition 4.2 (Barycentric quantum entropy) Let ρ be an $n \times n$ Hermitian density matrix (trace-1, positive semidefinite). Then the barycentric quantum entropy of ρ is defined (in our notation) as: $H_b(\rho) := \inf_{\nu} \int_{\mathcal{P}_1} \nu(X) \log(\nu(X)) d\mu_1(X)$ subject to

$$\nu(X) \geq 0 \quad \forall X \in \mathcal{P}_1 \quad \text{and} \quad \int_{\mathcal{P}_1} \nu(X) d\mu_1(X) = 1 \quad \text{and} \quad \int_{\mathcal{P}_1} X \nu(X) d\mu_1(X) = \rho,$$

where \mathcal{P}_1 denotes the set of pure states and μ_1 denotes the unitarily invariant measure on \mathcal{P}_1 .

Our results for computing max-entropy measures on \mathcal{P}_1 immediately imply efficient computability of the barycentric quantum entropy for density matrices that are polynomially in the interior.

Corollary 4.11 (Computability of barycentric quantum entropy) There exists an algorithm that, given a Hermitian density matrix ρ in the η -interior of the set of Hermitian density matrices and an $\varepsilon > 0$, returns a number \bar{H} such that $|\bar{H} - H_b(\rho)| < \varepsilon$. The running time of the algorithm is polynomial in n , $\frac{1}{\eta}$, $\log(\frac{1}{\varepsilon})$, and the number of bits needed to represent ρ .

Goemans-Williamson SDP rounding. In their seminal paper, Goemans-Williamson [15] gave a rounding scheme that gives a way to round a given PD matrix A to a vector. Their method goes by drawing a vector v from a particular distribution on \mathbb{R}^n based on the matrix A .

Definition 4.3 (Goemans-Williamson measure) Given $n \in \mathbb{N}$ and a real positive definite $n \times n$ matrix A , the Goemans-Williamson measure μ_{GW} can be defined via a sampling process on \mathbb{R}^n as follows.

1. Sample $g \in \mathbb{R}^n$ from the standard multivariate Gaussian distribution.
2. Compute $v := Vg$ where V is a square root of A , i.e., $A = VV^\top$.
3. v is a sample from μ_{GW} .

It is then straightforward to compute the marginals matrix associated to this distribution as follows:

$$\mathbb{E}[vv^\top] = \int_{\mathbb{R}^n} (vv^\top) d\mu_{\text{GW}}(v) = V \left[\int_{\mathbb{R}^n} gg^\top dg \right] V^\top = VV^\top = A.$$

Thus, if we map \mathbb{R}^n to \mathcal{V}_1 via $v \mapsto vv^\top$ and also pushforward the Lebesgue measure through this map, the above is precisely the marginal constraint in our max-entropy framework. This observation implies that the pushforward of the measure μ_{GW} is a (strictly) feasible solution to the max-entropy primal program on the domain \mathcal{V}_1 with the pushforward of the Lebesgue measure. We show that it is also the optimal solution to the max-entropy program.

Corollary 4.12 (Goemans-Williamson measure maximizes entropy) *For any positive definite matrix A , let μ_{GW} be the measure corresponding to the Goemans-Williamson rounding scheme for A . Then the pushforward of μ_{GW} to \mathcal{V}_1 is the max-entropy measure with marginals A on \mathcal{V}_1 with respect to the pushforward of Lebesgue measure.*

Entropic barrier function. Bubeck and Eldan in [7] prove that the entropic barrier of a convex body $K \subseteq \mathbb{R}^d$ is a $(1 + o(1))n$ -self-concordant barrier on K , improving a seminal result of Nesterov and Nemirovski [30]. In fact this gives the first explicit construction of a universal barrier for convex bodies with optimal self-concordance parameter.

Definition 4.4 (Entropic barrier) *Given a convex body $K \subseteq \mathbb{R}^d$, define the entropic barrier for K as the real-valued function on the interior of K defined as:*

$$B_K(v) := \sup_{y \in \mathbb{R}^d} \left[\langle y, v \rangle - \log \int_K e^{\langle y, x \rangle} dx \right].$$

Note that $-B_K(v)$ is precisely the maximum entropy dual program, up to negation of y in the expression.

Open questions still remain about the efficient computability of the entropic barrier. This is in particular true in the case where K is a polytope, given as a membership oracle. Towards this, the following is essentially a corollary to Theorem 4.2 (see Section 6.3 for a full proof), and can be used to efficiently compute the entropic barrier at points which are in the η -interior of K .

Corollary 4.13 (Bounding box for convex bodies) *Let $\Omega \subset \mathbb{R}^d$ be a convex body contained in a ball of radius R . Further, let A be an element of the η -interior of the convex hull of Ω . Then there is an optimal solution Y^* to the dual program such that $\|Y^*\| \leq \text{poly}(\eta^{-1}, d, \log(R))$.*

Details of how this implies computability of the entropic barrier are omitted from this paper.

5 Technical overview

In this section, we give overviews of the proofs of the main results of this paper and compare our techniques with those of previous work. We start by describing the approach of [36] in the case of discrete uniform measures μ with finite support $\Omega \subseteq \{0, 1\}^d$. In this case, the marginals vector A of a measure ν on Ω is defined by setting A_k to be the expected value of the k th entry of x when x is chosen according to ν . Note that the marginals vector A is always an element of $\text{hull}(\Omega)$. The problem the authors of [36] solve is described as follows: given a finite subset Ω and

a desired marginals vector A in the η -interior of $\text{hull}(\Omega)$, compute the probability measure on Ω with marginals A which maximizes entropy.

They consider the dual formulation

$$\inf_{y \in \mathbb{R}^d} F_A(y) := \langle y, A \rangle + \log \sum_{x \in \Omega} e^{-\langle x, y \rangle},$$

which gives rise to measures on Ω of the following succinct form for some real vector y^* :

$$\nu(x) \propto e^{-\langle x, y^* \rangle}.$$

By strong duality $\nu = \nu^*$ is the entropy maximizing measure, and they then use the ellipsoid method to approximate y^* .

We generalize their approach to continuous measures μ on continuous domains Ω . For the most part, the ellipsoid algorithm can be applied in the same way as in the discrete case once we have the three main pieces in hand: (1) strong duality, (2) a bound on Y^* , and (3) the strong counting oracle. Even in the continuous case, one can show that strong duality holds via a certain Slater-type condition (see Sections A.2 and A.3). What makes the passage from the discrete case to the continuous case much more interesting and nontrivial is proving the remaining two main pieces.

5.1 Proof overview: bounding box

The goal of this section is to explain the proofs of the main bounding box result and its corollaries. We first describe the approach of the discrete μ case discussed above. Note that for $B \in \text{hull}(\Omega)$, there exists some $X_0 \in \Omega$ such that

$$\langle -Y^*, X_0 - B \rangle \geq 0,$$

since every closed half-space containing B contains an extreme point $X_0 \in \text{hull}(\Omega)$. If A is in the η -interior of $\text{hull}(\Omega)$, we can choose $B = A - \eta \frac{Y^*}{\|Y^*\|}$ to get:

$$\langle -Y^*, X_0 - A \rangle \geq \eta \|Y^*\|.$$

Because μ is a discrete uniform measure, we have $\mu(\{X_0\}) = |\Omega|^{-1}$. This implies a bound on Y^* as follows, via the dual objective function $F_A(Y)$:

$$0 = F_A(0) \geq F_A(Y^*) = \log \int e^{\langle -Y^*, X - A \rangle} d\mu(X) \geq \log \left(e^{\eta \|Y^*\|} \cdot |\Omega|^{-1} \right) \implies \|Y^*\| \leq \frac{\log |\Omega|}{\eta}.$$

The lower bound on $F_A(Y^*)$ above follows from restricting the integral (which is a sum in the discrete case) to the single point X_0 . This demonstrates exactly why this argument fails in the continuous case, because in that case we have $\mu(\{X\}) = 0$ for all $X \in \Omega$.

This is the first difficulty we must overcome. We need a way to restrict the dual objective integral to a region of Ω which has positive mass, emulating the role of atoms in the discrete case.

We introduce a two-parameter interior for the measure μ . We say that A is in the (η, δ) -interior of μ if every half-space intersecting the η -ball about A contains at least δ mass of μ (Definition 6.2). Instead of restricting the dual integral to a single point of Ω , we restrict it to the appropriate δ -mass to obtain a bound on $\|Y^*\|$:

$$0 \geq \log \int e^{\langle -Y^*, X - A \rangle} d\mu(X) \geq \log \left(e^{\eta \|Y^*\|} \cdot \delta \right) \implies \|Y^*\| \leq \frac{1}{\eta} \log \frac{1}{\delta}.$$

We explain this formally in Lemma 6.1.

This leads to the second difficulty. Our bounding box theorem only refers to the η parameter, and so we need a way to handle or control δ in terms of η and d .

Here is where the key *balance* property comes into play. We say that a measure μ is balanced if for all $\varepsilon > 0$ and $X \in \Omega$, the ε -ball about X contains $\exp(-\text{poly}(\varepsilon^{-1}, d))$ of the mass of μ (Definition 4.1). This links the two interiority parameters: from any point of the ε -interior of $\text{hull}(\Omega)$, there will be at least $\exp(-\text{poly}(\varepsilon^{-1}, d))$ mass in the direction of any $X \in \Omega$ on the boundary.

The crucial feature of the balance property is then how this linking of the parameters allows one to transfer between them. Specifically for a balanced measure, the η -interior of $\text{hull}(\Omega)$ is contained in the $(\frac{\eta}{2}, \exp(-\text{poly}(\frac{\eta}{2}, d)))$ -interior of μ . To see this, let A be in the η -interior of $\text{hull}(\Omega)$. Hence, any half space which intersects the $\frac{\eta}{2}$ -ball about A contains another $\frac{\eta}{2}$ -ball in $\text{hull}(\Omega)$. By translating this ball toward a point of Ω , we can assume that the half-space contains an $\frac{\eta}{2}$ -ball about a point of Ω . Since μ is balanced, this implies A is in the $(\frac{\eta}{2}, \exp(-\text{poly}(\frac{\eta}{2}, d)))$ -interior of μ .

At this point, the rest of the proof of Theorem 4.2 is straightforward. For balanced μ and A in the η -interior of $\text{hull}(\Omega)$, we actually have that A is in the $(\frac{\eta}{2}, \exp(-\text{poly}(\frac{\eta}{2}, d)))$ -interior of μ . The two parameter bound described above then implies $\|Y^*\| \leq \text{poly}(\frac{1}{\eta}, d)$.

To obtain bounding boxes for μ_k on \mathcal{P}_k , $n \times n$ rank k projections, (Corollary 6.4) and to uniform measures on convex bodies (Corollary 6.6), we then demonstrate balance properties. In the case of μ_k , $\mathcal{P}_k \subset B_{\sqrt{k}}(0)$ can be covered by at most $\exp(\text{poly}(\log \delta^{-1}, n))$ balls of radius δ for any $\delta > 0$, morally because:

$$\frac{\text{vol}(B_{\sqrt{k}})}{\text{vol}(B_\delta)} = \frac{(\pi\sqrt{k})^n/n!}{(\pi\delta)^n/n!} = \left(\frac{\sqrt{k}}{\delta}\right)^n = \exp(\text{poly}(\log \delta^{-1}, n)).$$

Therefore a δ -ball about *some* point of \mathcal{P}_k must contain at least $\exp(-\text{poly}(\log \delta^{-1}, n))$ of the mass of μ_k , and unitary invariance then implies that this is actually true for *all* points of \mathcal{P}_k .

For uniform measures μ on convex bodies K contained in a ball of radius R , we prove the bounding box using similar arguments as follows. By the volume ratio computation above, every δ -ball contained in K contains at least $(\frac{\delta}{R})^d$ of the mass of μ . Therefore every A in the η -interior of $\text{hull}(\Omega)$ is also in the $(\frac{\eta}{2}, (\frac{\eta}{2R})^d)$ -interior of μ , since every half-space intersecting the $\frac{\eta}{2}$ -ball about A contains another $\frac{\eta}{2}$ -ball in K . The bounding box then follows from the two-parameter bound discussed above (Lemma 6.1).

5.2 Proof overview: counting oracle for \mathcal{P}_1 and \mathcal{V}_1

The goal of this section is to explain why we can efficiently evaluate and compute the gradient of

$$\mathcal{E}_\mu(Y) = \log \int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X)$$

in the case of $\Omega = \mathcal{P}_1$ and $\Omega = \mathcal{V}_1$

First consider the case of $\Omega = \mathcal{V}_1$, where μ is the pushforward of the Lebesgue measure through $x \mapsto xx^\top$. In this case we have a very explicit formula whenever Y is positive definite:

$$\mathcal{E}_\mu(Y) = \log \int e^{-\langle Y, X \rangle} d\mu(X) = \frac{n}{2} \log(\pi) - \frac{1}{2} \log \det(Y).$$

Since μ is the pushforward of the Lebesgue measure through $x \mapsto xx^\top$, this expression follows from the following classical Gaussian integral formula:

$$\int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) = \int_{\mathbb{R}^n} e^{-x^\top Y x} dx = \sqrt{\det(\pi Y)}. \quad (1)$$

This is demonstrated formally in Proposition 9.3. We show how leads to our optimality characterization of the Goemans-Williamson measure at the end of this section.

The above Gaussian formula for \mathcal{V}_1 suggests a natural approach for computing \mathcal{E}_1 on \mathcal{P}_1 . Allowing complex Hermitian matrices, note that \mathcal{P}_1 is the set of norm-1 elements of \mathcal{V}_1 . Hence, we “integrate out” the norm of the elements of \mathcal{V}_1 , in an attempt to obtain a similar formula for \mathcal{P}_1 . We do this via a standard change of variables (equalities are up to scalar):

$$\int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) = \int_{\mathcal{P}_1} \int_0^\infty e^{-\langle Y, r^2 X \rangle} r^{2n-1} dr d\mu_1(X) = \int_{\mathcal{P}_1} \langle Y, X \rangle^{-n} d\mu_1(X) \neq \mathcal{E}_1(Y).$$

This shows that this approach fails: that is, integrating out the norm does not provide us a formula for $\mathcal{E}_1(Y)$ (for more discussion see Section 9.2).

This demonstrates the first difficulty for constructing a counting oracle for \mathcal{P}_1 . Normalizing the max-entropy measure on \mathcal{V}_1 as above yields a measure on \mathcal{P}_1 which is *not* a max-entropy measure. Max-entropy measures on \mathcal{P}_1 and \mathcal{V}_1 are therefore fundamentally different objects, and thus constructing the associated counting oracles requires different techniques. In particular the well-known Gaussian integral formulas cannot help us in the case of \mathcal{P}_1 .

The remarkable fact is then that max-entropy measures on \mathcal{P}_1 *can* be translated into max-entropy measures on a very simple polytope: the standard simplex in \mathbb{R}^n . We have the following equality for real $Y = \text{diag}(y)$, where m is the Lebesgue measure on the simplex $\Delta_1 := \{p \in \mathbb{R}_+^n : \sum_{i=1}^n p_i = 1\}$:

$$\int_{\mathcal{P}_1} e^{-\langle Y, X \rangle} d\mu_1(X) = \int_{\Delta_1} e^{-\langle y, x \rangle} dm(x).$$

Put another way, max-entropy measures on \mathcal{P}_1 , a nonconvex manifold, correspond to max-entropy measures on Δ_1 , a convex polytope. To see this, first note the following for any m_1, \dots, m_n . The first equality is the Bombieri inner product formula (Lemma 7.2), and the second inequality is a basic induction after a change of variables:

$$\int_{\mathcal{P}_1} X_{11}^{m_1} \cdots X_{nn}^{m_n} d\mu_1(X) = \frac{m_1! \cdots m_n! (n-1)!}{(m_1 + \cdots + m_n + n-1)!} = \int_{\Delta_1} x_1^{m_1} \cdots x_n^{m_n} dm(x).$$

The exponential equality then follows from limiting, since \mathcal{P}_1 and Δ_1 are compact and since $e^{-\langle Y, X \rangle}$ and $e^{-\langle y, x \rangle}$ are limits of polynomials.

This argument also implies the more general fact: that m is the pushforward of μ_1 through the map $\phi : X \mapsto \text{diag}(X)$:

$$\int_{\mathcal{P}_1} f(\phi(X)) d\mu_1(X) = \int_{\Delta_1} f(x) dm(x).$$

This transfer to the simplex now leads to an explicit computation for $\mathcal{E}_1(Y)$ when $Y = \text{diag}(y)$. (Considering diagonal Y is actually without loss of generality, see the discussion in Section 7.) By making a change of variables, the simplex integral is an iterated convolution:

$$\frac{1}{(n-1)!} \int_{\mathcal{P}_1} e^{-\langle Y, X \rangle} d\mu_1(X) = \int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\cdots-x_{n-2}} e^{-\langle y, x \rangle} dx = (e^{-y_1 t} * \cdots * e^{-y_n t})|_{t=1}.$$

This is stated formally in Lemma 7.3. Applying the Laplace transform \mathcal{L} converts this convolution into a partial fraction decomposition problem for distinct values of y_i :

$$(e^{-y_1 t} * \cdots * e^{-y_n t})|_{t=1} = \mathcal{L}^{-1} \left[\frac{1}{\prod_i (s + y_i)} \right] (1) = \mathcal{L}^{-1} \left[\sum_i \frac{c_i}{s + y_i} \right] (1) = \sum_i c_i e^{-y_i}.$$

Computing the values of c_i via a standard partial fractions formula gives:

$$\frac{1}{(n-1)!} \int_{\mathcal{P}_1} e^{-\langle Y, X \rangle} d\mu_1(X) = \sum_{i=1}^n \frac{e^{-y_i}}{\prod_{j \neq i} (y_j - y_i)} = \frac{\det(M(-y))}{\prod_{i < j} (y_j - y_i)}.$$

This is stated formally in Proposition 7.4. Here $M(-y)$ is a Vandermonde-like matrix which arises when forming the common denominator of the last expression, given (for the case of distinct y_i s) as follows:

$$M(-y) := \begin{bmatrix} 1 & 1 & \dots & 1 \\ (-y_1) & (-y_2) & \dots & (-y_n) \\ (-y_1)^2 & (-y_2)^2 & \dots & (-y_n)^2 \\ \vdots & \vdots & \ddots & \vdots \\ (-y_1)^{n-2} & (-y_2)^{n-2} & \dots & (-y_n)^{n-2} \\ e^{-y_1} & e^{-y_2} & \dots & e^{-y_n} \end{bmatrix}.$$

We define this matrix formally in Definition 7.1.

This brings us to the second difficulty for constructing a counting oracle for \mathcal{P}_1 . When the values of y_i are not distinct, then the denominator vanishes and this formula cannot be used. Even though \mathcal{E}_1 is continuous, this could still be a major problem: if for example the gradient of $\mathcal{E}_1(Y)$ becomes large as y_i approaches y_j , then computing $\mathcal{E}_1(Y)$ could become computationally infeasible.

To handle this difficulty, we take limits by successively applying L'Hopital's rule. One iteration for $y_1 = y_2$ goes as follows:

$$\lim_{y_2 \rightarrow y_1} \frac{\det(M(-y))}{\prod_{i < j} (y_j - y_i)} = \frac{\partial_{y_2} \det(M(-y))}{\partial_{y_2} \prod_{i < j} (y_i - y_j)} \Big|_{y_2=y_1} = \frac{\det(M'(-y))}{\prod_{i < j} (y_i - y_1)^2 \prod_{i < j} (y_i - y_j)}.$$

The key observation here is the fact that the numerator is still a determinant, due to the fact that only one column of $M(-y)$ depends on y_i for all i . Applying L'Hopital's rule as many times as is necessary leads to the following, where λ_i represent the distinct values of y with multiplicities m_i :

$$\exp(\mathcal{E}_1(Y)) = \int_{\mathcal{P}_1} e^{-\langle Y, X \rangle} d\mu_1(X) = (n-1)! \frac{\det(M(-\lambda))}{\prod_{i < j} (\lambda_i - \lambda_j)^{m_i m_j}}.$$

Note that $M(-\lambda)$ is a matrix similar to $M(-y)$ above which handles the non-distinctness (we unify the notation of these matrices in Definition 7.1). A similar expression for the gradient is achieved using the same techniques, and so we state it here without further detail:

$$(\nabla \mathcal{E}_1(Y))_l = - \sum_{i \neq p} \frac{m_i}{\lambda_p - \lambda_i} - \frac{\det(M_p(-\lambda))}{\det(M(-\lambda))}.$$

$M_p(-\lambda)$ is another, similar Vandermonde-like matrix, see Proposition 7.6 and Definition 7.2.

Since the entries of $M(-\lambda)$ and $M_p(-\lambda)$ have bit complexity polynomial in n and the bit complexity of e^{-y_i} , their determinants have the same bit complexity. Therefore these formulas, for $\mathcal{E}_1(Y)$ and its gradient, lead to an efficient counting oracle for \mathcal{P}_1 .

The optimality of Goemans-Williamson measure. As a consequence of Equation (1), we now show briefly how this formula is used to prove that the Goemans-Williamson measure μ_{GW} with respect to a real symmetric positive definite matrix A is a max-entropy measure on \mathcal{V}_1 . For $A = VV^\top$, the measure μ_{GW} is defined to be distributed according to $xx^\top := (Vg)(Vg)^\top$ where g

is a standard Gaussian in \mathbb{R}^n . By the change of variables formula, xx^\top is distributed as follows on \mathbb{R}^n :

$$xx^\top \sim e^{-\frac{1}{2}\|V^{-1}x\|^2} \cdot \det(V^{-1})d\mu(xx^\top) \propto e^{-\langle \frac{1}{2}A^{-1}, xx^\top \rangle}d\mu(xx^\top).$$

We state this formally in Proposition 9.1. To prove that this is a max-entropy measure, we determine the critical point of the dual objective with respect to real symmetric positive definite Y :

$$0 = \nabla(\langle Y, A \rangle + \mathcal{E}_\mu(Y)) = A - \frac{1}{2}Y^{-1} \implies Y^* = \frac{1}{2}A^{-1}.$$

Therefore, the Goemans-Williamson measure $\mu_{\text{GW}}(X) \propto e^{-\langle \frac{1}{2}A^{-1}, X \rangle}d\mu(X)$ is the max-entropy measure on \mathcal{V}_1 with respect to A .

Proof overview: sampling for \mathcal{P}_1 . We now discuss how to sample from max-entropy distributions on \mathcal{P}_1 . Our main algorithm (Theorem 4.4) gives an efficient oracle for approximating the max-entropy density function:

$$\nu(X) \propto e^{-\langle Y^*, X \rangle}.$$

The main problem is that it is not at all clear how to use such a density function to sample from a manifold.

We avoid this difficulty by transferring the problem of sampling to the simplex Δ_1 for $Y^* = \text{diag}(y^*)$, using the following fact discussed in the previous section:

$$\frac{1}{(n-1)!} \int_{\mathcal{P}_1} e^{-\langle Y^*, X \rangle} d\mu_1(X) = \int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\cdots-x_{n-2}} e^{-\langle y^*, x \rangle} dx.$$

The sampling process for \mathcal{P}_1 then occurs in two parts.

First, we sample from the max-entropy distribution on the simplex, one coordinate at a time. We use the right-hand side of the above expression to compute the cumulative density function (CDF) for each coordinate, conditioned on the previously sampled coordinates. Formulas and computations for these conditioned CDFs are very similar to that of the counting oracle, and hence we omit them here (see Corollary 7.13).

Once we have a sample x on the simplex, we need to convert it into a sample on \mathcal{P}_1 by considering its inverse image under the map $\phi : X \mapsto \text{diag}(X)$. The difficulty that now arises is the fact that there are many elements of \mathcal{P}_1 which map to the same simplex element under ϕ .

Fortunately, there is a principled way to select from these possibilities. The fiber $\phi^{-1}(x)$ is an orbit of the action of diagonal unitary matrices on \mathcal{P}_1 by conjugation. Since Y^* is diagonal, this implies the max-entropy measure $\nu(X)$ is uniform when restricted to $\phi^{-1}(x)$. Given x , we then sample X from $\phi^{-1}(x)$ by picking an arbitrary $X_0 \in \phi^{-1}(x)$ and conjugating by a uniformly random diagonal unitary matrix.

Hence, to sample X from \mathcal{P}_1 we (1) sample x from the simplex, and then (2) sample X uniformly from $\phi^{-1}(x)$. This samples X from the correct measure due to the disintegration theorem, which says the following for any f :

$$\int_{\mathcal{P}_1} f(X) d\mu_1(X) = \int_{\Delta_1} \int_{\phi^{-1}(x)} f(X) d\mu_{\phi^{-1}(x)}(X) dm(x).$$

That is, the measure μ_1 can be split into measures on Δ_1 and on the fibers $\phi^{-1}(x)$ (see Proposition 7.11).

Therefore, the above sampling process efficiently samples the max-entropy measure on \mathcal{P}_1 with density $\nu(X)$.

5.3 Proof overview: extending the counting oracle for \mathcal{P}_1 to \mathcal{P}_k

For the case of \mathcal{P}_k and μ_k , we want to generalize the formulas of the $k = 1$ case. To do this, we make use of the famous Harish-Chandra-Itzykson-Zuber formula (Theorem 7.7) for integrals over the Haar measure of the unitary group $U(n)$. It is stated as follows for Hermitian Y, B with *distinct* eigenvalues y_i, β_i :

$$\int_{U(n)} e^{-\langle Y, UBU^* \rangle} dU = \left(\prod_{p=1}^{n-1} p! \right) \frac{\det([e^{-y_i \beta_j}]_{1 \leq i, j \leq n})}{\prod_{i < j} (y_i - y_j)(\beta_j - \beta_i)}.$$

For $B = \text{diag}(1, \dots, 1, 0, \dots, 0)$ with k 1s and $n - k$ 0s, notice that $\mathcal{P}_k = \{UBU^* : U \in U(n)\}$. This leads to the following:

$$\exp(\mathcal{E}_k(Y)) = \int_{\mathcal{P}_k} e^{-\langle Y, X \rangle} d\mu_k(X) = \int_{U(n)} e^{-\langle Y, UBU^* \rangle} dU.$$

To handle the issue of the denominator vanishing, and to compute the gradient, we apply all the same techniques which were required for the $k = 1$ case (see Corollaries 7.8 and 7.9). These formulas end up having the right bit complexity, and so they immediately imply an efficient strong counting oracle for \mathcal{P}_k .

Unlike in the case of $k = 1$, the problem of sampling in the case of $k > 1$ is more difficult as the image of \mathcal{P}_k under the map $\phi : X \mapsto \text{diag}(X)$ is much more complicated. Thus we leave as an open problem the question of sampling from the associated maximum entropy distributions in the case of \mathcal{P}_k for $k > 1$.

6 Bounding box

In this section, we prove the general bounding box result (Theorem 4.2). With this, we then specialize to the cases of rank- k projections and convex bodies.

6.1 General bounding box

In what follows we will discuss “interiors” of a probability distribution μ given by two parameters, (η, δ) . The η parameter will control how far we are from the boundary, and the δ parameter will control how well-distributed μ is on its support. At the end of the day, we will prove that for nice situations one only needs to consider the η parameter (as in the bounding box result of [36]).

We now define the two-parameter interior. In what follows, we will let $V_{\mathcal{L}}$ be the vector subspace given by $\mathcal{L}(X) = 0$, where $\mathcal{L}(X) = B$ is the maximal set of linearly independent equality constraints for Ω . More informally, $V_{\mathcal{L}}$ is the vector space corresponding to the minimal affine space in which $\mathcal{K} = \text{hull}(\Omega)$ lives (i.e., translate the affine space so that $0 \in V_{\mathcal{L}}$). The fact that $\mathcal{L}(X) = B$ is a maximal linearly independent set means that the optimal solution to the dual program is unique when restricted to $V_{\mathcal{L}}$. (Existence follows from Lemmas A.1 and A.3.) We discuss this further in Section 8.

Definition 6.1 *We define the $(0, \delta)$ -interior of μ to be the set of all $A \in \mathcal{K}$ such that for all $Y \in V_{\mathcal{L}}$ we have:*

$$\mu(\{X \in \Omega \mid \langle X - A, Y \rangle \geq 0\}) > \delta.$$

Morally, this says that every closed half-space containing A contains more than δ of the mass of μ . Note that this is not always an open set (which is perhaps a bit odd for something called the “interior”, but this will be our convention).

Definition 6.2 (Two-parameter interior) We define the (η, δ) -interior of μ to be the set of all $A \in \mathcal{K}$ such that the ball of radius η about A is contained in the $(0, \delta)$ -interior of μ . Note that this is not necessarily an open set.

The next lemma is then precisely how to combine the two parameters to get a bounding box for the optimal solution to the dual program.

Lemma 6.1 (Two-parameter bounding box) Given $A \in \mathcal{K}$, let $Y^* \in V_{\mathcal{L}}$ be the optimal solution to the dual program. Recall the dual objective:

$$\inf_Y F_A(Y) = \inf_Y \log \int_{\Omega} e^{-\langle Y, X - A \rangle} d\mu(X).$$

If A is in the (η, δ) -interior of μ , then $\|Y^*\| \leq \frac{1}{\eta} \log\left(\frac{1}{\delta}\right)$.

Proof: By definition, we have that $A - \eta \cdot \frac{Y^*}{\|Y^*\|}$ is in the $(0, \delta)$ -interior of \mathcal{K} . Therefore:

$$\delta \leq \mu(\{X \in \Omega \mid \langle X - (A - \eta \cdot Y^*/\|Y^*\|), -Y^* \rangle \geq 0\}) = \mu(\{X \in \Omega \mid \langle X - A, -Y^* \rangle \geq \eta \cdot \|Y^*\|\}).$$

This gives the bound:

$$\log \int e^{\langle -Y^*, X - A \rangle} d\mu(X) \geq \log\left(\delta \cdot e^{\eta \cdot \|Y^*\|}\right) = \log(\delta) + \eta \cdot \|Y^*\|.$$

On the other hand, plugging in $Y = 0$ gives an upper bound on the optimal value of the above dual program:

$$0 \geq \log \int e^{\langle -Y^*, X - A \rangle} d\mu(X) \geq \log(\delta) + \eta \cdot \|Y^*\|.$$

Rearranging this gives the result. ■

This gives us a good way of bounding solutions corresponding to interior points of \mathcal{K} . In general however, trying to get a bound on the δ parameter of the interior is much more difficult than that of the η parameter. To deal with this we define a property of μ which allows us to only have to consider the η parameter.

Definition 6.3 (δ -balanced measure) We say that μ is δ -balanced if for any $X \in \Omega$, we have that at least $\exp(-\text{poly}(\delta^{-1}, d))$ of the mass of μ is contained in the δ -ball about X (where d is the dimension of \mathcal{K}). If f is the polynomial in the exponent (i.e., $\exp(-f(\delta^{-1}, d))$), then we say that μ is δ -balanced with bound f .

We now prove the main bounding box theorem for such balanced measures. We then use this to obtain a bounding box for rank- k projections and for convex bodies in the following sections.

Theorem 6.2 (Bounding box for balanced measures) Suppose μ is $\frac{\eta}{2}$ -balanced with bound f . If A is in the $(\eta, 0)$ -interior of μ and $Y^* \in V_{\mathcal{L}}$ is the optimal solution to the corresponding dual program, then $\|Y^*\| \leq 2\eta^{-1} \cdot f(2\eta^{-1}, d) = \text{poly}(\eta^{-1}, d)$.

Proof: We first show that the $(\frac{\eta}{2}, 0)$ -interior of μ is contained in the $(0, \exp(-f(\frac{2}{\eta}, d)))$ -interior of μ . To see this, let A_0 be some element of the $(\frac{\eta}{2}, 0)$ -interior of μ . Then any closed half-space containing A_0 also contains an $\frac{\eta}{2}$ -ball about some $X \in \Omega$. That is, for every $Y \in V_{\mathcal{L}}$ there exists X such that:

$$B_{\eta/2}(X) \subseteq \{Z \in \Omega \mid \langle Z - A_0, Y \rangle \geq 0\}.$$

Since μ is $\frac{\eta}{2}$ -balanced, we have that $\exp(-f(\frac{2}{\eta}, d))$ of the mass of μ is contained in the $\frac{\eta}{2}$ -ball about X . This implies:

$$\exp(-f(2/\eta, d)) \leq \mu(B_{\eta/2}(X)) \leq \mu(\{Z \in \Omega \mid \langle Z - A_0, Y \rangle \geq 0\}).$$

That is, A_0 is in the $(0, \exp(-f(\frac{2}{\eta}, d)))$ -interior of μ .

Now for A in the $(\eta, 0)$ -interior of μ , we have that the $\frac{\eta}{2}$ -ball about A is contained in the $(\frac{\eta}{2}, 0)$ -interior of μ . Therefore A is in the $(\frac{\eta}{2}, \exp(-f(\frac{2}{\eta}, d)))$ -interior of μ . By Lemma 6.1, this implies $\|Y^*\| \leq 2\eta^{-1} \cdot f(2\eta^{-1}, d)$. \blacksquare

Remark 6.3 Note that Theorem 6.2 is immediately applicable to uniform discrete measures on (singly) exponentially sized sets S . In particular, such a measure is automatically balanced with constant bound $f = \log |S|$.

6.2 Rank- k projections

We now prove bounding box result for \mathcal{P}_k , by showing that μ_k is balanced and applying the previous theorem. Note that in this case $\mathcal{L}(X) = B$ reduces to $\text{Tr}(X) = k$, and so $V_{\mathcal{L}}$ is the set of traceless Hermitian matrices in this case.

Corollary 6.4 (Bounding box for \mathcal{P}_k) *Let μ_k be the uniform distribution on \mathcal{P}_k . Then given A in the $(\eta, 0)$ -interior of μ_k , the optimal traceless solution Y^* of the corresponding dual program is such that $\|Y^*\| \leq \frac{2n^2}{\eta} \log\left(\frac{8n\sqrt{k}}{\eta}\right)$.*

Proof: We prove that μ_k is balanced and then apply the previous proposition. The number of balls of size δ required to cover the unit ball in \mathbb{R}^{n^2} (with Euclidean/Frobenius norm) is at most $(2n/\delta)^{n^2}$. Since the set of projections of rank k is contained in the sphere of radius \sqrt{k} , we have that it requires at most $(2n\sqrt{k}/\delta)^{n^2}$ δ -balls to cover all such projections. With this, there exists some δ -ball (call it B_δ) in this cover which contains at least $(2n\sqrt{k}/\delta)^{-n^2}$ of the mass of μ_k . Pick some $X \in \mathcal{P}_k \cap B_\delta$, and let $B_{2\delta}(X)$ be the ball of radius 2δ which is centered at X . Thus, in fact $B_{2\delta}(X)$ contains at least $(2n\sqrt{k}/\delta)^{-n^2}$ of the mass of μ_k . By unitary invariance of μ_k , we have that the ball of radius 2δ about any point of \mathcal{P}_k contains at least $(2n\sqrt{k}/\delta)^{-n^2}$ of the mass of μ_k . That is, μ_k is δ -balanced with bound $f(\delta^{-1}, n) = n^2 \log(4n\sqrt{k} \cdot \delta^{-1})$ for all $\delta > 0$. Applying the previous proposition then gives the result. \blacksquare

Remark 6.5 In the discrete measure case, the authors of [38] were able to improve the dependence on η of the bounding box from η^{-1} to $\log(\eta^{-1})$. This leads to a max-entropy approximation algorithm which does not depend on η . One may then naturally ask whether or not this is possible for the bounding box for μ_k discussed here. The answer turns out to be “no”, and this can be seen by considering the optimal $Y^* = \text{diag}(y_1, y_2)$ in the case of $n = 2$ and $k = 1$. Specifically one can show that for $A = \text{diag}(\eta, 1 - \eta)$, the value of $|y_1 - y_2|$ is of the order η^{-1} as $\eta \rightarrow 0$. Since the relative entropy of the optimal distribution is unbounded as η approaches 0, approximation of Y^* cannot help us to improve the dependence of $|y_1 - y_2|$ on η^{-1} .

6.3 Convex bodies

We now prove bounding box result for convex bodies. Instead of applying the previous theorem directly, we make some simpler computations which are in the same spirit.

Corollary 6.6 (Bounding box for convex bodies) *Let μ be the uniform distribution on a d -dimensional convex body Ω contained in a ball of radius R . (Note that $\mathcal{K} = \text{hull}(\Omega) = \Omega$ in this case.) Then given α in the $(\eta, 0)$ -interior of μ , the optimal solution $y^* \in V_{\mathcal{L}}$ of the corresponding dual program is such that $\|y^*\| \leq \frac{2d}{\eta} \log(\frac{4R}{\eta})$.*

Proof: Note that α in the $(\eta, 0)$ -interior of μ is automatically in the $(\frac{\eta}{2}, (\frac{\eta}{4R})^d)$ -interior of μ , since:

$$\mu(B_{\eta/4}) \geq \frac{\text{vol}(B_{\eta/4})}{\text{vol}(B_R)} = \left(\frac{\eta}{4R}\right)^d.$$

By Lemma 6.1, this implies $\|y^*\| \leq \frac{2d}{\eta} \log(\frac{4R}{\eta})$. \blacksquare

7 Counting oracle for \mathcal{P}_k

In this section, we prove existence of a strong counting/integration oracle for the objective function of the dual program Dual_k . Recall the dual objective function:

$$F_A(Y) = \langle Y, A \rangle + \mathcal{E}_k(Y) = \langle Y, A \rangle + \log \int_{\mathcal{P}_k} e^{-\langle Y, X \rangle} d\mu_k(X).$$

We want to be able to efficiently compute this function and its gradient. In this case of rank- k projections, we make the simplifying assumption that Y and A are both diagonal. This simplification is actually without loss of generality, due to the Schur-Horn theorem (Corollary B.2) and unitary invariance of μ_k . Further, it is enough to consider only $\mathcal{E}_k(Y)$ (which is independent of A) since $\langle Y, A \rangle$ is linear and hence easy to handle. This leads to the main theorem of this section, stated originally as Theorem 4.5.

Theorem 7.1 (Counting oracle for \mathcal{P}_k) *There is an algorithm that, given $n \in \mathbb{N}$, $k \in [n]$, an $n \times n$ real diagonal matrix $Y = \text{diag}(y)$, and a $\delta > 0$, returns numbers \bar{E}, \bar{G} such that*

1. $|\bar{E} - \mathcal{E}_k(Y)| \leq \delta$
2. $|\bar{G} - \nabla \mathcal{E}_k(Y)| \leq \delta,$

where \mathcal{E}_k is the exponential integral defined above (and in Definition 2.2). The running time of the algorithm is polynomial in $n, \log(\frac{1}{\delta})$, and the number of bits needed to represent e^{-y_i} for any $i \in [n]$.

The main tool we use to prove this theorem is a collection of explicit formulas for computing \mathcal{E}_k and its gradient. We first discuss this in full detail for the case of $k = 1$. After that, we discuss how to generalize the arguments to the $k > 1$ case.

7.1 Algorithm for $k = 1$

In this section, we construct the strong counting/integration oracle for rank-1 projections by giving formulas for the function \mathcal{E}_1 and its gradient (Propositions 7.4 and 7.6). Specifically, for diagonal $Y = \text{diag}(y)$ with distinct entries $\lambda_1 > \dots > \lambda_k$ with multiplicities m_1, \dots, m_k , we can compute the following where $M(y)$ and $M_p(y)$ are matrices defined below:

$$\mathcal{E}_1(Y) = \sum_{i=1}^{n-1} \log(i) + \log \det(M(-y)) - \sum_{i < j} m_i m_j \log(\lambda_i - \lambda_j),$$

$$(\nabla \mathcal{E}_1(Y))_l = - \sum_{i \neq p} \frac{m_i}{\lambda_p - \lambda_i} - \frac{\det(M_p(-y))}{\det(M(-y))}.$$

The only potentially hard part of computing these expressions is computing the determinants of $M(-y)$ and $M_p(-y)$. It is a standard fact that one can compute a determinant in time polynomial in the number of bits needed to represent the matrix, so we just need to demonstrate that the matrices have the necessary bit complexity. Considering Definitions 7.1 (for $\gamma = 1$) and 7.2 below, we see that the matrix entries depend on computing e^{-y_i} , $n!$, and y_i^n . All of these can be computed in time polynomial in n and number of bits needed to represent e^{-y_i} , which is exactly what we need. That said, all we have left now is to prove the two formulas stated above, and we do this in the following sections.

7.1.1 Evaluating the dual integral

We now prove the main evaluation formulas for integrals on the manifold \mathcal{P}_1 . Throughout we will often consider integrals on the unit sphere in \mathbb{C}^n , denoted $S_{\mathbb{C}}^n$, instead of on \mathcal{P}_1 directly, and we will let $\mu_{S_{\mathbb{C}}^n}$ refer to the Haar measure on the unit sphere. Note that the transfer of formulas from the sphere to \mathcal{P}_1 is straightforward, as given by (2) of Proposition 7.4. First, we define a parameterized matrix of a particular form which will show up many times in our computations.

Definition 7.1 (Matrix for dual integral, $k = 1$) *Given $y_1, \dots, y_n \in \mathbb{R}$, let $\lambda_1 < \dots < \lambda_k$ denote the distinct values of y_i with multiplicities m_1, \dots, m_k . Given γ , we define an $n \times n$ matrix $M(y, \gamma)$ as follows:*

$$M(y, \gamma) := \begin{bmatrix} 1 & 0 & \dots & 0 & 1 & 0 & \dots \\ \lambda_1 & 1 & \dots & 0 & \lambda_2 & 1 & \dots \\ \lambda_1^2 & \binom{2}{1} \lambda_1 & \dots & 0 & \lambda_2^2 & \binom{2}{1} \lambda_2 & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \dots \\ \lambda_1^{n-2} & \binom{n-2}{1} \lambda_1^{n-3} & \dots & \binom{n-2}{m_1-1} \lambda_1^{n-m_1-1} & \lambda_2^{n-2} & \binom{n-2}{1} \lambda_2^{n-3} & \dots \\ \frac{e^{\gamma \lambda_1}}{0!} & \frac{\gamma e^{\gamma \lambda_1}}{1!} & \dots & \frac{\gamma^{m_1-1} e^{\gamma \lambda_1}}{(m_1-1)!} & \frac{e^{\gamma \lambda_2}}{0!} & \frac{\gamma e^{\gamma \lambda_2}}{1!} & \dots \end{bmatrix}.$$

We also define $M(y) := M(y, 1)$. Note that only one row of $M(y, \gamma)$ depends on γ .

We now state a lemma which gives the most basic result about integrals on \mathcal{P}_1 . Specifically, we state a well-known result for integrals of polynomial-like functions. This proof is very related to the unitarily invariant inner product on homogeneous polynomials, which has many names in the literature: Bombieri inner product, Fischer-Fock inner product, Segal-Bargmann inner product, etc. The following lemma is standard, see e.g. Lemma 3.2 of [33].

Lemma 7.2 (Bombieri inner product formula) *For $\alpha \in \{0, 1, 2, \dots\}^n$ such that $\sum_i \alpha_i = d$, we have:*

$$\int |v|^{2\alpha} d\mu_{S_{\mathbb{C}}^n}(v) = \int \prod_i |v_i|^{2\alpha_i} d\mu_{S_{\mathbb{C}}^n}(v) = \binom{d}{\alpha}^{-1} \binom{d+n-1}{n-1}^{-1} = \frac{\alpha_1! \cdots \alpha_n! (n-1)!}{(d+n-1)!}.$$

Here, $\binom{d}{\alpha}$ is the multinomial coefficient, and $\binom{d+n-1}{n-1}$ is the binomial coefficient.

The next lemma then shows the connection between the integrals we want to compute and the Laplace transform. As an immediately corollary, we obtain equality of (3) and (4) in Proposition 7.4 below in the case of distinct values of y_1, \dots, y_n .

Lemma 7.3 For $y_1, \dots, y_n \in \mathbb{R}$ and $x_n := 1 - x_1 - \dots - x_{n-1}$, we have the following where $*$ denotes the usual integral convolution:

$$\int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\dots-x_{n-2}} e^{\langle y, x \rangle} dx_{n-1} \cdots dx_1 = (e^{y_1 t} * \cdots * e^{y_n t}) \Big|_{t=1}.$$

If $y_1 < y_2 < \cdots < y_n$, then we further have:

$$\int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\dots-x_{n-2}} e^{\langle y, x \rangle} dx_{n-1} \cdots dx_1 = \sum_{i=1}^n \frac{e^{y_i}}{\prod_{j \neq i} (y_i - y_j)} = \frac{\det(M(y))}{\prod_{i < j} (y_j - y_i)}.$$

Proof: We first compute:

$$\begin{aligned} & \int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\dots-x_{n-2}} e^{\langle y, x \rangle} dx_{n-1} \cdots dx_1 \\ &= \int_0^1 e^{y_1 x_1} \cdots \int_0^{1-x_1-\dots-x_{n-2}} e^{y_{n-1} x_{n-1}} e^{y_n (1-x_1-\dots-x_{n-1})} dx_{n-1} \cdots dx_1 \\ &= \int_0^1 e^{y_1 x_1} \cdots \int_0^{1-x_1-\dots-x_{n-3}} e^{y_{n-2} x_{n-2}} (e^{y_{n-1} t} * e^{y_n t}) \Big|_{t=1-x_1-\dots-x_{n-2}} dx_{n-2} \cdots dx_1 \\ &= \cdots = \int_0^1 e^{y_1 x_1} (e^{y_2 t} * \cdots * e^{y_n t}) \Big|_{t=1-x_1} dx_1 \\ &= (e^{y_1 t} * \cdots * e^{y_n t}) \Big|_{t=1}. \end{aligned}$$

Using the Laplace transform, we have $\mathcal{L}[e^{y_i t}](s) = \frac{1}{s-y_i}$ which implies:

$$e^{y_1 t} * \cdots * e^{y_n t} = \mathcal{L}^{-1} \left[\frac{1}{(s-y_1)(s-y_2) \cdots (s-y_n)} \right] (t).$$

Assuming $y_1 < y_2 < \cdots < y_n$, we can use Lagrange interpolation to compute:

$$\mathcal{L}^{-1} \left[\frac{1}{(s-y_1)(s-y_2) \cdots (s-y_n)} \right] (t) = \mathcal{L}^{-1} \left[\frac{1}{Q'(y_1) \cdot (s-y_1)} + \cdots + \frac{1}{Q'(y_n) \cdot (s-y_n)} \right] (t).$$

Here, $Q(s) := (s-y_1)(s-y_2) \cdots (s-y_n)$. With this we have:

$$\mathcal{L}^{-1} \left[\frac{1}{Q'(y_1) \cdot (s-y_1)} + \cdots + \frac{1}{Q'(y_n) \cdot (s-y_n)} \right] (t) = \frac{e^{y_1 t}}{Q'(y_1)} + \cdots + \frac{e^{y_n t}}{Q'(y_n)} = \sum_{i=1}^n \frac{e^{y_i t}}{\prod_{j \neq i} (y_i - y_j)}.$$

Plugging in $t = 1$ gives the first equality in the second statement. To see the last equality, notice that because $y_1 < \cdots < y_n$, the expression for $\det(M(y))$ will be a sum of exponentials multiplied by Vandermonde determinants (expand along the last row of $M(y)$). The result follows, taking care to keep track of signs. \blacksquare

We now state and prove the full evaluation formula for \mathcal{P}_1 . The two most involved parts of the proof are showing equality of (1) and (4) on polynomials and showing equality of (4) and (5) for non-distinct values of y_1, \dots, y_n .

Proposition 7.4 (Evaluating the dual integral, $k = 1$) Fix $n \in \mathbb{N}$, and let $\mu_{S_{\mathbb{C}}^n}, \mu_1, \mu_{\Delta_1}$ be the uniform probability distributions on the complex unit sphere in \mathbb{C}^n , on \mathcal{P}_1 , and on the standard simplex in \mathbb{R}^n , respectively. For a given analytic function f on the standard simplex the following expressions are equal:

1. $\int_{S_{\mathbb{C}}^n} f(|v_1|^2, \dots, |v_n|^2) d\mu_{S_{\mathbb{C}}^n}(v),$
2. $\int_{\mathcal{P}_1} f(\text{diag}(X)) d\mu_1(X),$
3. $\int_{\Delta_1} f(x) d\mu_{\Delta_1}(x),$
4. $(n-1)! \int_0^1 \int_0^{1-x_1} \cdots \int_0^{1-x_1-\cdots-x_{n-2}} f(x_1, \dots, x_{n-1}, 1-x_1-\cdots-x_{n-1}) dx_{n-1} \cdots dx_1.$

If $f(x) = e^{\langle y, x \rangle}$ for some real y with distinct entries $\lambda_1 < \dots < \lambda_k$ with multiplicities m_1, \dots, m_k , then we have another equal expression:

$$5. (n-1)! \frac{\det(M(y))}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}}.$$

Proof: First, for the equality of (1) and (2), note that μ_1 is the pushforward measure of $\mu_{S_{\mathbb{C}}^n}$ through the map $\psi : S_{\mathbb{C}}^n \rightarrow \mathcal{P}_1$ given by $\psi : v \mapsto vv^*$. (To see this, note that ψ is unitarily invariant and $\mu_{S_{\mathbb{C}}^n}$ and μ_1 are the unique unitarily invariant measures on their respective domains.) With this, we then have:

$$\int_{\mathcal{P}_1} f(\text{diag}(X)) d\mu_1(X) = \int_{S_{\mathbb{C}}^n} f(\text{diag}(\psi(v))) d\mu_{S_{\mathbb{C}}^n}(v) = \int_{S_{\mathbb{C}}^n} f(|v_1|^2, \dots, |v_n|^2) d\mu_{S_{\mathbb{C}}^n}(v).$$

That is, (1) and (2) are equal.

Next, the equality of (3) and (4) follows from the fact that the map between the two domains of integration (both of which are simplices) is affine. Therefore the determinant of the Jacobian is a constant, and so we only need to integrate over a constant function to determine that constant. A simple induction shows that it is $(n-1)!$.

To prove the equality of (1) and (4), we compute the integrals on a given monomial $x^m := x_1^{m_1} \cdots x_{n-1}^{m_{n-1}} (1-x_1-\cdots-x_{n-1})^{m_n}$. First, by Lemma 7.2 we have:

$$\int |v|^{2m} d\mu_{S_{\mathbb{C}}^n}(v) = \frac{m_1! \cdots m_n! (n-1)!}{(|m| + n - 1)!} = \binom{|m| + n - 1}{m_1, \dots, m_n, n - 1}^{-1}.$$

Now, note:

$$\int_0^1 x^l (1-x)^m dx = \sum_{k=0}^m \binom{m}{k} (-1)^k \int_0^1 x^{l+k} dx = \frac{l! m!}{(l+m+1)!} \sum_{k=0}^m \frac{\prod_{j \neq k} (l+j+1)}{(-1)^k k! (m-k)!} = \frac{l! m!}{(l+m+1)!}.$$

The last equality is due to Lagrange interpolation, considering the sum as a function of n . We further have:

$$\frac{l! m!}{(l+m+1)!} = \binom{l+m+2-1}{l, m, 2-1}^{-1}.$$

That is, we have equality whenever $n = 2$, proving the base case. The rest of the proof goes by induction. First we compute for $\alpha = 1 - x_1 - \cdots - x_{n-2}$:

$$\int_0^\alpha x_{n-1}^j (\alpha - x_{n-1})^k dx_{n-1} = \int_0^1 (\alpha u)^j (\alpha - \alpha u)^k \alpha du = \alpha^{j+k+1} \int_0^1 u^j (1-u)^k du = \frac{\alpha^{j+k+1} \cdot j! k!}{(j+k+1)!}.$$

With this, we then compute the following by induction, letting $\beta = 1 - x_1 - \cdots - x_{n-3}$:

$$\begin{aligned}
(n-1)! \int_0^1 \cdots \int_0^\alpha x_1^{m_1} \cdots x_{n-1}^{m_{n-1}} (\alpha - x_{n-1})^{m_n} dx_{n-1} \cdots dx_1 \\
&= \frac{m_{n-1}!m_n!}{(m_{n-1} + m_n + 1)!} \cdot (n-1)! \int_0^1 \cdots \int_0^{\alpha+x_{n-2}} x_1^{m_1} \cdots x_{n-2}^{m_{n-2}} \alpha^{m_{n-1}+m_n+1} dx_{n-2} \cdots dx_1 \\
&= \frac{m_{n-1}!m_n!(n-1)}{(m_{n-1} + m_n + 1)!} \left[(n-2)! \int_0^1 \cdots \int_0^\beta x_1^{m_1} \cdots x_{n-2}^{m_{n-2}} (\beta - x_{n-2})^{m_{n-1}+m_n+1} dx_{n-2} \cdots dx_1 \right] \\
&= \frac{m_{n-1}!m_n!(n-1)}{(m_{n-1} + m_n + 1)!} \cdot \frac{m_1! \cdots m_{n-2}!(m_{n-1} + m_n + 1)!(n-1-1)!}{(m_1 + \cdots + m_n + 1 + n - 1 - 1)!} \\
&= \frac{m_1! \cdots m_n!(n-1)!}{(m_1 + \cdots + m_n + n - 1)!}.
\end{aligned}$$

This completes the proof of equality of (1) and (4).

Finally, we prove the equality of (4) and (5) for $f(x) = e^{\langle y, x \rangle}$. Note that if $y_1 < \cdots < y_n$, then the result follows from the previous lemma. Otherwise, the expression in (4) (for this function f) is continuous in y_1, \dots, y_n , and so we can limit the expression for distinct eigenvalues. That said, we let $y'_1 < \cdots < y'_n$ be distinct values near to the y_i , and we apply L'Hoptial's rule to $\frac{\det(M(y'))}{\prod_{i < j} (y'_j - y'_i)}$ based on the multiplicities of the y_i . Specifically, for each $i \in [k]$ we apply the following differential operator to numerator and denominator (let $\partial_i := \partial_{y'_i}$):

$$D_i := \prod_{j=1}^{m_i} \partial_{m_1+\cdots+m_{i-1}+j}^{j-1} = \partial_{m_1+\cdots+m_{i-1}+1}^0 \partial_{m_1+\cdots+m_{i-1}+2}^1 \cdots \partial_{m_1+\cdots+m_i}^{m_i-1}.$$

The powers here correspond to the number of terms of the denominator of $\frac{\det(M(y'))}{\prod_{i < j} (y'_j - y'_i)}$ which will vanish when the m_i values of y'_1, \dots, y'_n limit to λ_i . That said, we now want to compute:

$$\frac{D_1 \cdots D_k \det(M(y'))}{D_1 \cdots D_k \prod_{i < j} (y'_j - y'_i)} \Big|_{y'=\lambda}.$$

We first compute the denominator via the product rule, noting that the only nonzero term occurs whenever all derivatives from a given D_i are applied to differences of eigenvalues corresponding to λ_i :

$$\text{denominator} = \left(\prod_{i=1}^k \prod_{p=0}^{m_i-1} p! \right) \cdot \prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}.$$

We next compute the numerator using the fact that exactly one row of the matrix depends on any given y'_i , as so we can apply the derivatives to the appropriate rows. Further, this means the numerator can still be expressed as a determinant. We also incorporate the factorials in the denominator expression above, by dividing each column by the appropriate factorial:

$$\frac{\text{numerator}}{\prod_{i=1}^k \prod_{p=0}^{m_i-1} p!} = \frac{D_1 \cdots D_k \det(M(y'))}{\prod_{i=1}^k \prod_{p=0}^{m_i-1} p!} \Big|_{y'=\lambda} = \det(M(y)).$$

This gives the result. ■

Remark 7.5 Even though we assume the distinct values of y_1, \dots, y_n to be in increasing order in the previous result and in Definition 7.1, we actually don't need this. Note that swapping the order of λ_i and λ_{i+1} affects the numerator and denominator of the expression (4) in the same way, by multiplication by $(-1)^{m_i m_{i+1}}$.

7.1.2 Computing the gradient

We now compute the gradient of $\mathcal{E}_1(Y)$ for $Y = \text{diag}(y)$ using the above formulas. The first thing to note is that we can use an argument similar to what we used in the proof of the evaluation formula. Specifically, note the following expression where $\partial_{y_l} := \frac{\partial}{\partial y_l}$:

$$\partial_{y_l} \mathcal{E}_1(Y) = \partial_{y_l} \log \int e^{-\langle Y, X \rangle} d\mu_1(X) = \frac{\int -X_{ll} \cdot e^{-\langle Y, X \rangle} d\mu_1(X)}{\int e^{-\langle Y, X \rangle} d\mu_1(X)}.$$

In particular, we obtain the following bound where $y_1 \geq \dots \geq y_n$ are the entries of diagonal Y :

$$|\partial_{y_l} \mathcal{E}_1(Y)| \leq \frac{\int e^{-\langle Y, X \rangle} d\mu_1(X)}{\int e^{-\langle y_l I_n, X \rangle} d\mu_1(X)} = e^{y_l} \int e^{-\langle Y, X \rangle} d\mu_1(X).$$

From these observations, we have that $\partial_{y_l} \mathcal{E}_1(Y)$ is continuous on diagonal matrices Y . Therefore, to compute the gradient we can first assume that y_l is distinct from the other diagonal entries, and then limit via L'Hopital's rule (as in the proof of the evaluation formula). We do exactly this to prove the gradient formula, after defining another parameterized matrix.

Definition 7.2 (Matrix for gradient formula, $k = 1$) Given $y_1, \dots, y_n \in \mathbb{R}$, let $\lambda_1 < \dots < \lambda_k$ denote the distinct values of y_i with multiplicities m_1, \dots, m_k . Given $p \in [k]$, we define an $n \times n$ matrix $M_p(y)$ as the matrix which differs from $M(y)$ in one column, given as follows:

$$M_p(y) := \begin{bmatrix} \dots & 1 & 0 & \dots & 0 & 0 & \dots \\ \dots & \lambda_p & 1 & \dots & 0 & 0 & \dots \\ \dots & \lambda_p^2 & \binom{2}{1} \lambda_p & \dots & 0 & 0 & \dots \\ \dots & \vdots & \vdots & \ddots & \vdots & \vdots & \dots \\ \dots & \lambda_p^{n-2} & \binom{n-2}{1} \lambda_p^{n-3} & \dots & \binom{n-2}{m_p-2} \lambda_p^{n-m_p} & \binom{n-2}{m_p} \lambda_p^{n-m_p-2} & \dots \\ \dots & \frac{e^{\lambda_p}}{0!} & \frac{e^{\lambda_p}}{1!} & \dots & \frac{e^{\lambda_p}}{(m_p-2)!} & \frac{e^{\lambda_p}}{m_p!} & \dots \end{bmatrix}.$$

That is, $\frac{\partial \lambda_p}{m_p}$ is applied to the right-most column of $M(y)$ that depends on λ_p .

Proposition 7.6 (Gradient formula, $k = 1$) Assume y_1, \dots, y_n are the diagonal values of diagonal Y , with distinct values $\lambda_1 > \dots > \lambda_k$ and multiplicities m_1, \dots, m_k . Letting p be such that $y_l = \lambda_p$, we have the following expression:

$$(\nabla \mathcal{E}_1(Y))_l = - \sum_{i \neq p} \frac{m_i}{\lambda_p - \lambda_i} - \frac{\det(M_p(-y))}{\det(M(-y))}.$$

Proof: We first assume that y_l is distinct from λ_p , and then we limit at the end. Specifically, we assume the distinct values of y_1, \dots, y_n are $\lambda_1 > \dots > \lambda_p > y_l > \lambda_{p+1} > \dots > \lambda_k$, where now the multiplicity of λ_p is now one less than it was originally. We let y' denote these new values of y (with y_l possibly changed) and let m'_i denote these new multiplicities (only m_p decreased by 1). We now want to compute:

$$\begin{aligned} \partial_{y_l} \mathcal{E}_1(Y)|_{y_l=\lambda_p} &= \partial_{y_l} \log \left(\frac{(n-1)!}{\prod_{i < j} (\lambda_i - \lambda_j)^{m'_i m'_j}} \cdot \frac{\det(M(-y'))}{\prod_{i \leq p} (\lambda_i - y_l)^{m'_i} \prod_{i > p} (y_l - \lambda_i)^{m'_i}} \right) \Big|_{y_l=\lambda_p} \\ &= \partial_{y_l} \left[\log \left(\frac{\det(M(-y'))}{(\lambda_p - y_l)^{m_p-1}} \right) - \sum_{i < p} m_i \log(\lambda_i - y_l) - \sum_{i > p} m_i \log(y_l - \lambda_i) \right] \Big|_{y_l=\lambda_p}. \end{aligned}$$

It is at this point that we limit y_l to λ_p and use the L'Hopital's rule argument. (Recall the above discussion which describes why this argument is valid.) We want to apply this argument to the following part of the above expression:

$$\partial_{y_l} \log \left(\frac{\det(M(-y'))}{(\lambda_p - y_l)^{m_p-1}} \right) \Big|_{y_l=\lambda_p} = \frac{(\lambda_p - y_l) \cdot \partial_{y_l} \det(M(-y')) + \det(M(-y')) \cdot (m_p - 1)}{\det(M(-y')) \cdot (\lambda_p - y_l)} \Big|_{y_l=\lambda_p}.$$

The key is to notice that the denominator contains exactly $m'_p + 1 = m_p$ factors of $(\lambda_p - y_l)$ up to scalar, where m'_p factors come from the determinant. With this, we apply ∂_{y_l} to the numerator and denominator $2m'_p$ times and then set $y_l = \lambda_p$. Computing this for the denominator is straightforward:

$$\text{denominator} = \partial_{y_l}^{m_p} \det(M(-y')) \cdot (\lambda_p - y_l) \Big|_{y_l=\lambda_p} = (-1)^{m_p} m_p! \det(M(-y)).$$

The computation is easy here for the same reason as in the proof of Proposition 7.4: using the product rule for all the derivatives only leaves a single term which does not evaluate to zero once we set $y_l = \lambda_p$. A similar thing happens for the numerator, which yields:

$$\begin{aligned} \text{numerator} &= \partial_{y_l}^{m_p} [(\lambda_p - y_l) \cdot \partial_{y_l} \det(M(-y')) + \det(M(-y')) \cdot (m_p - 1)] \Big|_{y_l=\lambda_p} \\ &= -\partial_{y_l}^{m_p} \det(M(-y')) \Big|_{y_l=\lambda_p} \\ &= (-1)^{m_p+1} m_p! \det(M_p(-y)). \end{aligned}$$

With this, we have the following expression:

$$\partial_{y_l} \mathcal{E}_1(Y) \Big|_{y_l=\lambda_p} = -\sum_{i \neq p} \frac{m_i}{\lambda_p - \lambda_i} - \frac{\det(M_p(-y))}{\det(M(-y))}.$$

This completes the proof. ■

7.2 Algorithm for $k > 1$

We now discuss how to generalize the formulas and arguments from the rank-1 case to the rank- k case. The computations done here are very similar to those given above, and so we will be a bit less explicit in what follows. And although the matrices involved are a bit more complicated (see Definitions 7.3 and 7.4), we still achieve the required bit complexity bounds. Specifically, each of the entries of these matrices require a polynomial number of computations of $m!$, y_i^m , and e^{-y_i} for $m \leq n$, and so the determinants can still be computed as efficiently as is necessary for Theorem 7.1.

We now state the explicit integral formulas for \mathcal{E}_k and $\nabla \mathcal{E}_k$ which generalize those of the $k = 1$ case of the previous section. Our main tool to prove these formulas is the Harish-Chandra-Itzykson-Zuber formula ([21], [24]), given as follows.

Theorem 7.7 (HCIZ formula) *For $n \times n$ Hermitian matrices Y and B with distinct eigenvalues $y_1 < \dots < y_n$ and $\beta_1 < \dots < \beta_n$ respectively, we have the following where μ is the Haar measure on the unitary group $U(n)$:*

$$\int_{U(n)} e^{\langle Y, UBU^* \rangle} d\mu(U) = \left(\prod_{p=1}^{n-1} p! \right) \frac{\det([e^{y_i \beta_j}]_{1 \leq i, j \leq n})}{\prod_{i < j} (y_j - y_i)(\beta_j - \beta_i)}.$$

Using the L'Hoptial's rule argument used in the rank-1 case, we can limit B to the rank- k PSD projection $\text{diag}(1, \dots, 1, 0, \dots, 0)$ to obtain a formula for $\mathcal{E}_k(Y)$ for Y with distinct eigenvalues. Using again the same sort of argument, we can then limit Y to any real diagonal matrix (eigenvalues not necessarily distinct). First, we need to define a parameterized matrix as in the rank-1 case.

Definition 7.3 (Matrix for dual integral, $k > 1$) *Given $y_1, \dots, y_n \in \mathbb{R}$, let $\lambda_1 < \dots < \lambda_k$ denote the distinct values of y_i with multiplicities m_1, \dots, m_k . Let the polynomial $q_{i,j}(t)$ be defined as follows:*

$$q_{i,j}(t) = e^{-t} \frac{\partial_t^j}{j!} (t^i e^t) = \sum_{l=0}^{\min(i,j)} \binom{i}{l} \frac{t^{i-l}}{(j-l)!}.$$

We define an $n \times n$ matrix $M^{(k)}(y)$ as follows:

$$M^{(k)}(y) := \begin{bmatrix} 1 & 0 & \dots & 0 & 1 & \dots \\ \lambda_1 & 1 & \dots & 0 & \lambda_2 & \dots \\ \lambda_1^2 & \binom{2}{1} \lambda_1 & \dots & 0 & \lambda_2^2 & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \dots \\ \lambda_1^{n-k-1} & \binom{n-k-1}{1} \lambda_1^{n-k-2} & \dots & \binom{n-k-1}{m_1-1} \lambda_1^{n-k-m_1} & \lambda_2^{n-k-1} & \dots \\ e^{\lambda_1} q_{0,0}(\lambda_1) & e^{\lambda_1} q_{0,1}(\lambda_1) & \dots & e^{\lambda_1} q_{0,m_1-1}(\lambda_1) & e^{\lambda_2} q_{0,0}(\lambda_2) & \dots \\ e^{\lambda_1} q_{1,0}(\lambda_1) & e^{\lambda_1} q_{1,1}(\lambda_1) & \dots & e^{\lambda_1} q_{1,m_1-1}(\lambda_1) & e^{\lambda_2} q_{1,0}(\lambda_2) & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \dots \\ e^{\lambda_1} q_{k-1,0}(\lambda_1) & e^{\lambda_1} q_{k-1,1}(\lambda_1) & \dots & e^{\lambda_1} q_{k-1,m_1-1}(\lambda_1) & e^{\lambda_2} q_{k-1,0}(\lambda_2) & \dots \end{bmatrix}.$$

Also, any term of the form λ_i^m for $m < 0$ in the above matrix should be replaced by 0.

The matrix defined above and the arguments of the previous section then allow us to write down an explicit formula for $\mathcal{E}_k(Y)$.

Corollary 7.8 (Evaluating the dual integral, $k > 1$) *Let Y be an $n \times n$ Hermitian matrix Y with eigenvalues y_1, \dots, y_n and distinct eigenvalues $\lambda_1 > \dots > \lambda_k$ with multiplicities m_1, \dots, m_k . We have the following:*

$$\int_{\mathcal{P}_k} e^{-\langle Y, X \rangle} d\mu_k(X) = \frac{\prod_{p=1}^{n-1} p!}{\prod_{p=1}^{n-k-1} p! \cdot \prod_{p=1}^{k-1} p!} \cdot \frac{\det(M^{(k)}(-y))}{\prod_{i < j} (\lambda_i - \lambda_j)^{m_i m_j}}.$$

This leads to a formula for $\mathcal{E}_k(Y)$:

$$\mathcal{E}_k(Y) = \log \left(\frac{\prod_{p=1}^{n-1} p!}{\prod_{p=1}^{n-k-1} p! \cdot \prod_{p=1}^{k-1} p!} \right) + \log \det(M^{(k)}(-y)) - \sum_{i < j} m_i m_j \log(\lambda_i - \lambda_j).$$

Notice that this reduces to (5) of Proposition 7.4 whenever $k = 1$. As in the $k = 1$ case, we use the Schur-Horn theorem and unitary invariance to restrict the inputs of $\mathcal{E}_k(Y)$ to real diagonal matrices (see Section B). Therefore, we only need to compute the gradient on the diagonal entries of Y . The arguments are essentially the same as those of the $k = 1$ case, again via L'Hoptial's rule, and so we state the gradient formula for \mathcal{E}_k as a corollary without proof. First though, we need to define another parameterized matrix for the gradient formula, as in the $k = 1$ case.

Definition 7.4 (Matrix for gradient formula, $k > 1$) Given $y_1, \dots, y_n \in \mathbb{R}$, let $\lambda_1 < \dots < \lambda_k$ denote the distinct values of y_i with multiplicities m_1, \dots, m_k . Let $q_{i,j}(t)$ be defined as in Definition 7.3. Given $p \in [k]$, we define an $n \times n$ matrix $M_p^{(k)}(y)$ as the matrix which differs from $M^{(k)}(y)$ in one column, given as follows:

$$M_p^{(k)}(y) := \begin{bmatrix} \dots & 1 & 0 & \dots & 0 & 0 & \dots \\ \dots & \lambda_p & 1 & \dots & 0 & 0 & \dots \\ \dots & \lambda_p^2 & \binom{2}{1} \lambda_p & \dots & 0 & 0 & \dots \\ \dots & \vdots & \vdots & \ddots & \vdots & \vdots & \dots \\ \dots & \lambda_p^{n-k-1} & \binom{n-k-1}{1} \lambda_p^{n-k-2} & \dots & \binom{n-k-1}{m_p-2} \lambda_p^{n-k-m_p+1} & \binom{n-k-1}{m_p} \lambda_p^{n-k-m_p-1} & \dots \\ \dots & e^{\lambda_p} q_{0,0}(\lambda_p) & e^{\lambda_p} q_{0,1}(\lambda_p) & \dots & e^{\lambda_p} q_{0,m_p-2}(\lambda_p) & e^{\lambda_p} q_{0,m_p}(\lambda_p) & \dots \\ \dots & e^{\lambda_p} q_{1,0}(\lambda_p) & e^{\lambda_p} q_{1,1}(\lambda_p) & \dots & e^{\lambda_p} q_{1,m_p-2}(\lambda_1) & e^{\lambda_p} q_{1,m_p}(\lambda_p) & \dots \\ \dots & \vdots & \vdots & \ddots & \vdots & \vdots & \dots \\ \dots & e^{\lambda_p} q_{k-1,0}(\lambda_p) & e^{\lambda_p} q_{k-1,1}(\lambda_p) & \dots & e^{\lambda_p} q_{k-1,m_p-2}(\lambda_p) & e^{\lambda_p} q_{k-1,m_p}(\lambda_p) & \dots \end{bmatrix}.$$

That is, $\frac{\partial \lambda_p}{m_p}$ is applied to the right-most column of $M^{(k)}(y)$ that depends on λ_p . As in Definition 7.3, any term of the form λ_i^m for $m < 0$ in the above matrix should be replaced by 0.

Corollary 7.9 (Gradient formula, $k > 1$) Assume y_1, \dots, y_n are the diagonal values of diagonal Y , with distinct values $\lambda_1 > \dots > \lambda_k$ and multiplicities m_1, \dots, m_k . Letting p be such that $y_1 = \lambda_p$, we have the following expression:

$$(\nabla \mathcal{E}_k(Y))_l = - \sum_{i \neq p} \frac{m_i}{\lambda_p - \lambda_i} - \frac{\det(M_p^{(k)}(-y))}{\det(M^{(k)}(-y))}.$$

7.3 Sampling from \mathcal{P}_1

Given some real diagonal matrix Y as in the previous section, we want to be able to sample from the measure on \mathcal{P}_1 given by $e^{-\langle Y, X \rangle} d\mu_1(X)$. It is not immediately obvious how to do this on \mathcal{P}_1 itself, so we instead transfer the measure to a simpler domain.

Specifically, we use the Proposition 7.4 to transfer the sampling problem to the simplex. Once on the simplex, we can apply standard techniques via the coordinate-wise cumulative distribution function (CDF). That said, we now state the sampling process for \mathcal{P}_1 and then use the rest of the section to fill in the details.

Proposition 7.10 (Rank-one Sampling) Let $Y = \text{diag}(y)$ be a real diagonal $n \times n$ matrix. The following process produces samples from the measure $e^{-\langle Y, X \rangle} d\mu_1(X)$ on \mathcal{P}_1 .

1. Sample v from the measure $e^{-\langle y, v \rangle} d\mu_{\Delta_1}(v)$ on the simplex Δ_1 by iteratively sampling v_i conditioned on v_1, \dots, v_{i-1} .
2. Sample z_1, \dots, z_n independently uniformly from the complex unit circle.
3. Construct $X := (z\sqrt{v})(z\sqrt{v})^* \in \mathcal{P}_1$ where $z\sqrt{v}$ is the column vector $(z_1\sqrt{v_1}, \dots, z_n\sqrt{v_n})$.

Note that step (1) is nontrivial, but we discuss how to sample coordinate-wise from the simplex below.

Proof: We give a proof sketch here, leaving the details to the remainder of this section. First, the reason we are able to reduce to sampling on the simplex is due to Proposition 7.11. Specifically, let $\Phi : \mathcal{P}_1 \rightarrow \Delta_1$ be given by $\Phi : X \mapsto \text{diag}(X)$. Then for any $v \in \Delta_1$, we have the following where \mathbb{T} is the complex unit circle:

$$\Phi^{-1}(v) = \{X \in \mathcal{P}_1 : v = \text{diag}(X)\} = \{X \in \mathcal{P}_1 : X = (z\sqrt{v})(z\sqrt{v})^* \text{ for } z \in \mathbb{T}^n\}.$$

The fact that $e^{-\langle Y, X \rangle} d\mu_1(X)$ is invariant under the action of conjugating X by $\text{diag}(z)$ for $z \in \mathbb{T}^n$ then implies that we can uniformly sample z from \mathbb{T}^n via Proposition 7.11.

Second, sampling from the measure $e^{-\langle Y, X \rangle} d\mu_1(X)$ is nontrivial, but doable by sampling each coordinate conditioned on the previous coordinates sampled. To do this we need to be able to efficiently compute the cumulative density function (CDF) for the conditioned measures, and we discuss how to do this below. Once we have this, we can sample each conditioned coordinate using standard techniques; see [31], Section 4.5. \blacksquare

For the case of μ_k for $k > 1$, we leave the question of sampling from the associated maximum entropy distributions as an open problem.

Transferring to the simplex. To transfer sampling from \mathcal{P}_1 to sampling from the simplex, we need a way of applying pushforward to sampling. The way to do this is via *disintegration* (see [9]), which we discuss in the following result.

Proposition 7.11 (Pushforward sampling) *Let X, Y be separable complete metric spaces, and let μ, ν be probability measures on X, Y respectively. Let $\Phi : X \rightarrow Y$ be a map such that ν is the pushforward measure of μ . Further, for any $y \in Y$, let μ_y denote the measure on the fiber $\Phi^{-1}(y)$ given by disintegration: i.e., such that $\int_X f(x) d\mu(x) = \int_Y \int_{\Phi^{-1}(y)} f(x) d\mu_y(x) d\nu(y)$ for all measurable f (see [9]). Then the measure on X generated by sampling y from (Y, ν) , followed by sampling x from $(\Phi^{-1}(y), \mu_y)$, is equal to μ .*

Proof: Let γ denote the measure on X generated by the described two-step sampling process. For any measurable set A we have the following, where P_1 and P_2 denote the probabilities according to the first and second steps of the process respectively:

$$\begin{aligned} \gamma(A) &= P_1(y \in \Phi(A)) \cdot P_2(x \in \Phi^{-1}(y) \cap A \mid y \in \Phi(A)) \\ &= \left(\int_{\Phi(A)} d\nu(y) \right) \cdot \frac{\int_{\Phi(A)} \int_{\Phi^{-1}(y) \cap A} d\mu_y(x) d\nu(y)}{\int_{\Phi(A)} \int_X d\mu_y(x) d\nu(y)} \\ &= \int_{\Phi(A)} \int_{\Phi^{-1}(y) \cap A} d\mu_y(x) d\nu(y). \end{aligned}$$

The second equality is just by definition of conditional probability. We then further have:

$$\begin{aligned} \int_{\Phi(A)} \int_{\Phi^{-1}(y) \cap A} d\mu_y(x) d\nu(y) &= \int_{\Phi(A)} \int_{\Phi^{-1}(y)} 1_A(x) d\mu_y(x) d\nu(y) \\ &= \int_Y \int_{\Phi^{-1}(y)} 1_A(x) d\mu_y(x) d\nu(y) \\ &= \mu(A). \end{aligned}$$

That is, $\gamma(A) = \mu(A)$. \blacksquare

Computing the conditioned CDF. We now compute the conditioned CDF for each coordinate of the measure on the simplex in Corollary 7.13, after a necessary lemma. Note that the formula below in Corollary 7.13 is not given in full explicit detail. However, the formula is still a constant times a determinant of a matrix, and expressions are given for the entries of that matrix. They are in fact rational functions of polynomials in β , y_i , $e^{y_i\beta}$, and factorials at most n (see below). Therefore, the whole determinant is computable in time polynomial in n and the number of bits needed to represent $e^{y_i\beta}$.

Lemma 7.12 *Fix $y_1, \dots, y_n \in \mathbb{R}$, and let $\lambda_1 < \dots < \lambda_k$ be the distinct values of y_i with multiplicities m_i . For valid $\gamma > 0$ and $x_n := 1 - \gamma - x_1 - \dots - x_{n-1}$, we have the following:*

$$\int_0^{1-\gamma} \int_0^{1-\gamma-x_1} \cdots \int_0^{1-\gamma-x_1-\dots-x_{n-2}} e^{\langle y, x \rangle} dx_{n-1} \cdots dx_1 = \frac{\det(M(y, 1 - \gamma))}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}}.$$

Moreover, only one of the rows of the matrix $M(y, 1 - \gamma)$ depends on γ .

Proof: The proof follows from a simple substitution ($u_i = \frac{x_i}{1-\gamma}$), applying Proposition 7.4, and then multiplying and dividing factors of $(1 - \gamma)$ in the rows and columns of $M((1 - \gamma)y)$ to obtain $M(y, 1 - \gamma)$.

To see this, we apply the change of variables $u_i = \frac{x_i}{1-\gamma}$ and Proposition 7.4:

$$\begin{aligned} \text{left-hand side} &= \int_0^1 \int_0^{1-u_1} \cdots \int_0^{1-u_1-\dots-u_{n-2}} e^{(1-\gamma)\langle y, u \rangle} (1 - \gamma)^{n-1} du_{n-1} \cdots du_1 \\ &= (1 - \gamma)^{n-1} \frac{\det(M((1 - \gamma)y))}{\prod_{i < j} (1 - \gamma)^{m_i m_j} (\lambda_j - \lambda_i)^{m_i m_j}}. \end{aligned}$$

Note now that we can do the following to $M((1 - \gamma)y)$ to make it so that only one of its rows depends on γ (recall the definition of $M(y)$ from Definition 7.1). First, divide the i th row of the matrix by $(1 - \gamma)^{i-1}$ up to $i = n - 1$. Then, for any λ_p multiply the j th column depending on λ_p by $(1 - \gamma)^{j-1}$. Only the last row of the matrix obtained will depend on γ , and in fact this matrix is precisely $M(y, 1 - \gamma)$.

The process described above is equivalent to pulling out of the determinant a factor of $(1 - \gamma)$ with the following exponent:

$$\text{exponent of factor} = \sum_{i=0}^{n-2} i - \sum_{p=1}^k \sum_{j=0}^{m_p-1} j = \binom{n-1}{2} - \sum_{p=1}^k \binom{m_p}{2}.$$

With this have that

$$\text{left-hand side} = (1 - \gamma)^\xi \cdot \frac{\det(M(y, 1 - \gamma))}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}},$$

where $\xi = (n - 1) - \sum_{i < j} m_i m_j + \binom{n-1}{2} - \sum_{p=1}^k \binom{m_p}{2}$. Note that

$$\xi = \frac{n(n-1)}{2} - \frac{1}{2} \left(\sum_{p=1}^k m_p \right)^2 + \frac{1}{2} \sum_{p=1}^k m_p = 0,$$

since $\sum_{p=1}^k m_p = n$. The result follows. ■

Corollary 7.13 (Conditioned CDF formula) Fix $y_1, \dots, y_n \in \mathbb{R}$, and let $\lambda_1 < \dots < \lambda_p$ be the distinct values of y_{k+1}, \dots, y_n with multiplicities m_i . Further, fix $x_1 = \alpha_1, \dots, x_{k-1} = \alpha_{k-1}$ and let $\alpha := \sum_{i=1}^{k-1} \alpha_i$. Also, let $x_n := 1 - \alpha - x_k - \dots - x_{n-1}$. The CDF denoted $F_k(\beta)$ for the simplex distribution $e^{\langle y, x \rangle}$, conditioned on the given values of x_1, \dots, x_{k-1} , is given as follows for $\beta \in [0, 1 - \alpha]$ and $y' := (y_{k+1}, \dots, y_n)$:

$$\begin{aligned} \frac{F_k(\beta)}{(n-1)!} &:= \int_0^\beta \int_0^{1-\alpha-x_k} \cdots \int_0^{1-\alpha-x_k-\dots-x_{n-2}} e^{\langle y, x \rangle} dx_{n-1} \cdots dx_k \\ &= \frac{e^{y_1\alpha_1+\dots+y_{k-1}\alpha_{k-1}}}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}} \int_0^\beta e^{y_k x_k} \det(M(y', 1 - \alpha - x_k)) dx_k. \end{aligned}$$

Recall the definition of $M(y, \gamma)$ from Definition 7.1. Since only the last row of $M(y', 1 - \alpha - x_k)$ depends on x_k , the above integral can be passed to that row and computed explicitly when $y_k \neq \lambda_l$:

$$\int_0^\beta e^{y_k x_k} \cdot \frac{(1 - \alpha - x_k)^j e^{(1 - \alpha - x_k) \lambda_l}}{j!} dx_k = e^{(1 - \alpha) \lambda_l} \left[\sum_{i=0}^j \frac{(1 - \alpha - \beta)^i e^{(y_k - \lambda_l) \beta} - (1 - \alpha)^i}{i! (y_k - \lambda_l)^{j-i+1}} \right].$$

If $y_k = \lambda_l$, we have the simpler expression, $e^{(1 - \alpha) \lambda_l} \left[\frac{(1 - \alpha)^{j+1} - (1 - \alpha - \beta)^{j+1}}{(j+1)!} \right]$.

Proof: We have:

$$\frac{F_k(\beta)}{(n-1)!} = e^{y_1\alpha_1+\dots} \int_0^\beta e^{y_k x_k} \left[\int_0^{1-\alpha-x_k} \cdots \int_0^{1-\alpha-x_k-\dots-x_{n-2}} e^{y_{k+1} x_{k+1} + \dots} dx_{n-1} \cdots dx_{k+1} \right] dx_k.$$

We compute the inner expression using the previous lemma and $\gamma = \alpha + x_k$:

$$\int_0^{1-\alpha-x_k} \cdots \int_0^{1-\alpha-x_k-\dots-x_{n-2}} e^{y_{k+1} x_{k+1} + \dots} dx_{n-1} \cdots dx_{k+1} = \frac{\det(M(y', 1 - \alpha - x_k))}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}}.$$

This then implies:

$$\frac{F_k(\beta)}{(n-1)!} = \frac{e^{y_1\alpha_1+\dots+y_{k-1}\alpha_{k-1}}}{\prod_{i < j} (\lambda_j - \lambda_i)^{m_i m_j}} \int_0^\beta e^{y_k x_k} \det(M(y', 1 - \alpha - x_k)) dx_k.$$

Since only one row of $M(y', 1 - \alpha - x_k)$ depends on x_k , we can compute the above integral entrywise on that row by linearity (after multiplying that row by the $e^{y_k x_k}$ factor). We now compute the final expression of the result, removing subscripts to simplify notation. First we make the change of variables $t = 1 - \alpha - x$:

$$\int_0^\beta e^{yx} \frac{(1 - \alpha - x)^j e^{(1 - \alpha - x) \lambda}}{j!} dx = \frac{e^{y(1 - \alpha)}}{j!} \int_{1-\alpha-\beta}^{1-\alpha} t^j e^{t(\lambda-y)} dt.$$

If $\lambda = y$, then we simply obtain $e^{(1 - \alpha) \lambda} \left[\frac{(1 - \alpha)^{j+1} - (1 - \alpha - \beta)^{j+1}}{(j+1)!} \right]$. Otherwise, we use integration by parts to obtain:

$$\frac{e^{y(1 - \alpha)}}{j!} \int_{1-\alpha-\beta}^{1-\alpha} t^j e^{t(\lambda-y)} dt = e^{(1 - \alpha) \lambda} \left[\sum_{i=0}^j \frac{(1 - \alpha - \beta)^i e^{(y - \lambda) \beta} - (1 - \alpha)^i}{i! (y - \lambda)^{j-i+1}} \right].$$

■

8 Computing maximum entropy measures

In this section we describe the entire algorithm for computing the optimum Y^* for the dual program $\text{Dual}_\mu(A)$, given some $A \in \mathcal{K} = \text{hull}(\Omega)$. The algorithm is essentially an application of the ellipsoid algorithm, based on a bounding box and a strong counting/integration oracle. We first discuss this algorithm in general, and then apply it to specific cases based on results from the previous sections.

Before moving on, we discuss how the linear equality constraints $\mathcal{L}(X) = B$ come into play here. We want to restrict our search space to the vector space $V_{\mathcal{L}}$ defined as the set of all X such that $\mathcal{L}(X) = 0$. The main reason for this is, since the constraints given by $\mathcal{L}(X) = B$ pick out an affine space in which \mathcal{K} is full dimensional, restricting the search space to $V_{\mathcal{L}}$ causes the optimum Y^* to be unique. Further, the bounding box results above apply specifically to this particular Y^* .

Since we are given \mathcal{L} effectively and explicitly, we assume for the ellipsoid algorithm that we can project the gradient (given by the strong counting oracle) onto $V_{\mathcal{L}}$. That said, we will from now on assume $V_{\mathcal{L}}$ to be the domain in which we are optimizing.

8.1 The ellipsoid framework

Using the standard argument via Hölder's inequality, we have that the dual objective function is convex:

$$F_A(Y) := \langle Y, A \rangle + \mathcal{E}_\mu(Y) = \langle Y, A \rangle + \log \left(\int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X) \right).$$

With this, the main optimization tool we use to approximate the the dual optimum Y^* is the ellipsoid algorithm. Recall the following from [36] Theorem 2.13, which was essentially taken from [5].

Theorem 8.1 (Ellipsoid algorithm) *Given any $\beta > 0$ and $R > 0$, there is an algorithm which, given a strong first-order oracle for F_A , returns a $Y^\circ \in V_{\mathcal{L}}$ such that:*

$$F_A(Y^\circ) \leq \inf_{Y \in V_{\mathcal{L}}, \|Y\|_\infty \leq R} F_A(Y) + \beta \left(\sup_{Y \in V_{\mathcal{L}}, \|Y\|_\infty \leq R} F_A(Y) - \inf_{Y \in V_{\mathcal{L}}, \|Y\|_\infty \leq R} F_A(Y) \right).$$

The number of calls to the strong first-order oracle for F_A is bounded by a polynomial in d , $\log R$, and $\log(1/\beta)$. Here, d is the dimension of the ambient Hilbert space in which Ω lies.

We now prove the main theorem (Theorem 4.4) regarding the existence of an algorithm for approximating the optimum to the dual objective.

Theorem 8.2 (Main algorithm, general case) *Let μ be a balanced measure on a domain $\Omega \subseteq \mathbb{R}^d$ contained in a ball of radius r . There exists an algorithm that, given a maximal set of linearly independent equalities $\mathcal{L}(X) = B$, an A in the η -interior of $\mathcal{K} = \text{hull}(\Omega)$, an $\varepsilon > 0$, and a strong counting/integration oracle for the exponential integral $\mathcal{E}_\mu(Y)$, returns $Y^\circ \in V_{\mathcal{L}}$ such that*

$$F_A(Y^\circ) \leq F_A(Y^*) + \varepsilon,$$

where F_A is the objective function for the dual program $\text{Dual}_\mu(A)$, and $Y^ \in V_{\mathcal{L}}$ is the optimum of the dual program. The running time of the algorithm is polynomial in d , η^{-1} , $\log(\varepsilon^{-1})$, $\log(r)$, and the number of bits needed to represent A , \mathcal{L} , and B .*

Proof: To apply the ellipsoid algorithm, we need to choose the two parameters, β and R . Since μ is balanced with some polynomial bound f , we choose for R the bounding box given for balanced measures in Theorem 4.2:

$$R := 2\eta^{-1} \cdot f(2\eta^{-1}, d).$$

So, the set $\{Y \in V_{\mathcal{L}} : \|Y\| \leq R\} \subset \{Y \in V_{\mathcal{L}} : \|Y\|_{\infty} \leq R\}$ contains the optimal Y^* for the dual program. Next, we need to choose β . Note that for $\|Y\|_{\infty} \leq R$ we have:

$$|F_A(Y)| \leq |\langle Y, A \rangle| + \left| \log \int e^{-\langle Y, X \rangle} d\mu(X) \right| \leq r\|Y\|_{\infty} + r\|Y\|_{\infty} \leq 2r\sqrt{d}\|Y\| \leq 2rR\sqrt{d}.$$

Therefore, choosing $\beta := \frac{\varepsilon}{4rR\sqrt{d}}$ implies:

$$\beta = \frac{\varepsilon}{4rR\sqrt{d}} \leq \frac{\varepsilon}{\sup_{Y \in V_{\mathcal{L}}, \|Y\|_{\infty} \leq R} F_A(Y) - \inf_{Y \in V_{\mathcal{L}}, \|Y\|_{\infty} \leq R} F_A(Y)}.$$

The ellipsoid algorithm then guarantees a Y° such that:

$$F_A(Y^{\circ}) \leq \inf_{Y \in V_{\mathcal{L}}, \|Y\|_{\infty} \leq R} F_A(Y) + \varepsilon = F_A(Y^*) + \varepsilon.$$

The number of calls to the strong counting oracle is polynomial in d , $\log(R) = \log(2\eta^{-1} \cdot f(2\eta^{-1}))$ and $\log(1/\beta) = \log(4rR\sqrt{d}\varepsilon^{-1})$. Given the bounding box, each oracle call (now including computing $\langle Y, A \rangle$) can be implemented in time polynomial in d , η^{-1} , and the number of bits needed to represent A . This completes the proof. \blacksquare

8.2 Rank- k Projections

Next we apply the above result to the case of $\Omega = \mathcal{P}_k$ and $\mu = \mu_k$, i.e., the case of rank- k projections. To do so we make a few tweaks to the proof of the theorem for the general algorithm given in the previous section. In particular, even though our domain \mathcal{P}_k lies in the space of Hermitian matrices, our strong counting oracle for \mathcal{E}_k only applies to real diagonal matrices Y . That said, we now state the theorem for rank- k projections and discuss such issues in the proof.

Corollary 8.3 (Main algorithm, \mathcal{P}_k case) *There exists an algorithm that, given $n \in \mathbb{N}$, $k \in [n]$, A in the η -interior of \mathcal{P}_k , and any $\varepsilon > 0$, returns Hermitian Y° such that*

$$F_A(Y^{\circ}) \leq F_A(Y^*) + \varepsilon,$$

where F_A is the objective function for the dual program $\text{Dual}_k(A)$, and Y^* is an optimum of the dual program. The running time of the algorithm is polynomial in n , η^{-1} , $\log(\varepsilon^{-1})$, and the number of bits need to represent A .

Proof: The result essentially follows from the general case, with a few details that need to be dealt with. First, the maximal linear equalities for \mathcal{P}_k boils down to something very simple within the space of Hermitian matrices. It is simply given by $\text{Tr}(X) = k$. Thus, our search space $V_{\mathcal{L}}$ then becomes the set of traceless Hermitian matrices.

Next, by unitary invariance of μ_k we can assume A is diagonal by unitary conjugation. Once we obtain an approximate optimum Y° for the diagonalized A , we can obtain an approximate optimum for the original A via conjugation by this unitary. Next, by the Schur-Horn theorem (see §B and the discussion at the start of §7) we can further assume that Y^* is diagonal. That is, we can assume A is real diagonal and restrict the domain of $F_A(Y)$ to real diagonal matrices Y .

Once we make this simplifying assumption, we have access to a strong counting/integration oracle for $\mathcal{E}_k(Y)$ by Theorem 4.5. The proof for the general case then goes through (using this strong counting oracle and the bounding box result for rank- k projections), giving the desired result. \blacksquare

9 The Goemans-Williamson measure

We discuss two main features of the pushforward through $v \mapsto vv^\top$ of the Goemans-Williamson measure which are relevant to this paper. We abuse notation in this section by letting μ_{GW} refer to the pushforward measure on \mathcal{V}_1 . First, we prove that this measure is a max-entropy measure with respect to \mathcal{V}_1 . Second, we demonstrate that this measure *cannot* be interpreted as a max-entropy measure on \mathcal{P}_1 . This second point demonstrates the fundamental difference between mex entropy measures on \mathcal{V}_1 and \mathcal{P}_1 .

9.1 Goemans-Williamson measure on \mathcal{V}_1 maximizes entropy

In this section, we demonstrate how the measure associated to the Goemans-Williamson SDP rounding scheme can be interpreted as a max-entropy measure. We describe it formally as follows.

Definition 9.1 (Goemans-Williamson rounding scheme) *Given an $n \times n$ real symmetric positive definite matrix A , let V be a real $n \times n$ matrix such that $VV^\top = A$. The Goemans-Williamson rounding scheme proceeds as follows:*

1. Sample a random standard Gaussian vector g from \mathbb{R}^n .
2. Return the rank-1 PSD matrix $(Vg)(Vg)^\top$.

The measure associated to this sampling process we refer to as the Goemans-Williamson measure and denote it μ_{GW} . This measure is supported on the rank-1 real symmetric PSD matrices, which is the set of extreme points of the real symmetric PSD cone.

Now let m be the Lebesgue measure on \mathbb{R}^n , and let μ be the measure on the real symmetric PSD cone which is the pushforward of m through the map $\Phi : x \mapsto xx^\top$. With this we can also give an explicit description of the Goemans-Williamson measure.

Proposition 9.1 (Goemans-Williamson density function) *The Goemans-Williamson measure on the set of rank-1 real symmetric PSD matrices is given by*

$$d\mu_{\text{GW}}(X) \propto e^{-\langle \frac{1}{2}A^{-1}, X \rangle} d\mu(X),$$

where μ is the pushforward of Lebesgue measure through $x \mapsto xx^\top$.

Proof: Let $A = VV^\top$ as in the definition of μ_{GW} . Since a standard Gaussian g is distributed according to $e^{-\frac{1}{2}\|g\|^2} dm(g)$, we can apply the change of variables formula to determine how $x := Vg$ is distributed. We have:

$$x \sim e^{-\frac{1}{2}\|V^{-1}x\|^2} \cdot \det(V^{-1}) dm(x) = e^{-\langle \frac{1}{2}A^{-1}, xx^\top \rangle} \cdot \sqrt{\det(A^{-1})} dm(x).$$

Considering the pushforward of this measure through $x \mapsto xx^\top$ gives the desired result. \blacksquare

Note that strong duality then immediately implies μ_{GW} is a max-entropy measure with respect to μ , since its density function is of the correct form. To demonstrate this more concretely, we prove this explicitly below via an explicit formula $\mathcal{E}_\mu(Y)$. First, the following observation tells us that it is sufficient to restrict $\mathcal{E}_\mu(Y)$ to positive definite Y .

Lemma 9.2 *If Y is not PD, then $\int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) = +\infty$.*

Proof: Since X is PSD, we have that $Y \prec Z$ implies $-\langle Y, X \rangle \geq -\langle Z, X \rangle$. Hence, to prove the result, we only need to show it for singular PSD matrices Y . Further, unitary invariance means we can restrict to diagonal Y . So, assume $Y = \text{diag}(0, y_2, \dots, y_n)$ for $y_i \geq 0$. Now consider:

$$\begin{aligned} \int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) &= \int_{\mathbb{R}^n} e^{-\langle Y, xx^\top \rangle} dx = \int_{\mathbb{R}^n} e^{-\sum_{i=2}^n y_i |x_i|^2} dx \\ &= \int_{-\infty}^{\infty} \int_{\mathbb{R}^{n-1}} e^{-\sum_{i=2}^n y_i |x_i|^2} d(x_2, \dots, x_n) dx_1 \\ &= \int_{-\infty}^{\infty} C dx_1 = +\infty. \end{aligned}$$

Note that the inner integrand above does not depend on x_1 , and so the evaluation of the inner integral yields some positive (possibly infinite) constant C as written above. \blacksquare

We now give an explicit formula for $\mathcal{E}_\mu(Y)$ on positive definite Y .

Proposition 9.3 (Lebesgue evaluation formula) *We have the following explicit expression for $\mathcal{E}_\mu(Y)$ for $n \times n$ real symmetric positive definite Y :*

$$\mathcal{E}_\mu(Y) := \log \int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) = \frac{n}{2} \log(\pi) - \frac{1}{2} \log \det(Y).$$

Proof: Since μ is the pushforward measure of m through $x \mapsto xx^\top$, we have:

$$\log \int_{\mathcal{V}_1} e^{-\langle Y, X \rangle} d\mu(X) = \log \int_{\mathbb{R}^n} e^{-\langle Y, xx^\top \rangle} dm(x) = \log \left(\pi^{n/2} \det(Y)^{-1/2} \right).$$

The second equality is computed via the density function of the multivariate Gaussian. \blacksquare

This then leads to the main result of this section.

Corollary 9.4 (Max-entropy, SDP rounding) *Given an $n \times n$ real symmetric positive definite marginals matrix A , the Goemans-Williamson measure μ_{GW} is the max-entropy measure with respect to μ , the pushforward through $x \mapsto xx^\top$ of the Lebesgue measure on \mathbb{R}^n . That is, μ_{GW} is the optimal measure for $\text{Prim}_\mu(A)$.*

Proof: Proposition 9.3 gives the following explicit expression for $\mathcal{E}_\mu(Y)$ with $n \times n$ real symmetric positive definite input Y :

$$\mathcal{E}_\mu(Y) := \log \int e^{-\langle Y, X \rangle} d\mu(X) = \frac{n}{2} \log(\pi) - \frac{1}{2} \log \det(Y).$$

By a standard computation, we then have the following:

$$\nabla \mathcal{E}_\mu(Y) = -\frac{1}{2} \nabla \log \det(Y) = -\frac{1}{2} Y^{-1}.$$

This implies the following regarding the gradient of the dual program objective $\text{Dual}_\mu(A)$ for positive definite A :

$$0 = \nabla F_A(Y) = \nabla(\langle Y, A \rangle + \mathcal{E}_\mu(Y)) = A - \frac{1}{2} Y^{-1} \iff Y = \frac{1}{2} A^{-1}.$$

That is, $Y^* = \frac{1}{2} A^{-1}$ is the optimum for the dual program. By strong duality/Slater condition for μ (see Proposition A.6) and the density function for μ_{GW} given in Proposition 9.1 above, this implies the result. \blacksquare

9.2 Goemans-Williamson measure projected to the unit sphere does not maximize entropy

In this section we show that the Hermitian version of the measure μ_{GW} is not a max-entropy measure for \mathcal{P}_1 . We do not recompute the density function for μ_{GW} in the Hermitian case, but only say that Proposition 9.1 can be adapted to show that in this case it is of the same form: $\nu(X) \propto e^{-\langle A_0, X \rangle}$ for some positive definite A_0 .

We want to “project” the (Hermitian) SDP rounding measure onto \mathcal{P}_1 , and we want to compute the density with respect to μ_1 . To do this, we first project the Lebesgue measure onto the complex unit sphere $S_{\mathbb{C}}^n$ and then pushforward through $x \mapsto xx^*$. We first state a few standard lemmas.

Lemma 9.5 *Let $f(z)$ be a Lebesgue measurable function on \mathbb{C}^n . Then:*

$$\int_{\mathbb{C}^n} f(z) dm(z) = \frac{2\pi^n}{(n-1)!} \int_{S_{\mathbb{C}}^n} \int_0^\infty f(rv) r^{2n-1} dr d\mu_{S_{\mathbb{C}}^n}(v).$$

Proof: This is precisely the polar coordinates formula for Lebesgue measure in $\mathbb{C}^n \cong \mathbb{R}^{2n}$. The constant $\frac{2\pi^n}{(n-1)!}$ is the volume of the complex unit ball in \mathbb{C}^n . \blacksquare

This shows that the projected density g can be computed from the Lebesgue density f as follows:

$$g(v) = \frac{2\pi^n}{(n-1)!} \int_0^\infty f(rv) r^{2n-1} dr.$$

We will now use the following lemma, which is standard.

Lemma 9.6 *For $n \in \mathbb{N}$ and $a > 0$, we have:*

$$\int_0^\infty r^{2n-1} e^{-ar^2} dr = \frac{(n-1)!}{2a^n}.$$

With this, we compute the following for $f(x) \sim e^{-\langle A, xx^* \rangle}$:

$$\begin{aligned} g(v) &= \frac{2\pi^n}{(n-1)!} \int_0^\infty f(rv) r^{2n-1} dr \\ &\propto \frac{2\pi^n}{(n-1)!} \int_0^\infty r^{2n-1} e^{-r^2 \langle A, vv^* \rangle} dr \\ &= \frac{\pi^n}{\langle A, vv^* \rangle^n} \propto \langle A, vv^* \rangle^{-n}. \end{aligned}$$

That is, the projected density is proportional to $\langle A, vv^* \rangle^{-n}$ on the unit sphere. With this, we have the following interesting fact.

Proposition 9.7 *The “projection” of the (Hermitian) SDP rounding measure to \mathcal{P}_1 is not a max-entropy measure with respect to μ_1 on \mathcal{P}_1 .*

Proof: By strong duality, max-entropy densities in both contexts take the form $g(X) \propto e^{-\langle A, X \rangle}$. So, we just need to show that for all PD B we have:

$$\langle A, X \rangle^{-n} \not\propto e^{-\langle B, X \rangle}.$$

This is straightforward, e.g. using the fact that the left-hand side is a rational function in $\Re(v_i), \Im(v_i)$ but the right-hand side is not. \blacksquare

10 Generalization of the maximum entropy framework to Lie groups

Recent work (e.g., [10, 8]) has demonstrated interesting connections between Lie theory and TCS, and the max-entropy framework fits into this context as well. In what follows we will briefly discuss the case of $\Omega = \mathcal{P}_k$ and $\mu = \mu_k$, as well as how this can be generalized. However, a more detailed investigation of the computational aspects of the max-entropy framework in this context is outside the scope of this paper.

We first describe the case of $\Omega = \mathcal{P}_k$ and $\mu = \mu_k$ in a more general way. The unitary group $U(n)$ acts on the real vector space of $n \times n$ Hermitian matrices by conjugation. This group action partitions the vector space into orbits, with X and Y being in the same orbit if and only if they have the same eigenvalues. Given any Hermitian matrix F , we denote the orbit corresponding to F by $\mathcal{O}(F)$.

Consider now the matrix $P_k := \text{diag}(1, \dots, 1, 0, \dots, 0)$ where k denotes the number of 1s that appear in the matrix. Then the orbit $\mathcal{O}(P_k)$ is precisely the set of rank- k projections. That is, $\mathcal{O}(P_k) = \mathcal{P}_k$, and so the unitarily invariant measure μ_k on \mathcal{P}_k induces such a measure on $\mathcal{O}(P_k)$. In fact such a unitarily invariant measure μ_F exists for any orbit $\mathcal{O}(F)$ allowing us to extend our maximum entropy framework to such orbits of $U(n)$.

This can be generalized beyond the group $U(n)$, to the general setting of a Lie group G and its corresponding Lie algebra \mathfrak{g} upon which G naturally acts. The primal and dual programs for this generalized setting are the same as in the general case, with one exception. The element $F \in \mathfrak{g}$ is now an input, and any algorithm for approximating an optimum for $\text{Dual}_{\mu_F}(A)$ will necessarily depend on the complexity of F . That said, strong duality holds in this case whenever A is in the interior of $\mathcal{K} = \text{hull}(\mathcal{O}(F)) \subset \mathfrak{g}$, and so the bounding box and the strong counting oracle are the two main results needed to obtain the polynomial-time ellipsoid-based algorithm described in this paper. As an aside, in this case $\mathcal{K} = \text{hull}(\mathcal{O}(F))$ is called an *orbitope* (e.g., see [34, 4]).

Thus, the following optimization problem is a natural generalization of the (dual) maximum entropy problem considered in this paper. The G -invariant inner product used in the exponent here can be derived from the so-called Killing form of \mathfrak{g} when G is compact (e.g., see [29], Corollary 4.26).

$$\inf_{Y \in \mathfrak{g}} F_A(Y) = \inf_{Y \in \mathfrak{g}} \left[\langle Y, A \rangle + \log \int_{\mathcal{O}(F)} e^{-\langle Y, X \rangle} d\mu_F(X) \right]$$

Computability of this problem will be a subject of future work.

Acknowledgments

The authors would like to thank Sushant Sachdeva, Sebastien Bubeck, and Umesh Vazirani for useful discussions. They would also like to thank Simons Institute for the Theory of Computing where this work was initiated. This research was partially supported by NSF CCF-1908347 grant and by Vetenskapsrådet.

References

- [1] Zeyuan Allen Zhu, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Much faster algorithms for matrix scaling. In *FOCS'17: Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, 2017.
- [2] Arash Asadpour, Michel X Goemans, Aleksander Madry, Shayan Oveis Gharan, and Amin Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. *Operations Research*, 65(4):1043–1061, 2017.
- [3] William Band and James L. Park. New information-theoretic foundations for quantum statistics. *Foundations of Physics*, 6(3):249–262, Jun 1976.
- [4] Alexander Barvinok and Isabella Novik. A centrally symmetric version of the cyclic polytope. *Discrete & Computational Geometry*, 39(1-3):76–99, 2008.
- [5] Aharon Ben-Tal and Arkadi Nemirovski. Optimization III: Convex analysis, nonlinear programming theory, nonlinear programming algorithms. Lecture Notes, 2012.
- [6] Christopher Bingham. An antipodally symmetric distribution on the sphere. *Ann. Statist.*, 2(6):1201–1225, 11 1974.
- [7] Sébastien Bubeck and Ronen Eldan. The entropic barrier: a simple and optimal universal self-concordant barrier. In *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 279–279, Paris, France, 03–06 Jul 2015. PMLR.
- [8] Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Mendes de Oliveira, Michael Walter, and Avi Wigderson. Efficient algorithms for tensor scaling, quantum marginals, and moment polytopes. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 883–897, 2018.
- [9] Joseph T Chang and David Pollard. Conditioning as disintegration. *Statistica Neerlandica*, 51(3):287–317, 1997.
- [10] Matthias Christandl, Brent Doran, Stavros Kousidis, and Michael Walter. Eigenvalue distributions of reduced density matrices. *Communications in mathematical physics*, 332(1):1–52, 2014.
- [11] Michael B. Cohen, Aleksander Madry, Dimitris Tsipras, and Adrian Vladu. Matrix scaling and balancing via box constrained Newton’s method and interior point methods. In *FOCS'17: Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, 2017.
- [12] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [13] J. J. Duistermaat and G. J. Heckman. On the variation in the cohomology of the symplectic form of the reduced phase space. *Inventiones mathematicae*, 69(2):259–268, Jun 1982.
- [14] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Operator scaling: theory and applications. *Foundations of Computational Mathematics*, pages 1–68, 2015.

- [15] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, November 1995.
- [16] M. Gromov. *Convex sets and Kahler manifolds*, pages 1–38. 1990.
- [17] Osman Güler. Barrier functions in interior point methods. *Mathematics of Operations Research*, 21(4):860–885, 1996.
- [18] Osman. Güler. On the self-concordance of the universal barrier function. *SIAM Journal on Optimization*, 7(2):295–303, 1997.
- [19] Osman Güler and Levent Tunçel. Characterization of the barrier parameter of homogeneous convex cones. *Mathematical Programming*, 81(1):55–76, Mar 1998.
- [20] Leonid Gurvits and Alex Samorodnitsky. A deterministic polynomial-time algorithm for approximating mixed discriminant and mixed volume, and a combinatorial corollary. *Discrete & Computational Geometry*, 27:531–550, 2002.
- [21] Harish-Chandra. Differential operators on a semisimple lie algebra. *American Journal of Mathematics*, 79(1):87–120, 1957.
- [22] Peter D. Hoff. Simulation of the matrix bingham-von mises-fisher distribution, with applications to multivariate and relational data. *Journal of Computational and Graphical Statistics*, 18(2):438–456, 2009.
- [23] Alfred Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *American Journal of Mathematics*, 76(3):620–630, 1954.
- [24] C. Itzykson and J. Zuber. The planar approximation. ii. *Journal of Mathematical Physics*, 21(3):411–421, 1980.
- [25] Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620–630, May 1957.
- [26] Edwin T. Jaynes. Information theory and statistical mechanics. II. *Physical Review*, 108:171–190, October 1957.
- [27] C. G. Khatri and K. V. Mardia. The von mises-fisher matrix distribution in orientation statistics. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):95–106, 1977.
- [28] B. Klartag. On convex perturbations with a bounded isotropic constant. *Geometric & Functional Analysis GAFA*, 16(6):1274–1290, Dec 2006.
- [29] Anthony W Knapp. *Lie groups beyond an introduction*, volume 140. Springer Science & Business Media, 2013.
- [30] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*, volume 13. Siam, 1994.
- [31] Art B Owen. Monte carlo theory, methods and examples. *Monte Carlo Theory, Methods and Examples. Art Owen*, 2013.

- [32] James L. Park and William Band. Rigorous information-theoretic derivation of quantum-statistical thermodynamics. i. *Foundations of Physics*, 7(3):233–244, Apr 1977.
- [33] Damián Pinasco. Lower bounds for norms of products of polynomials via bombieri inequality. *Transactions of the American Mathematical Society*, 364(8):3993–4010, 2012.
- [34] Raman Sanyal, Frank Sottile, and Bernd Sturmfels. Orbitopes. *Mathematika*, 57(2):275–314, 2011.
- [35] Issai Schur. Über eine klasse von mittelbildungen mit anwendungen auf die determinantentheorie. *Sitzungsberichte der Berliner Mathematischen Gesellschaft*, 22(9-20):51, 1923.
- [36] Mohit Singh and Nisheeth K Vishnoi. Entropy, optimization and counting. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 50–59. ACM, 2014.
- [37] Paul B. Slater. Relations between the barycentric and von neumann entropies of a density matrix. *Physics Letters A*, 159(8):411 – 414, 1991.
- [38] Damian Straszak and Nisheeth K. Vishnoi. Maximum entropy distributions: Bit complexity and stability. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2861–2891, Phoenix, USA, 25–28 Jun 2019. PMLR.
- [39] Michèle Vergne. Convex polytopes and quantization of symplectic manifolds. *Proceedings of the National Academy of Sciences*, 93(25):14238–14242, 1996.
- [40] J. von Neumann and R.T. Beyer. *Mathematical Foundations of Quantum Mechanics*. Goldstine Printed Materials. Princeton University Press, 1955.

A The dual formulation and strong duality

A.1 The dual formulation

The dual formulation $\text{Dual}_\mu(A)$ is given as follows, for $A \in \mathcal{K} = \text{hull}(\Omega)$ and Y in the ambient real inner product space \mathbb{R}^d :

$$\inf_Y F_A(Y) := \inf_Y \left[\langle Y, A \rangle + \log \int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X) \right].$$

In $\text{Dual}_\mu(A)$ we also assume a linear constraint on Y : $\mathcal{L}(Y) = 0$ where $X \in \Omega$ is such that $\mathcal{L}(X) = B$. We ignore this constraint for now, and deal with it in Lemma A.1 below.

To prove the form of the dual formulation given above, we write:

$$L(\nu, Y, z) = - \int \nu(X) \log(\nu(X)) d\mu(X) + \langle Y, A \rangle - \int \langle Y, X \rangle \nu(X) d\mu(X) + z - z \int \nu(X) d\mu(X).$$

We now want to compute derivatives to connect this with the dual program. For any $f \in L^2(\mu)$, we compute:

$$\begin{aligned} 0 &= \partial_t|_{t=0} L(\nu + tf, Y, z) \\ &= - \int f(X) \log(\nu(X)) d\mu(X) - \int f(X) d\mu(X) - \int \langle Y, X \rangle f(X) d\mu(X) - z \int f(X) d\mu(X) \\ &= - \int f(X) [\log(\nu(X)) + 1 + \langle Y, X \rangle + z] d\mu(X) \\ &= - \langle f, [\log(\nu(X)) + 1 + \langle Y, X \rangle + z] \rangle. \end{aligned}$$

This immediately implies (almost everywhere, and we will suppress this caveat from now on):

$$\log(\nu(X)) + 1 + \langle Y, X \rangle + z = 0.$$

This then gives

$$\nu(X) = \exp(-1 - z - \langle Y, X \rangle).$$

and therefore:

$$- \int [\log(\nu(X)) + \langle Y, X \rangle + z] \nu(X) d\mu(X) = \int \nu(X) d\mu(X) = \int \exp(-1 - z - \langle Y, X \rangle) d\mu(X).$$

Combining these observations:

$$\begin{aligned} L(\nu, Y, z) &= \int \exp(-1 - z - \langle Y, X \rangle) d\mu(X) + \langle Y, A \rangle + z \\ &= z + \langle Y, A \rangle + e^{-1-z} \int e^{-\langle Y, X \rangle} d\mu(X). \end{aligned}$$

Now, we compute:

$$\begin{aligned} 0 &= \partial_z L(\nu, Y, z) = 1 - e^{-1-z} \int e^{-\langle Y, X \rangle} d\mu(X) \\ &\implies \int e^{-\langle Y, X \rangle} d\mu(X) = e^{1+z} \\ &\implies z = \log \left(\int e^{-\langle Y, X \rangle} d\mu(X) \right) - 1. \end{aligned}$$

And finally

$$\inf_Y L(\nu, Y, z) = \inf_Y \left[\langle Y, A \rangle + \log \left(\int e^{-\langle Y, X \rangle} d\mu(X) \right) \right],$$

where Y ranges over \mathbb{R}^d .

Lemma A.1 Let $\mathcal{L}(X) = B$ be a set of linear constraints satisfied by all $X \in \Omega$. There exists an optimal solution Y^* to the dual program $\text{Dual}_\mu(A)$ if and only if there exists a solution Z^* to $\text{Dual}_\mu(A)$ restricted to $\mathcal{L}(Z^*) = 0$.

Proof: For any Y , consider the decomposition $Y = Z + Z^\perp$ where $\mathcal{L}(Z) = 0$ and $\langle Z^\perp, Y' \rangle = 0$ for all Y' such that $\mathcal{L}(Y') = 0$. Note that $A \in \Omega$ implies $\mathcal{L}(X - A) = 0$ for all $X \in \Omega$, and so $\langle Z^\perp, X - A \rangle = 0$. Letting $F_A(Y)$ denote the dual objective, this implies:

$$F_A(Y) = \log \int e^{-\langle Y, X - A \rangle} d\mu(X) = \log \int e^{-\langle Z, X - A \rangle} d\mu(X) = F_A(Z).$$

This completes the proof. \blacksquare

A.2 Strong duality under Slater's condition

We now prove a general result above obtaining strong duality from a Slater-type condition. In the next section, we show that this Slater-type condition holds for the max-entropy program in general. We also give more concrete proofs for μ_k on \mathcal{P}_k and μ on \mathcal{V}_1 in the following section.

Proposition A.2 (Strong duality under Slater's condition) Let V be a real inner product space such that $\Omega \subseteq V$. If for any A in the relative interior of $\mathcal{K} = \text{hull}(\Omega)$ there exists ν_A in the relative interior of the constraints of $\text{Prim}_\mu(A)$, then we have strong duality for any A in the relative interior of \mathcal{K} .

Proof: Without loss of generality, we may assume that Ω and A have been translated such that $\mathcal{K} = \text{hull}(\Omega)$ is full dimensional within a subspace of V , and that A also lies in this subspace. We will let W refer to this subspace, and now \mathcal{K} has nonempty interior within W .

We now follow some standard proofs of strong duality. Fix A in the interior of \mathcal{K} , and let $\mathcal{D} := \{\nu : \text{supp}(\mu) \rightarrow \mathbb{R}_+\}$. Then we only have linear equality constraints, which we denote collectively by:

$$E(\nu) := \left(\int X\nu(X)d\mu(X) - A, \int \nu(X)d\mu(X) - 1 \right) \in W \times \mathbb{R}.$$

We also negate the objective function, denoting this negated function by F , and p^* denote its optimal value. Note that p^* is finite by the assumptions of the theorem. Now define:

$$\mathcal{W} := \{(B, y, t) : F(\nu) \leq t \text{ and } E(\nu) = (B, y) \text{ for some } \nu \in \mathcal{D}\} \subseteq W \times \mathbb{R} \times \mathbb{R}.$$

Note that this set is convex. Further, $(0, 0, p^*)$ is either on the boundary of \mathcal{W} or outside of \mathcal{W} by optimality. Hence, we can find a separating hyperplane, and therefore there exists some $0 \neq (C_0, z_0, s_0) \in W \times \mathbb{R} \times \mathbb{R}$ such that:

$$\langle (C_0, z_0, s_0), (B, y, t) \rangle \geq s \cdot p^* \quad \text{for all } (B, y, t) \in \mathcal{V}.$$

Note that $s_0 \geq 0$, or else we can pick t very large to get a contradiction. We now demonstrate that in fact $s_0 > 0$. To get a contradiction, we suppose $s_0 = 0$ which gives:

$$\langle (C_0, z_0), (B, y) \rangle \geq 0 \quad \text{for all } (B, y, t) \in \mathcal{W}.$$

For any $B = A + Z$ where $Z \in B_\varepsilon(0) \subset W$, we have an interior solution ν_B by assumption. Therefore for all such Z :

$$\langle C_0, Z \rangle = \langle C_0, Z \rangle + z_0 \cdot 0 = \langle (C_0, z_0), (B, y) \rangle \geq 0.$$

This is only possible if $C_0 = 0$. By scaling ν_A , we also see that $z_0 = 0$. This contradicts $(C_0, z_0, s_0) \neq 0$, and therefore $s_0 > 0$.

To finish the proof, we define $C'_0 := \frac{C_0}{s_0}$ and $z'_0 := \frac{z_0}{s_0}$. We can then write:

$$\langle (C'_0, z'_0, 1), (B, y, t) \rangle \geq p^* \quad \text{for all } (B, y, t) \in \mathcal{W}.$$

Note that, for the dual objective function $G(C, z) := \inf_{\nu \in \mathcal{D}} [F(\nu) + \langle (C, z), E(\nu) \rangle]$, we have:

$$d^* := \sup_{C, z} G(C, z) \geq G(C'_0, z'_0) = \inf_{(B, y, t) \in \mathcal{W}} \langle (C'_0, z'_0, 1), (B, y, t) \rangle \geq p^*.$$

On the other hand, we have for all C, z and ν_A satisfying constraints such that $F(\nu_A)$ is near optimal:

$$G(C, z) = \inf_{\nu \in \mathcal{D}} [F(\nu) + \langle (C, z), E(\nu) \rangle] \leq F(\nu_A) + \langle (C, z), E(\nu_A) \rangle = p^* + \varepsilon + 0.$$

Applying sup to $G(C, z)$ and letting $\varepsilon \rightarrow 0$ implies $d^* \leq p^*$. ■

A.3 Slater's condition holds for general Ω and μ

The dual formulation given in Section A.1 implies a succinct representation of the optimal density ν , given that we have strong duality. Here we prove strong duality for general Ω and μ by proving Slater's condition. The main thing needed for this is existence of an optimal Y^* for $\text{Dual}_\mu(A)$, given interiority of A . We prove this now.

Lemma A.3 (Existence of dual optimum) *If A is in the interior of $\mathcal{K} = \text{hull}(\Omega)$, then there exists Y^* which optimizes the dual program $\text{Dual}_\mu(A)$.*

Proof: By the previous lemma, we may assume that \mathcal{K} is of full dimension in its ambient inner product space V . That said, note now that there is no closed half-space $H \subset V$ such that $A \in H$ and $\mu(H \cap \Omega) = 0$. Otherwise this would imply that A is not in the interior of \mathcal{K} , since the interior of H would not intersect Ω , the support of μ . Now suppose A is in the η -interior of \mathcal{K} . Hence, for any $B \in B_{\eta/2}(A)$ and any half-space H with B on the boundary, there is an $\varepsilon_{B,H} > 0$ such that $\mu(H \cap \Omega) = \varepsilon_{B,H}$.

We now prove that $\varepsilon := \inf \varepsilon_{B,H} > 0$. If not, then there is some sequence (B_i, H_i) for which $\varepsilon_{B_i, H_i} \rightarrow 0$. By identifying half-spaces about a point with the unit sphere, we have that the set of all possible (B, H) pairs is compact. Thus we can assume (B_i, H_i) is convergent, with limit (B_0, H_0) . Since every point X of the interior of H_0 is eventually in the interior of H_i , we have that the measure in a small ball around any such X is 0. Therefore the interior of H_0 does not intersect Ω , the support of μ . This implies B_0 is not in the interior of \mathcal{K} , a contradiction.

Therefore, $\varepsilon := \inf \varepsilon_{B,H} > 0$. This in fact implies that A is in the $(\frac{\eta}{2}, \varepsilon)$ -interior of μ (see Definition 6.2).

Using the arguments of Lemma 6.1, for any $Y \in V$ we have:

$$\begin{aligned} \varepsilon &\leq \mu(\{X \in \Omega \mid \langle -Y, X - (A - (\eta/2) \cdot Y/\|Y\|) \rangle \geq 0\}) \\ &= \mu(\{X \in \Omega \mid \langle -Y, X - A \rangle \geq (\eta/2) \cdot \|Y\|\}). \end{aligned}$$

This implies:

$$F_A(Y) = \log \int e^{\langle -Y, X - A \rangle} d\mu(X) \geq \log \left(\varepsilon \cdot e^{(\eta/2) \cdot \|Y\|} \right) = \frac{\eta}{2} \|Y\| + \log(\varepsilon).$$

Hence, $\|Y\| > R$ implies a lower bound on the dual objective $F_A(Y)$, which goes to infinity as $R \rightarrow \infty$. Therefore F_A must be minimized at some bounded point $Y^* \in V$. ■

This lemma then implies Slater's condition in general.

Theorem A.4 (Strong duality) *Fix any μ with support Ω in a real Hilbert space V . If A is in the interior of $\mathcal{K} = \text{hull}(\Omega)$, then strong duality holds for A . In particular, the optimum density ν^* is of the form:*

$$\nu^*(X) \propto e^{-\langle Y^*, X \rangle}$$

Proof: Let Y^* be an optimum for $\text{Dual}_\mu(A)$, by Lemma A.3. Then:

$$0 = \nabla F_A(Y^*) = \nabla \left[\langle Y, A \rangle + \log \int e^{-\langle Y, X \rangle} d\mu(X) \right]_{Y=Y^*} = A - \frac{\int X e^{-\langle Y^*, X \rangle} d\mu(X)}{\int e^{-\langle Y^*, X \rangle} d\mu(X)}.$$

This precisely says that A is the marginals matrix of the measure $\nu^*(X) \propto e^{-\langle Y^*, X \rangle}$. Therefore strong duality holds, since ν^* is in the relative interior of the constraints of $\text{Prim}_\mu(A)$ for any A in the interior of $\text{hull}(\Omega)$. \blacksquare

A.4 Slater's condition for \mathcal{P}_k and \mathcal{V}_1

We now give more concrete and direct arguments for Slater's condition in the specific situations of \mathcal{P}_k and \mathcal{V}_1 that we consider in this paper.

Proposition A.5 (Slater's condition for \mathcal{P}_k) *Let A be in the interior of $\text{hull}(\mathcal{P}_k)$. Then there is a density function on \mathcal{P}_k which is in the interior of the constraints of $\text{Prim}_k(A)$.*

Proof: Define $P := \text{diag}(1, \dots, 1, 0, \dots, 0) \in \mathcal{P}_k$ and $P^\perp := I - P$. $Y = \alpha P^\perp - \beta P$ for some $\alpha, \beta > 0$ to be determined. Hence, Y is $U(k) \times U(n-k)$ -invariant, and therefore by the unitary invariance of μ_k we have the following for any $U \in U(k) \times U(n-k)$:

$$B := \int X e^{-\langle Y, X \rangle} d\mu_k(X) = \int U X U^* e^{-\langle Y, U X U^* \rangle} d\mu_k(X) = U \left(\int X e^{-\langle Y, X \rangle} d\mu_k(X) \right) U^* = U B U^*.$$

That is, B is $U(k) \times U(n-k)$ -invariant, and therefore $B = \gamma P^\perp + \delta P$ for some $\gamma, \delta \in \mathbb{R}$.

By picking $\alpha, \beta > 0$ large with $\beta \ll \alpha$, the mass of $e^{-\langle Y, X \rangle} d\mu_k(X)$ becomes concentrated at $X = P$. If we also normalize by multiplying the density by $e^{-c\langle I_n, X \rangle}$ for appropriate values of $c \in \mathbb{R}$, we in fact have that B approaches P as $\alpha, \beta \rightarrow \infty$. Combining this with the form that B must take means that for every $\varepsilon > 0$, there exist α_ε and β_ε such that the corresponding measure is normalized and the corresponding value of B is equal to $(1 - \varepsilon)P + \frac{k\varepsilon}{n}I_n$. We refer to this matrix as B_ε , and we refer to the corresponding Y matrix as Y_ε .

Note also that for any unitary $U \in U(n)$, the same argument holds for

$$U B_\varepsilon U^* = (1 - \varepsilon) U P U^* + \frac{k\varepsilon}{n} I_n$$

and $U Y_\varepsilon U^*$. This then proves the result for $A = U B_\varepsilon U^*$ for any $\varepsilon > 0$ and any $U \in U(n)$. For any fixed $\varepsilon > 0$, we further have:

$$\text{hull}(\{U B_\varepsilon U^* : U \in U(n)\}) = (1 - \varepsilon) \text{hull}(\mathcal{P}_k) + \frac{k\varepsilon}{n} I_n =: \mathcal{P}_{k,\varepsilon}.$$

Hence, for any $A \in \mathcal{P}_{k,\varepsilon}$, we can choose $U_1, \dots, U_m \in U(n)$ such that $A = \frac{1}{m} \sum_{i=1}^m U_i B_\varepsilon U_i^*$. Therefore:

$$\int X \left(\frac{1}{m} \sum_{i=1}^m e^{-\langle U_i Y U_i^*, X \rangle} \right) d\mu_k(X) = \frac{1}{m} \sum_{i=1}^m U_i B_\varepsilon U_i^* = A.$$

Since $\frac{1}{m} \sum_{i=1}^m e^{-\langle U_i Y U_i^*, X \rangle} d\mu_k(X)$ is a convex combination of measures in the interior of the constraints of $\text{Prim}_k(A)$, this proves the result for all $A \in \mathcal{P}_{k,\varepsilon}$. Letting $\varepsilon \rightarrow 0$ then proves the result in full generality. \blacksquare

Proposition A.6 (Slater's condition for \mathcal{V}_1) *Let A be in the interior of the PSD cone, and let μ be the pushforward of the Lebesgue measure m through $x \mapsto xx^\top$. Then there is a density function on the set of rank-one real symmetric PSD matrices which is in the interior of the constraints of $\text{Prim}_\mu(A)$.*

Proof: Let $\nu_0(x)dm(x)$ be a Gaussian probability measure on \mathbb{R}^n with covariance matrix A . This precisely means:

$$\int xx^\top \nu_0(x)dm(x) = A.$$

Let $\nu(X)d\mu(X)$ be the pushforward of $\nu_0(x)dm(x)$ through the map $x \mapsto xx^\top$. Then:

$$\int X\nu(X)d\mu(X) = A.$$

This $\nu(X)$ is the desired density function. \blacksquare

B The Schur-Horn theorem

We last discuss an idea that will be useful to us in a number of parts of this paper. Generally, the idea is that the unitary invariance of μ_k allows us to often restrict to looking at diagonal matrices when considering the dual objective. The main observation is a corollary of the famous Schur-Horn theorem [35, 23].

Proposition B.1 (Schur-Horn) *If D is a real diagonal matrix and U is unitary, then the diagonal of UDU^* is majorized by the diagonal of D .*

Corollary B.2 *Given two real diagonal $n \times n$ matrices D, D' , we have the following:*

$$\inf_{U \in U(n)} \langle UDU^*, D' \rangle = \inf_{\sigma \in S_n} \langle \sigma D \sigma^*, D' \rangle.$$

Here, $U(n)$ is the unitary group and S_n is the subgroup of permutation matrices.

Proof: Let σ_0 be the permutation matrix which minimizes $\langle \sigma D \sigma^*, D' \rangle$ over all permutation matrices. By majorization, for any U the diagonal of UDU^* can be written as a convex combination of the permutations of the diagonal of D . By linearity of $\langle \cdot, D' \rangle$, the value of $\langle UDU^*, D' \rangle$ must then be at least the value of $\langle \sigma_0 D \sigma_0^*, D' \rangle$. \blacksquare

Corollary B.3 *Let A be a diagonal trace- k PD matrix. Then:*

$$\inf_{Y, \text{Hermitian}} F_A(Y) = \inf_{Y, \text{real diagonal}} F_A(Y).$$

Proof: Recall:

$$F_A(Y) = \langle Y, A \rangle + \log \int e^{-\langle Y, X \rangle} d\mu_k(X).$$

To prove the result, we only need to show that for any fixed real diagonal matrix D we have:

$$\inf_{U \in U(n)} F_A(UDU^*) = \inf_{\sigma \in S_n} F_A(\sigma D \sigma^*).$$

Since the integration part of $F_A(Y)$ is unitarily invariant, this is then equivalent to:

$$\inf_{U \in U(n)} \langle UDU^*, A \rangle = \inf_{\sigma \in S_n} \langle \sigma D\sigma^*, A \rangle.$$

Since A is diagonal, this follows from the previous corollary. \blacksquare

C Closeness of the approximate distribution

Let μ_1 and μ_2 two probability measures on Ω , given as density functions with respect to a base measure μ . The *KL divergence* between μ_1 and μ_2 is defined as

$$D_{\text{KL}}(\mu_1 \| \mu_2) := \int_{\Omega} \mu_1(X) \log \left(\frac{\mu_1(X)}{\mu_2(X)} \right) d\mu(X).$$

With this we follow the proof of Lemma A.4 in [36] to obtain the following.

Lemma C.1 *Let Y^* be the optimal solution to the dual objective function*

$$F_A(Y) = \langle Y, A \rangle + \log \int_{\Omega} e^{-\langle Y, X \rangle} d\mu(X).$$

Further, let Y° be such that $F_A(Y^\circ) \leq F_A(Y^) + \varepsilon$. If μ^* and μ° are the probability distributions associated to Y^* and Y° respectively, then*

$$D_{\text{KL}}(\mu^* \| \mu^\circ) = F_A(Y^\circ) - F_A(Y^*) \leq \varepsilon.$$

Proof: By assumption we have $F_A(Y^\circ) - F_A(Y^*) \leq \varepsilon$, which implies

$$\langle Y^\circ - Y^*, A \rangle + \log \int_{\Omega} e^{-\langle Y^\circ, X \rangle} d\mu(X) - \log \int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X) \leq \varepsilon.$$

The density functions of the distributions associated to Y^* and Y° can be given as

$$\mu^\circ(X) := \frac{e^{-\langle Y^\circ, X \rangle}}{\int_{\Omega} e^{-\langle Y^\circ, X \rangle} d\mu(X)} \quad \text{and} \quad \mu^*(X) := \frac{e^{-\langle Y^*, X \rangle}}{\int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X)}.$$

Since Y^* is the optimal solution (and hence A is proportional to $\int_{\Omega} X e^{-\langle Y^*, X \rangle} d\mu(X)$), we can compute the KL divergence as

$$\begin{aligned} D_{\text{KL}}(\mu^* \| \mu^\circ) &= \frac{\int_{\Omega} e^{-\langle Y^*, X \rangle} \left[\log \left(\frac{\int_{\Omega} e^{-\langle Y^\circ, Z \rangle} d\mu(Z)}{\int_{\Omega} e^{-\langle Y^*, Z \rangle} d\mu(Z)} \right) + \langle Y^\circ - Y^*, X \rangle \right] d\mu(X)}{\int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X)} \\ &= \log \left(\frac{\int_{\Omega} e^{-\langle Y^\circ, X \rangle} d\mu(X)}{\int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X)} \right) + \left\langle Y^\circ - Y^*, \frac{\int_{\Omega} X e^{-\langle Y^*, X \rangle} d\mu(X)}{\int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X)} \right\rangle \\ &= \left[\langle Y^\circ, A \rangle + \log \int_{\Omega} e^{-\langle Y^\circ, X \rangle} d\mu(X) \right] - \left[\langle Y^*, A \rangle + \log \int_{\Omega} e^{-\langle Y^*, X \rangle} d\mu(X) \right] \\ &= F_A(Y^\circ) - F_A(Y^*) \leq \varepsilon. \end{aligned}$$

As in Corollary A.5 of [36], we use the previous result to obtain bounds on the approximate optimal distribution and on the marginals of this distribution.

Corollary C.2 Let Y^* be the optimal to the dual objective function $F_A(Y)$ with domain Ω and measure μ as in the previous lemma, and let Y° be such that $F_A(Y^\circ) \leq F_A(Y^*) + \varepsilon$. If μ^* and μ° are the probability distributions associated to Y^* and Y° respectively, then

$$\|\mu^* - \mu^\circ\|_{\text{TV}} \leq \sqrt{2\varepsilon}.$$

Proof: The result follows from the previous lemma and the following well-known inequality (see e.g. [12], Lemma 12.6.1, pp. 300-301) relating KL divergence and total variation distance:

$$\|\mu^* - \mu^\circ\|_{\text{TV}} \leq \sqrt{2 \cdot D_{\text{KL}}(\mu^* \parallel \mu^\circ)}.$$

■