

Maximum Entropy Distributions: Bit Complexity and Stability

Damian Straszak

Aleph Zero Foundation

Nisheeth K. Vishnoi

Yale University

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

Maximum entropy distributions with discrete support in m dimensions arise in machine learning, statistics, information theory, and theoretical computer science. While structural and computational properties of max-entropy distributions have been extensively studied, basic questions such as: *Do max-entropy distributions over a large support (e.g., 2^m) with a specified marginal vector have succinct descriptions (polynomial-size in the input description)?* and: *Are entropy maximizing distributions “stable” under the perturbation of the marginal vector?* have resisted a rigorous resolution. Here we show that these questions are related and resolve both of them. Our main result shows a $\text{poly}(m, \log 1/\varepsilon)$ bound on the bit complexity of ε -optimal dual solutions to the maximum entropy convex program – for very general support sets and with no restriction on the marginal vector. Applications of this result include polynomial time algorithms to compute max-entropy distributions over several new and old polytopes for any marginal vector in a unified manner, a polynomial time algorithm to compute the Brascamp-Lieb constant in the rank-1 case. The proof of this result allows us to show that changing the marginal vector by δ changes the max-entropy distribution in the total variation distance roughly by a factor of $\text{poly}(m, \log 1/\delta)\sqrt{\delta}$ – even when the size of the support set is exponential. Together, our results put max-entropy distributions on a mathematically sound footing – these distributions are robust and computationally feasible models for data.

Keywords: Maximum entropy distributions, Bit complexity, Stability, Convex optimization

1. Introduction

The central objects of study in this paper are maximum entropy (max-entropy) distributions over discrete domains. For a finite set of vectors $\mathcal{F} \subseteq \mathbb{Z}^m$ and a point $\theta \in \mathbb{R}^m$, the max-entropy distribution corresponding to θ is defined to be the one that maximizes the entropy among all distributions $\{q_\alpha\}_{\alpha \in \mathcal{F}}$ whose expectation is θ . This is formulated as the following convex program:¹

$$\begin{aligned} \sup \quad & \sum_{\alpha \in \mathcal{F}} q_\alpha \log \frac{1}{q_\alpha}, \\ \text{s.t.} \quad & \sum_{\alpha \in \mathcal{F}} q_\alpha \cdot \alpha = \theta, \quad \sum_{\alpha \in \mathcal{F}} q_\alpha = 1, \quad q \geq 0. \end{aligned} \tag{1}$$

The model of a max-entropy distribution is natural to consider when one has an unknown distribution q on the set \mathcal{F} and the only information available is its marginal vector θ . Then, according to the

1. For simplicity, in the introduction we only consider the case with a uniform prior.

max-entropy principle (Jaynes (1957a,b)), the best guess for q is the max-entropy distribution with expectation θ . The rationale behind this choice is that maximizing the entropy yields a distribution with the least amount of prior information built-in.

In many interesting cases, the support \mathcal{F} is a large set and may be specified only implicitly (e.g., as the set of all subsets of a universe or the set of all spanning trees or matchings in a graph). Hence, the program (1) may have a prohibitively large number of variables, $|\mathcal{F}|$, that makes solving it intractable. Thus, one often considers the following Lagrange dual of (1) that involves only m variables.

$$\inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha \in \mathcal{F}} e^{\langle \alpha - \theta, y \rangle} \right). \quad (2)$$

Two questions arise: is max-entropy a “stable” model? In other words, how much does the optimal distribution (in total variation distribution) change when one changes the marginal vector? And, given the marginal vector, can one compute the max-entropy distribution? Both these questions are important from the point of view of applicability of the max-entropy model. Stability is an important consideration in machine learning (Bousquet and Elisseeff (2002); Yu (2013); Huang and Vishnoi (2019)): in particular, for the max-entropy problem, when the marginal vector comes from averaging “observed” samples, stability plays a role in deciding the number of samples to learn the distribution; see Dudík et al. (2007). Computability of max-entropy distributions enables one to estimate the max-entropy distribution corresponding to the above marginal vector and has been studied extensively in theoretical computer science (Singh and Vishnoi (2014)) and machine learning (Dudík et al. (2007)).

Computability and bit complexity. Duality implies that the optimal solution q^* to (1) is succinctly represented by the optimal dual solution (2) as follows: $q_\alpha^* \propto e^{\langle \alpha, y^* \rangle}$ for every $\alpha \in \mathcal{F}$. While the above provides a way to represent the max-entropy distribution q^* using only a vector of m real numbers, the issue becomes the bit complexity of the optimal solution y^* . Indeed, any algorithm that tries to compute y^* requires the knowledge of an upper bound $\|y^*\| \leq R$. Further, for the representation above to be efficient one requires $\|y^*\|$ to be bounded by a polynomial in m , as (even assuming $\|\alpha\| = O(1)$ for all $\alpha \in \mathcal{F}$) the value $e^{\langle \alpha, y^* \rangle}$ can be as large as $e^{\Omega(\|y^*\|)}$ and thus requires $\Omega(\|y^*\|)$ bits to store.

The first step to understand the bit complexity of max-entropy distributions for general polytopes was taken in Singh and Vishnoi (2014) where it was proved that whenever the point θ is “deep” inside the marginal polytope $P := \text{conv}\{\alpha : \alpha \in \mathcal{F}\}$, then a polynomial bound on y^* holds. More precisely, whenever $\theta \in P$ is at a distance $\eta > 0$ away from the boundary of P then $\|y^*\| \leq \frac{\log |\mathcal{F}|}{\eta}$.

However, in many applications one has to deal with instances where the point θ lies very close or even on the boundary of P . For instance, in machine learning and statistics (Pietra et al. (1997); Nigam (1999); Soofi (2000); Shewry and Wynn (1987); Dudík et al. (2007)), the point θ is often obtained as the empirical average of observed samples. Here, it might happen that the point θ lies in the interior with only exponentially small probability or may end up on the boundary of P even if P is rather “round”; see Section C. In TCS, several problems can also be seen to be disguised forms of the max-entropy problem and polynomial time algorithms for them necessarily require a polynomial bound on the size of the optimal or near-optimal dual solution. Examples include the problems of computing Brascamp-Lieb constants (Garg et al. (2017)) in the rank-1 setting, deterministically approximating mixed discriminants (Gurvits and Samorodnitsky (2002)), solving convex relaxations for counting and optimization problems (Nikolov and Singh (2016)).

In all of the above applications, the bound given by [Singh and Vishnoi \(2014\)](#) is not sufficient as it necessarily deteriorates to $+\infty$ when η tends to 0 and one needs the following stronger result: *For a given P , if y_ε^* is an ε -approximate solution to the dual problem, then $\|y_\varepsilon^*\| \leq R_\varepsilon = \text{poly}(m, \log 1/\varepsilon)$ – where the bound holds for **every** $\theta \in P$.* Our first result shows that the above is true as long as all the faces describing P have integer coefficients that are polynomially bounded by m . This includes all the aforementioned examples as special cases and also covers a large class of combinatorial polytopes (e.g., matroid polytopes, matroid intersection polytopes such as bipartite matching, general matching, submodular polytopes, Brascamp-Lieb polytopes). This result is obtained by a geometrical reasoning that involves convex analysis and polyhedral geometry: it relies on the fact that the faces describing the convex hull of the discrete domain have polynomially small coefficients, and avoids combinatorial arguments based on the specific structure of the domain. We also complement this result by providing a matching bit-complexity lower-bound for polytopes whose faces have large coefficients. Surprisingly, we observe that the bit complexity of max-entropy distributions also is related to its stability.

Stability. Using the techniques in the proof of our bit complexity result, we show that max-entropy distributions are “stable”: If we consider two marginal vectors θ_1 and θ_2 , then the total-variation distance between the arising max-entropy distributions is bounded by $\text{poly}(m, \log 1/\varepsilon) \cdot \sqrt{\|\theta_1 - \theta_2\|}$. Note that there is no a priori reason for the stability bound to be $\text{poly}(m, \log 1/\varepsilon)$ – it could be as large as $|\mathcal{F}|$ that could be exponential in m . While stability of max-entropy distributions is assumed to hold in practice, and is crucial for making the model meaningful, ours is the first result to formally establish polynomial stability for exponentially-sized support sets under fairly general conditions.

2. Our results

We provide simplified and slightly informal variants of our theorems here.

2.1. Bit complexity of max-entropy distributions

We assume that the marginal polytope has a low “unary facet complexity”, as formalized below².

Definition 1 (Unary facet complexity – Informal; see Definition 6) *Let $P \subseteq \mathbb{R}^m$ be a polytope with integer vertices. The **unary facet complexity** of P ($\text{fc}(P)$) is the smallest number $M \in \mathbb{N}$, such that P can be described by linear inequalities with coefficients in $\{-M, \dots, 0, \dots, M\}$.*

The class of polytopes with $\text{fc}(P) = O(1)$ is very general and includes most polytopes of interest including matroid polytopes, matroid intersection polytopes, matching polytope, submodular polytopes and Newton polytopes of certain classes of polynomials (such as real stable polynomials).

The first result of this paper is the below theorem, where $h_\theta(y)$ denotes the objective function in the dual program (2) and $g(\theta) \in \mathbb{R}$ denotes the optimal value; i.e., $g(\theta) := \inf_{y \in \mathbb{R}^m} h_\theta(y)$.

Theorem 2 (Bit complexity of max-entropy - Simplified; see Theorem 7) *There exists a bound $R_\varepsilon = \text{poly}(m, \log \frac{1}{\varepsilon}) \in \mathbb{R}_{>0}$, such that for every set $\mathcal{F} \subseteq \mathbb{Z}^m$ with unary facet complexity of $\text{conv}(\mathcal{F})$ polynomially bounded³, for every $\varepsilon > 0$ we have*

$$\forall \theta \in \text{conv}(\mathcal{F}), \exists y \in B(0, R), \text{ s.t. } h_\theta(y) \leq g(\theta) + \varepsilon.$$

2. Note that this condition is different from the *binary facet complexity* defined in [Grötschel et al. \(1988\)](#).

3. Specifically, we require the bound on the unary complexity to be m^c for some constant $c > 0$ fixed in advance.

Here $B(0, R)$ denotes the unit ℓ_2 -ball of radius R around the origin. Note that the bound we provide is *uniform* over the whole polytope P , meaning that the radius R_ε does not depend on θ – this is crucial for some of the applications. The error of ε that we introduce when going from the optimal solution y^* to a near-optimal solution y (in Theorem 2) translates to an ℓ_1 distance of at most $\sqrt{\varepsilon}$ between the max-entropy distribution q^* and the distribution q obtained from the near-optimal dual solution y ; see Corollary 9. This means that we do not introduce a large error when using approximate dual solutions and, importantly, we have a good control over this error as the bound on R_ε in Theorem 2 depends on ε logarithmically.

Theorem 2 implies polynomial time computability of maximum entropy distributions for polytopes with polynomial unary facet complexity and for all points $\theta \in \text{conv}(\mathcal{F})$ (assuming existence of a suitable counting oracle) – a question that was left open in the work of Singh and Vishnoi (2014).⁴ For a formal statement of this result and a proof, we refer to Section D.

This result has both practical and theoretical consequences. In practical situations, when the size of the discrete domain is large enough to force one to use the dual formulation of the max-entropy program, our bound may give a fast algorithm with provable guarantees. We show in Section C that if the vector θ is an empirical mean then, even in the simplest of settings, it is exponentially close to the boundary of the marginal polytope justifying the need for our result. Theoretically, this result helps derive a few algorithmic results in a unified manner. For applications of Theorem 2 to the matrix scaling problem, computing Brascamp-Lieb constants, and computing polynomial-based relaxations for counting problems, we refer to Section E. We also prove the following theorem that complements Theorem 2.

Theorem 3 (Lower bound) *For every bound $R = R(m, \log \frac{1}{\varepsilon}) \in \mathbb{R}_{>0}$, polynomial in m and $\log \frac{1}{\varepsilon}$ there exists an $m \in \mathbb{N}$, an $\varepsilon > 0$ and a set $\mathcal{F} \subseteq \mathbb{Z}^m$ with $\|\alpha\| \leq O(m^{3/2})$ for every $\alpha \in \mathcal{F}$ such that*

$$\exists \theta \in \text{conv}(\mathcal{F}) \quad \forall y \in B(0, R) \quad \text{s.t.} \quad h_\theta(y) > g(\theta) + \varepsilon.$$

The above result relies on existence of “flat” 0 – 1 polytopes by Alon and Vu Alon and Vu (1997), i.e., full-dimensional polytopes whose vertices are in $\{0, 1\}^m$ and which fit between two $e^{-\Theta(m \log m)}$ -close hyperplanes. Its proof appears in Section B. We note that the lower bound in Theorem 3 is based on a “bad example” $\mathcal{F} \subseteq \mathbb{Z}^m$ with $\|\alpha\| = O(m^{3/2})$ for every $\alpha \in \mathcal{F}$. It is an interesting open question whether similar examples can be constructed in the 0 – 1 regime, i.e., when $\mathcal{F} \subseteq \{0, 1\}^m$.

2.2. Stability of max-entropy distributions

Theorem 4 (Stability of max-entropy distributions - Simplified; see Theorem 10) *There exists a bound $R_\varepsilon = \text{poly}(m, \log \frac{1}{\varepsilon}) \in \mathbb{R}_{>0}$, such that for every set $\mathcal{F} \subseteq \mathbb{Z}^m$ with $\text{fc}(\text{conv}(\mathcal{F}))$ is polynomially bounded and for every $\theta_1, \theta_2 \in P$, if $\|\theta_1 - \theta_2\|_2 \leq \varepsilon$ then $\|q^{\theta_1} - q^{\theta_2}\|_1 \leq R_\varepsilon \sqrt{\varepsilon}$.*

In machine learning and statistics applications, θ is typically computed as an average of observed samples and our stability result implies that even if there is inverse polynomial (in m) variance in this empirical marginal, we can still recover a distribution close to the max-entropy distribution.

4. A variant of Theorem 2 which applies to families $\mathcal{F} \subseteq \mathbb{N}^m$ that are spanning trees of undirected graphs was obtained in Asadpour et al. (2010) and, for matroids (and more generally – jump systems) in Anari and Oveis Gharan (2017). However, these arguments heavily rely on the combinatorial fact that such families admit greedy optimization algorithms and do not seem to generalize beyond this setting.

Thus, this result shows in a rigorous manner and in very general, exponential-sized domains, that max-entropy distributions provide a stable model for density estimation. The proof of Theorem 4 crucially relies on the techniques that go in the proof of Theorem 2 and is presented in Section A. It would be interesting to improve $\sqrt{\varepsilon}$ to ε in our stability bound.

2.3. Computability of Brascamp-Lieb constants in the rank-1 regime

As a corollary of our bit complexity result we present an application to computing Brascamp-Lieb constants in the rank-1 regime. Brascamp-Lieb inequalities are an ultimate generalization of many inequalities used in analysis and all of mathematics, such as the Hölder inequality and Loomis-Whitney (Brascamp and Lieb (2002); Lieb (1990)). Recently, these inequalities have been studied from the computational point of view (Garg et al. (2017)), where the main problem is to compute quantities of the following form – called Brascamp-Lieb constants. Given a collection of matrices $B = (B_1, B_2, \dots, B_m)$ with $B_j \in \mathbb{R}^{n_j \times n}$ and a point $p \in \mathbb{R}_{\geq 0}^m$ compute

$$\text{BL}(B, p) := \inf \left\{ \frac{\det(\sum_{j=1}^m p_j B_j^\top X_j B_j)}{\prod_{j=1}^m \det(X_j)^{p_j}} : X_j \in \mathbb{R}^{n_j \times n_j}, X_j \succeq 0, j = 1, 2, \dots, m \right\}. \quad (3)$$

The constant $\text{BL}(B, p)$ is non-zero whenever p belongs to the so-called Brascamp-Lieb polytope $P_B \subseteq \mathbb{R}^m$:

$$P_B := \left\{ p \in \mathbb{R}_{\geq 0}^m : \sum_{j=1}^m p_j \dim(B_j U) \geq \dim(U), \text{ for every lin. subspace } U \subseteq \mathbb{R}^n \right\}.$$

Recently Garg et al. (2017) gave a method for calculating the Brascamp-Lieb constant in polynomial time when the vector p is rational and given in unary. Note that this does not imply a polynomial time algorithm in the classical sense (when the vector p is given in binary). Here we consider the special, but already non-trivial, case when the matrices are of rank 1; i.e., $B_j \in \mathbb{R}^{1 \times n}$ for every $j = 1, 2, \dots, m$. By interpreting Brascamp-Lieb constants in the rank-1 regime as solutions to certain entropy-maximization problems we can deduce from Theorem 2 that they can be calculated, up to a multiplicative precision $\varepsilon > 0$, in time polynomial in m and $\log \frac{1}{\varepsilon}$.

Our entropy interpretation also leads to an algorithm for computing worst-case Brascamp-Lieb constants over the whole Brascamp-Lieb polytope, i.e., $C(B) := \sup_{p \in P_B} \text{BL}(B, p)$, in the rank-1 regime. This quantity can be used as a universal constant (for any $p \in P_B$) for the reverse Brascamp-Lieb inequality (Barthe (1998)).

Theorem 5 (Computability of BL-constants) *Consider a sequence of matrices $B_1, B_2, \dots, B_m \in \mathbb{R}^{1 \times n}$. There is an algorithm that*

1. *computes a multiplicative $(1 + \varepsilon)$ -approximation of the the Brascamp-Lieb constant $\text{BL}(B, p)$ in time polynomial in the description size of $(B_1, B_2, \dots, B_m, p)$, m and $\log \frac{1}{\varepsilon}$,*
2. *computes a multiplicative $(1 + \varepsilon)$ -approximation of the worst-case Brascamp-Lieb constant $\sup_{p \in P_B} \text{BL}(B, p)$ in time polynomial in the description size of (B_1, B_2, \dots, B_m) , m and $\log \frac{1}{\varepsilon}$.*

For the proof, we refer the reader to Section E.3.

3. Technical overview

Theorem 2 – previous work and possible approaches. In the context of Theorem 2, the result of Singh and Vishnoi (2014) can be restated as follows: whenever a ball of radius $\eta > 0$ centered at θ is contained in $P := \text{conv}(\mathcal{F})$ the optimal solution y^* of (2) exists and its length $\|y^*\|$ is bounded by $O\left(\frac{\log |\mathcal{F}|}{\eta}\right)$. The proof is simple: the $\frac{1}{\eta}$ term comes from the fact that after scaling by $\log |\mathcal{F}|$, the optimal solution y^* of the dual program belongs to the polar of the radius- η ball around θ . (The polar of radius- η ball is itself a ball of radius $1/\eta$.)

One could consider the following two approaches to prove Theorem 2 that try to take advantage of the above mentioned bound: *centering* and *projection*. Both are based on slightly moving θ to a new point θ' so that the result by Singh and Vishnoi (2014) is applicable to θ' and then reasoning that the small shift does not affect the optimal dual solution. Centering is based on moving the point more towards the interior, e.g., by taking the centroid of the polytope P (which is far from the boundary) and taking a small step from θ towards θ' . One can then prove that θ' is well in the interior of the polytope, and hence a suitable bound for it follows. The second idea is based on projections: start with any point θ , if θ is already (inverse-polynomially) far from the boundary of P , then the result of Singh and Vishnoi (2014) implies a suitable bound. Otherwise, project the point θ onto the closest facet of P and recurse. By doing this, either the resulting new point θ' will end up in a vertex of P , or on a lower-dimensional face of P , where it is far from the boundary.

In both approaches, proving the result for the new point θ' is easy, given Singh and Vishnoi (2014); what remains is to bound the error introduced when moving from θ to θ' . Proving such a bound on the error turns out to be a non-trivial task. In the case of centering one has to pick an appropriate direction along which the point θ is moved. Similarly, the first challenge in the projection approach is to even define a suitable “projection operator” on a facet which would behave as expected and do not cause the point θ to land outside of the polytope P (as Euclidean projections might do). An issue which concerns both approaches is to deal with the behavior of the function g close to the boundary, where it can be shown to be non-Lipschitz⁵ and hence very susceptible to local perturbations. More precisely: moving from θ to θ' guarantees that the gradient $\nabla_y h_\theta(y)$ at $y = y'^*$ – the dual optimal solution at θ' – is small, however, it is a challenge to derive a guarantee on the gap $g(\theta) - h_\theta(y'^*)$ as the function $y \mapsto h_\theta(y)$ is not strongly convex⁶. In fact, even a weaker claim that $g(\theta) \approx g(\theta')$ does not follow from Singh and Vishnoi (2014), since the convexity argument (as in Lemma 7.5 in the full version of Singh and Vishnoi (2014)) only allows one to prove that $g(\theta') \geq g(\theta) - O(\varepsilon)$, where ε denotes the distance between θ and its “centering” θ' , but not in the opposite direction, as required.

Our proof of Theorem 2 is based on a purely geometric reasoning and bypasses the above obstacles by working entirely in the dual space (working in which we believe is necessary). This allows us to appropriately capture the geometry of sub-level sets of $h_\theta(y)$ and as a consequence, understand effects of seemingly large perturbations in y which only lead to small changes in the function value. When tracked in the primal domain, the proof resembles the “centering” idea,

5. In the simple case when $\mathcal{F} := \{0, 1\}$, the function $g : [0, 1] \rightarrow \mathbb{R}$ is of the form $g(\theta) = -\theta \log \theta - (1-\theta) \log(1-\theta)$.

One can see that on every interval $(\varepsilon, 1-\varepsilon)$ for $\varepsilon > 0$ the function g is Lipschitz, but it is not Lipschitz on $(0, 1)$.

6. One can again consider the case of $\mathcal{F} := \{0, 1\}$ – the function $h(0, y)$ is then of the form $h(0, y) = \log(1 + e^y)$.

Note that $h(0, y)$ is a convex function of y , but $\frac{d^2}{dy^2} h(0, y) \rightarrow 0$ whenever $y \rightarrow \pm\infty$, hence the function is essentially “flat” at infinity, and not strongly convex.

however, the implicit direction to move θ along is not easy to come up with (or analyze) just from the primal perspective.

Theorem 2 – proof overview. At a high level, in the proof we consider the optimal dual solution y^* and first identify a vertex α^* such that y^* belongs to C_{α^*} – the normal cone at α^* (i.e., the set of directions z such that all $\alpha^* + tz \in P$ for sufficiently small $t > 0$). Subsequently, a projection operation of y^* with respect to the cone C_{α^*} is used to find a vector y° – a witness for a short and close to optimal solution of the dual problem. The proof can be decomposed naturally into the following three steps, which we subsequently explain in more detail.

- (a) Identify a “good” basis for the dual space with respect to θ .
- (b) Truncate the optimal dual solution with respect to the basis in (a).
- (c) Establish a bound on the length of the truncated solution.

In what follows, we assume for simplicity that the polytope $\text{conv}(\mathcal{F})$ is full-dimensional. Given any $\theta \in P$ and a point $y^* \in \mathbb{R}^m$ which satisfies $h_\theta(y^*) \leq g(\theta) + \varepsilon/2$ (i.e., is close to optimal⁷) we aim to find a point y° whose length is polynomial in m and $\log 1/\varepsilon$ such that $h_\theta(y^\circ) \leq h_\theta(y^*) + \varepsilon/2$.

Step (a) Good basis. We first identify a subset $I_0 \subseteq I$ such that $\{a_i : i \in I_0\}$ is a basis of \mathbb{R}^m and y^* is expressed as a nonnegative linear combination in this basis, i.e., $y^* = \sum_{i \in I_0} \beta_i a_i$ with $\beta_i \geq 0$. The basis I_0 is chosen as a basis of tight constraints at a point $\alpha^* \in \mathcal{F}$, i.e., α^* satisfies $\langle a_i, \alpha^* \rangle = b_i$ for $i \in I_0$. This follows by selecting an α that maximizes $\langle \alpha, y^* \rangle$ over all $\alpha \in \mathcal{F}$ and invoking Farkas’ lemma and Caratheodory’s theorem. This part of the reasoning does not make any assumptions on the polytope and only relies on the convexity of P .

Step (b) Truncation of coefficients. Subsequently, we prove that the point $y^\circ = \sum_{i \in I_0} \min(\Delta, \beta_i) a_i$ satisfies the claim stated above, for a suitable choice of Δ , polynomial in the considered parameters. The bound $h_\theta(y^\circ) \leq h_\theta(y^*) + \varepsilon/2$ is proved by replacing β_i by $\min(\Delta, \beta_i)$ one by one and showing that the value $h_\theta(\cdot)$ does not increase by more than $\varepsilon/2m$. This relies on a careful analysis of the effect such a perturbation has on the function and crucially uses the fact that the coefficients β_i for $i \in I_0$ are nonnegative. Most importantly, we rely on the fact that the coefficients of the inequalities defining P are integral; and hence for any point $\alpha \in \mathcal{F}$ which does not lie on a facet $\langle a_i, x \rangle = b_i$ for some $i \in I_0$ we have $\langle \alpha^* - \alpha, a_i \rangle \geq 1$.

Step (c) Concluding a polynomial bound. To bound the length of y° note first that all the vectors $\{a_i\}_{i \in I_0}$ are short, i.e., $\|a_i\|_2 \leq \|a_i\|_\infty \cdot m \leq \text{fc}(P) \cdot m = \text{poly}(m)$, because we assume that the unary facet complexity of P is polynomially bounded. Further, from the triangle inequality, we obtain $\|y^\circ\| \leq m \cdot \Delta \cdot \text{poly}(m) = \text{poly}(m, \log \frac{1}{\varepsilon})$. Thus, we conclude the proof of Theorem 2.

Theorem 3 – proof overview. The proof of Theorem 3 is based on existence of so-called *flat* 0-1 polytopes. These are polytopes of the form $\text{conv}\{\alpha_0, \dots, \alpha_m\}$ with $\alpha_0, \dots, \alpha_m \in \{0, 1\}^m$ such that the distance from $\alpha_0 = 0$ to the affine subspace H generated by $\alpha_1, \dots, \alpha_m$ is exponentially small. The existence of such configurations was proved in Alon and Vu (1997). Given such a polytope we consider the lattice generated on H by the points $\alpha_1, \dots, \alpha_m$ and construct a new polytope by taking a certain finite subset of such lattice points and the point 0. The vertices of the new polytope are still integral and have relatively small entries (polynomial in m). Moreover, the projection of 0 onto H lies within the opposite facet. The family \mathcal{F} is defined to be all the vertices of the newly constructed polytope, θ is chosen to be 0 and we pick $\varepsilon \approx e^{-m}$.

To prove that for every short $y \in \mathbb{R}^m$ the gap between $h_\theta(y) - g(\theta)$ is significant we consider the gradient $\nabla_y h_\theta(y)$. Intuitively, if the gradient is large (in magnitude) at a point y , then y cannot

7. Note that there might not exist a point y^* such that $g(\theta) = h_\theta(y^*)$ when θ is on the boundary of P ; that is why we allow a slack of $\varepsilon/2$.

be an approximately optimal solution, hence it is enough to show that every short vector y admits a suitably long gradient. For this one can show that the gradient at y is given by $\theta - \theta^y$ where θ^y is the expectation of a distribution defined by y .

In order to make θ^y ε -close to $\theta = 0$, one has to ensure that $\langle 0, y \rangle - \langle \alpha, y \rangle \gtrsim \Omega(1)$ for all $\alpha \in \mathcal{F} \setminus \{0\}$. However, by introducing an auxiliary optimization problem (see Fact B.1) we show that this can happen only when $\|y\|$ is roughly, inverse-proportional to the distance from 0 to $\text{conv}(\mathcal{F} \setminus \{0\})$. Since the distance is exponentially small, we arrive at a lower-bound on $\|y\|$.

Theorem 4 – proof overview. The proof of Theorem 4 crucially relies on the bound established in Theorem 2. One starts by picking any “close-by” points θ_1, θ_2 in the polytope P , we denote $\|\theta_1 - \theta_2\| = \varepsilon$. The goal is to show that the distributions q^{θ_1} and q^{θ_2} are also close (the distance can be upper-bounded by some function of ε).

To this end, we first find dual solutions $y_1, y_2 \in \mathbb{R}^m$ that have low bit complexity and, roughly, have dual suboptimality gap $\leq \varepsilon$. If we denote by q^{y_1} and q^{y_2} the distributions defined by these two dual solutions, then one can show by a simple calculation that

$$KL(q^{\theta_i}, q^{y_i}) = h_{\theta_i}(y_i) - g(\theta_i) \leq \varepsilon \quad (4)$$

for $i = 1, 2$, and further, using Pinsker’s inequality $\|q^{\theta_i} - q^{y_i}\|_1 \leq \sqrt{\varepsilon}$. While this gives us some handle on q^{θ_1} and q^{θ_2} with respect to the dual solutions y_1, y_2 , we still do not have any control on the distance between q^{y_1} and q^{y_2} . To bypass this problem, we again take advantage of the fact in (4) and consider the gap $h_{\theta_1}(y_2) - g(\theta_1)$, since this is, roughly, an upper bound on the distance between q^{θ_1} and q^{y_2} , thus showing that it is small would mean we are done. To show that this suboptimality gap is small we use the fact that $h_{\theta_1}(y_1) - g(\theta_1) \leq \varepsilon$ and that $h_{\theta_1}(y_2) \approx h_{\theta_1}(y_1)$, more precisely (by convexity)

$$h_{\theta_1}(y_2) - h_{\theta_1}(y_1) \leq \|\nabla_y h_{\theta_1}(y_1)\|_2 \cdot \|y_1 - y_2\|_2.$$

Here $\|\nabla_y h_{\theta_1}(y_1)\|_2$ can be easily bounded by a function of ε – because $y \mapsto h_{\theta}(y)$ is a smooth function and y_1 is ε -close to the optimal solution. Further, crucially, $\|y_1 - y_2\|$ is bounded by a polynomial in m and $\log \frac{1}{\varepsilon}$, as a consequence of Theorem 2, and this allows us to arrive at the claimed bound.

4. Bit complexity upper bound for max-entropy distributions

We start by reintroducing several notions and restating the structural result, as in Section 2 only simplified or informal definitions and theorem statements were given. Consider a finite subset $\mathcal{F} \subseteq \mathbb{Z}^m$ of the integer lattice and a positive function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$. We will use p_α to denote the value of the function at a point $\alpha \in \mathcal{F}$, thus treating p as an $|\mathcal{F}|$ -dimensional vector with coordinates indexed by \mathcal{F} . For any $\theta \in \mathbb{R}^m$ and $y \in \mathbb{R}^m$ we define the following generalized max-entropy program

$$\begin{aligned} \max \quad & \sum_{\alpha \in \mathcal{F}} q_\alpha \log \frac{p_\alpha}{q_\alpha}, \\ \text{s.t.} \quad & \sum_{\alpha \in \mathcal{F}} q_\alpha \cdot \alpha = \theta, \quad \sum_{\alpha \in \mathcal{F}} q_\alpha = 1, \quad q \geq 0. \end{aligned} \quad (5)$$

In the case when p is normalized so that $\sum_{\alpha \in \mathcal{F}} p_\alpha = 1$, (5) describes the distribution closest to p (in KL -distance), whose expectation is equal to θ . In the case when $p \equiv 1$ (corresponds to the uniform distribution) the above program asks simply for a max-entropy distribution with expectation θ , as

in Section 2. Let us also extend the definition of the dual program, as introduced in Section 2 to capture general functions p not only $p \equiv 1$.

$$g(\theta) := \inf_{y \in \mathbb{R}^m} h_\theta(y) := \inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha - \theta, y \rangle} \right). \quad (6)$$

The following notion of facet complexity of a polytope plays an important role in our main result.

Definition 6 (Unary facet complexity) *Let $P \subseteq \mathbb{R}^m$ be a convex polytope with integer vertices. Let $M \in \mathbb{N}$ be the smallest integer such that P has a description of the form*

$$P = \{x \in \mathbb{R}^m : \langle a_i, x \rangle \leq b_i, \text{ for } i \in I\} \cap H \quad (7)$$

where I is a finite index set, $a_i \in \mathbb{Z}^m$, $\|a_i\|_\infty \leq M$ and $b_i \in \mathbb{R}$ for $i \in I$, and H is a linear subspace of \mathbb{R}^m . Then we call M the unary facet complexity of P and denote $\text{fc}(P) = M$.

Observe that in a description of a polytope P as in the definition above we could have also included H in the first term of (7), by adding $\langle c, x \rangle \leq d$ and $\langle c, x \rangle \geq d$, for every equation $\langle c, x \rangle = d$ defining H . However, the unary facet complexity defined as above might be significantly lower than one measured with respect to all (including equality) constraints and, as it turns out, this is the right measure to study the bit complexity of close to optimal solutions of (5).

To state our main result also a notion of *bit complexity* of a function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$ is required. We denote it by L_p and define as $L_p := \max_{\alpha \in \mathcal{F}} |\log p_\alpha|$. Note that L_p is finite, because we assume that $p_\alpha > 0$ for every $\alpha \in \mathcal{F}$. It represents, roughly, the maximum⁸ number of bits required to store the binary representation of any p_α (for $\alpha \in \mathcal{F}$).⁹

Theorem 7 (Bit complexity upper bound) *Let \mathcal{F} be any finite subset of \mathbb{Z}^m and let M be the unary facet complexity of $\text{conv}(\mathcal{F})$. Then, for every function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$ and for every $\varepsilon > 0$ there exists a number $R = O(m^{3/2} \cdot M \cdot (L_p + \log |\mathcal{F}| + \log \frac{m}{\varepsilon}))$ such that*

$$\forall \theta \in P \exists y \in B(0, R) \quad h_\theta(y) \leq g(\theta) + \varepsilon,$$

where h and g are defined as in (6).

We start by a preliminary lemma which is then used in the proof of Theorem 7.

Lemma 8 (Good basis) *Let $P \subseteq \mathbb{R}^m$ be a polytope $P := \{x \in \mathbb{R}^m : \langle a_i, x \rangle \leq b_i \text{ for } i \in I\} \cap H$, where I is a finite index set, $H \subseteq \mathbb{R}^m$ is a linear subspace of \mathbb{R}^m and $a_i \in H$, $b_i \in \mathbb{R}$ for $i \in I$. Let $y \in H$ be any vector, then there exists a vertex $v \in P$ and a subset of the constraints $I_0 \subseteq I$ of size $|I_0| \leq \dim(H)$, such that $\langle a_i, v \rangle = b_i$ for $i \in I_0$ and there exist non-negative numbers $\{\beta_i\}_{i \in I_0}$ satisfying $\sum_{i \in I_0} \beta_i a_i = y$.*

We note that in the case when $H = \mathbb{R}^m$, Lemma 8 can be also geometrically interpreted as existence of a separating hyperplane for the convex set P . In the proof of Theorem 7 we require this slightly more general variant where H is an arbitrary linear subspace.

Proof We start by picking $v \in P$ as an x that maximizes $\langle x, y \rangle$ over $x \in P$. Note that since P is a polytope (and is compact) such a maximum exists and, we might assume that v is a vertex of P .

8. Importantly, the complexity measure L_p is the *maximum* – not the *total* – number of bits required to store any of the coefficients. The latter is always at least $|\mathcal{F}|$, hence typically exponential in m .

9. In this paper, we use $\|x\|$ to denote the Euclidean ℓ_2 norm of x . This choice of a norm is by no means essential.

Given such a $v \in P$ determine all inequalities which are tight at v , i.e., $I^* = \{i \in I : \langle a_i, v \rangle = b_i\}$. Note that $|I^*|$ might be arbitrarily large, not even polynomially bounded. We claim that there exist $\{\beta_i\}_{i \in I^*}$ with $\beta_i \geq 0$ for all $i \in I^*$, such that $y = \sum_{i \in I^*} \beta_i a_i$. Suppose it is not the case. Then from the Farkas lemma, there exists a vector $z \in \mathbb{R}^m$ such that:

$$\forall_{i \in I^*} \langle z, a_i \rangle \leq 0 \quad \text{and} \quad \langle z, y \rangle > 0.$$

Note also that we may assume that $z \in H$, by projecting z orthogonally onto H if necessary. Further, the above is true also for $\delta \cdot z$ (in place of z) for an arbitrarily small $\delta > 0$. In other words, we can take z of arbitrarily small norm. Hence we obtain that the cone $C = \{u \in H : \forall_{i \in I^*} \langle u, a_i \rangle \leq 0\}$ contains arbitrarily short vectors z with $\langle z, y \rangle > 0$. Note that since I^* is the collection of all inequalities tight at v , it follows that every point in H , which is sufficiently close to v and satisfies the inequalities in I^* belongs itself to P . In other words, there exists a $\delta > 0$ (might be exponentially small) such that $(v + C) \cap B(v, \delta) \subseteq P$, where $v + C = \{v + u : u \in C\}$ and $B(v, \delta) = \{x \in \mathbb{R}^m : \|x - v\| \leq \delta\}$. Combining this with our previous observation regarding the cone C it follows that there exists $z \in H$ such that $\mu := v + z \in P$ and $\langle z, y \rangle > 0$ and hence $\langle \mu, y \rangle > \langle v, y \rangle$. This contradicts our choice of $v \in P$; $\langle \alpha, y \rangle > \langle v, y \rangle$. Knowing that y belongs to the cone generated by $\{a_i\}_{i \in I^*}$ we can apply Caratheodory's theorem to reduce the number of nonzero coefficients in the resulting conic combination. Indeed there exists a set $I_0 \subseteq I^*$, such that $|I_0| \leq \dim(H)$ and non-negative $\{\beta_i\}_{i \in I_0}$ (possibly different β_i s than obtained above) such that $y = \sum_{i \in I_0} \beta_i a_i$. \blacksquare

Proof of Theorem 7: Before we proceed with the argument let us first observe that one can assume without loss of generality that $0 \in \mathcal{F}$. This follows from the ‘‘shift invariance’’ of our problem. Indeed if we consider \mathcal{F} and $\mathcal{F}' = \mathcal{F} + \gamma = \{\alpha + \gamma : \alpha \in \mathcal{F}\}$, then the corresponding functions h and h' satisfy $h_\theta(y) = h'_{\theta+\gamma}(y)$ for every $y \in \mathbb{R}^m$. Hence, by shifting \mathcal{F} by $\gamma = -\alpha$ for some $\alpha \in \mathcal{F}$ we obtain an equivalent instance of our problem with $0 \in \mathcal{F}$. It follows in particular, that the affine subspace H on which P is full-dimensional is now a *linear subspace* of \mathbb{R}^m .

Fix $\theta \in P$ and let y^* be such that $h_\theta(y^*) \leq g(\theta) + \frac{\varepsilon}{2}$. Note that we may assume that $y^* \in H$, by projecting it orthogonally onto H (which does not alter the value). Further, note that by denoting by $a'_i \in H$ (for $i \in I$) the orthogonal projection of a_i onto H , the polytope P can be equivalently written as $P = \{x \in H : \langle a'_i, x \rangle \leq b_i\}$. By applying Lemma 8 for y^* and P we obtain a vertex $\alpha^* \in \mathcal{F}$ and a subset of constraints I_0 of size at most $\dim(H) \leq m$, tight at α^* such that $y^* = \sum_{i \in I_0} \beta_i a'_i$, for some non-negative scalars $\{\beta_i\}_{i \in I_0}$. We prove that by modifying the coefficients in the above conic combination we can obtain a point $y' = \sum_{i \in I_0} \beta'_i a'_i$ (with $\beta'_i \geq 0$ for $i \in I_0$) such that the norm of y' is small (polynomial in $L_p, m, \log d, M$ and $\log \frac{1}{\varepsilon}$) and

$$h_\theta(y') \leq h_\theta(y^*) + \frac{\varepsilon}{2} \leq g(\theta) + \varepsilon.$$

Showing that will complete the argument. Let $\Delta > 0$ be a certain number (to be specified later), polynomial in L_p and $\log \frac{1}{\varepsilon}$. We define $\beta'_i := \min(\Delta, \beta_i)$ and prove that the point $y' = \sum_{i \in I_0} \beta'_i a'_i$ satisfies the above claim.

To this end, we prove that by changing one coordinate i_0 , from $\beta_{i_0} > \Delta$ to Δ we cause only a slight increase in the value of $h_\theta(y)$. In other words, by taking y^* as before and $y' = y^* - (\beta_{i_0} - \Delta)a_{i_0}$ we want to show that

$$\log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \right) \leq \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y^* \rangle - \langle \theta, y^* \rangle} \right) + \frac{\varepsilon}{2m}$$

Towards this, define $\mathcal{F}_0 = \{\alpha \in \mathcal{F} : \langle \alpha, a_{i_0} \rangle = b_{i_0}\}$. Below we analyze separately the effect of changing y to y' on the terms $p_\alpha e^{\langle \alpha, y \rangle}$ for $\alpha \in \mathcal{F}_0$ and for $\alpha \in \mathcal{F} \setminus \mathcal{F}_0$.

Case 1: $\mathcal{F} \setminus \mathcal{F}_0$. Consider any $\alpha \in \mathcal{F} \setminus \mathcal{F}_0$. We have

$$\begin{aligned} \langle \alpha, y' \rangle - \langle \alpha^*, y' \rangle &= \sum_{i \in I_0} \beta'_i \langle \alpha - \alpha^*, a'_i \rangle \\ &= \sum_{i \in I_0} \beta'_i \langle \alpha - \alpha^*, a_i \rangle \\ &\leq \Delta \langle \alpha - \alpha^*, a_{i_0} \rangle \\ &\leq -\Delta. \end{aligned}$$

In the above, we used the fact that a'_i is the projection of a_i onto H , $\langle \alpha - \alpha^*, a_i \rangle \leq 0$ for every $i \in I_0$ and that $\langle \alpha - \alpha^*, a_{i_0} \rangle$ is a negative integer. This implies in particular, that

$$\frac{p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle}}{p_{\alpha^*} e^{\langle \alpha^*, y' \rangle - \langle \theta, y' \rangle}} \leq \frac{p_\alpha}{p_{\alpha^*}} e^{-\Delta}$$

and hence:

$$\begin{aligned} \sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} &= \sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} + \sum_{\alpha \in \mathcal{F} \setminus \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \\ &\leq \left(\sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \right) \left(1 + |\mathcal{F} \setminus \mathcal{F}_0| \frac{\max_{\alpha \in \mathcal{F} \setminus \mathcal{F}_0} p_\alpha}{p_{\alpha^*}} e^{-\Delta} \right) \end{aligned}$$

Note that for any $\alpha \in \mathcal{F}$ we have $e^{-L_p} \leq p_\alpha \leq e^{L_p}$, hence we can pick $\Delta := \log |\mathcal{F}| + 2L_p + \log m + \log \frac{1}{\varepsilon}$ to guarantee

$$|\mathcal{F} \setminus \mathcal{F}_0| \frac{\max_{\alpha \in \mathcal{F} \setminus \mathcal{F}_0} p_\alpha}{p_{\alpha^*}} e^{-\Delta} \leq \frac{\varepsilon}{2m}.$$

For such a choice of Δ we have:

$$h_\theta(y') \leq \log \left(\sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \right) + \log \left(1 + \frac{\varepsilon}{2m} \right) \leq \log \left(\sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \right) + \frac{\varepsilon}{2m}.$$

Case 2: \mathcal{F}_0 . Consider now $\alpha \in \mathcal{F}_0$, we have

$$\langle \alpha, y' \rangle - \langle \theta, y' \rangle = \langle \alpha, y^* \rangle - \langle \theta, y^* \rangle - (\beta_i - \Delta) \langle \alpha - \theta, a_i \rangle \leq \langle \alpha, y^* \rangle - \langle \theta, y^* \rangle$$

as $\langle \theta, a_i \rangle \leq \langle \alpha, a_i \rangle = b_i$ (because $\theta \in P$). Consequently, we obtain

$$\begin{aligned} h_\theta(y') &\leq \log \left(\sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y' \rangle - \langle \theta, y' \rangle} \right) + \frac{\varepsilon}{2m} \\ &\leq \log \left(\sum_{\alpha \in \mathcal{F}_0} p_\alpha e^{\langle \alpha, y^* \rangle - \langle \theta, y^* \rangle} \right) + \frac{\varepsilon}{2m} \\ &\leq h_\theta(y^*) + \frac{\varepsilon}{2m}. \end{aligned}$$

It remains to argue that after performing the above procedure, the norm of y' is small. We have

$$\|y'\| = \left\| \sum_{i \in I_0} \beta'_i a'_i \right\| \leq \sum_{i \in I_0} \beta'_i \|a'_i\| \leq m \cdot \Delta \cdot (\sqrt{m} \cdot M) = m^{3/2} M \Delta.$$

In the above, we used the fact that since a'_i is a projection of a_i onto H (for any $i \in I_0$) we have $\|a'_i\| \leq \|a_i\| \leq \sqrt{m} \cdot M$. \blacksquare

Below we present a useful corollary of Theorem 7 which is often convenient in applications. To state it, denote by $d \in \mathbb{R}_{\geq 0}$ the diameter (in the Euclidean norm) of the set \mathcal{F} .

Corollary 9 *Under the assumptions of Theorem 7, for every function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$, for every $\varepsilon > 0$ there exists a number $R > 0$ which is polynomial in $m, \log d, M, L_p$ and $\log \frac{1}{\varepsilon}$ such that*

$$\forall \theta \in P \exists y \in B(0, R) \left\| \frac{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \cdot \alpha}{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle}} - \theta \right\| < \varepsilon.$$

Proof of Corollary 9: It is enough to establish it for $\theta \in \text{int}(P)$ (relative interior is meant here, i.e., interior of P when restricted to H). To this end, consider y^* to be $y^* = \text{argmin}_{y \in \mathbb{R}^m} h_\theta(y) = g(\theta)$. It is not hard to prove that such a y^* exists, i.e., the minimum is attained, for θ in the relative interior of P . Consider now the gradient of h with θ fixed:

$$\nabla_y h_\theta(y) = \frac{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \cdot \alpha}{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle}} - \theta. \quad (8)$$

To conclude the corollary from Theorem 7 one has to prove that if the value at a point is close to optimal, then the gradient is short. Towards this, we show that $y \mapsto h_\theta(y)$ is L -smooth (for some polynomially bounded L), i.e., $\|\nabla_y h_\theta(y_1) - \nabla_y h_\theta(y_2)\| \leq L \|y_1 - y_2\|$. This can be deduced from the fact that the Hessian matrix of h , $\nabla_y^2 h_\theta(y)$ has polynomially bounded entries (and the bound depends neither on y nor on θ). Indeed, under the notation that $\alpha' = \alpha - \theta$ for $\alpha \in \mathcal{F}$ we obtain

$$\left(\nabla_y^2 h_\theta(y) \right)_{i,j} = \frac{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \alpha'_i \alpha'_j}{\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle}} - \frac{\left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \alpha'_i \right) \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \alpha'_j \right)}{\left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \right)^2},$$

hence it follows

$$\left| \left(\nabla_y^2 h_\theta(y) \right)_{i,j} \right| \leq 2 \max_{\alpha \in \mathcal{F}} \|\alpha - \theta\|^2 \leq 2d^2.$$

Hence the function $y \mapsto h_\theta(y)$ is $L := 2d^2$ -smooth. Now, it is well known that for a convex, L -smooth function, we have (see e.g. [Nesterov \(2014\)](#)):

$$h \left(\theta, y + \frac{1}{2L} v \right) \leq h_\theta(y) - \frac{1}{4L} \|v\|^2,$$

where $v = \nabla_y h_\theta(y)$. Hence, if y is as in Theorem 7 then we obtain $\|\nabla_y h_\theta(y)\|^2 \leq 4L\varepsilon$ and consequently

$$\|\nabla_y h_\theta(y)\| \leq 4d\varepsilon^{1/2}.$$

The corollary follows by combining the above obtained bound with (8). \blacksquare

References

- 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.
- Noga Alon and Van H. Vu. Anti-Hadamard matrices, coin weighing, threshold gates, and indecomposable hypergraphs. *J. Comb. Theory, Ser. A*, 79(1):133–160, 1997. doi: 10.1006/jcta.1997.2780. URL <https://doi.org/10.1006/jcta.1997.2780>.
- Nima Anari and Shayan Oveis Gharan. A generalization of permanent inequalities and applications in counting and optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 384–396, 2017. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055469. URL <http://doi.acm.org/10.1145/3055399.3055469>.
- 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. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 379–389, 2010.
- Franck Barthe. On a reverse form of the brascamp-lieb inequality. *Inventiones mathematicae*, 134(2):335–361, Oct 1998.
- Aharon Ben-Tal and Arkadi Nemirovski. Optimization III: Convex analysis, nonlinear programming theory, nonlinear programming algorithms. Lecture Notes, 2012.
- André Bouchet and William H. Cunningham. Delta-matroids, jump systems, and bisubmodular polyhedra. *SIAM J. Discrete Math.*, 8(1):17–32, 1995. doi: 10.1137/S0895480191222926. URL <http://dx.doi.org/10.1137/S0895480191222926>.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- Petter Branden. Polynomials with the half-plane property and matroid theory. *Advances in Mathematics*, 216(1):302 – 320, 2007.
- Herm Jan Brascamp and Elliott H. Lieb. *Best constants in Young's inequality, its converse, and its generalization to more than three functions*, pages 417–439. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.
- L. Elisa Celis, Amit Deshpande, Tarun Kathuria, Damian Straszak, and Nisheeth K. Vishnoi. On the complexity of constrained determinantal point processes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA*, pages 36:1–36:22, 2017. doi: 10.4230/LIPIcs.APPROX-RANDOM.2017.36. URL <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2017.36>.
- L. Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. Fair and Diverse DPP-Based Data Summarization. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 715–724. PMLR, 2018.

- 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.
- Miroslav Dudík, Steven J Phillips, and Robert E Schapire. Maximum entropy density estimation with generalized regularization and an application to species distribution modeling. *Journal of Machine Learning Research*, 8(Jun):1217–1260, 2007.
- Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Algorithmic and optimization aspects of Brascamp-Lieb inequalities, via operator scaling. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing, Montreal, QC, Canada, June 19-23, 2017*, pages 397–409, 2017. doi: 10.1145/3055399.3055458. URL <http://doi.acm.org/10.1145/3055399.3055458>.
- Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 1988. ISBN 3-540-13624-X, 0-387-13624-X (U.S.).
- Osman Güler. Hyperbolic polynomials and interior point methods for convex programming. *Math. Oper. Res.*, 22(2):350–377, 1997. doi: 10.1287/moor.22.2.350. URL <https://doi.org/10.1287/moor.22.2.350>.
- Leonid Gurvits. Hyperbolic polynomials approach to Van der Waerden/Schrijver-Valiant like conjectures: sharper bounds, simpler proofs and algorithmic applications. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 417–426. ACM, 2006.
- 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.
- Leonid Gurvits and Alex Samorodnitsky. Bounds on the permanent and some applications. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 90–99, 2014. doi: 10.1109/FOCS.2014.18. URL <https://doi.org/10.1109/FOCS.2014.18>.
- Lingxiao Huang and Nisheeth K. Vishnoi. Stable and Fair Classification. In *ICML*, 2019.
- Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620–630, May 1957a. doi: 10.1103/PhysRev.106.620.
- Edwin T. Jaynes. Information theory and statistical mechanics. II. *Physical Review*, 108:171–190, October 1957b. doi: 10.1103/PhysRev.108.171.
- Bahman Kalantari, Isabella Lari, Federica Ricca, and Bruno Simeone. On the complexity of general matrix scaling and entropy minimization via the RAS algorithm. *Math. Program.*, 112(2): 371–401, 2008. doi: 10.1007/s10107-006-0021-4. URL <https://doi.org/10.1007/s10107-006-0021-4>.
- Alex Kulesza and Ben Taskar. *Determinantal point processes for machine learning*. Now Publishers Inc., Hanover, MA, USA, 2012. ISBN 1601986289, 9781601986283.

- Elliott H. Lieb. Gaussian kernels have only gaussian maximizers. *Inventiones mathematicae*, 102(1):179–208, Dec 1990.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*. Springer Publishing Company, Incorporated, 1 edition, 2014. ISBN 1461346916, 9781461346913.
- Kamal Nigam. Using maximum entropy for text classification. In *In IJCAI-99 Workshop on Machine Learning for Information Filtering*, pages 61–67, 1999.
- Aleksandar Nikolov and Mohit Singh. Maximizing determinants under partition constraints. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 192–201, 2016.
- Stephen Della Pietra, Vincent J. Della Pietra, and John D. Lafferty. Inducing features of random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(4):380–393, 1997. doi: 10.1109/34.588021. URL <https://doi.org/10.1109/34.588021>.
- M. C. Shewry and H. P. Wynn. Maximum entropy sampling. *Journal of Applied Statistics*, 14(2):165–170, 1987. doi: 10.1080/02664768700000020.
- 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.
- Ehsan S Soofi. Principal information theoretic approaches. *Journal of the American Statistical Association*, 95(452):1349–1353, 2000.
- Damian Straszak and Nisheeth K. Vishnoi. Real stable polynomials and matroids: Optimization and counting. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*, pages 370–383, 2017. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055457. URL <http://doi.acm.org/10.1145/3055399.3055457>.
- David Wagner. Multivariate stable polynomials: theory and applications. *Bulletin of the American Mathematical Society*, 48(1):53–84, 2011.
- Bin Yu. Stability. *Bernoulli*, 19(4):1484–1500, 09 2013. doi: 10.3150/13-BEJSP14. URL <https://doi.org/10.3150/13-BEJSP14>.

Appendix A. Stability of max-entropy distributions

The purpose of this section is to first restate Theorem 4 in a more general (weighted) form and then provide a proof of this generalized theorem. We let $\mathcal{F} \subseteq \mathbb{Z}^m$ to be any finite family of integer vectors and fix any function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$. As usually, we denote the convex hull of \mathcal{F} by P . For brevity, we introduce the following handy notation:

1. Given $\theta \in P$ we denote by q^θ the optimal solution of the max-entropy convex program (5).
2. Given $y \in \mathbb{R}^m$ let q^y be the distribution over \mathcal{F} defined as

$$q_\alpha^y \propto p_\alpha e^{\langle y, \alpha \rangle} \quad \text{for } \alpha \in \mathcal{F}.$$

Moreover, we let $\theta^y := \sum_{\alpha \in \mathcal{F}} q_\alpha^y \alpha$.

Note in particular that if the dual program (6) has an optimal solution y then the KKT conditions imply that $q^y = q^\theta$ and it follows from feasibility that $\theta^y = \theta$. We are thus ready to state the generalized version of Theorem 4.

Theorem 10 (Polynomial stability of max-entropy distributions) *Let \mathcal{F} be any finite subset of \mathbb{Z}^m and let $d \in \mathbb{R}_{\geq 0}$ be its diameter (in the ℓ_2 -norm) and let M be the unary facet complexity of $\text{conv}(\mathcal{F})$. Then, for every function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$ and for every $\varepsilon > 0$ there exists a number $R = O(m^{3/2} \cdot M \cdot (L_p + \log |\mathcal{F}| + \log \frac{dm}{\varepsilon}))$ such that for every $\theta_1, \theta_2 \in P$, if $\|\theta_1 - \theta_2\|_1 \leq \varepsilon$ then*

$$\left\| q^{\theta_1} - q^{\theta_2} \right\|_1 \leq \sqrt{R} \sqrt{\varepsilon}.$$

Before we proceed to the proof we need to show a few simple properties of the max-entropy convex program and especially how do the primal and dual variables relate with each other. We begin by an elementary, yet important technical lemma. Below, whenever we refer to $h_\theta(y)$ or $g(\theta)$ we mean the dual objective, and its minimum respectively, as defined in (6), where \mathcal{F} and p will be clear from the context.

Lemma 11 (KL-distance between feasible distributions of the max-entropy program) *Consider any finite family $\mathcal{F} \subseteq \mathbb{Z}^m$ along with a positive function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$ then, for any $\theta \in P$, any $y \in \mathbb{R}^m$ and any distribution q that is feasible for the max-entropy program (5), we have*

$$KL(q, q^y) = h_\theta(y) - \sum_{\alpha \in \mathcal{F}} q_\alpha \log \frac{p_\alpha}{q_\alpha}.$$

Proof The proof follows by a direct calculation, it is enough to use the fact that $\sum_{\alpha \in \mathcal{F}} q_\alpha \alpha = \theta$. ■

The next step is to show some basic properties of close-to-optimal solutions to the dual program, and more precisely how does a small suboptimality gap in the dual sense translate to the primal world.

Lemma 12 (Basic properties of close-to-optimal solutions to the dual program) *Let \mathcal{F} be any finite subset of \mathbb{Z}^m and let $d \in \mathbb{R}_{\geq 0}$ be its diameter, let M be the unary facet complexity of $\text{conv}(\mathcal{F})$. Then, for every function $p : \mathcal{F} \rightarrow \mathbb{R}_{>0}$, for every $\theta \in P$ and for every $\varepsilon > 0$ if for $y \in \mathbb{R}^m$ we have $h_\theta(y) \leq g(\theta) + \varepsilon$ then*

1. $\|q^\theta - q^y\|_1 \leq \sqrt{2\varepsilon}$.
2. $\|\theta^y - \theta\|_2 \leq 2d\sqrt{\varepsilon}$.

Proof

Part 1. We start by applying Lemma 11 for $q := q^\theta$ to obtain

$$h_\theta(y) - g(\theta) = KL(q^\theta, q^y)$$

and thus by Pinsker inequality

$$\left\| q^\theta - q^y \right\|_1 \leq \sqrt{2KL(q^\theta, q^y)}.$$

Part 2. One can start by observing that

$$\theta^{y'} - \theta = \nabla_y h_\theta(y).$$

Further, since the function $y \mapsto h_\theta(y)$ is $2d^2$ -smooth, it is well known that one step of the gradient descent algorithm started at y (with step size $\frac{1}{4d^2}$) produces a point y' with value

$$h_\theta(y') \leq h_\theta(y) - \frac{\|\nabla_y h_\theta(y)\|_2^2}{4d^2},$$

and consequently $\|\nabla_y h_\theta(y)\|_2 \leq 2d\sqrt{\varepsilon}$. ■

We are now ready to prove the theorem.

Proof of Theorem 10: Using Theorem 7 we can find an $R = O(m^{3/2} \cdot M \cdot (L_p + \log |\mathcal{F}| + \log \frac{m}{\delta}))$ and $y_1, y_2 \in B(0, R)$ such that

$$h_{\theta_1}(y_1) - g(\theta_1) \leq \delta \quad \text{and} \quad h_{\theta_2}(y_2) - g(\theta_2) \leq \delta.$$

The value of $\delta > 0$ will be chosen later (as a function of $\varepsilon > 0$). We prove that $h_{\theta_1}(y_2) - g(\theta_1)$ is small and thus by Fact 12 we deduce that q^{y_2} and q^{θ_1} are close, thus also q^{θ_1} and q^{θ_2} are close. The details follow.

Since the function $y \mapsto h_{\theta_1}(y)$ is convex, the first order convexity condition yields

$$h_{\theta_1}(y_1) \geq h_{\theta_1}(y_2) + \langle \nabla_y h_{\theta_1}(y_2), y_1 - y_2 \rangle \tag{9}$$

and thus

$$\begin{aligned} h_{\theta_1}(y_2) &\leq h_{\theta_1}(y_1) + \langle \nabla_y h_{\theta_1}(y_2), y_2 - y_1 \rangle \\ &\leq h_{\theta_1}(y_1) + \|\nabla_y h_{\theta_1}(y_2)\|_2 \cdot \|y_2 - y_1\|_2 \\ &\leq g(\theta_1) + \delta + 2R \cdot \|\nabla_y h_{\theta_1}(y_2)\|_2 \end{aligned}$$

By a simple calculation:

$$\begin{aligned} \nabla_y h_{\theta_1}(y_2) &= \theta^{y_2} - \theta_1 \\ &= (\theta^{y_2} - \theta_2) + (\theta_2 - \theta_1) \end{aligned}$$

and further, by taking the norm and using Lemma 12

$$\|\nabla_y h_{\theta_1}(y_2)\|_2 \leq \|\theta^{y_2} - \theta_2\|_2 + \|\theta_2 - \theta_1\|_2 \leq 4d\sqrt{\delta} + \varepsilon.$$

Thus consequently we obtain:

$$h_{\theta_1}(y_2) \leq g(\theta_1) + 10Rd\sqrt{\varepsilon}.$$

Now, by Lemma 12 it follows that

$$\left\| q^{\theta_1} - q^{y_2} \right\|_1 \leq O(\sqrt{R}) \sqrt{4d\sqrt{\delta} + \varepsilon}$$

And further, again by Lemma 12 we obtain

$$\left\| q^{\theta_2} - q^{y_2} \right\|_1 \leq \sqrt{2\delta}$$

and hence by the triangle inequality

$$\left\| q^{\theta_1} - q^{\theta_2} \right\|_1 \leq O(\sqrt{R})\sqrt{4d\sqrt{\delta} + \varepsilon} + \sqrt{2\delta}$$

Now, by taking $\delta > 0$ such that $4d\sqrt{\delta} = \varepsilon$ we obtain that

$$\left\| q^{\theta_1} - q^{\theta_2} \right\|_1 \leq O(\sqrt{R\varepsilon}).$$

■

Appendix B. Proof of the bit complexity lower bound

We start by proving a technical fact which will be useful in establishing the large bit complexity example.

Fact B.1 *Let $v_1, v_2, \dots, v_N \subseteq \mathbb{R}^m$ be a set of vectors and denote $\delta := \text{dist}(0, \text{conv}(v_1, v_2, \dots, v_N))$. Assume that $\delta > 0$ and consider the optimal value of the following optimization problem:*

$$\tau = \min \{ \|y\| : y \in \mathbb{R}^m, \langle y, v_i \rangle \leq -1 \text{ for all } i = 1, 2, \dots, N \},$$

then $\tau \geq \frac{1}{\delta}$.

Proof We formulate the problem as a convex, quadratic program with linear constraints:

$$\begin{aligned} \min \quad & \|y\|^2, \\ \text{s.t.} \quad & \langle y, \alpha \rangle \leq -1 \text{ for all } \alpha \in \mathcal{F}'. \end{aligned} \tag{10}$$

To derive a lower bound on the optimal value of (10) we consider the dual program:

$$\begin{aligned} \max \quad & \sum_{i=1}^N \lambda_i - \frac{1}{4} \left\| \sum_{i=1}^N \lambda_i v_i \right\|^2, \\ \text{s.t.} \quad & \lambda_i \geq 0 \text{ for all } i = 1, 2, \dots, N. \end{aligned} \tag{11}$$

From weak duality we know that the optimal value of (11) is a lower bound to (10) thus we just need to provide a feasible solution to the dual program. To this end, let v be the shortest vector in the convex hull of v_1, v_2, \dots, v_N , i.e., $v \in \text{conv}(v_1, v_2, \dots, v_N)$ and $\|v\| = \delta$. v can be written as

$$v = \sum_{i=1}^N \mu_i v_i$$

for some $\mu \geq 0$ with $\sum_{i=1}^N \mu_i = 1$. Consider now $\lambda := \frac{2}{\delta^2} \mu \in \mathbb{R}_{\geq 0}^N$. The dual objective value for λ is

$$\sum_{i=1}^N \lambda_i - \frac{1}{4} \left\| \sum_{i=1}^N \lambda_i v_i \right\|^2 = \frac{2}{\delta^2} - \frac{1}{4} \cdot \frac{4}{\delta^4} \left\| \sum_{i=1}^N \mu_i v_i \right\|^2 = \frac{1}{\delta^2}.$$

This provides us with a lower bound of $\frac{1}{\delta^2}$ on the optimal value of (10) and thus a lower bound of $\frac{1}{\delta}$ on the optimization problem as in the statement of the Fact. \blacksquare

Remark 13 *It is not hard to prove that in Fact B.1 the value $\frac{1}{\tau}$ is not only a lower bound but is actually equal to the optimal value of the considered optimization problem. This can be established by plugging in an appropriate scaling of the shortest vector in $\text{conv}(v_1, v_2, \dots, v_N)$ for y .*

Proof of Theorem 3: Our construction of \mathcal{F} is based on existence of “flat” 0-1 polytopes, as established by Alon and Vu (1997). There exist $m + 1$ affinely independent points $\alpha_0, \alpha_1, \alpha_2, \dots, \alpha_m \in \{0, 1\}^m$ such that if we let $H = \alpha_1 + \text{span}\{\alpha_2 - \alpha_1, \dots, \alpha_m - \alpha_1\}$ (the $(m - 1)$ -dimensional affine subspace containing all points $\alpha_1, \dots, \alpha_m$) then

$$\text{dist}(\alpha_0, H) = e^{-\Omega(m \log m)}.$$

Without loss of generality we assume that $\alpha_0 = 0$. Let $y \in H$ be the projection of $\alpha_0 = 0$ onto H , i.e., a point such that

$$\langle y, \alpha_i - y \rangle = 0 \quad \text{for every } i = 1, 2, \dots, m.$$

Consider the lattice

$$L = \alpha_1 + \sum_{i=2}^m (\alpha_i - \alpha_1) \cdot \mathbb{Z},$$

with the origin at α_1 , with basis $\{(\alpha_i - \alpha_1)\}_{2 \leq i \leq m}$. Since the hyperplane H is covered by disjoint copies of the fundamental parallelepiped

$$F := \left\{ \sum_{i=2}^m \beta_i (\alpha_i - \alpha_1) : \beta_2, \dots, \beta_m \in [0, 1) \right\},$$

there is an integer translation of F which contains the point y . More formally, there exists an integer vector $\gamma \in \mathbb{Z}^m$ such that

$$y \in \gamma + F \subseteq H.$$

Note now that by denoting $F' = \gamma + F$ we obtain

$$\text{diam}(F') = \text{diam}(F) \leq m^{3/2},$$

since every vertex of F has integer coordinates in the range $[0, m]$. Let $\alpha \in \mathbb{Z}^m$ be now any of the 2^m vertices of F' . Since y belongs to F' we have

$$\|\alpha\| \leq \|y\| + \text{diam}(F') \leq O(1) + m^{3/2}.$$

Let \mathcal{F}' be a subset of \mathbb{Z}^m consisting of all the vertices of F' and let $\mathcal{F} := \mathcal{F}' \cup \{0\}$. Further, define $\theta := 0$. We prove that the conclusion of the Lemma holds under such a choice of \mathcal{F} and θ .

Towards this, we first note that affine hull of \mathcal{F}' is equal to H , as $F' \subseteq H$ and its vertices clearly span H . Moreover the point y , which is the projection of 0 onto H belongs to the convex hull of \mathcal{F}' . Let $\delta > 0$ be the distance between 0 and H and let $a \in \mathbb{R}^m$ (with $\|a\| = 1$) be the normal vector of H , in other words

$$H = \{x \in \mathbb{R}^m : a^\top x = \delta\}.$$

Note that the gradient of $h_\theta(y)$ with respect to y is given by:

$$\nabla_y h_\theta(y) = \frac{\sum_{\alpha \in \mathcal{F}} \alpha \cdot e^{\langle y, \alpha \rangle}}{\sum_{\alpha \in \mathcal{F}} e^{\langle y, \alpha \rangle}}.$$

Since h is L -smooth for some $L = \text{poly}(m)$ (see the proof of Corollary 9), we know that points with a large-magnitude gradient cannot be close to optimal. Quantitatively, we have

$$|g(\theta) - h_\theta(y)| \geq \frac{\|\nabla_y h_\theta(y)\|^2}{L}.$$

Thus to prove that $|g(\theta) - h_\theta(y)| \geq \varepsilon$ (for some $\varepsilon > 0$) it is enough to prove that $\|\nabla_y h_\theta(y)\| \geq \sqrt{\varepsilon L}$. We pick ε to be $\frac{\delta^2}{e^{4L}}$, then $\varepsilon = e^{-O(m \log m)}$. Moreover, the condition $|g(\theta) - h_\theta(y)| < \varepsilon$ implies that $\|\nabla_y h_\theta(y)\| < \frac{\delta}{e^2}$. We prove that the latter is possible only when $\|y\| \geq e^{\Omega(m \log m)}$.

Indeed, assume that $\|\nabla_y h_\theta(y)\| < \frac{\delta}{e^2}$. By the Cauchy-Schwarz inequality we have

$$\langle a, \nabla_y h_\theta(y) \rangle \leq \|\nabla_y h_\theta(y)\| \cdot \|a\| = \|\nabla_y h_\theta(y)\|$$

and moreover

$$\langle a, \nabla_y h_\theta(y) \rangle = \frac{0 + \sum_{\alpha \in \mathcal{F}'} \langle a, \alpha \rangle e^{\langle y, \alpha \rangle}}{\sum_{\alpha \in \mathcal{F}} e^{\langle y, \alpha \rangle}} = \delta \cdot \frac{\sum_{\alpha \in \mathcal{F}'} e^{\langle y, \alpha \rangle}}{1 + \sum_{\alpha \in \mathcal{F}'} e^{\langle y, \alpha \rangle}}.$$

It follows that

$$\frac{\sum_{\alpha \in \mathcal{F}'} e^{\langle y, \alpha \rangle}}{1 + \sum_{\alpha \in \mathcal{F}'} e^{\langle y, \alpha \rangle}} < \frac{1}{e^2},$$

and consequently $\sum_{\alpha \in \mathcal{F}'} e^{\langle y, \alpha \rangle} < \frac{1}{e}$, which implies in particular, that

$$\forall \alpha \in \mathcal{F}' \quad \langle y, \alpha \rangle \leq -1. \quad (12)$$

The question we would like to answer is: what is the shortest $y \in \mathbb{R}^m$ which satisfies condition (12)? This will give us a lower bound on $\|y\|$ satisfying $|g(\theta) - h_\theta(y)| < \varepsilon$. To answer this question, we apply Fact B.1 and conclude that every such y has length at least $\frac{1}{\delta}$. As $\delta = e^{-\Omega(m \log m)}$ we conclude that the optimal solution y^* to (10) satisfies $\|y^*\| = e^{\Omega(m \log m)}$ and the Lemma follows by contraposition. \blacksquare

Appendix C. Example when the empirical mean lies on the boundary

In this section, we show that for certain well-behaved classes of polytopes the empirical mean of a sequence of uniform samples will yield a θ that is very likely to lie on the boundary, and hence the result of Singh and Vishnoi (2014) is not applicable in such cases.

Theorem 14 (Existence of distributions whose empirical mean arbitrarily close to the boundary)

There exists a family of vectors $\mathcal{F} \subseteq \{0, 1\}^m$ such that $P := \text{conv}(\mathcal{F})$ contains a ball of radius $\Omega(1)$, has unary facet complexity 1, and if $X_1, X_2, X_3, \dots \in \mathcal{F}$ are independent uniform samples from \mathcal{F} , then for every $N \leq 2^m$ it holds

$$\Pr \left[\frac{1}{N} \sum_{i=1}^N X_i \notin \partial P \right] = O \left(\frac{N}{2^m} \right),$$

where ∂P denotes the boundary of the polytope P .

Proof Let \mathcal{F} be

$$\mathcal{F} := (\{0\} \times \{0, 1\}^{m-1}) \cup \{(1, 0, 0, \dots, 0)\}.$$

The first step is now to observe that the polytope $P := \text{conv}(\mathcal{F})$ can be simply described by:

$$P := \{x \in \mathbb{R}^m : 0 \leq x_1 \leq 1, 0 \leq x_i \leq 1 - x_1, \text{ for } i = 1, 2, \dots, m\}.$$

From this description, it already follows that P has unary facet complexity of 1, because all coefficients used to describe this polytope have magnitude 1. Moreover, another simple consequence of the above description of P is that

$$[0, 1/2]^m \subseteq P.$$

In particular, it follows that P contains a Euclidean ball of radius $\Omega(1)$ (for instance the ball centered at $(1/4, 1/4, \dots, 1/4)$ with radius $1/4$).

It remains to establish a bound on the probability that the empirical mean of $N \leq 2^m$ independent samples lies on the boundary of P . For this, observe first that

$$F = \{0\} \times [0, 1]^{m-1}$$

is an $(m - 1)$ -dimensional facet of P and in particular, is a subset of the boundary. Moreover

$$\Pr[X_i \in F] = \frac{1}{|\mathcal{F}|} = \frac{1}{2^{m-1} + 1}.$$

Consequently, we have

$$\Pr\left[\sum_{i=1}^N X_i \in \partial P\right] \geq \Pr\left[\frac{1}{m} \sum_{i=1}^N X_i \in F\right] = \prod_{i=1}^N \Pr[X_i \in F] = \left(1 - \frac{1}{2^{m-1} + 1}\right)^N.$$

Thus,

$$\Pr\left[\sum_{i=1}^N X_i \notin \partial P\right] \leq 1 - \left(1 - \frac{1}{2^{m-1} + 1}\right)^N = O\left(\frac{N}{2^m}\right),$$

the latter follows because $\frac{N}{2^m} = O(1)$. ■

Appendix D. Computability of max-entropy distributions

This section is devoted to proving computability of maximum entropy distributions, assuming that an appropriate counting oracle is provided. The main component in the argument is the bound on the magnitude of almost optimal dual solutions established in Theorem 7.

Theorem 15 *Let \mathcal{F} be any finite subset of \mathbb{Z}^m and let $d \in \mathbb{R}_{\geq 0}$ be its diameter, let M be the unary facet complexity of $\text{conv}(\mathcal{F})$. Then, there exists an algorithm such that given a probability distribution p on \mathcal{F} (via an evaluation oracle for g_p), $\theta \in P$ and an $\varepsilon > 0$, computes a vector $y \in \mathbb{R}^m$ with $\|y\| \leq \text{poly}(m, M, \log d, L_p, \log \frac{1}{\varepsilon})$ such that*

$$\|q^y - q^*\|_1 < \varepsilon,$$

where q^* is the optimal solution to (5) and q^y is a distribution over \mathcal{F} defined as $q_\alpha^y = \frac{p_\alpha e^{\langle \alpha, y \rangle}}{\sum_{\beta \in \mathcal{F}} p_\beta e^{\langle \beta, y \rangle}}$, for $\alpha \in \mathcal{F}$. The running time of the algorithm is polynomial in $m, M, d, L_p, \log \frac{1}{\varepsilon}$.

In the above g_p is a generalized counting function. An oracle for g_p is then simply defined as a procedure that given an $x > 0$ outputs

$$g_p(x) := \sum_{\alpha \in \mathcal{F}} p_\alpha x^\alpha, \quad \text{where } x^\alpha := \prod_{i=1}^m x_i^{\alpha_i}.$$

It was shown in [Singh and Vishnoi \(2014\)](#) that efficient (polynomial time) computability of max-entropy distribution essentially implies the existence of polynomial time oracles for g_p . Hence, in principle, the existence of such oracles is assumed without loss of generality.

Proof For convenience without loss of generality we might assume that $\mathcal{F} \subseteq \mathbb{N}^m$ and in fact even $\mathcal{F} \subseteq [0, 2d]^m \cap \mathbb{N}^m$ where d is the diameter of \mathcal{F} . This is because one can always shift the set \mathcal{F} (together with θ) and this operation does not affect the problem nor its parameters. To obtain a vector y , as required, we solve the dual program up to the desired precision (below we explain why is this enough). Recall that the dual program to (5) is given by

$$\begin{aligned} \inf \quad & h(\theta, y) = \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha - \theta, y \rangle} \right), \\ \text{s.t.} \quad & y \in \mathbb{R}^m. \end{aligned} \tag{13}$$

By a direct calculation one can show that for every feasible solution q of the primal problem (5)

$$KL(q, q^y) = h(\theta, y) - \sum_{\alpha \in \mathcal{F}} q_\alpha \log \frac{p_\alpha}{q_\alpha},$$

where $KL(\cdot, \cdot)$ denotes the KL -divergence. In particular for $q := q^*$

$$h(\theta, y) - g(\theta) = KL(q^*, q^y).$$

This means that in order to obtain a distribution q^y being ε -close in the KL -distance to the max-entropy distribution q^* it is enough to find an ε -optimal solution to the dual program. Moreover, from Pinsker's inequality we have

$$\|q - q^*\|_1 \leq \sqrt{2 \cdot KL(q^*, q)}$$

hence it suffices to find a solution y to the dual program which is $\delta := \Theta(\varepsilon^2)$ -optimal, to guarantee that $\|q^* - q^y\|_1 < \varepsilon$.

To find a δ -optimal solution to the dual problem we apply the ellipsoid method. First of all we note that $h(\theta, y)$ is a convex function of y (this follows from Hölder's inequality), which is the first requirement for the ellipsoid method to be applicable. It follows from [Theorem 15](#) that in order to find a δ -optimal solution to $\inf_{y \in \mathbb{R}^m} h(\theta, y)$ it is sufficient to solve

$$\min_{y \in B(0, R)} h(\theta, y),$$

where R is a certain bound, polynomial in $m, M, \log d, L_p$ and $\log \frac{1}{\delta}$. Now, following the treatment of the ellipsoid method in [Ben-Tal and Nemirovski \(2012\)](#) ([Theorem 8.2.1](#)) it remains to address the following issues.

1. Construct a first order oracle for $y \mapsto h(\theta, y)$, i.e., an efficient way to evaluate values $h(\theta, y)$ and gradients $\nabla_y h(\theta, y)$ of this function.
2. Provide an enclosing ball – containing the domain and a ball contained in the domain.
3. Provide a bound on the gap between the maximum and minimum value of $h(\theta, \cdot)$ in $B(0, R)$. More precisely, defining $D := \max_{y \in B(0, R)} h(\theta, y) - \min_{y \in B(0, R)} h(\theta, y)$ we would like $\log D$ to be polynomially bounded.
4. Provide a separation oracle for the domain $B(0, R)$.

Point (1) We first note that h can be equivalently written as

$$h(\theta, y) = \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha, y \rangle} \right) - \langle \theta, y \rangle.$$

Thus $h(\theta, y) = \log g_p(e^y) - \langle \theta, y \rangle$, where $e^y = (e^{y_1}, e^{y_2}, \dots, e^{y_m})$ and consequently $h(\theta, y)$ can be evaluated using just the evaluation oracle for g_p . For the case of gradients observe first that

$$\nabla_y h(\theta, y) = \frac{1}{g_p(e^y)} \sum_{\alpha \in \mathcal{F}} \alpha p_\alpha e^{\langle \alpha, y \rangle} - \theta.$$

Since computing $g_p(e^y)$ is easy, it remains to deal with $\sum_{\alpha \in \mathcal{F}} \alpha p_\alpha e^{\langle \alpha, y \rangle}$. For this note that the i th coordinate of the above is

$$\sum_{\alpha \in \mathcal{F}} \alpha_i p_\alpha e^{\langle \alpha, y \rangle} = \frac{d}{dt} g_p(e^{y_1}, \dots, e^{y_{i-1}}, e^{y_i} + t, e^{y_{i+1}}, \dots, e^{y_m}).$$

The right hand side above is a univariate polynomial $h(t)$ of degree at most $2d$ (since $\mathcal{F} \subseteq [0, 2d]$). To compute its derivative it is enough to learn all its coefficients (in fact it is enough to learn the coefficient of t in $h(t)$). Towards this, note that the evaluation oracle for g_p implies an evaluation oracle for $t \mapsto h(t)$, and hence we can simply evaluate h at $2d + 1$ different points and recover its coefficients using polynomial interpolation. The running time of such a procedure is polynomial in d – as required.

Note also that, importantly, to implement the first order oracle for $h(\theta, \cdot)$ the oracle g_p is queried only on inputs of polynomial bit complexity, as for every $y \in B(0, R)$ the vector e^y has polynomial bit complexity (since $R = \text{poly}(m, \log \frac{1}{\delta})$). Thus the running time of these procedures is polynomial in $m, \log \frac{1}{\delta}$ and d .

Point (2) This is not an issue since the domain is already a ball.

Point (3) Note that from the Cauchy-Schwarz inequality, for every $y \in B(0, R)$

$$\langle y, \alpha - \theta \rangle \leq \|y\| \|\alpha - \theta\| \leq R \cdot d,$$

hence, we have

$$\max_{y \in B(0, R)} h(\theta, y) \leq \log \left(e^{L_p} |\mathcal{F}| \cdot e^{R \cdot d} \right) \leq L_p \cdot R \cdot d \cdot \log |\mathcal{F}|.$$

Similarly, for the minimum

$$\min_{y \in B(0,R)} h(\theta, y) \geq -L_p \cdot R \cdot d \cdot \log |\mathcal{F}|.$$

Clearly the logarithm of the gap:

$$\log \left(\max_{y \in B(0,R)} h(\theta, y) - \min_{y \in B(0,R)} h(\theta, y) \right) \leq \log(2L_p \cdot R \cdot d \cdot \log |\mathcal{F}|)$$

is polynomially bounded (even without taking the logarithm it is still true), as $|\mathcal{F}| \leq (2d)^m$.

Point (4) This is clear - given a vector $y_0 \in \mathbb{R}^m$, if $\|y\| \leq R$ we just report y_0 to be in the domain and if $\|y_0\| = R' > R$ then $\{y \in \mathbb{R}^m : \langle y, y_0 \rangle = \frac{R+R'}{2}\}$ is the required separating hyperplane. ■

Appendix E. Applications to TCS

We present several new and old application of our main result: Theorem 7 in TCS. Since some of them rely on the fact that Newton polytopes of real stable polynomials have low unary facet complexity, below we briefly state the necessary background.

Preliminaries on real stability

We first define the concept of real stable polynomials that appear in some applications of our results. For a survey on real stable polynomials we refer the reader to [Wagner \(2011\)](#).

Definition 16 *A polynomial $p \in \mathbb{C}[x_1, \dots, x_m]$ is called real stable if all its coefficients are real numbers and the following condition holds*

$$\forall z \in \mathbb{C}^m \quad (\Im(z_i) > 0 \text{ for all } i = 1, 2, \dots, m) \Rightarrow p(z) \neq 0.$$

The following fact shows that polytopes arising from such polynomials have low facet complexity.

Fact E.1 (Bouchet and Cunningham (1995); Branden (2007)) *Let $p \in \mathbb{R}[x_1, \dots, x_m]$ be a real stable polynomial with non-negative coefficients. Then, there exists a “rank function” $r : \{-1, 0, 1\}^m \rightarrow \mathbb{Z}$ is, such that the convex hull of $\mathcal{F} \subseteq \mathbb{N}^m$ – the support of p can be described as:*

$$\text{conv}(\mathcal{F}) = \{x \in \mathbb{R}^m : \forall c \in \{-1, 0, 1\}^m \quad \langle x, c \rangle \leq r(c)\}.$$

Proof The proof is a simple consequence of two results. It was proved in [Branden \(2007\)](#) that the support of a real stable polynomial is a *jump system*. Such sets were studied previously in [Bouchet and Cunningham \(1995\)](#), where a polyhedral characterization, as in the conclusion, was shown. ■

Another important, yet classical fact about real stable polynomial is that they are log-concave in the positive orthant.

Fact E.2 (Güler (1997)) *Let $p \in \mathbb{R}[x_1, \dots, x_m]$ be a real stable polynomial with non-negative coefficients. Then the function $x \mapsto \log p(x)$ is concave over $x \in \mathbb{R}_{>0}^m$.*

E.1. Bounds for the matrix scaling problem

Consider the (r, c) -matrix scaling problem, where one is given a nonnegative square matrix $A \in \mathbb{R}^{n \times n}$ and two vectors $r, c \in \mathbb{N}^n$ (with $\|r\|_1 = \|c\|_1$) and the goal is to find a scaling: two positive vectors $x, y \in \mathbb{R}_{>0}^n$ such that for B defined as $B := XAY$ (with $X = \text{Diag}(x)$ and $Y = \text{Diag}(y)$) it holds that $B\mathbb{1} = r$ and $B^\top \mathbb{1} = c$ where $\mathbb{1} \in \mathbb{R}^n$ is the all-one vector. In other words, we want the row-sums of the matrix B to be equal to r and column sums to be equal to c .

For applications of matrix scaling, we refer to [Allen Zhu et al. \(2017\)](#); [Cohen et al. \(2017\)](#), where recently fast algorithms for matrix scaling were recently derived. Here let us only mention the problem of approximating the permanent of a nonnegative matrix. One can show that $\text{per}(XAY) = \text{per}(A) \cdot \prod_{i=1}^n x_i \cdot \prod_{i=1}^n y_i$ and thus, by scaling (with $r = c = \mathbb{1}$), one can reduce the problem of computing the permanent of a nonnegative matrix to the problem of computing the permanent of a doubly-stochastic matrix which is better understood. In particular several useful bounds, such as the Van der Waerden's bound (see [Gurvits \(2006\)](#)), and others ([Gurvits and Samorodnitsky \(2014\)](#)), are known for the permanent of doubly stochastic matrices.

One can prove that if such a scaling exists even *asymptotically* (i.e., a sequence of scalings exists such that they satisfy the scaling condition in the limit), then it can be recovered from the optimal solution to the following convex program

$$\inf_{z \in \mathbb{R}^n} \sum_{i=1}^n r_i \log \left(\sum_{j=1}^n A_{i,j} e^{z_j} \right) - \langle c, z \rangle; \quad (14)$$

see [Gurvits \(2006\)](#); [Kalantari et al. \(2008\)](#); [Allen Zhu et al. \(2017\)](#); [Cohen et al. \(2017\)](#). Indeed, the scaling is recovered as $x_i := r_i \cdot \left(\sum_{l=1}^n A_{i,l} e^{z_l^*} \right)^{-1}$ for $i = 1, 2, \dots, n$ and $y_j := e^{z_j^*}$ for $j = 1, 2, \dots, n$, where z^* stands for the optimal (or approximately optimal) solution to (14). The question which arises naturally is: does the optimal (or approximately optimal) scaling (x, y) have polynomial bit complexity? Or in other words: can we prove that the vector z^* has polynomially bounded entries: $\|z\|_\infty = \text{poly}(n, L_A)$? Here L_A denotes the bit complexity of the matrix A , i.e., $L_A := \max_{i,j \in [n]} \lceil \log A_{i,j} \rceil$.

We interpret the optimization problem (14) as an instance of entropy maximization which will allow us to apply Theorem 2 and to deduce appropriate bounds. To this end we rewrite the problem (14) as $\inf_{z \in \mathbb{R}^n} \log \prod_{i=1}^n \left(\sum_{j=1}^n A_{i,j} e^{z_j} \right)^{r_i} - \langle c, z \rangle$. Hence, there is a polynomial $p \in \mathbb{R}_{\geq 0}[x_1, x_2, \dots, x_n]$ such that the above optimization problem is equivalent to

$$\inf_{z \in \mathbb{R}^n} \log \frac{p(e^{z_1}, e^{z_2}, \dots, e^{z_n})}{e^{\langle z, c \rangle}} = \inf_{z \in \mathbb{R}^n} \log \left(\sum_{\alpha} p_{\alpha} e^{\langle \alpha - c, z \rangle} \right),$$

where the summation is over the support of p , which in this case (if, say, all the entries $A_{i,j}$ are positive) is equal to $\mathcal{F} = \left\{ \alpha \in \mathbb{N}^n : \sum_{j=1}^n \alpha_j = \|r\|_1 \right\}$. Hence (14) is essentially a generalized max-entropy program over the set \mathcal{F} with expectation $c \in \text{conv}(\mathcal{F})$. This allows us to analyze the scaling problem by applying Theorem 7 (a generalization of Theorem 2) and prove polynomial bounds on the bit complexity.

Below we use L_A for the bit complexity of the matrix A , i.e., $L_A := \max\{\lceil \log |A_{i,j}| \rceil : i, j \in \{1, 2, \dots, n\}, A_{i,j} \neq 0\}$.

Corollary 17 *Let $A \in \mathbb{R}^{n \times n}$ be a nonnegative matrix which is asymptotically (r, c) -scalable with $\|r\|_1 = \|c\|_1 = h$. Then for every $\varepsilon > 0$ there exists a scaling (x, y) such that:*

$$\begin{aligned} \left| \sum_{k=1}^n x_i y_k A_{i,k} - r_i \right| &< \varepsilon \quad \text{for all } i \in [n], \\ \left| \sum_{l=1}^n x_l y_j A_{l,j} - c_j \right| &< \varepsilon \quad \text{for all } j \in [n], \end{aligned}$$

and the bit complexities of all entries of x and y , i.e., $\max_{i \in [n]} |\log x_i|$ and $\max_{i \in [n]} |\log y_i|$ are bounded by $\text{poly}(n, L_A, h, \log \frac{1}{\varepsilon})$.

Proof Our strategy is to derive the result from Corollary 9. We first rewrite the objective (14) so that it matches the form as in Theorem 7.

$$\begin{aligned} \sum_{i=1}^n r_i \log \left(\sum_{j=1}^n A_{i,j} e^{z_j} \right) - \langle c, z \rangle &= \log \left(\prod_{i=1}^n \left(\sum_{j=1}^n A_{i,j} e^{z_j} \right)^{r_i} \right) - \langle c, z \rangle \\ &= \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha - c, z \rangle} \right) \\ &= h(c, z) \end{aligned}$$

For some set $\mathcal{F} \subseteq \mathbb{N}^m$ and some family of positive numbers $\{p_\alpha\}_{\alpha \in \mathcal{F}}$.

The next step is to obtain bounds on the various quantities m, d, M, L_p which appear in Corollary (9). Clearly we choose $m := n$. Next observe that $\|\alpha\|_\infty \leq h$ for $\alpha \in \mathcal{F}$, hence $d \leq h$ and that $|\log p_\alpha| \leq \text{poly}(h, L_A)$ for all $\alpha \in \mathcal{F}$, hence $L_p = \text{poly}(h, L_A)$.

Let us now prove that the polytope $\text{conv}(\mathcal{F})$ can be described by inequalities with small integer coefficients. To this end note that p can be naturally treated as a polynomial (p_α is the coefficient of $x^\alpha := \prod_{i \in [n]} x_i^{\alpha_i}$).

$$p(x) = \sum_{\alpha \in \mathcal{F}} p_\alpha x^\alpha = \prod_{i=1}^n \left(\sum_{j=1}^n A_{i,j} x_j \right)^{r_i},$$

and the support of p is equal to \mathcal{F} . Since p is a product of linear polynomials with nonnegative coefficients, it is a *real stable polynomial* (see preliminaries at the beginning of this section for some background). As shown in the preliminaries (Fact E.1), the Newton polytope of a real stable polynomial can be described by inequalities of the form

$$\langle a, x \rangle \leq b$$

where $a \in \{-1, 0, 1\}^n$. Hence we can take $M = O(1)$ in the statement of Corollary 9.

We now apply Corollary 9 to obtain a point z^* and use it to define a scaling by the following formulas

$$x_i = r_i \cdot \left(\sum_{l=1}^n A_{i,l} e^{z_l^*} \right)^{-1} \quad \text{and} \quad y_j = e^{z_j^*}, \quad \text{for } i, j = 1, 2, \dots, n. \quad (15)$$

Note that such a pair (x, y) has bit complexity which is polynomial in $n, L_A, h, \log \frac{1}{\varepsilon}$, hence it remains to reason about the precision of the resulting scaling.

By a direct calculation one obtains that

$$\left| \sum_{k=1}^n x_i y_k A_{i,k} - r_i \right| = 0 \quad \text{for all } i \in [n],$$

hence it remains to prove a bound on the precision of the scaling with respect to columns. For this, note that

$$(\nabla_z h(c, z))_j = \sum_{l=1}^n x_l y_j A_{l,j} - c_j.$$

Since Corollary 9 implies that

$$\|\nabla_z h(c, z^*)\|_2 < \varepsilon$$

the bound on the scaling precision follows. ■

E.2. Computability of recent continuous relaxations for counting and optimization Problems

The recent works [Straszak and Vishnoi \(2017\)](#) and [Anari and Oveis Gharan \(2017\)](#) study counting and optimization problems involving polynomials and provide approximation algorithms for them.

In the discussion below, for concreteness, we follow [Straszak and Vishnoi \(2017\)](#). – the setting of [Anari and Oveis Gharan \(2017\)](#) is similar.

[Straszak and Vishnoi \(2017\)](#) consider the following approach to construct approximation algorithms for counting and optimization problems: consider a multiaffine polynomial $p \in \mathbb{R}[x_1, x_2, \dots, x_m]$ with nonnegative coefficients and a family of sets $\mathcal{B} \subseteq 2^{[m]}$.¹⁰ Let us denote by p_α the coefficient of the monomial $x^\alpha := \prod_{i=1}^m x_i^{\alpha_i}$ in p . Then the problems considered are to compute

$$p_{\mathcal{B}} := \sum_{\alpha \in \mathcal{B}} p_\alpha \quad \text{and} \quad p_{\mathcal{B}}^{\max} := \max_{\alpha \in \mathcal{B}} p_\alpha. \quad (16)$$

One particular application where one is required to solve such problems is when dealing with constrained Determinantal Point Processes [Kulesza and Taskar \(2012\)](#); [Celis et al. \(2017, 2018\)](#). There, the polynomial $p(x)$ is of the form $\sum_{S \subseteq [m]} \det(L_{S,S}) x^S$, where $L \in \mathbb{R}^{m \times m}$ is a PSD matrix and $L_{S,S}$ denotes the submatrix of L corresponding to rows and columns in $S \subseteq [m]$. Solving such counting and optimization problems allows one to draw samples from constrained DPPs, i.e., distributions where the probability of a set S is proportional to $\det(L_{S,S})$ when $S \in \mathcal{B}$ and is 0 when $S \notin \mathcal{B}$. Such distributions have various interesting uses in data summarization and fair and diverse sampling (see [Celis et al. \(2017\)](#)).

To tackle the counting problem of computing $p_{\mathcal{B}}$ (as in (16)) the following relaxation is considered:

$$\text{Cap}_{\mathcal{B}}(p) := \sup_{\theta \in P(\mathcal{B})} \inf_{x > 0} \frac{p(x)}{\prod_{i=1}^m x_i^{\theta_i}}. \quad (17)$$

10. The polynomial in this problem is provided as an evaluation oracle. Similarly the family \mathcal{B} is given implicitly in the input, as a separation oracle for the convex hull of \mathcal{B} .

The relaxation for maximization is derived as a small modification of the above, hence we focus on $\text{Cap}_{\mathcal{B}}(p)$ in this discussion. $\text{Cap}_{\mathcal{B}}(p)$ approximates $p_{\mathcal{B}}$ provably well under the assumption that (roughly) the polynomial p is *real stable* (see [Wagner \(2011\)](#)) and that the family \mathcal{B} has a *matroid* structure. However, the question of efficient computability of these relaxations was not established in [Straszak and Vishnoi \(2017\)](#). The results of this paper allow us to deduce polynomial time algorithms for this relaxation in a fairly general setting. We interpret this relaxation as a variant of a max-entropy program which then allows us to apply [Theorem 15](#) to reason about its computability.

After taking the logarithm and replacing the variables $x_i > 0$ by e^{y_i} with $y_i \in \mathbb{R}$ for $i = 1, 2, \dots, m$ in $\text{Cap}_{\mathcal{B}}(p)$ we get

$$\log \text{Cap}_{\mathcal{B}}(p) = \sup_{\theta \in P(\mathcal{B})} \inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha} p_{\alpha} e^{\langle y, \alpha - \theta \rangle} \right),$$

where the summation in the inner optimization problem runs over $\alpha \in \text{supp}(p)$, i.e., the set of all monomials $\alpha \in \mathbb{N}^m$ whose coefficient p_{α} is non-zero. Note that the inner optimization problem is the dual of a generalized max-entropy program, as in [\(5\)](#). Hence, the inner optimization problem searches for a probability distribution over monomials of the polynomial p , whose expectation is θ and has the smallest possible *KL*-distance to the distribution given by the coefficients of p .

By taking into account the outer optimization over θ , one can see that $\log \text{Cap}_{\mathcal{B}}(p)$ minimizes the *KL*-distance of a distribution over monomials of p restricting its expectation to be in the convex hull of \mathcal{B} . This makes it suitable to apply [Theorem 15](#) and deduce polynomial time computability. Finally, we note that having a max-entropy solver for all θ is crucial here. If $\mathcal{F} \subseteq \mathbb{N}^m$ denotes the support of p , then for any $\theta \notin \text{conv}(\mathcal{F})$ the value of the entropy maximization problem is $-\infty$. Hence, in the outer optimization problem, the variable θ is constrained to be in $P(\mathcal{B}) \cap \text{conv}(\mathcal{F})$. Thus, the optimal solution θ^* of such an optimization problem might lie at the boundary of $\text{conv}(\mathcal{F})$ or very close to it, hence the result of [Singh and Vishnoi \(2014\)](#) does not yield polynomial bit complexity bounds in this setting. We prove the following

Theorem 18 *Let $p \in \mathbb{R}[x_1, x_2, \dots, x_m]$ be a polynomial with nonnegative coefficients with support $\mathcal{F} \subseteq \mathbb{N}^m$ and let $\mathcal{B} \subseteq \mathbb{N}^m$. Assume that the unary facet complexity of $\text{conv}(\mathcal{F})$ is M and denote $d = \max(\text{diam}(\mathcal{F}), \text{diam}(\mathcal{B}))$. There is an algorithm which given an evaluation oracle for p , a separation oracle for $\text{conv}(\mathcal{F})$, a separation oracle for $\text{conv}(\mathcal{B})$ and an $\varepsilon > 0$ computes a number X such that $1 - \varepsilon < \frac{X}{\text{Cap}_{\mathcal{B}}(p)} < 1 + \varepsilon$ in time $\text{poly}(m, L_p, \log d, M, \log \frac{1}{\varepsilon})$.*

Proof Denote by $P \subseteq \mathbb{R}^m$ the convex hull of \mathcal{F} – the support of p . As discussed above, $\text{Cap}_{\mathcal{B}}(p)$ can be rewritten equivalently as

$$\log \text{Cap}_{\mathcal{B}}(p) = \sup_{\theta \in P(\mathcal{B})} \inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha \in \mathcal{F}} p_{\alpha} e^{\langle \alpha - \theta, y \rangle} \right) = \sup_{\theta \in P(\mathcal{B})} \inf_{y \in \mathbb{R}^m} h(\theta, y),$$

where $\mathcal{F} \subseteq \mathbb{N}^m$ denotes the support of p . Thus we obtain a form of $\log \text{Cap}_{\mathcal{B}}(p)$ in terms of a function h as in [\(6\)](#). Let us denote

$$g(\theta) = \inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha \in \mathcal{F}} p_{\alpha} e^{\langle \alpha - \theta, y \rangle} \right).$$

Hence the goal is to solve $\sup_{\theta \in P(\mathcal{B})} g(\theta)$. In fact, since $g(\theta) = -\infty$ whenever $\theta \notin P$, we can rewrite it equivalently as

$$\sup_{\theta \in P(\mathcal{B}) \cap P} g(\theta).$$

Importantly, since we are only interested in an additive ε -approximation of the above quantity we can apply Theorem 7 to replace g by the following Lipschitz proxy:

$$\tilde{g}(\theta) := \inf_{y \in B(0, R)} h(\theta, y)$$

for an appropriate number R – polynomial in the input size and $\log \frac{1}{\varepsilon}$ (as in Theorem 7).

We now apply the ellipsoid method to find an additive ε -approximation to

$$\sup_{\theta \in P(\mathcal{B}) \cap P} \tilde{g}(\theta). \tag{18}$$

Firstly, observe that the function \tilde{g} is concave – as a pointwise infimum of affine functions. Now, following the treatment of the ellipsoid method in [Ben-Tal and Nemirovski \(2012\)](#) (Theorem 8.2.1) we have to address the following requirements to obtain polynomial running time

1. Construct a first order oracle for $\theta \mapsto \tilde{g}(\theta)$, i.e., an efficient way to evaluate values $\tilde{g}(\theta)$ and (sub)gradients of this function for $\theta \in P(\mathcal{B}) \cap P$.
2. Provide an outer ball – containing the domain of the considered optimization problem and an inner ball – contained in the domain. The radii of them should be of polynomial bit complexity.
3. A bound on the gap between the maximum and minimum value of $\tilde{g}(\theta)$ in $P(\mathcal{B}) \cap P$. More precisely, defining $D := \max_{\theta \in P(\mathcal{B}) \cap P} \tilde{g}(\theta) - \min_{\theta \in P(\mathcal{B}) \cap P} \tilde{g}(\theta)$ we would like $\log D$ to be polynomially bounded.
4. Provide a separation oracle for the domain $P(\mathcal{B}) \cap P$.

Point (1). To obtain an evaluation oracle for \tilde{g} note that Theorem 15 gives an algorithm to compute \tilde{g} up to δ additive error in time polynomial in $\log \frac{1}{\delta}$. This provides an approximate (weak) evaluation oracle, which is still enough to run the ellipsoid method (using shallow cuts, see [Grötschel et al. \(1988\)](#)).

Let us now discuss the oracle for the gradient of \tilde{g} . It is a standard fact in convex programming that for any $\theta \in P$ it holds that the point

$$y^* := \operatorname{argmin}_{y \in B(0, R)} h(\theta, y)$$

is a subgradient of \tilde{g} at θ . Using Theorem 15 we can efficiently compute a vector y which is a δ -approximation to y^* in the sense that $h(\theta, y) - h(\theta, y^*) \leq \delta$. We use this algorithm to implement an approximate subgradient oracle for \tilde{g} , i.e., for a given θ we output the vector y as above, using the algorithm in Theorem 15.

It remains to justify why is such an approximate gradient enough to run the ellipsoid method. Note that if $\theta \in P$ and $y \in B(0, R)$ is such that $h(\theta, y) \leq \tilde{g}(\theta) + \delta$ then for any $\theta' \in P$ we have

$$g(\theta') \leq h(\theta', y) = h(\theta, y) + \langle \theta' - \theta, y \rangle \leq g(\theta) + \delta + \langle \theta' - \theta, y \rangle,$$

where the equality above holds because h is an affine function of θ . Hence, by applying such a y as a gradient oracle, we obtain a separation hyperplane which never cuts out a point $\theta' \in P$ of value $g(\theta') > g(\theta) + \delta$. This, along with the fact that y can be found in time polynomial in $\log \frac{1}{\delta}$, allows us to apply the shallow cut ellipsoid method. We also refer the reader to the full version of [Singh and Vishnoi \(2014\)](#) and the proof of Theorem 2.11 therein where a detailed discussion of an ellipsoid algorithm for a related problem is provided.

Point (2). An outer ball is easy to obtain since $P(\mathcal{B}) \subseteq [0, 1]^m$. For the inner ball, suppose first that $P(\mathcal{B}) \cap P$ is full-dimensional. Then, since the vertices of both these polytopes have small integer entries, one can show using standard techniques that a ball of radius $\Omega(e^{-\text{poly}(m)})$ can be fit inside this polytope. In the non-full-dimensional case one has to work in a linear subspace $H \subseteq \mathbb{R}^m$ on which $P(\mathcal{B}) \cap P$ is full-dimensional; H can be found given separation oracles of $P(\mathcal{B})$ and P using standard subspace identification techniques (see e.g. [Grötschel et al. \(1988\)](#)).

Point (3). As in the proof of Theorem 15 we observe that for every $y \in B(0, R)$

$$-L_p \cdot R \cdot m \cdot d \log d \leq \log \left(\sum_{\alpha \in \mathcal{F}} p_\alpha e^{\langle \alpha - \theta, y \rangle} \right) \leq L_p \cdot R \cdot m \cdot d \log d,$$

and hence the gap (and hence clearly its logarithm as well) is polynomially bounded.

Point (4). This is clear as we are given separation oracles for both $P(\mathcal{B})$ and P . ■

E.3. Proof of computability of worst-case Brascamp-Lieb constants

Proof [of Theorem 5] Let us denote by $v_j \in \mathbb{R}^n$ the only row of the matrix B_j and by $V \in \mathbb{R}^{m \times n}$ the matrix collecting all the v_j 's as rows. Then the the Brascamp-Lieb constant (for the rank-1 case) can be computed as

$$\text{BL}(B, p) = \inf_{x > 0} \frac{\det \left(\sum_{j=1}^m p_j x_j v_j v_j^\top \right)}{\prod_{j=1}^m x_j^{p_j}}. \quad (19)$$

Note that the numerator is simply a polynomial $r(x) = \sum_{S \subseteq [m], |S|=n} p^S x^S \det(V_S V_S^\top)$, where V_S is the submatrix of V corresponding to the index set $S \subseteq [m]$ and $x^S := \prod_{i \in S} x_i$. Therefore, computability of $\text{BL}(B, p)$ directly follows from our general result (Theorem 15) on computability of maximum entropy distributions.

The problem of computing the (logarithm of the) worst-case constant can be reformulated as

$$\sup_{p \in P_B} \inf_{y \in \mathbb{R}^m} \log \left(\sum_{\alpha \in \mathcal{F}} r_\alpha(p) e^{\langle \alpha - p, y \rangle} \right),$$

where $\mathcal{F} \subseteq \{0, 1\}^m$ is the support of r and $r_\alpha(p)$ is the coefficient of α in r , as a function of $p \in P_B$. Thus, the inner optimization problem is the dual of a certain max-entropy program. Moreover, the

function

$$p \mapsto \log \left(\sum_{\alpha \in \mathcal{F}} r_\alpha(p) e^{\langle \alpha - p, y \rangle} \right) = \log \left(\sum_{S \subseteq [m], |S|=n} p^S x^S \det(V_S V_S^\top) \right)$$

is concave over $\mathbb{R}_{>0}^m$, as the polynomial $p \mapsto \sum_{S \subseteq [m], |S|=n} p^S x^S \det(V_S V_S^\top)$ (treating $x > 0$ as a vector of positive constants) is real stable, see Fact E.2. From this point on, the argument follows along the same lines as the proof of Theorem 18. ■