

12:2 2021

AStat

Algebraic Statistics

**TORIC INVARIANT THEORY
FOR MAXIMUM LIKELIHOOD ESTIMATION IN LOG-LINEAR MODELS**

CARLOS AMÉNDOLA, KATHLÉN KOHN, PHILIPP REICHENBACH AND ANNA SEIGAL



TORIC INVARIANT THEORY FOR MAXIMUM LIKELIHOOD ESTIMATION IN LOG-LINEAR MODELS

CARLOS AMÉNDOLA, KATHLÉN KOHN, PHILIPP REICHENBACH AND ANNA SEIGAL

We establish connections between invariant theory and maximum likelihood estimation for discrete statistical models. We show that norm minimization over a torus orbit is equivalent to maximum likelihood estimation in log-linear models. We use notions of stability under a torus action to characterize the existence of the maximum likelihood estimate, and discuss connections to scaling algorithms.

1. Introduction

Fruitful, sometimes unexpected, connections between algebra and statistics are constantly being discovered in the field of algebraic statistics. In this paper we unveil a connection between toric invariant theory and maximum likelihood estimation for log-linear models. Log-linear models are widespread in statistics and play a fundamental role in categorical data analysis, with a wide range of applications [5]. They consist of discrete probability distributions whose coordinatewise logarithm lies in a fixed linear space and include, for example, independence models and discrete graphical models [21]. There is a long history of the study of log-linear models in statistics, with an emphasis on understanding their maximum likelihood inference [13]. This concerns the existence of the maximum likelihood estimate (MLE), which maximizes the likelihood function given sample data, and statistical procedures for its computation.

Log-linear models play a prominent role in algebraic statistics [27]. The key link to algebra is that the Zariski closure of a log-linear model is a toric variety, defined by a monomial parametrization. Toric varieties have a foundational place among the algebraic varieties studied in algebraic geometry [7].

In our companion work [1], we establish a connection between finding the MLE and norm minimization along an orbit under a group action. We focus there on the setting of Gaussian group models, centered multivariate Gaussian models whose concentration matrices are of the form $g^T g$, where g lies in a group. In this paper, we study the connection between invariant theory and maximum likelihood estimation in the setting of discrete exponential families. We find remarkable similarities and differences between the discrete and Gaussian settings.

The paper is organized as follows. We introduce maximum likelihood estimation and our toric invariant theory setting in the expository Sections 2 and 3. Our main results are in Section 4. We give a characterization of MLE existence in terms of the existence of a vector of minimal norm in an orbit under a torus action (see Theorem 4.3), and an explicit way to compute the MLE from such a vector (see Theorem 4.7). We provide an alternative characterization in terms of null cones in Propositions 4.4

MSC2020: 14L24, 20G45, 62F10, 62H22, 62R01.

Keywords: maximum likelihood estimation, log-linear models, graphical models, torus actions, null cone, scaling algorithms.

and 4.5. We compare iterative proportional scaling (IPS), a classical method to find the MLE for log-linear models, with approaches to norm minimization in Section 5. We conclude the paper with a comparison with the multivariate Gaussian setting of [1] in Section 6 and outline a possible generalization for future research.

2. Maximum likelihood estimation

In this section we describe our statistical set-up: maximum likelihood estimation for discrete probability distributions. A distribution on m states is determined by its probability mass function p , where p_j is the probability that the j -th state occurs. Such a probability mass function is a point in the $(m-1)$ -dimensional probability simplex:

$$\Delta_{m-1} = \left\{ p \in \mathbb{R}^m \mid p_j \geq 0 \text{ for all } j \text{ and } \sum_{j=1}^m p_j = 1 \right\}.$$

A statistical model \mathcal{M} of distributions with m states is a subset of Δ_{m-1} .

The data for a discrete distribution is a *vector of counts* $u \in \mathbb{Z}_{\geq 0}^m$, where the coordinate u_j is the number of times that the j -th state occurs, and $n = u_+ := \sum_{j=1}^m u_j$ is the total number of observations. The corresponding empirical distribution is $\bar{u} = \frac{1}{n}u \in \Delta_{m-1}$.

Maximum likelihood estimation in \mathcal{M} given data u finds a point in the model most likely to give rise to the observed data. That is, an MLE given u is a maximizer \hat{p} of the likelihood function over the model \mathcal{M} . The likelihood function is

$$L_u(p) = p_1^{u_1} \cdots p_m^{u_m}. \tag{1}$$

For example, if the model fills Δ_{m-1} , the likelihood is maximized uniquely at $\hat{p} = \bar{u}$.

An MLE given u is, equivalently, a point in \mathcal{M} that minimizes the *Kullback-Leibler* (KL) divergence to the empirical distribution \bar{u} . The KL divergence from $q \in \mathbb{R}_{>0}^m$ to $p \in \mathbb{R}_{>0}^m$ is

$$\text{KL}(p \parallel q) = \sum_{j=1}^m p_j \log \frac{p_j}{q_j}.$$

Although the KL divergence is not a metric, for $p, q \in \Delta_{m-1}$ it satisfies $\text{KL}(p \parallel q) \geq 0$, and $\text{KL}(p \parallel q) = 0$ if and only if $p = q$. The logarithm of the likelihood given u can be written, up to additive constant, as

$$\ell_u(p) = -n \sum_{j=1}^m \bar{u}_j \log \frac{\bar{u}_j}{p_j} = -n \text{KL}(\bar{u} \parallel p).$$

We see that maximizing the log-likelihood is equivalent to minimizing the KL divergence to the empirical distribution.

3. Toric invariant theory

In this section we describe the invariant theory of a torus action that we will use. We begin by introducing notions of stability under a group action.

fibers are copies of \mathbb{C} . Given a group action on X , a linearization is an action on L that agrees with the original action under projection under p , and that is a linear action on each fiber [10, Chapter 7]. For example, the following projection map is a line bundle,

$$p : \{(x, v) \in \mathbb{P}_{\mathbb{C}}^{m-1} \times \mathbb{C}^m \mid v \in \ell_x\} \rightarrow \mathbb{P}_{\mathbb{C}}^{m-1},$$

where the fiber over each $x \in \mathbb{P}_{\mathbb{C}}^{m-1}$ corresponds to the line ℓ_x in \mathbb{C}^m that the point represents. That way, a linearization lifts an action on $\mathbb{P}_{\mathbb{C}}^{m-1}$ to an action on \mathbb{C}^m .

We now consider the notions of stability in Definition 3.1 for this torus action. They specialize to give polyhedral conditions. The convex hull of the columns $a_j \in \mathbb{Z}^d$ of the matrix A is the polytope

$$P(A) := \text{conv}\{a_1, \dots, a_m\} \subseteq \mathbb{R}^d.$$

The set $P(A)$ consists of vectors in \mathbb{R}^d of the form Au for some $u \in \Delta_{m-1}$. We define subpolytopes that depend on an indexing set $J \subseteq [m]$

$$P_J(A) := \text{conv}\{a_j \mid j \in J\}.$$

For $v \in \mathbb{C}^m$, let $\text{supp}(v) := \{j \mid v_j \neq 0\} \subseteq [m]$. We abbreviate $P_{\text{supp}(v)}(A)$ to $P_v(A)$, i.e., we define

$$P_v(A) := \text{conv}\{a_j \mid j \in \text{supp}(v)\}.$$

For a polytope $P \subseteq \mathbb{R}^d$, we denote its interior by $\text{int}(P)$ and its relative interior by $\text{relint}(P)$.

Theorem 3.4 (Hilbert–Mumford criterion for a torus). *Let $v \in \mathbb{C}^m$ and consider the action of the complex torus GT_d on \mathbb{C}^m given by matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. We have*

- (a) v unstable $\Leftrightarrow b \notin P_v(A)$,
- (b) v semistable $\Leftrightarrow b \in P_v(A)$,
- (c) v polystable $\Leftrightarrow b \in \text{relint}(P_v(A))$,
- (d) v stable $\Leftrightarrow b \in \text{int}(P_v(A))$.

We give an elementary proof of Theorem 3.4 in Appendix A. Alternative proofs can be found in [10, Theorem 9.2] or [29, Theorem 1.5.1].

3C. The moment map. We introduce the moment map and state the Kempf–Ness theorem, for a torus action. As before, GT_d denotes the d -dimensional complex torus, and we consider its action on \mathbb{C}^m via the matrix $A \in \mathbb{Z}^{d \times m}$. We first consider the trivial linearization $b = 0$ and later a general linearization $b \in \mathbb{Z}^d$.

Fix $v \in \mathbb{C}^m$ and consider a torus element $\lambda = (\lambda_1, \dots, \lambda_d)$ in GT_d . The j -th coordinate of the vector $\lambda \cdot v \in \mathbb{C}^m$ is

$$(\lambda \cdot v)_j = \lambda_1^{a_{1j}} \cdots \lambda_d^{a_{dj}} v_j.$$

Next, we consider the map that sends λ to the squared norm of $\lambda \cdot v$:

$$\gamma_v : \text{GT}_d \rightarrow \mathbb{R}, \quad \lambda \mapsto \|\lambda \cdot v\|^2 = \sum_{j=1}^m |\lambda_1|^{2a_{1j}} \cdots |\lambda_d|^{2a_{dj}} |v_j|^2.$$

The infimum of γ_v over $\lambda \in \text{GT}_d$ is the capacity of v .

More generally, for an algebraic group G acting linearly on a space V we can consider the map $\gamma_v : G \rightarrow \mathbb{R}$ that sends $g \mapsto \|g \cdot v\|^2$, for fixed $v \in V$. The derivative is a map $D_I \gamma_v : T_I G \rightarrow \mathbb{R}$, where $T_I G$ is the tangent space to G at I . The *moment map* μ assigns to $v \in V$ the derivative of the map γ_v at $I \in G$.

For the group GT_d , the tangent space at I is equal to \mathbb{C}^d , and the derivative is a map $\mathbb{C}^d \rightarrow \mathbb{R}$. Recall that the map $f : \mathbb{C} \rightarrow \mathbb{C}$, $z \mapsto |z|^2$ is not complex differentiable. We identify \mathbb{C} with \mathbb{R}^2 , writing $z = x + iy$. Then the differential of f is given in terms of x and y , and their tangent directions \dot{x} and \dot{y} , as $\dot{z} \mapsto 2x\dot{x} + 2y\dot{y}$. In particular, the differential at $z = 1$ is the map $\dot{z} \mapsto 2\Re(\dot{z})$, where $\Re(\cdot)$ denotes the real part of a complex scalar. Extending to the multivariate setting, we obtain the derivative map

$$D_I \gamma_v : \mathbb{C}^d \rightarrow \mathbb{R}, \quad \dot{\lambda} \mapsto \sum_{i=1}^d \left(\sum_{j=1}^m 2a_{ij} |v_j|^2 \right) \Re(\dot{\lambda}_i) = 2\Re \left(\sum_{i=1}^d (Av^{(2)})_i \dot{\lambda}_i \right),$$

where $v^{(2)}$ is the vector with j -th coordinate $|v_j|^2$. We can identify \mathbb{C}^d with the space of \mathbb{R} -linear functionals $\text{Hom}(\mathbb{C}^d, \mathbb{R})$, by associating to $u \in \mathbb{C}^d$ the map $w \mapsto \Re(\sum_{i=1}^d u_i w_i)$. Under this identification, the linear map $D_I \gamma_v$ corresponds to the vector $2Av^{(2)} \in \mathbb{C}^d$. Hence the moment map, for linearization $b = 0$, is

$$\mu : \mathbb{C}^m \rightarrow \mathbb{C}^d, \quad v \mapsto 2Av^{(2)}.$$

For a general linearization $b \in \mathbb{Z}^d$, we replace the columns a_j of A by $a_j - b$. This replaces the vector $Av^{(2)}$ by $Av^{(2)} - b\|v\|^2$, where $\|v\|^2 = \sum_{j=1}^m |v_j|^2$. We obtain

$$\mu : \mathbb{C}^m \rightarrow \mathbb{C}^d, \quad v \mapsto 2(Av^{(2)} - \|v\|^2 b).$$

The Kempf–Ness theorem relates points of minimal norm in an orbit, or orbit closure, to the vanishing of the moment map. It was first proven in [19]. Nowadays, several statements are referred to as Kempf–Ness theorem; see [1, Section 2] for a summary. For our torus action, we obtain the following.

Theorem 3.5 (Kempf–Ness theorem for a torus). *Consider the torus action of GT_d given by matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. A vector is semistable (resp. polystable) if and only if there is a nonzero v in its orbit closure (resp. orbit) with $Av^{(2)} = \|v\|^2 b$. This v is a vector of minimal norm in the orbit closure (resp. orbit).*

We give two proofs of Theorem 3.5 in Appendix B. The first proof is a translation from the original paper of Kempf and Ness [19]; the second proof uses Theorem 3.4.

3D. The null cone. The set of unstable points under a group action on a vector space is the null cone, see Definition 3.1. In many settings of interest, the null cone is a Zariski closed set, the vanishing locus of all homogeneous invariants of positive degree. It is a classical object of interest, studied by Hilbert [18].

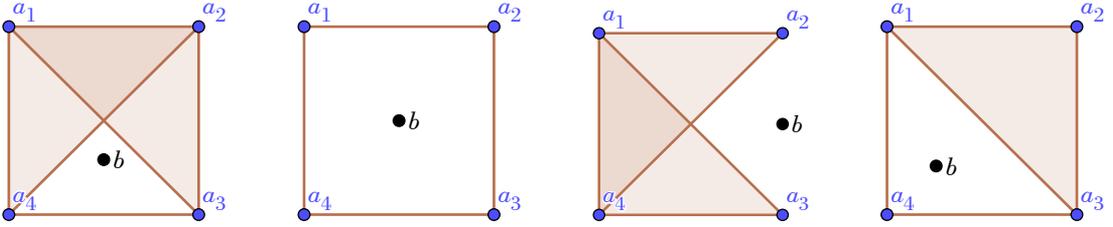


Figure 1. The maximal subpolytopes of $P(A)$ not containing b , for four different $b \in \mathbb{Z}^2$. For example, the leftmost picture displays $P_J(A)$ for J equal to $\{1, 2, 3\}$, $\{1, 2, 4\}$, and $\{3, 4\}$. Each subpolytope corresponds to an irreducible component of the null cone, see Proposition 3.6. For b on the boundary of $P(A)$, all maximal subpolytopes intersect, see Proposition 4.5.

We recall the setting of the action of the complex torus GT_d on \mathbb{C}^m given by a matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. The stability of $v \in \mathbb{C}^m$ is determined by its support $\text{supp}(v)$; see Theorem 3.4. In particular the null cone, as a set, is a union of coordinate linear spaces. We describe it in terms of the standard basis vectors in \mathbb{C}^m , denoted e_1, \dots, e_m . The linear space spanned by $\{e_j \mid j \in J\}$ is denoted $\langle e_j \mid j \in J \rangle$.

A vector b in $P(A)$ can be written as a convex combination of the m columns of A . We consider the maximal subpolytope of $P(A)$ that does not contain b , as well as the minimal subpolytope of $P(A)$ that contains b . Both minimality and maximality are taken with respect to inclusion in the set $[m]$.

In Proposition 3.6, we see the connection between irreducible components of the null cone and maximal subpolytopes of $P(A)$ not containing b , see Figure 1. Then, in Proposition 3.7, we see that minimal subpolytopes containing b give set-theoretic defining equations for the null cone, see Figure 2.

Proposition 3.6. Consider the action of GT_d on \mathbb{C}^m given by matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. The irreducible components of the null cone are the linear spaces $\langle e_j \mid j \in J \rangle$, where $P_J(A)$ is a maximal subpolytope of $P(A)$ with $b \notin P_J(A)$.

Proof. Assume that a point $v \in \mathbb{C}^m$ lies in a linear space $\langle e_j \mid j \in J \rangle$ where $b \notin P_J(A)$. Then $\text{supp}(v) \subseteq J$, hence $b \notin P_v(A)$, and v is unstable by Theorem 3.4(a). Conversely, assume that $v \in \mathbb{C}^m$ is not contained in any linear space $\langle e_j \mid j \in J \rangle$ as in the statement. Since the $P_J(A)$ are maximal with $b \notin P_J(A)$, we have $b \in P_v(A)$ and v is semistable. \square

Proposition 3.7. Consider the action of GT_d on \mathbb{C}^m given by matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. A vector $v \in \mathbb{C}^m$ is in the null cone if and only if all products $\prod_{j \in J} v_j$ vanish, where $J \subseteq [m]$ indexes a minimal subpolytope of $P(A)$ containing b .

Proof. Denote $v_J := \prod_{j \in J} v_j$. If some v_J is nonzero, i.e., $J \subseteq \text{supp}(v)$, then $b \in P_J(A)$ implies $b \in P_v(A)$, hence v is semistable by Theorem 3.4(b). Conversely, if v is semistable then $b \in P_v(A)$. By minimality of the minimal subpolytopes $P_J(A)$ containing b we have, for some J in the statement, the containment $P_J(A) \subseteq P_v(A)$, i.e., $J \subseteq \text{supp}(v)$, hence $v_J \neq 0$. \square

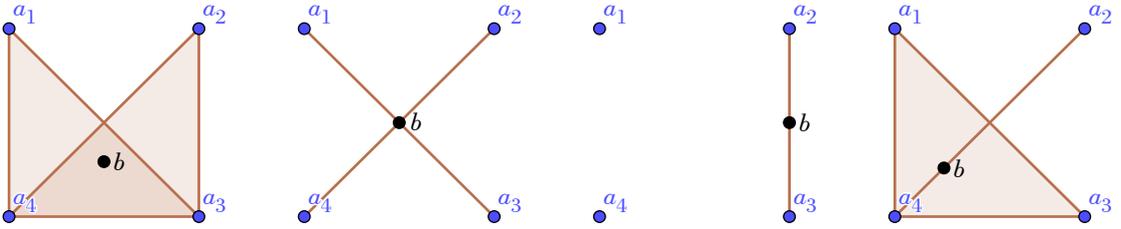


Figure 2. The minimal subpolytopes of $P(A)$ containing b , for four choices of $b \in \mathbb{Z}^2$. For example, the leftmost picture displays $P_J(A)$ for J equal to $\{1, 3, 4\}$ and $\{2, 3, 4\}$. Each subpolytope corresponds to a generator of the null cone, see Proposition 3.7.

The null cone is defined by the vanishing of all homogeneous invariants of positive degree. The monomials from Proposition 3.7 give the square-free part of the generators of the null cone. We describe how to take powers of the indeterminates appearing in the monomials, in order to turn them into invariants. Let $J \subseteq [m]$ index a minimal subpolytope of $P(A)$ containing b . Then 0 can be written as a strictly positive convex combination of $\{(a_j - b) \mid j \in J\}$. Since the entries of the matrix A and the vector b are integers, the convex combination is rational. Multiplying by the lowest common denominator gives a positive integer linear combination

$$\sum_{j \in J} r_j (a_j - b) = 0, \quad r_j \in \mathbb{Z}_{>0}. \tag{4}$$

The monomials $\prod_{j \in J} x_j^{r_j}$ are invariants under the group action, since

$$\lambda \cdot \left(\prod_{j \in J} x_j^{r_j} \right) = \prod_{j \in J} (\lambda^{-(a_j - b)} x_j)^{r_j} = \prod_{j \in J} x_j^{r_j} \cdot \lambda^{-\sum_{j \in J} r_j (a_j - b)} = \prod_{j \in J} x_j^{r_j},$$

where the first equality follows from Remark 3.2 and the last equality follows from (4).

4. Main Results

We begin this section by introducing maximum likelihood estimation in log-linear models, pointing out connections to torus actions. We relate stability under a torus action to maximum likelihood estimation for log-linear models in Section 4B, and describe how to compute the MLE from the moment map in Section 4C.

4A. Log-linear models. A log-linear model consists of distributions whose logarithms lie in a fixed linear space. The log-linear model corresponding to a matrix $A \in \mathbb{Z}^{d \times m}$ is

$$\mathcal{M}_A = \{p \in \Delta_{m-1} \mid \log p \in \text{rowspan}(A)\}. \tag{5}$$

The coordinatewise logarithm $\log p$ applies to p with strictly positive entries, and we therefore have $\mathcal{M}_A \subseteq \text{relint}(\Delta_{m-1})$. A parametrization of the model \mathcal{M}_A is given by

$$\phi^A : \mathbb{R}_{>0}^d \rightarrow \Delta_{m-1}, \quad \theta \mapsto \left(\frac{1}{Z(\theta)} \prod_{i=1}^d \theta_i^{a_{ij}} \right)_{1 \leq j \leq m}, \tag{6}$$

see, e.g., [11, Proposition 2.1.5] or [27, Corollary 7.3.9].

Since the model \mathcal{M}_A is not closed, the MLE may not exist. Birch [3] was the first to rigorously study MLE existence in the context of multiway tables, where he observed that u having all entries strictly positive is a sufficient condition for the MLE to exist and derived condition (8), sometimes known as Birch’s Theorem, see [23, Theorem 1.10]. The fact that some entries could still be 0 without affecting MLE existence was not fully understood until the work of Haberman, who gave the first characterization of MLE existence in her paper [17]. A modern necessary and sufficient condition is the following.

Proposition 4.2 [27, Theorem 8.2.1]. *Let $A \in \mathbb{Z}^{d \times m}$ be such that $\mathbb{1} \in \text{rowspan}(A)$ and let \mathcal{M}_A be the corresponding log-linear model. Suppose we observe a vector of counts $u \in \mathbb{Z}_{\geq 0}^m$. Then the MLE given u exists in \mathcal{M}_A if and only if $A\bar{u}$ lies in the relative interior of the polytope $P(A)$.*

In particular, we see that, indeed, if u has all entries positive, the MLE always exists. However, if u has some entries zero, the MLE may or may not exist.

Following [21, Section 4.2.3], we define the *extended log-linear model* $\overline{\mathcal{M}}_A$ to be the closure of \mathcal{M}_A in the Euclidean topology on \mathbb{R}^m . The extended model allows for distributions that have some zero coordinates. The MLE always exists for the extended model, because it is compact and the likelihood function is continuous. If the MLE given u does not exist in \mathcal{M}_A , we refer to the MLE given u in the extended model $\overline{\mathcal{M}}_A$ as the *extended MLE* given u .

Since the likelihood function (1) is strictly concave for log-linear models, the MLE is unique if it exists, and similarly for the extended MLE.

4B. Relating stability to the MLE. We now describe connections between existence of the MLE for a log-linear model and stability under a torus action.

Theorem 4.3. *Consider a vector of counts $u \in \mathbb{Z}_{\geq 0}^m$ with sample size $u_+ = n$, matrix $A \in \mathbb{Z}^{d \times m}$ with $\mathbb{1} \in \mathbb{C}^m$ in the rowspan, and vector $b = Au \in \mathbb{Z}^d$. The stability under the action of the complex torus GT_d given by matrix nA with linearization b is related to ML estimation in \mathcal{M}_A as follows:*

- (a) $\mathbb{1}$ unstable does not happen,
- (b) $\mathbb{1}$ semistable \Leftrightarrow extended MLE exists and is unique,
- (c) $\mathbb{1}$ polystable \Leftrightarrow MLE exists and is unique,
- (d) $\mathbb{1}$ stable does not happen.

Proof. We refer to the conditions for the different notions of stability, coming from the Hilbert–Mumford criterion in Theorem 3.4. By Proposition 4.2, the MLE of u exists if and only if b lies in the relative interior of the polytope $P(nA)$, which is the condition for polystability in Theorem 3.4.

It remains to see that the cases of unstable and stable do not occur. The all-ones vector $\mathbb{1}$ can never be unstable with respect to the action in Theorem 4.3, because $b = Au$ is in the polytope $P(nA)$. Finally, the stable case also cannot arise, due to the assumption that the vector $\mathbb{1}$ lies in the row span of A , as follows. Writing $\mathbb{1}$ as a linear combination of the rows, i.e., $r^T A = \mathbb{1}$, we have that all vectors a_j lie on the hyperplane $r_1 x_1 + \dots + r_d x_d = 1$ and the polytope $P(A)$ has empty interior in \mathbb{R}^d . \square

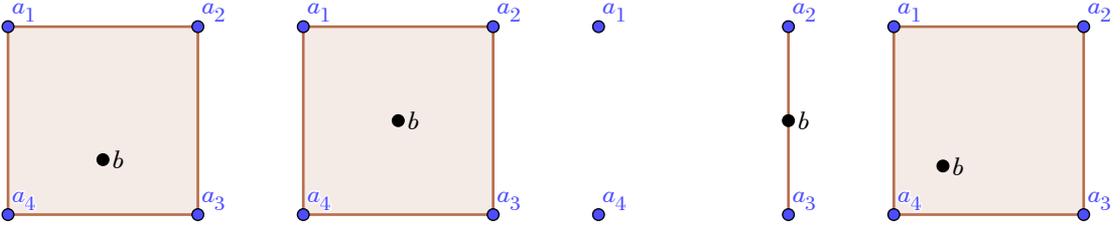


Figure 3. The face $F_b(A)$ of $P(A)$, for four choices of $b \in \mathbb{Z}^2$. For example, the leftmost picture displays the face $P_J(A)$ for $J = \{1, 2, 3, 4\}$. The vectors a_i outside of the face are in the intersection of all the irreducible components of the null cone, see Proposition 4.5. For the corresponding components of the null cone, see Figure 1.

We remark that we could take any other vector of full support in Theorem 4.3. The theorem shows that MLE existence can be tested by checking polystability under the group action. We now give alternative characterizations that involve semistability, which has advantages over polystability. The semistability of v can be checked by evaluating generators of the null cone at v . If all generators vanish then v is unstable, otherwise it is semistable.

Proposition 4.4. *For a vector of counts $u \in \mathbb{Z}_{\geq 0}^m$ with $u_+ = n$ and $A \in \mathbb{Z}^{d \times m}$, the MLE given u exists if and only if there is some $b \in \mathbb{Z}^d$, of the form $b = Av$ for $v \in \mathbb{R}_{>0}^m$, such that u is semistable for the torus action given by matrix nA with linearization b .*

Proof. We first assume that the MLE given u exists. Since the vector Au lies in the polytope $P_u(nA)$, the vector $u \in \mathbb{Z}_{\geq 0}^m$ is semistable for the action given by matrix nA with linearization Au , by Theorem 3.4(b). Moreover, since Au is in the relative interior of $P(nA)$, by Proposition 4.2, the vector Au is of the form Av for some $v \in \mathbb{R}_{>0}^m$.

Conversely, if the MLE given u does not exist, then Au lies on the boundary of the polytope $P(nA)$. Hence the whole polytope $P_u(nA)$ is contained in the boundary. Thus, for every $b \in \mathbb{Z}^d$ of the form $b = Av$ for $v \in \mathbb{R}_{>0}^m$, we have $b \notin P_u(nA)$. Then u is unstable under the torus action given by matrix nA with linearization b , by Theorem 3.4(a). □

To test MLE existence with Proposition 4.4, we need to test null cone membership for multiple linearizations. We now discuss a different approach, involving one null cone. For a vector $b \in P(A)$ we denote by $F_b(A)$ the minimal face of the polytope $P(A)$ that contains b ; see Figure 3.

Proposition 4.5. *Consider a vector of counts $u \in \mathbb{Z}_{\geq 0}^m$ with $u_+ = n$ and $A \in \mathbb{Z}^{d \times m}$. The intersection of the irreducible components of the null cone for the torus action given by matrix nA with linearization $b = Au$ is $\langle e_j \mid na_j \notin F_b(nA) \rangle$.*

In particular, the MLE given u exists in \mathcal{M}_A if and only if the intersection of the irreducible components of the null cone is $\{0\}$.

Proof. Define $A' := nA$ and consider the polytope $P(A')$, the convex hull of $a'_j := na_j$. We consider the null cone under the torus action given by matrix A' with linearization b . A linear space $\langle e_j \mid j \in J \rangle$ is in the null cone if and only if $b \notin P_J(A')$, by Proposition 3.6.

We will show that e_j is contained in every irreducible component of the null cone if and only if $a'_j \notin F_b(A')$. From this, the second paragraph of the statement follows because the MLE given u exists if and only if $b = Au$ is in the relative interior of the polytope $P(A')$, i.e., b does not lie on a proper face, and $F_b(A') = P(A')$.

Consider an index j with $a'_j \notin F_b(A')$. All possible expressions for b as $b = Av$ for some $v \geq 0$ have $v_j = 0$, since $F_b(A')$ is a face of $P(A')$. Let $J \subseteq [m]$ be such that $b \notin P_J(A')$, i.e., $\langle e_j \mid j \in J \rangle$ is in the null cone. Taking $J' = J \cup \{j\}$, the polytope $P_{J'}(A')$ still does not contain b . Hence, e_j lies in an irreducible component of the null cone that contains $\langle e_{j'} \mid j' \in J' \rangle$; so e_j lies in every irreducible component.

Conversely, consider an index j with $a'_j \in F_b(A')$. We show that there exists an irreducible component of the null cone that does not contain e_j . For each facet $F \subseteq F_b(A')$, let v_F be a vector with $\text{supp}(v_F) = \{k \mid a'_k \in F\}$, and take w_F with $\text{supp}(w_F) = \text{supp}(v_F) \cup \{j\}$. The union of $P_{w_F}(A')$ over facets $F \subseteq F_b(A')$ is the whole polytope $F_b(A')$, so $b \in P_{w_F}(A')$ for some facet F . By the minimality of $F_b(A')$, we have $b \notin P_{v_F}(A')$. Hence $\langle e_k \mid a'_k \in F \rangle$ is contained in an irreducible component of the null cone but, since $b \in P_{w_F}(A')$, the irreducible component does not contain e_j . \square

Example 4.6. We illustrate Proposition 4.5 for the log-linear model \mathcal{M}_A , where

$$A = \begin{matrix} & p_{000} & p_{001} & p_{010} & p_{011} & p_{100} & p_{101} & p_{110} & p_{111} \\ \begin{matrix} p_{00+} \\ p_{01+} \\ p_{10+} \\ p_{11+} \\ p_{+00} \\ p_{+01} \\ p_{+10} \\ p_{+11} \end{matrix} & \begin{pmatrix} 1 & 1 & & & & & & & \\ & & 1 & 1 & & & & & \\ & & & & 1 & 1 & & & \\ & & & & & & 1 & 1 & \\ 1 & & & & 1 & & & & \\ & 1 & & & & 1 & & & \\ & & 1 & & & & 1 & & \\ & & & 1 & & & & 1 & \\ & & & & 1 & & & & 1 \end{pmatrix} \end{matrix}$$

This is the graphical model on three binary random variables X_i given by the path graph 1—2—3, defined by the conditional independence relation $X_1 \perp\!\!\!\perp X_3 \mid X_2$. To identify the graphical model with \mathcal{M}_A , we identify \mathbb{R}^8 with $\mathbb{R}^{2 \times 2 \times 2}$ and label the columns of A by entries p_{ijk} . The sufficient statistics of the model are the eight marginals $p_{ij+} := p_{ij0} + p_{ij1}$ and $p_{+ij} := p_{0ij} + p_{1ij}$, where $(i, j) \in \{0, 1\}^2$.

We compute the irreducible components of the null cone for the torus action given by matrix nA with linearization Au , for various $u \in \mathbb{Z}^8$. The null cone is the zero locus of those monomials in the ring $\mathbb{C}[x_1, \dots, x_8]$ such that the supports of their exponent vectors index minimal subpolytopes of $P(nA)$ that contain b , as in Proposition 3.7.

Let $u = [1 \ 0 \ 1 \ 0 \ 0 \ 1 \ 0 \ 1]^T$. Then $b = \mathbb{1} \in \mathbb{R}^8$ and the null cone is the vanishing locus of $x_1x_3x_6x_8$, $x_1x_4x_6x_7$, $x_2x_3x_5x_8$, and $x_2x_4x_5x_7$. The irreducible components only intersect at $\{0\}$, hence the MLE given u exists in \mathcal{M}_A .

Let $u = [1 \ 1 \ 1 \ 1 \ 0 \ 1 \ 1 \ 0]^T$. Then $b = [2 \ 2 \ 1 \ 1 \ 1 \ 2 \ 2 \ 1]$ and the null cone is the vanishing locus of $x_1x_2x_3x_4x_6x_7$, $x_2^2x_3x_4x_5x_7$, $x_1x_2x_3^2x_6x_8$, and $x_2^2x_3^2x_5x_8$. The irreducible components only intersect at $\{0\}$, so the MLE given u exists in \mathcal{M}_A .

When $u = [1 \ 0 \ 1 \ 0 \ 0 \ 1 \ 0 \ 0]^\top$, the null cone is the vanishing locus of $x_1x_3x_6$ and $x_2x_3x_5$. The irreducible components intersect at $\langle e_4, e_7, e_8 \rangle$, hence the MLE given u does not exist in \mathcal{M}_A . We can also see this from Proposition 4.2, as follows. The vector $b = Au = [1 \ 1 \ 1 \ 0 \ 1 \ 1 \ 1 \ 0]^\top$ has some zero entries. Since $P(A)$ only contains nonnegative points, b must lie on the boundary of $P(A)$.

4C. The moment map gives the MLE. In Theorem 4.3, we compared two optimization problems: finding the MLE in a log-linear model, and norm minimization in an orbit under a related torus action. More specifically, we have seen that one problem attains its optimum if and only if the other one does. We now describe how these two optima are related via the moment map. For this, we consider two possible closures of the log-linear model \mathcal{M}_A . The Euclidean closure of the model is the extended log-linear model $\overline{\mathcal{M}}_A$. We also consider the smallest Zariski closed subset of Δ_{m-1} containing \mathcal{M}_A , denoted $\overline{\mathcal{M}}_A^Z$, i.e., the Zariski closure of \mathcal{M}_A in \mathbb{R}^m intersected with the simplex Δ_{m-1} . Its defining ideal is given in (7). In the proof of the following theorem, we use the result that these two closures are equal [15, Theorem 3.2].

Theorem 4.7. *Let $u \in \mathbb{Z}_{\geq 0}^m$ be a vector of counts with $u_+ = n$. Consider a matrix $A \in \mathbb{Z}^{d \times m}$ with $\mathbb{1} \in \text{rowspan}(A)$, and let $b = Au \in \mathbb{Z}^d$. Consider the orbit closure of $\mathbb{1}$ under the torus action of GT_d given by matrix nA with linearization b . Let $q \in \mathbb{C}^m$ be a point in the orbit closure where the moment map vanishes. Then the extended MLE given u for the model \mathcal{M}_A has j -th entry*

$$\frac{|q_j|^2}{\|q\|^2}. \tag{9}$$

If $\mathbb{1}$ is polystable, then this vector is the MLE.

Proof. At a point $q \in \mathbb{C}^m$ where the moment map vanishes, we have $nAq^{(2)} = \|q\|^2 b$ by Theorem 3.5. Consider the vector q' with j -th entry as in (9). We show that q' is the extended MLE given u in \mathcal{M}_A . Since $q' \in \Delta_{m-1}$ and $Aq' = A \frac{u}{n}$, it remains to show that $q' \in \overline{\mathcal{M}}_A$. Using the equality $\overline{\mathcal{M}}_A^Z = \overline{\mathcal{M}}_A$, it suffices to show that q' satisfies the equations in (7). Since q lies in the orbit closure of $\mathbb{1}$, it satisfies the equations in (7), where A is replaced by the matrix with (i, j) entry $na_{ij} - b_j$. That is, it satisfies $q^v - q^w = 0$ for all $v, w \in \mathbb{Z}_{\geq 0}^m$ with $nAv - bw_+ = nAw - bw_+$. We show that this covers all pairs of vectors v, w with $Av = Aw$. Indeed, if $Av = Aw$ then $v_+ = w_+$, because $\mathbb{1}$ is in the row span of A . Hence $q \in \overline{\mathcal{M}}_A$. Now we conclude that q' also satisfies the equations in (7), as follows. For each equation $q^v = q^w$, we can take norms on both sides and square both sides. The equality $v_+ = w_+$ then shows that $(q')^v = (q')^w$, since the denominator is present on both sides with equal power $\|q\|^{2v_+}$.

In the polystable case, the vector q is in the orbit of $\mathbb{1}$, hence has all entries positive. Thus the entries of q' are also positive, so q' is the MLE given u in \mathcal{M}_A . □

Example 4.8. Consider the log-linear model \mathcal{M}_A , and vector of counts u , where

$$A = \begin{bmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \end{bmatrix}, \quad u = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}, \quad b = Au = \begin{bmatrix} 5 \\ 3 \end{bmatrix}.$$

This model is the plane conic $x_2^2 = x_1x_3$. The existence of the MLE given u in \mathcal{M}_A can be characterized by the torus action given by matrix nA with linearization b , by Theorem 4.3, where $n = u_+ = 4$. Since b

is a positive combination of the columns of A , the vector $\mathbb{1}$ is polystable under this action and the MLE given u exists. The MLE relates to a point of minimal 2-norm in the orbit of $\mathbb{1}$ under the torus action given by matrix nA with linearization b , by [Theorem 4.7](#). We show how to obtain the MLE from a point q of minimal norm in the orbit of $\mathbb{1}$.

Since q lies in the orbit of $\mathbb{1}$, its entries are $q_j = \lambda_1^{na_{1j}-b_1} \lambda_2^{na_{2j}-b_2}$, where λ_i are nonzero complex numbers, i.e., $q = [\lambda^3 \ \lambda^{-1} \ \lambda^{-5}]^T$ where $\lambda = \lambda_1/\lambda_2$. Moreover, the moment map vanishes at q , so we have $nAq^{(2)} = \|q\|^2 b$. Combining these, gives the condition $3v^2 - v - 5 = 0$, where $v = |\lambda|^8$, and we obtain that the MLE is

$$\hat{p} = \frac{1}{v^2+v+1} \begin{bmatrix} v^2 \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{31+\sqrt{61}}{4\sqrt{61}+52} \\ \frac{3+3\sqrt{61}}{4\sqrt{61}+52} \\ \frac{9}{2\sqrt{61}+26} \end{bmatrix} \sim \begin{bmatrix} 0.4662 \\ 0.3175 \\ 0.2162 \end{bmatrix}.$$

[Theorem 4.7](#) shows that the MLE can be obtained from norm minimization on an orbit. It suggests the possibility of using algorithms from invariant theory to compute the MLE, as we describe in [Section 5](#). In the next example, we motivate the study of these algorithms, returning to the independence model.

The independence model on a pair of discrete random variables is a log-linear model, as we saw in a special case in [Example 4.1](#). In the following example, we apply [Theorem 4.7](#) to obtain the MLE given u for the independence model from the point of minimal 2-norm in the orbit of $\mathbb{1} \otimes \mathbb{1}$ under a torus action.

Example 4.9. The independence model on a pair of random variables, each with m states, is the log-linear model \mathcal{M}_A , where

$$A = \begin{bmatrix} I_m \otimes \mathbb{1} \\ \mathbb{1} \otimes I_m \end{bmatrix} \in \mathbb{Z}^{2m \times m^2}. \tag{10}$$

The first m rows are $I_m \otimes \mathbb{1}$ and second m rows are $\mathbb{1} \otimes I_m$, where I_m is the $m \times m$ identity matrix, and $\mathbb{1}$ is the all-ones vector of length m . The Kronecker product $A_1 \otimes A_2$ of two matrices $A_k \in \mathbb{R}^{m_k \times n_k}$ is a matrix of size $m_1 m_2 \times n_1 n_2$. We index its rows by (i_1, i_2) where i_k ranges from 1 to m_k , and its columns by (j_1, j_2) , where j ranges from 1 to n_k . The entry of $A_1 \otimes A_2$ at index $((i_1, i_2), (j_1, j_2))$ is $(A_1)_{i_1 j_1} (A_2)_{i_2 j_2}$. See [Example 4.1](#) for (10) in the case $m = 3$.

The model is the orbit of the all-ones matrix under the left-right action of $\text{GT}_m(\mathbb{R}) \times \text{GT}_m(\mathbb{R})$ on the space of $m \times m$ matrices, after restricting to positive entries that sum to one. Equivalently, the model is the orbit under the action in (2) of the torus $\text{GT}_{2m}(\mathbb{R})$ on $\mathbb{C}^{m \times m}$ given by the matrix A in (10), again after restricting to positive entries that sum to one. Equivalently, the model consists of all rank one matrices with positive entries summing to one.

Given a data matrix $u \in \mathbb{Z}^{m \times m}$, we consider the orbit of the all-ones matrix $\mathbb{1} \otimes \mathbb{1} \in \mathbb{C}^{m \times m}$, under the action of $\text{GT}_{2m}(\mathbb{C})$ given by the matrix nA with linearization b , where A is (10), the sample size is

$n = u_{++}$, and $b = Au \in \mathbb{Z}^{2m}$. We seek a matrix in the orbit closure of $\mathbb{1} \otimes \mathbb{1}$ at which the infimum norm is attained. By Kempf–Ness, such matrices are those at which the moment map vanishes. The vanishing of the moment map at $q \in \mathbb{C}^{m \times m}$ gives, by [Theorem 3.5](#),

$$n \begin{bmatrix} |q_{1+}|^2 \\ \vdots \\ |q_{m+}|^2 \\ |q_{+1}|^2 \\ \vdots \\ |q_{+m}|^2 \end{bmatrix} = \|q\|^2 \begin{bmatrix} u_{1+} \\ \vdots \\ u_{m+} \\ u_{+1} \\ \vdots \\ u_{+m} \end{bmatrix}.$$

We relate the MLE to the matrix p with entries $p_{ij} = |q_{ij}|^2 / \|q\|^2$. The matrix p has nonnegative entries summing to one, and the same row and column sums as the empirical distribution \bar{u} . It remains to show that p is in \mathcal{M}_A . For $g = [\lambda_1 \dots \lambda_m \mu_1 \dots \mu_m]$ in $\text{GT}_{2m}(\mathbb{C})$, we have

$$(g \cdot (\mathbb{1} \otimes \mathbb{1}))_{ij} = \lambda_i^n \mu_j^n \left(\prod_{k=1}^m \lambda_k^{-b_k} \right) \left(\prod_{l=1}^m \mu_l^{-b_{m+l}} \right).$$

Hence p is a scalar multiple of the matrix with (i, j) entry $|\lambda_i|^{2n} |\mu_j|^{2n}$, and all such matrices have rank at most one. The latter is a closed condition, so any nonzero p obtained from the orbit closure of $\mathbb{1} \otimes \mathbb{1}$ has rank one. Hence p lies in the closure of the independence model. If the orbit is closed, all entries are positive and it is the MLE. Otherwise, it is the extended MLE; see [Theorem 4.3](#).

5. Scaling algorithms

We saw in [Theorem 4.7](#) that the MLE in a log-linear model can be obtained from a point of minimal norm in an orbit. This connects two problems:

- (1) norm minimization in a complex torus orbit,
- (2) maximum likelihood estimation in a log-linear model.

Algorithms exist for both problems: the former can be approached with convex optimization methods, and the latter with an algorithm called iterative proportional scaling. In fact, both families of algorithms can be thought of as generalizations of two sides of a classical scaling algorithm due to Sinkhorn [\[25\]](#). We explain these different generalizations, and how [Theorem 4.7](#) completes the circle of algorithms, see [Figure 4](#).

5A. Sinkhorn scaling. The classical scaling algorithm of Sinkhorn [\[25\]](#) scales a square matrix M with positive entries to a doubly stochastic matrix. That is, one finds diagonal matrices D_1 and D_2 such that $D_1 M D_2$ has all row sums and all column sums equal to one. The doubly stochastic matrix is obtained by alternately scaling the row and column marginals to one. A natural extension is to scale the matrix M to other fixed row sums and column sums [\[26\]](#). Both versions of Sinkhorn scaling are depicted on the left of [Figure 4](#). These algorithms involve the left-right action of a pair of tori $\text{GT}_{m_1} \times \text{GT}_{m_2}$ on an $m_1 \times m_2$ matrix: the algorithms iterate between updates via the left torus and via the right torus.

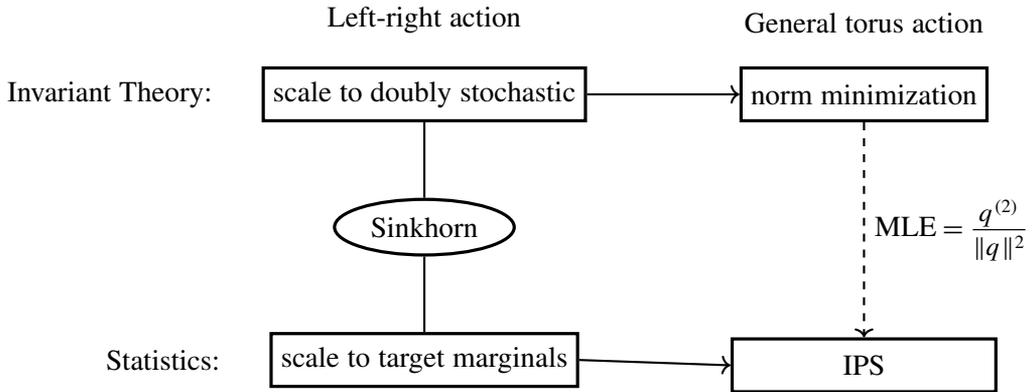


Figure 4. Overview of different scaling algorithms. The historical progression is from left to right, starting with two different sides to Sinkhorn scaling.

Alternating scaling of the rows and columns of a matrix to fixed marginals is an instance of a scaling algorithm which, in the statistics literature, goes back to Deming and Stephan in [9]. For the independence model on two variables, the algorithm finds the MLE by alternating between scaling the row sums and the column sums to match the marginals of the empirical distribution. Given an observed matrix of counts $u \in \mathbb{Z}_{\geq 0}^{m \times m}$ with sample size $u_{++} = n$, and initialized at the uniform distribution, the algorithm has two steps. The (i, j) entry for the two steps is

$$\frac{1}{m^2} \mapsto \frac{1}{m} \cdot \frac{u_{i+}}{n} \mapsto \frac{u_{i+}}{n} \cdot \frac{u_{+j}}{n}. \tag{11}$$

If all entries are positive, the output is the MLE to the independence model given u , otherwise it is the extended MLE. This is the first example of *iterative proportional scaling* (IPS), which we describe in the next subsection.

5B. Iterative proportional scaling. In the previous section, we saw that alternating scaling of a matrix to fixed row and column sums gives the MLE to the independence model, when initialized at the uniform distribution. This is scaling under a product of tori $GT_m \times GT_m$. We saw in Examples 4.1 and 4.9 how the independence model fits into the framework of log-linear models. In terms of the group action, this replaces the left-right action of a pair of tori $GT_m \times GT_m$ with the action of a single torus GT_{2m} , acting via (2), where A is the matrix in (10).

In this section, we explain how Sinkhorn scaling extends to algorithms for maximum likelihood estimation for a general log-linear model, the bottom arrow of Figure 4.

Alternating between matching row and column sums can be extended to hierarchical models, which summarize data by contingency tables [12], by iteratively updating the various marginals. The approach was extended to more general log-linear models by Darroch and Ratcliff in [8].

For the log-linear model \mathcal{M}_A , the MLE \hat{p} must satisfy the equation $A\hat{p} = A\bar{u}$ from Birch’s theorem, where $\bar{u} = \frac{u}{n}$ is the empirical distribution. IPS finds the extended MLE in \mathcal{M}_A given an empirical distribution $\bar{u} \in \Delta_{m-1}$. We define IPS for a log-linear model given by a matrix $A \in \mathbb{Z}_{\geq 0}^{d \times m}$ whose column

sums are all equal. Starting at the uniform distribution $p^{(0)} = \frac{1}{m}\mathbb{1}$, we iterate until the k -th update $p^{(k)}$ has sufficient statistics $b^{(k)} = Ap^{(k)}$ close to the target sufficient statistics $b = A\bar{u}$, i.e., until (8) holds approximately. The update step is

$$p_j^{(k+1)} = \prod_{i=1}^d \left(\frac{(A\bar{u})_i}{(Ap^{(k)})_i} \right)^{a_{ij}/\alpha} p_j^{(k)}, \quad (12)$$

where α is the common column sum of A ; see [27, Algorithm 7.3.11]. This is the action of a torus element (obtained by componentwise division of $A\bar{u}$ by $Ap^{(k)}$ and then componentwise exponentiation by $1/\alpha$) on the vector $p^{(k)}$. Here the torus action is given by the matrix A with linearization $b = 0$, see (3).

We can view maximum likelihood estimation as a norm minimization problem in a different way to Theorem 4.7, by interpreting IPS as minimizing KL divergence.

Proposition 5.1. *Consider the log-linear model \mathcal{M}_A where $A \in \mathbb{Z}^{d \times m}$ has $\mathbb{1}$ in its row span. Then there exists a matrix $\tilde{A} \in \mathbb{Q}_{\geq 0}^{(d+1) \times m}$, with all column sums equal, such that $\mathcal{M}_A = \mathcal{M}_{\tilde{A}}$, iterative proportional scaling in (12) with matrix \tilde{A} converges, and at each update step the KL divergence to the MLE decreases.*

Proof. The proof of convergence of IPS is given in [8, Theorem 1] in the case where the entries of A are real and nonnegative with each column of A summing to one. There, the authors show that each step of IPS decreases the KL divergence $\text{KL}(\hat{p} \| p^{(k)})$ from the k -th iterate $p^{(k)}$ to the MLE \hat{p} . Since replacing A by $\frac{1}{\alpha}A$ does not change the update step (12), the KL divergence also decreases for any matrix with real and nonnegative entries and all column sums equal.

We explain how this covers log-linear models defined by integer matrices with $\mathbb{1}$ in the row span. We modify A without changing its row span, i.e., without changing the model \mathcal{M}_A . First, we add a sufficiently large positive integer to every entry of A . For a general choice of integer, this does not change $\text{rowspan}(A)$ since it adds a multiple of the vector $\mathbb{1}$, which belongs to $\text{rowspan}(A)$, to every row. Second, let α be the maximum of the column sums a_{+j} . Add another row to the matrix, with entries $\alpha - a_{+j}$. The extra row is a linear combination of $\mathbb{1}$ and the rows of A , so the augmented matrix has the same row span as A . The column sums of the augmented matrix \tilde{A} are all α . \square

Remark 5.2. We saw in Section 2 that $\hat{p} = \text{argmin}_{p \in \mathcal{M}} \text{KL}(\bar{u} \| p)$. Here, we use KL divergence differently, measuring the KL divergence from iterate $p^{(k)}$ to the MLE, $\text{KL}(\hat{p} \| p^{(k)})$.

Curiously, when IPS for log-linear models in (12) is applied to the independence model, we do not recover the classical IPS with Sinkhorn updates, because the column sums of the integer matrix A for the independence model in (10) are $\alpha = 2$, hence there is a square root in the update step. If, instead, we did IPS with the same matrix A but $\alpha = 1$ in (12) we would recover the two steps in (11) in a single step. This leads naturally to the question of which exponents α achieve convergence, and how the choice of α affects the convergence rate. This is the essence of an open problem in algebraic statistics, see [11, Section 7.3].

5C. Norm minimization. In this section, we explain how Sinkhorn scaling generalizes to norm minimization in invariant theory; see the top arrow of Figure 4.

The condition that a matrix can be scaled to a doubly stochastic matrix is dual to testing membership in the null cone under a group action, as follows. We consider pairs of diagonal matrices (D_1, D_2) of

determinant one that act on square matrices M via $M \mapsto D_1 M D_2$. A matrix does not lie in the null cone under this action if and only if its orbit closure contains a matrix M such that the matrix with (i, j) entry $|m_{ij}|^2$ is a nonzero scalar multiple of a doubly stochastic matrix [14, Corollary 3.6]. This is an instance of Kempf–Ness. Norm minimization on the orbit of a square matrix either converges to zero, or to a nonzero matrix M at which the moment map vanishes. The condition $\mu(M) = 0$ translates to the matrix with entries $|m_{ij}|^2$ being a scalar multiple of a doubly stochastic matrix. So we see that norm minimization scales to a nonzero multiple of a doubly stochastic matrix, if such a matrix exists in the orbit.

Norm minimization on an orbit can be considered for a wide range of groups and their actions. If a group G can be expressed as a product of groups, then the alternating minimization idea from Sinkhorn’s algorithm generalizes. An important example of this is operator scaling, which solves the scaling problem for the left-right action of $\mathrm{SL}_{m_1}(\mathbb{C}) \times \mathrm{SL}_{m_2}(\mathbb{C})$ on the space of matrix tuples $(\mathbb{C}^{m_1 \times m_2})^n$. We discuss connections between operator scaling and statistics in our companion paper [1].

We consider the norm minimization problem for the action of the torus $\mathrm{GT}_d(\mathbb{C})$ given by the matrix nA with linearization b . We take $b = Au$, where u is a vector of counts. By Kempf–Ness, Theorem 3.5, a vector is not in the null cone if and only if there is a nonzero vector q in its orbit closure satisfying $Aq^{(2)} = \|q\|^2 b$. Hence, the problem of scaling a vector by acting with the torus to such a nonzero vector q is dual to testing membership in the null cone under the torus action. This duality generalizes the discussion of doubly stochastic matrices above.

Since the vector $\mathbb{1}$ is semistable, see Theorem 4.3, norm minimization converges to such a nonzero vector q . This is a convex optimization problem, as follows. Consider the action of $\mathrm{GT}_d(\mathbb{C})$ given by matrix $A' = nA - b \otimes \mathbb{1} \in \mathbb{Z}^{d \times m}$. For a torus element $(\lambda_1, \dots, \lambda_d)$, the coordinate change $y_i := \log |\lambda_i|^2$ gives

$$\mathrm{cap}(\mathbb{1}) = \inf_{\lambda \in \mathrm{GT}_d(\mathbb{C})} \|\lambda \cdot \mathbb{1}\|^2 = \inf_{\lambda \in \mathrm{GT}_d(\mathbb{C})} \sum_{j=1}^m \prod_{i=1}^d |\lambda_i|^{2a'_{ij}} = \inf_{y \in \mathbb{R}^d} \sum_{j=1}^m \exp\langle y, a'_j \rangle.$$

Convexity then follows from the fact that each exponential function is convex and a sum of convex functions is convex. This minimization problem is known as geometric programming. Hence, common algorithms from the vast literature on convex optimization can be used to compute the capacity and find the MLE, e.g., interior point methods [6] or ellipsoid methods.

5D. Comparison of algorithms. We have seen in the previous two subsections that IPS and norm minimization are generalizations of Sinkhorn scaling that have emerged in different communities. Theorem 4.7 closes the cycle of algorithms from different communities, by showing how to obtain the (extended) MLE from a complex point of minimal norm in an orbit (or orbit closure); see Figure 4.

This bridges several differences between IPS and norm minimization. We summarize these differences here. First, when computing the capacity, the norm is minimized along a *complex* orbit closure (see Theorem 4.7), whereas every step in IPS involves *real* numbers. Secondly, the torus action given by matrix nA that is used for computing the capacity is linearized by $b = Au$ (see Theorem 4.7), whereas IPS uses the action given by matrix A with trivial linearization $b = 0$. Finally, the objective functions differ: the capacity is defined in terms of the Euclidean norm, which does not appear in IPS; instead

IPS minimizes KL divergence (see [Proposition 5.1](#)). In the following example we see that, while IPS decreases the KL divergence to the MLE, it may increase the Euclidean norm.

Example 5.3. Consider the matrix A and vector of counts u from [Example 4.8](#). We can use IPS to compute the MLE in \mathcal{M}_A . We start at the uniform distribution $p^{(0)} = \frac{1}{3}\mathbb{1}$ and do update steps as in [\(12\)](#) with matrix A . These IPS steps converge by [Proposition 5.1](#), since the matrix A has real nonnegative entries and all column sums equal. We obtain

$$p^{(1)} = \left[\frac{5}{12} \quad \frac{\sqrt{15}}{12} \quad \frac{1}{4} \right]^T.$$

Note that the sum of the entries of $p^{(1)}$ is strictly less than one. The KL divergence from the uniform distribution to the MLE is $\text{KL}(\hat{p} \| p^{(0)}) \sim 0.047$, and after the first update it is $\text{KL}(\hat{p} \| p^{(1)}) \sim 0.016$. However, we have $\|p^{(1)}\|^2 = \frac{49}{144}$, which exceeds $\|p^{(0)}\|^2 = \frac{1}{3}$.

6. Comparison with multivariate Gaussian models

We highlight similarities and differences with the multivariate Gaussian setting studied in [\[1\]](#). For this, we compare results from this paper with the related results in [\[1\]](#). We start by comparing the two statistical settings.

In the discrete setting, a model is given as a subset of the $(m-1)$ -dimensional probability simplex $\Delta_{m-1} \subseteq \mathbb{R}^m$. In comparison, in the multivariate Gaussian setting, a model is given by a set of concentration matrices Ψ in the cone of positive definite matrices. For a discrete model $\mathcal{M} \subseteq \Delta_{m-1}$ the data is a vector of counts $u \in \mathbb{Z}_{\geq 0}^m$ with $u_+ = n$ the total numbers of observations. The log-likelihood given u at $p \in \mathcal{M}$ is $\sum_{j=1}^m u_j \log(p_j)$. In comparison, for a Gaussian model the data is summarized by the sample covariance matrix

$$S_Y = \frac{1}{n} \sum_{i=1}^n Y_i Y_i^T$$

and the log-likelihood given a tuple of samples $Y \in (\mathbb{R}^m)^n$ is $\log \det(\Psi) - \text{tr}(\Psi S_Y)$.

6A. Stability. In both papers we link notions of stability under group actions to maximum likelihood estimation of certain statistical models: for log-linear models in [Theorem 4.3](#) and for Gaussian group models in [\[1, Section 3\]](#). For log-linear models it is enough to consider actions of complex tori on \mathbb{C}^m . In contrast, in [\[1\]](#) we work with actions of (reductive algebraic) groups over \mathbb{R} or \mathbb{C} , depending on whether we consider multivariate Gaussian distributions on \mathbb{R}^m or \mathbb{C}^m . In the log-linear case we study stability of the all-ones vector, while in [\[1\]](#) we consider the stability notions for the observed tuple of samples.

For log-linear models, the log-likelihood is always bounded from above and the all-ones vector cannot be unstable. In contrast, in the Gaussian setting a tuple of samples is unstable if and only if the log-likelihood is not bounded from above. In both cases, semistability is equivalent to the log-likelihood being bounded from above and polystability is equivalent to the existence of an MLE. In the log-linear case, the MLE is unique if it exists, while for Gaussian group models there may be infinitely many. In

fact, the existence of a unique MLE for Gaussian group models relates to stability of a tuple of samples. In contrast, for log-linear models the all-ones vector is never stable.

6B. MLE computation. An important similarity between the log-linear and Gaussian settings is that norm minimizers under the respective group actions give an MLE (if it exists), see [Theorem 4.7](#) and [\[1, Section 3\]](#). For log-linear models, we compute real MLEs from complex torus orbits. However, for Gaussian group models, we compute the MLE over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ from orbits over the same field \mathbb{K} . If the all-ones vector is semistable but not polystable, [Theorem 4.7](#) yields the extended MLE. However, in the Gaussian case, if a tuple of samples Y is semistable but not polystable there is no corresponding notion of extended MLE.

6C. Scaling. From the point of view of scaling algorithms, Sinkhorn’s algorithm is a common origin to both the log-linear and the Gaussian settings. As we described in [Section 5](#), Sinkhorn scaling to target marginals is IPS for the independence model and this extends to IPS for a general log-linear model. On the Gaussian side, Sinkhorn scaling generalizes to alternating minimization procedures for computing MLEs of matrix normal models and tensor normal models. This algorithm is used both in invariant theory for norm minimization and in statistics to compute the MLE; see [\[1\]](#). In contrast, for log-linear models the algorithms from invariant theory and statistics are not the same; see [Figure 4](#).

We conclude this paper by pointing out that log-linear models and the Gaussian group models in [\[1\]](#) are examples of exponential transformation families. Hence, it is an interesting and natural question to ask whether there is a unifying concept that links invariant theory to maximum likelihood estimation for exponential families.

Appendix A: Hilbert–Mumford for a torus action

Almost all the results in this paper use the polyhedral characterization of stability under a torus action, given by the Hilbert–Mumford criterion ([Theorem 3.4](#)). In this appendix we present a proof of [Theorem 3.4](#), in what we hope is an elementary and accessible style.

Remark A.1 (disregarding the linearization). The setting of [Theorem 3.4](#) is a torus action of GT_d on \mathbb{C}^m given by a matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$. For our statistical connections, it is important to separate the role of A (which is fixed by the model) from that of b (which depends on the data). However, we can remove the need for a linearization by altering the matrix A , as follows. The action of GT_d on \mathbb{C}^m given by matrix $A \in \mathbb{Z}^{d \times m}$ with linearization $b \in \mathbb{Z}^d$ is the action given by matrix $A' \in \mathbb{Z}^{d \times m}$ with linearization $0 \in \mathbb{Z}^d$, where the matrix A' has j -th column $a_j - b$, see [\(3\)](#). Hence we assume without loss of generality that the linearization is zero for proving [Theorems 3.4](#) and [3.5](#). The effect of the linearization on the moment map is outlined in [Section 3C](#).

The classical statement of the Hilbert–Mumford criterion, see [\[22, page 53\]](#), concerns one parameter subgroups. For the group GT_d , a one parameter subgroup is given by a map

$$\sigma : \mathbb{C}^\times \rightarrow \text{GT}_d, \quad \lambda \mapsto (\lambda^{\sigma_1}, \dots, \lambda^{\sigma_d}), \tag{13}$$

for some fixed $(\sigma_1, \dots, \sigma_d) \in \mathbb{Z}^d$. We use σ to denote both the map and the vector $(\sigma_1, \dots, \sigma_d)$. For $v \in \mathbb{C}^m$, the j -th entry of $\sigma(\lambda) \cdot v$ is

$$(\sigma(\lambda) \cdot v)_j = \lambda^{(a_j, \sigma)} v_j.$$

We consider $\lim_{\lambda \rightarrow 0} \sigma(\lambda) \cdot v$. The j -th entry of the limiting vector is zero for $j \notin \text{supp}(v)$. For $j \in \text{supp}(v)$, we have three possibilities:

$$\left(\lim_{\lambda \rightarrow 0} \sigma(\lambda) \cdot v \right)_j = \begin{cases} 0 & \text{if } \langle \sigma, a_j \rangle > 0, \\ v_j & \text{if } \langle \sigma, a_j \rangle = 0, \\ \infty & \text{if } \langle \sigma, a_j \rangle < 0. \end{cases} \tag{14}$$

The classical statement of the Hilbert Mumford criterion for a torus action is as follows.

Theorem A.2. *Consider the action of GT_d on \mathbb{C}^m via the matrix $A \in \mathbb{Z}^{d \times m}$. Given a nonzero $v \in \mathbb{C}^m$, with zero in its orbit closure, there exists a one-parameter subgroup of GT_d that scales v to zero in the limit.*

We give a proof of [Theorem A.2](#) following [\[28\]](#). Other references for the statement of the theorem include [\[30, Proposition 5.3; 4, Lemma 3.4\]](#) for a torus, and [\[4, Theorem 4.1; 22, page 53; 30, Theorem 5.2\]](#) for a general reductive group.

Proof of Theorem A.2. We seek a one parameter subgroup $\sigma : \mathbb{C}^\times \rightarrow \text{GT}_d$ such that $\lim_{\lambda \rightarrow 0} \sigma(\lambda) \cdot v$ is zero. From the form of a one parameter subgroup from [\(13\)](#) and the limiting behavior from [\(14\)](#), we see that this is equivalent to showing that

$$\text{there exists } \sigma \in \mathbb{Z}^d \text{ such that } \langle \sigma, a_j \rangle > 0 \text{ for all } j \in \text{supp}(v). \tag{15}$$

Reordering the entries of v , we can assume without loss of generality that $\text{supp}(v) = [k]$ for some $k \leq m$. Then the existence of such a $\sigma \in \mathbb{Z}^d$ as in [\(15\)](#) is equivalent to the following statement about $A \in \mathbb{Z}^{m \times d}$:

if $t = (t_1, \dots, t_k) \in \mathbb{R}^k \setminus \{0\}$ is such that $t_1 a_{i1} + \dots + t_k a_{ik} = 0$ for all $i \in [d]$
 then at least two entries of t are of opposite sign. (16)

The equivalence of [\(15\)](#) and [\(16\)](#) is [\[28, Lemma 1.1\]](#), and is an analogue of Gordan’s theorem [\[24, Section 7.8 Equation \(31\)\]](#) over the rational numbers. Thus it remains to prove [\(16\)](#).

Since v has 0 in its orbit closure, there exists a sequence $\lambda^{(n)} = (\lambda_1^{(n)}, \dots, \lambda_d^{(n)}) \in \text{GT}_d$ with $\lambda^{(n)} \cdot v \rightarrow 0$ as $n \rightarrow \infty$. In coordinates,

$$(\lambda_1^{(n)})^{a_{1j}} \dots (\lambda_d^{(n)})^{a_{dj}} \rightarrow 0 \quad \text{as } n \rightarrow \infty \text{ for all } j \in [k]. \tag{17}$$

The hypothesis of [\(16\)](#) is that we have $t \in \mathbb{R}^k \setminus \{0\}$ with $t_1 a_{i1} + \dots + t_k a_{ik} = 0$ for all $i \in [d]$. Without loss of generality, we can assume t_1 is nonzero and therefore

$$-a_{i1} = \frac{t_2}{t_1} a_{i2} + \dots + \frac{t_k}{t_1} a_{ik} \quad \text{for all } i \in [d], \tag{18}$$

which implies

$$\prod_{i=1}^d (\lambda_i^{(n)})^{-a_{i1}} = \left(\prod_{i=1}^d (\lambda_i^{(n)})^{a_{i2}} \right)^{t_2/t_1} \dots \left(\prod_{i=1}^d (\lambda_i^{(n)})^{a_{ik}} \right)^{t_k/t_1}. \tag{19}$$

If $t_j/t_1 \geq 0$ for all $j \in \{2, \dots, k\}$, then the right-hand side of (19) either equals one (if all t_j/t_1 are zero) or tends to zero (if there exists some j with $t_j/t_1 > 0$). But the left-hand side of (19) tends to infinity as $n \rightarrow \infty$, since it is the inverse of (17) for $j = 1$. Hence t_j/t_1 must be strictly negative for some j , i.e., t_1 and t_j have opposite signs. \square

Theorem A.2 has the following generalization. It can be proved via a polyhedral geometry argument similar to the proof of Theorem A.2.

Theorem A.3 [20, page 173]. *Consider the action of GT_d on \mathbb{C}^m given by matrix $A \in \mathbb{Z}^{d \times m}$, and fix $v \in \mathbb{C}^m$. If $w \in \overline{\text{GT}_d \cdot v} \setminus \text{GT}_d \cdot v$, then there exists a one-parameter subgroup that scales v to an element of $\text{GT}_d \cdot w$ in the limit.*

Remark A.4. Theorem A.3 has an analogue for a general reductive group G , but it only applies with the additional assumption that $G \cdot w$ is the unique closed orbit in $\overline{G \cdot v}$, see [30, Section 6.8].

Equipped with Theorems A.2 and A.3, we now prove Theorem 3.4. The proof mostly rests on Theorem A.2; we only use the stronger statement of Theorem A.3 for one direction of one of the four cases.

Proof of Theorem 3.4. We first prove parts (a) and (b). If $v = 0$, then the polytope $P_v(A)$ is empty, hence $0 \notin P_v(A)$. If $v \neq 0$ is unstable, then there exists some $\sigma \in \mathbb{Z}^d$ such that $\langle \sigma, a_j \rangle > 0$ for all $j \in \text{supp}(v)$, by combining Theorem A.2 with (14). Hence σ defines a hyperplane

$$H_\sigma = \{x \in \mathbb{R}^d \mid \langle \sigma, x \rangle = 0\}$$

that separates zero from $P_v(A)$. By Farkas' lemma, see [24, Section 7.3], such a hyperplane exists if and only if $0 \notin P_v(A)$.

For (c), we first prove that if 0 is on the boundary of $P_v(A)$, then v is not polystable. We construct a point in the orbit closure of v , with support strictly smaller than that of v , and hence deduce that the orbit of v is not closed. Since 0 lies on the boundary of $P_v(A)$, it is contained in a minimal face $F \subsetneq P_v(A)$. Since A has integer entries, there is a hyperplane H_σ , with $\sigma \in \mathbb{Z}^d$, such that $F = H_\sigma \cap P_v(A)$. We choose the sign of σ so that it has nonnegative inner product with all of $P_v(A)$. This ensures that the limit $w := \lim_{\lambda \rightarrow 0} \sigma(\lambda) \cdot v$ exists. The limit w has $\text{supp}(w) \subsetneq \text{supp}(v)$, since $P_w(A) \subseteq F$. Hence $w \in \overline{\text{GT}_d \cdot v} \setminus \text{GT}_d \cdot v$, and $\text{GT}_d \cdot v$ is not closed.

For the converse direction of (c), we show that if v is semistable but not polystable, then $0 \notin \text{relint}(P_v(A))$. Let $w' \in \overline{\text{GT}_d \cdot v} \setminus \text{GT}_d \cdot v$. There exists $\sigma \in \mathbb{Z}^d$ such that $w := \lim_{\lambda \rightarrow 0} \sigma(\lambda) \cdot v \in \text{GT}_d \cdot w'$, by Theorem A.3. We have $\text{supp}(w) \subseteq \text{supp}(v)$ and, moreover, $\text{supp}(w) \subsetneq \text{supp}(v)$ (otherwise $w = v$ by (14), a contradiction). Hence $\langle \sigma, a_j \rangle > 0$ for all $j \in \text{supp}(v) \setminus \text{supp}(w)$, while $\langle \sigma, a_j \rangle = 0$ for all $j \in \text{supp}(w)$, by (14). We obtain $P_v(A) \not\subseteq H_\sigma$ and $P_w(A) = H_\sigma \cap P_v(A)$, i.e., $P_w(A)$ is a proper face of $P_v(A)$. We have $\text{GT}_d \cdot w = \text{GT}_d \cdot w' \subseteq \overline{\text{GT}_d \cdot v}$ and so w is semistable as v is semistable. By (b), $0 \in P_w(A)$ and hence 0 is on the boundary of $P_v(A)$.

It remains to prove (d). We can assume v is polystable, i.e., $0 \in \text{relint}(P_v(A))$. We want to show that the dimension of the stabilizer $\{\lambda \in \text{GT}_d \mid \lambda \cdot v = v\}$ is zero if and only if the interior of $P_v(A)$ equals its relative interior (i.e., if and only if $P_v(A)$ is full-dimensional). Since $0 \in P_v(A)$, the equality of the interior and relative interior holds if and only if $U := \text{span}\{a_j \mid j \in \text{supp}(v)\}$ equals \mathbb{R}^d . If the stabilizer is

positive dimensional, it must contain a one-parameter subgroup, i.e., some $\sigma \in \mathbb{Z}^d \setminus \{0\}$ with $\sigma(\lambda) \cdot v = v$ for all $\lambda \in \mathbb{C}^\times$. Then $\langle \sigma, a_j \rangle = 0$ for all $j \in \text{supp}(v)$, so the orthogonal complement $U^\perp \subseteq \mathbb{R}^d$ contains a line, and $U \neq \mathbb{R}^d$. Conversely, if $U \neq \mathbb{R}^d$ then there exists $\sigma \in U^\perp$, which can be chosen to have integer entries, since A has integer entries. The one parameter subgroup $\sigma(\lambda)$ then lies in the stabilizer, which is therefore positive-dimensional. \square

Appendix B: Kempf–Ness for a torus action

Many of the results in this paper use the Kempf–Ness theorem for a torus action, as stated in [Theorem 3.5](#). An elementary proof of the Kempf–Ness theorem for torus actions can be found in Sections 1 and 2 of the original paper of Kempf and Ness [\[19\]](#). In this appendix, we translate between the setting of [\[19\]](#) and our setting, to explain how the results of [\[19\]](#) give [Theorem 3.5](#). In addition, we present an alternative proof of [Theorem 3.5](#), obtaining it as a consequence of [Theorem 3.4](#).

As before, we consider the action of GT_d on \mathbb{C}^m given by a matrix $A \in \mathbb{Z}^{d \times m}$. We can assume without loss of generality that the linearization is $b = 0$, by [Remark A.1](#). We first describe how [Theorem 3.5](#) follows from [\[19\]](#).

First proof of [Theorem 3.5](#). Let $v \in \mathbb{C}^m$ and consider the action of GT_d via the matrix $A \in \mathbb{Z}^{d \times m}$. Recall that the moment map at v is the derivative $D_I \gamma_v$, where

$$\gamma_v : \text{GT}_d \rightarrow \mathbb{R}, \quad \lambda \mapsto \|\lambda \cdot v\|^2.$$

Identifying the space of \mathbb{R} -linear functionals $\text{Hom}(\mathbb{C}^d, \mathbb{R})$ with \mathbb{C}^d gives the moment map

$$\mu : \mathbb{C}^m \rightarrow \mathbb{C}^d, \quad v \mapsto 2Av^{(2)},$$

where $v^{(2)}$ is the vector with j -th coordinate $|v_j|^2$.

We translate between our notation and that in [\[19\]](#). Most importantly, our notion of “polystable” is called “stable” by Kempf and Ness, see [\[19, page 234\]](#). Our function γ_v is denoted p_v in [\[19\]](#). Moreover, “a critical point $g \in G$ of p_v ” in [\[19\]](#) means the vanishing of the moment map at $g \cdot v$:

$$0 = D_g p_v = D_I p_{g \cdot v} = D_I \gamma_{g \cdot v} = \mu(g \cdot v).$$

Thus the *polystable* part of [Theorem 3.5](#) is a direct consequence of [\[19, Theorems 0.1\(a\) and 0.2\]](#) for $G = \text{GT}_d$. The *semistable* part follows from the polystable part, using the fact that any orbit closure for the group GT_d contains a unique closed orbit, see, e.g., [\[20, Bemerkung 1 on page 96\]](#). That unique closed orbit is not the zero orbit if and only if the vector is semistable. \square

We end with an alternative proof of [Theorem 3.5](#), which uses an important connection between the polytope $P_v(A)$ and the moment map μ , see [\[2; 16\]](#). The connection relates $P_v(A)$ to the image of the orbit $\text{GT}_d \cdot v$ under the moment map. It is a first example of a moment polytope, an important object of study in invariant theory.

Second proof of Theorem 3.5. To link the moment map μ to the polytope $P_v(A)$, we need to rescale μ as follows:

$$\tilde{\mu} : \mathbb{C}^m \setminus \{0\} \rightarrow \mathbb{R}^d, \quad v \mapsto \frac{\mu(v)}{2\|v\|^2} = \frac{Av^{(2)}}{\|v\|^2}. \tag{20}$$

The vector $v^{(2)}/\|v\|^2$ consists of nonnegative numbers that sum up to one. Hence $\tilde{\mu}(v)$ is a convex combination of the columns of A , therefore $\tilde{\mu}(v) \in P_v(A)$. The stronger statement

$$\text{relint } P_v(A) = \tilde{\mu}(\text{GT}_d \cdot v), \tag{21}$$

was proven independently in [2, Theorem 2; 16, Theorem 4], see Remark B.1. Given (21), the polystable case of Theorem 3.5 is a direct consequence of Theorem 3.4(c), as follows. Polystability is equivalent to $0 \in \text{relint } P_v(A)$ which, by (21), implies $0 \in \tilde{\mu}(\text{GT}_d \cdot v)$, i.e., that $0 = Aw^{(2)}/\|w\|^2$ for some $w \in \text{GT}_d \cdot v$. As in the first proof, the semistable case of Theorem 3.5 can be deduced from the polystable case, since any orbit closure $\overline{\text{GT}_d \cdot v}$ contains a unique closed orbit. \square

Remark B.1. The statements in [2, Theorem 2; 16, Theorem 4] discuss a moment map whose domain is a projective space $\mathbb{P}_{\mathbb{C}}^{m-1}$, rather than the space $\mathbb{C}^m \setminus \{0\}$ in (20). However, the projective results still allow us to obtain (21), as follows. For nonzero $v \in \mathbb{C}^m$, let $[v]$ be the point in $\mathbb{P}_{\mathbb{C}}^{m-1}$ that represents the line $\mathbb{C}v$. We consider the action of GT_d on $\mathbb{P}_{\mathbb{C}}^{m-1}$ given by the matrix $A \in \mathbb{Z}^{d \times m}$. We denote the GT_d -orbit of $[v]$ in $\mathbb{P}_{\mathbb{C}}^{m-1}$ by $\text{GT}_d \cdot [v]$ (and denote the orbit closure by $\overline{\text{GT}_d \cdot [v]}$). The map $\tilde{\mu}$ factors through the projective space $\mathbb{P}_{\mathbb{C}}^{m-1}$ via a unique map $\bar{\mu} : \mathbb{P}_{\mathbb{C}}^{m-1} \rightarrow \mathbb{R}^d$. In fact, $\bar{\mu}$ is the moment map for the action of GT_d on $\mathbb{P}_{\mathbb{C}}^{m-1}$ given by $A \in \mathbb{Z}^{d \times m}$. This action fits the setting of [2; 16] because $\mathbb{P}_{\mathbb{C}}^{m-1}$ is a compact Kähler manifold. The results [2, Theorem 2; 16, Theorem 4] give

$$P_v(A) = \bar{\mu}(\overline{\text{GT}_d \cdot [v]}).$$

For (21), we require a statement for the orbit of v rather than the orbit closure of $[v]$. The closure $\overline{\text{GT}_d \cdot [v]}$ is the disjoint union of finitely many GT_d orbits. The orbits relate to $P_v(A)$ as follows. For each open face F of $P_v(A)$ the set $\bar{\mu}^{-1}(F) \cap \overline{\text{GT}_d \cdot [v]}$ is a single GT_d orbit in $\mathbb{P}_{\mathbb{C}}^{m-1}$, see [2, Theorem 2]. In particular, when $F = \text{relint } P_v(A)$ we obtain the orbit $\text{GT}_d \cdot [v]$. This yields (21), since $\bar{\mu}(\text{GT}_d \cdot [v]) = \tilde{\mu}(\text{GT}_d \cdot v)$.

Acknowledgements

We are grateful to Peter Bürgisser, Mathias Drton, and Bernd Sturmfels for fruitful discussions. We are also grateful to an anonymous referee for helpful comments that improved the paper. Améndola was partially supported by the Deutsche Forschungsgemeinschaft (DFG) in the context of the Emmy Noether junior research group KR 4512/1-1. Kohn was partially supported by the Knut and Alice Wallenberg Foundation within their WASP (Wallenberg AI, Autonomous Systems and Software Program) AI/Math initiative. Research of Reichenbach is funded by the European Research Council (ERC) under the European’s Horizon 2020 research and innovation programme (grant agreement no. 787840).

References

- [1] C. Améndola, K. Kohn, P. Reichenbach, and A. Seigal, “Invariant theory and scaling algorithms for maximum likelihood estimation”, *SIAM J. Appl. Algebra Geom.* **5**:2 (2021), 304–337.
- [2] M. F. Atiyah, “Convexity and commuting Hamiltonians”, *Bull. London Math. Soc.* **14**:1 (1982), 1–15.
- [3] M. W. Birch, “Maximum likelihood in three-way contingency tables”, *J. Roy. Statist. Soc. Ser. B* **25**:1 (1963), 220–233.
- [4] D. Birkes, “Orbits of linear algebraic groups”, *Ann. of Math. (2)* **93** (1971), 459–475.
- [5] Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland, *Discrete multivariate analysis: theory and practice*, reprint of the 1975 original ed., Springer, 2007.
- [6] P. Bürgisser, Y. Li, H. Nieuwboer, and M. Walter, “Interior-point methods for unconstrained geometric programming and scaling problems”, preprint, 2020. [arXiv 2008.12110](https://arxiv.org/abs/2008.12110)
- [7] D. A. Cox, J. B. Little, and H. K. Schenck, *Toric varieties*, Graduate Studies in Mathematics **124**, Amer. Math. Soc., Providence, RI, 2011.
- [8] J. N. Darroch and D. Ratcliff, “Generalized iterative scaling for log-linear models”, *Ann. Math. Statist.* **43** (1972), 1470–1480.
- [9] W. E. Deming and F. F. Stephan, “On a least squares adjustment of a sampled frequency table when the expected marginal totals are known”, *Ann. Math. Statistics* **11**:4 (1940), 427–444.
- [10] I. Dolgachev, *Lectures on invariant theory*, London Mathematical Society Lecture Note Series **296**, Cambridge University Press, 2003.
- [11] M. Drton, B. Sturmfels, and S. Sullivant, *Lectures on algebraic statistics*, Oberwolfach Seminars **39**, Birkhäuser, 2009.
- [12] S. E. Fienberg, “An iterative procedure for estimation in contingency tables”, *Ann. Math. Statist.* **41** (1970), 907–917.
- [13] S. E. Fienberg and A. Rinaldo, “Maximum likelihood estimation in log-linear models”, *Ann. Statist.* **40**:2 (2012), 996–1023.
- [14] A. Garg and R. Oliveira, “Recent progress on scaling algorithms and applications”, *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **125** (2018), 14–49.
- [15] D. Geiger, C. Meek, and B. Sturmfels, “On the toric algebra of graphical models”, *Ann. Statist.* **34**:3 (2006), 1463–1492.
- [16] V. Guillemin and S. Sternberg, “Convexity properties of the moment mapping”, *Invent. Math.* **67**:3 (1982), 491–513.
- [17] S. J. Haberman, “Log-linear models for frequency tables derived by indirect observation: maximum likelihood equations”, *Ann. Statist.* **2** (1974), 911–924.
- [18] D. Hilbert, “Ueber die Theorie der algebraischen Formen”, *Math. Ann.* **36**:4 (1890), 473–534.
- [19] G. Kempf and L. Ness, “The length of vectors in representation spaces”, pp. 233–243 in *Algebraic geometry* (Univ. Copenhagen, 1978), Lecture Notes in Math. **732**, Springer, 1979.
- [20] H. Kraft, *Geometrische Methoden in der Invariantentheorie*, Band D1, Friedr. Vieweg & Sohn, Braunschweig, 1984.
- [21] S. L. Lauritzen, *Graphical models*, Oxford Statistical Science Series **17**, Clarendon Press, 1996.
- [22] D. Mumford, J. Fogarty, and F. Kirwan, *Geometric invariant theory*, 3rd ed., Ergebnisse der Mathematik (2) **34**, Springer, 1994.
- [23] L. Pachter and B. Sturmfels, “Biology”, pp. 125–159 in *Algebraic statistics for computational biology*, edited by L. Pachter and B. Sturmfels, Cambridge Univ. Press, New York, 2005.
- [24] A. Schrijver, *Theory of linear and integer programming*, Wiley, 1986.
- [25] R. Sinkhorn, “A relationship between arbitrary positive matrices and doubly stochastic matrices”, *Ann. Math. Statist.* **35** (1964), 876–879.
- [26] R. Sinkhorn and P. Knopp, “Concerning nonnegative matrices and doubly stochastic matrices”, *Pacific J. Math.* **21** (1967), 343–348.
- [27] S. Sullivant, *Algebraic statistics*, Graduate Studies in Mathematics **194**, Amer. Math. Soc., Providence, RI, 2018.
- [28] B. Sury, “An elementary proof of the Hilbert–Mumford criterion”, *Electron. J. Linear Algebra* **7** (2000), 174–177.
- [29] G. Székelyhidi, *Extremal metrics and K-stability*, Ph.D. thesis, Imperial College London, 2006.

- [30] E. B. Vinberg and V. L. Popov, “Invariant theory”, pp. 123–278 in *Algebraic geometry, IV*, edited by A. N. Parshin and I. R. Shafarevich, Encyclopaedia of Mathematical Sciences **55**, Springer, 1994.

Received 2021-01-29. Revised 2021-07-28. Accepted 2021-09-27.

CARLOS AMÉNDOLA: carlos.amendola@tum.de
Department of Mathematics, Technical University of Munich, Germany

KATHLÉN KOHN: kathlen@kth.se
Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

PHILIPP REICHENBACH: reichenbach@tu-berlin.de
Institut für Mathematik, Technische Universität Berlin, Germany

ANNA SEIGAL: seigal@maths.ox.ac.uk
Mathematical Institute, University of Oxford, United Kingdom
Current address: Department of Mathematics, Harvard University, Cambridge, MA, United States



msp.org/astat

MANAGING EDITORS

Thomas Kahle Otto-von-Guericke-Universität Magdeburg, Germany
Sonja Petrovic Illinois Institute of Technology, United States

ADVISORY BOARD

Mathias Drton Technical University of Munich, Germany
Peter McCullagh University of Chicago, United States
Giorgio Ottaviani University of Florence, Italy
Bernd Sturmfels University of California, Berkeley, and Max Planck Institute, Leipzig
Akimichi Takemura University of Tokyo, Japan

EDITORIAL BOARD

Marta Casanellas Universitat Politècnica de Catalunya, Spain
Alexander Engström Aalto University, Finland
Hisayuki Hara Doshisha University, Japan
Jason Morton Pennsylvania State University, United States
Uwe Nagel University of Kentucky, United States
Fabio Rapallo Università del Piemonte Orientale, Italy
Eva Riccomagno Università degli Studi di Genova, Italy
Yuguo Chen University of Illinois, Urbana-Champaign, United States
Caroline Uhler Massachusetts Institute of Technology, United States
Ruriko Yoshida Naval Postgraduate School, United States
Josephine Yu Georgia Institute of Technology, United States
Piotr Zwiernik Universitat Pompeu Fabra, Barcelona, Spain

PRODUCTION

Silvio Levy (Scientific Editor)
production@msp.org

See inside back cover or msp.org/astat for submission instructions.

Algebraic Statistics (ISSN 2693-3004 electronic, 2693-2997 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840 is published continuously online. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices.

AStat peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY
 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2021 Mathematical Sciences Publishers

Algebraic Statistics

2021

12:2

An algebraic Monte-Carlo algorithm for the partition adjacency matrix realization problem	115
ÉVA CZABARKA, LÁSZLÓ SZÉKELY, ZOLTÁN TOROCZKAI and SHANISE WALKER	
Bimonotone subdivisions of point configurations in the plane	125
ELINA ROBEVA and MELINDA SUN	
A generalization for the expected value of the earth mover's distance	139
WILLIAM Q. ERICKSON	
Likelihood equations and scattering amplitudes	167
BERND STURMFELS and SIMON TELEN	
Toric invariant theory for maximum likelihood estimation in log-linear models	187
CARLOS AMÉNDOLA, KATHLÉN KOHN, PHILIPP REICHENBACH and ANNA SEIGAL	
Discrete max-linear Bayesian networks	213
BENJAMIN HOLLERING and SETH SULLIVANT	