

Likelihood-based inference for exponential-family random graph models via linear programming

Pavel N. Krivitsky¹, Alina R. Kuvelkar² and David R. Hunter²

¹*Department of Statistics, University of New South Wales*
e-mail: p.krivitsky@unsw.edu.au

²*Department of Statistics, Pennsylvania State University*
e-mail: ark336@psu.edu; dhunter@stat.psu.edu

Abstract: The problem of determining whether a given point, or set of points, lies within the convex hull of another set of points in d dimensions arises naturally in the context of certain exponential family models in statistics. This article discusses the general convex hull problem and its application to the particular problem of modelling network data using an exponential-family random graph model (ERGM). While the convex hull question may be solved via a simple linear program, this approach is not well known in the statistical literature. The article also details several substantial improvements to the convex hull-testing algorithm currently implemented in the widely used `ergm` package for network modeling. It provides direct numerical comparisons of two linear programming packages for R that can be called by `ergm` and offers several illustrative examples.

MSC2020 subject classifications: Primary 62-04, 62-08; secondary 62F30.

Keywords and phrases: Convex hull, duality, MCMC, missing data.

Received February 2022.

1. Monte Carlo maximum likelihood estimation for exponential families

Suppose that we observe a complex data structure Y from a sample space \mathcal{Y} , and we believe that the process that produced Y may be captured sufficiently by the d -dimensional vector statistic $g : \mathcal{Y} \mapsto \mathbb{R}^d$. A natural probability model for such an object—one that minimizes the additional assumptions made in the sense of maximizing entropy—is the exponential family class of models [16]. If Y modeled via a discrete or a continuous distribution, the probability mass function or density has the form, parameterized by the *canonical parameter* $\theta \in \mathbb{R}^d$ [1, among others],

$$p_{\theta; \mathcal{Y}, h, g}(Y) = h(Y) \frac{\exp\{\theta^\top g(Y)\}}{\kappa_{\mathcal{Y}, h, g}(\theta)}, \quad Y \in \mathcal{Y}.$$

The form is then a product of a function $h(\cdot)$ of the data alone specifying the distribution under $\theta = 0$, a function $\kappa(\cdot)$ of the parameter alone, defined below,

and an exponential with the interaction between data and parameter. The exact form of $\kappa(\theta)$ depends on the type of distribution:

$$\text{discrete:} \quad \kappa_{\mathcal{Y},h,g}(\theta) = \sum_{Y' \in \mathcal{Y}} h(Y') \exp\{\theta^\top g(Y')\} \quad (1a)$$

$$\text{continuous:} \quad \kappa_{\mathcal{Y},h,g}(\theta) = \int_{\mathcal{Y}} h(Y') \exp\{\theta^\top g(Y')\} dY'. \quad (1b)$$

A full measure-theoretic formulation is also possible for distributions that are neither discrete nor continuous; and a $\theta \in \mathbb{R}^q$ for $q \leq d$ may be mapped through a vector function $\eta : \mathbb{R}^q \mapsto \mathbb{R}^d$. Likelihood-based inference for exponential-family models centers on the log-likelihood function

$$\ell(\theta) = \theta^\top g(Y_{\text{obs}}) - \log \kappa_{\mathcal{Y},h,g}(\theta). \quad (2)$$

Likelihood calculations can be highly computationally challenging when the sum or integral (1) is intractable [e.g., 10].

For the sake of brevity, we will omit “ \mathcal{Y}, h, g ” from the subscript of $p_{\theta;\mathcal{Y},h,g}(Y)$ and $\kappa_{\mathcal{Y},h,g}(\theta)$ for the remainder of this paper, unless they differ from these defaults.

Most commonly, exponential families are used as a tool for deriving inferential properties of families that belong to this class, albeit with a different parametrization. For a common example, the canonical parameter of the Normal(μ, σ^2) distribution in its exponential family form is $\theta = [\mu\sigma^{-2}, \sigma^{-2}]$, which is far less convenient and interpretable.

However, there are domains in which an exponential family model is specified directly through its sufficient statistics, often with the help of the Hammersley–Clifford Theorem [3, among others]. These include the Strauss spatial point processes [25], the Conway–Maxwell–Poisson model for count data [22], and exponential-family random graph models (ERGMs) [13, 6, for original derivations] for networks and other relational data. While our development is motivated by and focuses on ERGMs, it applies to other scenarios involving exponential-family models with intractable normalizing constants because the methods we discuss operate on the sufficient statistic $g(Y)$ rather than on the original data structure.

Consider a graph Y on n vertices, with the vertices being labelled $1, \dots, n$. We will focus on binary undirected graphs with no self-loops and no further constraints—a discrete distribution. Thus, we can express $\mathcal{Y} = 2^{\{\{i,j\} \in \{1,\dots,n\}^2 : i \neq j\}}$, the power set of the set of all distinct unordered pairs of vertex indices.

An exponential family on such a sample space is called an exponential-family random graph model (ERGM). Substantively, elements of $g(Y)$ then represent features of the graph—e.g., the number of edges or other structures, or connections between exogenously defined groups—whose prevalence is hypothesised to influence the relative likelihoods of different graphs.

For example, an edge count statistic would, through its corresponding parameter, control the relative probabilities of sparser and denser graphs, and in

turn the expected density of the graph, conditional on other statistics. Some of the graph statistics, the most familiar being the number of triangle, or “friend of a friend is a friend,” configurations, induce stochastic dependence among the edges. The triangle statistic in particular is problematic for reasons reviewed by a number of authors [e.g., Section 3.1 of 21] and has been largely superseded by statistics proposed by [24] and others.

For our special case, $\kappa(\theta)$ has the form (1a), a summation over all possible graphs. For the population of binary, undirected, vertex-labelled graphs with no self-loops, the cardinality of \mathcal{Y} is $2^{\binom{n}{2}}$, a number exponential in the square of the number of vertices. Thus, even for small networks, this sum is intractable. For example, for $n = 10$, summation is required of $|\mathcal{Y}| \approx 3.5 \times 10^{13}$ elements—too many to compute by “brute force.” Under some choices of $g(\cdot)$, the summation (1a) simplifies, but for most interesting models—those involving complex dependence among relationships in the network represented by the graph—maximization of $\ell(\theta)$ can be an enormous computational challenge.

Various authors have proposed methods for approximately maximizing the log-likelihood function. For instance, [23] introduced a Robbins–Monro algorithm [19] based on the fact that when $\hat{\theta}$ denotes the maximum likelihood estimator (MLE),

$$\mathbb{E}_{\hat{\theta}}[g(Y)] = g(Y_{\text{obs}}).$$

Alternatively, the MCMC MLE idea of [10] was adapted to the ERGM framework by, among others, [14]. This idea is based on the fact that the log-likelihood-ratio

$$\lambda(\theta, \theta_0) \equiv \ell(\theta) - \ell(\theta_0) = (\theta - \theta_0)^\top g(Y_{\text{obs}}) - \log \mathbb{E}_{\theta_0}[\exp\{(\theta - \theta_0)^\top g(Y)\}],$$

suggesting that if we sample r networks Y_1, \dots, Y_r from the approximate distribution $p_{\theta_0}(\cdot)$ via MCMC, we can employ an estimator

$$\hat{\lambda}(\theta, \theta_0) \equiv (\theta - \theta_0)^\top g(Y_{\text{obs}}) - \log \left[\frac{1}{r} \sum_{i=1}^r \exp\{(\theta - \theta_0)^\top g(Y_i)\} \right] \quad (3)$$

as an approximation to $\lambda(\theta, \theta_0)$.

In some cases, the network may be partially unobserved, i.e., Y_{obs} is not a complete network in \mathcal{Y} . Letting $\mathcal{Y}(Y_{\text{obs}})$ denote the set of all networks in \mathcal{Y} that coincide with Y_{obs} wherever Y_{obs} is observed, we may generalize (2) as in [11] by writing

$$\ell(\theta) = \log \sum_{Y' \in \mathcal{Y}(Y_{\text{obs}})} p_\theta(Y') = \log \kappa_{\mathcal{Y}(Y_{\text{obs}})}(\theta) - \log \kappa_{\mathcal{Y}}(\theta). \quad (4)$$

Equation (4) generalizes (2) because $\mathcal{Y}(Y_{\text{obs}})$ consists of the singleton $\{Y_{\text{obs}}\}$ when the network is fully observed.

This likelihood may be approximated using the approach of [7] and [8]: For a known θ_0 , draw samples Y_1, \dots, Y_r and Z_1, \dots, Z_s from \mathcal{Y} and $\mathcal{Y}(Y_{\text{obs}})$, respectively via MCMC, such that the stationary distributions are $p_{\theta_0, \mathcal{Y}}(\cdot)$ and $p_{\theta_0, \mathcal{Y}(Y_{\text{obs}})}(\cdot)$, respectively. The generalization of (3) is then

$$\hat{\lambda}(\theta, \theta_0) \equiv \log \left[\frac{1}{s} \sum_{i=1}^s \exp\{(\theta - \theta_0)^\top g(Z_i)\} \right] - \log \left[\frac{1}{r} \sum_{i=1}^r \exp\{(\theta - \theta_0)^\top g(Y_i)\} \right]. \quad (5)$$

The idea of Markov chain Monte Carlo maximum likelihood estimation, or MCMC-MLE, is to use maximize (5) instead of (4) as a function of θ . One potential problem with this approach is the focus of this article: It is not always the case that (5) has a maximizer, nor even an upper bound. Thus, we address the question of when there exists an MCMC-MLE based on the samples Y_1, \dots, Y_r and Z_1, \dots, Z_s . This does not directly address the question of whether an MLE itself exists, which is an important question in its own right.

To characterize precisely when the approximation of (5) has a maximizer, we first define a term that will be important throughout this article: The *convex hull* of any set of points in d -dimensional Euclidean space is the smallest convex set containing that set, i.e., the intersection of all convex sets containing that set. According to exponential family theory [e.g., Theorem 9.13 of 1], $\hat{\lambda}(\theta, \theta_0)$ in (3) has a maximizer if and only if $g(Y_{\text{obs}})$ is contained in the *interior* of the convex hull of $\{g(Y_1), \dots, g(Y_r)\}$. Indeed, it is straightforward to show that $\hat{\lambda}(\theta, \theta_0)$ in the more general expression (5) does not have a maximum if *any* of $\{g(Z_1), \dots, g(Z_s)\}$ lies outside the convex hull of $\{g(Y_1), \dots, g(Y_r)\}$. We prove this in Section 2.

The question of determining when a point lies within the convex hull of another set of points is thus relevant when using approximations (3) and (5). The remainder of this article shows how to determine whether a given point or set of points lies inside a given convex hull using linear programming. It also describes various improvements to the linear programming algorithm originally proposed and implemented in the `ergm` package by [14]. These improvements are all implemented in the most recently released version of `ergm`, version 4.5.0 [12].

2. Convex hull testing as a linear program

We introduce the terms *target set* and *test set* to refer to the set $T = \{g(Y_1), \dots, g(Y_r)\}$ and the set $S = \{g(Z_1), \dots, g(Z_s)\}$, respectively. To reiterate, the convex hull of any set of d -dimensional points is the smallest convex set containing that set, and the convex hull is always closed in \mathbb{R}^d . We let $\mathcal{C}(T)$ denote the convex hull of a set $T \subset \mathbb{R}^d$. The interior of this convex hull, denoted $\mathcal{C}^\circ(T)$, will play a special role in this article. If the points in T satisfy a linear constraint, $\mathcal{C}^\circ(T)$ is empty since in this case the convex hull lies in an affine subspace of dimension smaller than d . We assume here that $\mathcal{C}^\circ(T)$ is nonempty.

As mentioned near the end of Section 1, the approximation in (5) fails to admit a maximizer whenever $S \not\subset \mathcal{C}(T)$, which may be proved directly:

Proposition 1. If $S \not\subset \mathcal{C}(T)$, then $\sup_{\theta} \hat{\lambda}(\theta, \theta_0) = \infty$.

Proof. Suppose there exists an element of S , say $g(Z_1)$, in the open set $\mathbb{R}^d \setminus \mathcal{C}(T)$. By the convexity of $\mathcal{C}(T)$, this means there exists a hyperplane \mathcal{H} —that is, an affine $(d - 1)$ -dimensional subspace—containing $g(Z_1)$ and such that $\mathcal{H} \cap \mathcal{C}(T) = \emptyset$. In other words, there exist a scalar z_0 and a d -vector z such that $z_0 + z^\top g(Z_1) = 0$ and $z_0 + z^\top g(Y_i) < 0$ for $i = 1, \dots, r$.

If we now let $\theta = \theta_0 + \alpha z$ for a real number α , (5) gives

$$\begin{aligned} \exp\{\hat{\lambda}(\theta, \theta_0)\} &= \frac{r}{s} \times \frac{\exp\{\alpha z^\top g(Z_1)\} + \sum_{i=2}^s \exp\{\alpha z^\top g(Z_i)\}}{\sum_{i=1}^r \exp\{\alpha z^\top g(Y_i)\}} \\ &\geq \frac{r}{s} \times \frac{1}{\sum_{i=1}^r \exp[\alpha z^\top \{g(Y_i) - g(Z_1)\}]}. \end{aligned}$$

Since $z_0 + z^\top g(Z_1) = 0$ and $z_0 + z^\top g(Y_i) < 0$ imply that $z^\top \{g(Y_i) - g(Z_1)\} < 0$ for $i = 1, \dots, r$, every term in the denominator goes to zero as $\alpha \rightarrow \infty$. \square

However, $S \subset \mathcal{C}(T)$ is merely a necessary, not sufficient, condition for (5) to have a unique maximizer. Another necessary condition is that $\frac{\partial^2 \hat{\lambda}(\theta, \theta_0)}{\partial \theta^2}$ must be negative-definite for all θ and θ_0 , but it is not sufficient either, since a direction of recession [9, for example] may exist. We are thus not aware of a necessary and sufficient condition when S contains more than one point. As we see below, for our purposes, the conditions that $S \subset \mathcal{C}^\circ(T)$ and that $\text{Var}(T) - \text{Var}(S)$ is positive-definite suffice.

Below in LP (6)—where ‘LP’ stands for ‘linear program’—we describe how linear programming can provide a method of testing definitively whether $S \subset \mathcal{C}(T)$ but *not* whether $S \subset \mathcal{C}^\circ(T)$. As a practical matter, we can apply this method to a slightly perturbed version of S in which each test point is expanded away from the centroid of T by a small amount; if the perturbed S is contained in $\mathcal{C}(T)$, then $S \subset \mathcal{C}^\circ(T)$. This approach has the added benefit of ensuring not only that $S \subset \mathcal{C}^\circ(T)$ but that each point in S is bounded away from the boundary of $\mathcal{C}^\circ(T)$ by a small amount, which can prevent some computational challenges in using the approximation in (5). This ‘slightly perturbed’ algorithm had been implemented in previous versions of the `ergm` package. On the other hand, the recently released `ergm` version 4.5.0 uses an improved linear program, explained in Section 3, that transforms the linear program so that it tests directly whether $S \subset \mathcal{C}^\circ(T)$.

In the particular case in which Y_{obs} is a full network, i.e., there are no missing data, S is the singleton $\{g(Y_{\text{obs}})\}$. In this case, checking whether $S \subset \mathcal{C}^\circ(T)$ is the same as checking whether S and T are separable by a hyperplane, a problem that has a lengthy history in the linear programming literature [e.g., 5]. In general, the linear separability check does not tell us whether $S \subset \mathcal{C}^\circ(T)$, and we discuss this more general case in Section 5.

We first frame the check of whether $p \in \mathcal{C}(T)$, for an arbitrary column vector $p \in \mathbb{R}^d$, as a linear program. Let M be the $(r \times d)$ -dimensional matrix whose rows M_1, \dots, M_r are the points in the target set T ; furthermore, let $\mathcal{C}(M)$ denote $\mathcal{C}(T)$. Because $\mathcal{C}(M)$ is convex in \mathbb{R}^d , for any $p \notin \mathcal{C}(M)$, we may find an affine $(d - 1)$ -dimensional subspace, which we call a *separating hyperplane*,

that separates the points in T from the point p in the sense that p lies in one (open) half-space defined by the hyperplane and the points in T —i.e., the rows of M —all lie in the other (closed) half-space. Mathematically, this separating hyperplane is determined by the affine subspace $\{x \in \mathbb{R}^d : z_0 + x^\top z = 0\}$ for some scalar z_0 and d -vector z . Thus, $z \in \mathbb{R}^d$ and $z_0 \in \mathbb{R}$ determine a separating hyperplane whenever $z_0 + p^\top z < 0$ while $z_0 + M_i z \geq 0$ for each row M_i of M .

In other words, a hyperplane separating p from $\mathcal{C}(M)$ exists whenever the minimum value of $z_0 + p^\top z$ is strictly negative, where the minimum is taken over all $z_0 \in \mathbb{R}$ and $z = (z_1, \dots, z_d)^\top \in \mathbb{R}^d$ such that $M_i z \geq -z_0$ for $i = 1, \dots, n$. Notationally, we will use inequality symbols to compare vectors componentwise, so for example instead of “ $M_i z \geq -z_0$ for $i = 1, \dots, n$ ”, we may write “ $Mz \geq -z_0 \mathbf{1}$ ”. The existence of a separating hyperplane can thus be determined using the following linear program:

$$\begin{aligned} \text{minimize:} \quad & z_0 + p^\top z \\ \text{subject to:} \quad & Mz \geq -z_0 \mathbf{1} \\ & z \geq -1 \\ & z \leq 1. \end{aligned} \tag{6}$$

If the minimum value of the objective function $z_0 + p^\top z$ is exactly 0—it can never be strictly positive because $(z_0, z) = 0_{d+1}$ determines a feasible point—then p is in $\mathcal{C}(M)$. If the minimum value is strictly negative, then a separating hyperplane not containing p exists and thus $p \notin \mathcal{C}(M)$. The bounds $-1 \leq z \leq 1$, which we call *box constraints*, ensure that the linear program has a finite minimizer when $p \notin \mathcal{C}(M)$: If z is unconstrained and there exists a feasible z_0, z with $z_0 + p^\top z < 0$, then $(cz_0) + (cp)^\top z < z_0 + p^\top z$ for any $c > 1$ and no finite minimizer exists.

This case $S = \{g(Y_{\text{obs}})\}$ is considered by [14], who propose an approximation method that checks $S \subset \mathcal{C}^\circ(T)$ and then replaces $g(Y_{\text{obs}})$ by some other point contained in $\mathcal{C}^\circ(T)$ whenever $S \not\subset \mathcal{C}^\circ(T)$. In particular, the method defines a “pseudo-observation”, $\hat{\xi}$, as the convex combination $\gamma g(Y_{\text{obs}}) + (1 - \gamma)\bar{t}$, where \bar{t} is the centroid of T and $\gamma \in (0, 1]$. To implement this idea, earlier versions of `ergm` relied on a simplistic function called `is.inCH()` that simply returned a TRUE or FALSE value, depending on whether $g(Y_{\text{obs}}) \in \mathcal{C}(T)$, after the ‘slight perturbation’ described earlier. This trial-and-error approach used either a grid search or a bisection search to choose the largest $\gamma \in (0, 1]$ such that $\hat{\xi} \in \mathcal{C}^\circ(T)$. This implementation required multiple `is.inCH()` function calls, resulting in considerable inefficiency relative to the methods described in Section 3. In `ergm` version 4.5.0, `is.inCH()` is deprecated in favor of a new function called `shrink_into_CH()` that need only be called once. Online documentation for this function is obtained in R by typing `help(shrink_into_CH)`.

3. Reformulating the single-test-point algorithm

The trial-and-error algorithm of [14], which applies only to the case where the test set S consists of the single point p , never exploits the fact that the separating

hyperplane, if it exists, is the minimizer of a function. We propose to improve this algorithm by reformulating the linear program and then using the minimizers found. First, we establish a few basic facts that help us simplify the problem.

Since translating p and all points in T by the same constant d -vector does not change whether $p \in \mathcal{C}(M)$, we assume throughout this section that all points are translated so that the origin is a point in $\mathcal{C}^\circ(M)$ that we refer to as the “centroid” of M . In practice, we often take the centroid of M to be the sample mean of the points in the target set. Yet none of the results of this section rely on any particular definition of “centroid”; here, we assume only that the centroid $0 \in \mathbb{R}^d$ lies in the interior of the convex hull of the target set.

Since p may be assumed not to coincide with the centroid, else the question of whether $p \in \mathcal{C}^\circ(M)$ is answered immediately, at least one of the coordinates of p is nonzero. We assume without loss of generality that $p_1 \neq 0$; otherwise, we may simply permute the coordinates of all the points without changing whether $p \in \mathcal{C}(M)$. Below, we define a simple invertible linear transformation that maps p to the unit vector $e_1 = (1, 0, \dots, 0)^\top$. Applying this transformation allows us to simplify the original linear program and thereby clarifies our exposition while facilitating establishing results; then, once we have proved our results, we will apply the inverse transformation to restore the original coordinates, achieving a simplification of the linear program in the process.

One important fact about any invertible linear transformation, when applied to p and M , is that it does not change whether or not $p \in \mathcal{C}(M)$. For if $p \in \mathcal{C}(M)$, there exists $a \in \mathbb{R}^r$ such that $p^\top = a^\top M$ and $\sum_{i=1}^r a_i = 1$. Since $(Rp)^\top = a^\top MR^\top$, we conclude that $Rp \in \mathcal{C}(MR^\top)$. The converse follows via the same argument because R^{-1} is full rank whenever R is.

Consider the full-rank linear transformation defined by

$$R = \begin{bmatrix} \frac{1}{p_1} & 0 & \cdots & 0 \\ \frac{-p_2}{p_1} & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ \frac{-p_d}{p_1} & 0 & \cdots & 1 \end{bmatrix} = \begin{bmatrix} p_1 & 0 & \cdots & 0 \\ p_2 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ p_d & 0 & \cdots & 1 \end{bmatrix}^{-1} \quad (7)$$

that maps p to the standard basis vector $e_1 \in \mathbb{R}^d$. After applying this linear transformation, we want to know whether a hyperplane exists that separates the point e_1 from the rows of MR^\top .

Such a hyperplane, if it exists, may be written as $\{x \in \mathbb{R}^d : z_0 + x^\top z = 0\}$ for some $z_0 \in \mathbb{R}$ and $z \in \mathbb{R}^d$. Since the origin is in $\mathcal{C}^\circ(MR^\top)$, no separating hyperplane may pass through the origin. Therefore, we may limit our search to those cases where $z_0 \neq 0$, which means that we may divide by z_0 the equation defining the hyperplane. Stated differently, we may take $z_0 = 1$ and rewrite our hyperplane as $\{x \in \mathbb{R}^d : 1 + x^\top z = 0\}$ without loss of generality.

Summarizing the arguments above, we now seek a d -vector z satisfying $1 + e_1^\top z < 0$ and $1 + (MR^\top)_i z \geq 0$ for $i = 1, \dots, r$. Furthermore, fixing $z_0 = 1$ means that no box constraints are necessary in the reformulation of the linear program designed to search for a separating hyperplane. Therefore, our new

linear program, after transformation by R , becomes

$$\begin{aligned} & \text{minimize:} && z_1 \\ & \text{subject to:} && [MR^\top]z \geq -1. \end{aligned} \tag{8}$$

The zero vector is a feasible point in both LP (6) and LP (8). However, unlike LP (6), zero cannot be the optimum of LP (8) because 0 does not lie on the boundary of the set of feasible points of LP (8). In other words, a solution z^* of LP (8), if it exists, is not zero so it defines a hyperplane $\mathcal{H}^* = \{x \in \mathbb{R}^d : x^\top z^* = -1\}$. The constraints of LP (8) imply that \mathcal{H}^* is a supporting hyperplane of $\mathcal{C}(MR^\top)$, whether or not it separates $\mathcal{C}^\circ(MR^\top)$ from e_1 . We now demonstrate the existence of a solution to LP (8) and show how to determine whether $p \in \mathcal{C}(M)$:

Proposition 2. Let M be a full-rank matrix so that $\mathcal{C}^\circ(M)$ is nonempty, with $0 \in \mathcal{C}^\circ(M)$, let $p \in \mathbb{R}^d$ satisfy $p_1 \neq 0$, and let R be defined as in (7). Then a (nonzero) minimizer z^* of LP (8) exists, and the ray with endpoint at the origin and passing through e_1 intersects the boundary of $\mathcal{C}(MR^\top)$ at $-(1/z_1^*)e_1$. In particular, $e_1 \in \mathcal{C}(MR^\top)$ if and only if $-1 \leq z_1^*$.

Proof. Since 0 is contained in the open set $\mathcal{C}^\circ(MR^\top)$, there exists $\delta > 0$ such that $\delta e_1 \in \mathcal{C}^\circ(MR^\top)$. Thus, any feasible point z of LP (8) must satisfy $1 + \delta e_1^\top z = 1 + \delta z_1 \geq 0$, which implies $z_1 \geq -1/\delta$. Since the objective function has a lower bound, it must attain its minimum on the closed set of feasible points, so a solution z^* to LP (8) exists.

Now define the point $a = -(1/z_1^*)e_1$. If a lies in the open set $\mathcal{C}^\circ(MR^\top)$, then there exists $\epsilon > 0$ such that $(1 + \epsilon)a$ also lies in $\mathcal{C}^\circ(MR^\top)$; but this is impossible since $(1 + \epsilon)a^\top z^* = -(1 + \epsilon) < -1$ and thus $(1 + \epsilon)a$ violates the constraint that must be satisfied by every point in $\mathcal{C}^\circ(MR^\top)$ due to convexity. On the other hand, a must lie in $\mathcal{C}(MR^\top)$ since otherwise there exists $\epsilon > 0$ such that $(1 - \epsilon)a \notin \mathcal{C}(MR^\top)$, which in turn means that there exists a hyperplane separating $\mathcal{C}(MR^\top)$ from $(1 - \epsilon)a$. But this means there exists a supporting hyperplane of $\mathcal{C}(MR^\top)$ that intersects the positive x_1 -axis at a point between the origin and $(1 - \epsilon)a$, contradicting the fact that $-(1/z_1^*)$ is the smallest coordinate among all such points of intersection. We conclude that $a = -(1/z_1^*)e_1$ must lie on the boundary of $\mathcal{C}(MR^\top)$, so $e_1 \in \mathcal{C}(MR^\top)$ if and only if $-1 \leq z_1^*$. \square

Applying the full-rank transformation R^{-1} , in order to transform back to the original coordinates, establishes the following corollary:

Corollary 1. Let z^* denote a minimizer of LP (8). Then the ray with endpoint at the origin and passing through p intersects the boundary of $\mathcal{C}^\circ(M)$ at $-(1/z_1^*)p$. In particular, p is in $\mathcal{C}^\circ(M)$ if and only if $z_1^* > -1$.

This result gives a way to determine when p lies in $\mathcal{C}^\circ(M)$, not merely $\mathcal{C}(M)$, which is an improvement over LP (6). It also suggests that we may reconsider LP (8) after transforming by R^{-1} back to the original coordinates. Indeed, since z in LP (8) may take any value in \mathbb{R}^d and R has full rank, there is no loss of

generality in replacing z by $(R^\top)^{-1}z$, which allows us to rewrite LP (8) as

$$\begin{aligned} \text{minimize:} \quad & p^\top z \\ \text{subject to:} \quad & Mz \geq -1. \end{aligned} \tag{9}$$

Taking z^* and z^{**} as minimizers of LP (8) and LP (9), respectively, we must have $z_1^* = p^\top z^{**}$ since the minimum objective function value must be the same for the two equivalent linear programs. We conclude by Corollary 1 that p is in $\mathcal{C}^\circ(M)$ if and only if $p^\top z^{**} > -1$ and, furthermore, the point $-p/(p^\top z^{**})$ is the point in $\mathcal{C}(M)$ closest to the origin in the direction of p .

According to standard linear program theory [28, Section 5.8], LP (9) is precisely what is known as the *dual* of the following linear program:

$$\begin{aligned} \text{maximize:} \quad & -1^\top y \\ \text{subject to:} \quad & M^\top y = p \\ & y \geq 0. \end{aligned} \tag{10}$$

LP (10) may be derived from first principles, since $\mathcal{C}(M)$ consists precisely of those points that may be expressed as a convex combination of the rows of M , i.e., $p \in \mathcal{C}(M)$ if and only if $p = M^\top y$ for some $y \geq 0$ such that $1^\top y = 1$. In other words, $p \in \mathcal{C}(M)$ if and only if there exists a solution y^* of LP (10) satisfying $-1^\top y^* \geq -1$. (The maximum value of $-1^\top y^*$ could be strictly greater than -1 because zero is in $\mathcal{C}^\circ(M)$.)

Furthermore, if $z \in \mathbb{R}^d$ satisfies the r constraints in LP (9) and $y \in \mathbb{R}^r$ is nonnegative, we may multiply the i th constraint by y_i and then sum over i to obtain $y^\top Mz \geq -y^\top 1$. If the constraints in LP (10) are completely satisfied, we may replace $y^\top M$ by p^\top . Minimizing with respect to z and maximizing with respect to y yields

$$\min_z p^\top z \geq \max_y -1^\top y. \tag{11}$$

Inequality (11) is sometimes called the *weak duality theorem*. The *strong duality theorem*, which we do not prove here, says that equality holds in (11) [28, Section 5.4]. In our convex hull problem, strong duality implies among other things that p is on the boundary of $\mathcal{C}(M)$ if and only if the maximum objective function value in LP (10) equals -1 .

4. Applications and benchmarks

To illustrate the results of Section 3, we first consider problems in 2 dimensions since they are easy to visualize. Such low-dimensional examples are also helpful since they sometimes lend themselves to closed-form solutions of linear programs.

For the following benchmarks and examples, the following relevant R [20] package versions were used: `ergm` 4.5.0, `Rglpk` 0.6.5, and `lpSolveAPI` 5.5.2.0.17.9.

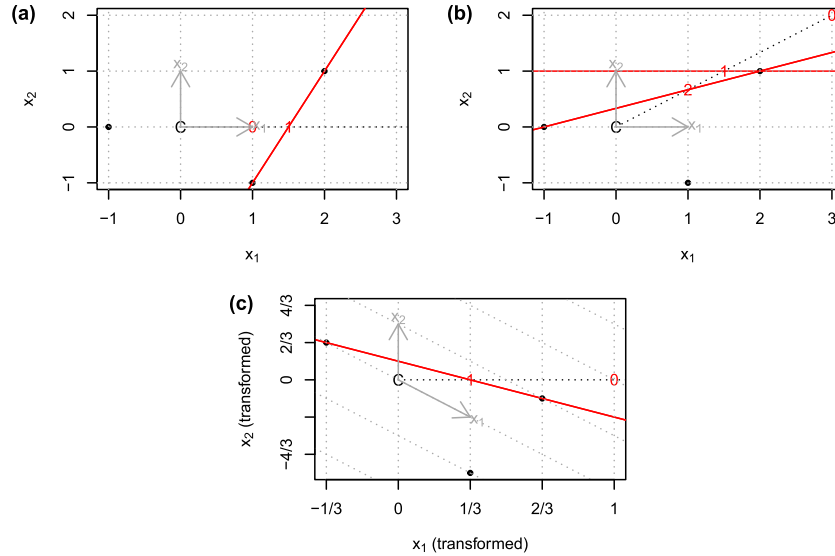


FIG 1. A toy example with three points. The centroid is marked with C; arrows indicate standard basis vectors in each dimension on the original scale, target set points are dots, supporting lines are colored red, the initial test point is 0, and its successive iterations are labeled with integers. Clockwise from upper left, we see (a) LP (8) with test point on the x_1 -axis, where a simple closed-form solution exists; (b) LP (6) on the original scale with box constraints, where a suboptimal first solution results; and (c) LP (8) on transformed scale and without box constraints, where the intersection point is found immediately.

4.1. A three-point example

Let us consider a particularly simple example in which the target set consists of the 2-vectors $(-1, 0)^T$, $(a, 1)^T$, and $(b, -1)^T$ for $a > 0$ and $b > 0$. This choice guarantees that the convex hull is a triangle containing the origin, so we shall consider the origin to be the centroid in this example, regardless of the actual value of the mean of the three points.

When the test point p equals $(1, 0)^T$, we are in a situation that could arise after transforming by R . This leads to LP (8), which may be solved in closed form because in this case the constraints become

$$\max \left\{ -\frac{1}{a} - \frac{z_2}{a}, -\frac{1}{b} + \frac{z_2}{b} \right\} \leq z_1 \leq 1.$$

Since the lines $x_2 = -1/a - x_1/a$ and $x_2 = -1/b + x_1/b$ have one positive slope and one negative slope, we minimize the maximum at the point of intersection, i.e., when $-1/a - z_2/a = -1/b + z_2/b$, which implies $z_2^* = (a - b)/(a + b)$, which in turn implies $z_1^* = -2/(a + b)$. By simple examination, we know that $(1, 0)^T$ is interior to the convex hull exactly when the line through $(a, 1)^T$ and $(b, -1)^T$ intersects the x_1 -axis at a value larger than 1. Thus, $p \in \mathcal{C}^\circ(M)$ if and only if


```

return(ans$optimum)
}

```

A test case simulates $n = 100,000$ points uniformly distributed in the $d = 20$ -dimensional unit cube, then verifies that the point $p = (1, \dots, 1)^\top$ is not inside the convex hull by finding the smallest $\gamma > 0$ such that γp is on the boundary:

```

set.seed(123)
M <- matrix(runif(1e5 * 20), ncol=20) ; Mbar <- colMeans(M)
M <- sweep(M, 2, Mbar) # Translate M so sample mean is at origin
p <- rep.int(1, 20) - Mbar
lpstime <- system.time(lpsstep<--1/LPmod1(M, p))[1]
glpktime <- system.time(glpkstep<--1/LPmod2(M, p))[1]
stopifnot(isTRUE(all.equal(lpsstep, glpkstep)))

```

Both implementations produce the same result: $\gamma \approx 0.4801$, but `Rglpk` takes 5 seconds, whereas `lpSolveAPI` takes 24. However, `Rglpk` requires GLPK to be installed separately on some platforms; for this reason, version 4.5.0 of `ergm` first tests to see whether `Rglpk` is installed and, if so, uses that package; otherwise, it uses `lpSolveAPI`.

5. Testing multiple points

As explained at the beginning of Section 2, the approximated difference of log-likelihoods in (5) requires that every $g(Z_i) \in S$ be contained in $\mathcal{C}^\circ(T)$ in order for the approximation to have a maximizer. We might therefore consider a strategy of checking, prior to using (5), whether $g(Z_i) \in \mathcal{C}^\circ(T)$ for all $i = 1, \dots, s$ using the single-test-point methods discussed earlier. This strategy has two potential drawbacks: First, the computational burden might be quite high if any of the dimension d , the number of test points s , or the number of target points r is large. Second, we must decide what to do in the case where one or more of the points in S is found to be outside $\mathcal{C}^\circ(T)$. This section addresses each of these questions and then presents an illustrative example using a network dataset with missing edges.

5.1. Reducing computational burden

There are evidently many possible ways to approach the question of how to efficiently decide which test points, if any, lie outside $\mathcal{C}^\circ(T)$. Even if we only consider ideas for reducing the size of the set T in such a way as to have little or no influence on the answer, we might attempt to search for and eliminate target set points either that lie entirely inside $\mathcal{C}^\circ(T)$, since eliminating such points from T does not change $\mathcal{C}^\circ(T)$ at all, or that lie close to other points in T , since eliminating such points does not change $\mathcal{C}^\circ(T)$ very much. Here, we merely suggest a simplistic version of the first approach, as a thorough exploration of methods for reducing the size of T without altering $\mathcal{C}^\circ(T)$ is well beyond the scope of this article.

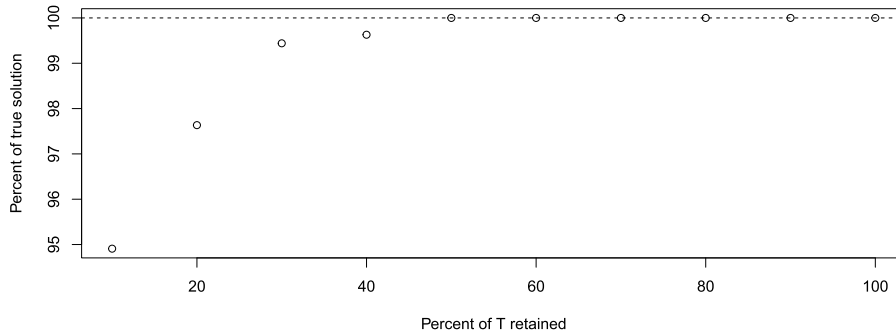


FIG 2. In this example where $d = 20$ and $n = 100,000$, half of the target set can be eliminated using a simplistic Mahalanobis distance-based algorithm without appreciably degrading the quality of the convex hull scaling factor.

Whether a point in T is interior to $\mathcal{C}^\circ(T)$ or whether it lies on the boundary of $\mathcal{C}^\circ(T)$ has to do with that point's so-called *data depth*, a concept often attributed (in the two-dimensional case) to [27]. Because the deepest points are the ones that can be eliminated without changing $\mathcal{C}^\circ(T)$, the sizable literature on data depth [see, e.g., the discussion by 29] may be relevant to our problem. For the $d = 2$ -dimensional case, several authors have developed efficient methods for identifying exactly the points lying on the convex hull boundary; this is a particular case of the so-called *convex layers* problem [4].

Among the most simplistic ideas is to use Mahalanobis distance from the centroid, defined for any point $x \in T$ as

$$d(x) = \sqrt{x^\top \hat{\Sigma}^{-1} x},$$

where $\hat{\Sigma}$ is the sample covariance matrix, as a measure of data depth. For the $|T| = 100,000$ -point example of Section 4.2, we also sample $|S| = 5$ corners of the unit cube in \mathbb{R}^{20} and try eliminating the fraction $f = i/10$ for $i = 1, \dots, 10$ of the deepest points as measured by Mahalanobis distance:

```
p <- sweep(matrix(rbinom(5 * 20, 1, 0.5), ncol = 20), 2, Mbar)
d <- rowSums((M %*% solve(cov(M))) * M) # Find squared Mahalanobis distances
M <- M[order(d, decreasing=TRUE),]
b <- rep(Inf, 10) # Vector for storing LP minima
for (i in 10:1) {
  for (j in 1:nrow(p)) {
    b[i] <- min(b[i], LPmod2(M[1:(i*nrow(M)/10),], p[j,]))
  }
}
```

In this 20-dimensional problem, we see that we can effectively disregard half the points in T , using a simple Mahalanobis ordering, without degrading the solution appreciably. Since computing effort scales linearly with the number of points in T , this seems like a useful tool. We nonetheless recommend caution, as our experiments indicate that the fraction of points that may be discarded

by their Mahalanobis distances varies considerably for different choices of d . For example, we find for $d = 50$ that each of the 100,000 points sampled uniformly randomly in the unit cube lies on the boundary of the convex hull.

5.2. Rescaling observed statistics

Recalling the case of a completely observed network, where S consists of a single point $g(Y_{\text{obs}})$, the strategy used by [14] is to find a point, say $\hat{\xi}$, between 0 and $g(Y_{\text{obs}})$ that lies inside $\mathcal{C}^\circ(T)$ yet close to the boundary. By treating $\hat{\xi}$ as though it is $g(Y_{\text{obs}})$ in constructing approximation (3), the approximation will have a well-defined maximizer, and this maximizer is then used to generate a new target set in the next iteration of the algorithm.

In analogous fashion, when S consists of multiple points and some of them lie outside $\mathcal{C}^\circ(T)$, we propose to shrink each of these points toward 0 using the same scaling factor, say γ , where $\gamma \in (0, 1]$ is chosen so that every element of S lies within $\mathcal{C}^\circ(T)$ after the scaling is applied. Since LP (9) yields the optimal step length for each point p , we obtain this step length by iterating through the points in S and selecting the least of the resulting step lengths. This transformation also has the effect of shifting the sample mean of the points in S to a point somewhere between 0 and the sample mean of the untransformed points in S .

5.3. Example

This section illustrates the idea of Section 5.2 using a network representing a school friendship network based on a school community in the rural western United States. The synthetic `faux.mesa.high` network in the `ergm` package includes 205 students in grades 7 through 12. We may create a copy of this network and then randomly select 10% of the edge observations (whether the edge is present or absent) to be categorized as missing:

```
library(ergm)
data(faux.mesa.high)
fmh <- faux.mesa.high
m <- as.matrix(fmh)
e1 <- cbind(row(m)[lower.tri(m)], col(m)[lower.tri(m)])
set.seed(123)
s <- sample(nrow(e1), round(nrow(e1) * 0.1))
fmh[e1[s,]] <- NA
```

If we now define an ERGM using a few statistics related to those originally used to create the `faux.mesa.high` network—details are found via `help(faux.mesa.high)`—we begin with an estimator that can be derived using straightforward logistic regression. We denote this maximum pseudo-likelihood estimator or MPLE, details of which may be found in Section 5.2 of [15], by `theta0` or θ_0 .

```
fmhFormula <- fmh ~ edges + nodematch("Grade") +
  gwesp(log(3/2), fixed=TRUE)
theta0 <- coef(ergm(fmhFormula, estimate="MPL"))
```

To construct approximation (5), we need two samples of random networks, Y_1, \dots, Y_r and Z_1, \dots, Z_s from \mathcal{Y} and $\mathcal{Y}(Y_{\text{obs}})$, respectively. For the latter, we employ the `constraints` capability of the `ergm` package.

```
gY <- simulate(fmhFormula, coef = theta0, nsim = 500, output = "stats",
  control = snctrl(MCMC.interval = 1e4))
gZ <- simulate(fmhFormula, coef = theta0, nsim = 100, output = "stats",
  constraints = ~observed,
  control = snctrl(MCMC.interval = 1e4))
```

We can use the code developed earlier in this article to show that the $g(Z_j)$ statistics are not interior to the convex hull of the $g(Y_i)$ statistics:

```
centroid <- colMeans(gY)
gY <- sweep(gY, 2, centroid) # Translate the g(Y) statistics
gZ <- sweep(gZ, 2, centroid) # Translate the g(Y) statistics
scale <- Inf
for (j in 1:nrow(gZ)){
  scale <- min(scale, -1/LPmod2(gY, gZ[j,]))
}
```

The code above finds that 0.733 is the largest scaling factor that, when multiplied by each $g(Z_j)$ vector, ensures that the result is on or inside the boundary of the convex hull of the target points $g(Y_1), \dots, g(Y_{100})$. Since this scaling factor is less than 1, at least one element in the test set lies outside $\mathcal{C}^\circ(T)$ and so the approximation in (5) has no maximizer. Figure 3 depicts the target points, the test points, and the test points after scaling by 0.733.

With our samples in place, we may now construct approximation (5). Here, we negate the function so that our objective is to minimize it:

```
l <- function(ThetaMinusTheta0, gY, gZ, scale) {
  -log(mean(exp(scale * gZ %*% ThetaMinusTheta0))) +
  log(mean(exp(gY %*% ThetaMinusTheta0))) }
}
```

Finally, we employ the `optim` function in R [20] to minimize the objective function. We use a scale value of 90% of the value that places one of the test values exactly on the boundary of the convex hull so that all of the scaled test points are interior to the convex hull. We may repeat the whole process iteratively until the entire set of test points is well within the convex hull boundary:

```
multipliers <- scale; NewTheta <- theta0
while (scale < 1.11) { # We want 0.9 * scale > 1 when finished
  NewTheta <- NewTheta + optim(0 * NewTheta, l, gY=gY, gZ=gZ,
    scale=0.9 * scale)$par
  theta0 <- rbind(theta0, NewTheta)
  gY <- simulate(fmhFormula, coef=NewTheta, nsim=500, output="stats",
    control=snctrl(MCMC.interval = 1e4))
  gZ <- simulate(fmhFormula, coef=NewTheta, nsim=100, output="stats",
```




FIG 3. Pairwise scatterplots of three-dimensional network statistics generated using the MPLE, with statistics on the full sample space of networks as black dots and those on the sample space constrained to coincide with the observed network as orange dots. The constrained points are rescaled as blue points so that none lies outside the convex hull of the black points. In each plot, there are 500 black, 100 orange, and 100 blue points.

```

constraints=~observed,
control=snctrl(MCMC.interval = 1e4))
centroid <- colMeans(gY)
gY <- sweep(gY, 2, centroid) # Translate the g(Y) statistics
gZ <- sweep(gZ, 2, centroid) # Translate the g(Y) statistics
scale <- Inf
for (j in 1:nrow(gZ)){
  scale <- min(scale, -1/LPmod2(gY, gZ[j,]))
}
multipliers <- c(multipliers, scale)
}

```

Table 1 gives successive values of θ_0 that are determined as maximizers of (5) after the test set points $g(Z_1), \dots, g(Z_s)$ are rescaled to lie within the convex hull of $g(Y_1), \dots, g(Y_r)$. The maximum pseudo-likelihood estimate (MPLE) in the first row of the table is obtained using logistic regression and is often used as an initial approximation to the maximum likelihood estimator when employing MCMC MLE [15]. However, the MPLE fails to take the missing network observations into account and, as demonstrated by [14], it is not the case that the MPLE generates sample network statistics near the observed statistics.

Figure 4 shows that the final value of θ_0 produces samples such that the whole test set lies on the interior of the target set, which allows (5) to be maximized to

TABLE 1
 Values of θ_0 and their associated test point multipliers, starting with the maximum pseudo-likelihood estimator.

Iteration	edges	nodematch.Grade	gwesp.fixed.0.405	Multiplier
0	-6.302	2.264	1.249	0.733
1	-6.236	2.098	1.276	1.094
2	-6.223	1.981	1.302	2.043

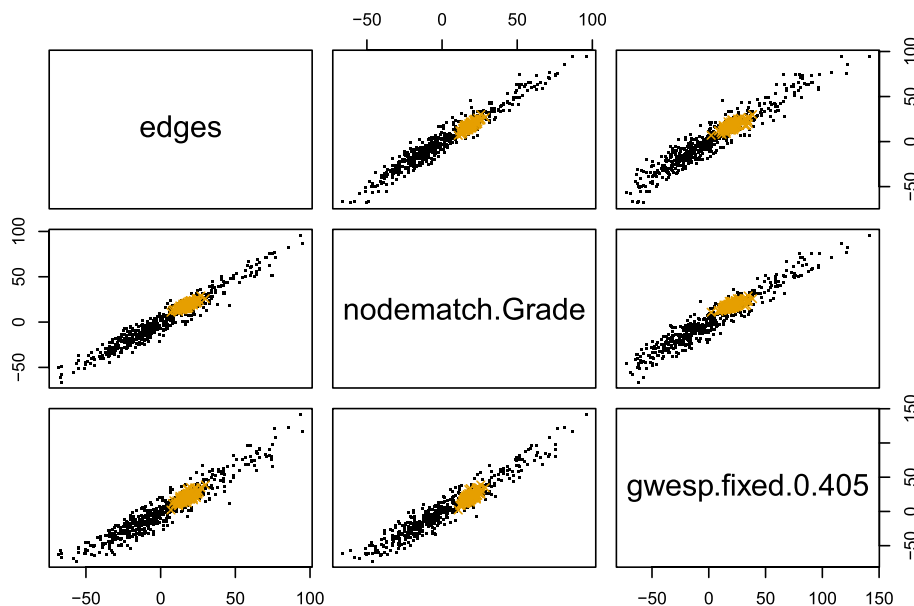


FIG 4. At the final iteration, the test set (orange) is entirely within the convex hull of the target set (black).

produce an approximate MLE. As recommended by [14], we might use moderate-sized samples to obtain a viable θ_0 value using the idea here, then use much larger samples once θ_0 has been found in order to improve the accuracy of Approximation (5).

6. Discussion

This article discusses the problem of determining whether a given point, or set of points, lies within the convex hull of another set of points in d dimensions. While this problem, along with its solution via linear programming, is known, we are not aware of any work that discusses it in the context of a statistical problem such as one discussed here, namely, the maximization of an approximation to the loglikelihood function for an intractable exponential-family model.

Here, we provide multiple improvements on the simplistic implementation of the linear programming solution to the yes-or-no question involving a single

test point that exists in the `ergm` package [15] as a means of implementing the idea of [14]: First, we eliminate the need for the “box constraints” of `ergm` and show how the dual linear program may be derived from first principles. Second, we render the “trial-and-error” approach obsolete by showing how to find the exact point of intersection between the convex hull boundary and the ray originating at the origin and passing through the test point. Third, we test the `lpSolveAPI` package that is currently used by `ergm` against the `Rglpk` package, finding that the latter appears to be far more efficient at solving the particular linear programs we encounter. Fourth, we discuss the statistical case of missing network observations, in which the test set may consist of multiple points, establish an important necessary condition, and suggest a method for handling this case.

In addition, we point out several ways in which this work might be extended, particularly in the case of multiple test points. For one, the question of how to streamline computations is wide open, particularly since it is not necessary to find the exact maximum scaling factor that maps each test point into the convex hull. For the purposes of approximating a maximum likelihood estimator, we seek only an upper bound on the acceptable scaling factors; indeed, in practice we want to scale all test points so that they are inside the boundary. This means that it might be possible to eliminate from the target set any points that are sufficiently close to another target set point and that doing so would not change the needed scaling factor too much.

We might also consider how to optimize the size of the sample chosen for the test set in the first place. For instance, if the scaling factor needed is considerably smaller than one, there might be an advantage in sampling just a handful of points, possibly just a single point in order to move the initial value of θ_0 closer to the true MLE, when a larger sample of test points could be drawn. Another question is whether it might be possible, if one finds $s \in S$ with $s \notin \mathcal{C}^\circ(T)$, to add more sampled points to T , particularly near s , in order to diagnose whether the lack of an MCMC-MLE might be correctable or whether an MLE does not exist at all.

One thing that is clear is that the extensions described here would be much more difficult, if not impossible, to consider without the improvements to the `ergm` package’s convex-hull testing procedure outlined in this article.

Acknowledgments

The authors would like to thank the reviewers, the associate editor, and the editor, who all provided helpful feedback during the preparation of this paper.

Funding

The first author was supported in part by US Army Research Office Award W911NF-21-1-0335 (79034-NS).

References

- [1] BARNDORFF-NIELSEN, O. E. (1978). *Information and Exponential Families in Statistical Theory*. Wiley, New York. [MR0489333](#)
- [2] BERKELAAR, M., DIRKS, J. et al. (2021). `lp_solve` reference guide.
- [3] BESAG, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *Journal of the Royal Statistical Society, Series B* **36** 192–236. [MR0373208](#)
- [4] CHAZELLE, B. (1985). On the convex layers of a planar set. *IEEE Transactions on Information Theory* **31** 509–517. <https://doi.org/10.1109/tit.1985.1057060>. [MR0798557](#)
- [5] ELIZONDO, D. (2006). The linear separability problem: some testing methods. *IEEE Transactions on Neural Networks* **17** 330–344. <https://doi.org/10.1109/TNN.2005.860871>.
- [6] FRANK, O. and STRAUSS, D. (1986). Markov graphs. *Journal of the American Statistical Association* **81** 832–842. <https://doi.org/10.1080/01621459.1986.10478342>. [MR0860518](#)
- [7] GELFAND, A. E. and CARLIN, B. P. (1993). Maximum-likelihood estimation for constrained- or missing-data models. *Canadian Journal of Statistics* **21** 303–311. <https://doi.org/10.2307/3315756>. [MR1241062](#)
- [8] GEYER, C. J. (1994). On the convergence of Monte Carlo maximum likelihood calculations. *Journal of the Royal Statistical Society: Series B* **56** 261–274. <https://doi.org/10.1111/j.2517-6161.1994.tb01976.x>. [MR1257812](#)
- [9] GEYER, C. J. (2009). Likelihood inference in exponential families and directions of recession. *Electronic Journal of Statistics* **3** 259–289. <https://doi.org/10.1214/08-EJS349>. [MR2495839](#)
- [10] GEYER, C. J. and THOMPSON, E. A. (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society: Series B* **54** 657–683. <https://doi.org/10.1111/j.2517-6161.1992.tb01443.x>. [MR1185217](#)
- [11] HANDCOCK, M. S. and GILE, K. J. (2010). Modeling social networks from sampled data. *Annals of Applied Statistics* **4** 5–25. <https://doi.org/10.1214/08-AOAS221>. [MR2758082](#)
- [12] HANDCOCK, M. S., HUNTER, D. R., BUTTS, C. T., GOODREAU, S. M., KRIVITSKY, P. N. and MORRIS, M. (2023). `ergm`: Fit, simulate and diagnose exponential-family models for networks The Statnet Project (<https://statnet.org>) R package version 4.5.0.
- [13] HOLLAND, P. W. and LEINHARDT, S. (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association* **76** 33–65. [MR0608176](#)
- [14] HUMMEL, R. M., HUNTER, D. R. and HANDCOCK, M. S. (2012). Improving simulation-based algorithms for fitting ERGMs. *Journal of Computational and Graphical Statistics* **21** 920–939. <https://doi.org/10.1080/10618600.2012.679224>. [MR3005804](#)
- [15] HUNTER, D. R., HANDCOCK, M. S., BUTTS, C. T., GOODREAU, S. M.

- and MORRIS, M. (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software* **24** 1–29. <https://doi.org/10.18637/jss.v024.i03>.
- [16] JAYNES, E. T. (1957). Information theory and statistical mechanics. *Physical Review* **106** 620–630. <https://doi.org/10.1103/physrev.106.620>. MR0087305
- [17] MAKHORIN, A. O. et al. (2020). GLPK (GNU Linear Programming Kit).
- [18] LP_SOLVE, KONIS, K., SCHWENDINGER, F. and HORNIK, K. (2022). lpSolveAPI: R interface to lp_solve version 5.5.2.0 R package version 5.5.2.0-17.9.
- [19] ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics* **22** 400–407. <https://doi.org/10.1214/aoms/1177729586>. MR0042668
- [20] R CORE TEAM (2023). R: A language and environment for statistical computing R Foundation for Statistical Computing, Vienna, Austria.
- [21] SCHWEINBERGER, M., KRIVITSKY, P. N., BUTTS, C. T. and STEWART, J. R. (2020). Exponential-family models of random graphs: Inference in finite, super and infinite population scenarios. *Statistical Science* **35** 627–662. <https://doi.org/10.1214/19-STS743>. MR4175389
- [22] SHMUELI, G., MINKA, T. P., KADANE, J. B., BORLE, S. and BOATWRIGHT, P. (2005). A useful distribution for fitting discrete data: Revival of the Conway–Maxwell–Poisson distribution. *Journal of the Royal Statistical Society: Series C* **54** 127–142. <https://doi.org/10.1111/j.1467-9876.2005.00474.x>. MR2134602
- [23] SNIJDERS, T. A. B. (2002). Markov chain Monte Carlo estimation of exponential random graph models. *Journal of Social Structure* **3** 1–40.
- [24] SNIJDERS, T. A. B., PATTISON, P. E., ROBINS, G. L. and HANDCOCK, M. S. (2006). New specifications for exponential random graph models. *Sociological Methodology* **36** 99–153. <https://doi.org/10.1111/j.1467-9531.2006.00176.x>.
- [25] STRAUSS, D. J. (1975). A model for clustering. *Biometrika* **62** 467–475. <https://doi.org/10.1093/biomet/62.2.467>. MR0383493
- [26] THEUSSL, S. and HORNIK, K. (2023). Rglpk: R/GNU linear programming kit interface R package version 0.6-5.
- [27] TUKEY, J. W. (1975). Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians, Vancouver, 1974* (R. D. JAMES, ed.) **2** 523–531. MR0426989
- [28] VANDERBEI, R. J. (1997). *Linear Programming: A Modern Integrated Analysis*. Kluwer Academic Publishers, Boston. MR1433802
- [29] ZUO, Y. and SERFLING, R. (2000). General notions of statistical depth function. *The Annals of Statistics* **28** 461–482. <https://doi.org/10.1214/aos/1016218226>. MR1790005