

BOUNDS ON THE NUMBER OF INFERENCE FUNCTIONS OF A GRAPHICAL MODEL

Sergi Elizalde and Kevin Woods

Dartmouth College and Oberlin College

Abstract: Directed and undirected graphical models, also called Bayesian networks and Markov random fields, respectively, are important statistical tools in a wide variety of fields, ranging from computational biology to probabilistic artificial intelligence. We give an upper bound on the number of inference functions of any graphical model. This bound is polynomial on the size of the model, for a fixed number of parameters. We also show that our bound is tight up to a constant factor, by constructing a family of hidden Markov models whose number of inference functions agrees asymptotically with the upper bound. This paper elaborates and expands on results of the first author from Elizalde (2005).

Key words and phrases: Graphical models, hidden Markov models, inference functions, polytopes.

1. Introduction

Many statistical models seek, given a set of observed data, to find the *hidden* (unobserved) data which best explains these observations. In this paper we consider graphical models (both directed and undirected), a broad class that includes many useful models, such as hidden Markov models (HMMs), pairwise-hidden Markov models, hidden tree models, Markov random fields, and some language models (background on graphical models will be given in Section 2.1). These graphical models relate the hidden and observed data probabilistically, and a natural problem is to determine, given a particular observation, what is the most likely hidden data (which is called the *explanation*). These models rely on parameters that are the probabilities relating the hidden and observed data. Any fixed values of the parameters determine a way to assign an explanation to each possible observation. This gives us a map, called an *inference function*, from observations to explanations. We will define an “inference function” precisely in Definition 8.

An example of an inference function is the popular “*Did you mean*” feature from **google**¹ which could be implemented as a hidden Markov model, where the

¹For example, if we search for “grafical modl” (<http://www.google.com/search?q=grafical+modl>), we are kindly asked “Did you mean: *graphical model?*”

observed data is what we type into the computer, and the hidden data is what we were meaning to type. Graphical models are frequently used in these sorts of probabilistic approaches to machine learning, pattern recognition, and artificial intelligence (see Jensen (2001) for an introduction).

Inference functions for graphical models are also important in computational biology (Pachter and Sturmfels (2005, Sec. 1.5)), from which we originally drew inspiration for this paper. For example, consider the *gene-finding functions*, which were discussed in (Pachter and Sturmfels (2006, Sec. 5)). These inference functions (corresponding to a particular HMM) are used to identify gene structures in DNA sequences. An observation in such a model is a sequence of nucleotides in the alphabet $\Sigma' = \{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{T}\}$, and an explanation is a sequence of 1's and 0's which indicate whether the particular nucleotide is in a gene or is not. We seek to use the information in the observed data (which we can find via DNA sequencing) to decide on the hidden information of which nucleotides are part of genes (which is hard to figure out directly). Another class of examples is that of sequence alignment models (Pachter and Sturmfels (2005, Sec. 2.2)). In such models, an inference function is a map from a pair of DNA sequences to an optimal alignment of those sequences. If we change the parameters of the model, the alignments that are optimal may change, and so the inference functions may change.

A surprising conclusion of this paper is that there cannot be *too many* different inference functions, though the parameters may vary continuously over all possible choices. For example, in the homogeneous binary HMM of length 5 (see Section 2.1 for some definitions; they are not important at the moment), the observed data is a binary sequence of length 5, and the explanation will also be a binary sequence of length 5. At first glance, there are 32^{32} possible maps from observed sequences to explanations. In fact, Christophe Weibel has computed that only 5,266 of these possible maps are actually inference functions. Indeed, for an arbitrary graphical model, the number of possible maps from observed sequences to explanations is, at first glance, doubly exponential in the size of the model. The following theorem, which we call the *Few Inference Functions Theorem*, states that, if we fix the number of parameters, the number of inference functions is actually bounded by a polynomial in the size of the model.

Theorem 1.(The Few Inference Functions Theorem) *Let d be a fixed positive integer. Consider a graphical model with d parameters (see Definitions 3 and 5 for directed and undirected graphs, respectively). Let M be the complexity of the graphical model, where complexity is given by Definitions 4 and 6, respectively. Then, the number of inference functions of the model is $O(M^{d(d-1)})$.*

As we shall see, the complexity of a graphical model is often linear in the number of vertices or edges of the underlying graph.

The Few Inference Functions Theorem for the particular case of undirected graphical models appears in Elizalde (2005). Here we extend it to the case of directed graphical models, and we prove that the bound is asymptotically sharp, up to a constant factor.

Different inference functions represent different criteria to decide what is the most likely explanation for each observation. A bound on the number of inference functions is important because it indicates how badly a model may respond to changes in the parameter values (which are generally known with very little certainty and only guessed at). Also, the polynomial bound given in Section 3 suggests that it might be feasible to precompute all the inference functions of a given graphical model, which would yield an efficient way to provide an explanation for each given observation.

This polynomial bound with exponent $d(d-1)$ is asymptotically sharp for a sequence alignment model with 2 parameters that is actually used in computational biology. This example is given in (Elizalde (2005, Sec. 9.3)) and, in that case, the bound is quadratic on the length of the input DNA sequences.

This paper is structured as follows. In Section 2 we introduce some preliminaries about graphical models and inference functions, as well as some facts about polytopes. In Section 3 we prove Theorem 1. The main ideas in that section appeared in (Elizalde (2005)). In Section 4 we prove that our upper bound on the number of inference functions of a graphical model is sharp, up to a constant factor, by constructing a family of HMMs whose number of inference functions asymptotically matches the bound. We conclude with a few remarks and possible directions for further research.

2. Preliminaries

2.1. Graphical models

A *statistical model* is a family of joint probability distributions for a collection of discrete random variables $\mathbf{W} = (W_1, \dots, W_m)$, where each W_i takes on values in some finite state space Σ_i . A *graphical model* is represented by a graph where each vertex v_i corresponds to a random variable W_i . The edges of the graph represent the dependencies between the variables. There are two major classes of graphical models depending on whether G is a directed or an undirected graph.

We start by discussing *directed graphical models*, also called *Bayesian networks*, which are those represented by a finite directed acyclic graph G . Each

vertex v_i has an associated probability map

$$p_i : \left(\prod_{j: v_j \text{ a parent of } v_i} \Sigma_j \right) \longrightarrow [0, 1]^{|\Sigma_i|}. \tag{1}$$

Given the states of each W_j such that v_j is a parent of v_i , the probability that v_i has a given state is independent of all other vertices that are not descendants of v_i , and this map p_i gives that probability. In particular, we have the equality

$$\begin{aligned} \text{Prob}(\mathbf{W} = \rho) &= \prod_i \text{Prob}(W_i = \rho_i, \text{ given that } W_j = \rho_j \text{ for all parents } v_j \text{ of } v_i) \\ &= \prod_i \left([p_i(\rho_{j_1}, \dots, \rho_{j_k})]_{\rho_i} \right), \end{aligned}$$

where v_{j_1}, \dots, v_{j_k} are the parents of v_i . Sources in the digraph (which have no parents) are generally given the uniform probability distribution on their states, though more general distributions are possible. See (Pachter and Sturmfels (2005, Sec. 1.5)) for general background on graphical models.

Example 2. The hidden Markov model (HMM) is a model with random variables $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{Y} = (Y_1, \dots, Y_n)$. Edges go from X_i to X_{i+1} and from X_i to Y_i .




Figure 1. The graph of an HMM for $n = 3$.

Generally, each X_i has the same state space Σ and each Y_i has the same state space Σ' . An HMM is called *homogeneous* if the p_{X_i} , for $1 \leq i \leq n$, are identical and the p_{Y_i} are identical. In this case, the p_{X_i} each correspond to the same $|\Sigma| \times |\Sigma|$ matrix $T = (t_{ij})$ (the *transition matrix*) and the p_{Y_i} each correspond to the same $|\Sigma| \times |\Sigma'|$ matrix $S = (s_{ij})$ (the *emission matrix*).

In the example, we have partitioned the variables into two sets. In general graphical models, we also have two kinds of variables: observed variables $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ and hidden variables $\mathbf{X} = (X_1, X_2, \dots, X_q)$. Generally, the observed variables are the sinks of the directed graph, and the hidden variables are the other vertices, but this does not need to be the case. To simplify the

notation, we make the assumption, which is often the case in practice, that all the observed variables take their values in the same finite alphabet Σ' , and that all the hidden variables are on the finite alphabet Σ .

Notice that for given Σ and Σ' the homogeneous HMMs in this example depend only on a fixed set of parameters, t_{ij} and s_{ij} , even as n gets large. These are the sorts of models we are interested in.

Definition 3. A *directed graphical model with d parameters, $\theta_1, \dots, \theta_d$* , is a directed graphical model such that each probability $[p_i(\rho_{j_1}, \dots, \rho_{j_k})]_{\rho_i}$ in (1) is a monomial in $\theta_1, \dots, \theta_d$.

In what follows we denote by E the number of edges of the underlying graph of a graphical model, by n the number of observed random variables, and by q the number of hidden random variables. The observations, then, are sequences in $(\Sigma')^n$ and the explanations are sequences in Σ^q . Let $l = |\Sigma|$ and $l' = |\Sigma'|$.

For each observation τ and hidden variables \mathbf{h} , $\text{Prob}(\mathbf{X} = \mathbf{h}, \mathbf{Y} = \tau)$ is a monomial $f_{\mathbf{h}, \tau}$ in the parameters $\theta_1, \dots, \theta_d$. Then for each observation $\tau \in (\Sigma')^n$, the observed probability $\text{Prob}(\mathbf{Y} = \tau)$ is the sum over all hidden data \mathbf{h} of $\text{Prob}(\mathbf{X} = \mathbf{h}, \mathbf{Y} = \tau)$, and so $\text{Prob}(\mathbf{Y} = \tau)$ is the polynomial $f_\tau = \sum_{\mathbf{h}} f_{\mathbf{h}, \tau}$ in the parameters $\theta_1, \dots, \theta_d$.

Definition 4. The *complexity, M* , of a directed graphical model is the maximum, over all τ , of the degree of the polynomial f_τ .

In many graphical models, M will be a linear function of n , the number of observed variables. For example, in the homogeneous HMM, $M = E = 2n - 1$.

Note that we have not assumed that the appropriate probabilities sum to 1. It turns out that the analysis is much easier if we do not place that restriction on our probabilities. At the end of the analysis, these restrictions may be added if desired (there are many models in use, however, which never place that restriction; these can no longer be properly called “probabilistic” models, but in fact belong to a more general class of “scoring” models which our analysis also encompasses).

The other class of graphical models are those that are represented by an undirected graph. They are called *undirected graphical models* and are also known as *Markov random fields*. As for directed models, the vertices of the graph G correspond to the random variables, but the joint probability is now represented as a product of local functions defined on the maximal cliques of the graph, instead of transition probabilities p_i defined on the edges.

Recall that a *clique* of a graph is a set of vertices with the property that there is an edge between any two of them. A clique is *maximal* if it cannot be extended to include additional vertices without losing the property of being a clique (see Figure 2).




Figure 2. An undirected graph with maximal cliques $\{v_1, v_2\}$, $\{v_2, v_3\}$, $\{v_2, v_4, v_5\}$, $\{v_3, v_6\}$, and $\{v_5, v_6\}$.

Each maximal clique C of the graph G has an associated *potential function*

$$\psi_C : \left(\prod_{j: v_j \in C} \Sigma_j \right) \longrightarrow \mathbb{R}. \quad (2)$$

Given the states ρ_j of each W_j such that v_j is a vertex in the clique C , if we denote by ρ_C the vector of such states, then $\psi_C(\rho_C)$ is a nonnegative real number. We denote by \mathcal{C} the set of all maximal cliques C .

Then, the joint probability distribution of all the variables W_i is given by

$$\text{Prob}(\mathbf{W} = \rho) = \frac{1}{Z} \prod_{C \in \mathcal{C}} \psi_C(\rho_C),$$

where Z is the normalization factor

$$Z = \sum_{\rho} \prod_{C \in \mathcal{C}} \psi_C(\rho_C),$$

obtained by summing over all assignments of values to the variables ρ .

The value of the function $\psi_C(\rho_C)$ for each possible choice of the states ρ_i is given by the parameters of the model. We will be interested in models in which the set of parameters is fixed, even as the size of the graph gets large.

Definition 5. An *undirected graphical model with d parameters, $\theta_1, \dots, \theta_d$* , is an undirected graphical model such that each probability $\psi_C(\rho_C)$ in (2) is a monomial in $\theta_1, \dots, \theta_d$.

As in the case of directed models, the variables can be partitioned into observed variables $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ (which can be assumed to take their values in the same finite alphabet Σ') and hidden variables $\mathbf{X} = (X_1, X_2, \dots, X_q)$ (which can be assumed to be on the finite alphabet Σ). For each observation τ and hidden variables \mathbf{h} , $Z \cdot \text{Prob}(\mathbf{X} = \mathbf{h}, \mathbf{Y} = \tau)$ is a monomial $f_{\mathbf{h}, \tau}$ in the parameters $\theta_1, \dots, \theta_d$. Then for each observation $\tau \in (\Sigma')^n$, the observed probability

$\text{Prob}(\mathbf{Y} = \tau)$ is the sum over all hidden data \mathbf{h} of $\text{Prob}(\mathbf{X} = \mathbf{h}, \mathbf{Y} = \tau)$, and so $Z \cdot \text{Prob}(\mathbf{Y} = \tau)$ is the polynomial $f_\tau = \sum_{\mathbf{h}} f_{\mathbf{h},\tau}$ in the parameters $\theta_1, \dots, \theta_d$.

Definition 6. The *complexity*, M , of an undirected graphical model is the maximum, over all τ , of the degree of the polynomial f_τ .

It is usually the case for undirected models, as in directed, that M is a linear function of n .

2.2. Inference functions

For fixed values of the parameters, the basic inference problem is to determine, for each given observation τ , a value $\mathbf{h} \in \Sigma^q$ of the hidden data that maximizes $\text{Prob}(\mathbf{X} = \mathbf{h} \mid \mathbf{Y} = \tau)$. A solution to this optimization problem is denoted $\hat{\mathbf{h}}$ and is called an *explanation* of the observation τ . Each choice of parameter values $(\theta_1, \theta_2, \dots, \theta_d)$ defines an *inference function* $\tau \mapsto \hat{\mathbf{h}}$ from the set of observations $(\Sigma')^n$ to the set of explanations Σ^q .

Example 7. Consider the *binary* homogeneous HMM with $n=3$ (see Example 2) with hidden states $\Sigma = \{A, B\}$ and observed states $\Sigma' = \{0, 1\}$. Suppose the transition matrix and emission matrix are

$$\begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix} \text{ and } \begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix},$$

respectively, and that the source X_1 has the uniform probability distribution.

If, for example, the string 010 is observed ($Y_1 = 0, Y_2 = 1, Y_3 = 0$), then the most likely values of the hidden variables are AAA , with

$$\text{Prob}(\mathbf{X} = AAA, \mathbf{Y} = 010) = 0.5 \cdot 0.6 \cdot 0.6 \cdot 0.4 \cdot 0.6 \cdot 0.6 = 0.02592.$$

Therefore our inference function should map 010 to AAA . On the other hand, if the string 011 is observed, then there are actually two possibilities for the explanation: ABB and BBB are equally likely and are also more likely than any other string of hidden variables. One possible solution is to say that the inference function maps 011 to the set of all possible explanations, that is, $011 \mapsto \{ABB, BBB\}$. Repeating this process for each possible string of observed values, we get the inference function given by

$$\begin{array}{ll} 000 \mapsto \{AAA\} & 100 \mapsto \{AAA, BAA\} \\ 001 \mapsto \{AAA, AAB\} & 101 \mapsto \{BBB\} \\ 010 \mapsto \{AAA\} & 110 \mapsto \{BBA, BBB\} \\ 011 \mapsto \{ABB, BBB\} & 111 \mapsto \{BBB\}. \end{array}$$

For simplicity, we would like to pick only one such explanation for each possible observed sequence, according to some consistent tie-breaking rule decided ahead of time (this will not affect the results of the paper, merely ease exposition). For example, we could pick the lexicographically first among the possibilities. This would give us the inference function, Φ , mapping

$$\begin{array}{ll} 000 \mapsto AAA & 100 \mapsto AAA \\ 001 \mapsto AAA & 101 \mapsto BBB \\ 010 \mapsto AAA & 110 \mapsto BBA \\ 011 \mapsto ABB & 111 \mapsto BBB. \end{array}$$

In general, we fix an order of the hidden states Σ , that is, if $\Sigma = \{\sigma_1, \dots, \sigma_l\}$, we say that $\sigma_1 < \sigma_2 < \dots < \sigma_l$.

Definition 8. An *inference function* is a map $\Phi : (\Sigma')^n \rightarrow \Sigma^q$ that assigns to each observation $\tau \in (\Sigma')^n$ an explanation $\hat{\mathbf{h}} \in \Sigma^q$ that maximizes $\text{Prob}(\mathbf{X} = \mathbf{h} | \mathbf{Y} = \tau)$. For definiteness, if there is more than one such explanation, we define $\Phi(\tau)$ to be the minimum of all such $\hat{\mathbf{h}}$ in lexicographic order.

It is interesting to observe that the total number of maps $(\Sigma')^n \rightarrow \Sigma^q$ is $(l^q)^{(l')^n} = l^{q(l')^n}$, which is doubly-exponential in the length n of the observations. However, the vast majority of these maps are not inference functions for any values of the parameters. Before our results, the best upper bound in the literature is an exponential bound given in (Pachter and Sturmfels (2004, Cor. 10)) Theorem 1 gives a polynomial upper bound on the number of inference functions of a graphical model.

2.3. Polytopes

Here we review some facts about convex polytopes, and we introduce some notation. Recall that a polytope is a bounded intersection of finitely many closed halfspaces, or equivalently, the convex hull of a finite set of points. For the basic definitions about polytopes we refer the reader to (Ziegler (1995)).

Given a polynomial $f(\theta) = \sum_{i=1}^N \theta_1^{a_{1,i}} \theta_2^{a_{2,i}} \dots \theta_d^{a_{d,i}}$, its *Newton polytope*, denoted by $\text{NP}(f)$, is defined as the convex hull in \mathbb{R}^d of the set of points $\{(a_{1,i}, a_{2,i}, \dots, a_{d,i}) : i = 1, \dots, N\}$.

For example, if $f(\theta_1, \theta_2) = 2\theta_1^3 + 3\theta_1^2\theta_2^2 + \theta_1\theta_2^2 + 3\theta_1 + 5\theta_2^4$, then its Newton polytope $\text{NP}(f)$ is given in Figure 3.




Figure 3. The Newton polytope of $f(\theta_1, \theta_2) = 2\theta_1^3 + 3\theta_1^2\theta_2^2 + \theta_1\theta_2^2 + 3\theta_1 + 5\theta_2^4$.

Given a polytope $P \subset \mathbb{R}^d$ and a vector $w \in \mathbb{R}^d$, the set of all points in P at which the linear functional $x \mapsto x \cdot w$ attains its maximum determines a *face* of P . It is denoted

$$\text{face}_w(P) = \{x \in P : x \cdot w \geq y \cdot w \text{ for all } y \in P\}. \tag{3}$$

Faces of dimension 0 (consisting of a single point) are called *vertices*, and faces of dimension 1 are called *edges*. If d is the dimension of the polytope, then faces of dimension $d - 1$ are called *facets*.

Let P be a polytope and F a face of P . The *normal cone* of P at F is

$$N_P(F) = \{w \in \mathbb{R}^d : \text{face}_w(P) = F\}.$$

The collection of all cones $N_P(F)$ as F runs over all faces of P is denoted $\mathcal{N}(P)$ and is called the *normal fan* of P . Thus the normal fan $\mathcal{N}(P)$ is a partition of \mathbb{R}^d into cones. The cones in $\mathcal{N}(P)$ are in bijection with the faces of P , and if $w \in N_P(F)$, then the linear functional $x \cdot w$ is maximized on F . Figure 4 shows the normal fan of the polytope from Figure 3.




Figure 4. The normal fan of a polytope.

The *Minkowski sum* of two polytopes P and P' is defined as

$$P + P' := \{x + x' : x \in P, x' \in P'\}.$$

Figure 5 shows an example in 2 dimensions. The Newton polytope of the map $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^{(l')^n}$ is defined as the Minkowski sum of the individual Newton polytopes of its coordinates, namely $\text{NP}(\mathbf{f}) := \sum_{\tau \in (\Sigma')^n} \text{NP}(f_\tau)$.




Figure 5. Two polytopes and their Minkowski sum.

The *common refinement* of two or more normal fans is the collection of cones obtained as the intersection of a cone from each of the individual fans. For polytopes P_1, P_2, \dots, P_k , the common refinement of their normal fans is denoted $\mathcal{N}(P_1) \wedge \dots \wedge \mathcal{N}(P_k)$. Figure 6 shows the normal fans for the polytopes P and P' from Figure 5, together with the common refinement. Comparing $\mathcal{N}(P) \wedge \mathcal{N}(P')$ to the polytope $P + P'$ in Figure 5, we see an illustration of the well-known fact that the normal fan of a Minkowski sum of polytopes is the common refinement of their individual fans (see Ziegler (1995), Proposition 7.12, or Gritzmann and Sturmfels (1993), Lemma 2.1.5). To be precise:

Lemma 9. $\mathcal{N}(P_1 + \dots + P_k) = \mathcal{N}(P_1) \wedge \dots \wedge \mathcal{N}(P_k)$.

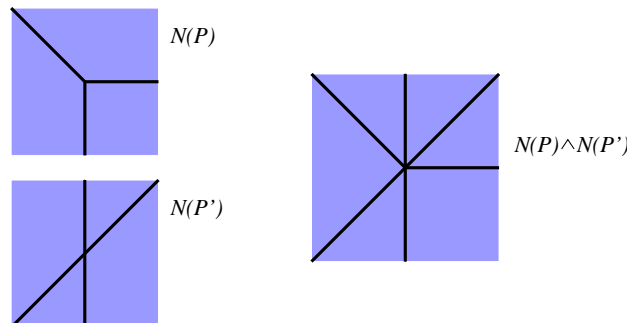


Figure 6. The normal fans of the polytopes in Figure 5 and their common refinement.

We finish with a result of Gritzmann and Sturmfels (1993) that will be useful later. It gives a bound on the number of vertices of a Minkowski sum of polytopes.

Theorem 10. *Let P_1, \dots, P_k be polytopes in \mathbb{R}^d , and let m denote the number of non-parallel edges of P_1, \dots, P_k . Then the number of vertices of $P_1 + \dots + P_k$ is at most*

$$2 \sum_{j=0}^{d-1} \binom{m-1}{j}.$$

Note that this bound is independent of the number k of polytopes.

3. An upper Bound on the Number of Inference Functions

For fixed parameters, the inference problem of finding the explanation $\hat{\mathbf{h}}$ that maximizes $\text{Prob}(\mathbf{X} = \mathbf{h} | \mathbf{Y} = \tau)$ is equivalent to identifying the monomial $f_{\hat{\mathbf{h}}, \tau} = \theta_1^{a_{1,i}} \theta_2^{a_{2,i}} \dots \theta_d^{a_{d,i}}$ of f_τ with maximum value. Since the logarithm is a monotonically increasing function, the desired monomial also maximizes the quantity

$$\begin{aligned} \log(\theta_1^{a_{1,i}} \theta_2^{a_{2,i}} \dots \theta_d^{a_{d,i}}) &= a_{1,i} \log(\theta_1) + a_{2,i} \log(\theta_2) + \dots + a_{d,i} \log(\theta_d) \\ &= a_{1,i} v_1 + a_{2,i} v_2 + \dots + a_{d,i} v_d, \end{aligned}$$

where we replace $\log(\theta_i)$ with v_i . This is equivalent to the fact that the corresponding point $(a_{1,i}, a_{2,i}, \dots, a_{d,i})$ maximizes the linear expression $v_1 x_1 + \dots + v_d x_d$ on the Newton polytope $\text{NP}(f_\tau)$. Thus, the inference problem for fixed parameters becomes a linear programming problem.

Each choice of the parameters $\theta = (\theta_1, \theta_2, \dots, \theta_d)$ determines an inference function. If $\mathbf{v} = (v_1, v_2, \dots, v_d)$ is the vector in \mathbb{R}^d with coordinates $v_i = \log(\theta_i)$, then we denote the corresponding inference function by

$$\Phi_{\mathbf{v}} : (\Sigma')^n \longrightarrow \Sigma^q.$$

For each observation $\tau \in (\Sigma')^n$, its explanation $\Phi_{\mathbf{v}}(\tau)$ is given by the vertex of $\text{NP}(f_\tau)$ that is maximal in the direction of the vector \mathbf{v} . Note that for certain values of the parameters (if \mathbf{v} is perpendicular to a positive-dimensional face of $\text{NP}(f_\tau)$) there may be more than one vertex attaining the maximum. It is also possible that a single point $(a_{1,i}, a_{2,i}, \dots, a_{d,i})$ in the polytope corresponds to several different values of the hidden data. In both cases, when there is more than one possible explanation attaining the maximal probability, we pick the explanation according to Definition 8. This simplification does not affect the asymptotic number of inference functions.

Different values of θ yield different directions \mathbf{v} , which can result in distinct inference functions. We are interested in bounding the number of different inference functions that a graphical model can have. Theorem 1 gives an upper bound

which is polynomial in the size of the graphical model. In other words, extremely few of the $l^{q(l')^n}$ functions $(\Sigma')^n \rightarrow \Sigma^q$ are actually inference functions.

Before proving Theorem 1, observe that usually M , the complexity of the graphical model, is linear in n . For example, in the case of directed models, consider the common situation where M is bounded by E , the number of edges of the underlying graph (this happens when each edge “contributes” at most degree 1 to the monomials $f_{\mathbf{h},\tau}$, as in the homogeneous HMM). In most graphical models of interest, E is a linear function of n , so the bound becomes $O(n^{d(d-1)})$. For example, the homogeneous HMM has $M = E = 2n - 1$.

Also, in the case of undirected models, if each $\psi_C(\rho_C)$ is a parameter of the model, then $f_{\mathbf{h},\tau} = Z \cdot \text{Prob}(\mathbf{X} = \mathbf{h}, \mathbf{Y} = \tau)$ is a product of potential functions for each maximal clique of the graph, so M is bounded by the number of maximal cliques, which in many cases is also a linear function of the number of vertices of the graph. For example, this is the situation in language models where each word depends on a fixed number of previous words in the sentence.

Proof of Theorem 1. In the first part of the proof we reduce the problem of counting inference functions to the enumeration of the vertices of a certain polytope. We have seen that an inference function is specified by a choice of the parameters, which is equivalent to choosing a vector $\mathbf{v} \in \mathbb{R}^d$. The function is denoted $\Phi_{\mathbf{v}} : (\Sigma')^n \rightarrow \Sigma^q$, and the explanation $\Phi_{\mathbf{v}}(\tau)$ of a given observation τ is determined by the vertex of $\text{NP}(f_{\tau})$ that is maximal in the direction of \mathbf{v} . Thus, cones of the normal fan $\mathcal{N}(\text{NP}(f_{\tau}))$ correspond to sets of vectors \mathbf{v} that give rise to the same explanation for the observation τ . Non-maximal cones (i.e., those contained in another cone of higher dimension) correspond to directions \mathbf{v} for which more than one vertex is maximal. Since ties are broken using a consistent rule, we disregard this case for simplicity. Thus, in what follows we consider only maximal cones of the normal fan.

Let $\mathbf{v}' = (v'_1, \dots, v'_d)$ be another vector corresponding to a different choice of parameters (see Figure 7). By the above reasoning, $\Phi_{\mathbf{v}}(\tau) = \Phi_{\mathbf{v}'}(\tau)$ if and only if \mathbf{v} and \mathbf{v}' belong to the same cone of $\mathcal{N}(\text{NP}(f_{\tau}))$. Thus, $\Phi_{\mathbf{v}}$ and $\Phi_{\mathbf{v}'}$ are the same inference function if and only if \mathbf{v} and \mathbf{v}' belong to the same cone of $\mathcal{N}(\text{NP}(f_{\tau}))$ for all observations $\tau \in (\Sigma')^n$. Consider the common refinement of all these normal fans, $\bigwedge_{\tau \in (\Sigma')^n} \mathcal{N}(\text{NP}(f_{\tau}))$. Then, $\Phi_{\mathbf{v}}$ and $\Phi_{\mathbf{v}'}$ are the same function exactly when \mathbf{v} and \mathbf{v}' lie in the same cone of this common refinement.

This implies that the number of inference functions equals the number of cones in

$$\bigwedge_{\tau \in (\Sigma')^n} \mathcal{N}(\text{NP}(f_{\tau})).$$

By Lemma 9, this common refinement is the normal fan of $\text{NP}(\mathbf{f}) = \sum_{\tau \in (\Sigma')^n} \text{NP}(f_{\tau})$, the Minkowski sum of the polytopes $\text{NP}(f_{\tau})$ for all observations τ . It

follows that enumerating inference functions is equivalent to counting vertices of $\text{NP}(\mathbf{f})$. In the remaining part of the proof we give an upper bound on the number of vertices of $\text{NP}(\mathbf{f})$.




Figure 7. Two different inference functions, $\Phi_{\mathbf{v}}$ (left column) and $\Phi_{\mathbf{v}'}$ (right column). Each row corresponds to a different observation. The respective explanations are given by the marked vertices in each Newton polytope.

Note that for each τ , the polytope $\text{NP}(f_\tau)$ is contained in the hypercube $[0, M]^d$, since by definition of M , each parameter θ_i appears in f_τ with exponent at most M . Also, the vertices of $\text{NP}(f_\tau)$ have integral coordinates, because they are exponent vectors. Polytopes whose vertices have integral coordinates are called *lattice polytopes*. It follows that the edges of $\text{NP}(f_\tau)$ are given by vectors where each coordinate is an integer between $-M$ and M . There are only $(2M + 1)^d$ such vectors, so this is an upper bound on the number of different directions that the edges of the polytopes $\text{NP}(f_\tau)$ can have.

This property of the Newton polytopes of the coordinates of the model will allow us to give an upper bound on the number of vertices of their Minkowski sum $\text{NP}(\mathbf{f})$. The last ingredient that we need is Theorem 10. In our case we have a sum of polytopes $\text{NP}(f_\tau)$, one for each observation $\tau \in (\Sigma')^n$, having at most $(2M + 1)^d$ non-parallel edges in total. Hence, by Theorem 10, the number of vertices of $\text{NP}(\mathbf{f})$ is at most

$$2 \sum_{j=0}^{d-1} \binom{(2M + 1)^d - 1}{j}.$$

As M goes to infinity, the dominant term of this expression is

$$\frac{2^{d^2-d+1}}{(d-1)!} M^{d(d-1)}.$$

Thus, we get an $O(M^{d(d-1)})$ upper bound on the number of inference functions of the graphical model.

In the next section we show that the bound given in Theorem 1 is tight up to a constant factor.

4. A Lower Bound

As before, we fix d , the number of parameters in our model. The Few Inferences Function Theorem tells us that the number of inference functions is bounded from above by some function $cM^{d(d-1)}$, where c is a constant (depending only on d) and M is the complexity of the model. Here we show that that bound is tight up to a constant, by constructing a family of graphical models whose number of inference functions is at least $c'M^{d(d-1)}$, where c' is another constant. In fact, we will construct a family of hidden Markov models with this property. To be precise, we have the following theorem.

Theorem 11. *Fix d . There is a constant $c' = c'(d)$ such that, given $n \in \mathbb{Z}_+$, there exists an HMM of length n , with d parameters, $4d + 4$ hidden states, and 2 observed states, such that there are at least $c'n^{d(d-1)}$ distinct inference functions. (For this HMM, M is a linear function of n , so this also gives us the lower bound in terms of M).*

In Section 4.1 we prove Theorem 11. This proof requires several lemmas that will be proved in Section 4.2. Lemma 15, which is interesting in its own right as a statement in the geometry of numbers, is proved in Elizalde and Woods (2006).

4.1. Proof of Theorem 11

Given n , we first construct the appropriate HMM, \mathcal{M}_n , using the following lemma.

Lemma 12. *Given $n \in \mathbb{Z}_+$, there is an HMM, \mathcal{M}_n , of length n , with d parameters, $4d + 4$ hidden states, and 2 observed states, such that for any $a \in \mathbb{Z}_+^d$ with $\sum_i a_i < n$, there is an observed sequence which has one explanation if $a_1 \log(\theta_1) + \dots + a_d \log(\theta_d) > 0$ and another explanation if $a_1 \log(\theta_1) + \dots + a_d \log(\theta_d) < 0$.*

This means that, for the HMM \mathcal{M}_n , the decomposition of (log-)parameter space into inference cones includes all of the hyperplanes $\{x : \langle a, x \rangle = 0\}$ such that $a \in \mathbb{Z}_+^d$ with $\sum_i a_i < n$. Call the arrangement of these hyperplanes \mathcal{H}_n . It

suffices to show that the arrangement \mathcal{H}_n consists of at least $c'n^{d(d-1)}$ chambers (full dimensional cones determined by the arrangement). There are c_1n^d ways to choose one of the hyperplanes from \mathcal{H}_n , for some constant c_1 . Therefore there are $c_1^{d-1}n^{d(d-1)}$ ways to choose $d - 1$ of the hyperplanes; their intersection is, in general, a 1-dimensional face of \mathcal{H}_n (that is, the intersection is a ray which is an extreme ray for the cones it is contained in). It is quite possible that two different ways of choosing $d - 1$ hyperplanes give the same extreme ray. The following lemma says that some constant fraction of these choices of extreme rays are actually distinct.

Lemma 13. *Fix d . Given n , let \mathcal{H}_n be the hyperplane arrangement consisting of the hyperplanes of the form $\{x : \langle a, x \rangle = 0\}$ with $a \in \mathbb{Z}_+^d$ and $\sum_i a_i < n$. Then the number of 1-dimensional faces of \mathcal{H}_n is at least $c_2n^{d(d-1)}$, for some constant c_2 .*

Each chamber will have a number of these extreme rays on its boundary. The following lemma gives a constant bound on this number.

Lemma 14. *Fix d . Given n , define \mathcal{H}_n as above. Each chamber of \mathcal{H}_n has at most $2^{d(d-1)}$ extreme rays.*

Conversely, each ray is an extreme ray for at least 1 chamber. Therefore there are at least $(c_2/2^{d(d-1)})n^{d(d-1)}$ chambers, and Theorem 11 is proved.

In proving Lemma 13, we will need one more lemma. This lemma is interesting in its own right as a probabilistic statement about integer lattices, and so is proved in the companion paper Elizalde and Woods (2006). Given a set $S \subset \mathbb{Z}^d$ of integer vectors, $span_{\mathbb{R}}(S)$ is a linear subspace of \mathbb{R}^d and $span_{\mathbb{R}}(S) \cap \mathbb{Z}^d$ is a sublattice of \mathbb{Z}^d . We say that S is *primitive* if S is a \mathbb{Z} -basis for the lattice $span_{\mathbb{R}}(S) \cap \mathbb{Z}^d$. Equivalently, a set S is primitive if and only if it may be extended to a \mathbb{Z} -basis of all of \mathbb{Z}^d (see Lekkerkerker (1969)).

We imagine picking each vector in S uniformly at random from some large box in \mathbb{R}^d . As the size of the box approaches infinity, the following lemma will tell us that the probability that S is primitive approaches

$$\frac{1}{\zeta(d)\zeta(d-1)\cdots\zeta(d-m+1)},$$

where $|S| = m$ and $\zeta(a)$ is the Riemann Zeta function $\sum_{i=1}^{\infty} 1/i^a$.

Lemma 15. *Let d and m be given, with $m < d$. For $n \in \mathbb{Z}_+$, $1 \leq k \leq m$, and $1 \leq i \leq d$, let $b_{n,k,i} \in \mathbb{Z}$. For a given n , choose integers s_{ki} uniformly (and independently) at random from the set $b_{n,k,i} \leq s_{ki} \leq b_{n,k,i} + n$. Let $s_k = (s_{k1}, \dots, s_{kd})$ and let $S = \{s_1, s_2, \dots, s_m\}$.*

If $|b_{n,k,i}|$ is bounded by a polynomial in n , then, as n approaches infinity, the probability that S is a primitive set approaches

$$\frac{1}{\zeta(d)\zeta(d-1)\cdots\zeta(d-m+1)},$$

where $\zeta(a)$ is the Riemann Zeta function $\sum_{i=1}^{\infty} 1/i^a$.

When $m = 1$, this lemma gives the probability that a d -tuple of integers are relatively prime as $1/\zeta(d)$. For $m = 1, d = 2$, this is a classic result in number theory (see Apostol (1976)), and for $m = 1, d > 2$, this was proven in Nyman (1972). Note also that, if $m = d$ and we choose S of size m , then the probability that S is primitive (i.e., that it is a basis for \mathbb{Z}^d) approaches zero. This agrees with the lemma in the sense that we would expect the probability to be

$$\frac{1}{\zeta(d)\zeta(d-1)\cdots\zeta(1)},$$

but $\zeta(1)$ does not converge.

4.2. Proofs of Lemmas

Proof of Lemma 12. Given d and n , define a length n HMM with parameters $\theta_1, \dots, \theta_d$, as follows. The observed states will be S and C (for “start of block,” and “continuing block,” respectively). The hidden states will be s_i, s'_i, c_i , and c'_i , for $1 \leq i \leq d + 1$ (think of s_i and s'_i as “start of the i th block” and c_i and c'_i as “continuing the i th block”).

Here is the idea of what we want this HMM to do: if the observed sequence has S’s in position $1, a_1 + 1, a_1 + a_2 + 1, \dots$, and $a_1 + \dots + a_d + 1$ and C’s elsewhere, then there will be only two possibilities for the sequence of hidden states, either

$$t = s_1 \underbrace{c_1 \cdots c_1}_{a_1-1} s_2 \underbrace{c_2 \cdots c_2}_{a_2-1} \cdots s_d \underbrace{c_d \cdots c_d}_{a_d-1} s_{d+1} \underbrace{c_{d+1} \cdots c_{d+1}}_{n-a_1-\dots-a_d-1}$$

or

$$t' = s'_1 \underbrace{c'_1 \cdots c'_1}_{a_1-1} s'_2 \underbrace{c'_2 \cdots c'_2}_{a_2-1} \cdots s'_d \underbrace{c'_d \cdots c'_d}_{a_d-1} s'_{d+1} \underbrace{c'_{d+1} \cdots c'_{d+1}}_{n-a_1-\dots-a_d-1}.$$

We also make sure that t has a priori probability $\theta_1^{a_1} \cdots \theta_d^{a_d}$ and t' has a priori probability 1. Then t is the explanation if $a_1 \log(\theta_1) + \dots + a_d \log(\theta_d) > 0$ and t' is the explanation if $a_1 \log(\theta_1) + \dots + a_d \log(\theta_d) < 0$. Remember that we are not constraining our probability sums to be 1. A very similar HMM could be constructed that obeys that constraint, if desired. To simplify notation it will be more convenient to treat the transition probabilities as parameters that do

not necessarily sum to one at each vertex, even if this forces us to use the term “probability” somewhat loosely.

Here is how we set up the transitions/emmissions. Let s_i and s'_i , for $1 \leq i \leq d + 1$, all emit S with probability 1 and C with probability 0. Let c_i and c'_i emit C with probability 1 and S with probability 0. Let s_i , for $1 \leq i \leq d$, transition to c_i with probability θ_i , transition to s_{i+1} with probability θ_i , and transition to everything else with probability 0. Let s_{d+1} transition to c_{d+1} with probability 1 and to everything else with probability 0. Let s'_i , for $1 \leq i \leq d$, transition to c'_i with probability 1, to s'_{i+1} with probability 1, and to everything else with probability 0. Let s'_{d+1} transition to c'_{d+1} with probability 1 and to everything else with probability 0. Let c_i , for $1 \leq i \leq d$, transition to c_i with probability θ_i , to s_{i+1} with probability θ_i , and to everything else with probability 0. Let c_{d+1} transition to c_{d+1} with probability 1 and to everything else with probability 0. Let c'_i , for $1 \leq i \leq d$, transition to c'_i with probability 1, to s_{i+1} with probability 1, and to everything else with probability 0. Let c'_{d+1} transition to c'_{d+1} with probability 1 and to everything else with probability 0.

Starting with the uniform probability distribution on the first hidden state, this does exactly what we want it to: given the correct observed sequence, t and t' are the only explanations, with the correct probabilities.

Proof of Lemma 13. We are going to pick $d - 1$ vectors $a^{(1)}, \dots, a^{(d-1)}$ which correspond to the $d - 1$ hyperplanes $\{x : \langle a^{(i)}, x \rangle = 0\}$ that will intersect to give us extreme rays of our chambers. We will restrict the region from which we pick each $a^{(i)} \in \mathbb{Z}^d$. Let

$$b^{(i)} = (1, 1, \dots, 1) - \frac{1}{2}e_i,$$

for $1 \leq i \leq d - 1$, where e_i is the i th standard basis vector. Let $s = 1/(4d + 4)$. For $1 \leq i \leq d - 1$, we choose $a^{(i)} \in \mathbb{Z}^d$ such that

$$\left\| \frac{n}{d}b^{(i)} - a^{(i)} \right\|_\infty < \frac{n}{d}s. \tag{4}$$

Note that $\sum_j a_j^{(i)} < n$, so there are observed sequences which give us the hyperplanes $\{x : \langle a^{(i)}, x \rangle = 0\}$. Note also that there are $(2s/d)^{d(d-1)}n^{d(d-1)}$ choices for the $(d - 1)$ -tuple of vectors $(a^{(1)}, \dots, a^{(d-1)})$. To prove this lemma, we must then show that a positive fraction of these actually give rise to distinct extreme rays $\bigcap_{i=1}^{d-1} \{x : \langle a^{(i)}, x \rangle = 0\}$.

First, we imagine choosing the $a^{(i)}$ uniformly at random in the range given by (4), this probability distribution meets the condition in the statement of Lemma 15, as n approaches infinity. Therefore, there is a positive probability that

$$\{a^{(i)} : 1 \leq i \leq d-1\} \text{ form a basis for the lattice } \mathbb{Z}^d \cap \text{span}\{a^{(i)} : 1 \leq i \leq d-1\}, \tag{5}$$

and this probability approaches

$$\frac{1}{\zeta(d)\zeta(d-1)\cdots\zeta(2)}.$$

Second, we look at all choices of $a^{(i)} \in \mathbb{Z}^d$ such that (4) and (5) hold. There are $c_2 n^{d(d-1)}$ of these, for some constant c_2 . We claim that these give distinct extreme rays $\bigcap_{i=1}^{d-1} \{x : \langle a^{(i)}, x \rangle = 0\}$. Indeed, say that $a^{(i)}$ and $c^{(i)}$ are both chosen such that (4) and (5) hold and such that

$$\bigcap_{i=1}^{d-1} \{x : \langle a^{(i)}, x \rangle = 0\} = \bigcap_{i=1}^{d-1} \{x : \langle c^{(i)}, x \rangle = 0\}.$$

We argue that $a^{(i)}$ and $c^{(i)}$ are “so close” that they must actually be the same.

Let j , for $1 \leq j \leq d - 1$ be given. We prove that $a^{(j)} = c^{(j)}$. Since

$$\bigcap_{i=1}^{d-1} \{x : \langle a^{(i)}, x \rangle = 0\} \subset \{x : \langle c^{(j)}, x \rangle = 0\},$$

we know that $c^{(j)}$ is in $\text{span}\{a^{(i)} : 1 \leq i \leq d - 1\}$, and therefore

$$c^{(j)} \in \mathbb{Z}^d \cap \text{span}\{a^{(i)} : 1 \leq i \leq d\}.$$

Let $g = c^{(j)} - a^{(j)}$. Then $\|g\|_\infty < 2(n/d)s$, by Condition (4) for $a^{(i)}$ and $c^{(i)}$, and $g = \alpha_1 a^{(1)} + \cdots + \alpha_{d-1} a^{(d-1)}$, for some $\alpha_i \in \mathbb{Z}$, by Condition (5) for $a^{(i)}$. We must show that $g = 0$. By reordering indices and possibly considering $-g$, we may assume that $\alpha_1, \dots, \alpha_k \geq 0$, for some k , $\alpha_{k+1}, \dots, \alpha_{d-1} \leq 0$, and $|\alpha_1|$ is maximal over all $|\alpha_i|$, $1 \leq i \leq d - 1$.

Examining the first coordinate of g , we have that

$$\begin{aligned} -2\frac{n}{d}s &< g_1 \\ &= \alpha_1 a_1^{(1)} + \cdots + \alpha_{d-1} a_1^{(d-1)} \\ &< \alpha_1 \frac{n}{d}(b_1^{(1)} + s) + \cdots + \alpha_k \frac{n}{d}(b_1^{(k)} + s) + \alpha_{k+1} \frac{n}{d}(b_1^{(k+1)} - s) + \cdots \\ &\quad + \alpha_{d-1} \frac{n}{d}(b_1^{(d-1)} - s) \\ &= \frac{n}{d}[\alpha_1 + \cdots + \alpha_{d-1} - \frac{1}{2}\alpha_1 + s(|\alpha_1| + \cdots + |\alpha_{d-1}|)] \\ &\hspace{15em} \text{(using } b^{(i)} = (1, \dots, 1) - \frac{1}{2}e_i) \\ &\leq \frac{n}{d}[\alpha_1 + \cdots + \alpha_{d-1} - \frac{1}{2}\alpha_1 + (d-1)s\alpha_1]. \end{aligned}$$

Negating and dividing by $\frac{n}{d}$,

$$-(\alpha_1 + \dots + \alpha_{d-1}) + \frac{1}{2}\alpha_1 - (d-1)s\alpha_1 < 2s. \tag{6}$$

Similarly, examining the $(k+1)$ -st coordinate of g , we have

$$\begin{aligned} 2\frac{n}{d}s &> g_{k+1} \\ &= \alpha_1 a_{k+1}^{(1)} + \dots + \alpha_{d-1} a_{k+1}^{(d-1)} \\ &> \alpha_1 \frac{n}{d}(b_{k+1}^{(1)} - s) + \dots + \alpha_k \frac{n}{d}(b_{k+1}^{(k)} - s) + \alpha_{k+1} \frac{n}{d}(b_{k+1}^{(k+1)} + s) + \dots \\ &\quad + \alpha_{d-1} \frac{n}{d}(b_{k+1}^{(d-1)} + s) \\ &= \frac{n}{d}[\alpha_1 + \dots + \alpha_{d-1} - \frac{1}{2}\alpha_{k+1} - s(|\alpha_1| + \dots + |\alpha_{d-1}|)] \\ &\geq \frac{n}{d}[\alpha_1 + \dots + \alpha_{d-1} - \frac{1}{2}\alpha_{k+1} - (d-1)s\alpha_1], \end{aligned}$$

and so

$$(\alpha_1 + \dots + \alpha_{d-1}) - \frac{1}{2}\alpha_{k+1} - (d-1)s\alpha_1 < 2s. \tag{7}$$

Adding the equations (6) and (7),

$$\frac{1}{2}\alpha_1 - \frac{1}{2}\alpha_{k+1} - 2(d-1)s\alpha_1 < 4s,$$

and so, since $s = 1/(4d+4)$,

$$\frac{1}{d+1}\alpha_1 - \frac{1}{2}\alpha_{k+1} < \frac{1}{d+1}.$$

Therefore, since $\alpha_{k+1} \leq 0$, we have that $\alpha_1 < 1$ and so $\alpha_1 = 0$. Since $|\alpha_1|$ was maximal over all $|\alpha_i|$, we have that $g = 0$. Therefore $a^{(j)} = c^{(j)}$, and the lemma follows.

Proof of Lemma 14. Suppose $N > 2^{d(d-1)}$, and suppose $a^{(i,j)}$, for $1 \leq i \leq N$ and $1 \leq j \leq d-1$, are such that $a^{(i,j)} \in \mathbb{Z}_+^d$, $\sum_{k=1}^d a_k^{(i,j)} < n$, and the N rays

$$r^{(i)} = \bigcap_{j=1}^{d-1} \{x : \langle a^{(i,j)}, x \rangle = 0\}$$

are the extreme rays for some chamber. Then, since $N > 2^{d(d-1)}$, there are some i and i' such that $a_k^{(i,j)} \equiv a_k^{(i',j)} \pmod{2}$, for $1 \leq j \leq d-1$ and $1 \leq k \leq d$ (i.e., all of the coordinates in all of the vectors have the same parity). Then let

$$c^{(j)} = \frac{a^{(i,j)} + a^{(i',j)}}{2},$$

for $1 \leq j \leq d-1$. Then $c^{(j)} \in \mathbb{Z}_+^d$ and $\sum_{k=1}^d c_k^{(j)} < n$, and the ray

$$r = \bigcap_{j=1}^{d-1} \{x : \langle c^{(j)}, x \rangle = 0\} = \frac{r^{(i)} + r^{(i')}}{2}$$

is in the chamber, which is a contradiction.

5. Final remarks

An interpretation of Theorem 1 is that the ability to change the values of the parameters of a graphical model does not give as much freedom as it may appear. There is a very large number of possible ways to assign an explanation to each observation. However, only a tiny proportion of these come from a consistent method for choosing the most probable explanation for a certain choice of parameters. Even though the parameters can vary continuously, the number of different inference functions that can be obtained is at most polynomial in the number of edges of the model, assuming that the number of parameters is fixed.

Having shown that the number of inference functions of a graphical model is polynomial in the size of the model, an interesting next step would be to find an efficient way to precompute all the inference functions for given models. This would allow us to give the answer (the explanation) to a query (an observation) very quickly. It follows from this paper that it is computationally feasible to precompute the polytope $\text{NP}(\mathbf{f})$, whose vertices correspond to the inference functions. However, the difficulty arises when we try to describe a particular inference function efficiently. The problem is that the characterization of an inference function involves an exponential number of observations.

Acknowledgements

The authors are grateful to Graham Denham, Lior Pachter, Carl Pomerance, Bernd Sturmfels, and Ravi Kannan for helpful discussions. They would also like to thank the anonymous referees for their valuable comments. The first author was partially supported by the J. William Fulbright Association of Spanish Fulbright Alumni.

References

- Apostol T. M. (1976). *Introduction to Analytic Number Theory*. Springer-Verlag, New York.
- Elizalde S. (2005). Inference functions. In *Algebraic Statistics for Computational Biology* (Pachter and Sturmfels, editors), Chapter 9, Cambridge University Press.
- Elizalde S. and Woods K. (2006). The probability of choosing primitive sets, *J. Number Theory*, to appear, <http://dx.doi.org/10.1016/j.jnt.2006.11.001>.

- Gritzmann, P. and Sturmfels, B. (1993). Minkowski addition of polytopes: Computational complexity and applications to Gröbner bases. *SIAM J. Discrete Math.* **6**, 246-269.
- Jensen, F. (2001). *Bayesian Networks and Decision Graphs*. Springer.
- Lekkerkerker, C. G. (1969). *Geometry of Numbers*. Wolters-Noordhoff, Groningen.
- Nymann, J. E. (1972). On the probability that k positive integers are relatively prime. *J. Number Theory* **4**, 469-473.
- Pachter, L. and Sturmfels, B. (2004). Tropical Geometry of Statistical Models. *Proc. Natl. Acad. Sci.* **101**, 16132-16137.
- Pachter, L. and Sturmfels, B. (2004). Parametric Inference for Biological Sequence Analysis. *Proc. Natl. Acad. Sci.* **101**, 16138-16143.
- Pachter, L. and Sturmfels, B. (2005). *Algebraic Statistics for Computational Biology*. Cambridge University Press.
- Pachter, L. and Sturmfels, B. (2006). The Mathematics of Phylogenomics. *SIAM Review*, in press.
- Ziegler, G. M. (1995). *Lectures on Polytopes*. Graduate Texts in Mathematics **152**, Springer, New York.

Department of Mathematics, Dartmouth College, Hanover, NH 03755, U.S.A.

E-mail: sergi.elizalde@dartmouth.edu

Department of Mathematics, Oberlin College, Oberlin, OH 44074, U.S.A.

E-mail: Kevin.Woods@oberlin.edu

(Received June 2006; accepted November 2006)