

Approximating faces of marginal polytopes in discrete hierarchical models

Nanwei Wang*, Johannes Rauh† and H el ene Massam‡

April 29, 2016

Abstract

The existence of the maximum likelihood estimate (abbreviated mle) in hierarchical loglinear models has important consequences for inference. Determining whether this estimate exists is equivalent to finding whether the data belongs to the boundary of the marginal polytope of the model. For higher-dimensional problems, it may not be possible to solve this problem exactly. Massam and Wang (2015) found an outer approximation to the smallest face, that is a face containing the smallest face. The purpose of this paper is to refine the methodology to find this outer approximation and to devise a new methodology to find an inner approximation. The inner approximation is given not in terms of a face of the marginal polytope, but in terms of a subset of the vertices of the smallest face. While outer approximations can be found by looking at submodels of the original hierarchical model, inner approximations correspond to larger models. To obtain larger models that are easier to understand than the original model, we suggest to complete separators in order to obtain a reducible model.

Using real-world data and simulated data we illustrate that the two approximations often agree, and in this case we have found the smallest face. Even if the two approximations do not agree, they

*Department of Statistics, York University, Toronto, M3J 1P3, Canada

†Department of Statistics, York University, Toronto, M3J 1P3, Canada

‡Department of Statistics, York University, Toronto, M3J 1P3, Canada. This author gratefully acknowledges support from NSERC Discovery Grant A8947.

provide us with information that can be used to better estimate the parameters of the mle (or the extended mle whenever the mle does not exist).

Keywords: existence of the maximum likelihood estimate, marginal polytope, faces, facial sets, extended maximum likelihood estimate.

1 Introduction

Discrete hierarchical models are an essential tool for the analysis of categorical data given under the form of a contingency table. The study of these models goes back more than a century, and a detailed history of their development is given in [Fienberg and Rinaldo \(2007\)](#). Nowadays, discrete hierarchical models are used for the analysis of large sparse contingency tables where many, if not most, of the entries are small or zero counts. It is well-known that in such cases, the maximum likelihood estimate (henceforth abbreviated mle) of the parameters may not exist. The non existence of the mle has problematic consequences for inference, clearly for estimation, but also for testing and model selection. The reader is referred to [Fienberg and Rinaldo \(2012\)](#) for a list of statistical implications of the non existence of the mle. The most serious of these implications is that, when the mle does not exist, the degrees of freedom used to approximate various measures of fit are incorrect. [Fienberg and Rinaldo \(2012\)](#) also give necessary and sufficient conditions for the existence of the mle. [Geyer \(2009\)](#) describes the problems attached to the nonexistence of the mle and presents an R program that yields meaningful confidence intervals and tests. [Letac and Massam \(2012\)](#) study the statistical implications of the nonexistence of the mle on model selection in Bayesian inference.

The necessary and sufficient conditions for the existence of the mle given by [Fienberg and Rinaldo \(2012\)](#) are extensions of those given earlier by [Haberman \(1974\)](#), [Barndorff-Nielsen \(1978\)](#) and [Eriksson et al. \(2006\)](#) and are essentially as follows. Let V be a finite index set, and let $X = (X_v, v \in V)$ be a vector of discrete random variables representing $|V|$ characteristics of an object. We consider N objects classified in a $|V|$ -dimensional contingency table according to these characteristics. Let I denote the set of cells in the contingency table and let $n(i)$, $i \in I$, be the number of observations falling

in cell $i \in I$. We assume that $(n(i), i \in I)$ follows a multinomial distribution and that the cell probabilities are modelled according to a hierarchical model. As we recall in Section 2, this latter condition implies that the expression $\prod_{i \in I} p(i)^{n(i)}$ in the multinomial density can be written as the density of an exponential family distribution

$$f(t)dt = \exp\{\langle \theta, t \rangle - Nk(\theta)\} \nu(dt), \quad (1)$$

where θ is a vector of parameters and t is the vector of sufficient statistic determined by the hierarchical model and $\nu(dt)$ is a discrete measure with support on a finite number of vectors $f_i, i \in I$. The domain of the means of this natural exponential family is the convex hull of the support of ν and forms a convex polytope \mathbf{P} with extreme points $f_i, i \in I$, called the marginal polytope (because t is a vector of marginal counts of subsets of the random variables X_v). Then the mle exists if and only if t belongs to the relative interior of the marginal polytope \mathbf{P} . If the mle does not exist, then t belongs to the relative interior of a proper face \mathbf{F}_t of \mathbf{P} . Thus, determining whether, for a given data set, the mle of a discrete hierarchical loglinear model exists is equivalent to determining whether t belongs to a proper face of \mathbf{P} .

To the best of our knowledge, the papers that have tackled this topic so far are Eriksson et al. (2006), Geyer (2009), Fienberg and Rinaldo (2012) and Massam and Wang (2015). The first three seek to identify \mathbf{F}_t in order to compute the extended mle as defined by Barndorff-Nielsen (1978) or Lauritzen (1996). To do so, they solve a sequence of linear programming problems aimed at finding the f_i in \mathbf{F}_t . Unfortunately this method becomes computationally infeasible in large dimensions. In our experience, it is not possible to use the linear programming approach when the number of nodes is larger than 16. Massam and Wang (2015) show how to find an approximation \mathbf{F}_2 to \mathbf{F}_t . They consider subsets $V_i, i = 1, \dots, k$, of V containing less than 16 nodes and compute the facial set of the smallest face containing the corresponding sufficient statistic in each of the corresponding models using linear programming. They prove that the intersection of the smallest faces corresponding to $V_i, i = 1, \dots, k$, is a face \mathbf{F}_2 of \mathbf{P} containing \mathbf{F}_t . The face \mathbf{F}_2 gives us some information in the sense that if \mathbf{F}_2 is a proper face of \mathbf{P} , then \mathbf{F}_t is a proper face and therefore the mle does not exist. We call \mathbf{F}_2 an outer approximation to \mathbf{F}_t .

The purpose of this paper is to add to this outer approximation \mathbf{F}_2 , an inner approximation \mathbf{F}_1 that is included in the face \mathbf{F}_t and to explore in detail

the most efficient ways to obtain an inner and an outer approximations to \mathbf{F}_t . The inner approximation \mathbf{F}_1 we find is not a face of the marginal polytope \mathbf{P} , but it is a face of the marginal polytope of a model containing the original model. We have the relationship

$$\mathbf{F}_1 \subset \mathbf{F}_t \subset \mathbf{F}_2.$$

Clearly, if $\mathbf{F}_1 = \mathbf{F}_2$, then we have found \mathbf{F}_t . In this paper, we will present our methodology for finding \mathbf{F}_1 and \mathbf{F}_2 in a unified way.

We now sketch the main lines of our argument. We will work in terms of facial sets, that is, for a given face \mathbf{F} , in terms of $F = \{i \in I \mid f_i \in \mathbf{F}\}$ rather than in terms of \mathbf{F} . A hierarchical model for the discrete random variable $X = (X_v, v \in V)$ is determined by the set of possible interactions between the variables $X_v, v \in V$. This set of interactions is represented by a simplicial complex Δ , that is, Δ is a set of subsets $D \subset V$ such that $D \in \Delta$ and $D' \subset D$ implies $D' \in \Delta$. For any $T \subseteq I$, denote by $F_\Delta(T)$ the smallest facial set containing T . Thus $F_\Delta(T)$ is the facial set of the smallest face of the marginal polytope \mathbf{P}_Δ of the hierarchical model given by Δ which contains $f_i, i \in T$.

The data is given by the cell counts $(n(i), i \in I)$. We let $I_+ = \{i \in I \mid n(i) > 0\}$ denote the subset of I indexing the cells with positive counts. As we shall see in Section 2, the vector t in (1) is

$$t = \sum_{i \in I} \frac{n(i)}{N} f_i = \sum_{i \in I_+} \frac{n(i)}{N} f_i,$$

where $N = \sum_{i \in I} n(i)$. Thus \mathbf{F}_t is the smallest face of \mathbf{P}_Δ containing $\{f_i, i \in I_+\}$, and so its facial set is $F_t = F_\Delta(I_+)$. Identifying \mathbf{F}_t is therefore equivalent to identifying $F_\Delta(I_+)$.

To find inner and outer approximations to $F_\Delta(I_+) := F_t$, we rely on the following facts which will be proved in Section 3. If $\Delta_1 \subset \Delta_2$ are two simplicial complexes on the same vertex set V , then

$$F_{\Delta_2}(T) \subseteq F_{\Delta_1}(T). \tag{2}$$

From (2), any sub-complex Δ_1 of our original simplicial complex Δ yields an outer approximation to F_t . We can improve this outer approximation by taking the intersection of several such approximations. Namely, given $\Delta_{1,1}, \Delta_{1,2}, \dots, \Delta_{1,r} \subseteq \Delta$, then $F_\Delta(I_+) \subseteq F_{\Delta_{1,k}}(I_+)$ for all k , and so $F_\Delta(I_+) \subseteq \bigcap_{k=1}^r F_{\Delta_{1,k}}(I_+)$.

Similarly, for every simplicial complex Δ_2 that contains the original simplicial complex Δ , by (2) we obtain an inner approximation $F_{\Delta_2}(I_+) \subseteq F_{\Delta}(I_+)$. We can improve this approximation by looking at several simplicial complexes in turn and iterating. Namely, if $\Delta \subseteq \Delta_{2,1}$ and $\Delta \subseteq \Delta_{2,2}$, then $F_{\Delta_{2,1}}(I_+) \subseteq F_{\Delta}(I_+) = F_t$, and thus $F_{\Delta_{2,2}}(F_{\Delta_{2,1}}(I_+)) \subseteq F_{\Delta}(F_t) = F_t$.

Moreover, for any simplicial complex Γ on V and any $T \subset I$, we have $F_{\Gamma}(T) \supset T$ by the very definition of a facial set. Applying this inclusion to $T = F_{\Delta_{2,1}}(I_+)$ and $\Gamma = \Delta_{2,2}$ yields $I_+ \subseteq F_{\Delta_{2,1}}(I_+) \subseteq F_{\Delta_{2,2}}(F_{\Delta_{2,1}}(I_+)) \subseteq F_t$, that is, $F_{\Delta_{2,2}}(F_{\Delta_{2,1}}(I_+))$ is a better approximation of F_t than $F_{\Delta_{2,1}}(I_+)$. We can further improve our approximation by applying $F_{\Delta_{2,1}}$ once more: in general, $F_{\Delta_{2,1}}(F_{\Delta_{2,2}}(F_{\Delta_{2,1}}(I_+)))$ will again be a superset of $F_{\Delta_{2,2}}(F_{\Delta_{2,1}}(I_+))$ contained in F_t . We carry on this iteration until the sets do not increase any more and we arrive at a set F_1 that satisfies $I_+ \subseteq F_1 \subseteq F_t$; that is, F_1 is an inner approximation of F_t . F_1 is not necessarily facial for Δ (unless $F_1 = F_t$), but F_1 is facial for $\Delta_{2,1}$ and $\Delta_{2,2}$. This procedure easily generalizes to the case of more than two simplicial complexes containing Δ .

The various approximating facial sets will be computed using the linear programming technique of [Fienberg and Rinaldo \(2012\)](#) and [Massam and Wang \(2015\)](#). The question that remains is how to find suitable simplicial complexes that are contained in Δ or that contain Δ and for which it is easy to compute faces.

To obtain a simplicial complex containing Δ , we complete separators. That is, we find sets $V_1, V_2 \subseteq V$ such that Δ can be written as a union $\Delta_1 \cup \Delta_2$ of simplicial complexes Δ_1 on V_1 and Δ_2 on V_2 . We call $S := V_1 \cap V_2$ a *separator*. If $S \in \Delta$, then the hierarchical model is *reducible*, and facial sets of Δ can easily be computed by computing facial sets with respect to the hierarchical models of Δ_1 and Δ_2 (see Lemma 3.6 below). If $S \notin \Delta$, we can use $\Delta_S = \Delta \cup \{M : M \subseteq S\}$ as a simplicial complex containing Δ to find an inner approximation to F_t .

To find an outer approximation, we show that the approach of [Massam and Wang \(2015\)](#) of looking at a small subset of nodes can be seen as a special case of looking at a simplicial sub-complex.

Our results apply to not only hierarchical models, but also more general discrete exponential families. While all of our examples are hierarchical models, our theoretical results are best understood from a more general point of view.

The remainder of this paper is organized as follows. In Section 2, we give

preliminaries on hierarchical models, and faces and facial sets. In Section 3, we develop our tools for the inner and outer approximation F_1 and F_2 to F_t . In Section 4, we show how to use F_1 and F_2 to identify the parameters of the hierarchical models that can be estimated and those that cannot be estimated. This extends in some ways the work of Fienberg and Rinaldo (2012) and that of Geyer (2009). In Section 5, we illustrate our methodology with simulated data and a real world example using the NLTCs data. This data set has been studied in Dobra and Lenkoski (2011) and Dobra et al. (2003). Both of these examples have 16 nodes. In Section 6, we apply our methodology to larger models, one with simulated data and the other on a real-world data set. Our simulated data set is obtained from a graphical model of the 5×10 grid, while the real-world data set uses the voting records in the US Senate during the portion of 2015 year that was available at the time of our study. A similar data set but for a different year was used in Banerjee et al. (2008).

2 Preliminaries

In the following four subsections, we recall some basic facts about hierarchical models, discrete exponential families, polytopes and the closure of exponential families, and we also define the extended mle.

2.1 Hierarchical models

For details and proofs on the material in this subsection, we refer to Letac and Massam (2012) and Rauh et al. (2011). Let $X = (X_v, v \in V)$ be a discrete random vector with components indexed by $V = \{1, \dots, p\}$, a finite set. Each variable X_v takes values in a finite set $I_v, v \in V$. The vector X takes its values in

$$I = \prod_{v \in V} I_v,$$

the set of cells $i = (i_v, v \in V)$ of the p -dimensional contingency table. Let Δ be a set of subsets of V which is a simplicial complex, that is, Δ is a set of subsets $D \subset V$ such that $D \in \Delta$ and $D' \subset D$ implies $D' \in \Delta$. We say that the joint distribution of X is *hierarchical* with underlying simplicial complex Δ (or generating set Δ) if the probability $p(i) = P(X = i)$ of a single cell

$i = (i_v, v \in V)$ is of the form

$$\log(p(i)) = \sum_{D \in \Delta} \theta_D(i_D) \quad (3)$$

where $\theta_D(i_D)$ is a function of the marginal cell $i_D = (i_v, v \in D)$ only. If we need to make precise the dependence on θ , then we write $p_\theta(i)$ instead of $p(i)$. The set of all such distributions $\mathcal{E}_\Delta := \{p_\theta\}$ is called the *hierarchical model* of Δ .

The parametrization (3) is not identifiable; that is, for any joint distribution p from the hierarchical model there are different choices for the functions θ_D that satisfy (3). One way to make the parameters unique is to choose a special element within each set I_v , which we denote by 0. The choice of 0 is arbitrary, and a different choice of 0 leads to a simple affine change of parameters. With this choice, the functions θ_D become unique if one requires $\theta_D(i_D) = 0$ whenever $i_v = 0$ for some $v \in D$. Thus, we arrive at the identifiable parametrization

$$\log p_\theta(i) = \theta_0 + \sum_{D \in \Delta \setminus \{\emptyset\}, i_v \neq 0, \forall v \in D} \theta_D(i_D), \quad (4)$$

where $\theta_0 := \theta_\emptyset$. We separate the parameter θ_0 corresponding to the empty set, since it has a special role. It does not depend on the cell index i and acts as a normalizing constant: When all other parameters are chosen freely, θ_0 is determined by the requirement $\sum_{i \in I} p_\theta(i) = 1$. To make it clear that we consider θ_0 as a dependent parameter, we also write

$$-\theta_0 = k(\theta) = \log \left(\sum_{i \in I} \exp \left(\sum_{D \in \Delta \setminus \{\emptyset\}, i_v \neq 0, \forall v \in D} \theta_D(i_D) \right) \right).$$

The parametrization (4) can be further reformulated using the definitions

$$\begin{aligned} S(i) &= \{v \in V ; i_v \neq 0\} \\ J &= \{j \in I \setminus \{0\}, S(j) \in \Delta\}. \end{aligned}$$

For a given $D \in \Delta$ and for given $\theta_D(i_D)$ such that $i_\gamma \neq 0, \forall \gamma \in D$, there is only one $j \in J$ such that $S(j) = D$ and $j_D = j_{S(j)} = i_D$ and conversely. We can therefore write

$$\theta_D(i_D) = \theta_j \text{ for the unique } j \in J \text{ with } S(j) = D, i_D = j_D.$$

To simplify the notation, we write $j \triangleleft i$ whenever $S(j) \subseteq S(i)$ and $j_{S(j)} = i_{S(j)}$. Then the parametrization (4) in terms of the *free* parameters $\theta = \{\theta_j, j \in J\}$ becomes

$$\log p_\theta(i) = \sum_{j \in J: j \triangleleft i} \theta_j - k(\theta). \quad (5)$$

It is convenient to introduce the vectors

$$f_i = \sum_{j \in J: j \triangleleft i} e_j, \quad i \in I$$

where $e_j, j \in J$ are the unit vectors in R^J . Moreover, let A be the $J \times I$ matrix with columns $f_i, i \in I$, and let \tilde{A} be the $(1 + |J|) \times I$ matrix with columns equal to $\begin{pmatrix} 1 \\ f_i \end{pmatrix}, i \in I$. The representation (5) becomes

$$\log p_\theta(i) = \langle \theta, f_i \rangle - k(\theta) = A^t \theta - k(\theta) = \tilde{A}^t \tilde{\theta}, \quad (6)$$

where $\tilde{\theta} = (\theta_0, \theta)$ as a column vector. Both A and \tilde{A} are called *design matrices* of the model.

From the definition of $f_i, i \in I$, it follows immediately that if $n = (n(i), i \in I)$ denotes the I -dimensional column vector of cell counts, then

$$\tilde{A}n = \begin{pmatrix} N \\ t \end{pmatrix} \quad \text{and} \quad An = t, \quad (7)$$

where $N = \sum_{i \in I} n(i)$ is the total cell counts and t is the column vector of $j_{S(j)}$ -marginal counts $n(j_{S(j)})$, i.e. $t = (t_j, j \in J)$ where $t_j = n(j_{S(j)}) = \sum_{i | i_{S(j)} = j_{S(j)}} n(i), j \in J$.

It follows from (7) that $\frac{t}{N} = \sum_{i \in I} \frac{n(i)}{N} f_i$. Therefore, t belongs to the convex polytope with extreme points $f_i, i \in I$. This polytope is called the *marginal polytope* of the hierarchical model, and we denote it by \mathbf{P}_Δ .

Example 2.1. For the model defined by $V = \{a, b, c\}, I_a = \{0, 1\} = I_b = I_c$ and $\Delta = \{a, b, c, ab, bc\}$, we have $I = (000, 100, 010, 110, 001, 101, 011, 111)$ and $J = \{(100), (010), (001), (110), (011)\}$. Then

$$\tilde{A} = \begin{pmatrix} \overbrace{1}^{f_{000}} & \overbrace{1}^{f_{001}} & \overbrace{1}^{f_{010}} & \overbrace{1}^{f_{011}} & \overbrace{1}^{f_{100}} & \overbrace{1}^{f_{101}} & \overbrace{1}^{f_{110}} & \overbrace{1}^{f_{111}} \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \begin{matrix} \theta_{000} \\ \theta_{100} \\ \theta_{010} \\ \theta_{001} \\ \theta_{110} \\ \theta_{011} \end{matrix}$$

An important subclass of hierarchical model is the class of graphical models. Let $G = (V, E)$ be an undirected graph with vertex set V and edge set E . A subset $D \subseteq V$ is a clique of G if any $i, j \in D, i \neq j$, define an edge $(i, j) \in E$. The set of cliques of G , denoted by $\Delta(G)$, is a simplicial complex. The *graphical model* of G is defined as the hierarchical model of $\Delta(G)$. Graphical models are important because of their interpretation in terms of conditional independence, see [Lauritzen \(1996\)](#).

Another class of hierarchical models related to graphs is the class of Ising models. For an undirected graph $G = (V, E)$, the corresponding *Ising model* is the hierarchical model of the simplicial complex Δ that contains the singletons $\{x\}, x \in V$, and the pairs $\{x, y\}$ with $(x, y) \in E$.

In Sections [5.1](#) and [6.2](#), we will consider models with underlying graphs that are grid graphs, i.e., four-neighbour lattices. In such cases, the cliques are at most of size two, and the corresponding graphical models are the same as the corresponding Ising models. The hierarchical model for the NLTC data studied in Section [5.2](#) is a graphical model, while the hierarchical model used for the US Senate voting data studied in Section [6.1](#) is an Ising model.

2.2 Discrete exponential families

Hierarchical models are examples of discrete exponential families. Let I and J be finite sets, and let $A \in \mathbf{R}^{J \times I}$ be a real matrix. Denote the columns of A by $f_i, i \in I$. The discrete exponential family corresponding to A , denoted by \mathcal{E}_A , consists of all probability distributions on I that are of the form

$$p_\theta(i) = \exp(\langle \theta, f_i \rangle - k(\theta)), \quad \theta \in \mathbf{R}^J,$$

where $k(\theta) = \log \sum_{i \in I} \exp(\langle \theta, f_i \rangle)$. Define $\tilde{J} = J \cup \{0\}$, $\theta_0 = -k(\theta)$ and $\tilde{\theta} = (\theta_0, \theta)$, and let \tilde{A} be the matrix A with one additional row of ones; that is, $\tilde{A} \in \mathbf{R}^{\tilde{J} \times I}$ is the matrix with columns $\begin{pmatrix} 1 \\ f_i \end{pmatrix}, i \in I$. Then \mathcal{E}_A consists of the probability distributions p_θ that satisfy $\log p_\theta = \tilde{A}^t \tilde{\theta}$ for some $\tilde{\theta} \in \mathbf{R}^{\tilde{J}}$. The convex hull of the columns $f_i, i \in I$, is called the *convex support polytope*, denoted by \mathbf{P}_A . It generalizes the marginal polytope.

The parametrization $\theta \rightarrow p_\theta$ is identifiable if and only if \tilde{A} has full rank. If \tilde{A} does not have full rank, then one can drop certain rows of A to obtain a submatrix A' such that \tilde{A}' has full rank. This is equivalent to setting certain parameters to zero until the remaining parameters are identifiable.

Later, the following reparametrization will be useful: Select an element of I , denoted by 0 in the following. Let A_0 be the matrix with columns $f_i - f_0$, $i \in I \setminus \{0\}$. It is not difficult to see that A and A_0 define the same exponential family (since \tilde{A} and \tilde{A}_0 have the same row span). Let $h' = \text{rank}(A_0) = \text{rank}(\tilde{A}_0) - 1$, and select a set L of h' linearly independent vectors among the columns of A_0 . For $i \in L$ let $\mu_i = \mu_i(\theta) := \langle \theta, f_i - f_0 \rangle$, and let $\mu_L = (\mu_i, i \in L)$. It is not difficult to see that the μ_L are identifiable parameters on \mathcal{E}_A : In fact, their number is equal to h' , and they are independent by construction.

It is possible to extend the definition of $\mu_i(\theta)$ to all $i \in I$. Note that only the parameters μ_i with $i \in L$ are free parameters, while the parameters μ_i with $i \in I \setminus L$ are linear functions of μ_L .

The parameters μ_i can be interpreted as log-likelihood ratios:

$$\mu_i(\theta) = \log \frac{p_\theta(i)}{p_\theta(0)}, \quad \mu_0(\theta) = 0, \quad p_\theta(0) = \frac{1}{\sum_{i \in I \setminus \{0\}} \exp(\mu_i)}.$$

For a discrete exponential family, $\prod_{i \in I} p_\theta(i)^{n(i)}$ can be written under the form of a natural exponential family as given in (1). Indeed,

$$\begin{aligned} \prod_{i \in I} p_\theta(i)^{n(i)} &= \exp\left(\sum_{i \in I} n(i) \log p_\theta(i)\right) = \exp(\langle n, \log(p_\theta) \rangle) = \exp(\langle n, \tilde{A}^t \tilde{\theta} \rangle) \\ &= \exp(\langle \tilde{A}n, \tilde{\theta} \rangle) = \exp\left(\sum_{j \in J} \theta_j t_j - Nk(\theta)\right). \end{aligned}$$

The log-likelihood function in θ for the loglinear parameter of the multinomial distribution with the given hierarchical model is therefore

$$l(\theta) = \sum_{j \in J} \theta_j t_j - Nk(\theta). \quad (8)$$

It is well-known that $l(\theta)$ is concave. If the parameters are identifiable, then it is strictly concave.

We can also express the log-likelihood as a function of $\mu = (\mu_i, i \in I)$:

$$\begin{aligned} l(\mu) &= \sum_{i \in I} n(i) \log p(i) = \sum_{i \in I \setminus \{0\}} n(i) \log \frac{p(i)}{p(0)} + N \log p(0) \\ &= \sum_{i \in I} n(i) \mu_i - N \log\left(\sum_{i \in I} \exp \mu_i\right). \quad (9) \end{aligned}$$

As stated before, only a subset μ_L of the parameters μ are independent, and the remaining μ_i , $i \notin L$, can be expressed as linear functions of μ_L .

2.3 Polytopes

We next recall some general facts about polytopes and their faces. We refer to [Ziegler \(1998\)](#) for details and more information.

Definition 2.2. A set $\mathbf{P} \subset \mathbf{R}^h$ is a (*convex*) *polytope* if \mathbf{P} is the convex hull of a finite subset of \mathbf{R}^h . Equivalently, a polytope can be defined as a bounded subset of \mathbf{R}^h defined by linear inequalities.

Definition 2.3. For any vector $g \in \mathbf{R}^h$ and any constant $c \in \mathbf{R}$, define three sets $H_{g,c} = \{x \in \mathbf{R}^h : \langle g, x \rangle = c\}$, $H_{g,c}^+ = \{x \in \mathbf{R}^h : \langle g, x \rangle \geq c\}$ and $H_{g,c}^- = \{x \in \mathbf{R}^h : \langle g, x \rangle \leq c\}$. If $g \neq 0$, then $H_{g,c}$ is an (*affine*) *hyperplane*, and $H_{g,c}^+$ and $H_{g,c}^-$ are the *positive* and *negative halfspace* defined by g and c .

Let $\mathbf{P} \subseteq \mathbf{R}^h$ be a polytope, let $g \in \mathbf{R}^h$ and $c \in \mathbf{R}$, and suppose that $\mathbf{P} \subset H_{g,c}^+$ or $\mathbf{P} \subset H_{g,c}^-$. Then $\mathbf{F} := H_{g,c} \cap \mathbf{P}$ is called a *face* of \mathbf{P} . If $g \neq 0$, then $H_{g,c}$ is called a *supporting hyperplane* of \mathbf{P} . If $\mathbf{F} \neq \mathbf{P}$ and $\mathbf{F} \neq \emptyset$, then \mathbf{F} is a *proper face* of \mathbf{P} .

The dimension of a face \mathbf{F} is the dimension of the smallest affine subspace of \mathbf{R}^h that contains it. Its co-dimension is $\dim(\mathbf{P}) - \dim(\mathbf{F})$. A *facet* of a polytope \mathbf{P} is a proper face that is maximal with respect to inclusion and is thus of co-dimension 1. A minimal proper face of a polytope is a singleton $\{p\} \subseteq \mathbf{P}$; in this case, p is a *vertex*.

Intersections of faces are again faces: If $g_1, g_2 \in \mathbf{R}^h$ and $c_1, c_2 \in \mathbf{R}$ define faces $\mathbf{F}_1, \mathbf{F}_2$ of \mathbf{P} and if $\mathbf{P} \subset H_{g_1, c_1}^+ \cap H_{g_2, c_2}^+$, then $\mathbf{P} \subset H_{g_1+g_2, c_1+c_2}^+$, and $\mathbf{F}_1 \cap \mathbf{F}_2 = \mathbf{P} \cap H_{g_1+g_2, c_1+c_2}$. Any face is an intersection of facets.

By definition, every face \mathbf{F} of a polytope $\mathbf{P} \subset \mathbf{R}^h$ is characterized by a linear inequality $\langle g, x \rangle \geq c$ that is valid on \mathbf{P} and that holds as an equality on \mathbf{F} . This linear inequality is unique only if \mathbf{F} is a facet. Sometimes it is convenient to give all linear equations that hold on a face \mathbf{F} . These linear equations determine the smallest affine subspace of \mathbf{R}^h containing \mathbf{F} .

When a polytope is defined as the convex hull of a finite number of points f_i , $i \in I$, then it is of interest to know, which subsets of $\{f_i\}_{i \in I}$ lie on a common face. Indeed, it is the purpose of this paper to compute the smallest face of the marginal polytope containing the data vector t , and we will determine this face by identifying which vectors f_i belong to it.

Definition 2.4. For a finite set I let $\{f_i\}_{i \in I} \subset \mathbf{R}^h$, and let \mathbf{P} be the convex hull of $\{f_i\}_{i \in I}$. A subset $F \subseteq I$ is called *facial* (with respect to \mathbf{P}), if there exists a face \mathbf{F} of \mathbf{P} with $F = \{i : f_i \in \mathbf{F}\}$. For any subset $S \subseteq I$, denote by $F_{\mathbf{P}}(S)$ the smallest facial set that contains S .

Since the intersection of facial sets is again facial, $F_{\mathbf{P}}(S)$ is well-defined.

Lemma 2.5. Let $\{f_i\}_{i \in I} \subset \mathbf{R}^h$, let $\phi : \mathbf{R}^h \rightarrow \mathbf{R}^{h'}$, $x \mapsto Bx + d$ be an affine map, and let $f'_i = \phi(f_i)$. If \mathbf{P} is the convex hull of the f_i , then $\mathbf{P}' := \phi(\mathbf{P})$ is the convex hull of the f'_i . The faces and facial sets of \mathbf{P} and \mathbf{P}' are related as follows:

1. Any inequality $\langle g', x' \rangle \geq c'$ that is valid on \mathbf{P}' corresponds to an inequality $\langle g, x \rangle \geq c$ that is valid on \mathbf{P} , where $g = B^t g'$ and $c = c' - \langle g', d \rangle$. Thus, if \mathbf{F}' is a face of \mathbf{P}' , then $\phi^{-1}(\mathbf{F}')$ is a face of \mathbf{P} .
2. A subset of I that is facial with respect to \mathbf{P}' is also facial with respect to \mathbf{P} . Thus, $F_{\mathbf{P}}(S) \subseteq F_{\mathbf{P}'}(S)$ for any $S \subseteq I$.

Proof. The first statement follows from

$$c \leq \langle g', \phi(f_i) \rangle = \langle g', Bf_i + d \rangle = \langle B^t g', f_i \rangle + \langle g', d \rangle,$$

which holds for any $i \in I$. The second statement follows immediately from the equation above and the fact that $F_{\mathbf{P}}(S)$ is the smallest facial set containing S . \square

We note that in Lemma 2.5, the dimension of $\phi(\mathbf{P})$ is at most equal to h . We will only apply Lemma 2.5 to coordinate projections ϕ with $h' < h$.

Remark 2.6. Sometimes it is convenient to embed the polytope in a vector space that has one additional dimension using a map $\mathbf{R}^h \rightarrow \mathbf{R}^{h+1}$, $x \mapsto \tilde{x} := (1, x)$. This has the advantage that all defining inequalities can be brought into a homogeneous form with vanishing constant c : Note that $\langle g, f_i \rangle - c = \langle \tilde{g}_c, \tilde{f}_i \rangle$, where $\tilde{g}_c := (c, g)$.

When a defining inequality of a face \mathbf{F} is given, its facial set F can be obtained by checking whether $f_i \in \mathbf{F}$ for each $i \in I$. In the other direction, when a facial set F is given, it is much more difficult to compute a defining inequality of the corresponding face \mathbf{F} . However, it is straightforward to compute the linear equations defining \mathbf{F} : The set of such equations $0 = \langle g, x \rangle - c = \langle \tilde{g}, \tilde{x} \rangle$ corresponds to the set of vectors $\tilde{g} \in \ker \tilde{A}_F^t$, where \tilde{A}_F is the matrix obtained from A by adding a row of ones and dropping the columns not in F .

2.4 The closure of an exponential family and existence of the mle

We fix a discrete exponential family \mathcal{E}_A . While our main interest lies in hierarchical models, the results that we need are more naturally formulated in the language of discrete exponential families. We assume that a vector of observed counts $n = (n(i) : i \in I)$ is given.

Definition 2.7. A parameter value θ^* is a *maximum likelihood estimate* (mle) if it is a global maximum of $l(\theta)$.

The function $l(\theta)$ is always bounded (clearly, it is never positive). As stated above, $l(\theta)$ is strictly concave (if the parameters are identifiable), and so the maximum is unique (up to identifiability), if it exists. However, a maximum need not exist, since the domain of the parameters θ is unbounded. To understand this, it is convenient to interpret the likelihood as a function of probabilities. Let \tilde{l} be the function that assigns to any probability distribution p on I the value

$$\tilde{l}(p) = \log\left(\prod_{i \in I} p(i)^{n(i)}\right)$$

Then $l(\theta) = \tilde{l}(p_\theta)$, and θ^* is an mle if and only if p_{θ^*} maximizes \tilde{l} subject to the constraint that p belongs to the hierarchical model (and thus is of the form p_θ for some θ). While the set of all probability distributions on I is compact, the hierarchical model itself is not closed and therefore not compact, and so there is no guarantee that \tilde{l} attains its maximum on the hierarchical model. However, things become better when we pass from the hierarchical model to its topological closure, where the topology comes from interpreting a probability distribution as a vector $p = (p(i))_{i \in I} \in \mathbf{R}^I$ of real numbers (this choice of the topology is canonical since we are dealing with a finite set I ; for infinite sample spaces see [Csiszár and Matúš \(2005\)](#)). The closure is sometimes also called *completion* ([Barndorff-Nielsen, 1978](#), p. 154). Since the closure of the hierarchical model is again compact, the continuous function \tilde{l} always attains its maximum.

Theorem 2.8. *The closure of a discrete exponential family can be written as a union*

$$\overline{\mathcal{E}_A} = \bigcup_F \mathcal{E}_{F,A},$$

where F runs over all facial sets of the convex support polytope \mathbf{P}_A and where $\mathcal{E}_{F,A}$ consists of all probability distributions of the form $p_{F,\theta}$, with

$$p_{F,\theta} = \begin{cases} \exp(\langle \theta, f_i \rangle - k_F(\theta)), & \text{if } i \in F, \\ 0, & \text{otherwise,} \end{cases}$$

where $k_F(\theta) = \log \sum_{i \in F} \exp(\langle \theta, f_i \rangle)$.

Proof. See [Barndorff-Nielsen \(1978\)](#). For self-containedness we provide a proof in our notation in [Appendix A.1](#). \square

[Theorem 2.8](#) shows that $\overline{\mathcal{E}}_A$ is a finite union of sets $\mathcal{E}_{F,A}$ that are exponential families themselves with a very similar parametrization, using the same number of parameters and the same design matrix A (or, rather, the submatrix A_F consisting of those columns of A indexed by F). However, for any proper facial set F , the parametrization $\theta \mapsto p_{F,\theta}$ is not injective, i.e. the parameters θ are not identifiable on $\mathcal{E}_{F,\Delta}$. The reason is that the matrix \tilde{A}_F does not have full rank, even if \tilde{A} has full rank, since all columns of \tilde{A}_F lie on a supporting hyperplane defining F .

A second thing to note is that although the parameters θ on \mathcal{E}_A and the parameters θ on $\mathcal{E}_{F,A}$ play similar roles, they are very different in the following sense: If $\theta^{(s)}$ is a sequence of parameters with $p_{\theta^{(s)}} \rightarrow p_{F,\theta}$ for some θ , then, in general, $\lim_{s \rightarrow \infty} \theta_j^{(s)} \neq \theta_j$ for all $j \in J$.

Theorem 2.9. *For any vector of observed counts n , there is a unique maximum p^* of \tilde{l} in $\overline{\mathcal{E}}_A$. For t as defined in [\(7\)](#), this maximum p^* satisfies:*

- $Ap^* = \frac{t}{N}$.
- $\text{supp}(p^*) = F_t$.

Proof. See [Barndorff-Nielsen \(1978\)](#). For self-containedness we provide a proof in our notation in [Appendix A.2](#). \square

Definition 2.10. The maximum in [Theorem 2.9](#) is called the *extended* maximum likelihood estimate (EMLE).

Clearly, if the mle θ^* exists, then $p^* = p_{\theta^*}$.

3 Approximations of Facial sets

We consider a hierarchical model with simplicial complex Δ and marginal polytope \mathbf{P}_Δ . In this section, we develop the details of our methodology to obtain an inner and an outer approximation to the facial set F_t of the smallest face \mathbf{F}_t of \mathbf{P}_Δ containing the data vector t . Our main tool is Lemma 3.1, which is an application of Lemma 2.5 to hierarchical models of simplicial complexes that are contained in each other. For any $S \subseteq I$, we abbreviate the facial set $F_{\mathbf{P}_\Delta}(S)$ by $F_\Delta(S)$.

Lemma 3.1. *Let Δ and Δ' be simplicial complexes on the same vertex set with $\Delta' \subseteq \Delta$, and denote by f_i, f'_i ($i \in I$) the columns of the design matrices of the corresponding hierarchical models. Then there is a linear map $\phi : \mathbf{R}^h \rightarrow \mathbf{R}^{h'}$ with $\phi(f_i) = f'_i$. In fact, ϕ is a coordinate projection. In particular, the marginal polytope $\mathbf{P}_{\Delta'}$ is a coordinate projection of \mathbf{P}_Δ . Thus, for any $S \subseteq I$, we have $F_\Delta(S) \subseteq F_{\Delta'}(S)$.*

Proof. The matrix A_Δ has one row for each parameter θ_j , $j \in J_\Delta$. Removing sets from Δ leads to a smaller set $J_{\Delta'}$ and thus leads to a matrix $A_{\Delta'}$ with less rows. The definition of each row that remains does not change. The lemma now clearly follows from Lemma 2.5. \square

Next we discuss marginal polytopes of decomposable (or reducible) models. Then, in Sections 3.2 and 3.3, we explain how to use Lemma 3.1 to obtain inner and outer approximations to $F_\Delta(S)$.

3.1 Decomposable models

Definition 3.2. Let $V' \subset V$. The *restriction* or *induced sub-complex* is $\Delta|_{V'} = \{S \in \Delta \mid S \subseteq V'\}$. The sub-complex $\Delta|_{V'}$ is *complete*, if $\Delta|_{V'}$ contains V' (and thus all subsets of V'). For brevity, in this case we say that V' is *complete* in Δ .

Definition 3.3. A subset $S \subset V$ is a *separator* of Δ if there exist $V_1, V_2 \subset V$ with $V_1 \cap V_2 = S$, $\Delta = \Delta|_{V_1} \cup \Delta|_{V_2}$ and $V_1 \neq S \neq V_2$. A simplicial complex that has a complete separator is called *reducible*. By extension, we also call the hierarchical model reducible.

Definition 3.4. A hierarchical model is *decomposable* if Δ can be written as a union $\Delta = \Delta_1 \cup \Delta_2 \cup \dots \cup \Delta_r$ of induced sub-complexes $\Delta_i = \Delta|_{V_i}$ in such a way that

1. each Δ_i is a complete simplex: $\Delta_i = \{S \subseteq V_i\}$; and
2. $(\Delta_1 \cup \dots \cup \Delta_i) \cap \Delta_{i+1}$ is a complete simplex.

In other words, Δ arises by iteratively gluing simplices along complete sub-simplices.

Faces of a reducible hierarchical model are combinations of the faces of its two parts:

Proposition 3.5. *Suppose that Δ has a complete separator S that separates V into V_1 and V_2 . Each face of $\mathbf{P}_{\Delta|_{V_1}}$ corresponds to an inequality*

$$\sum_{j \in J_{\Delta|_{V_1}}} g_j^{(1)} t_j \geq c_1.$$

The same inequality also defines a face of \mathbf{P}_{Δ} . Similarly, each face of $\mathbf{P}_{\Delta|_{V_2}}$ defines a face of \mathbf{P}_{Δ} . Each face of \mathbf{P}_{Δ} either arises in this way, or it is the intersection of two such faces, one induced by $\mathbf{P}_{\Delta|_{V_1}}$ and one induced by $\mathbf{P}_{\Delta|_{V_2}}$.

Proof. See [Eriksson et al. \(2006\)](#), Lemma 8. □

In the sequel, for any $V' \subseteq V$ and $i \in I = \prod_{v \in V} I_v$, it will be convenient to use the seemingly more complicated notation $\pi_{V'}(i) = (i_v, v \in V')$ for the marginal cell $i_{V'} \in I_{V'} := \prod_{v \in V'} I_v$. Similarly, for a set $S \subseteq I$, the restriction to V' is $\pi_{V'}(S) := \{\pi_{V'}(i) : i \in S\}$. For $T \subset I_{V'}$, the opposite action yields $\pi_{V'}^{-1}(T) = \{i \in I \mid i_{V'} \in T\}$.

We next translate Proposition 3.5 to the language of facial sets:

Lemma 3.6. *Suppose that Δ has a complete separator S that separates V into V_1 and V_2 .*

1. *If $F \subseteq I$ is facial with respect to Δ , then $\pi_{V_1}(F)$ and $\pi_{V_2}(F)$ are facial with respect to $\Delta|_{V_1}$ and $\Delta|_{V_2}$.*
2. *Conversely, if $F_1 \subseteq I_{V_1}$ and $F_2 \subseteq I_{V_2}$ are facial with respect to $\Delta|_{V_1}$ and $\Delta|_{V_2}$, then $\pi_{V_1}^{-1}(F_1) \cap \pi_{V_2}^{-1}(F_2)$ is facial with respect to Δ .*

Thus, for any $T \subseteq I$, let $T_1 = \pi_{V_1}(T)$ and $T_2 = \pi_{V_2}(T)$.

$$F_{\Delta}(T) = \pi_{V_1}^{-1}(F_{\Delta|_{V_1}}(T_1)) \cap \pi_{V_2}^{-1}(F_{\Delta|_{V_2}}(T_2)).$$

Proof. Consider an inequality as in Proposition 3.5 that defines a face \mathbf{F} of \mathbf{P}_Δ as well as a face \mathbf{F}_1 of \mathbf{P}_{Δ_1} . Then the corresponding facial sets F and F_1 satisfy $F = \pi_{V_1}^{-1}(F_1)$; because in order to check whether some f_i , $i \in I$, satisfies the inequality, we only need to look at the components involving V_1 ; that is, we only need to look at $\pi_{V_1}(i)$. \square

Lemma 3.6 easily generalizes to more than one separator and thus to more than two components and it becomes particularly simple when these components are complete. Indeed, in that case, $F_{\Delta|V_1}(T_1) = T_1$ and taking the preimage we obtain

$$\pi_{V_1}^{-1}(\pi_{V_1}(T)) = \{i \in I : \exists i' \in T \text{ such that } \pi_{V_1}(i) = \pi_{V_1}(i')\} \supseteq T.$$

The following lemma is an immediate consequence of Lemma 3.6.

Lemma 3.7. *Let Δ be a decomposable model with decomposition $\Delta = \Delta_1 \cup \Delta_2 \cup \dots \cup \Delta_r$ where Δ_i is a complete simplex on V_i , and let $\pi_i = \pi_{V_i}$ be the corresponding marginalization map. Then, for any $T \subseteq I$,*

$$F_\Delta(T) = \pi_1^{-1}(\pi_1(T)) \cap \pi_2^{-1}(\pi_2(T)) \cap \dots \cap \pi_r^{-1}(\pi_r(T)).$$

3.2 Inner approximations

To obtain an inner approximation, our strategy is to find a separator S of Δ and to complete it. To be precise, we augment Δ by adding all subsets of S . Thus, we obtain a simplicial complex $\Delta_S = \Delta \cup \{M : M \subseteq S\}$ in which S is a complete separator. We can apply Lemma 3.6 to find the facial set $F_{\Delta_S}(I_+)$, and this will be our inner approximation of $F_\Delta(I)$.

An even simpler approximation is obtained by not only completing the separator itself, but also the two parts V_1, V_2 separated by S : The simplicial complex $\Delta_{V_1, V_2} := \{M : M \subseteq V_1\} \cup \{M : M \subseteq V_2\}$ is decomposable and contains Δ . Its facial sets can be computed from Lemma 3.7.

In general, the approximation obtained from a single separator (or, in general, a single super-complex) is not good; that is, $F_t = F_\Delta(I_+)$ tends to be much larger than $F_{\Delta_S}(I_+)$ or $F_{\Delta_{V_1, V_2}}(I_+)$. Thus we need to combine information from several separators. For example, given two separators $S, S' \subseteq V$,

we find a chain of approximations

$$\begin{aligned}
G'_0 &:= I_+, \\
G_1 &:= F_{\Delta_S}(G'_0), \quad G'_1 := F_{\Delta_{S'}}(G_1), \\
G_2 &:= F_{\Delta_S}(G'_1), \quad G'_2 := F_{\Delta_{S'}}(G_2), \\
&\vdots
\end{aligned}$$

that satisfy

$$I_+ \subseteq G_1 \subseteq G'_1 \subseteq G_2 \subseteq \cdots \subseteq F_t,$$

where all inclusions except the last one are due to the definition of $F_{\Delta_S}(T)$ or $F_{\Delta_{S'}}(T)$ as the smallest facial sets containing T in Δ_S or $\Delta_{S'}$. The last inclusion is a consequence of Lemma 3.1 since both Δ_S and $\Delta_{S'}$ contain Δ . This chain of approximations has to stabilize at a certain point; that is, after a certain number of iterations, the approximations will not improve any more. The limit, which we denote by $F_{S,S'}(I^+) := \bigcup_i G_i = \bigcup_i G'_i$, can be characterized as the smallest subset of I that contains I^+ and is facial both with respect to Δ_S and $\Delta_{S'}$. The same iteration can be done replacing Δ_S and $\Delta_{S'}$ by Δ_{V_1, V_2} and $\Delta_{V'_1, V'_2}$. Applying in turn $F_{\Delta_{V_1, V_2}}$ and $F_{\Delta_{V'_1, V'_2}}$ gives another approximation $\tilde{F}_{S,S'}(I^+)$, namely the smallest subset of I that contains I^+ and is facial both with respect to Δ_{V_1, V_2} and $\Delta_{V'_1, V'_2}$. This latter approximation will be used in Section 5.1. Clearly, $\tilde{F}_{S,S'}(I^+)$ is a worse approximation than $F_{S,S'}(I^+)$, since $\tilde{F}_{S,S'}(I^+) \subseteq F_{S,S'}(I^+) \subseteq F_t$, but it is easier to compute.

We use the following strategies:

1. If possible, use all separators of a graph.

There are two problems with this strategy: First, if S is such that either V_1 or V_2 is large, then it is almost as difficult to compute $F_{\Delta|_{V_1}}$ and $F_{\Delta|_{V_2}}$ as $F_{\Delta|_V}$. Such “bad” separators always exist: namely, each node $i \in V$ is separated by its neighbours from all other nodes. In this case, V_1 consists of i its neighbours, and V_2 consists of $V \setminus \{i\}$. For such a “bad” separator we can only compute $F_{\Delta_{V_1, V_2}}$, but not F_{Δ_S} . Second, the number of separators may be large. Since we have to iterate over this set until the approximation converges, it may take a long time to compute the inner approximation.

A faster alternative strategy is the following:

2. Look at separators such that both $V_1 \setminus S$ and $V_2 \setminus S$ are not too small (for example, $\min\{|V_1 \setminus S|, |V_2 \setminus S|\} \geq 3$).

We illustrate the first strategy in Section 5.2, using a graphical model associated with the NLTCs data set. In the case of the grids studied in Sections 5.1 and 6.2, which have a lot of regularity, we use an adapted strategy:

3. In a grid, use the horizontal, vertical and diagonal separators.

In the case of grids, the vertical separators form a family of pairwise disjoint separators. In Section 6 we show how we can make use of such a family to study faces of hierarchical models, even if the facial sets are so large that they become computationally intractable.

3.3 Outer approximations

By Lemma 3.1, when we compute $F_{\Delta'}(S)$ for a simplicial complex $\Delta' \subseteq \Delta$, then we obtain an outer approximation of $F_{\Delta}(S)$. Removing sets from Δ decreases the dimension of the marginal polytope, so it is often easier to compute $F_{\Delta'}(S)$ than to compute $F_{\Delta}(S)$. Our main strategy is to look at subcomplexes induced by subset $V' \subset V$.

Let $\Delta_{V'}$ be the simplicial complex induced by V' . Let $J \subset I$ be its set of interactions. When comparing Δ with $\Delta_{V'}$, we have to be precise about whether we consider $\Delta_{V'}$ as a simplex on V or on V' . When we consider it on V , let A be its $I \times J$ design matrix with rows $f_i, i \in I$. When we consider it on V' , the design matrix A' is an $I_{V'} \times J$ matrix with columns $f'_{i'}, i' \in I_{V'}$. Because we have the same set of interactions whether we are on V or V' , we have for $i \in I$ and $i' \in I_{V'}$,

$$f_i = f'_{i'} \Leftrightarrow i \in \pi_{V'}^{-1}(i'). \quad (10)$$

Therefore the marginal polytopes of the two models are the same since they are the convex hull of the same set of vectors $\{f_i, i \in I\} = \{f'_{i'}, i' \in I_{V'}\}$. The relationship between the facial sets on V and V' is as follows:

Lemma 3.8. *Let $V' \subseteq V$. For $K \subset I$, we have*

$$F_{\Delta_{V'}}(K) = \pi_{V'}^{-1}(F'_{\Delta_{V'}}(\pi_{V'}(K))).$$

Here, $F'_{\Delta_{V'}}$ denotes the facial set when $\Delta_{V'}$ is considered as a simplicial complex on V' , and $F_{\Delta_{V'}}$ denotes the facial set when $\Delta_{V'}$ is considered as a simplicial complex on V .

Proof. For $K \subset I$, the two sets $\mathcal{A} = \{a_i, i \in K\}$ and $\mathcal{B} = \{b_{i'}, i' \in \pi_{V'}(K)\}$ are identical and therefore the smallest faces of the marginal polytopes for $\Delta_{V'}$ on V or V' containing \mathcal{A} and \mathcal{B} respectively are the same.

By definition of $F'_{\Delta_{V'}}(\pi_{V'}(K))$, the smallest face containing \mathcal{B} is defined by $\{b_{i'}, i' \in F'_{\Delta_{V'}}(\pi_{V'}(K))\}$. By definition of $F_{\Delta_{V'}}(K)$, the smallest face containing \mathcal{A} is $\{a_i, i \in F_{\Delta_{V'}}(K)\}$. Also by (10), we have that $\{a_i, i \in \pi_{V'}^{-1}(F'_{\Delta_{V'}}(\pi_{V'}(K)))\} = \{b_{i'}, i' \in F'_{\Delta_{V'}}(\pi_{V'}(K))\}$. Therefore $F_{\Delta_{V'}}(K) = \pi_{V'}^{-1}(F'_{\Delta_{V'}}(\pi_{V'}(K)))$. \square

In general, $F_{\Delta|_{V'}}(I_+)$ is not a good approximation of $F_{\Delta}(I_+)$. We can improve this approximation by considering several subsets of V . To be precise, if $V_1, \dots, V_r \subseteq V$, then $F_{\Delta}(I_+) \subseteq F_{\Delta|_{V_i}}(I_+)$ for $i = 1, \dots, r$, and thus $F_{\Delta}(I_+) \subseteq \bigcap_{i=1}^r F_{\Delta|_{V_i}}(I_+) =: F_{V_1, \dots, V_r; \Delta}(I_+)$.

The question is now how to choose the subsets V_i . Clearly, the subsets V_i should cover V , and, more precisely, they should cover Δ , in the sense that for any $D \in \Delta$ there should be one V_i with $D \subseteq V_i$. The larger the sets V_i , the better the approximation becomes, but the more difficult it is to compute $F_{V_1, \dots, V_r; \Delta}(I_+)$. One generic strategy is the following:

1. Use all subsets of V of fixed cardinality k plus all facets $D \in \Delta$ with $|D| \geq k$.

This choice of subsets indeed covers Δ . The parameter k should be chosen as large as possible such that computing $F_{V_1, \dots, V_r; \Delta}(I_+)$ is still feasible. Note that computing $F_{\Delta|_D}(I_+)$ for $D \in \Delta$ is trivial, since $\mathbf{P}_{\Delta|_D}$ is a simplex. Another natural strategy due to [Massam and Wang \(2015\)](#) is the following:

2. For fixed k , use balls $B_k(v) = \{w : d(v, w) \leq k\}$ around the nodes $v \in V$, where $d(\cdot, \cdot)$ denotes the edge distance in the graph.

In general, our philosophy is that the subsets V_i should be large enough to preserve some of the structure of Δ . For example, for the grid graphs, we suggest to use 3×3 -subgrids. These graphs have two nice properties: First, they already have the appearance of a small grid. Second, for any vertex $v \in V$, there is a 3×3 subgrid that contains v and all neighbours of v . We will compare two different strategies:

3. For a grid, use all 3×3 -subgrids.
4. Cover a grid by 3×3 -subgrids.

In Section 6.2 we compare these two methods, and we observe that, in the example of the 5×10 grid, it suffices to only look at a covering.

In general, it is not enough to look at induced sub-complexes, unless Δ has a complete separator (see Section 3.1). However, the approximation tends to be good and gives the correct facial set in many cases.

3.4 Comparing the two approximations

Suppose that we have computed two approximations F_1, F_2 of F_t such that $F_1 \subseteq F_t \subseteq F_2$. If we are in the lucky case that $F_1 = F_2$, then we know that $F_t = F_1 = F_2$. In general, the cardinality of $F_2 \setminus F_1$ indicates the quality of our approximations.

F_1, F_2 and F_t can also be compared by the ranks of the matrices $\tilde{A}_{F_1}, \tilde{A}_{F_2}$ and \tilde{A}_{F_t} obtained from \tilde{A} by keeping only the columns indexed by F_1, F_2 and F_t , respectively. Clearly, $\text{rank } \tilde{A}_{F_1} \leq \text{rank } \tilde{A}_{F_t} \leq \text{rank } \tilde{A}_{F_2}$. Note that $\text{rank } \tilde{A}_{F_2}$ equals the dimension of the corresponding face \mathbf{F}_2 of \mathbf{P} , and $\text{rank } \tilde{A}_{F_t}$ equals the dimension of \mathbf{F}_t . But F_1 does not necessarily correspond to a face of \mathbf{P} . Nevertheless, we can bound the codimension of \mathbf{F}_t in \mathbf{F}_2 by

$$\dim \mathbf{F}_2 - \dim \mathbf{F}_t \leq \text{rank } A_{F_2} - \text{rank } A_{F_1}.$$

In particular, if $\text{rank } A_{F_2} = \text{rank } A_{F_1}$, then we know that $F_t = F_2$. In this case, our approximations give us a precise answer, even if $F_1 \neq F_2$ and the lower approximation F_1 is not tight.

4 Parameter Estimation when the MLE does not exist

4.1 Computing the extended MLE

If the mle θ^* exists, then it can be computed by finding the unique maximum of the log-likelihood function $l(\theta)$ given in (8). As mentioned before, $l(\theta)$ is concave (or even strictly concave, if the parameters θ are identifiable), and thus the maximum is, at least in principle, easy to find (in practice, for larger models, it may be difficult to evaluate the function $k(\theta)$, which involves a sum over I ; but we will not discuss this problem here). In general, the maximum cannot be found symbolically, but there are efficient numerical algorithms to

maximize concave functions. Any reasonable hill-climbing algorithm should be capable of finding the mle. An example of an algorithm commonly used is *iterative proportional fitting* (IPF), which can be thought of as an algorithm of Gauss-Seidel type.

When the mle does not exist but the facial set $F = F_t$ of the data is known, then it is straight forward to compute the extended mle p^* . In this case, we know that p^* lies in $\mathcal{E}_{F,A}$. To find p^* , we need to optimize the log-likelihood \tilde{l} over $\mathcal{E}_{F,A} = \{p_{F,\theta} : \theta \in \mathbf{R}^h\}$. Plugging the parametrization $p_{F,\theta}$ into \tilde{l} tells us that we need to optimize the restricted log-likelihood function

$$l_F(\theta) = \log\left(\prod_{i \in I_+} p_{F,\theta}(i)^{n(i)}\right) = \sum_{j \in J} \theta_j t_j - Nk_F(\theta). \quad (11)$$

This problem is of a similar type as the problem to maximize l in the case that the mle exists, and the same algorithms as discussed above can be used. The problem here is slightly easier, since F is smaller than I . However, as stated above, the parametrization $\theta \mapsto p_{F,\theta}$ is never identifiable. Of course, this problem is easy to solve by selecting a set of independent parameters among the θ_j , as explained in Section 2.2. However, depending on the choice of the independent subset, the values of the parameters change, and in particular, it is meaningless to compare the values of the parameters θ_j with parameter values of any other distribution in \mathcal{E}_A or in the closure $\overline{\mathcal{E}_A}$.

Before explaining how to find better parameters on $\mathcal{E}_{F,A}$, let us discuss what happens if the facial set F_t of the data is not known. As mentioned before, whether or not the mle exists, the log-likelihood function $l(\theta)$ is always strictly concave (assuming that the parametrization is identifiable). When the mle does not exist, then the maximum is not at a finite value θ^* , but lies “at infinity.” Still, as observed by Geyer (2009, Section 3.15), any reasonable numerical “hill-climbing” algorithm that tries to maximize the likelihood will tend towards the right direction. Such a numeric algorithms generates a sequence of parameter values $\theta^{(1)}, \theta^{(2)}, \theta^{(3)}, \dots$ with increasing log-likelihood values $l(\theta^{(1)}) \leq l(\theta^{(2)}) \leq \dots$. Since $l(\theta)$ is concave, our optimization problem is numerically easy (at least in theory), and for any reasonable such algorithms, the limit $\lim_{s \rightarrow \infty} l(\theta^{(s)})$ will equal $\sup_{\theta} l(\theta) = \max_{p \in \overline{\mathcal{E}_A}} \tilde{l}(p)$. The algorithm will stop when the difference $l(\theta^{(s+1)}) - l(\theta^{(s)})$ becomes negligibly small. The output, $\theta^{(s)}$, then gives a good approximation of the EMLE, in the sense that p^* and $p_{\theta^{(s)}}$ are close to each other. For many applications, such as in machine learning, where it is more important to have good values

of the parameters instead of trying to model the “true underlying distribution,” or when doing a likelihood test, where the value of the likelihood is more important than the parameter values, this may be good enough.

However, in this numerical optimization, some of the parameters θ_j will tend to $\pm\infty$, which may lead to numerical problems. For example, it may happen that one parameter goes to $+\infty$ and a second parameter to $-\infty$ in such a way that their sum remains finite. This implies that a difference between two large numbers has to be computed, which is numerically unstable. Also, it is not clear, which parameters tend to infinity numerically. In fact, this may depend on the chosen algorithm; i.e. different algorithms may yield approximations of the EMLE that are qualitatively different in the sense that different parameters diverge. We give an example of this in Appendix B.

To avoid such problems, we propose a change of coordinates that allows us to control which parameters diverge, at least in the case where we know the facial set F_t . If we don't know F_t , but if we know approximations $F_1 \subseteq F_t \subseteq F_2$, we can use this knowledge to identify some parameters that definitely remain finite, while some parameters definitely diverge. While we cannot control the behaviour of the remaining parameters, the hope is that the more information we have about the facial set F_t , the better control we have about the above mentioned pathologies.

4.2 An identifiable parametrization

We have seen that when we use the parametrization $\theta \mapsto p_{F_t, \theta}$ of \mathcal{E}_{A, F_t} in the case where $F_t \neq I$, we have to expect the following (interrelated) issues:

1. The parametrization is not identifiable, i.e. there are parameters θ, θ' with $p_{F_t, \theta} = p_{F_t, \theta'}$.
2. While the parametrization $\theta \mapsto p_{F_t, \theta}$ looks similar to the parametrization $\theta \mapsto p_\theta$ of \mathcal{E}_A , the values of the parameters in both parametrizations are not related to each other.
3. When $p_{\theta^{(s)}} \rightarrow p_{F_t, \theta}$ as $s \rightarrow \infty$ for some parameter values $\theta^{(s)}, \theta$, then some of the parameter values $\theta^{(s)}$ diverge to $\pm\infty$. When computing probabilities, there may be linear combinations of these diverging parameters that remain finite.

Next we show that if F_t is known, then, with a convenient choice of L , the parameters μ_L (introduced in Section 2.2) solve 1 and 2 and improves 3.

Afterwards, we discuss what can be done if F_t is not known. We briefly discuss the general solution towards 3 in Appendix D. In any case, the choice of the parameters will depend on the facial set F_t ; i.e. it is not possible to define a single parametrization that works for all facial sets simultaneously.

Suppose that F_t is known. We consider the parameters μ_i as in Section 2.2, and we make sure that we choose the zero element 0 in I_+ . Recall that

$$\mu_i(\theta) = \langle \theta, f_i - f_0 \rangle = \log p(i)/p(0), i \in I.$$

As mentioned in Section 2.2, the parameters μ_i are not independent, and we need to choose an independent subset L . We will do this in two steps:

1. Choose a maximal subset L_t of F_t such that the parameters $\mu_i, i \in L_t$ are independent.
2. Then extend L_t to a maximal subset $L \subseteq I$ such that the parameters $\mu_i, i \in L$ are independent by adding elements $i \in I \setminus F_t$.

It follows from Theorem 2.9 that the following holds:

1. The subset $\mu_i, i \in L_t$, of the parameters μ_L gives an identifiable parametrization of $\mathcal{E}_{F_t, A}$.
2. Let $\mu_i^*, i \in L_t$, be the parameter values that maximize l_{F_t} (and thus give the EMLE). When the likelihood $l(\mu)$ in (9) is maximized numerically on I , then in successive iterations of the maximization, the estimates $\mu_i^{(s)}$ are such that

$$\mu_i^{(s)} \rightarrow \begin{cases} \mu_i^*, & i = 1, \dots, h_t, \\ -\infty, & \text{otherwise.} \end{cases}$$

In particular, no parameter tends to $+\infty$.

The last property ensures a consistency of the parameters μ_i on \mathcal{E}_A and on $\mathcal{E}_{F_t, A}$. This is important in those cases where the parameters have an interpretation and where it is of interest to know the value of some parameters, if it is well-defined. For example, in hierarchical models, the parameters correspond to “interactions” of the random variables, and it may be of interest to know, which of these interactions are important. Thus, it is of interest to know the size of the corresponding parameter. Usually, it is not the parameter μ_i , but the original parameters θ_i that have an interpretation. But when

we understand the parameters μ_i , we can also tell which of the parameters θ_i or which combinations of the parameters θ_i have finite well-defined values and can be computed, and which parameters diverge:

Lemma 4.1. *Suppose that $\theta^{(s)}$, $s \in \mathbb{N}$, are parameter values such that $p_{\theta^{(s)}} \rightarrow p^*$ as $s \rightarrow \infty$. For any $i \in L_t$, the linear combination*

$$\mu_i^{(s)} = \langle \theta^{(s)}, f_i \rangle$$

has a well-defined finite limit as $s \rightarrow \infty$. Any linear combination of the $\theta_i^{(s)}$ that has a well-defined finite limit (that is, a limit that is independent of the choice of the sequence $\theta^{(s)}$) is itself a linear-combination of the $\mu_i^{(s)}$ with $i \in L_t$.

Proof. The first statement follows from

$$\mu_i^{(s)} = \log p_{\theta^{(s)}}(i)/p_{\theta^{(s)}}(0) \rightarrow \log p^*(i)/p^*(0).$$

For the second statement, note that any linear combination of the θ is also a linear combination of the μ , since the linear map $\theta \mapsto \mu(\theta)$ is invertible. We now show that if a linear combination $\sum_i a_i \mu_i$ involves some μ_j with $j \notin L_t$, then there exist sequences $\mu^{(s)}$, $\mu'^{(s)}$ of parameters with

$$\lim_{s \rightarrow \infty} p_{\mu^{(s)}} = \lim_{s \rightarrow \infty} p_{\mu'^{(s)}} \quad \text{and} \quad \lim_{s \rightarrow \infty} \sum_i a_i \mu_i^{(s)} \neq \lim_{s \rightarrow \infty} \sum_i a_i \mu_i'^{(s)}.$$

So suppose that $\mu^{(s)}$ is a sequence of parameters such that $\lim_{s \rightarrow \infty} p_{\mu^{(s)}}$ exists and such that $\lim_{s \rightarrow \infty} \sum_i a_i \mu_i^{(s)}$ is finite. Define

$$\mu_i'^{(s)} = \begin{cases} \mu_j^{(s)} + 1, & \text{if } i=j, \\ \mu_i^{(s)}, & \text{otherwise.} \end{cases}$$

An easy computation shows that

$$\lim_{s \rightarrow \infty} p_{\mu'^{(s)}} = \lim_{s \rightarrow \infty} p_{\mu^{(s)}} \quad \text{and} \quad \lim_{s \rightarrow \infty} \sum_i a_i \mu_i'^{(s)} = \lim_{s \rightarrow \infty} \sum_i a_i \mu_i^{(s)} + a_j. \quad \square$$

Suppose now that we do not know F_t , but that instead we have approximations F_1, F_2 that satisfy

$$I_+ \subseteq F_1 \subseteq F_t \subseteq F_2 \subseteq I.$$

In this case, we proceed as follows to obtain an independent subset L among the parameters μ_i :

1. Choose a maximal subset L_1 of F_1 such that the parameters $\mu_i, i \in L_1$ are independent.
2. Then extend L_1 to a maximal subset $L_2 \subseteq F_2$ such that the parameters $\mu_i, i \in L_2$ are independent by adding elements $i \in F_2 \setminus F_1$.
3. Finally, extend L_2 to a maximal subset $L \subseteq I$ such that the parameters $\mu_i, i \in L$ are independent by adding elements $i \in I \setminus F_2$.

These parameters have the following properties that follow directly from Lemma 4.1:

Lemma 4.2. *Suppose that $\theta^{(s)}, s \in \mathbb{N}$, are parameter values such that $p_{\theta^{(s)}} \rightarrow p^*$ as $s \rightarrow \infty$, and let $\mu_i^{(s)} = \langle \theta^{(s)}, f_i \rangle$.*

1. *For any $i \in L_1$, the linear combination*

$$\mu_i^{(s)} = \langle \theta, f_i \rangle$$

has a well-defined finite limit as $s \rightarrow \infty$. Thus, any linear combination of the $\mu_i^{(s)}$ with $i \in L_1$ has a well-defined limit as $s \rightarrow \infty$.

2. *Any linear combination $\sum_i a_i \mu_i^{(s)}$ that has a well-defined limit as $s \rightarrow \infty$ is in fact a linear combination of the $\mu_i^{(s)}$ with $i \in L_2$. Thus, a linear combination that involves at least one $\mu_j^{(s)}$ with $j \in L \setminus L_2$ does not have a well-defined limit.*

5 Simulation study and applications to real data

In this section, we illustrate our methodology. In 5.1, we simulate data for the graphical model of the 4×4 grid and show how to exploit the various types of separators in order to obtain good inner and outer approximations. We find that our method gives very accurate result in this model of modest size. In 5.2, we work with the NLTCs data set, a real-world data set. We compare different inner approximations F_1 and find that most of the time, F_1 and F_2 are equal, and thus they are both equal to F_t . We also compute the EMLE and compare the result to what happens when maximizing the likelihood functions l and l_{F_2} .

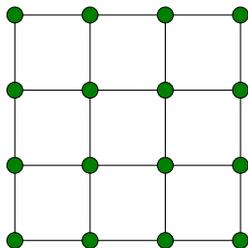


Figure 1: 4×4 grid graph

5.1 4×4 grid graph

We generated random samples of varying sizes for the graphical model of the 4×4 grid graph (Fig. 1). For each sample, we compute inner and outer approximations F_1 and F_2 , and we compare them to the true facial set F_t , which we can obtain using linear programming. To obtain an inner approximation, we use two strategies. Either, we iterate over all possible separators of which there are 106 (Strategy (1) in Section 3.2) or we iterate over the 3 horizontal, 3 vertical and 8 diagonal separators only (Strategy (3) in Section 3.2). We obtain the same result with either strategy. Clearly, Strategy (3) is much faster. To compute the outer approximation, we cover the 4×4 grid by four 3×3 -grids (Strategy (3) in Section 3.3).

We first generate random samples from the uniform distribution, that is from the probability distribution P_θ in the hierarchical model where all parameters $\theta_j, j \in J$ are set to zero. The results are given in Table 1. For each sample size, 1000 samples were obtained. As the table shows, for larger samples the probability that our random sample lies on a proper face becomes very small. If $F_t = I$, then clearly $F_t = F_2$. But we also found $F_t = F_2$ for all samples with t lying on a proper face, which shows that F_2 is an excellent approximation of F_t in this model. For the inner approximation, we observed some samples with $F_1 \neq F_t$, but they seem to be very rare.

Second, to better understand what happens for large samples, we change our sampling scheme. Instead of sampling from the uniform distribution, we generate samples from the hierarchical model P_θ , where the vector of parameters θ is drawn from a multivariate standard normal distribution (for each sample, new parameters were drawn). The results are given in Table 2. Again, for each sample size, 1000 samples were obtained. One can see that in this sampling scheme, we are much more likely to find that $F_t \neq I$. Observe

Table 1: facial set approximation of 4×4 grid graph (sample from uniform distribution)

sample size	data on face	$F_1 = F_t$	$F_2 = F_t$
10	98.5%	96.3%	100.0%
15	68.9%	99.9%	100.0%
20	29.0%	100.0%	100.0%
50	0.0%	100.0%	100.0%

Table 2: facial set approximation of 4×4 grid graph(hierarchical log-linear model with parameters from standard normal distribution)

sample size	data on face	$F_1 = F_t$	$F_2 = F_t$
10	100.0%	97.7%	100.0%
50	89.5%	100.0%	100.0%
100	71.0%	100.0%	100.0%
150	52.0%	100.0%	100.0%

that the squared length of the parameter vector θ is χ^2 -distributed with 39 degrees of freedom (since the number of parameters is 40). Thus, the expected length of θ is 39, which is large enough to move the distribution p_θ close to the boundary of the model. Indeed, we observed that when the mle does not exist, the length of the numerical estimate of the mle vector is of the order of magnitude of 40 (see also the next example in Section 5.2). Again, in all the samples that we generated, $F_t = F_2$, and $F_1 = F_2$ in the vast majority of cases. Thus, for this graph of relatively modest size, our approximations are very good.

5.2 NLTCS data set

To illustrate how approximate knowledge of the facial set allows us to say which parameters can be estimated (as explained in Section 4), we study the NLTCS data set, which consists of 21 574 observations on 16 binary variables, called ADL1, ..., ADL6, IADL1, ..., IADL10. The reader is referred to [Dobra and Lenkoski \(2011\)](#) for a detailed description of the data set. To associate a hierarchical to this data, we rely on the results of [Dobra](#)

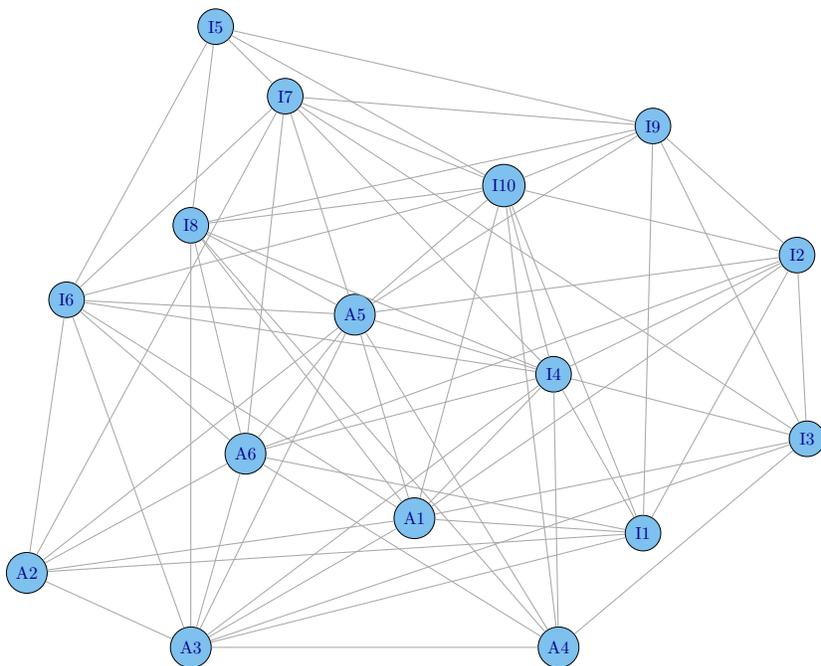


Figure 2: Graphical model for NLTCS data set. The label “An” abbreviates ADLn, “In” abbreviates IADLn.

and Lenkoski (2011) who use a Bayesian approach to estimate the posterior inclusion probabilities of edges. We construct a graph by saying that (x, y) is an edge if and only if the posterior inclusion probability of (x, y) is at least 0.40; see Figure 2. Then we take the corresponding clique complex of this graph so that our hierarchical model is a graphical model. There are 314 parameters in this model, including up to 6-way interactions. In total, the graph has 40 separators.

In order to compare the maximum likelihood estimates obtained with or without worrying about its existence and with or without approximation to F_t , we maximize the loglikelihood given in terms of μ (rather than θ) as in (9). There are 314 independent parameters.

First we ignore the fact that the mle might not exist and compute the mle of μ using the standard minfunc optimization software in Matlab: we call this estimate $\hat{\mu}^{\text{MLE}}$. Second, we find F_t and compute the EMLE with parameters denoted $\hat{\mu}^{\text{EMLE}}$. Third, we obtain an inner and outer approximation to F_t and consider the resulting information on the mle of the parameters. We call

the resulting estimate $\hat{\mu}^{F'_1/F'_2}$.

To compute $\hat{\mu}^{\text{EMLE}}$, we first compute the inner approximation F_1 that makes use of all the separators in the graph (Strategy 1 in Section 3.2). We also compute an outer approximation F_2 from all $\binom{16}{5} = 4368$ size five local models and the cliques of size six (Strategy 1 in Section 3.3). We obtain $F_1 = F_2$ and thus deduce that $F_t = F_1 = F_2$. We find $|F_t| = 49\,536$, and so $|F_t^c| = 2^{16} - 49\,536 = 16\,000$. Therefore, 16 000 cell probabilities are zero in the EMLE. We can obtain the mle by maximizing the loglikelihood function l_{F_t} as in (11). Since $\text{rank}(A_{F_t}) = 302$, the dimension of \mathbf{F}_t is 302, and there are only 302 parameters in l_{F_t} .

To show how to use the inner and outer approximations when F_t is not known, we choose to find coarser inner and outer approximations to F_t , respectively denoted F'_1 and F'_2 , and use them to compute the other approximation $\hat{\mu}^{F'_1/F'_2}$ to the mle. To compute F'_1 , we just use 10 random separators. We find $|F'_1| = 36\,954$ and $\dim \mathbf{F}'_1 = \text{rank } A_{F'_1} = 300$. To compute the outer approximation F'_2 , we consider the 4368 local size-five induced models and select among them the 1000 with the facial sets of smallest cardinality, which we glue together. We find $|F'_2| = 50\,688$ and $\dim \mathbf{F}'_2 = \text{rank } A_{F'_2} = 310$. Thus, we know that at least $|I \setminus F'_2| = 2^{16} - 50\,688 = 14\,848$ cell probabilities vanish in the extended mle. Since we pretend not to know F_t , we replace l_{F_t} by

$$l_{F'_2}(\mu) = \sum_{i \in I_+} \mu_i n(i) - N \sum_{i \in F'_2} \exp(\mu_i). \quad (12)$$

We know that μ_i is estimable, for $i \in F'_1$, μ_i goes to negative infinity for $i \in F'_2^c$, and we cannot say anything for μ_i with $i \in F'_2 \setminus F'_1$.

As explained in Section 4.2, the components of μ are not functionally independent. We choose $L_1 \subseteq F'_1$, $L_2 \subseteq F'_2$ and $L \subseteq I$ as in Section 4.2 (we note that the zero cell belongs to I_+). Then any μ_i , $i \in F'_2$, can be written as a linear combination of $\mu_{L_2} = (\mu_i, i \in L_2)$, and we can write $\mu_i = \langle b_i, \mu_L \rangle$ for an appropriate vector b_i . Thus, $l_{F'_2}(\mu)$ only depends on $\mu_{L_2} = (\mu_i, i \in L_2)$, and (12) can be rewritten as

$$l_{F'_2}(\mu_L) = \sum_{i \in I_+} \langle b_i, \mu_L \rangle n(i) - N \sum_{i \in F'_2} \exp\langle b_i, \mu_L \rangle. \quad (13)$$

Of course, the maximum of $l_{F'_2}$ does not exist but, as for the maximization of l , the computer still gives us a numerical approximation, $\hat{\mu}_L$, and thus also a numerical estimate $\hat{\mu}_i = \langle b_i, \hat{\mu}_L \rangle$, $i \in F'_2$.

Table 3: The MLE estimation from 3 methods compared with the relative frequency in the NLTCs data. Here, each $i = (i_1, \dots, i_{16}) \in I = \{0, 1\}^{16}$ is represented by the natural number $\sum_{j=1}^{16} i_j 2^{j-1} \in \{0, \dots, 2^{16} - 1\}$.

Parameter		naive estimate	maximum likelihood estimates		
		$\log n_i/n_0$	$\hat{\mu}_i^{\text{MLE}}$	$\hat{\mu}_i^{\text{EMLE}}$	$\hat{\mu}_i^{F'_1/F'_2}$
$i \in F'_1$	μ_{512}	-1.2472	-1.2482	-1.2482	-1.2482
	μ_{65536}	-1.7644	-1.7976	-1.7975	-1.7975
	μ_{16}	-2.3958	-2.3844	-2.3846	-2.3846
	μ_{528}	-2.5429	-2.6504	-2.6504	-2.6504
	μ_{2048}	-2.8813	-2.7246	-2.7243	-2.7243
$i \in F_t \setminus F'_1$	μ_{32960}	$-\infty$	-13.8205	-13.8207	-13.8205
	μ_{34881}	$-\infty$	-14.3693	-14.3693	-14.3692
$i \in F'_2 \setminus F_t$	μ_{36864}	$-\infty$	-30.8729	$-\infty$	-34.9805
	μ_{36880}	$-\infty$	-39.6536	$-\infty$	-45.2229
	μ_{388}	$-\infty$	-28.9090	$-\infty$	-29.4525
	μ_{32769}	$-\infty$	-32.3799	$-\infty$	-36.9537
	μ_{385}	$-\infty$	-37.1365	$-\infty$	-35.9399
	μ_{449}	$-\infty$	-38.9673	$-\infty$	-44.9405
	μ_{32785}	$-\infty$	-40.1221	$-\infty$	-45.8318
	μ_{389}	$-\infty$	-43.7297	$-\infty$	-40.0158
$i \in I \setminus F'_2$	μ_{256}	$-\infty$	-35.5482	$-\infty$	$-\infty$
	μ_{320}	$-\infty$	-42.5454	$-\infty$	$-\infty$
	μ_{257}	$-\infty$	-52.9224	$-\infty$	$-\infty$
	μ_{321}	$-\infty$	-60.2208	$-\infty$	$-\infty$

In total, there are $|L_2| = \text{rank}(A_{F'_2}) = 310$ independent parameters in the loglikelihood function (13). Among them, we find $|L_2| = \text{rank}(A_{F'_2}) = 300$ estimable parameters $\mu_i, i \in L_2$. We cannot say anything about the 10 parameters indexed by $L_2 \setminus L_1$. If we know F_t , we can identify two more estimable parameters.

In Table 3, we give the three estimates of μ_i that we mentioned above, namely, $\hat{\mu}_i^{\text{MLE}}, \hat{\mu}_i^{\text{EMLE}}$ and $\hat{\mu}_i^{F'_1/F'_2}$. We also list the naive estimator $\log \frac{n_i}{n_0}$. We list estimates for 19 of the 310 possible parameters. In the first column of the table, we indicate which category the index i belongs to, that is, whether it belongs to F'_1, F_t or F'_2 . In the second column, we list the particular parameters considered.

Table 4: Expected cell counts for the top six largest counts cells in the NLTCS data estimated according to Grade of Membership models (GoM), Latent class models (LC), copula Gaussian graphical models (CGGM) and maximum likelihood (MLE).

Support of Cell	Observed	GoM	LC	CGGMs	MLE on facial set
\emptyset	3853	3269	3836.01	3767.76	3647.4
$\{10\}$	1107	1010	1111.51	1145.86	1046.9
$\{1 : 16\}$	660	612	646.39	574.76	604.4
$\{5\}$	351	331	360.52	452.75	336
$\{5, 10\}$	303	273	285.27	350.24	257.59
$\{12\}$	216	202	220.47	202.12	239.24

In Table 4, we list the estimates of the top five cell counts obtained using our method and compare them with those obtained by other methods in [Dobra and Lenkoski \(2011\)](#).

6 Computing faces for large complexes

If our statistical model contains many variables and is not reducible, the problem of determining \mathbf{F}_t quickly becomes infeasible. Not only does the marginal polytope become very complicated, but also the size of the objects that one has to store or compute grows exponentially. Consider for example a 10×10 grid of binary random variables. This hierarchical model has 280 parameters, and the total sample space has cardinality $|I| = 2^{100} \approx 1.27 \times 10^{30}$. If F_t is close to I , we cannot even list the elements of F_t , which consists of approximately 10^{30} elements. Therefore, we take a local approach and look for separators.

If the simplicial complex Δ contains a complete separator separating V into V_1 and V_2 , we can identify a facial set F implicitly without listing it explicitly. We only need the two projections $F_{V_1} = \pi_{V_1}(F)$ and $F_{V_2} = \pi_{V_2}(F)$. Since $F = \pi_{V_1}^{-1}(F_{V_1}) \cap \pi_{V_2}^{-1}(F_{V_2})$ (by Lemma 3.6), these two projections identify F , and they allow us to do most of the operations that we would want to do with F . For example, for any $i \in I$, we can check whether $i \in F$ by checking whether $\pi_{V_1}(i) \in F_{V_1}$ and $\pi_{V_2}(i) \in F_{V_2}$, and we can check whether $F = I$ by checking whether $F_{V_1} = I_{V_1}$ and $F_{V_2} = I_{V_2}$. In particular, we can check whether the MLE exists by looking only at the two subsets V_1 and V_2 .

If Δ contains a separator that is not complete, we can use similar ideas when computing inner and outer approximations to F_t , and also when comparing these two approximations. Suppose that S separates V_1 from V_2 in Δ . We want to use $F_2 := F_{\Delta|_{V_1}}(I_+) \cap F_{\Delta|_{V_2}}(I_+)$ as an outer approximation and $F_1 := F_{\Delta_S}(I_+)$ as an inner approximation to F_t . Due to the problems mentioned above, we do not directly compute F_1 and F_2 , but we compute their projections on V_1 and V_2 . Instead of F_2 , we compute the facial set $F_{2,V_1} := F_{\Delta|_{V_1}}(\pi_{V_1}(I_+))$ of the V_1 -marginal $\pi_{V_1}(I_+)$ with respect to $\Delta|_{V_1}$, and similarly we compute $F_{2,V_2} := F_{\Delta|_{V_2}}(\pi_{V_2}(I_+))$. Instead of F_1 , we compute $F_{1,V_1} := F_{\Delta_S|_{V_1}}(\pi_{V_1}(I_+))$ and $F_{1,V_2} := F_{\Delta_S|_{V_2}}(\pi_{V_2}(I_+))$. Then we could recover F_1 and F_2 from the equations

$$F_2 = \pi_{V_1}^{-1}(F_{2,V_1}) \cap \pi_{V_2}^{-1}(F_{2,V_2}) \quad \text{and} \quad F_1 = \pi_{V_1}^{-1}(F_{1,V_1}) \cap \pi_{V_2}^{-1}(F_{1,V_2}).$$

For any $x \in I$, we can check whether $x \in F_1$ by checking whether $\pi_{V_1}(x) \in F_{1,V_1}$ and $\pi_{V_2}(x) \in F_{1,V_2}$. More importantly, we can check whether $F_1 = F_2$ by checking whether $F_{1,V_1} = F_{2,V_1}$ and $F_{1,V_2} = F_{2,V_2}$. This idea can be applied iteratively when $\Delta|_{V_1}$ or $\Delta|_{V_2}$ has a separator.

The next two subsections illustrate these ideas. In Section 6.1, we consider a graph with no particular regularity pattern on 100 nodes and identify two convenient separators. In Section 6.2, we consider a grid graph and work with two families of “parallel” separators that can be used to iteratively improve the inner approximation.

6.1 US Senate Voting Records Data

We consider the voting record of all 100 US senators on 309 bills from January 1 to November 19 2015. Similar data for the years 2004–2006 was analyzed by [Banerjee et al. \(2008\)](#). The votes are recorded as “yea,” “nay” or “not voting.” We transformed the “not voting” into “nay” and consequently have a 100-dimensional binary data set. To fit a hierarchical model to this data set, we use the ℓ_1 -regularized logistic regression method proposed by [Ravikumar et al. \(2010\)](#) to identify the neighbours of each variable and construct an Ising model. We set the penalty parameter to $\lambda = 32\sqrt{\log p/n} \approx 0.35$, resulting in the sparse graph in Figure 3. There are 277 parameters in this model (the number of vertices plus the number of edges). The graph consists of two large connected components and 14 independent nodes.

There are 309 sample points, and $|I_+| = 278$. We want to know whether the data lies on a proper face of the marginal polytope to see if the mle of the

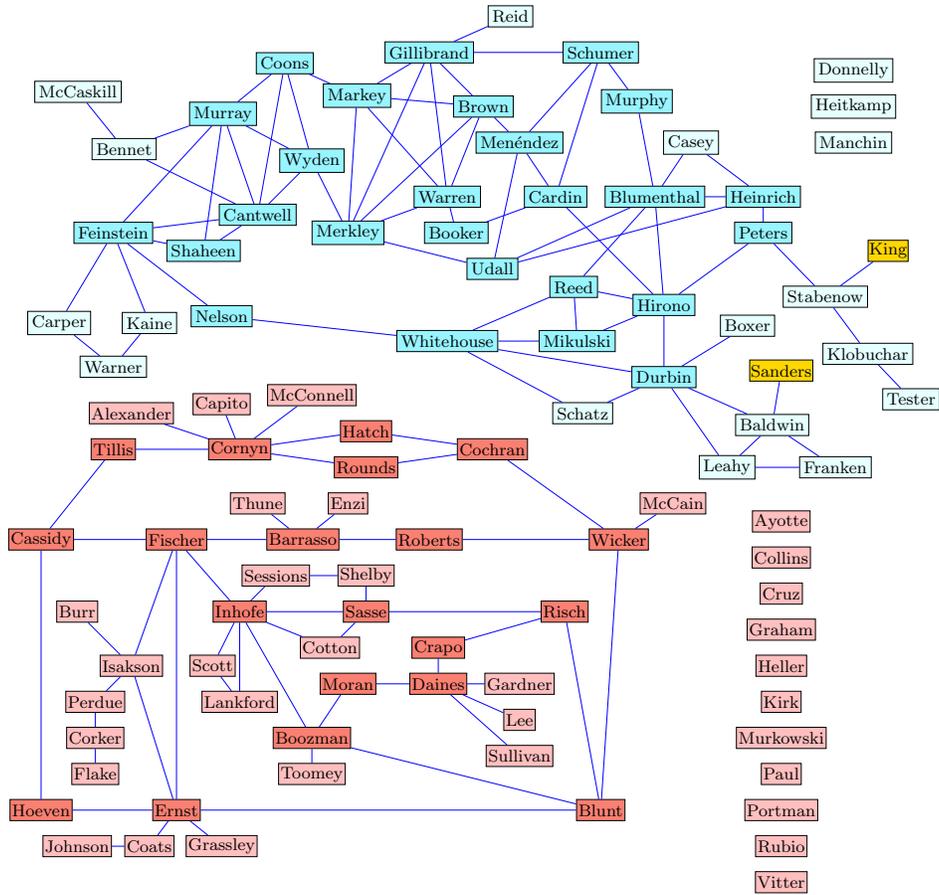


Figure 3: The graph for the US senate voting records data. Golden nodes are independent senators, blue nodes are democratic, and red nodes are republican.

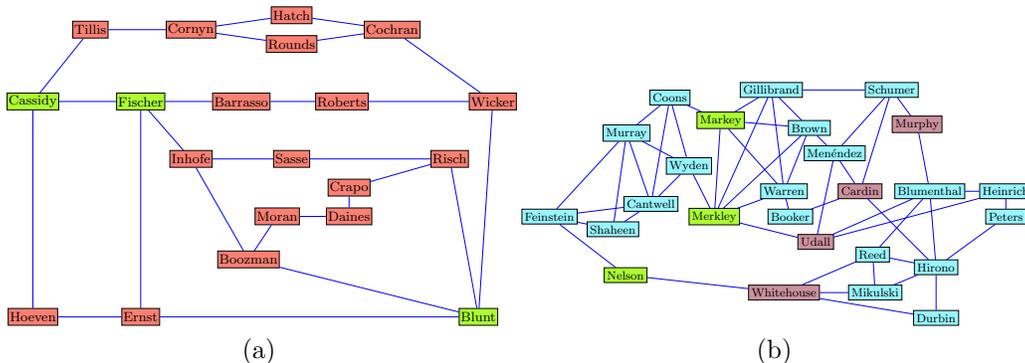


Figure 4: The simplicial complexes after cutting off the small prime components: (a) the republican party prime component Δ_r . (b) the democratic party prime component Δ_d . The yellow and pink nodes are the two separator sets we found to compute the facial set.

parameters exist. From Lemma 3.6, we know that if we find complete separators, we need only work with each of the irreducible simplicial complexes defined by these separators. We easily “cut-off” a number of relatively small prime components and verify that the data does not lie on a proper face of their corresponding marginal polytopes. We are left with one irreducible prime component in each of the two connected subgraphs, i.e. for each of the two parties as shown in Figure 4.

The democratic party simplicial complex Δ_d consists of 26 variables, and the model induced from Δ_d contains 77 parameters. The size of the design matrix A_{Δ_d} is $2^{26} \times 77$, which is too large to use linear programming to compute the facial set of the face \mathbf{P}_{Δ_d} containing the vector t_d . Therefore we look for separators that will help us obtain good inner and outer approximations. In Figure 4b, we indicate in yellow and pink two separators, which separate Δ_d into three simplicial complexes denoted (from top to bottom) by Δ_α , Δ_β and Δ_γ . The number of vertices of the three simplicial complexes are 9, 13, 11, respectively, and so we can apply the linear programming method to the three corresponding marginal polytopes.

The dimension of the model induced by Δ_α is 24. The corresponding data vector t_α lies in the relative interior of $\mathbf{P}_{\Delta_\alpha}$.

The dimension of the model induced by Δ_β is 34, and the data vector t_β lies on a facet \mathbf{F}_{t_β} of $\mathbf{P}_{\Delta_\beta}$. To simplify our notation, we denote the 100 senators not by their name but by an integer between 1 and 100. We only

ID	Senator	ID	Senator	ID	Senator	ID	Senator
22	Nelson	37	Cardin	52	Murphy	61	Whitehouse
23	Reed	41	Markey	53	Hirono	87	Warren
26	Schumer	47	Udall	56	Gillibrand		

Table 5: Numbering of some senators

need to identify a few and their numbers are given in Table 5. The inequality of \mathbf{F}_{t_β} is

$$t_{87} - t_{56,87} \geq 0, \quad (14)$$

where t_{87} denotes the marginal count of senator Warren voting “yea” and $t_{56,87}$ denotes the marginal counts of both senators Gillibrand and Warren voting “yea.”

The dimension of the model induced by Δ_γ is 27. The data vector t_γ lies on the facet of $\mathbf{P}_{\Delta_\gamma}$ with inequality

$$t_{23} - t_{23,53} \geq 0. \quad (15)$$

The intersection of the two facets (14) and (15) gives the outer approximation \mathbf{F}_2 to F_t .

To get an inner approximation, we complete each separator, i.e. the yellow vertices are completed and the pink vertices are completed in Figure 4b. Denote the three simplicial complexes with complete separators as $\Delta_{\tilde{\alpha}}$, $\Delta_{\tilde{\beta}}$, $\Delta_{\tilde{\gamma}}$ respectively. Then $\Delta_{\tilde{d}} = \Delta_{\tilde{\alpha}} \cup \Delta_{\tilde{\beta}} \cup \Delta_{\tilde{\gamma}}$ is a simplicial complex with two complete separators. The smallest face $\mathbf{F}_{t_{\tilde{d}}}$ of the marginal polytope $\mathbf{P}_{\Delta_{\tilde{d}}}$ containing the data vector $t_{\tilde{d}}$ is our inner approximation. Now the models of $\Delta_{\tilde{\alpha}}$, $\Delta_{\tilde{\beta}}$, $\Delta_{\tilde{\gamma}}$ and $\Delta_{\tilde{d}}$ are not models with main effects and two-way interactions only; they also include parameters for third and fourth order interactions. The dimension of the model induced by $\Delta_{\tilde{d}}$ is 91: we added 14 parameters to the original model by completing the two separators. Again, we apply the linear programming method to the three marginal polytopes $\mathbf{P}_{\Delta_{\tilde{\alpha}}}$, $\mathbf{P}_{\Delta_{\tilde{\beta}}}$ and $\mathbf{P}_{\Delta_{\tilde{\gamma}}}$.

The dimension of the model of $\Delta_{\tilde{\alpha}}$ is 27, and $\mathbf{F}_{t_{\tilde{\alpha}}}$ is a facet with equation

$$\langle g_1, t_{\tilde{\alpha}} \rangle = t_{41} - t_{22,41} - t_{41,70} + t_{22,41,70} = 0. \quad (16)$$

It follows that $\{g_1\}$ is a basis of the kernel of $A_{F_{\tilde{\alpha}}}^t$.

The dimension of the model for $\Delta_{\tilde{\beta}}$ is 48. The face $\mathbf{F}_{t_{\tilde{\beta}}}$ has codimension 5, with defining equations

$$\begin{cases} \langle g_2, t_{\tilde{\beta}} \rangle = t_{87} - t_{56,87} = 0 \\ \langle g_3, t_{\tilde{\beta}} \rangle = t_{47,52,61} + t_{37,52} - t_{37,52,61} - t_{37,47,52} = 0 \\ \langle g_4, t_{\tilde{\beta}} \rangle = t_{37,47,52,61} - t_{47,52,61} = 0 \\ \langle g_5, t_{\tilde{\beta}} \rangle = t_{37,52} + t_{26} - t_{26,52} - t_{26,37} = 0 \\ \langle g_6, t_{\tilde{\beta}} \rangle = t_{41} - t_{22,41} - t_{41,70} + t_{22,41,70} = 0 \end{cases} . \quad (17)$$

Again, $\{g_2, g_3, g_4, g_5, g_6\}$ is a basis of the kernel of $A_{F_{\tilde{\beta}}}$.

The dimension of the model for $\Delta_{\tilde{\gamma}}$ is 38. The face $\mathbf{F}_{t_{\tilde{\gamma}}}$ has codimension 3. It is defined by the equations

$$\begin{cases} \langle g_7, t_{\tilde{\gamma}} \rangle = t_{47,52,61} + t_{37,52} - t_{37,52,61} - t_{37,47,52} = 0 \\ \langle g_8, t_{\tilde{\gamma}} \rangle = t_{37,47,52,61} - t_{47,52,61} = 0 \\ \langle g_9, t_{\tilde{\gamma}} \rangle = t_{23} - t_{23,53} = 0 \end{cases} . \quad (18)$$

Again, $\{g_7, g_8, g_9\}$ is a basis of the kernel of $A_{F_{\tilde{\gamma}}}$.

From Lemma 3.6, we know that $\mathbf{F}_{t_{\tilde{d}}} = \mathbf{F}_{\tilde{\alpha}} \cap \mathbf{F}_{\tilde{\beta}} \cap \mathbf{F}_{\tilde{\gamma}}$, and the equations for $\mathbf{F}_{t_{\tilde{d}}}$ are

$$\begin{cases} \langle g'_1, t_{\tilde{d}} \rangle = t_{41} - t_{22,41} - t_{41,70} + t_{22,41,70} = 0 \\ \langle g'_2, t_{\tilde{d}} \rangle = t_{87} - t_{56,87} = 0 \\ \langle g'_3, t_{\tilde{d}} \rangle = t_{47,52,61} + t_{37,52} - t_{37,52,61} - t_{37,47,52} = 0 \\ \langle g'_4, t_{\tilde{d}} \rangle = t_{37,47,52,61} - t_{47,52,61} = 0 \\ \langle g'_5, t_{\tilde{d}} \rangle = t_{37,52} + t_{26} - t_{26,52} - t_{26,37} = 0 \\ \langle g'_9, t_{\tilde{d}} \rangle = t_{23} - t_{23,53} = 0 \end{cases} , \quad (19)$$

where the vectors g'_1, \dots, g'_9 are the vectors g_1, \dots, g_9 extended to R^{91} by adding zeros on the corresponding complementary coordinates. Note that since $g'_1 = g'_6$, $g'_3 = g'_7$, $g'_4 = g'_8$, we only need six of the nine equations. Thus, $\mathbf{F}_1 := \mathbf{F}_{t_{\tilde{d}}}$, defined by (19), is a strict subset of the face \mathbf{F}_2 defined by (14) and (15). Next, we will refine our argument and show that indeed $\mathbf{F}_{t_{\tilde{d}}} = \mathbf{F}_2$.

From what we know, it follows that the orthogonal complement of the subspace generated by $\mathbf{F}_{t_{\tilde{d}}}$ is

$$G = \{g' \in R^{91} \mid g' = k_1 g'_1 + k_2 g'_2 + k_3 g'_3 + k_4 g'_4 + k_5 g'_5 + k_9 g'_9\}.$$

To describe \mathbf{F}_{t_d} , we want to describe the defining equations of \mathbf{F}_{t_d} . Each such equation is of the form $\langle g, t_d \rangle = 0$, where g is orthogonal to \mathbf{F}_{t_d} . For any such g , let g' be its extension to a vector in R^{91} by adding zero components. Then $g' \perp \mathbf{F}_{t_d}$, which implies that $g' \in G$. Therefore, we can find g by finding all vectors $g' \in G$ that vanish on all added components. This yields a system of linear equations in k_1, \dots, k_5, k_9 . We claim that all solution must satisfy $k_1 = k_3 = k_4 = k_5 = 0$. Indeed, the coefficient of any triple or quadruple interaction must vanish (since these don't belong to the original Ising model), which implies $k_1 = k_3 = k_4 = 0$, and also the coefficient of $t_{37,52}$ must vanish, which implies $k_5 = 0$. On the other hand, the vectors g'_2 and g'_9 only contain interactions that are already present in Δ , and so the coefficients k_2 and k_9 are free. Thus the equations for \mathbf{F}_{t_d} are

$$\begin{cases} \langle g_2, t_{\tilde{\beta}} \rangle = t_{87} - t_{56,87} = 0, \\ \langle g_9, t_{\tilde{\gamma}} \rangle = t_{23} - t_{23,53} = 0. \end{cases} \quad (20)$$

This is the same as the outer approximation \mathbf{F}_2 .

The republican simplicial complex Δ_r consists of 20 variables, and the model induced from Δ_r contains 46 parameters. The size of the design matrix A_{Δ_r} is $2^{20} \times 46$, which is also too large to directly compute F_t . The yellow nodes in Figure 4a separate Δ_r into two simplicial complexes denoted (from left to right) by Δ_a and Δ_b . To compute the inner approximation, we complete the yellow separators and we get two new simplicial complexes $\Delta_{\tilde{a}}$ and $\Delta_{\tilde{b}}$. With the linear programming algorithm, we find that the corresponding data $t_{\tilde{a}}$ and $t_{\tilde{b}}$ lie in the relative interior of the polytopes $\mathbf{P}_{\Delta_{\tilde{a}}}$ and $\mathbf{P}_{\Delta_{\tilde{b}}}$, respectively. Therefore we have $\mathbf{F}_1 = \mathbf{P}_{\Delta_r}$. Since $\mathbf{F}_1 \subseteq \mathbf{F}_t \subseteq \mathbf{P}_{\Delta_r}$, we conclude that the corresponding data vector t_r lies in the relative interior of \mathbf{P}_{Δ_r} .

6.2 The 5×10 -grid

Let Δ be the simplicial complex of the 5×10 grid graph. We exploit the regularity of this graph and make use of the vertical separators in the grid to obtain inner and outer approximations of the facial sets. The graph has 50 nodes, which is too many to directly compute a facial set or even to store it. However, the 5×10 grid has 8 vertical separators marked in red and blue in Figure 5, and we can use these to approximate F_t . Since facial sets for 5×3 -grids can be computed reasonably fast (3 to 4 seconds on a laptop with

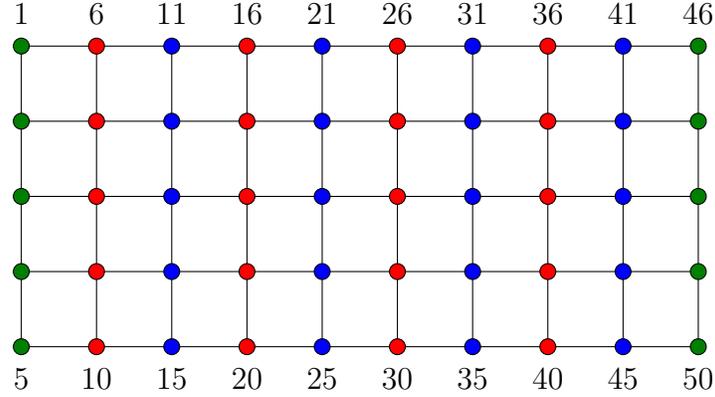


Figure 5: 5×10 grid graph, the red and blue nodes are the set of separators we use to compute F_1 , they are used iteratively to get a better lower approximation

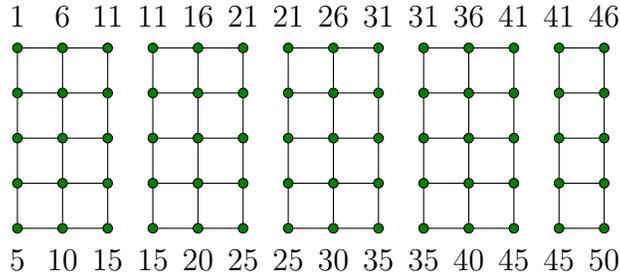


Figure 6: Five induced sub-grids

2.50 GHz processor and 12 GB memory), we only use three of these vertical separators at a time, say the blue separators

$$S_2 = \{11, \dots, 15\}, S_4 = \{21, \dots, 25\}, S_6 = \{31, \dots, 35\}, S_8 = \{41, \dots, 45\}.$$

These separators separate the vertex sets

$$V_1 = \{1, \dots, 15\}, V_3 = \{11, \dots, 25\}, V_5 = \{21, \dots, 35\}, \\ V_7 = \{31, \dots, 45\}, V_9 = \{41, \dots, 50\}.$$

Adding the blue separators to Δ gives a simplicial complex

$$\Delta_{S_2; S_4; S_6; S_8} := \Delta \bigcup_{j=2,4,6,8} \{F : F \subseteq S_j\}$$

with five irreducible components supported on the vertex sets V_1, V_3, V_5, V_7 and V_9 (Figure 7). To compute a facial set with respect to $\Delta_{S_2;S_4;S_6;S_8}$, according to Lemma 3.6 applied four times, we need to compute

$$\begin{aligned} G_{1,V_1} &:= F_{\Delta_{S_2}|_{V_1}}(\pi_{V_1}(I_+)), & G_{1,V_3} &:= F_{\Delta_{S_2;S_4}|_{V_3}}(\pi_{V_3}(I_+)), \\ G_{1,V_5} &:= F_{\Delta_{S_4;S_6}|_{V_5}}(\pi_{V_5}(I_+)), & G_{1,V_7} &:= F_{\Delta_{S_6;S_8}|_{V_7}}(\pi_{V_7}(I_+)), \\ G_{1,V_9} &:= F_{\Delta_{S_8}|_{V_9}}(\pi_{V_9}(I_+)). \end{aligned}$$

Then $G_1 := \bigcap_i \pi_{V_i}^{-1}(G_{1,V_i})$ is equal to $F_{\Delta_{S_2;S_4;S_6;S_8}}(I_+)$, and thus an inner approximation of F_t . As stated before, we do not need to compute G_1 explicitly, but we represent it by means of the G_{1,V_i} .

We can improve the approximations by also considering the red separators

$$S_1 = \{6, \dots, 10\}, \quad S_3 = \{16, \dots, 20\}, \quad S_5 = \{26, \dots, 30\}, \quad S_7 = \{36, \dots, 40\},$$

that separate

$$\begin{aligned} V_0 &= \{1, \dots, 10\}, \quad V_2 = \{6, \dots, 20\}, \quad V_4 = \{16, \dots, 30\}, \\ V_6 &= \{26, \dots, 40\}, \quad V_8 = \{36, \dots, 50\}. \end{aligned}$$

As explained in Section 3.2, we want to compute $G_1^{(2)} := F_{\Delta_{S_1;S_3;S_5;S_7}}(G_1)$. Again, instead of computing $G_1^{(2)}$ directly, we need only compute the much smaller sets $G_{1,V_0}^{(2)} := \pi_{V_0}(G_1^{(2)})$, $G_{1,V_2}^{(2)} := \pi_{V_2}(G_1^{(2)})$, \dots , $G_{1,V_8}^{(2)} := \pi_{V_8}(G_1^{(2)})$. So the question is: Is it possible to compute $G_{1,V_0}^{(2)}$, $G_{1,V_2}^{(2)}$, \dots , $G_{1,V_8}^{(2)}$ from $G_{1,V_1}, G_{1,V_3}, \dots, G_{1,V_9}$, without computing G_1 in between?

It turns out that this is indeed possible: By Lemma 3.6, all we need to compute $G_{1,V_i}^{(2)}$ is $G_{1,V_j} := \pi_{V_j}(G_1)$, $j = i - 1, i + 1$. For $i = 0$, since $V_0 \subset V_1$, we can compute G_{1,V_0} from $\pi_{V_1}(G_1) = G_{1,V_1}$. For $i = 2, 4, 6, 8$, since $V_i \subset V_{i-1} \cup V_{i+1}$, we can compute G_{1,V_i} from $\pi_{V_{i-1} \cup V_{i+1}}(G_1)$, which itself can be obtained by “gluing” $\pi_{V_{i-1}}(G_1) = G_{1,V_{i-1}}$ and $\pi_{V_{i+1}}(G_1) = G_{1,V_{i+1}}$:

$$\pi_{V_{i-1} \cup V_{i+1}}(G_1) = \left(\pi_{V_{i-1}}^{V_{i-1} \cup V_{i+1}} \right)^{-1} (G_{1,V_{i-1}}) \cap \left(\pi_{V_{i+1}}^{V_{i-1} \cup V_{i+1}} \right)^{-1} (G_{1,V_{i+1}}),$$

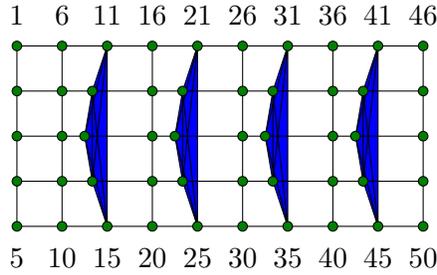
where $\pi_{V''}^{V'}$ for $V'' \subseteq V'$ denotes the marginalization map from $I_{V'}$ to $I_{V''}$ and where $\left(\pi_{V''}^{V'} \right)^{-1}$ denotes the lifting from $I_{V''}$ to $I_{V'}$.

As explained in Section 3.2, we have to iterate this procedure: From $G_1^{(2)}$ we want to compute $G_1^{(3)} := F_{\Delta_{S_2;S_4;S_6;S_8}}(G_1^{(2)})$ or, more precisely, we want to

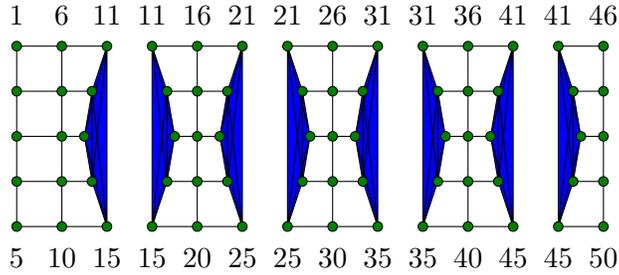
compute $G_{1,V_i}^{(3)} = \pi_{V_i}(G_1^{(3)})$ for $i = 1, 3, \dots, 9$. Again, we do this without looking at $G_1^{(2)}$ directly by just using the information available through the $G_{1,V_i}^{(3)}$. Iterating this procedure, we obtain a sequence of sets $G_{1,V_i}^{(k)}, G_{1,V_j}^{(k)}$ (with odd i and even j), which stabilizes after a finite number of steps. Let

$$F_{1,V_i} := \bigcup G_{1,V_i}^{(k)},$$

Our best inner approximation is then $F_1 = \bigcap_{i=0}^9 \pi_{V_i}^{-1}(F_{1,V_i})$. Again, we do not compute F_1 explicitly, but we represent it in terms of the F_{1,V_i} .



(a)



(b)

Figure 7: (a) The 5×10 -grid with the blue separators completed. (b) The five irreducible subcomplexes after completing the separators.

The process is visualized in Figure 8.

Let us now consider the outer approximation F_2 . We adapt Strategy 3 of Section 3.3 and cover the graph with 5×3 grid subgraphs, since the facial sets for such graphs can easily be computed. These subgrids are supported on the same vertex subsets $V_i, i = 1, \dots, 8$ as used when computing F_1 . This makes it possible to compare F_1 and F_2 . For $i = 1, 3, \dots, 8$ we compute

Table 6: facial set approximation of 5×10 grid graph

sample size	$F_2 \neq I$	$F_1 = F_2$
50	100.0%	94.3%
100	100.0%	82.5%
150	99.9%	76.5%
200	99.6%	81.2%
300	96.4%	87.7%
400	92.9%	91.5%
500	84.8%	93.9%
1000	44.7%	99.9%

$F_{2,V_i} = F_{\Delta|V_i}(\pi_{V_i}(I_+))$. Our outer approximation is then $F_2 = \bigcap_i \pi_{V_i}^{-1}(F_{2,V_i})$. Again, we don't compute F_2 explicitly, but we only store F_{2,V_i} in a computer as a representation of F_2 . To compare the two approximations F_1 and F_2 , we need only compare their projections F_{1,V_i} and F_{2,V_i} pairwise, $i = 1, \dots, 8$.

We generated random data of varying sample size. For each fixed sample size, we generated 100 data samples. The simulation results are shown in Table 6. For each simulated sample, we compute the sets F_{1,V_i} and F_{2,V_i} as described above. When computing F_{1,V_i} , we found that 2 iterations actually suffice. Then we checked whether F_2 is a proper subset of I (second column), and we checked whether $F_1 = F_2$ (third column). enough. Both for small and large sample sizes, we found that the $F_1 = F_2$ quite often.

We also investigated what happens when the outer approximation is not computed using all 3×5 -subgrids, but only a cover of four 3×5 -subgrids and one 2×5 -subgrid (as in Figure 6). In all our simulations, this easier approximation gave the same result. The same is not true for the inner approximation: When using just one of the two families of parallel separators we obtain an inner approximation that is much too small.

References

Onureena Banerjee, Laurent El Ghaoui, and Alexandre d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *The Journal of Machine Learning Research*, 9: 485–516, 2008.

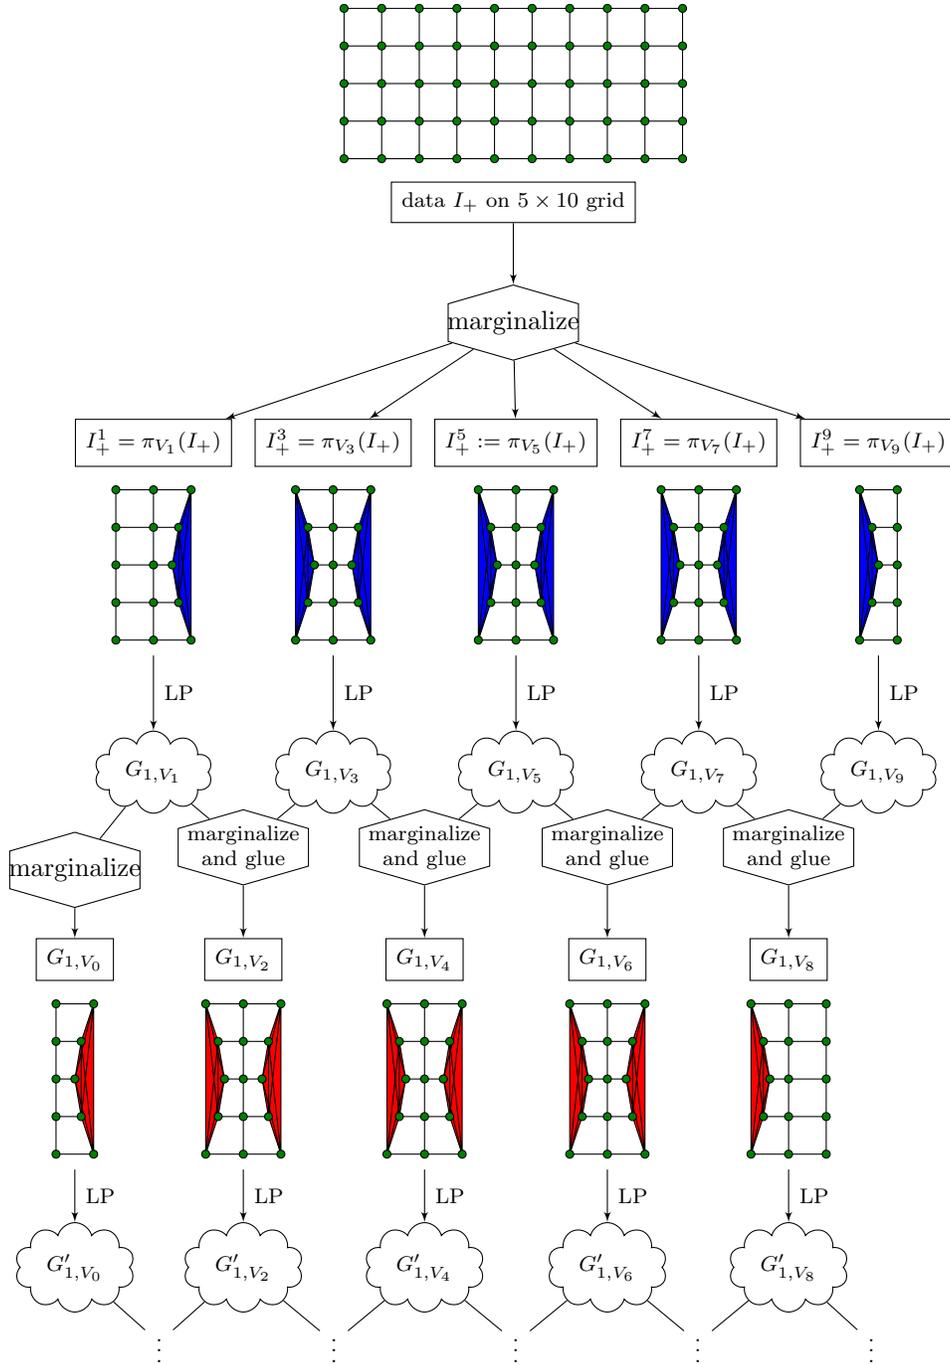


Figure 8: Flow chart

- O.E. Barndorff-Nielsen. *Information and Exponential Families*. Wiley, Chichester, first edition, 1978.
- Imre Csiszár and František Matúš. Closures of exponential families. *Annals of Probability*, 33:582–600, 2005.
- A. Dobra, E.A. Erosheva, and S.E. Fienberg. Disclosure limitation methods based on bounds for large contingency tables with application to disability data. In *Proceedings of conference on the new frontiers of statistical data mining*, pages 93–116, 2003.
- Adrian Dobra and Alex Lenkoski. Copula Gaussian graphical models and their application to modeling functional disability data. *Ann. Appl. Stat.*, 5(2A):969–993, 06 2011.
- N. Eriksson, S. Fienberg, A. Rinaldo, and S. Sullivant. Polyhedral conditions for the nonexistence of the MLE for hierarchical log-linear models. *J. Symbolic Comput.*, 41:222–233, 2006.
- S. E. Fienberg and A. Rinaldo. Three centuries of categorical data analysis: log-linear models and maximum likelihood estimation. *J. of Statistical Planning and Inference*, 137:3430–3445, 2007.
- S. E. Fienberg and A. Rinaldo. Maximum likelihood estimation in log-linear models. *Annals of Statistics*, 40:9961023, 2012.
- Charles J. Geyer. Likelihood inference in exponential families and directions of recession. *Electron. J. Statist.*, 3:259–289, 2009.
- S. J. Haberman. *The Analysis of Frequency Data*. Univ. Chicago Press, Chicago, IL, 1974.
- S.L. Lauritzen. *Graphical Models*. Oxford Science Publications, first edition, 1996.
- G. Letac and H. Massam. Bayes regularization and the geometry of discrete hierarchical loglinear models. *Annals of Statistics*, 40:861–890, 2012.
- H. Massam and N. Wang. A local approach to estimation in discrete loglinear models. *Preprint*, 2015. arXiv:1504.05434.

Johannes Rauh, Thomas Kahle, and Nihat Ay. Support sets of exponential families and oriented matroids. *International Journal of Approximate Reasoning*, 52(5):613–626, 2011.

Pradeep Ravikumar, Martin J Wainwright, John D Lafferty, et al. High-dimensional ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.

Günter Ziegler. *Lectures on Polytopes*. Springer, second edition, 1998.

A Some proofs

A.1 Proof of Theorem 2.8

Theorem 2.8 goes back to [Barndorff-Nielsen \(1978\)](#), who studies the closure of much more general exponential families. The case of a discrete exponential family is much easier.

The theorem follows from the following lemmas:

Lemma A.1. *Let $p \in \overline{\mathcal{E}}_A$. Then $p \in \mathcal{E}_{A, \text{supp}(p)}$.*

Lemma A.2. *Let $p \in \overline{\mathcal{E}}_A$. Then $\mathcal{E}_{A, \text{supp}(p)} \subseteq \overline{\mathcal{E}}_A$.*

Lemma A.3. *Let $p \in \overline{\mathcal{E}}_A$. Then $\text{supp}(p)$ is facial.*

Lemma A.4. *If F is facial, then there exists $p \in \overline{\mathcal{E}}_A$ with $\text{supp}(p) = F$.*

Indeed, Lemma A.1 shows that $\overline{\mathcal{E}}_A \subseteq \bigcup_F \mathcal{E}_{A,F}$, where the union is over all support sets F . Lemma A.2 shows the converse containment, so that $\overline{\mathcal{E}}_A = \bigcup_F \mathcal{E}_{A,F}$. It remains to see that a subset $F \subseteq I$ is a support set if and only if F is facial. This follows from Lemmas A.3 and A.4.

In the proofs of Lemmas A.1 to A.4, we need the following easy lemma of which we omit the proof:

Lemma A.5. *$p \in \mathcal{E}_A$ if and only if $\log(p) \perp \ker A$.*

Proof of Lemma A.1. Let $p = \lim_{k \rightarrow \infty} p_k$, where $p_k \in \mathcal{E}_A$, and let $F = \text{supp}(p)$. Then $\mathcal{E}_{A,F}$ is the exponential family \mathcal{E}_{A_F} , where A_F consists of

the columns of A indexed by F . Any $v \in \ker A_F$ can be extended by zeros to $v' \in \ker A$. By Lemma A.5,

$$0 = \langle \log(p_k), v' \rangle = \sum_{i \in F} \log(p_k(i))v(i) \rightarrow \langle \log(p), v \rangle.$$

Thus, $\log(p) \perp \ker A_F$, which implies $p \in \mathcal{E}_{A,F}$. \square

Proof of Lemma A.2. Let $p = \lim_{k \rightarrow \infty} p_k$, where $p_k \in \mathcal{E}_A$, let $F = \text{supp}(p)$, and let $q \in \mathcal{E}_{A,F}$. Then there exists parameters θ with $\log(q(i)) - \log(p(i)) = \langle \theta, f_i \rangle$ for all $i \in F$. For any k , there exists a positive constant c_k such that $q_k := c_k p_k \exp(\langle \theta, A \rangle) \in \mathcal{E}_A$. Then $q_k \rightarrow q$ as $k \rightarrow \infty$, and so $q \in \overline{\mathcal{E}_A}$. \square

Proof of Lemma A.3. Let $p = \lim_{k \rightarrow \infty} p_k$, where $p_k \in \mathcal{E}_A$, and let $F = F_A(\text{supp}(p))$. Then $x = \frac{1}{|\text{supp}(p)|} \sum_{i \in \text{supp}(p)} f_i$ is an interior point of the face corresponding to F , and thus there exist positive coefficients $\lambda_i > 0$, $i \in F$, with $x = \sum_{i \in F} \lambda_i f_i$. The vector $v = (v_i, i \in I)$ defined by

$$v_i = \begin{cases} \frac{1}{|\text{supp}(p)|} - \lambda_i, & i \in \text{supp}(p), \\ -\lambda_i, & i \in F \setminus \text{supp}(p), \\ 0, & i \notin F, \end{cases}$$

satisfies $Av = x - x = 0$. By Lemma A.5, $\log(p_k) \perp v$ for all k . In particular,

$$\sum_{i \in F \setminus \text{supp}(p)} \lambda_i \log(p_k(i)) = \sum_{i \in \text{supp}(p)} \log(p_k(i))v_i \rightarrow \sum_{i \in \text{supp}(p)} \log(p(i))v_i.$$

On the other hand, note that each coefficient λ_i for $i \in F \setminus \text{supp}(p)$ on the left hand side is positive, while $\log(p_k(i)) \rightarrow -\infty$ for $i \notin \text{supp}(p)$. This shows that $F \setminus \text{supp}(p) = \emptyset$. \square

Proof of Lemma A.4. If F is facial, there exist $g \in \mathbf{R}^h$ and $c \in \mathbf{R}$ with $\langle g, f_i \rangle \geq c$ for all $i \in I$ and $\langle g, f_i \rangle = c$ if and only if $i \in F$. Let $\theta^{(s)} = -s \cdot g$. Then

$$k_F(\theta^{(s)}) + sc = \log \sum_{i \in I} \exp(-s \langle g, f_i \rangle + sc) \rightarrow \log |F|,$$

and so

$$\begin{aligned} \log p_{\theta^{(s)}}(i) &= -s \langle g, f_i \rangle - k_F(\theta^{(s)}) = (sc - s \langle g, f_i \rangle) - (k_F(\theta^{(s)}) + sc) \\ &\rightarrow \begin{cases} -\log |F|, & \text{if } i \in F, \\ -\infty, & \text{if } i \notin F, \end{cases} \end{aligned}$$

as $s \rightarrow \infty$. Thus, $p_{\theta^{(s)}}$ converges to the uniform distribution on F . \square

A.2 Proof of Theorem 2.9

By definition, any EMLE p_* belongs to the closure of the model. According to Theorem 2.8, the support of p_* is facial. If $\text{supp}(p)$ does not contain $\text{supp}(n)$, then the log-likelihood goes to minus infinity, $l(p) = -\infty$, and so p does not maximize the likelihood. Therefore, $\text{supp}(p_*)$ is a facial set containing $\text{supp}(n)$. Thus, $F_t \subseteq \text{supp}(p_*)$.

By Lemma A.1, p_* belongs to $\mathcal{E}_{\Delta, \text{supp}(p_*)}$, which is parametrized by a vector θ , see Theorem 2.8. On $\mathcal{E}_{\Delta, \text{supp}(p_*)}$, the log-likelihood function in terms of this parameter θ is

$$l_F(\theta) = \sum_{j \in J} \theta_j t_j - Nk_F(\theta).$$

l_F is strictly concave, and so it has a unique maximum. The critical equations are

$$Ap_* = \frac{t}{N},$$

proving the first property. Note that these equations are independent of the parameters and the support of p_* . We now show that any solution to these equations is supported on the same face of \mathbf{P} as $\frac{t}{N}$.

Let p be a probability distribution on I such that $\text{supp}(p)$ does not contain F_t . This means that there is a linear inequality $\langle g, t \rangle \geq c$ that is valid on \mathbf{P} and such that

- $\langle g, f_i \rangle = c$ for all $i \in F_t$;
- $\langle g, f_i \rangle > c$ for some $i \in \text{supp}(p)$.

Then

$$\langle g, Ap \rangle = \sum_i \langle g, f_i \rangle p(i) > c = \frac{1}{N} \sum_i n(i) \langle g, f_i \rangle = \langle g, \frac{t}{N} \rangle,$$

which implies $Ap \neq \frac{t}{N}$. This shows $\text{supp}(p_*) \subseteq F_t$ and finishes the proof of $\text{supp}(p_*) = F_t$.

We have now shown the two properties, and it remains to argue that the EMLE is unique. But this follows from the fact that $\text{supp}(p_*)$ is equal to F_t , and l_F is strictly convex, such that the likelihood has a unique maximizer on $\mathcal{E}_{\Delta, F_t}$.

B Example: Two binary random variables

Consider two binary random variables, and let $\Delta = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$. The hierarchical model \mathcal{E}_Δ is the *saturated model*; that is, it contains all possible probability distributions with full support. Then

$$\tilde{A} = \begin{pmatrix} \overbrace{1}^{f_{00}} & \overbrace{1}^{f_{01}} & \overbrace{1}^{f_{10}} & \overbrace{1}^{f_{11}} \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{matrix} \theta_{00} \\ \theta_{01} \\ \theta_{10} \\ \theta_{11} \end{matrix}$$

The marginal polytope is a 3-simplex (a tetrahedron) with facets

$$\begin{aligned} \mathbf{F}_{00} : 1 - t_{01} - t_{10} + t_{11} &\geq 0, & \mathbf{F}_{01} : t_{01} - t_{11} &\geq 0, \\ \mathbf{F}_{10} : t_{10} - t_{11} &\geq 0, & \mathbf{F}_{11} : t_{11} &\geq 0. \end{aligned}$$

Each of the corresponding facets contains three columns of \tilde{A} . In fact, the facet \mathbf{F}_i in the above list does not contain the column f_i of A .

The EMLE of the saturated model is just the empirical distribution; that is, $p_* = \frac{1}{N}n$. Suppose that t lies on the facet \mathbf{F}_{00} (i.e. $n = (0, n_{01}, n_{10}, n_{11})$ with $n(01), n(10), n(11) > 0$). If $p_{\theta^{(s)}} \rightarrow p_*$, then $p_{\theta^{(s)}}(00) \rightarrow 0$, while all other probabilities converge to a non-zero value. It follows that

$$\begin{aligned} \theta_{00}^{(s)} &= \log p_{\theta^{(s)}}(00) \rightarrow -\infty, \\ \theta_{01}^{(s)} &= \log \frac{p_{\theta^{(s)}}(01)}{p_{\theta^{(s)}}(00)} \rightarrow +\infty, \\ \theta_{10}^{(s)} &= \log \frac{p_{\theta^{(s)}}(10)}{p_{\theta^{(s)}}(00)} \rightarrow +\infty, \\ \theta_{11}^{(s)} &= \log \frac{p_{\theta^{(s)}}(11)p_{\theta^{(s)}}(00)}{p_{\theta^{(s)}}(01)p_{\theta^{(s)}}(10)} \rightarrow -\infty. \end{aligned}$$

On the other hand, $\theta_{01}^{(s)} + \theta_{00}^{(s)} = \log p_{\theta^{(s)}}(01)$ converges to a finite value, as do $\theta_{10}^{(s)} + \theta_{00}^{(s)} = \log p_{\theta^{(s)}}(10)$ and $\theta_{11}^{(s)} + \theta_{01}^{(s)} = \log p_{\theta^{(s)}}(11)/p_{\theta^{(s)}}(10)$.

Proceeding similarly for the other facets, one can show for the limits $\theta_{ij} := \lim_{s \rightarrow \infty} \theta_{ij}^{(s)}$:

	θ_{00}	θ_{01}	θ_{10}	θ_{11}	finite parameter combinations:
\mathbf{F}_{00}	$-\infty$	$+\infty$	$+\infty$	$-\infty$	$\theta_{01}^{(s)} + \theta_{00}^{(s)}, \theta_{10}^{(s)} + \theta_{00}^{(s)}, \theta_{11}^{(s)} + \theta_{01}^{(s)}$
\mathbf{F}_{01}	finite	$-\infty$	finite	$+\infty$	$\theta_{00}^{(s)}, \theta_{10}^{(s)}, \theta_{01}^{(s)} + \theta_{11}^{(s)}$
\mathbf{F}_{10}	finite	finite	$-\infty$	$+\infty$	$\theta_{00}^{(s)}, \theta_{01}^{(s)}, \theta_{10}^{(s)} + \theta_{11}^{(s)}$
\mathbf{F}_{11}	finite	finite	finite	$-\infty$	$\theta_{00}^{(s)}, \theta_{10}^{(s)}, \theta_{01}^{(s)}$

Each line of the last column contains three combinations of the parameters $\theta_i^{(s)}$ that converge to a finite value. Any other parameter combination that converges is a linear combination of these three. This can be seen by using the coordinates μ_i introduced in Section 4.2 and applying Lemma 4.1. For example, on the facet \mathbf{F}_{01} , consider the parameters

$$\begin{aligned}\mu_{10} &= \log p(10)/p(00) = \theta_{10}, & \mu_{11} &= \log p(11)/p(00) = \theta_{10} + \theta_{01} + \theta_{11}, \\ \mu_{01} &= \log p(01)/p(00) = \theta_{01}.\end{aligned}$$

Then μ_{10} and μ_{11} are identifiable parameters on $\mathcal{E}_{F_{01}}$, and μ_{01} diverges close to \mathbf{F}_{01} . By Lemma 4.1, the linear combinations that are well-defined are $\mu_{10} = \theta_{10}$ and $\mu_{11} = \theta_{10} + (\theta_{01} + \theta_{11})$. The above table also lists θ_{00} , which is not a linear combination of those but that is fine because it is not free.

We obtain similar results for the facets \mathbf{F}_{01} and \mathbf{F}_{11} . The results are summarized in the following table:

facet	μ_{01}	μ_{10}	μ_{11}
\mathbf{F}_{01}	$-\infty$	finite	finite
\mathbf{F}_{10}	finite	$-\infty$	finite
\mathbf{F}_{11}	finite	finite	$-\infty$

Of course, by definition of the μ_i s, we cannot consider the facet \mathbf{F}_{00} where $n(00) = 0$. To study \mathbf{F}_{00} , we have to choose another zero cell and redefine the parameters μ_i .

The situation is more complicated for faces smaller than facets, because sending a single parameter to plus or minus infinity can be enough to send the distribution to a face F of higher codimension, as we will see below. The remaining parameters then determine the position within $\mathcal{E}_{\Delta, F}$. Thus, in this case there are more remaining parameters than the dimension of $\mathcal{E}_{\Delta, F}$.

For example, the data vector $n = (n_{00}, 0, n_{10}, 0)$ (with $n_{00}, n_{10} > 0$) lies

on the face $\mathbf{F} = \mathbf{F}_{01} \cap \mathbf{F}_{11}$ of codimension two. If $p_{\theta^{(s)}} \rightarrow p_*$, then

$$\begin{aligned}\theta_{00}^{(s)} &= \log p_{\theta^{(s)}}(00) \rightarrow \log \frac{n_{00}}{N}, \\ \theta_{01}^{(s)} &= \log \frac{p_{\theta^{(s)}}(01)}{p_{\theta^{(s)}}(00)} \rightarrow -\infty, \\ \theta_{10}^{(s)} &= \log \frac{p_{\theta^{(s)}}(10)}{p_{\theta^{(s)}}(00)} \rightarrow \log \frac{n_{10}}{n_{00}}.\end{aligned}$$

However, the limit of $\theta_{11}^{(s)} = \log \frac{p_{\theta^{(s)}}(11)p_{\theta^{(s)}}(00)}{p_{\theta^{(s)}}(01)p_{\theta^{(s)}}(10)}$ is not determined. The only constraint is that $\theta_{11}^{(s)}$ cannot go to $+\infty$ faster than $\theta_{01}^{(s)}$ goes to $-\infty$, since $p_{\theta^{(s)}} = \exp(\theta_{00}^{(s)} + \theta_{01}^{(s)} + \theta_{10}^{(s)} + \theta_{11}^{(s)})$ has to converge to zero.

With the same data vector $n = (n_{00}, 0, n_{10}, 0)$, suppose we use a numerical algorithm to optimize the likelihood function by optimizing the parameters θ_j in turn. To be precise, we order the parameters θ_j in some way. For simplicity, say that the parameters are $\theta_1, \theta_2, \dots, \theta_h$. Then we let

$$\theta_j^{(k+1)} = \arg \max_{y \in \mathbf{R}} l(\theta_1^{(k+1)}, \dots, \theta_{j-1}^{(k+1)}, y, \theta_{j+1}^{(k)}, \dots, \theta_h^{(k)})$$

(this is called the *non-linear Gauss-Seidel method*). Let us choose the ordering $\theta_{01}, \theta_{10}, \theta_{11}$ (note that $\theta_{00} = -k(\theta)$ is not a free parameter). We start at $\theta_{01}^{(0)} = \theta_{10}^{(0)} = \theta_{11}^{(0)} = 0$. In the first step, we only look at θ_{01} . That is, we want to solve

$$\begin{aligned}0 &= \frac{\partial}{\partial \theta_{01}} l(\theta) = -\frac{\exp(\theta_{01}^{(1)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(0)} + \theta_{11}^{(0)})}{1 + \exp(\theta_{01}^{(1)}) + \exp(\theta_{10}^{(0)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(0)} + \theta_{11}^{(0)})} \\ &= -\frac{2 \exp(\theta_{01}^{(1)})}{1 + 2 \exp(\theta_{01}^{(1)})}.\end{aligned}\quad (21)$$

Clearly, the derivative is negative for any finite value of $\theta_{01}^{(1)}$, and thus the critical equation has no finite solution. If we try to solve this equation numerically, we will find that $\theta_{01}^{(1)}$ will be a large negative number. Next, we look at θ_{10} . We fix the other variables and try to solve

$$\begin{aligned}0 &= \frac{\partial}{\partial \theta_{10}} l(\theta) = \frac{n_{10}}{N} - \frac{\exp(\theta_{10}^{(1)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(1)} + \theta_{11}^{(0)})}{1 + \exp(\theta_{01}^{(1)}) + \exp(\theta_{10}^{(1)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(1)} + \theta_{11}^{(0)})} \\ &\approx \frac{n_{10}}{N} - \frac{\exp(\theta_{10}^{(1)})}{1 + \exp(\theta_{10}^{(1)})},\end{aligned}$$

where we have used that $\theta_{01}^{(1)}$ is a large negative number. This equation always has a unique solution

$$\theta_{10}^{(1)} \approx \log \frac{n_{10}}{N - n_{10}}.$$

Finally, we look at θ_{11} . We have to solve

$$0 = \frac{\partial}{\partial \theta_{11}} l(\theta) = - \frac{\exp(\theta_{01}^{(1)} + \theta_{10}^{(1)} + \theta_{11}^{(1)})}{1 + \exp(\theta_{01}^{(1)}) + \exp(\theta_{10}^{(1)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(1)} + \theta_{11}^{(1)})} \approx 0.$$

Actually, this equation again has no solution, and the numerical solution for $\theta_{11}^{(1)}$ should be close to numerical minus infinity. However, since $\theta_{01}^{(1)}$ is already close to $-\infty$, the equation is already approximately satisfied. Thus, there is no need to change θ_{11} . In simulations, we observed that usually $\theta_{11}^{(1)}$ will be negative, but not as negative as $\theta_{01}^{(1)}$. In theory, we would have to iterate and now optimize θ_{01} again. But the values will not change much, since the critical equations are already satisfied to a high numerical precision after one iteration.

It is not difficult to see that the result is different if we change the order of the variables. If θ_{11} is optimized before θ_{01} , then θ_{11}^1 will in any case be a large negative number.

For general data, the derivative of with respect to θ_{01} (equation (21)) takes the form

$$\frac{\partial}{\partial \theta_{01}} l(\theta) = \frac{t_{01}}{N} - \frac{\exp(\theta_{01}^{(1)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(0)} + \theta_{11}^{(0)})}{1 + \exp(\theta_{01}^{(1)}) + \exp(\theta_{10}^{(0)}) + \exp(\theta_{01}^{(1)} + \theta_{10}^{(0)} + \theta_{11}^{(0)})}.$$

Setting this derivative to zero and solving for $\theta_{01}^{(1)}$ leads to a linear equation in $\theta_{01}^{(1)}$ with symbolic solution

$$\theta_{01}^{(1)} = \log \frac{1 + \exp(\theta_{10}^{(0)})}{1 + \exp(\theta_{10}^{(0)} + \theta_{11}^{(0)})} \frac{\frac{t_{01}}{N}}{1 - \frac{t_{01}}{N}}.$$

In fact, for any hierarchical model, the likelihood equation is linear in any single parameter θ_j , as long as all other parameters are kept fixed (more generally this is true when the design matrix A is a 0-1-matrix). Instead of optimizing the likelihood numerically with respect to one parameter, it is possible to use these symbolic solutions. This leads to the Iterative Proportional Fitting Procedure (IPFP). In our example, the IPFP would lead to a division by zero right in the first step, indicating that the MLE does not exist.

C A Linear Programming algorithm to compute facial set

Denote A as the design matrix, A_+ as the sub-matrix with columns indexed by the positive cells and A_0 as the sub-matrix indexed by the empty cells.

Lemma C.1. *Solution g^* of the non-linear problem*

$$\begin{aligned} \max \quad & z = \|Ag\|_0 \\ \text{s.t.} \quad & A_+g = 0 \\ & A_0g \geq 0 \end{aligned} \tag{22}$$

is a perpendicular vector to the smallest face containing t . The corresponding facial set is $F_t = I \setminus \text{supp}(Ag^*)$.

The optimization problem (22) is highly non-linear and non-convex: it can be solved by repeatedly solving the associated ℓ_1 -norm optimization problem:

$$\begin{aligned} \max \quad & z = \|A_0g\|_1 \\ \text{s.t.} \quad & A_+g = 0 \\ & A_0g \geq 0 \\ & A_0g \leq 1 \end{aligned} \tag{23}$$

Problem (23) is a linear programming problem: we can solve it repeatedly until we get the smallest facial set F_t . The process is as follows:

The algorithm is introduced in the supplementary material to (Fienberg and Rinaldo, 2012), where it is also proved that it outputs the correct result.

D Parametrizations adapted to facial sets

Let us briefly discuss how to remedy problems 1. to 3. from the beginning of Section 4.2. The idea to remedy 1. and 2. is to define parameters μ_i , $i \in L$, of \mathcal{E}_A , such that a subset $L_t \subseteq L$ of the parameters parametrizes $\mathcal{E}_{F_t, A}$ in a consistent way. Denote by $A^\mu = (a_{j,i}^\mu, j \in L, i \in I)$ the design matrix of \mathcal{E}_A corresponding to the new parameters μ . Then the necessary conditions are:

- (*) Let $A_{L_t, F_t}^\mu := (a_{j,i}^\mu, j \in L_t, i \in F_t)$ be the submatrix of A^μ with rows indexed by L_t and columns indexed by L_t , and denote by \tilde{A}_{L_t, F_t}^μ the same matrix with an additional row of ones. The rank of \tilde{A}_{L_t, F_t}^μ is equal to $|L_t| + 1$, the number of its rows (and thus, A_{L_t, F_t}^μ has rank $|L_t|$).

Algorithm 1 Face computation by Linear programming method

Require: Design matrix A and positive cell index I_+

INITIALIZE $A_+ = A(I_+, :)$, $A_0 = A \setminus A_+$

Solve problem 23, get the solution g^* and the corresponding maximum z^*

while $A_0 \neq \emptyset$ and $z^* \neq 0$ **do**

Let matrix B be the submatrix of A_0 , by taking columns of A_0 which satisfy $\langle f_i, g^* \rangle > 0$, update $A_0 = A_0 \setminus B$,

Solve problem 23, get the solution g^* and the corresponding maximum z^*

end while

if $A_0 = \emptyset$ **then**

$F_t = I_+$

end if

if $Z^* = 0$ **then**

$F_t = I_+ \cup \{i | i \text{ is the index of } A_0\}$

end if

(**) $a_{j,i}^\mu = 0$ for all $i \in F_t$ and $j \in L \setminus L_t$.

In fact, (**) implies that A_{L_t, F_t}^μ is the design matrix of \mathcal{E}_{A, F_t} , since the parameters μ_i with $i \notin L_t$ do not play a role in the parametrization $\mu \mapsto p_{F_t, \mu}$. Moreover, (*) implies that the parametrization $\mu \mapsto p_{F_t, \mu}$ is identifiable. In this sense, we have remedied problem 1. from the beginning of the section.

Since the matrix \tilde{A}_{L_t, F_t}^μ has full row rank, it has a right inverse matrix \tilde{C} , such that $\tilde{A}_{L_t, F_t}^\mu \tilde{C} = I_{|L_t|+1}$ equals the identity matrix of size $|L_t| + 1$. Recall that

$$\begin{aligned} \log p_{F_t, \mu}(i) &= \langle \mu^t, f_i^\mu \rangle - k_F(\mu), \\ \log p_\mu(i) &= \langle \tilde{\mu}^t, f_i^\mu \rangle - k(\mu), \end{aligned}$$

for any parameter vector μ and all $i \in F_t$. Since f_i^μ are the columns of A^μ and since the components of f_i^μ corresponding to $L \setminus L_t$ vanish by (**), we may apply the matrix C obtained from \tilde{C} by dropping the row corresponding to k_F or k and obtain

$$(\log p_\mu)C = \mu_{L_t} \quad \text{and} \quad (\log p_{F_t, \mu})C = \mu_L. \quad (24)$$

When $p_{\mu^{(s)}}$ is a sequence in \mathcal{E}_A with limit p_μ in $\mathcal{E}_{F_t, A}$, then (24) shows that $\mu_i^{(s)} \rightarrow \mu_i$ for $i \in L_t$. In this sense, we have remedied problem 2.

Finally, we solve problem 3. Suppose that we have chosen parameters μ_L as in Section 4.2, and let A^{μ_L} be the design matrix with respect to these parameters. Then $(A^{\mu_L})_{j,i} = 0$ if $i \in F_t$ and $j \notin L_t$. Moreover, for $j \in L_t$, the j th column of A_{μ_L} has a single non-vanishing entry (equal to one) at position j . Suppose that F_t corresponds to a face \mathbf{F}_t of codimension c . Then there are c facets of \mathbf{P} whose intersection is \mathbf{F}_t . Thus, following the notation introduced in Remark 2.6, there exist c inequalities

$$\langle \tilde{g}_1, \tilde{x} \rangle \geq 0, \quad \dots, \quad \langle \tilde{g}_c, \tilde{x} \rangle \geq 0 \quad (25)$$

that together define \mathbf{F}_t . In this case, the vectors $\tilde{g}_1, \dots, \tilde{g}_c$ are linearly independent and satisfy $\langle \tilde{g}_j, \tilde{f}_i \rangle = 0$ (thus, they are a basis of the kernel of $(\tilde{A}_{F_t}^{\mu_L})^t$). It follows that the k th component of g_j , denoted by $g_{j,k}$, vanishes if $k \in L_t$; that is, the inequalities (25) do not involve the variables corresponding to L_t . Let G be the square matrix, indexed by $L \setminus L_t$ with entries $g_{j,k}$, $j, k \in L \setminus L_t$. Then the square matrix

$$\tilde{G} = \begin{pmatrix} 1 & 0 \\ 0 & G \end{pmatrix}$$

is invertible. We claim that the parameters $\lambda = \tilde{G}^{-1}\mu_L$ are what we are looking for.

The design matrix with respect to the parameters λ is $A^\lambda = \tilde{G}A^{\mu_L}$. For any $j \notin L_t$,

$$A_{j,i}^\lambda = 0, \quad \text{if } i \in F_t, \quad \text{and} \quad A_{j,i}^\lambda = \langle \tilde{g}_j, \tilde{f}_i \rangle \geq 0, \quad \text{if } i \notin F_t.$$

This implies the following properties:

1. If all parameters λ_j with $j \notin L_t$ are sent to $-\infty$, then p_λ tends towards a limit distribution with support F_t .
2. The coefficient of λ_j in any log-probability is non-negative, so there is no cancellation of $\pm\infty$.

So far, we only used the fact that the vectors \tilde{g}_j define valid inequalities for the face \mathbf{F}_t . Suppose that we choose \tilde{g}_j in such a way that each inequality $\langle \tilde{g}_j, \tilde{x} \rangle \geq 0$ defines a facet. The intersection of less than c facets is a face that strictly contains \mathbf{F}_t . This implies that for each j , there exists $i_j \in I \setminus F_t$ such that f_{i_j} satisfies

$$\langle \tilde{g}_j, \tilde{f}_{i_j} \rangle > 0, \quad \text{and} \quad \langle \tilde{g}_{j'}, \tilde{f}_{i_j} \rangle = 0 \text{ for all } j' \neq j.$$

This implies

$$A_{j,i_j}^\lambda > 0, \quad \text{and} \quad A_{j',i_j}^\lambda = 0 \text{ for all } j' \neq j.$$

This implies the following:

3. If $\lambda_j^{(s)}$ are sequences of parameters such that $p_{\lambda^{(s)}}$ tends towards a limit distribution with support F_t , then $\lambda_j^{(s)} \rightarrow -\infty$ for all $j \notin L_t$.

It is not difficult to see that, conversely, any parametrization that satisfies these three properties comes from facets defining the face \mathbf{F}_t .