

Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig

Dimension of Marginals of Kronecker Product
Models; Geometry of hidden-visible products of
exponential families

by

Guido Montúfar and Jason Morton

Preprint no.: 75

2015



Dimension of Marginals of Kronecker Product Models

Geometry of hidden-visible products of exponential families

Guido Montúfar^{*1} and Jason Morton^{†2}

¹Max Planck Institute for Mathematics in the Sciences, Inselstraße 22, 04103 Leipzig, Germany

²Department of Mathematics, Pennsylvania State University, PA 16802, University Park, USA

November 10, 2015

Abstract

A Kronecker product model is the set of visible marginal probability distributions of an exponential family whose sufficient statistics matrix factorizes as a Kronecker product of two matrices, one for the visible variables and one for the hidden variables. We estimate the dimension of these models by the maximum rank of the Jacobian in the limit of large parameters. The limit is described by the tropical morphism; a piecewise linear map with pieces corresponding to slicings of the visible matrix by the normal fan of the hidden matrix. We obtain combinatorial conditions under which the model has the expected dimension, equal to the minimum of the number of natural parameters and the dimension of the ambient probability simplex. Additionally, we prove that the binary restricted Boltzmann machine always has the expected dimension.

Keywords: expected dimension, tropical geometry, secant variety, restricted Boltzmann machine, inference function, Kronecker product, Hadamard product, Khatri-Rao product, exponential family harmonium

MSC2010: 14T05, 52B05

1 Introduction

Simple probability distributions are often composed in order to obtain more interesting or complex probability distributions. Natural compositions include tensor products, convex combinations, and renormalized entrywise products. Stochastic neural networks, for example, define compositions of elementary probability distributions on the states of individual neurons, which result in interesting joint probability distributions on the states of subsets of neurons of the entire network. We study probability models defined by building the Kronecker product of the sufficient statistics matrices of

^{*}montufar@mis.mpg.de

[†]morton@math.psu.edu

two exponential families and marginalizing over one of the two. We call these probability models hidden-visible Kronecker products of exponential families, or simply Kronecker product models. Examples include mixtures of exponential families, restricted Boltzmann machines [10], and, more generally, hierarchical models on Cartesian products of simplicial complexes of hidden and visible variables. A related class of probability models is known in machine learning under the name exponential family harmonium [27].

We are interested in the dimension of these compositions, depending on the properties of the visible and hidden factors. Marginalization is in general a non-injective map which may collapse the dimension of the probability model that is being marginalized. Well known examples of this behavior are mixture models, which may have a dimension that is strictly smaller than the dimension of the corresponding non-marginalized models and the ambient probability simplex. Many interesting mixture models correspond to secant varieties and as such their dimension has been subject of numerous works in algebraic geometry. Examples include the secants of Segre products [2, 5, 21] and Veronese varieties [1, 21]. Reference texts are [28, 16]. Computing the dimension of secant varieties is a notoriously difficult problem, even in cases where the basis variety is of striking simplicity, such as the Segre products of one-dimensional projective spaces (corresponding to binary independence models in statistics or to the decomposable $2 \times \dots \times 2$ tensors in signal processing). In turn, computing the dimension of marginals of general exponential families seems near to hopeless. In certain cases, however, the dimension of secant varieties can be estimated by solving linear programs. The main idea has been illuminated in the tropical approach to secant dimensions by Draisma [8]. Depending on the basis variety, the dimension estimates resulting from this approach give in fact the exact dimension. For Segre varieties, a sufficient condition is the existence of error correcting codes of a certain cardinality. Following these ideas, the dimension of binary and discrete restricted Boltzmann machines (Hadamard products of secant varieties of Segre products) have been studied in [6] and [19], respectively. The goal of the present paper is to elucidate the approach from the perspective of exponential families, especially hierarchical models, and to provide explicit results for Kronecker product models (Definition 7) generalizing some of the above mentioned work on secants and restricted Boltzmann machine dimensions.

Consider two finite integers $N, M \in \mathbb{N}$ and a probability model $\{p_\theta : \theta \in \mathbb{R}^M\} \subseteq \Delta_{N-1}$, parametrized by a function $\phi: \mathbb{R}^M \rightarrow \Delta_{N-1}; \theta \mapsto p_\theta$ from M -dimensional Euclidean space \mathbb{R}^M to the $(N-1)$ -dimensional probability simplex $\Delta_{N-1} := \{p \in \mathbb{R}^N : p(x) \geq 0 \text{ for all } x \text{ and } \sum_{x=1}^N p(x) = 1\}$. The dimension of $\phi(\mathbb{R}^M)$ is given by the maximum rank of the Jacobian J_ϕ of the map ϕ over the points of \mathbb{R}^M where ϕ is smooth. Hence the problem reduces to computing the maximum of the rank function over a parametric set of matrices. In general, computing the rank of these matrices is difficult. A possible approach is to maximize the rank only over a subset of parameters where the Jacobian is simpler. When the result matches the maximum possible value, equal to the dimension M of the parameter domain \mathbb{R}^M or to the dimension $N-1$ of the codomain Δ_{N-1} , we can be sure that we have attained the global maximum.

The rank becomes tractable when there is an obvious way to transform the Jacobian into a suitable block matrix by elementary matrix operations. Consider the case

where ϕ parametrizes an exponential family, i.e., $\phi(\theta) = \exp(\langle \theta, F \rangle - \psi(\theta))$ for some matrix $F \in \mathbb{R}^{M \times N}$ of sufficient statistics, where ψ is the normalizing (log-partition) function. Note that $\phi(\alpha\theta) \propto \phi^\alpha(\theta)$ for each $\alpha \in \mathbb{R}$. Hence, in the limit of large parameters ($\alpha \rightarrow \infty$), the resulting probability vector is proportional to the indicator of $\operatorname{argmax}(\theta, F)$. The optimizing set is piecewise constant on θ and usually has a small cardinality. Accordingly, if ϕ parametrizes marginals of an exponential family, then, in the limit of large α , the matrix $J_\phi(\alpha\theta)$ has a block structure and its rank may be easy to determine. The mathematical formalism describing this limit of large parameters is known as tropical geometry. The tropicalization of ϕ produces a piecewise linear version of ϕ that captures its combinatoric properties. In this approach, the dimension of the models under consideration can be related to polyhedral optimization problems. This leads to combinatorial problems about optimal slicings of point configurations by polyhedral fans of point configurations. Characterizing the sets of maximizers of vectors in the row span of F and the combinatorics of the convex support of the exponential family (the convex hull of the columns of F) is not an easy task in general. However, we do not need to fully solve that problem in order to estimate the maximum rank.

As mentioned above, we will place emphasis on Kronecker product models with factors given by hierarchical models. Hierarchical models are ubiquitous in statistics applications [26]. An algebraic perspective on hierarchical graphical models was given in [11]. Some works have discussed the convex support polytopes of hierarchical models, showing relations to linear codes and oriented matroids [14, 22] and describing properties such as neighborliness and simpliciality [13, 18].

This paper is organized as follows. Section 2 provides basic definitions of marginals of discrete exponential families and discusses the Jacobian of their natural parametrization. It also introduces the associated inference functions and tropical models. Section 3 defines the Kronecker product model and discusses its general properties. The sections that follow are devoted to specific types of Kronecker product models. Section 4 gives a brief introduction to hierarchical models and defines related types of error correcting codes. Section 5 studies the tropical dimension of mixtures of hierarchical models. Section 6 studies the tropical dimension of Hadamard products of mixtures of hierarchical models. Section 7 studies the general case of Kronecker product hierarchical models. Theorem 24 from Section 7 includes the Theorems 16 and Theorem 22 from Sections 5 and 6 as special cases. Section 8 proves that binary restricted Boltzmann machines always have the expected dimension, thereby solving the dimension question for cases that were left open in [6]. Section 9 discusses the results.

2 Marginals of Exponential Families

In this section we present basic definitions of exponential families and their marginals. We discuss the Jacobian of the natural parametrization and relate its behavior in the limit of large parameters with the inference functions of the model. This leads us to the definition of the tropical morphism and a simplified rank estimation problem.

Consider two finite sets \mathcal{X} and \mathcal{Y} . A probability distribution on $\mathcal{X} \times \mathcal{Y}$ is a real-valued vector $p \in \mathbb{R}^{\mathcal{X} \times \mathcal{Y}}$ with entries $p(x, y) \geq 0$, $(x, y) \in \mathcal{X} \times \mathcal{Y}$, satisfying

$\sum_{x,y} p(x,y) = 1$. Consider a function $F: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$.

Definition 1. The exponential family \mathcal{E}_F with sufficient statistics F consists of all probability distributions on $\mathcal{X} \times \mathcal{Y}$ of the form

$$p_\theta(x,y) = \frac{1}{Z(\theta)} \exp(\langle \theta, F(x,y) \rangle), \quad (x,y) \in \mathcal{X} \times \mathcal{Y}, \quad \theta \in \mathbb{R}^d,$$

where $\langle \cdot, \cdot \rangle$ denotes the standard inner product and $Z: \theta \mapsto \sum_{x',y'} \exp(\langle \theta, F(x',y') \rangle)$ is a normalization function. Note that each distribution in the exponential family has strictly positive entries. We will regard F as a matrix with columns $F(x,y) \in \mathbb{R}^d$, $(x,y) \in \mathcal{X} \times \mathcal{Y}$. We note that the exponential family \mathcal{E}_F is fully characterized by the row space $\langle \mathbb{R}^d, F \rangle \subseteq \mathbb{R}^{\mathcal{X} \times \mathcal{Y}}$ of the sufficient statistics matrix F . More precisely, two matrices F and G produce the same exponential family if and only if $(F; \mathbb{1})$ and $(G; \mathbb{1})$ have the same row span, where $\mathbb{1}$ is a row of ones. From now on we will always assume, without loss of generality, that F includes $\mathbb{1}$ in its row span. The dimension of the exponential family is then $\dim(\mathcal{E}_F) = \text{rank}(F) - 1$. This is the same as the dimension of the convex support polytope $\text{conv}\{F(x,y): (x,y) \in \mathcal{X} \times \mathcal{Y}\}$.

Definition 2. The marginal model \mathcal{M}_F on \mathcal{X} of the exponential family \mathcal{E}_F is the set of all probability distributions of the form

$$p_\theta(x) = \sum_{y \in \mathcal{Y}} \frac{1}{Z(\theta)} \exp(\langle \theta, F(x,y) \rangle), \quad x \in \mathcal{X}, \quad \theta \in \mathbb{R}^d.$$

The marginal model \mathcal{M}_F is the image of \mathcal{E}_F by the marginalization map, which is the linear map represented by the matrix with rows equal to the indicators $\mathbb{1}_x \in \mathbb{R}^{\mathcal{X} \times \mathcal{Y}}$ of $\{x\} \times \mathcal{Y}$, for each $x \in \mathcal{X}$. We are interested in the dimension of \mathcal{M}_F . When $\dim(\mathcal{M}_F) = \min\{\dim(\mathcal{E}_F), |\mathcal{X}| - 1\}$, we say that \mathcal{M}_F has the expected dimension, meaning that marginalization does not collapse the dimension of \mathcal{E}_F , or that the marginal \mathcal{M}_F is full dimensional, having the same dimension as the simplex $\Delta_{|\mathcal{X}|-1}$ of all probability distributions on \mathcal{X} .

The dimension of \mathcal{M}_F is equal to the maximum rank of the Jacobian matrix of the parametrization $\theta \mapsto (p_\theta(x))_x$. The Jacobian is given by

$$J_{\mathcal{M}_F}(\theta) = \left(\sum_y p_\theta(x,y) F(x,y) - \sum_y p_\theta(x,y) \sum_{x',y'} p_\theta(x',y') F(x',y') \right)_x. \quad (1)$$

The second term corresponds to the normalization function Z . For the rank we have

$$\begin{aligned} \text{rank}(J_{\mathcal{M}_F}(\theta)) &= \text{rank} \left(\sum_y p_\theta(x,y) F(x,y) \right)_x - 1 \\ &= \text{rank} \left(\sum_y p_\theta(y|x) F(x,y) \right)_x - 1. \end{aligned} \quad (2)$$

The first equality follows from the assumption that F has a constant row. The second one is because $p_\theta(x) = \sum_{y'} p_\theta(x, y') > 0$ for all x . Here $p_\theta(y|x) := p_\theta(x, y) / \sum_{y'} p_\theta(x, y')$ denotes the conditional probability of y given x . A geometric interpretation is that $\text{rank}(J_{\mathcal{M}_F}(\theta))$ is the dimension of the polytope defined as the convex hull of

$$\sum_y p_\theta(y|x) F(x, y), \quad x \in \mathcal{X}.$$

Evaluating Equation 2 is difficult, in general. The problem is easier in the limit of large parameters, where the sum over y almost always reduces to a single term. To see that this is the case, note that multiplicative factors of the parameter θ correspond to exponential factors of the probability distribution, such that $p_{\alpha\theta}(\cdot|x) \propto p_\theta(\cdot|x)^\alpha$. Therefore, for any $\theta \in \mathbb{R}^d$, the limit $\lim_{\alpha \rightarrow \infty} p_{\alpha\theta}(y|x)$ is non-zero only for $y \in \text{argmax}_y p_\theta(y|x) = \text{argmax}_y \langle \theta, F(x, y) \rangle$. Following this line of thought, it is convenient to define the function that outputs the most likely value of y to any given x :

Definition 3. The inference function of \mathcal{M}_F with parameter $\theta \in \mathbb{R}^d$ is given by

$$h_\theta: \mathcal{X} \rightarrow 2^{\mathcal{Y}}; \quad x \mapsto h_\theta(x) = \text{argmax}_y \langle \theta, F(x, y) \rangle.$$

Here $2^{\mathcal{Y}}$ denotes the power set of \mathcal{Y} . Geometrically, $h_\theta(x)$ is the set of $y \in \mathcal{Y}$ for which $F(x, y)$ lies in the supporting hyperplane of $F(x, y)$, $y \in \mathcal{Y}$, with normal θ . The situation is illustrated in Figure 1.

We have the following dimension bounds:

Proposition 4. *The dimension of the marginal model \mathcal{M}_F satisfies*

$$\text{rank}(F) - 1 = \dim(\mathcal{E}_F) \geq \dim(\mathcal{M}_F) \geq \max_\theta \text{rank}(\bar{F}(x, h_\theta(x)))_x - 1,$$

where $\bar{F}(x, h_\theta(x)) := \frac{1}{|h_\theta(x)|} \sum_{y \in h_\theta(x)} F(x, y)$.

Proof. We have

$$\begin{aligned} \max_{\theta \in \mathbb{R}^d} \text{rank}(J_{\mathcal{M}_F}(\theta)) &= \max_{\theta \in \mathbb{R}^d} \text{rank} \left(\sum_y p_\theta(y|x) F(x, y) \right)_x - 1 \\ &\geq \max_{\theta \in \mathbb{R}^d} \text{rank} \left(\lim_{\alpha \rightarrow \infty} \sum_y p_{\alpha\theta}(y|x) F(x, y) \right)_x - 1 \\ &= \max_{\theta \in \mathbb{R}^d} \text{rank} \left(\sum_{y \in h_\theta(x)} \frac{1}{|h_\theta(x)|} F(x, y) \right)_x - 1. \end{aligned}$$

The first line is Equation 2. The second line follows from the continuity of the parametrization of the exponential family (see Definition 2) and the lower semicontinuity of the rank function. The third line is because $\lim_{\alpha \rightarrow \infty} p_{\alpha\theta}(y|x)$ is positive and constant on $h_\theta(x) = \text{argmax}_y p_\theta(y|x)$ and zero on $\mathcal{Y} \setminus h_\theta(x)$. \square




Figure 1: Illustration of an inference function for a model with $\mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{0, \dots, 5\}$. Each dot is a vector $F(x, y)$. Small dots correspond to $x = 0$ and large ones to $x = 1$. The value of y is indicated next to each dot. For the illustrated choice of θ the inference function is $h_\theta(x = 0) = \{2\}$, $h_\theta(x = 1) = \{1, 2, 3\}$ and the tropical morphism is given by $\mathcal{A}_\theta = (F(0, 2), \frac{1}{3}(F(1, 1) + F(1, 2) + F(1, 3)))$. For generic choices of θ the vector $F(x, h_\theta(x))$ is a vertex of the polytope $\text{conv}\{F(x, y) : y \in \mathcal{Y}\}$, for all x .

Proposition 4 shows that we can estimate the dimension of the marginal model by studying the maximum rank over θ of the piecewise constant matrix-valued function

$$\mathcal{A}_\theta := (\bar{F}(x, h_\theta(x)))_x.$$

For each $x \in \mathcal{X}$, the column $\bar{F}(x, h_\theta(x))$ is the average of the maximizers of the linear form $\langle \theta, \cdot \rangle$ over $F(x, y)$, $y \in \mathcal{Y}$. For generic choices of $\theta \in \mathbb{R}^d$, the set of maximizers $F(x, y)$, $y \in h_\theta(x)$, consists of one single element, for each x . In particular, when all vectors $F(x, y)$ are different, for generic choices of θ , the inference function h_θ maps each x to a single y .

The vectors $\bar{F}(x, h_\theta(x))$, $x \in \mathcal{X}$, do not necessarily lie on the same supporting hyperplane of $F(x, y)$, $(x, y) \in \mathcal{X} \times \mathcal{Y}$, although the converse is true in the following sense. If there is supporting hyperplane that intersects $F(x, y)$, $(x, y) \in \mathcal{X} \times \mathcal{Y}$, exactly at $F(x, f(x))$, $x \in \mathcal{X}$, for some $f : \mathcal{X} \rightarrow \mathcal{Y}$, then f is an inference function. However, since these vectors lie on a supporting hyperplane (which usually defines a proper face of the convex support polytope), they are not suited for estimating the maximum rank.

The matrix \mathcal{A}_θ defines a geometric object called the tropical version of the original marginal model:

Definition 5. The tropical version of \mathcal{M}_F , denoted $\mathcal{M}_F^{\text{tropical}}$, is the set of all vectors in

$\mathbb{R}^{\mathcal{X}}/\mathbb{R}\mathbf{1}$ (modulo addition of constant vectors) of the form

$$\Phi_{\theta}(x) = \langle \theta, \bar{F}(x, h_{\theta}(x)) \rangle, \quad x \in \mathcal{X}, \quad \text{parametrized by } \theta \in \mathbb{R}^d.$$

Proposition 4 can be regarded as a version of the Bieri-Groves theorem [3, 8], stating that the dimension of the marginal model is bounded below by the dimension of its tropical version:

$$\dim(\mathcal{M}_F) \geq \dim(\mathcal{M}_F^{\text{tropical}}).$$

In particular, when $\mathcal{M}_F^{\text{tropical}}$ has the expected dimension, then \mathcal{M}_F also has the expected dimension. This is the central idea of [8] and subsequently [6, 19] for estimating the dimension of secant varieties and restricted Boltzmann machines.

We note that the marginal model and its tropical version are independent of the sufficient statistics matrix used to parametrize the underlying exponential family:

Proposition 6. *If $E = Q^{\top}F$ is a non-singular linear transformation of F , then $\mathcal{E}_E = \mathcal{E}_F$, $\mathcal{M}_E = \mathcal{M}_F$, and $\mathcal{M}_E^{\text{tropical}} = \mathcal{M}_F^{\text{tropical}}$. More generally, if $G = R^{\top}F$ is a linear transformation of F , then $\mathcal{E}_G \subseteq \mathcal{E}_F$, $\mathcal{M}_G \subseteq \mathcal{M}_F$, and $\mathcal{M}_G^{\text{tropical}} \subseteq \mathcal{M}_F^{\text{tropical}}$.*

Proof. The equality of the exponential families follows from the equality of the row spaces of F and E . For the tropical models note that for each θ there is a $\vartheta = Q^{-1}\theta$ with $\langle \theta, F(x, h_{\theta}(x)) \rangle = \langle \vartheta, E(x, h_{\vartheta}(x)) \rangle$. For the inclusions note that the row space of $G = R^{\top}F$ is a linear subspace of the row space of F . \square

3 The Kronecker Product Model

Kronecker product models are marginals of exponential families whose sufficient statistics matrix factorizes as the Kronecker product of a sufficient statistics matrix over the visible states and one over the hidden states. Recall the definition of the Kronecker product $(A_{i,j})_{i,j} \otimes (B_{k,l})_{k,l} := (A_{i,j}B_{k,l})_{(i,k),(j,l)}$.

Definition 7. A *Kronecker product model* is a marginal model \mathcal{M}_F , where F factorizes as $F(x, y) = A(x) \otimes B(y)$, $x \in \mathcal{X}$, $y \in \mathcal{Y}$, for some $A \in \mathbb{R}^{a \times \mathcal{X}}$ and $B \in \mathbb{R}^{b \times \mathcal{Y}}$.

We will use the notations $A = (A_x)_x \in \mathbb{R}^{a \times \mathcal{X}}$, $B = (B_y)_y \in \mathbb{R}^{b \times \mathcal{Y}}$, and assume that the row space of each matrix contains a constant non-zero vector. For simplicity in the following we assume that all columns of B are different. The general case with repeated columns is very similar, but needs more complicated notations.

Consider a generic choice of the parameter θ , such that the inference function h_{θ} maps each $x \in \mathcal{X}$ to a single $y \in \mathcal{Y}$. The tropical morphism $\Phi_{\theta}(x) = \langle \theta, \mathcal{A}_{\theta} \rangle$ of a Kronecker product model has the form

$$\mathcal{A}_{\theta} = (F(x, h_{\theta}(x)))_x = (A_x \otimes B_{h_{\theta}(x)})_x = A \odot B_{h_{\theta}}, \quad (3)$$

where $(A_{i,j})_{i,j} \odot (B_{k,l})_{k,l} := (A_{i,j}B_{k,l})_{(i,k),j}$ denotes the column-wise Kronecker product or Khatri-Rao product [15]. Alternatively, after rearranging columns, we can write this as

$$\mathcal{A}_{\theta} = (B_y \otimes A_{h_{\theta}^{-1}(y)})_y = B \odot A_{h_{\theta}^{-1}},$$

where now \odot denotes the column-block-wise Kronecker product. Here the column-blocks, indexed by y , are B_y and $A_{h_\theta^{-1}(y)}$. When $h_\theta^{-1}(y) = \emptyset$, we simply omit the block $B_y \otimes A_{h_\theta^{-1}(y)}$.

Unfortunately there is no formula for expressing the rank of a Khatri-Rao product in terms of the ranks of its factors. A simple lower bound for matrices consisting of non-zero columns is $\text{rank}(A \odot B_h) \geq \max\{\text{rank}(A), \text{rank}(B_h)\}$. More elaborate lower bounds can be given in terms of Kruskal ranks [24], also for column-block partitioned matrices [17]. Our analysis seeks to characterize pairs of matrices A and B for which the upper bound $\text{rank}(A \odot B_h) \leq \text{rank}(A \otimes B_h) = \text{rank}(A) \cdot \text{rank}(B_h)$ is attained for some inference function h . For this it is critical to study the possible inference functions.

The factorization property $F = A \otimes B$ leads to highly structured inference functions. We explain this in the following. For a given parameter vector $\theta = (\theta_{(i,j)})_{(i,j)} \in \mathbb{R}^{ab}$ let $\Theta = (\theta_{(i,j)})_{j,i} \in \mathbb{R}^{b \times a}$ denote the matrix with column-by-column vectorization equal to θ . By Roth's lemma [23] we have the following equalities:

$$\langle \theta, (A \otimes B)_{(x,y)} \rangle = \langle \Theta A_x, B_y \rangle = \langle \Theta^\top B_y, A_x \rangle \quad \text{for all } x \in \mathcal{X}, y \in \mathcal{Y}.$$

In turn, the same inner product describes the following distributions over (x, y) , y , and x :

$$\begin{aligned} p_\theta(\cdot, \cdot) &= \frac{1}{Z(\theta)} \exp(\langle \theta, A \otimes B \rangle) \in \mathcal{E}_{A \otimes B}, \\ p_\theta(\cdot | x) &= \frac{1}{Z(\Theta A_x)} \exp(\langle \Theta A_x, B \rangle) \in \mathcal{E}_B, \\ p_\theta(\cdot | y) &= \frac{1}{Z(\Theta^\top B_y)} \exp(\langle \Theta^\top B_y, A \rangle) \in \mathcal{E}_A. \end{aligned} \tag{4}$$

In particular, $p_\theta(\cdot | x)$ appears in the Equation 2 of the Jacobian rank. Geometrically, ΘA is the linear projection of the columns of A by the matrix Θ to the parameter space of the hidden exponential family \mathcal{E}_B . Similarly, $\Theta^\top B$ is the projection of the columns of B by the matrix Θ^\top to the parameter space of the visible exponential family \mathcal{E}_A .

The inference function of a Kronecker model satisfies

$$\begin{aligned} h_\theta(x) &= \operatorname{argmax}_y p_\theta(y | x) = \operatorname{argmax}_y \langle \theta, A_x \otimes B_y \rangle \\ &= \operatorname{argmax}_y \langle \Theta A_x, B_y \rangle = \{y \in \mathcal{Y} : \Theta A_x \in N_B(y)\} \\ &= \operatorname{argmax}_y \langle A_x, \Theta^\top B_y \rangle = \{y \in \mathcal{Y} : A_x \in N_{\Theta^\top B}(y)\}. \end{aligned}$$

Here the normal cones of B at B_y and of $\Theta^\top B$ at $\Theta^\top B_y$ are defined, respectively, as

$$\begin{aligned} N_B(y) &:= \{r \in \mathbb{R}^b : \langle r, B_y - B_{y'} \rangle \geq 0 \text{ for all } y' \in \mathcal{Y} \setminus \{y\}\}, \\ N_{\Theta^\top B}(y) &:= \{r \in \mathbb{R}^a : \langle r, \Theta^\top B_y - \Theta^\top B_{y'} \rangle \geq 0 \text{ for all } y' \in \mathcal{Y} \setminus \{y\}\}. \end{aligned}$$

In turn, the inference function h_θ can be interpreted as a slicing of ΘA by the normal fan of B , or, equivalently, a slicing of A by the normal fan of $\Theta^\top B$:

Definition 8. A B -slicing of A is a partition of the column set of A into the blocks $A_{C_y} := (A_x)_{x \in C_y}$ with $C_y := h_\theta^{-1}(y) = \{x \in \mathcal{X} : h_\theta(x) = y\}$, for all $y \in \mathcal{Y}$. Here we assume that h_θ maps each x to a single y , which is the generic case when all columns of B are different. Note that some of the sets $C_y = h_\theta^{-1}(y)$ may be empty.

Related to the statements of Proposition 6 we have the following observations: If $B' = DB$ for some matrix D , then any B' -slicing is also a B -slicing. Furthermore, if B' is a matrix consisting of all columns of B that lie in a common supporting hyperplane of B , then any B' -slicing is a B -slicing. Note also that, by the mixed-product property of the Kronecker product, $(C \otimes D)(A \otimes B) = (CA) \otimes (DB)$, the Kronecker product of two linearly transformed matrices CA and DB is given by a linear transformation of the Kronecker product of the two original matrices.

4 Hierarchical Models

We are interested in Kronecker product models for which the factor exponential families \mathcal{E}_A and \mathcal{E}_B are hierarchical models. In this section we provide the necessary definitions.

Consider n random variables with finite state sets $\mathcal{X}_i = \{0, 1, \dots, |\mathcal{X}_i| - 1\}$, $i \in [n] := \{1, \dots, n\}$. We write $x = (x_1, \dots, x_n)$ for an element of $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$. Given some $x \in \mathcal{X}$ and $\lambda \subseteq [n]$, we write $x_\lambda = (x_i)_{i \in \lambda}$ for the natural projection of x to the λ -coordinates. Unless otherwise stated, in all that follows Λ will denote an inclusion closed set of subsets of $[n]$, such that $\lambda \in \Lambda$ and $\lambda' \subseteq \lambda$ imply $\lambda' \in \Lambda$. Consider the linear subspace of $\mathbb{R}^{\mathcal{X}}$ consisting of all linear combinations of real-valued functions that depend only on x_λ , $\lambda \in \Lambda$,

$$V_\Lambda(\mathcal{X}) := \left\{ \sum_{\lambda \in \Lambda} f_\lambda : f_\lambda(x) = f_\lambda(x_\lambda) \right\} \subseteq \mathbb{R}^{\mathcal{X}}.$$

Definition 9. The hierarchical model on \mathcal{X} with interactions Λ is the exponential family \mathcal{E}_A , where $A \in \mathbb{R}^{a \times \mathcal{X}}$ is a matrix with row span $V_\Lambda(\mathcal{X})$. We will denote this model by \mathcal{E}_Λ . An important special case is the k -interaction model \mathcal{E}_{Λ_k} , defined by some $1 \leq k \leq n$ and $\Lambda_k = \{\lambda \subseteq [n] : |\lambda| \leq k\}$.

An important example is the independence model \mathcal{E}_{Λ_1} , which is the k -interaction model with $k = 1$. The independence model consists of probability distributions that factorize as $p(x_1, \dots, x_n) = p_1(x_1) \otimes \dots \otimes p_n(x_n)$, where, for each $i \in [n]$, p_i is a probability distribution over \mathcal{X}_i . This model corresponds to the Segre embedding of $\mathbb{P}^{|\mathcal{X}_1|-1} \times \dots \times \mathbb{P}^{|\mathcal{X}_n|-1}$ into $\mathbb{P}^{\prod_{i \in [n]} |\mathcal{X}_i|-1}$.

The sufficient statistics matrix of a hierarchical model can be constructed in the following simple way. The first row, with index \emptyset , is the constant vector of ones $A_{\emptyset, x} := 1$, $x \in \mathcal{X}$. The other rows are indexed by pairs $(\lambda, \tilde{x}_\lambda)$, with $\lambda \in \Lambda$ and $\tilde{x}_\lambda \in \tilde{\mathcal{X}}_\lambda := \times_{i \in \lambda} (\mathcal{X}_i \setminus \{0\})$. The $((\lambda, \tilde{x}_\lambda), x)$ -th entry of the matrix is

$$A_{(\lambda, \tilde{x}_\lambda), x} := \begin{cases} 1, & \text{if } x_\lambda = \tilde{x}_\lambda \\ 0, & \text{otherwise} \end{cases}. \quad (5)$$

Proposition 10. *The matrix A from Equation 5 has row space $V_\Lambda(\mathcal{X})$. Furthermore, $\dim(\mathcal{E}_\Lambda) = \text{rank}(A) - 1 = \dim(V_\Lambda(\mathcal{X})) - 1 = \sum_{\lambda \in \Lambda \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1)$.*

The statement of Proposition 10 is well known in the context of hierarchical models. One way of proving it is as follows:

Proof. Let $\lambda \in \Lambda$ and let $f_\lambda \in \mathbb{R}^{\mathcal{X}}$ be a function with $f_\lambda(x) = f_\lambda(x_\lambda)$ for all $x \in \mathcal{X}$. Note that f_λ can be written as a linear combination of indicator functions of cylinder sets as $f_\lambda(x) = \sum_{x_\lambda^* \in \mathcal{X}_\lambda} f_\lambda(x_\lambda^*) \mathbb{1}_{\{x' \in \mathcal{X}: x'_\lambda = x_\lambda^*\}}(x)$ for all $x \in \mathcal{X}$. Hence we only need to show that the row span of A contains the indicator functions $\mathbb{1}_{\{x' \in \mathcal{X}: x'_\lambda = x_\lambda^*\}}$ for all $x_\lambda^* \in \mathcal{X}_\lambda$. For any x_λ^* and $\lambda^* = \text{supp}(x^*) \cap \lambda$ we have that

$$\mathbb{1}_{\{x' \in \mathcal{X}: x'_\lambda = x_\lambda^*\}}(x) = \sum_{\lambda^* \subseteq \lambda' \subseteq \lambda} (-1)^{|\lambda' \setminus \lambda^*|} \sum_{\substack{\tilde{x}_{\lambda'} \in \tilde{\mathcal{X}}_{\lambda'} \\ \tilde{x}_{\lambda^*} = x_{\lambda^*}^*}} A_{(\lambda', \tilde{x}_{\lambda'}), x}, \quad \text{for all } x \in \mathcal{X}.$$

Hence the row span of A contains the indicator function $\mathbb{1}_{\{x' \in \mathcal{X}: x'_\lambda = x_\lambda^*\}}$. Since x_λ^* was arbitrary, this shows that the row span contains f_λ . Since λ was arbitrary in Λ , this shows that the row span contains $V_\Lambda(\mathcal{X})$. The reverse inclusion is direct. The matrix A has $1 + \sum_{\lambda \in \Lambda \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1)$ linearly independent rows, including a row of ones. This implies the dimension statement. \square

We now introduce the concept of Λ -balls, which we will use for constructing slicings of hierarchical models and formulating our theorems later on.

Definition 11. Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$. The Λ -ball in \mathcal{X} centered at $x \in \mathcal{X}$ is the set of vectors that differ from x exactly at the entries from some $\lambda \in \Lambda$,

$$K_{\mathcal{X}}(x, \Lambda) := \{x' \in \mathcal{X}: \{i \in [n]: x'_i \neq x_i\} = \lambda \in \Lambda\}.$$

Note that all Λ -balls in \mathcal{X} have the same cardinality, regardless of their center,

$$|K_{\mathcal{X}}(\Lambda)| := |K_{\mathcal{X}}(x, \Lambda)| = 1 + \sum_{\lambda \in \Lambda \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1).$$

We will drop the subscript \mathcal{X} when it is clear from the context.

An important special case of Λ -balls are Hamming balls. Recall that the Hamming distance between two vectors $x, x' \in \mathcal{X}$ is defined as $d_H(x, x') := |\{i \in [n]: x_i \neq x'_i\}|$. The radius- k Hamming ball in \mathcal{X} centered at x is the set of vectors that differ from x at most in k entries, $K(x, \Lambda_k) = \{x' \in \mathcal{X}: d_H(x, x') \leq k\}$.

In later sections we will consider slicings of a matrix A with row span $V_\Lambda(\mathcal{X})$ into blocks corresponding to Λ -balls. We will use the following lemma:

Lemma 12. *Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, let A be a matrix with row span $V_\Lambda(\mathcal{X})$, let $x \in \mathcal{X}$, and let $K(x, \Lambda)$ be the Λ -ball in \mathcal{X} centered at x . Then $A_{K(x, \Lambda)}$ has full rank,*

$$\text{rank}(A_{K(x, \Lambda)}) = \dim(V_\Lambda(\mathcal{X})) = 1 + \sum_{\lambda \in \Lambda \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1) = K(\Lambda).$$

Proof. Consider the matrix entries $A_{(\lambda, \tilde{x}_\lambda), x}$ defined in Equation 5. For any $\lambda \notin \Lambda$ and $\tilde{x}_\lambda \in \tilde{\mathcal{X}}_\lambda = \times_{i \in \lambda} \mathcal{X}_i \setminus \{0\}$, we have that $A_{(\lambda, \tilde{x}_\lambda), x} = 0$ for all $x \in K(0, \Lambda)$. On the other hand, by Proposition 14 the full matrix $(A_{(\lambda, \tilde{x}_\lambda), x})_{(\lambda \in 2^{[n]}, \tilde{x}_\lambda \in \tilde{\mathcal{X}}_\lambda), x \in \mathcal{X}}$ has rank equal to $\dim(V_{2^{[n]}}(\mathcal{X})) = |\mathcal{X}|$ and thus it has row span $\mathbb{R}^{\mathcal{X}}$. This implies that the matrix $(A_{(\lambda, \tilde{x}_\lambda), x})_{(\lambda \in \Lambda, \tilde{x}_\lambda \in \tilde{\mathcal{X}}_\lambda), x \in K(0, \Lambda)}$ has row span $\mathbb{R}^{K(0, \Lambda)}$. This proves the claim for the case of a Λ -ball centered at the zero vector. The other cases follow from this after relabeling the states. \square

Lemma 12 states that certain collections of columns of the sufficient statistics matrix of a hierarchical model are linearly independent. This property is related to the notion of Kruskal rank, defined as the largest r for which any r columns of a matrix are linearly independent, which has been used before to study the rank of Khatri-Rao products [17].

Kronecker products of hierarchical models correspond to hierarchical models with interaction sets given by the Cartesian product of the interaction sets of the factors. More precisely, if A has row span $V_\Lambda(\mathcal{X}_1 \times \dots \times \mathcal{X}_n)$ and B has row span $V_{\Lambda'}(\mathcal{Y}_1 \times \dots \times \mathcal{Y}_m)$, then $A \otimes B$ has row span $V_{\Lambda \times \Lambda'}(\mathcal{X} \times \mathcal{Y})$, where $\Lambda \times \Lambda' = \{\lambda \times \lambda' \subseteq [n] \times [m] : \lambda \in \Lambda, \lambda' \in \Lambda'\}$. Figure 2 shows various types of examples. In the top row of the figure, the visible factor \mathcal{E}_A is an independence model. In the bottom row \mathcal{E}_A is an interaction model. The first column shows examples where the hidden factor \mathcal{E}_B is the set of all strictly positive distributions on \mathcal{Y} . These correspond to mixture models of \mathcal{E}_A . We will cover them in Section 5. The second column shows examples where \mathcal{E}_B is an independence model. These correspond to Hadamard products of mixture models of \mathcal{E}_A . We will cover them in Section 6. The third column shows examples where \mathcal{E}_B is an interaction model. We will cover them in Section 7.

5 Mixture Models

Definition 13. The k -mixture of \mathcal{E}_A consists of all possible convex combinations of k probability distributions from \mathcal{E}_A ; that is, the probability distributions

$$p(x) = \sum_{i \in [k]} \alpha(i) p^{(i)}(x), \quad x \in \mathcal{X},$$

where

$$p^{(1)}, \dots, p^{(k)} \in \mathcal{E}_A, \quad \alpha(1), \dots, \alpha(k) \geq 0, \quad \sum_{i \in [k]} \alpha(i) = 1.$$

In algebraic geometry one usually considers secants instead of mixtures, in which case the weights $\alpha(i)$ in Definition 13 add to one but are not required to be non-negative. The resulting set is the union of possible affine hulls of k points in \mathcal{E}_A . The Zariski closure of the k -mixture of \mathcal{E}_A is k -th secant variety of \mathcal{E}_A . A standard reference on secant varieties is [28].

Mixtures of exponential families can be expressed as Kronecker product models:




Figure 2: Examples of Kronecker product hierarchical models. The dark and light circles represent hidden and visible variables, respectively. The dark and light squares represent interactions among the adjacent hidden and visible variables, respectively. The edges between squares represent interactions between the hidden and visible variables adjacent to those squares. The graph between squares is full bipartite. (a) mixture model of an independence model, (b) mixture model of a pairwise interaction model, (c) Hadamard product of three mixture models of an independence model (restricted Boltzmann machine with three hidden units), (d) Hadamard product of three mixture models of a pairwise interaction model, (e) pairwise interaction mixture model of an independence model, (f) pairwise interaction mixture model of a pairwise interaction model. From top to bottom and from left to right the models are more general. From left to right these examples are covered in Theorems 16, 22 and 24.

Proposition 14. *Let B have row span $\mathbb{R}^{\mathcal{Y}}$ and let A include a row of ones. Then the Kronecker product model $\mathcal{M}_{A \otimes B}$ is the $|\mathcal{Y}|$ -mixture of \mathcal{E}_A .*

Proof. Without loss of generality let B be the $|\mathcal{Y}| \times |\mathcal{Y}|$ identity matrix I . The distributions of the exponential family $\mathcal{E}_{A \otimes B}$ have the form $p_{\theta}(x, y) = \frac{1}{Z(\theta)} \exp(\langle \Theta^{\top} B_y, A_x \rangle) = \frac{1}{Z(\theta)} \exp(\langle \Theta_y^{\top}, A_x \rangle)$. We may assume that the first row of A is a vector of ones. Hence, adding a suitable $\tilde{\theta}$ to the parameter vector, we obtain $p_{\theta + \tilde{\theta}}(x, y) = \frac{1}{Z(\theta + \tilde{\theta})} \exp(\tilde{\Theta}_{1,y}^{\top}) \exp(\langle \Theta_y^{\top}, A_x \rangle)$. The first term can be adjusted to obtain the mixture weights α from Equation 13 and the second term can be chosen independently for each value of y . \square

If B has row span $\mathbb{R}^{\mathcal{Y}}$, we may assume that B is the $|\mathcal{Y}| \times |\mathcal{Y}|$ identity matrix I . By Equation 3, the tropical morphism has the following form:

$$\mathcal{A}_{\theta} = I \odot A_{h_{\theta}^{-1}} = \begin{pmatrix} \boxed{A_{C_1}} & & & & \\ & \boxed{A_{C_2}} & & & \\ & & \ddots & & \\ & & & & \boxed{A_{C_{|\mathcal{Y}|}}} \end{pmatrix}.$$

In particular, the rank is just the sum of the ranks of the individual diagonal blocks,

$$\text{rank}(\mathcal{A}_\theta) = \sum_{y \in \mathcal{Y}} \text{rank}(A_{C_y}).$$

In order to estimate the maximum of $\text{rank}(\mathcal{A}_\theta)$ over θ , an obvious strategy is to search for a slicing that produces as many full-rank blocks A_{C_y} as possible. When B is an identity matrix, the slicings can be constructed using any matrix $\Theta^\top B = \Theta^\top \in \mathbb{R}^{a \times b}$. By Lemma 12, if A is the sufficient statistics matrix of \mathcal{E}_Λ and C_y contains, or is contained in, a Λ -ball, then the block A_{C_y} has full rank. As we will show below, inference regions C_y containing interaction balls can be obtained as slicings of A by matrices of the form $\Theta^\top B = (A_{c_y})_y$, where c_y are the centers of disjoint Λ -balls. We focus on k -interaction models.

Lemma 15. *Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, let A have row span $V_k(\mathcal{X})$, and let B have row span $\mathbb{R}^{\mathcal{Y}}$. Given $|\mathcal{Y}|$ disjoint Λ_k -balls in \mathcal{X} , denoted $K(c_y, \Lambda_k)$, $y \in \mathcal{Y}$, there is a B -slicing of A with $C_y \supseteq K(c_y, \Lambda_k)$ for all $y \in \mathcal{Y}$.*

Proof. Without loss of generality we choose a matrix A with entries $A_{(\lambda, \tilde{x}_\lambda), x}$ equal to 1 if $x_\lambda = \tilde{x}_\lambda$ and -1 otherwise, for $\lambda \in \Lambda_k$, $\tilde{x}_\lambda \in \mathcal{X}_\lambda$, and $x \in \mathcal{X}$. Denoting the number of rows by $a = \sum_{\lambda \in \Lambda_k} \prod_{i \in \lambda} |\mathcal{X}_i|$ we have that $\langle A_x, A_{x'} \rangle = a - 2d_H(A_x, A_{x'})$. Furthermore, $d_H(A_x, A_{x'}) = 2|\{\lambda \in \Lambda_k : x_\lambda \neq x'_\lambda\}|$. If $x' \in K(x, \Lambda_k)$, then $d_H(A_x, A_{x'}) \leq 2(2^k - 1)$. On the other hand, if $x' \notin K(x, \Lambda_k)$, then $d_H(A_x, A_{x'}) \geq 2(2^k - 1) + 1$. Hence choosing Θ such that $\Theta^\top B_y = A_{c_y}$ for all $y \in \mathcal{Y}$, yields $C_y \supseteq K(c_y, \Lambda_k)$ for all $y \in \mathcal{Y}$. \square

Theorem 16. *Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, let A have row span $V_k(\mathcal{X})$, and let B have row span $\mathbb{R}^{\mathcal{Y}}$.*

- *If \mathcal{X} contains $|\mathcal{Y}|$ disjoint radius- k Hamming balls, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{Y}| \left(1 + \sum_{\lambda \in \Lambda_k \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1) \right) - 1.$$

- *If \mathcal{X} can be covered by $|\mathcal{Y}|$ radius- k Hamming balls, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{X}| - 1.$$

Proof. If $K(c_y, \Lambda_k)$, $y \in \mathcal{Y}$, are disjoint, then by Lemma 15 we obtain a slicing with $C_y \supseteq K(c_y, \Lambda_k)$ for all $y \in \mathcal{Y}$. By Lemma 12 we have $\text{rank}(A_{C_y}) = \text{rank}(A)$ for all $y \in \mathcal{Y}$. This yields $\text{rank}(\mathcal{A}_\theta) = \sum_y \text{rank}(A_{C_y}) = |\mathcal{Y}| \text{rank}(A)$. If $K(c_y, \Lambda_k)$, $y \in \mathcal{Y}$, cover \mathcal{X} , then $C_y \subseteq K(c_y, \Lambda_k)$ and $\text{rank}(A_{C_y}) = |C_y|$ for all $y \in \mathcal{Y}$. This yields $\text{rank}(\mathcal{A}_\theta) = \sum_y \text{rank}(A_{C_y}) = \sum_y |C_y| = |\mathcal{X}|$. \square

We note the following special case where \mathcal{E}_A is an independence model. This case has been covered previously in Draisma's tropical approach to secant dimensions [8]. The corresponding implications for the dimension of (not tropical) mixtures of independence models have also been studied before in algebraic geometry and tensor analysis; see [4, 2, 16].

Corollary 17. Let $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$, let A have row span $V_1(\mathcal{X})$, and let B have row span $\mathbb{R}^{\mathcal{Y}}$.

- If \mathcal{X} contains $|\mathcal{Y}|$ disjoint radius-one Hamming balls, then

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{Y}| \left(1 + \sum_{i \in [n]} (|\mathcal{X}_i| - 1) \right) - 1.$$

- If \mathcal{X} can be covered by $|\mathcal{Y}|$ radius-one Hamming balls, then

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{X}| - 1.$$

6 Hadamard Products

Definition 18. The Hadamard product $\mathcal{M}_1 * \mathcal{M}_2$ of two probability models $\mathcal{M}_1, \mathcal{M}_2$ on \mathcal{X} is the set of all probability distributions of the form

$$(p * q)(x) = \frac{p(x)q(x)}{\sum_{x' \in \mathcal{X}} p(x')q(x')}, \quad x \in \mathcal{X}, \quad \text{where } p \in \mathcal{M}_1 \text{ and } q \in \mathcal{M}_2.$$

The Hadamard product of $\mathcal{M}_1, \dots, \mathcal{M}_m$ is defined in an analogous way.

Consider m independent hidden variables, each with an associated exponential family. In other words, let $B^j \in \mathbb{R}^{b_j \times \mathcal{Y}_j}$, $j \in [m]$, and let $B \in \mathbb{R}^{b \times \mathcal{Y}}$ be the matrix with columns $B_y = (B_{y_1}^1; \dots; B_{y_m}^m)$, $y \in \mathcal{Y}$. The corresponding exponential family factorizes as $\mathcal{E}_B = \mathcal{E}_{B^1} \otimes \dots \otimes \mathcal{E}_{B^m}$, where \mathcal{E}_{B^j} is an exponential family on \mathcal{Y}_j , for each $j \in [m]$. A Kronecker product model with independent hidden variables is the Hadamard product of the marginal models for the individual hidden variables:

Proposition 19. Let $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_m$ and let $B = (B^1; \dots; B^m)$ with $B_y^j = B_{y_j}^j$, $y \in \mathcal{Y}$. Then $\mathcal{M}_{A \otimes B} = \mathcal{M}_{A \otimes B^1} * \dots * \mathcal{M}_{A \otimes B^m}$.

Proof. Consider a parameter vector $\theta = (\theta^1; \dots; \theta^m) \in \mathbb{R}^{a \cdot b}$ with blocks corresponding to the blocks of $B = (B^1; \dots; B^m)$. We have

$$\begin{aligned} p(x) &= \sum_{y \in \mathcal{Y}} \frac{1}{Z(\theta)} \exp(\langle \theta, A_x \otimes B_y \rangle) \\ &= \sum_{y_1 \in \mathcal{Y}_1} \dots \sum_{y_m \in \mathcal{Y}_m} \frac{1}{Z(\theta)} \exp\left(\sum_{j \in [m]} \langle \theta^j, A_x \otimes B_{y_j}^j \rangle\right) \\ &= \frac{1}{Z(\theta)} \prod_{j \in [m]} \sum_{y_j \in \mathcal{Y}_j} \exp(\langle \theta^j, A_x \otimes B_{y_j}^j \rangle). \end{aligned}$$

This proves the claim. \square

The tropical morphism of a Hadamard product decomposes into individual factor parts:

Now consider the matrix $B''' = (B'_1 + cB''_2, \dots, B'_N + cB''_2, cB''_1)$. Let E_1, \dots, E_{N+1} denote the B''' -slicing of A . We have $\langle A_x, B'''_{N+1} - B'''_i \rangle = \langle A_x, cB''_1 - (B'_i + cB''_2) \rangle = c\langle A_x, B''_1 - B''_2 \rangle - \langle A_x, B'_i \rangle$ for $i \leq N$. Choosing $c > 0$ large enough, this is positive iff $x \in D_1$. This means that $E_{N+1} = D_1$.

For any $i \leq N$, we have $\langle A_x, B'''_i - B'''_j \rangle = \langle A_x, (B'_i + cB''_2) - (B'_j + cB''_2) \rangle = \langle A_x, B'_i - B'_j \rangle$ for all $j \leq N$, and $\langle A_x, B'''_i - B'''_{N+1} \rangle = \langle A_x, (B'_i + cB''_2) - cB''_1 \rangle = \langle A_x, B'_i \rangle + c\langle A_x, B''_2 - B''_1 \rangle$. Choosing $c > 0$ large enough, all of these expressions are positive iff $x \in C_i \cap D_2$. This means that $E_i = D_2 \cap C_i$ for all $i \leq N$. \square

Theorem 22. *Let $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$ and let A have row span $V_k(\mathcal{X})$ for some $1 \leq k \leq n$. Let $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_m$ and let B have row span $V_1(\mathcal{Y})$.*

- *If \mathcal{X} contains m disjoint Hamming balls K^1, \dots, K^m , whereby K^j contains $|\mathcal{Y}_j| - 1$ disjoint radius- k Hamming balls for $j = 1, \dots, m$, and $\text{rank}(A_{\mathcal{X} \setminus (\cup_j K^j)}) = \text{rank}(A)$, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = \left(1 + \sum_{\lambda \in \Lambda_k \setminus \emptyset} \prod_{i \in \lambda} (|\mathcal{X}_i| - 1)\right) \left(1 + \sum_{j \in [m]} (|\mathcal{Y}_j| - 1)\right) - 1.$$

- *If \mathcal{X} can be covered by Hamming balls K^1, \dots, K^m , such that K^j can be covered by $|\mathcal{Y}_j| - 1$ radius- k Hamming balls, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{X}| - 1.$$

Proof. Have Equation 6 in mind. Consider some $j \in [m]$. We use Lemma 21 with $N = |\mathcal{Y}_j| - 1$. Let $D_2 = K^j$ and let C_1, \dots, C_N be the slicing of A by the matrix $(A_{c_i^j})_{i \in [N]}$, where c_i^j are the centers of $|\mathcal{Y}_j| - 1$ disjoint radius- k Hamming balls in K^j . This shows that there is a B^j -slicing of A with blocks $C_1^j, \dots, C_{|\mathcal{Y}_j|-1}^j, C_{|\mathcal{Y}_j|}^j$, where C_i^j is contained in K^j and contains a radius- k Hamming ball, for all $1 \leq i \leq |\mathcal{Y}_j| - 1$, and $C_{|\mathcal{Y}_j|}^j = \mathcal{X} \setminus K^j$.

Since all K^j are disjoint, we have that all C_i^j , $1 \leq i \leq |\mathcal{Y}_j| - 1$, $j \in [m]$, are disjoint. Since each of these sets contains a radius- k Hamming ball, Lemma 12 implies $\text{rank}(A_{C_i^j}) = \text{rank}(A)$, for all $1 \leq i \leq |\mathcal{Y}_j| - 1$, $j \in [m]$. The remaining columns of the tropical morphism have rank at least $\text{rank}(A)$, since it is assumed that $\text{rank}(A_{\mathcal{X} \setminus (\cup_j K^j)}) = \text{rank}(A)$. Hence we have $\text{rank}(\mathcal{A}_\theta) = \sum_{j \in [m]} (|\mathcal{Y}_j| - 1) \text{rank}(A) + \text{rank}(A)$. This proves the first item. The second item follows from similar arguments as the second item of Theorem 16. \square

We note the following special case, where both \mathcal{E}_B and \mathcal{E}_A are independence models, which is known as a restricted Boltzmann machine.

Corollary 23. *Let $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$ and let A have row span $V_1(\mathcal{X})$. Let $\mathcal{Y} = \mathcal{Y}_1 \times \dots \times \mathcal{Y}_m$ and let B have row span $V_1(\mathcal{Y})$.*

- If \mathcal{X} contains m disjoint Hamming balls K^1, \dots, K^m , whereby K^j contains $|\mathcal{Y}_j| - 1$ disjoint radius-one Hamming balls for $j = 1, \dots, m$, and $\text{rank}(A_{\mathcal{X} \setminus (\cup_j K^j)}) = \text{rank}(A)$, then

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = \left(1 + \sum_{i \in [n]} (|\mathcal{X}_i| - 1)\right) \left(1 + \sum_{j \in [m]} (|\mathcal{Y}_j| - 1)\right) - 1.$$

- If \mathcal{X} can be covered by Hamming balls K^1, \dots, K^m , such that K^j can be covered by $|\mathcal{Y}_j| - 1$ radius- k Hamming balls, then

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{X}| - 1.$$

A weaker version of Corollary 23 was obtained previously in [19]. That result was based on slicings by parallel hyperplanes, which are less efficient than our construction with truncated slicings. One should note, however, that in order to realize slicings by parallel hyperplanes it is not required that each B^j has row space $\mathbb{R}^{\mathcal{Y}_j}$, but only that B^j can be projected into an arbitrary set of collinear points. The special case of Corollary 23 where all variables are binary, $\mathcal{X} = \{0, 1\}^n$, $\mathcal{Y} = \{0, 1\}^m$, was obtained previously in [6]. That case is not improved by the present analysis, since for binary variables the truncated slicings are just slicings by hyperplanes.

7 Interacting Hidden Variables

Consider a matrix B of rank b in reduced row echelon form. In this case the tropical morphism has the form

$$\mathcal{A}_\theta = \begin{pmatrix} \boxed{A_{C_{p_1}}} & \cdots & \boxed{\phantom{A_{C_{p_1}}}} & & \boxed{\phantom{A_{C_{p_1}}}} & \cdots & \boxed{\phantom{A_{C_{p_1}}}} & & \boxed{\phantom{A_{C_{p_1}}}} & \cdots & \boxed{\phantom{A_{C_{p_1}}}} \\ & & & \boxed{A_{C_{p_2}}} & \boxed{\phantom{A_{C_{p_2}}}} & \cdots & \boxed{\phantom{A_{C_{p_2}}}} & & \boxed{\phantom{A_{C_{p_2}}}} & \cdots & \boxed{\phantom{A_{C_{p_2}}}} \\ & & & & & & & \ddots & & & \\ & & & & & & & & \boxed{A_{C_{p_b}}} & \boxed{\phantom{A_{C_{p_b}}}} & \cdots & \boxed{\phantom{A_{C_{p_b}}}} \end{pmatrix}. \quad (7)$$

From this we see that $\text{rank}(\mathcal{A}_\theta) \geq \sum_{r=1}^{\text{rank}(B)} \text{rank}(A_{C_{p_r}})$. Rearranging the columns of B suitably, any subset $P \subseteq \mathcal{Y}$ with $|P| = \text{rank}(B_P) = \text{rank}(B)$ can be obtained as the set of pivots $p_1, \dots, p_{\text{rank}(B)}$ of the reduced row echelon form. For instance, Lemma 12 shows that, if $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_m$ and B has row span $V_{\Lambda'}(\mathcal{Y})$, then $|P| = \text{rank}(B_P) = \text{rank}(B)$ whenever P is a Λ' -ball in \mathcal{Y} . However, it may be difficult or even impossible to find a B -slicing of A such that $\text{rank}(A_{C_{p_r}}) = \text{rank}(A)$, for all $1 \leq r \leq \text{rank}(B)$. Nevertheless, in order to show that the tropical model $\mathcal{M}_{A \otimes B}^{\text{tropical}}$ has dimension equal to $\text{rank}(A) \cdot \text{rank}(B) - 1$, it suffices to show that $\text{rank}(A_{C_{p_r, 1}, \dots, C_{p_r, s_r}}) = \text{rank}(A)$ for all $1 \leq r \leq \text{rank}(B)$, where $(p_r, 1), \dots, (p_r, s_r)$ index the columns of (the reduced row echelon form of) B with a non-zero entry at the r -th position and zeros in all the next entries.

The following theorem addresses the tropical dimension of a Kronecker product model with an arbitrary hierarchical model \mathcal{E}_B and a k -interaction model \mathcal{E}_A . This result includes Theorems 16 and 22 as special cases. In fact it relaxes the hypothesis of Theorem 22.

Theorem 24. *Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ and let A have row span $V_k(\mathcal{X})$. Let $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_m$ and let B have row span $V_{\Lambda'}(\mathcal{Y})$.*

- *If \mathcal{X} contains $\text{rank}(B)$ disjoint radius- k Hamming balls, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = \text{rank}(B) \cdot \text{rank}(A) - 1.$$

- *If \mathcal{X} can be covered by $\text{rank}(B)$ disjoint radius- k Hamming balls, then*

$$\dim(\mathcal{M}_{A \otimes B}^{\text{tropical}}) = |\mathcal{X}| - 1.$$

Proof. The first item is as follows. Let $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_m$ and let B be the matrix defined in Proposition 10 with row span $V_{\Lambda}(\mathcal{Y})$. We can group the columns according to their largest non-zero entry. This yields one block of columns for each possible $(\lambda, \tilde{y}_{\lambda})$. The columns in the block $(\lambda, \tilde{y}_{\lambda})$ have a 1 in the $(\lambda, \tilde{y}_{\lambda})$ -th entry and zeros all the next entries.

Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ and let A be the matrix with row span $V_{\Lambda}(\mathcal{X})$, with entries $A_{(\lambda, \tilde{x}_{\lambda}), x}$ equal to 1 if $x_{\lambda} = \tilde{x}_{\lambda}$ and -1 otherwise, for $\lambda \in \Lambda$, $\tilde{x}_{\lambda} \in \mathcal{X}_{\lambda}$, and $x \in \mathcal{X}$. We consider the sufficient statistics matrix given by $(\mathbf{1}; A) \in \mathbb{R}^{a \times \mathcal{X}}$, where $a = 1 + \sum_{\lambda \in \Lambda \setminus \emptyset} |\mathcal{X}_{\lambda}|$.

Let $c_1, \dots, c_b \in \mathcal{X}$ be the centers of b disjoint Λ -balls in \mathcal{X} . Consider a parameter matrix $\Theta^{\top} \in \mathbb{R}^{a \times b}$ with columns $\Theta_j^{\top} = \kappa_j(2^{k+1} - a; A_{c_j})$, $j = 1, \dots, b$, where $\kappa_j \gg \sum_{j' < j} \|\Theta_{j'}^{\top}\|_1$. Then $\Theta^{\top} B_y = \kappa_j((2^{k+1} - a; A_{c_j}) + V_y)$ for all y in the j -th block of columns of B , where $V_y \in \mathbb{R}^a$ is some vector with $\|V_y\|_1 \ll 1$. In this case, we have that $\langle (1; A_x), \Theta^{\top} B_y \rangle = \kappa_j(2^{k+1} - a + \langle A_x, A_{c_j} \rangle + \epsilon)$. This is positive if $\langle A_x, A_{c_j} \rangle > a - 2^{k+1}$ and negative if $\langle A_x, A_{c_j} \rangle < a - 2^{k+1}$. Now, note that if $d_H(x, x') \leq k$, then $\langle A_x, A_{x'} \rangle \geq (a - 1) - 2(2^k - 1) = a - 2^{k+1} + 1$, and if $d_H(x, x') > k$, then $\langle A_x, A_{x'} \rangle \leq (a - 1) - 2(2^k) = a - 2^{k+1} - 1$.

In turn, we have that $\cup_{y: l(y)=j} C_y \supseteq K(c_j, \Lambda_k)$, for all $1 \leq j \leq b$, where $l(y)$ denotes the largest non-zero entry of B_y . Lemma 12 then yields the claim.

The second item is as follows. In this case $\cup_{y: l(y)=j} C_y \subseteq K(c_j, \Lambda_k)$, for all $1 \leq j \leq b$. Since $C_y \cap C_{y'} = \emptyset$ for all $y \neq y'$, the matrix $(A_{C_y})_{y: l(y)=j}$ has linearly independent columns, for all $1 \leq j \leq b$. \square

8 Binary Restricted Boltzmann Machines are Not Defective

In [6] it was shown that the restricted Boltzmann machine with n visible and m hidden binary units has the expected dimension $\min\{2^n - 1, (n + 1)(m + 1) - 1\}$ whenever $\{0, 1\}^n$ contains $m + 1$ disjoint radius-one Hamming balls, $m + 1 \leq \mathfrak{A}(n, 3)$, or

when $\{0, 1\}^n$ can be covered by $m + 1$ radius-one Hamming balls, $m + 1 \geq \mathfrak{K}(n, 1)$. This also follows from Theorems 22 and 24 as the special case where both \mathcal{E}_A and \mathcal{E}_B are independence models with binary variables. This leaves open the cases where $\mathfrak{A}(n, 3) < m + 1 < \mathfrak{K}(n, 1)$. An additional problem is that in general the functions \mathfrak{A} and \mathfrak{K} can only be bounded but not evaluated exactly. In [6] it was conjectured that the restricted Boltzmann machine always has the expected dimension. In this section we resolve that conjecture affirmatively.

Our strategy is to bound the dimension of the RBM from below by the dimension of a mixture model. By the results from [5], mixtures of binary independence models are not defective whenever the number of visible variables is at least 5. In general RBMs do not contain mixture models of independence models with the same number of parameters, as was shown in [20]. However, it is possible to relate the dimension of the two models. Here we consider the dimension of the actual model, not of its tropical version.

The following lemma lower bounds the dimension of a Kronecker product model with independent binary hidden variables by the dimension of a mixture model with the same number of parameters.

Lemma 25. *Let B have row span $V_1(\{0, 1\}^m)$ and let \tilde{B} have row span \mathbb{R}^{m+1} . Then $\dim(\mathcal{M}_{A \otimes B}) \geq \dim(\mathcal{M}_{A \otimes \tilde{B}})$.*

Proof. Let B be given by $B_y = (1; y_1; \dots; y_m)$, $y \in \{0, 1\}^m$, and let $\tilde{B} = I_{m+1}$ be the $(m + 1) \times (m + 1)$ identity matrix. For any $j \in [m]$ let $e_j \in \{0, 1\}^m$ denote the vector with a single one at the j -th position and zeros elsewhere. Any $y \in \{0, 1\}^m$ can be written as $y = \sum_{j \in [m]: e_j \leq y} e_j$. Here $y' \leq y$ if and only if $y'_j \leq y_j$ for all $j \in [m]$. As discussed in Equation 2, the Jacobian of the natural parametrization of a marginal model \mathcal{M}_F has rank $\text{rank}(J_{\mathcal{M}_F}(\theta)) = \text{rank}\left(\sum_y p_\theta(y|x) F(x, y)\right)_x - 1$. We have

$$\begin{aligned}
\sum_{y \in \{0, 1\}^m} p_\theta(y|x) (A_x \otimes B_y) &= A_x \otimes \sum_{y \in \{0, 1\}^m} p_\theta(y|x) B_y \\
&= A_x \otimes \left(1; \sum_{y \in \{0, 1\}^m} p_\theta(y|x) y\right) \\
&= A_x \otimes \left(1; \sum_{y \in \{0, 1\}^m} p_\theta(y|x) \sum_{j \in [m]: e_j \leq y} e_j\right) \\
&= A_x \otimes \left(1; \sum_{j \in [m]} \sum_{y \in \{0, 1\}^m: y \geq e_j} p_\theta(y|x) e_j\right) \\
&= A_x \otimes \left(1; \sum_{j \in [m]} p_\theta(y_j = 1|x) e_j\right).
\end{aligned}$$

The conditional distributions are given by

$$p_\theta(y_j = 1|x) = \frac{\exp(\Theta_j A_x)}{1 + \exp(\Theta_j A_x)}, \quad j \in [m], x \in \mathcal{X}.$$

On the other hand, for a hidden unit with sufficient statistics matrix $\tilde{B} = I_{m+1}$ we have

$$\sum_{j \in \{0, 1, \dots, m\}} \tilde{p}_{\tilde{\theta}}(j|x)(A_x \otimes \tilde{B}_j) = A_x \otimes \left(\tilde{p}_{\tilde{\theta}}(0|x); \sum_{j \in [m]} \tilde{p}_{\tilde{\theta}}(j|x)e_j \right).$$

In this case the conditional distributions are given by

$$\tilde{p}_{\tilde{\theta}}(j|x) = \frac{\exp(\tilde{\Theta}_j A_x)}{\sum_j \exp(\tilde{\Theta}_j A_x)}, \quad j \in \{0, 1, \dots, m\}, x \in \mathcal{X}.$$

Without loss of generality let the first row of A be $\mathbf{1}$. Consider the parameters $\Theta_{j,i} = \tilde{\Theta}_{j,i} - \tilde{\Theta}_{0,i}$, $i = 2, \dots, a$ and $\Theta_{j,1} = \tilde{\Theta}_{j,1} - \tilde{\Theta}_{0,1} - \gamma$. For any $\epsilon > 0$ we can choose γ large enough such that

$$\sum_{j=1}^m \left| \exp(\gamma) p_{\theta}(y_j = 1|x) - \frac{\tilde{p}_{\tilde{\theta}}(j|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \right| \leq \epsilon, \quad \text{for all } x \in \mathcal{X}.$$

This implies that

$$\begin{aligned} \dim(\mathcal{M}_{A \otimes B}) + 1 &= \max_{\theta} \text{rank} \left(A_x \otimes \left(\mathbf{1}; \sum_j p_{\theta}(y_j = 1|x)e_j \right) \right)_x \\ &\geq \max_{\tilde{\theta}} \text{rank} \left(A_x \otimes \left(\mathbf{1}; \exp(-\gamma) \sum_{j=1}^m \frac{\tilde{p}_{\tilde{\theta}}(j|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} e_j \right) \right)_x \\ &= \max_{\tilde{\theta}} \text{rank} \left(A_x \otimes \left(\tilde{p}_{\tilde{\theta}}(0|x); \sum_{j=1}^m \tilde{p}_{\tilde{\theta}}(j|x)e_j \right) \right)_x \\ &= \dim(\mathcal{M}_{A \otimes I_{m+1}}) + 1. \end{aligned}$$

This completes the proof. See Example 27 in the Appendix for a more explicit formulation of this proof in the case $m = 2$. \square

The dimension bound from Lemma 25 is not always tight. For example, the 3-mixture of 4 independent binary variables is defective and has dimension 13, whereas the RBM with 4 visible and 2 hidden binary units has the expected dimension 14.

Corollary 26. *Let n and m be non-negative integers. The restricted Boltzmann machine with n visible and m hidden binary units has dimension $\min\{2^n - 1, (n+1)(m+1) - 1\}$.*

Proof. Let A, B, \tilde{B} be matrices with row span $V_1(\{0, 1\}^n)$, $V_1(\{0, 1\}^m)$, \mathbb{R}^{m+1} , respectively. Then $\mathcal{M}_{A \otimes B}$ is the restricted Boltzmann machine with n visible and m hidden binary variables and $\mathcal{M}_{A \otimes \tilde{B}}$ is the $(m+1)$ -mixture of the independence model of n binary variables.

The work [5] shows that $\dim(\mathcal{M}_{A \otimes \tilde{B}}) = \min\{2^n - 1, (m + 1)(n + 1) - 1\}$ unless $(n, m) = (4, 2)$. Lemma 25 then implies $\dim(\mathcal{M}_{A \otimes B}) \geq \dim(\mathcal{M}_{A \otimes \tilde{B}}) = \min\{2^n - 1, (m + 1)(n + 1) - 1\}$ whenever $(n, m) \neq (4, 2)$. Since this is also the maximum possible dimension, the bound is tight. That the RBM with $(n, m) = (4, 2)$ has the expected dimension is shown in [9, 7]. \square

An RBM with 5 visible units is the first case with a gap between the largest RBM previously known to have dimension equal to the number of parameters (4 hidden units and 29 parameters) and the smallest RBM previously known to be full dimensional (7 hidden units and 47 parameters suffice to define a 31-dimensional subset of the 31-simplex) as described by [6], which lists a collection of such gaps in Table 4.1 for $5 \leq n \leq 512$. Corollary 26 closes all such gaps. Thus the smallest RBM with 5 visible units which could possibly define a full-dimensional subset of the simplex, with 5 hidden units and 35 parameters, does so. These examples can be tested by computing Jacobians with the Matlab code provided at <http://personal-homepages.mis.mpg.de/montufar>.

Corollary 26 proves that a binary RBM always has the expected dimension. The conjecture posed in [6] goes further and states that the tropical binary RBM always has the expected dimension. That question remains unsettled at this point.

9 Conclusion

In this work we study the dimension of marginals of exponential families whose sufficient statistics matrix factorizes as the Kronecker product of a visible and a hidden sufficient statistics matrix, called Kronecker product models. The Jacobian of these models factorizes as a Khatri-Rao product of the visible sufficient statistics matrix and the expectation parameters of the hidden exponential family given the visible variables. The tropical morphism arises as the limit of the Jacobian when the natural parameters of the model are scaled by an infinitely large number. It is described by the inference functions of the hidden variables given the visible variables, which correspond to slicings of the visible sufficient statistics matrix by the normal fan of the hidden sufficient statistics matrix.

Based on these geometric and combinatorial descriptions, we computed the tropical dimension of mixtures of interaction models and Hadamard products of mixtures of interaction models. These results extend previous work on secant dimensions, which are most often focused on Segre and Veronese varieties (corresponding to independence models and multinomial models). These results also generalize previous work on Hadamard products, which were focused on products of mixtures of independence models. Theorem 24 generalizes this further to the case of Kronecker products of arbitrary hierarchical models and k -interaction models. Additionally, we showed that binary restricted Boltzmann machines always have the expected dimension, thus completing the dimension description of these models from [6].

Our analysis leaves many questions unanswered. In this work we have focused on the case where the visible exponential family is a k -interaction model. The generalization to arbitrary visible hierarchical models is left for future work. Furthermore, similarly to [8], the tropical approach leads in many cases to combinatorial conditions

that can be very difficult to verify outside of well established cardinality bounds for error correcting codes. We think that a promising direction is the formulation of dimension bounds in terms of simpler models, as done in Lemma 25 for the case of a binary independence hidden model. Extending that result one could ask: when is the dimension of $\mathcal{M}_{A \otimes B}$ bounded below by the dimension of $\mathcal{M}_{A \otimes I}$, where I is an identity matrix with the same rank as B ?

The factorization property of the Jacobian of Kronecker product models suggests to study the models of conditional probability distributions of the hidden variables given the visible variables in more detail. This is a manifold of tuples of exponential family distributions with natural parameters given by the linear projection of the sufficient statistics of the visible model. An analysis of the Kruskal ranks for these sets can be used to obtain bounds on the Kruskal rank of the Jacobian.

Another interesting line of investigation is the classification of the support sets of distributions in the closure of Kronecker product models. For mixture models the problem is simple, when the support sets of the visible exponential family are known. For Hadamard products the problem has been studied in [20] based on linear threshold codes, which are the images of inference functions. We think that studying the combinatorics of Kronecker product polytopes could yield helpful insights. Given two polytopes $P_A = \text{conv}\{A_x: x \in \mathcal{X}\}$ and $P_B = \text{conv}\{B_y: y \in \mathcal{Y}\}$, we define the Kronecker product polytope as $P_{A \otimes B} = \text{conv}\{A_x \otimes B_y: (x, y) \in \mathcal{X} \times \mathcal{Y}\}$. Although this appears as a rather natural composition of polytopes, we are not aware of works studying such objects explicitly or in a principled way.

Appendix

A Examples

Example 27. Here we give a more comprehensive version of the proof of Lemma 25 for the special case where $m = 2$. Consider the matrices

$$B = \begin{pmatrix} & 00 & 01 & 10 & 11 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \tilde{B} = \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

We have

$$\begin{aligned} \sum_{y \in \{0,1\}^2} p_\theta(y|x)(A_x \otimes B_y) &= A_x \otimes \sum_{y \in \{0,1\}^2} p_\theta(y|x)B_y \\ &= A_x \otimes \begin{pmatrix} 1 \\ p_\theta(10|x) + p_\theta(11|x) \\ p_\theta(01|x) + p_\theta(11|x) \end{pmatrix} \\ &= A_x \otimes \begin{pmatrix} 1 \\ p_\theta(y_1 = 1|x) \\ p_\theta(y_2 = 1|x) \end{pmatrix}. \end{aligned}$$

By Equation 4 the conditional distributions are given by

$$p_\theta(y|x) = \frac{1}{Z(\Theta A_x)} \exp(\langle \Theta A_x, B_y \rangle)$$

such that

$$\begin{aligned} p_\theta(y_1 = 1|x) &= \frac{\exp(\langle \Theta A_x, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \rangle)}{\exp(\langle \Theta A_x, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \rangle) + \exp(\langle \Theta A_x, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \rangle)} \\ &= \frac{\exp(\Theta_1 A_x)}{1 + \exp(\Theta_1 A_x)}, \end{aligned}$$

and similarly for $p_\theta(y_2 = 1|x)$.

On the other hand, for a hidden unit with sufficient statistics matrix \tilde{B} we have

$$\sum_{j \in \{0,1,2\}} \tilde{p}_{\tilde{\theta}}(A_x \otimes \tilde{B}_j) = A_x \otimes \begin{pmatrix} \tilde{p}_{\tilde{\theta}}(0|x) \\ \tilde{p}_{\tilde{\theta}}(1|x) \\ \tilde{p}_{\tilde{\theta}}(2|x) \end{pmatrix}.$$

In this case the conditional distributions are given by

$$\tilde{p}_{\tilde{\theta}}(1|x) = \frac{\exp(\tilde{\Theta}_1 A_x)}{\exp(\tilde{\Theta}_0 A_x) + \exp(\tilde{\Theta}_1 A_x) + \exp(\tilde{\Theta}_2 A_x)},$$

and similarly for $j = 0$ and $j = 2$.

For any given

$$\tilde{\Theta} = \begin{pmatrix} \tilde{\Theta}_{0,1} & \tilde{\Theta}_{0,2} & \cdots & \tilde{\Theta}_{0,a} \\ \tilde{\Theta}_{1,1} & \tilde{\Theta}_{1,2} & \cdots & \tilde{\Theta}_{1,a} \\ \tilde{\Theta}_{2,1} & \tilde{\Theta}_{2,2} & \cdots & \tilde{\Theta}_{2,a} \end{pmatrix}$$

we can define

$$\Theta = \begin{pmatrix} \tilde{\Theta}_{0,1} - \tilde{\Theta}_{0,1} - \gamma & \tilde{\Theta}_{0,2} - \tilde{\Theta}_{0,2} & \cdots & \tilde{\Theta}_{0,a} - \tilde{\Theta}_{0,a} \\ \tilde{\Theta}_{1,1} - \tilde{\Theta}_{0,1} - \gamma & \tilde{\Theta}_{1,2} - \tilde{\Theta}_{0,2} & \cdots & \tilde{\Theta}_{1,a} - \tilde{\Theta}_{0,a} \\ \tilde{\Theta}_{2,1} - \tilde{\Theta}_{0,1} - \gamma & \tilde{\Theta}_{2,2} - \tilde{\Theta}_{0,2} & \cdots & \tilde{\Theta}_{2,a} - \tilde{\Theta}_{0,a} \end{pmatrix}.$$

Without loss of generality assume that the first row of A is a row of ones. We have then

$$\begin{pmatrix} 1 \\ p_\theta(y_1 = 1|x) \\ p_\theta(y_2 = 1|x) \end{pmatrix} = \begin{pmatrix} 1 \\ \frac{\exp(\Theta_1 A_x)}{1 + \exp(\Theta_1 A_x)} \\ \frac{\exp(\Theta_2 A_x)}{1 + \exp(\Theta_2 A_x)} \end{pmatrix} = \begin{pmatrix} 1 \\ \frac{\exp(-\gamma + (\tilde{\Theta}_1 - \tilde{\Theta}_0) A_x)}{1 + \exp(-\gamma + (\tilde{\Theta}_1 - \tilde{\Theta}_0) A_x)} \\ \frac{\exp(-\gamma + (\tilde{\Theta}_2 - \tilde{\Theta}_0) A_x)}{1 + \exp(-\gamma + (\tilde{\Theta}_2 - \tilde{\Theta}_0) A_x)} \end{pmatrix}.$$

For γ large enough we obtain an arbitrarily accurate approximation of the form

$$\begin{pmatrix} 1 \\ p_\theta(y_1 = 1|x) \\ p_\theta(y_2 = 1|x) \end{pmatrix} \approx \begin{pmatrix} 1 \\ \exp(-\gamma) \exp((\tilde{\Theta}_1 - \tilde{\Theta}_0) A_x) \\ \exp(-\gamma) \exp((\tilde{\Theta}_2 - \tilde{\Theta}_0) A_x) \end{pmatrix} = \begin{pmatrix} 1 \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(1|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(2|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \end{pmatrix}.$$

Hence for any $\tilde{\theta}$ there is a θ with

$$\begin{aligned}
\text{rank} \left(A_x \otimes \begin{pmatrix} \tilde{p}_{\tilde{\theta}}(0|x) \\ \tilde{p}_{\tilde{\theta}}(1|x) \\ \tilde{p}_{\tilde{\theta}}(2|x) \end{pmatrix} \right)_x &= \text{rank} \left(L \cdot \left(A_x \otimes \begin{pmatrix} 1 \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(1|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(2|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \end{pmatrix} \right) \cdot R \right) \\
&= \text{rank} \left(A_x \otimes \begin{pmatrix} 1 \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(1|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(2|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \end{pmatrix} \right)_x \\
&\leq \text{rank} \left(\left(A_x \otimes \begin{pmatrix} 1 \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(1|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \\ \exp(-\gamma) \frac{\tilde{p}_{\tilde{\theta}}(2|x)}{\tilde{p}_{\tilde{\theta}}(0|x)} \end{pmatrix} \right) + \epsilon \right) \\
&= \text{rank} \left(A_x \otimes \begin{pmatrix} 1 \\ p_{\theta}(y_1 = 1|x) \\ p_{\theta}(y_2 = 1|x) \end{pmatrix} \right)_x,
\end{aligned}$$

where

$$L = \begin{pmatrix} 1 & & \\ & \ddots & \\ & & 1 \end{pmatrix} \otimes \begin{pmatrix} 1 & & \\ & \exp(\gamma) & \\ & & \exp(\gamma) \end{pmatrix} \quad \text{and} \quad R = \begin{pmatrix} \tilde{p}_{\tilde{\theta}}(0|1) & & \\ & \ddots & \\ & & \tilde{p}_{\tilde{\theta}}(0|a) \end{pmatrix}.$$

In turn

$$\max_{\tilde{\theta}} \text{rank} \left(A_x \otimes \begin{pmatrix} \tilde{p}_{\tilde{\theta}}(0|x) \\ \tilde{p}_{\tilde{\theta}}(1|x) \\ \tilde{p}_{\tilde{\theta}}(2|x) \end{pmatrix} \right)_x \leq \max_{\theta} \text{rank} \left(A_x \otimes \begin{pmatrix} 1 \\ p_{\theta}(y_1 = 1|x) \\ p_{\theta}(y_2 = 1|x) \end{pmatrix} \right)_x$$

and

$$\dim(\mathcal{M}_{A \otimes \tilde{B}}) \leq \dim(\mathcal{M}_{A \otimes B}).$$

B Simple Bounds for Error Correcting Codes

Some of our results are formulated in terms of the maximal number of disjoint Λ -balls that can be fit in some $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$. Let $\mathfrak{A}(\mathcal{X}, d)$ denote the maximal cardinality of a subset of \mathcal{X} of minimum Hamming distance at least d . If $\mathcal{X} = \{0, 1, \dots, q-1\}^n$, we write $\mathfrak{A}_q(n, d)$. Closed forms for these functions are known only in special cases. We recall the Gilbert-Varshamov bound [12, 25]:

$$\mathfrak{A}_q(n, d) \geq \frac{q^n}{\sum_{j=0}^{d-1} \binom{n}{j} (q-1)^j}.$$

If q is a prime power, then

$$\mathfrak{A}_q(n, d) \geq q^{n-1 - \lfloor \log_q(\sum_{j=0}^{d-2} \binom{n-1}{j} (q-1)^j) \rfloor}.$$

A simple upper bound is the sphere packing bound, $\mathfrak{A}_q(n, d) \leq q^n / K_q(t)$, $t = \lfloor \frac{d-1}{2} \rfloor$.

Theorem 22 is formulated in terms of the maximal number of disjoint Hamming balls that fit in a larger Hamming ball. An estimate for this number can be given as follows.

Proposition 28. *Let $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, and let $k \leq l \leq n$. Then it is possible to fit $|K(l-k)|/|K(2k)|$ disjoint radius- k Hamming balls in a radius- l Hamming ball.*

Proof. Denote by C the set of centers of a largest possible collection of disjoint radius- k Hamming balls contained in a radius- l Hamming ball $K(0, l)$. Consider the radius- $(l+k+1)$ sphere S . Let $d = 2k+1$. For every $x \in K(0, l+k+1)$ there is a $c_x \in C \cup S$ such that $d_H(x, c_x) \leq d-1$. This implies that $K(0, l+k+1) \subseteq \cup_{c \in C} K(c, d-1) \cup (K(0, l+k+1) \setminus K(0, l-k))$ and $K(0, l-k) \subseteq \cup_{c \in C} K(c, d-1)$. Therefore, $|C| \geq |K(l-k)|/|K(2k)|$. \square

Acknowledgments This work was supported in part by DARPA grant FA8650-11-1-7145. Parts of this work were carried out while G.M. was at the Pennsylvania State University.

References

- [1] H. Abo and M. Brambilla. On the dimensions of secant varieties of Segre-Veronese varieties. *Annali di Matematica Pura ed Applicata*, 192(1):61–92, 2013.
- [2] H. Abo, G. Ottaviani, and C. Peterson. Induction for secant varieties of Segre varieties. *Transactions of the American Mathematical Society*, 361(2):767–792, 2008.
- [3] R. Bieri and J. Groves. The geometry of the set of characters induced by valuations. *Journal für die reine und angewandte Mathematik*, 347:168–195, 1984.
- [4] M. Catalisano, A. Geramita, and A. Gimigliano. Ranks of tensors, secant varieties of segre varieties and fat points. *Linear Algebra and its Applications*, 355(13):263–285, 2002.
- [5] M. Catalisano, A. Geramita, and A. Gimigliano. Secant varieties of $\mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ (n -times) are not defective for $n \geq 5$. *J. Algebraic Geometry*, 20:295–327, 2011.
- [6] M. A. Cueto, J. Morton, and B. Sturmfels. Geometry of the restricted Boltzmann machine. In M. A. G. Viana and H. P. Wynn, editors, *Algebraic methods in statistics and probability II, AMS Special Session*, volume 2. American Mathematical Society, 2010.
- [7] M. A. Cueto, E. A. Tobis, and J. Yu. An implicitization challenge for binary factor analysis. *J. Symb. Comput.*, 45(12):1296–1315, 2010.
- [8] J. Draisma. A tropical approach to secant dimensions. *J. Pure Appl. Algebra*, 212(2):349–363, 2008.

- [9] M. Drton, B. Sturmfels, and S. Sullivant. *Lectures on Algebraic Statistics*. Oberwolfach Seminars. Springer Verlag, 2009.
- [10] Y. Freund and D. Haussler. Unsupervised learning of distributions of binary vectors using 2-layer networks. In J. E. Moody, S. J. Hanson, and R. Lippmann, editors, *Advances in Neural Information Processing Systems 4*, pages 912–919. Morgan Kaufmann, 1991.
- [11] D. Geiger, C. Meek, and B. Sturmfels. On the toric algebra of graphical models. *Annals of Statistics*, 34(3):1463–1492, 2006.
- [12] E. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31:504–522, 1952.
- [13] T. Kahle. Neighborliness of marginal polytopes. *Beiträge zur Algebra und Geometrie*, 51(1):45–56, 2010.
- [14] T. Kahle, W. Wenzel, and N. Ay. Hierarchical models, marginal polytopes, and linear codes. *Kybernetika*, 45(2):189–207, 2009.
- [15] C. G. Khatri and C. R. Rao. Solutions to some functional equations and their applications to characterization of probability distributions. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)*, 30(2):pp. 167–180, 1968.
- [16] J. Landsberg. *Tensors: Geometry and Applications*. Graduate studies in mathematics. American Mathematical Soc., 2011.
- [17] L. D. Lathauwer. Decompositions of a higher-order tensor in block terms—part I: Lemmas for partitioned matrices. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1022–1032, 2008.
- [18] G. Montúfar. Mixture decompositions of exponential families using a decomposition of their sample spaces. *Kybernetika*, 49(1):23–39, 2013.
- [19] G. Montúfar and J. Morton. Discrete restricted Boltzmann machines. *Journal of Machine Learning Research*, 16:653–672, 2015.
- [20] G. Montúfar and J. Morton. When does a mixture of products contain a product of mixtures? *SIAM Journal on Discrete Mathematics*, 29(1):321–347, 2015.
- [21] C. Raicu. Secant varieties of segre–veronese varieties. *Algebra & Number Theory*, 6(8):1817–1868, 2012.
- [22] J. Rauh, T. Kahle, and N. Ay. Support sets of exponential families and oriented matroids. *International Journal of Approximate Reasoning*, 52(5):613–626, 2011.
- [23] W. E. Roth. On direct product matrices. *Bulletin of the American Mathematical Society*, 40:461–468, 1934.
- [24] N. Sidiropoulos and X. Liu. Identifiability results for blind beamforming in incoherent multipath with small delay spread. *Signal Processing, IEEE Transactions on*, 49(1):228–236, Jan 2001.

- [25] R. Varshamov. Estimate of the number of signals in error correcting codes. *Doklady Akad. Nauk SSSR*, 117:739–741, 1957.
- [26] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 1(1-2):1–305, 2008.
- [27] M. Welling, M. Rosen-Zvi, and G. E. Hinton. Exponential family harmoniums with an application to information retrieval. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1481–1488. MIT Press, Cambridge, MA, 2005.
- [28] F. Zak. *Tangents and Secants of Algebraic Varieties*, volume 127 of *Translations of mathematical monographs*. American Mathematical Soc., Providence, RI, 1993.