



Classical iterative proportional scaling of log-linear models with rational maximum likelihood estimator

Jane Ivy Coons^{a,b,*}, Carlotta Langer^c, Michael Ruddy^d

^a St John's College, University of Oxford, United Kingdom

^b Mathematical Institute, University of Oxford, United Kingdom

^c Hamburg University of Technology, Hamburg, Germany

^d Qventus, United States of America

ARTICLE INFO

Keywords:

Maximum likelihood estimation
Iterative proportional scaling
Log-linear models
Hierarchical models
Staged tree models
Toric fiber product

ABSTRACT

In this work we investigate *multipartition models*, the subset of log-linear models for which one can perform the classical iterative proportional scaling (IPS) algorithm to numerically compute the maximum likelihood estimate (MLE). Multipartition models include families of models such as hierarchical models and balanced, stratified staged trees. We define a sufficient condition, called the Generalized Running Intersection Property (GRIP), on the matrix representation of a multipartition model under which the classical IPS algorithm produces the exact MLE in one cycle. In this case, the MLE is a rational function of the data. Additionally we connect the GRIP to the toric fiber product and to previous results for hierarchical models and balanced, stratified staged trees. This leads to a characterization of balanced, stratified staged trees in terms of the GRIP.

1. Introduction

The iterative proportional scaling (IPS) algorithm is a simple and efficient numerical algorithm for computing the maximum likelihood estimate (MLE) for certain families of log-linear statistical models, which we call *multipartition models*. The IPS algorithm is widely used in survey statistics, such as in [9,28,43], and has been researched in connection to optimal transport problems in its more restricted form as Sinkhorn's algorithm [6,36]. Multipartition models include hierarchical models, which have been the subject of much study in connection with the IPS algorithm [20,26,42]. Maximum likelihood estimation for log-linear models and its relationship to algebraic and combinatorial objects is also of interest in algebraic statistics. For example in [2], the authors connect invariant theory and algorithms to compute the MLE including a form of the IPS algorithm. Many recent works have also explored the number of critical points of the complex log-likelihood function, also known as the ML-degree (see [1,8,18,29,30]).

From an information-geometric perspective, calculating the MLE for a multipartition model can be described as projecting to linear families defined by the partitions of the matrix representing the model (see Definitions 2.1 and 2.3). Given a data vector and an estimate on the model, the IPS algorithm updates this estimate at each step by projecting onto a different linear family, converging towards the MLE. We say that the IPS algorithm has completed *one cycle* after it has iterated through each linear family exactly once.

At each step of the IPS algorithm, the estimate is a rational function of the data vector. Thus in the case of one-cycle convergence, the MLE itself is a rational function of the data vector. Models for which the MLE is a rational function of the data vector are called

* Corresponding author.

E-mail addresses: jane.coons@maths.ox.ac.uk (J.I. Coons), carlotta.langer@tuhh.de (C. Langer), mruddy@qventus.com (M. Ruddy).

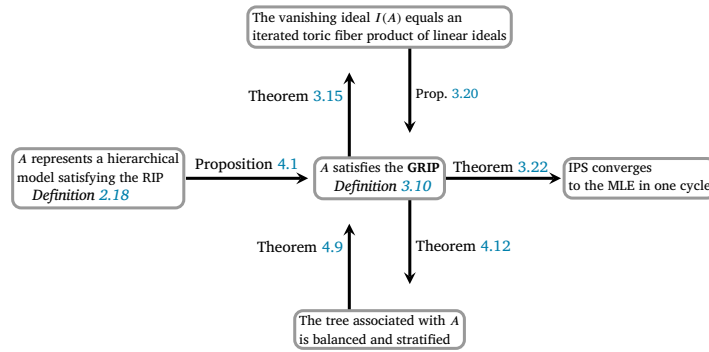


Fig. 1. Diagram of the main results of this work.

rational and are a subject of recent interest [11,18,31]. This is in contrast to what is sometimes called the *generalized*¹ IPS algorithm, another numerical algorithm for computing the MLE for log-linear models, which does not produce a rational function at each step, yet can be applied to any log-linear model.

In this work we are interested in the question, “When does the IPS algorithm exactly produces the MLE after one cycle?” We first note that the matrix representation of a multipartition model heavily influences the outcome of the IPS algorithm (see Example 2.11). In [26] the author defines the Running Intersection Property (RIP) for hierarchical models, which gives sufficient conditions on the matrix representation of a hierarchical model so that the IPS algorithm always produces the MLE exactly in one cycle. The author shows that for decomposable hierarchical models, there always exists such a representation. Drawing inspiration from the RIP, we define the Generalized Running Intersection Property (GRIP) on the matrix representations of general multipartition models and show that it gives sufficient conditions for the IPS algorithm to always produce the MLE exactly in one cycle. In the case of hierarchical models we show that the RIP is a special case of the GRIP (see Proposition 4.1). We also connect previous work on the rationality of hierarchical models with the GRIP.

We additionally investigate another family of multipartition models: balanced and stratified staged trees. Staged tree models, also known as chain event graphs, are probability tree models that encode conditional independence relationships among events. They have been the subject of much study from an algebraic perspective in recent years [4,18,25,38]. In this work we connect the GRIP to balanced and stratified staged trees. One can associate a tree to any matrix representation of a multipartition model, as described in Appendix A. We show that if a matrix representation satisfies the GRIP, then the associated tree must be a balanced and stratified staged tree. Moreover we show that the unique matrix with no repeated columns associated to a balanced and stratified staged tree satisfies the GRIP. In [19], the authors show that decomposable hierarchical models are also balanced staged tree models and claim that balanced staged trees are a natural generalization of such models. Our work also implies this result and supports this claim.

We make use of the toric geometry of a multipartition model in order to determine whether IPS achieves one-cycle convergence on a given parametrization of the model. In particular, we show that when a parametrization of the model satisfies the GRIP, the model is a *toric fiber product* of smaller rational multipartition models [39]. This in turn allows us to write the MLE of the larger multipartition model as a normalized product of the MLEs of the smaller models [3]. Our proof uses this rational form of the MLE and relates it to convergence of the IPS algorithm in one step.

Our work draws novel connections between the geometry of a rational multipartition model and the performance of the iterative proportional scaling algorithm applied to it. Moreover, it gives sufficient conditions for a multipartition model to be rational. If the multipartition model has a parametrization which satisfies the GRIP, then it is rational and one can read the MLE directly from the matrix of this parametrization (Corollary 3.18). These results highlight the significant role that model representation plays in the performance of numerical computations of MLEs. We include a diagram of our main results in Fig. 1, in which the matrix A is a multipartition matrix representing a log-linear statistical model.

The work is structured as follows. In Section 2 we introduce some basic notions and previous results surrounding log-linear models, the MLE, the IPS algorithm, and some important multipartition model families. In particular, we define partition matrices, multipartition matrices, and multipartition models in 2.1; describe the IPS algorithm in 2.2; introduce the toric fiber product in 2.3; and discuss hierarchical models in Section 2.4. In Section 3 we introduce the GRIP and show this implies that the IPS algorithm produces the MLE after one cycle (Theorem 3.22). In Section 4.2 we investigate the connection between the GRIP and staged tree models, showing that the GRIP characterizes a subset of such models. The necessary background on staged tree models is discussed in Appendix A. In Section 4.1 we show that the GRIP is indeed a generalization of the RIP. We end the work with a discussion of open problems stemming from our results in Section 5.

¹ Many works will refer to the generalized IPS algorithm as simply “the IPS algorithm” while referring to its predecessor, the algorithm we consider in this work, as the “classical IPS algorithm.”

2. Preliminaries

2.1. Log-linear models

In this section we introduce the class of models to which we may apply the iterative proportional scaling algorithm in order to compute the maximum likelihood estimate. We start with some background on log-linear models before defining our main object of interest, *multipartition models*.

Consider an $n \times m$ matrix $A = (a_{ij})$ with the vector of all ones in its rowspan. Then the matrix A defines a homogeneous polynomial map $\phi_A: \mathbb{R}^n \rightarrow \mathbb{R}^m$ where

$$\phi_A(t_1, \dots, t_n) = \left(\prod_{i=1}^n t_i^{a_{i1}}, \prod_{i=1}^n t_i^{a_{i2}}, \dots, \prod_{i=1}^n t_i^{a_{im}} \right). \quad (1)$$

Let $\Delta_{m-1} \subset \mathbb{R}^m$ denote the open $(m-1)$ -dimensional probability simplex. The **support** of a row a of A , denoted $\text{supp}(a)$ is the set of column indices in which a has a nonzero entry.

Definition 2.1. The **log-linear** model, denoted \mathcal{M}_A , associated to the integer matrix A described above is the intersection of the image of ϕ_A with the open probability simplex; that is, $\mathcal{M}_A = \text{Im}(\phi_A) \cap \Delta_{m-1}$. This is the discrete exponential family whose sufficient statistics are given by the rows of A . Hence it is exactly

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

We note that several integer matrices with the vector of all ones in their rowspan may give rise to the same log-linear model. In this work we are concerned with the relationship between the IPS algorithm and a particular choice of matrix A . This matrix should be thought of as a particular representation of the log-linear model \mathcal{M}_A .

Note that the map ϕ_A is a monomial map. Hence the Zariski closure of \mathcal{M}_A is a *toric variety*; for this reason the model \mathcal{M}_A is also referred to as a *toric model* in the algebraic statistics literature. Let $I(A)$ denote the ideal of polynomials that vanishes on \mathcal{M}_A . The following proposition, adapted from [40, Prop 6.2.4] describes the relationship between the matrix A and the ideal of polynomials vanishing on \mathcal{M}_A .

Proposition 2.2. Let $A \in \mathbb{Z}^{n \times m}$. Then the vanishing ideal of \mathcal{M}_A ,

$$I(A) = \langle p^x - p^y \mid x, y \in \mathbb{Z}_{\geq 0}^m \text{ and } Ax = Ay \rangle \quad (2)$$

is the **toric ideal** of A and \mathcal{M}_A is the intersection of the variety $V(I(A))$ with the open simplex Δ_{m-1} . Moreover, if $A' \in \mathbb{Z}^{n' \times m}$ satisfies $\text{rowspan}(A) = \text{rowspan}(A')$, then $I(A) = I(A')$ and $\mathcal{M}_A = \mathcal{M}_{A'}$.

The integer matrices A that we can apply IPS to must satisfy several conditions, which we outline below.

Definition 2.3. A matrix $A \in \{0, 1\}^{n \times m}$ is a **multipartition matrix** if it consists of consecutive blocks of rows A^1, \dots, A^k such that in each A^ℓ , $\ell \in \{1, \dots, k\}$, the entries of each column sum to 1. The matrices A^1, \dots, A^k are called the **partition matrices** of A . Such a matrix gives rise to a **multipartition model** \mathcal{M}_A .

The monomial map that parametrizes a multipartition model as described in (1) is homogeneous and multilinear.

Remark 2.4. In [37] the author uses the name “partition model” for the log-linear model associated to one partition matrix. He studies these models as an example of a family of log-linear models for which there exists a low constant that bounds the KL-divergence from an arbitrary point in the simplex to these models. Thereby he proves additional properties of the models defined by one partition matrix.

Let A^1, A^2 be partition matrices, each with m columns. Let $A^1 A^2$ denote the matrix obtained by stacking A^1 above A^2 ; that is,

$$A^1 A^2 := \begin{bmatrix} A^1 \\ A^2 \end{bmatrix}.$$

Suppose A is a multipartition matrix with m columns and partition matrices denoted by A^1, \dots, A^k . From the partition matrices of A we build a new matrix, by stacking the partitions as defined above $A^{1, \dots, k} := A^1 A^2 \dots A^k$. Permuting the rows of $A^{1, \dots, k}$ in any way, or indeed, replacing $A^{1, \dots, k}$ with any matrix with the same rowspan, yields the same toric model. However, as we show in the next subsection, a different representation of the same model may affect the convergence of the iterative proportional scaling algorithm. We assume throughout that A is a multipartition matrix of the form $A^{1, \dots, k}$ and we refer to the partition matrices of a multipartition matrix as “partitions” for short.

$$j \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14\}$$

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 \\ 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 \end{pmatrix} \begin{matrix} a_1^1 \\ a_2^1 \\ \hline a_1^2 \\ a_2^2 \\ \hline a_1^3 \\ a_2^3 \\ a_3^3 \\ a_4^3 \\ a_5^3 \end{matrix} \left. \begin{matrix} \left. \begin{matrix} a_1^1 \\ a_2^1 \end{matrix} \right\} A^1 \\ \left. \begin{matrix} a_1^2 \\ a_2^2 \end{matrix} \right\} A^2 \\ \left. \begin{matrix} a_1^3 \\ a_2^3 \\ a_3^3 \\ a_4^3 \\ a_5^3 \end{matrix} \right\} A^3 \end{matrix} \right\}$$

Fig. 2. A multipartition matrix with three partitions and the used row and column notation.

Example 2.5. The matrix A in Fig. 2 is an example of a multipartition matrix with three partitions and $m = 14$.

Definition 2.6. Let a_i^ℓ denote the i -th row of the ℓ -th partition, as illustrated in Example 2.5. Let I_i^ℓ denote the set of column indices $j \in [m]$ such that the j -th entry of a_i^ℓ is equal to 1. That is, I_i^ℓ is exactly the support of a_i^ℓ . Let n_ℓ denote the number of rows of A^ℓ .

For a fixed partition A^ℓ and index $j \in \{1, \dots, m\}$ there is exactly one row a_i^ℓ with j in its support. We define the function $S(\ell, j) \in \{1, \dots, n_\ell\}$ where $S(\ell, j)$ is the index such that $j \in I_{S(\ell, j)}^\ell$. Thus $a_{S(\ell, j)}^\ell$ is the row of A^ℓ that has j in its support.

In Example 2.5 the matrix has three partitions and the number of rows of these partitions are $n_1 = n_2 = 2$ and $n_3 = 5$. For instance, the support of the second row of the second partition is $I_2^2 = \{4, 5, 6, 7, 11, 12, 13, 14\}$ and the row of the first partition that has the tenth column in its support is $S(1, 10) = 2$.

Definition 2.7. For a matrix A the **column weight** of the j -th column is the number of times the column is repeated and is denoted $c_j \in \mathbb{Z}_+$.

For a multipartition matrix A , let \bar{A} denote the $n \times \bar{m}$ matrix obtained by removing all repeated columns. This is then a k -way **quasi-independence model**. In fact what we call multipartition models can also be thought of as k -way quasi-independence models with repeated columns. We refer the reader to [11] for more information on the MLE of 2-way quasi-independence models.

We are able to apply the iterative proportional scaling algorithm as described in the next section to a multipartition model \mathcal{M}_A in order to compute the maximum likelihood estimate of a data vector in \mathbb{R}_+^m . These are in fact the only integer matrices defining log-linear models for which we can apply the classical IPS algorithm.

2.2. Iterative proportional scaling

The iterative proportional scaling (IPS) algorithm is a method to calculate the maximum likelihood estimator with respect to the model \mathcal{M}_A . Before we discuss the algorithm in more detail, we first define the maximum likelihood method.

Broadly speaking, maximum likelihood estimation is a way to find a distribution in a statistical model, \mathcal{M}_A that best fits some observed data. Let u be the vector of counts of the observed data. Assuming that the observations are independent and identically distributed (i.i.d.), the **likelihood function** for u with respect to the model is the function $L(p | u) : \mathcal{M}_A \rightarrow \mathbb{R}$ defined by

$$L(p | u) = \prod_{j=1}^m p_j^{u_j}.$$

Evaluating the likelihood function at $p \in \mathcal{M}_A$ yields the probability of observing the data u from the distribution p . The **maximum likelihood estimate** (or MLE) for u is then

$$\arg \max_{p \in \mathcal{M}_A} L(p | u).$$

That is, the MLE for u is the distribution in the model that maximizes the probability of observing u . We assume throughout that the entries of u are positive so that the MLE is guaranteed to exist, see for example [40, Theorem 8.2.1]. We note that solving this optimization problem is equivalent to maximizing the **log-likelihood function**, $\log(L(p | u))$, as the logarithm is monotonic. We use \bar{u} throughout to refer to *normalized* data, which is the data vector divided by the sum of its entries. We denote the MLE for \bar{u} by $p^*(\bar{u})$, or simply by p^* when \bar{u} is understood. One can find an introduction to maximum likelihood estimation in most statistics resources, such as [27, 33, 34]. The MLE problem is also described from an algebraic perspective in [40, Chapter 6].

Definition 2.8. If $p^*(\bar{u})$ is a rational function in the entries of \bar{u} whenever it exists, then we say that \mathcal{M}_A has **rational MLE**. If moreover \mathcal{M}_A is a multipartition model, then we call it a **rational multipartition model**.




Fig. 3. Sketch of the IPS algorithm in the case of three linear families.

The maximum likelihood estimate of any log-linear model is characterized by the following result. This result is sometimes referred to as Birch's theorem [16,35].

Proposition 2.9 (Corollary 7.3.9 in [40]). *Let A be a matrix corresponding to a log-linear model and \bar{u} be the normalized data, which is strictly positive. Then the maximum likelihood estimate is the unique solution to the equations $A\bar{u} = Ap$ for $p \in \mathcal{M}_A$.*

The iterative proportional scaling, or iterative proportional fitting algorithm is a method for approximating the MLE in multipartition models. If this algorithm converges in finitely many steps, then it reached the MLE directly and in these cases we also refer to it as calculating the MLE. This is a well-known algorithm, first defined in the statistics literature by Deming and Stephan in 1940 in [15] and analyzed further by for example Csizár in [12] and Brown in [7]. There are various types of iterative scaling algorithms. In [16] and [40] the authors describe a variant of the original algorithm, proposed as generalized iterative scaling by Darroch and Ratcliff in [14]. We will focus on the earlier version defined below, because in this case each step is guaranteed to produce a rational function.

We denote by $a_i^{\ell}(\bar{u})$ the product of the i -th row of the ℓ -th partition of a multipartition matrix with the normalized data vector \bar{u} , $a_i^{\ell}(\bar{u}) = a_i^{\ell} \cdot \bar{u}$. As a starting point for the algorithm we choose the uniform distribution $p^0 = (\frac{1}{n}, \dots, \frac{1}{n})^T$ since this vector lies in every multipartition model. The ℓ -th-step of the algorithm is then defined as

$$p_j^{\ell} = p_j^{\ell-1} \cdot \frac{(a_{S(\ell',j)}^{\ell'}(\bar{u}))}{(a_{S(\ell',j)}^{\ell'}(p^{\ell-1}))} \quad (3)$$

for all $j \in [m]$ and $\ell' = \ell \bmod k$. If p^r is the MLE then, by Proposition 2.9, $A\bar{u} = Ap^r$ and the fraction in (3) equals one for $\ell > r$. Thus if p^r is the true MLE, then $p^{\ell} = p^r$ for all $\ell \geq r$.

Note that it follows from (3) and the assumption that u is strictly positive that the entries of p_j^{ℓ} are strictly positive.

Moreover, we note that every step is an information projection to the linear family L^i defined by the i th partition A^i as defined below. Indeed, the **information projection** of p to a non-empty, closed, convex set $M \subset \Delta_{m-1}$ is defined as the element $p' \in M$ that minimizes the KL-divergence between M and p , that is,

$$p' = \arg \min_{q \in M} D(q \parallel p) = \arg \min_{q \in M} \sum_{j=1}^m q_j \log \frac{q_j}{p_j}.$$

The information projection of p to the linear family

$$L^i = \{p \in \Delta_{m-1} \mid A^i p = A^i \bar{u}\} \quad (4)$$

is well-known and given by $p'_j = p_j \cdot \frac{a_{S(i,j)}^i \bar{u}}{a_{S(i,j)}^i p}$. This can be found, for instance, in [13, Lemma 4.1].

Hence the algorithm defined in (3) performs an information projection to the linear family associated to a partition matrix A^i at each step. This is sketched in Fig. 3.

The IPS algorithm converges to the unique point in the intersection between $\overline{\mathcal{M}}_A$, the Euclidean closure of \mathcal{M}_A , and $\mathcal{L}_A = \bigcap_i L_i$ as shown in Theorem 2.8 in [5]. Hence the result of the IPS algorithm minimizes the KL-Divergence with respect to the first argument

$$p^* = \arg \inf_{p \in \mathcal{L}_A} D(p \parallel q) \text{ for all } q \in \overline{\mathcal{M}}_A$$

as well as regarding the second argument

$$p^* = \arg \inf_{q \in \overline{\mathcal{M}}_A} D(p \parallel q) \text{ for all } p \in \mathcal{L}_A.$$

The last minimization is equivalent to maximizing the log-likelihood in case of a discrete model as discussed, for example, in [16].

Note that since the uniform distribution is in $\overline{\mathcal{M}}_A$ minimizing with respect to the first argument results in the maximum entropy estimate. This is interesting because the IPS algorithm can also be used in the context of the maximum entropy method proposed by Jaynes [32].

Since the data vector \bar{u} represents a normalized vector of counts, the only statistically relevant data vectors have non-negative entries. However, from an algebraic perspective, it can be useful to allow for non-positive, or even complex entries. In these cases, one can still find critical points of the complex log-likelihood function. The analysis in the present work still applies to complex data vectors, but in this case, the output of IPS is the point $\mathcal{L}_A \cap \overline{\mathcal{M}}_A$, which need not lie in the probability simplex.

Example 2.10 (The Independence model). A multipartition matrix for the 2x2 independence model is given below on the right where each dot represents a 0. Let $\bar{u} = (\bar{u}_1, \bar{u}_2, \bar{u}_3, \bar{u}_4)^T$ denote a vector of normalized data. Performing the first step of the algorithm results in:

$$p_1^0 = p_2^0 = \frac{1}{4} \cdot \frac{\bar{u}_1 + \bar{u}_2}{\frac{2}{4}} = \frac{1}{2}(\bar{u}_1 + \bar{u}_2),$$

$$p_3^0 = p_4^0 = \frac{1}{4} \cdot \frac{\bar{u}_3 + \bar{u}_4}{\frac{2}{4}} = \frac{1}{2}(\bar{u}_3 + \bar{u}_4)$$

$$A = \begin{pmatrix} 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 \\ 1 & \cdot & 1 & \cdot \\ \cdot & 1 & \cdot & 1 \end{pmatrix}$$

In this case the algorithm converges to the MLE with the second step:

$$p_1^1 = \frac{1}{2}(\bar{u}_1 + \bar{u}_2) \cdot \frac{\bar{u}_1 + \bar{u}_3}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \frac{1}{2}(\bar{u}_3 + \bar{u}_4)} = (\bar{u}_1 + \bar{u}_2)(\bar{u}_1 + \bar{u}_3)$$

$$p_2^1 = \frac{1}{2}(\bar{u}_1 + \bar{u}_2) \cdot \frac{\bar{u}_2 + \bar{u}_4}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \frac{1}{2}(\bar{u}_3 + \bar{u}_4)} = (\bar{u}_1 + \bar{u}_2)(\bar{u}_2 + \bar{u}_4)$$

$$p_3^1 = \frac{1}{2}(\bar{u}_3 + \bar{u}_4) \cdot \frac{\bar{u}_1 + \bar{u}_3}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \frac{1}{2}(\bar{u}_3 + \bar{u}_4)} = (\bar{u}_3 + \bar{u}_4)(\bar{u}_1 + \bar{u}_3)$$

$$p_4^1 = \frac{1}{2}(\bar{u}_3 + \bar{u}_4) \cdot \frac{\bar{u}_2 + \bar{u}_4}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \frac{1}{2}(\bar{u}_3 + \bar{u}_4)} = (\bar{u}_3 + \bar{u}_4)(\bar{u}_2 + \bar{u}_4)$$

Note that at each step of the algorithm the approximation of the MLE is a rational function of the data \bar{u} . Thus if after finitely many steps the algorithm results in the MLE of \bar{u} with respect to \mathcal{M}_A , the model has a rational maximum likelihood estimator. The question naturally arises: for a multipartition model with rational MLE, does iterative proportional scaling always result in the MLE after finitely many steps?

Example 2.11. Here we consider the two matrices A and \tilde{A} depicted below.

$$A = \begin{pmatrix} 1 & 1 & \cdot \\ \cdot & \cdot & 1 \\ 1 & \cdot & \cdot \\ \cdot & 1 & 1 \end{pmatrix}$$

$$\tilde{A} = \begin{pmatrix} 1 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{pmatrix}$$

Both matrices have rowspan equal to \mathbb{R}^3 , so by Proposition 2.2, $\mathcal{M}_A = \mathcal{M}_{\tilde{A}}$. Although they represent the same model, the convergence of the IPS algorithm is heavily influenced by the chosen representation. Let $\bar{u} = (\bar{u}_1, \bar{u}_2, \bar{u}_3)^T$ be normalized data. Using matrix \tilde{A} the IPS algorithm converges in one step to the MLE, $p^* = \bar{u}$.

However, the IPS algorithm on A with normalized data vector \bar{u} results in

$$p^0 = \left(\frac{1}{2}(\bar{u}_1 + \bar{u}_2), \frac{1}{2}(\bar{u}_1 + \bar{u}_2), \bar{u}_3 \right)$$

$$p^1 = \left(\bar{u}_1, \frac{1}{2}(\bar{u}_1 + \bar{u}_2) \frac{(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3}, \bar{u}_3 \frac{(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3} \right)$$

$$p^2 = \left(\bar{u}_1 \frac{(\bar{u}_1 + \bar{u}_2)}{\bar{u}_1 + \frac{\frac{1}{2}(\bar{u}_1 + \bar{u}_2)(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3}}, \frac{1}{2}(\bar{u}_1 + \bar{u}_2) \frac{(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3} \frac{(\bar{u}_1 + \bar{u}_2)}{\bar{u}_1 + \frac{\frac{1}{2}(\bar{u}_1 + \bar{u}_2)(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3}}, \bar{u}_3 \frac{(\bar{u}_2 + \bar{u}_3)}{\frac{1}{2}(\bar{u}_1 + \bar{u}_2) + \bar{u}_3} \right).$$

In general the different projections have the following form:

$$p^k = \left(\bar{u}_1, \quad p_2^{k-1} \frac{(\bar{u}_2 + \bar{u}_3)}{p_2^{k-1} + \bar{u}_3}, \bar{u}_3 \frac{(\bar{u}_2 + \bar{u}_3)}{p_2^{k-1} + \bar{u}_3} \right), \quad k \text{ odd} \quad (5)$$

$$p^k = \left(\bar{u}_1 \frac{(\bar{u}_1 + \bar{u}_2)}{\bar{u}_1 + p_2^{k-1}}, p_2^{k-1} \frac{(\bar{u}_1 + \bar{u}_2)}{\bar{u}_1 + p_2^{k-1}}, \bar{u}_3 \right), \quad k \text{ even} \quad (6)$$

Suppose that there exists an index k such that the second entry of p^k is exactly \bar{u}_2 in (5) as well as (6). Then we deduce that $\bar{u}_2 = p_2^{k-1}$. Hence IPS can only result in the exact result if $\bar{u}_1 = \bar{u}_2$. Note that while the model is symmetric, the application of the IPS is not. Hence if we would reorder the partitions, then exact convergence would require $\bar{u}_2 = \bar{u}_3$.

Additionally, we generated 500000 integer vectors of counts with entries between 0 and 200000 with the pseudo-random number generator from NumPy. After normalization we used them as input distributions \bar{u} for the IPS with respect to the matrix A . The arithmetic mean of the iteration steps taken to get a step size smaller than 10^{-8} was approx 84 with a minimum value of 4, a maximum value of 315848 and a standard error of approx. 1.73. Recall that in case of \bar{A} the necessary number of iteration steps is exactly 1. This example demonstrates that the chosen matrix representation of a multipartition model can greatly influence its behavior under the IPS algorithm.

2.3. Toric fiber products

In this section, we review an algebraic construction, known as the *toric fiber product*, which will facilitate the proofs in the following sections. Toric fiber products were first introduced by Sullivant, and this exposition follows the notation of [39]. For the purposes of this article, we may broadly think of the toric fiber product as describing a way to concatenate the parametrizations of two log-linear models that preserves some of their properties.

For any positive integer r , let $[r] := \{1, \dots, r\}$. Fix positive integers r and s_i, t_i for each $i \in [r]$. Define three polynomial rings,

$$\mathbb{C}[x] = \mathbb{C}[x_j^i \mid i \in [r], j \in [s_i]],$$

$$\mathbb{C}[y] = \mathbb{C}[y_k^i \mid i \in [r], k \in [t_i]], \text{ and}$$

$$\mathbb{C}[z] = \mathbb{C}[z_{jk}^i \mid i \in [r], j \in [s_i], k \in [t_i]].$$

Recall that a **multigrading** on a polynomial ring $\mathbb{C}[t] := \mathbb{C}[t_1, \dots, t_d]$ is an assignment of a multidegree vector to each monomial, $\deg(t^\mu) \in \mathbb{Z}^s$ that satisfies $\deg(t^\mu t^\nu) = \deg(t^\mu) + \deg(t^\nu)$. An ideal $I \in \mathbb{C}[t]$ is homogeneous with respect to this multigrading if it has a generating set $\{f_1, \dots, f_s\}$ such that for each i , every term of f_i has the same multidegree.

Fix multigradings on $\mathbb{C}[x]$ and $\mathbb{C}[y]$ given by $\deg(x_j^i) = \deg(y_k^i) = d^i$ for integer vectors d^i . Note that the multigrading of each indeterminate depends only on its superscript, i . Let D be the matrix with columns d^i for $i \in [r]$.

Let $I \subset \mathbb{C}[x]$ and $J \subset \mathbb{C}[y]$ be homogeneous ideals with respect to this multigrading. We define a ring homomorphism,

$$\begin{aligned} \psi_{I,J} : \mathbb{C}[z] &\rightarrow (\mathbb{C}[x]/I) \otimes_{\mathbb{C}} (\mathbb{C}[y]/J) \\ z_{jk}^i &\mapsto x_j^i \otimes_{\mathbb{C}} y_k^i \end{aligned}$$

Definition 2.12. The **toric fiber product** of I and J with multigrading D is the ideal $I \times_D J = \ker(\psi_{I,J})$ in $\mathbb{C}[z]$.

We assume throughout that I and J are toric, as the ideals associated to rational multipartition models are always toric. We say that a polynomial map θ between polynomial rings is a *parametrization* of an ideal I if $I = \ker \theta$, or equivalently, if θ is the pullback of a parametrization of the algebraic variety $V(I)$. In order to parametrize the toric fiber product of I and J , we may simply take the product of their parametrizations. More precisely, if ϕ_I and ϕ_J are monomial parametrizations of I and J , then we obtain a parametrization for their toric fiber product according to multigrading D via $z_{jk}^i \mapsto \phi_I(x_j^i) \phi_J(y_k^i)$. In the case where the columns of D are linearly independent, there is an explicit description of a generating set for $I \times_D J$ using the generators of I and J . For more details on this construction, we refer the reader to Section 2 of [39].

When I and J are toric, the toric fiber product construction can also be defined using the parametrizing matrices [21, Section 2.3]. Let $I = I(B)$ and $J = I(C)$ for some integer matrices B and C . Then the multigrading given by D assigns a vector in D to each column of B and C . The matrix parametrizing their toric fiber product is obtained by concatenating each pair of columns of B and C with the same degree vector.

Example 2.13. Consider the matrices

$$B = \left[\begin{array}{cc|cc} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{array} \right] \quad \text{and} \quad C = \left[\begin{array}{cc|cc} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{array} \right]$$

The columns of B correspond in order to variables $x_1^1, x_2^1, x_1^2, x_2^2$ and the columns of C correspond in order to variables $y_1^1, y_2^1, y_1^2, y_2^2$. Let I be the vanishing ideal of the log-linear model specified by B and let J be that of C . We have $I = \langle x_1^1 - x_2^1 \rangle$ and $J = \langle y_1^1 y_2^1 - y_1^2 y_2^2 \rangle$. Consider the multigrading with respect to the 2×2 identity matrix D , where $\deg(x_j^1) = \deg(y_j^1) = e_1$ and $\deg(x_j^2) = \deg(y_j^2) = e_2$ for $j = 1, 2$. Observe that both I and J are homogeneous with respect to this multigrading. The matrix defining the parametrization of their toric fiber product, $I \times_D J$, is

$$\begin{array}{c|cccc} z_{11}^1 & z_{12}^1 & z_{21}^1 & z_{22}^1 & z_{11}^2 & z_{12}^2 & z_{21}^2 & z_{22}^2 \\ \hline 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ \hline 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{array}.$$

We conclude this introduction to the toric fiber product by describing the MLE of the model associated to the toric fiber product in terms of the MLE of its factors. Let B, C be matrices representing the log-linear models associated to I and J . Let p_B^*, p_C^* and p_D^* denote the maximum likelihood estimators for the log-linear models given by B, C and D , respectively. The following is Theorem 5.5 of [3] and plays a critical role in our proofs of the results in Section 3.

Theorem 2.14 ([3]). *If D has linearly independent columns, then the (i, j, k) th coordinate function of the maximum likelihood estimator for the log-linear model associated to the toric fiber product $I \times_D J$ is*

$$\frac{(p_B^*)_j (p_C^*)_k}{(p_D^*)_i}.$$

2.4. Hierarchical models

A hierarchical model is a log-linear model for which the independence structure between random variables is described by a simplicial complex. We will introduce hierarchical models using the notation in Chapter 9 in [40]. They comprise another interesting class of multipartition models.

Let $2^{[n]}$ be the powerset of $[n]$. A **simplicial complex** Γ with the ground set $[n]$ is a subset of $2^{[n]}$ with the property that if $F \in \Gamma$ and $F' \subseteq F$, then $F' \in \Gamma$. The elements of a simplicial complex are the **faces** of Γ , the inclusion-maximal faces of Γ are called **facets**, and the set of all facets of Γ is denoted $\text{facet}(\Gamma)$. A simplicial complex is completely identified by its set of facets. So we may, for example, write $\Gamma = [12][13]$ for the simplicial complex $\Gamma = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}\}$.

Up to this point we did not assume any additional structure on Δ_{m-1} , but now we will consider a collection of random variables X_1, \dots, X_ℓ with discrete state spaces R_1, \dots, R_ℓ . A simplicial complex with the ground set $[\ell]$ describes the interactions among the different random variables on the joint state space $R = R_1 \times \dots \times R_\ell$. The corresponding log-linear model is called hierarchical model and defined as follows. A description of the partition matrices that define a hierarchical model can be found in Section 4.1.

Definition 2.15. Let Γ be a simplicial complex. For every facet $F = \{f_1, f_2, \dots\}$ in $\text{facet}(\Gamma)$, let $r_F = \prod_{f \in F} |R_f|$. For $i = (i_1, \dots, i_\ell) \in R$ let $i_F = (i_{f_1}, i_{f_2}, \dots)$. Additionally, we introduce a set of r_F positive parameters $\theta_{i_F}^F$ for every facet. Now the **hierarchical model** corresponding to Γ is

$$\mathcal{M}_\Gamma = \left\{ p \in \Delta_{|R|-1} \mid p_i = \frac{1}{Z(\theta)} \prod_{F \in \text{facet}(\Gamma)} \theta_{i_F}^F \right\}$$

with the normalizing factor

$$Z(\theta) = \sum_{i \in R} \prod_{F \in \text{facet}(\Gamma)} \theta_{i_F}^F.$$

Consider the binary case where each random variable X_1, \dots, X_ℓ has exactly two states. We define the matrix A_Γ that realizes the binary hierarchical model on Γ as a log-linear model in the case where each random variable has exactly two states. The columns of the matrix A_Γ are indexed by subsets of $[n]$. We could equivalently index the columns by 0/1 strings of length n by taking each $S \subset [n]$ to be the set of positions in the string equal to 1.

Let F_1, \dots, F_k be the facets of Γ . The rows of A_Γ are divided into blocks A^1, \dots, A^k each with $2^{|F_i|}$ rows. The rows of the partition matrix A^i are of the form a_S^i where $S \subset F_i$ with

$$a_S^i(T) = \begin{cases} 1, & \text{if } T \cap F_i = S \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$




Fig. 4. Three different examples of hierarchical models.

for each $T \subset [n]$. The **hierarchical model** \mathcal{M}_Γ is the closure of the image of the monomial map defined by A_Γ intersected with the probability simplex, Δ_{2^n} .

Definition 2.16. A simplicial complex Γ is called **decomposable** if it is a simplex or if it is the union of two decomposable simplicial complexes Γ_1, Γ_2 such that there exist $F_1 \in \text{facet}(\Gamma_1)$ and $F_2 \in \text{facet}(\Gamma_2)$ with

$$\bigcup_{F \in \text{facet}(\Gamma_1)} F \cap \bigcup_{\tilde{F} \in \text{facet}(\Gamma_2)} \tilde{F} = F_1 \cap F_2.$$

Example 2.17.

- The 2x2 independence model in Example 2.10 is a hierarchical model on two disjoint vertices where each random variable has state space $\{1, 2\}$. This is depicted in Fig. 4 (1). This simplicial complex is decomposable.
- The decomposable simplicial complex depicted in Fig. 4 (2) is $\Gamma = [1, 2, 3][3, 4, 5]$. The facets of Γ are $\{1, 2, 3\}$ and $\{3, 4, 5\}$.
- In the simplicial complex $\Gamma = [12][13][23]$, the set $\{1, 2, 3\}$ is not a face, hence the associated graph in Fig. 4 (3) lacks the grayed space between the nodes. This simplicial complex is not decomposable.

The simplicial complex in Example 2.17 (2) is decomposable; we can see this by letting Γ_1 have the unique maximal face $\{1, 2, 3\}$ and Γ_2 have the unique maximal face $\{3, 4, 5\}$. The example in 2.17 (3) is not decomposable.

Now we will introduce a special ordering to the facets of a simplicial complex.

Definition 2.18. Let $E = \{F_1, \dots, F_s\}$ be an ordering of the facets of a simplicial complex Γ . Then this ordering satisfies the **running intersection property (RIP)**, if for each $r \in [s]$ there exists a k_r such that

$$\left(\bigcup_{k=1}^r F_k \right) \cap F_{r+1} = F_{k_r} \cap F_{r+1}$$

The connection between decomposable simplicial complexes and the RIP, as stated in the lemma below, is proven for example as Lemma 5.10 in [26].

Lemma 2.19. A simplicial complex Γ is decomposable if and only if there exists an ordering of the facets of Γ that satisfies the RIP.

The following is Theorem 5.3 in [26]. It draws a connection between decomposability of a simplicial complex and one-cycle convergence of the IPS algorithm for the corresponding hierarchical model. One-cycle convergence can be achieved by requiring that the multipartition matrix is ordered so that the facets corresponding to the partition matrices satisfy the RIP. An additional condition for one-cycle convergence was proven in [41].

Theorem 2.20 ([26]). Suppose the simplicial complex Γ is decomposable. Let $A = A^{1, \dots, k}$ be the multipartition matrix corresponding to an ordering of the facets of Γ that satisfies the RIP. Then the iterative proportional scaling algorithm applied to A yields the maximum likelihood estimate in one cycle.

The hierarchical model associated to a non-decomposable graphical model has no parametrization for which IPS converges in one cycle. In particular, a hierarchical model has rational MLE if and only if it is decomposable. This fact follows from [22, Theorem 4.4] and [11, Theorem 3.2]. In the following section, we define a generalization of Theorem 2.20 that can be applied to a broader range of multipartition models.

3. Generalized running intersection property

In this section, we define the generalized running intersection property, or GRIP. We show that a multipartition model satisfies the GRIP if and only if it can be written as an iterated toric fiber product of multipartition models. This allows us to give a formula for the maximum likelihood estimate of a model that satisfies the GRIP. We use this formula to show that iterative proportional scaling exhibits one-cycle convergence on models that satisfy the GRIP.




Fig. 5. The matrix A with its connection ratios shown on the right.

Definition 3.1. Let B and C be two partition matrices with the same number of columns and with rows a_μ^B and $a_{\mu'}^C$. Two rows a_μ^B and $a_{\mu'}^C$ are **connected** if their supports intersect nontrivially; that is, if $I_\mu^B \cap I_{\mu'}^C \neq \emptyset$.

For a matrix to satisfy the GRIP, we require that rows from one partition to the next are connected in a way that satisfies certain conditions. The following definition is motivated by the way column weights (see Definition 2.7) appear in the expressions produced by each step of the IPS algorithm.

Definition 3.2. For a multipartition matrix $A = A^{1 \dots k}$, define c_j^ℓ to be j -th column weight for the matrix $A^{1 \dots \ell}$ obtained by only considering the first ℓ partitions of A . Then A is **well-connected** if for any row vector a_μ^ℓ with $\ell > 1$, we have that $(c_j^\ell / c_j^{\ell-1}) = (c_{j'}^\ell / c_{j'}^{\ell-1})$ for all $j, j' \in I_\mu^\ell$. We call this quantity the **connection ratio** for a_μ^ℓ and denote this quantity as C_μ^ℓ with the convention that $C_\mu^1 = |I_\mu^1|$.

Remark 3.3. Suppose that A is a well-connected matrix. Then for any index $j \in [m]$, we see that the column weight c_j can be expressed as the product of the connection ratios, $c_j = \prod_{\ell=1}^k C_{S(\ell, j)}^\ell$.

Example 3.4. Consider the matrix A in Fig. 5. We note that it is well-connected and display the connection ratios of each row to the right of the matrix.

Let A^1, \dots, A^ℓ be a set of partition matrices with m columns. Let β be the number of *distinct* columns of $A^{1 \dots \ell}$. Define a labeling of the columns of $A^{1 \dots \ell}$, $\lambda: [m] \rightarrow [\beta]$ such that $\lambda(j) = \lambda(j')$ if and only if the j th and j' th columns of $A^{1 \dots \ell}$ are equal.

Definition 3.5. We define the partition matrix $B = \Psi_{n=1}^\ell A^n$ to be the $\beta \times m$ matrix with j th column $e_{\lambda(j)}$. Since the labeling λ of the columns of $A^{1 \dots \ell}$ simply permutes the rows of $\Psi_{n=1}^\ell A^n$, we omit the specification of λ from this notation.

In order to define the GRIP, we require that the tree associated to the multipartition matrix be staged. The word “floret” in the definition below is deliberately suggestive of the terminology for staged trees discussed in Appendix A.

Definition 3.6. Let B and C be two partition matrices with the same number of columns and with rows a_μ^B and a_ν^C . The matrices B and C satisfy the **floret condition** if for every two rows of B , a_μ^B and $a_{\mu'}^B$, the sets of rows of C that are connected to a_μ^B and $a_{\mu'}^B$ are disjoint or equal. In this case, the set of rows of B connected to a row a_ν^C is called a **floret** of B and the set of rows of C connected to a row a_μ^B is a **floret** of C .

If the pair of matrices B and C satisfy the floret condition, we also say that the matrix BC satisfies the floret condition. Suppose that B and C are two partition matrices that satisfy the floret condition, and let F_1^C, \dots, F_f^C be the distinct florets of C with respect to the matrix B . For each $i \in [f]$, let F_i^B be the set of all rows of B that are connected to the rows of F_i^C .

Remark 3.7. The floret condition in Definition 3.6 is symmetric in B and C ; that is, BC satisfies the floret condition if and only if CB does. Hence the floret condition implies that F_i^B is well-defined for any $i \in [f]$ and that every floret in B is of the form F_i^B for some i . Indeed the florets of B are in one-to-one correspondence with the florets of C in this way.

Let $t(\mu, \nu)$ be the index in $[f]$ such that $a_\mu^B \in F_{t(\mu, \nu)}^B$ and $a_\nu^C \in F_{t(\mu, \nu)}^C$. In fact, $t(\mu, \nu)$ is determined by the row a_μ^B or the row a_ν^C ; it is not necessary to provide both indices. Hence, we write $t(\mu, \bullet)$ to be the index such that $a_\mu^B \in F_{t(\mu, \bullet)}^B$. Similarly, $t(\bullet, \nu)$ is the index such that $a_\nu^C \in F_{t(\bullet, \nu)}^C$. Thus a_μ^B and a_ν^C are connected if and only if $t(\mu, \bullet) = t(\bullet, \nu)$. Note that the function $t(\mu, \nu)$ implicitly depends on the matrices B and C .




Fig. 6. The matrices $A^1 \mathbin{\lrcorner} A^2$, $A^1 \mathbin{\lrcorner} A^2$ and $(A^1 \mathbin{\lrcorner} A^2) \mathbin{\lrcorner} A^3$ on the left and the tree associated to A on the right.

Example 3.8. Consider the matrix in Example 3.4, and let $B = A^1 \mathbin{\lrcorner} A^2$, which results in the following matrix:

$$B = A^1 \mathbin{\lrcorner} A^2 = \begin{pmatrix} 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 \end{pmatrix}.$$

Now let $C = A^3$. Note that B and C satisfy the floret condition and that the florets of C are given by $\mathcal{F}_1^C = \{a_1^C, a_2^C, a_3^C\}$ and $\mathcal{F}_2^C = \{a_4^C, a_5^C\}$. We see that $t(\cdot, 1) = t(\cdot, 2) = t(\cdot, 3) = 1$ and $t(\cdot, 4) = t(\cdot, 5) = 2$. The corresponding florets of B are given by $\mathcal{F}_1^B = \{a_1^B, a_3^B\}$ and $\mathcal{F}_2^B = \{a_2^B, a_4^B\}$. This implies that $t(1, \cdot) = t(3, \cdot) = 1$ and $t(2, \cdot) = t(4, \cdot) = 2$. At this point the function $t(\mu, \nu)$ is completely defined and it can be used to answer questions like, “Which floret of C is connected to the row a_3^B ?” This is given by $\mathcal{F}_{t(3, \cdot)}^C = \mathcal{F}_1^C$.

For a set of partition matrices A^1, \dots, A^ℓ with m columns, let $B = \mathbb{W}_{n=1}^\ell A^n$. Let $C = A^{\ell+1}$ have γ rows and m columns. Then the matrix BC has m columns of the form $[e_\mu e_\nu]^T$ where $\mu \in [\beta]$ and $\nu \in [\gamma]$. Let $\omega_{\mu\nu}$ denote the column weight of $[e_\mu e_\nu]^T$ in BC . Then the columns of BC are indexed by triples of the form (μ, ν, s) for $s \in \{0, \dots, \omega_{\mu\nu} - 1\}$ where column (μ, ν, s) is the $(s+1)$ st column of BC of the form $[e_\mu e_\nu]^T$. We adopt this column indexing scheme for the remainder of this section.

Definition 3.9. Suppose that B and C satisfy the floret condition with f florets, $\mathcal{F}_1^C, \dots, \mathcal{F}_f^C$. We define the matrix $B \mathbin{\lrcorner} C$ to be the $f \times m$ matrix whose columns are indexed by the triples (μ, ν, s) such that the (μ, ν, s) entry of the r th row of $B \mathbin{\lrcorner} C$ is equal to 1 if $t = t(\mu, \nu)$ and 0 otherwise. Note that any two columns with indices (μ, ν, s) and (μ, ν, s') are identical.

In other words, a column of $B \mathbin{\lrcorner} C$ are indicator vectors for the floret that the corresponding column's nonzero rows belong to in BC . Note that $B \mathbin{\lrcorner} C$ is only defined when B and C satisfy the floret condition.

Definition 3.10. Let A^1, \dots, A^k be partition matrices. For each ℓ , let B_ℓ denote $\mathbb{W}_{n=1}^\ell A^n$. Then the multipartition matrix $A^{1, \dots, k}$ satisfies the **generalized running intersection property**, or **GRIP** if for each $\ell \in [k-1]$,

1. the matrix $B_\ell A^{\ell+1}$ is well-connected,
2. $B_\ell A^{\ell+1}$ satisfies the floret condition, and
3. the rows of $B_\ell \mathbin{\lrcorner} A^{\ell+1}$ lie in the rowspan of $A^{1, \dots, \ell}$.

Remark 3.11. The first point in the above definition is true for each $\ell \in [k-1]$ if and only if $A^{1, \dots, k}$ is well-connected. Additionally it is easy to see that the second point is satisfied if and only if $\mathcal{T}_{A^{1, \dots, k}}$ is a staged tree (see Appendix A). We frame these conditions above in an iterative way to facilitate the proofs in this section.

Example 3.12. We show that the matrix A from Example 3.4 satisfies the GRIP, as depicted in Figure 6. We also illustrate the connection between the GRIP and staged trees, as discussed in Appendix A. In Example 3.4, we saw that A is well-connected and in Example 3.8, we saw that it satisfies the floret condition. Equivalently, in order to check that A satisfies the floret condition, it suffices to check that its associated tree \mathcal{T}_A , depicted below, is a staged tree. This is the content of Example A.4. Note that both when $B = A^1$ and $C = A^2$ and when $B = A^{1,2}$ and $C = A^3$, it is the case that the florets of the matrix C correspond exactly to the second-level florets of the tree associated with the multipartition matrix BC . This is true in general; BC satisfies the floret condition if and only if \mathcal{T}_{BC} is a staged.

Finally, we check the third condition of the GRIP. For $\ell = 1$, $B_1 = A^1$ and $B_1 \mathbin{\lrcorner} A^2 = A^2$, so (3) is trivially satisfied. Thus all that is left is to check that the rows of $B^2 \mathbin{\lrcorner} A^3$ lie in the rowspan of A . First consider $B^2 = A^1 \mathbin{\lrcorner} A^2$. The rows of this matrix are indexed by the distinct columns of the multipartition matrix $A^{1,2}$ and thus, all the distinct paths in $\mathcal{T}_{A^{1,2}}$.

Note the rows of the matrix $B^2 \mathbin{\lrcorner} A^3 = (A^1 \mathbin{\lrcorner} A^2) \mathbin{\lrcorner} A^3$ index the level three florets of the staged tree \mathcal{T}_A . Since $B^2 \mathbin{\lrcorner} A^3 = A^2$, we easily see that (3) is satisfied.

We now describe the iterated toric fiber product structure on the multipartition model defined by a matrix $A^{1,\dots,k}$ that satisfies the GRIP. The following lemma is crucial to our construction of the toric fiber product.

Lemma 3.13. *Suppose that B and C are partition matrices such that BC is well-connected and satisfies the floret condition. Then there exist positive integers x_1, \dots, x_β and y_1, \dots, y_γ such that whenever $\omega_{\mu\nu}$ is nonzero, we have $\omega_{\mu\nu} = x_\mu y_\nu$ for all $\mu \in [\beta]$ and $\nu \in [\gamma]$.*

Proof. By Remark 3.3, for all $\mu \in [\beta], \nu \in [\gamma]$ such that $\iota(\mu, \bullet) = \iota(\bullet, \nu)$, we have $\omega_{\mu\nu} = C_\mu^1 C_\nu^2$ where C_μ^1 is an integer and C_ν^2 is a rational number. Let $C_\nu^2 = a_\nu/b_\nu$ for positive coprime integers a_ν and b_ν . Since $(C_\mu^1 a_\nu)/b_\nu$ is an integer for all μ and ν , we have that b_ν divides C_μ^1 for all μ and ν . Let $x_\mu = C_\mu^1 / \text{lcm}(b_h)_{h \in [\gamma]}$ and let $y_\nu = (a_\nu \text{lcm}(b_h)_{h \in [\gamma]})/b_\nu$. Each of these is an integer and satisfies $x_\mu y_\nu = C_\mu^1 C_\nu^2 = \omega_{\mu\nu}$. Applying this result to each of the florets yields the desired integers. \square

We assume throughout that $A^{1,\dots,k}$ satisfies the GRIP. Consider the submatrix $A^{1,\dots,\ell+1}$, where $\ell \in [k-1]$. Fix $B = \bigcup_{n=1}^\ell A^n$ and $C = A^{\ell+1}$ as in the paragraph preceding Definition 3.10.

Define the polynomial ring,

$$R = \mathbb{C}[p_{\mu,\nu,s}^{(\mu,\nu)} \mid \mu \in [\beta], \nu \in [\gamma], s \in \{0, \dots, \omega_{\mu\nu} - 1\}].$$

Let $x_1, \dots, x_\beta, y_1, \dots, y_\gamma$ be positive integers such that $\omega_{\mu\nu} = x_\mu y_\nu$; these are guaranteed to exist by Lemma 3.13. Define two other polynomial rings,

$$R_B = \mathbb{C}[q_{\mu,s'}^{(\mu,\bullet)} \mid \mu \in [\beta], s' \in \{0, \dots, x_\mu - 1\}], \text{ and}$$

$$R_C = \mathbb{C}[r_{\nu,s''}^{(\bullet,\nu)} \mid \nu \in [\gamma], s'' \in \{0, \dots, y_\nu - 1\}].$$

Recall that λ maps the columns of $A^{1,\dots,\ell}$ onto $[\beta]$ by labeling the columns of $A^{1,\dots,\ell}$ so that $\lambda(\mu) = \lambda(\nu)$ if and only if the μ th and ν th columns of $A^{1,\dots,\ell}$ are equal. Let $\underline{A}^{1,\dots,\ell}$ be the matrix defined as follows. Its columns are indexed by ordered pairs (μ, s') such that $\mu \in [\beta]$ and $s' \in \{0, \dots, x_\mu - 1\}$. The (μ, s') column of $\underline{A}^{1,\dots,\ell}$ is equal to the columns of $A^{1,\dots,\ell}$ with label μ . In other words, the matrix $\underline{A}^{1,\dots,\ell}$ has the same set of columns as $A^{1,\dots,\ell}$; however, the column with label μ is repeated in $\underline{A}^{1,\dots,\ell}$ only x_μ times. Hence, the column weights of \underline{A} are bounded by the column weights of A . In general, we say that a matrix \underline{A} is a **compression** of A if A and \underline{A} have the same set of columns, and \underline{A} contains at most as many copies of each column as A .

Similarly, let $\underline{A}^{\ell+1}$ have columns indexed by ordered pairs (ν, s'') such that $\nu \in [\gamma]$ and $s'' \in \{0, \dots, y_\nu - 1\}$. The (ν, s'') column of $\underline{A}^{\ell+1}$ is equal to the ν th standard basis vector. The matrix $\underline{A}^{\ell+1}$ has the same underlying set of columns as $A^{\ell+1}$, and this underlying set is equal to $\{e_1, \dots, e_\gamma\}$; however, the ν th standard basis vector is repeated in $\underline{A}^{\ell+1}$ only y_ν times.

Recall that for any integer matrix A , $I(A)$ is the vanishing ideal of the log-linear model \mathcal{M}_A as described in Proposition 2.2. We define a ring homomorphism by

$$\begin{aligned} \psi : R &\rightarrow R_B/I(\underline{A}^{1,\dots,\ell}) \otimes_{\mathbb{C}} R_C/I(\underline{A}^{\ell+1}) \\ p_{\mu,\nu,s}^{(\mu,\nu)} &\mapsto q_{\mu,s'}^{(\mu,\bullet)} \otimes_{\mathbb{C}} r_{\nu,s''}^{(\bullet,\nu)}, \end{aligned} \tag{8}$$

where s', s'' are the nonnegative integers such that $s = s' y_\nu + s''$. Note that each $s \in \{0, \dots, \omega_{\mu\nu} - 1\}$ has a unique representation in this way. Define a multigrading on R_B by $\deg(q_{\mu,s'}^{(\mu,\bullet)}) = e_{\iota(\mu,\bullet)}$ and a multigrading on R_C by $\deg(r_{\nu,s''}^{(\bullet,\nu)}) = e_{\iota(\bullet,\nu)}$. Since each indeterminate $q_{\mu,s'}^{(\mu,\bullet)}$ corresponds to the column of $\underline{A}^{1,\dots,\ell}$ with index (μ, s') , we also say that this column has degree $e_{\iota(\mu,\bullet)}$. Similarly, we say that the (ν, s'') column of $\underline{A}^{\ell+1}$ has degree $e_{\iota(\bullet,\nu)}$. We will show that $\ker(\psi)$ is the toric fiber product of $I(\underline{A}^{1,\dots,\ell})$ and $I(\underline{A}^{\ell+1})$ with respect to this multigrading and is equal to $I(A^{1,\dots,\ell+1})$.

Proposition 3.14. *Let $\underline{A}^{1,\dots,\ell}$ and $\underline{A}^{\ell+1}$ be any compressions of $A^{1,\dots,\ell}$ and $A^{\ell+1}$, respectively. The toric ideals $I(\underline{A}^{1,\dots,\ell})$ and $I(\underline{A}^{\ell+1})$ are multihomogeneous with respect to the multigradings $\deg(q_{\mu,s'}^{(\mu,\bullet)}) = e_{\iota(\mu,\bullet)}$ and $\deg(r_{\nu,s''}^{(\bullet,\nu)}) = e_{\iota(\bullet,\nu)}$.*

Proof. Let $D = B \# C$. The rows of D are in the rowspan of $A^{1,\dots,\ell}$ by the generalized running intersection property. Moreover, the row of D corresponding to a floret F in C is $\sum_{a_v^C \in F} a_v^C$. Hence, the rows of D are in the rowspan of $A^{\ell+1}$ as well.

The matrix $\underline{A}^{1,\dots,\ell}$ is obtained by deleting some columns of $A^{1,\dots,\ell}$. Let \bar{D} be obtained from $D = B \# C$ by deleting the same columns. Then the matrix \bar{D} defines the multigrading $\deg(q_{\mu,s'}^{(\mu,\bullet)}) = e_{\iota(\mu,\bullet)}$ on R_B . The rows of \bar{D} remain in the rowspan of $\underline{A}^{1,\dots,\ell}$. Thus $I(\underline{A}^{1,\dots,\ell})$ is multihomogeneous with respect to the grading given by \bar{D} since $\ker(\underline{A}^{1,\dots,\ell}) \subset \ker(\bar{D})$. The argument for the linear ideal $I(\underline{A}^{\ell+1})$ is analogous. \square

Theorem 3.15. *Let $A^{1,\dots,\ell+1}$ be a multipartition matrix that satisfies the generalized running intersection property. Then $I(A^{1,\dots,\ell+1})$ is a toric fiber product of toric ideals. In particular this toric fiber product is given by $\ker(\psi)$, defined in (8).*

Proof. By Proposition 3.14, $I(\underline{A}^{1,\dots,\ell})$ and $I(\underline{A}^{\ell+1})$ are both multihomogeneous with respect to the grading specified by D . Consider the (μ, ν, s) column of $A^{1,\dots,\ell+1}$. Since this is an index for a column of $A^{1,\dots,\ell+1}$, we must have $\iota(\mu, \bullet) = \iota(\bullet, \nu)$. Let x_μ and y_ν be the positive integers given by Lemma 3.13 and let s', s'' be the unique nonnegative integers such $s = s'y_\nu + s''$. We obtain the (μ, ν, s) column of $A^{1,\dots,\ell+1}$ by concatenating the (μ, s') column of $\underline{A}^{1,\dots,\ell}$ and the (ν, s'') column of $\underline{A}^{\ell+1}$.

Note that for each $\iota \in [f]$, we obtain a unique column of $A^{1,\dots,\ell}$ with degree e_ι for every pair of columns of $\underline{A}^{1,\dots,\ell}$ and $\underline{A}^{\ell+1}$ with degree e_ι . Moreover, this process only concatenates a pair of columns of $\underline{A}^{1,\dots,\ell}$ and $\underline{A}^{\ell+1}$ if they have the same degree. Hence $A^{1,\dots,\ell+1}$ is matrix that parametrizes the toric fiber product $I(\underline{A}^{1,\dots,\ell}) \times_D I(\underline{A}^{\ell+1})$. \square

We are now ready to examine the maximum likelihood estimators of multipartition matrices that satisfy the GRIP. Recall that the number of copies of column e_μ in B is

$$\sum_{\nu: a_\nu^C \in F_{I(\mu, \bullet)}^C} \omega_{\mu\nu} = x_\mu \sum_{\nu: a_\nu^C \in F_{I(\mu, \bullet)}^C} y_\nu.$$

For each μ , let

$$Y_\mu = \sum_{\nu: a_\nu^C \in F_{I(\mu, \bullet)}^C} y_\nu.$$

Then we index columns of $A^{1,\dots,\ell}$ by (μ, s) where $\mu \in [\beta]$ and $s \in \{0, \dots, x_\mu Y_\mu - 1\}$. Let \bar{u} be an m -dimensional normalized data vector for $A^{1,\dots,\ell}$ with this indexing. We form an associated data vector \tilde{u} for $\underline{A}^{1,\dots,\ell}$ by setting

$$\tilde{u}_{(\mu, s')} = \sum_{n=0}^{Y_\mu-1} \bar{u}_{(\mu, s'Y_\mu+n)} \quad (9)$$

for each $\mu \in [\beta]$ and $s' \in \{0, \dots, x_\mu - 1\}$. When we compress the matrix $A^{1,\dots,\ell}$ to $\underline{A}^{1,\dots,\ell}$, we have a single column with label (μ, s') representing Y_μ copies of an identical column in $A^{1,\dots,\ell}$. In order to form the (μ, s') coordinate of \tilde{u} , we add the Y_μ coordinates of \bar{u} that are represented by column (μ, s') in $\underline{A}^{1,\dots,\ell}$. Let $p_{(\mu, s)}^*(\bar{u})$ denote the (μ, s) component of the MLE for \bar{u} in $\mathcal{M}(A^{1,\dots,\ell})$.

Proposition 3.16. *Let \bar{u} be a data vector for $\mathcal{M}(A^{1,\dots,\ell})$ and let \tilde{u} be as defined in Equation (9). Then the MLE for \tilde{u} in $\mathcal{M}(\underline{A}^{1,\dots,\ell})$ has (μ, s') component $Y_\mu \cdot p_{(\mu, 0)}^*(\bar{u})$, which is independent of the value of s' for each $s' \in \{0, \dots, x_\mu - 1\}$.*

Proof. Let q^* be the $(\sum x_\mu)$ -dimensional vector with (μ, s') component equal to $Y_\mu \cdot p_{(\mu, 0)}^*(\bar{u})$ for each $\mu \in [\beta]$ and $s' \in \{0, \dots, x_\mu - 1\}$. We show that q^* satisfies the conditions of Proposition 2.9.

First, we show that $\underline{A}^{1,\dots,\ell} q^* = \underline{A}^{1,\dots,\ell} \tilde{u}$. Let \underline{a} be a row of $\underline{A}^{1,\dots,\ell}$ and let a be the corresponding row in $A^{1,\dots,\ell}$. Then we have

$$\underline{a} \cdot q^* = \sum_{(\mu, s') \in \text{supp}(\underline{a})} Y_\mu \cdot p_{(\mu, 0)}^*(\bar{u}).$$

Note that if $(\mu, s') \in \text{supp}(\underline{a})$ for some s' , then we have $(\mu, h) \in \text{supp}(\underline{a})$ for all h in $\{0, \dots, x_\mu - 1\}$. Moreover, by Proposition 2.9 applied to $\mathcal{M}(A^{1,\dots,\ell})$, $p_{(\mu, s)}^* = p_{(\mu, 0)}^*$ for all $s \in \{0, \dots, x_\mu Y_\mu - 1\}$. Indeed, since the $(\mu, 0)$ and (μ, s) columns of $A^{1,\dots,\ell}$ are equal, any point in $\mathcal{M}(A^{1,\dots,\ell})$ must have equal $(\mu, 0)$ and (μ, s) coordinates. Thus

$$\underline{a} \cdot q^* = \sum_{\mu: (\mu, 0) \in \text{supp}(\underline{a})} x_\mu Y_\mu \cdot p_{(\mu, 0)}^*(\bar{u}) = \sum_{\mu: (\mu, 0) \in \text{supp}(\underline{a})} \sum_{n=0}^{x_\mu Y_\mu-1} p_{(\mu, n)}^*(\bar{u}) = a \cdot p^*.$$

Similarly, we have that

$$\begin{aligned} \underline{a} \cdot \tilde{u} &= \sum_{(\mu, s') \in \text{supp}(\underline{a})} \tilde{u}_{(\mu, s')} = \sum_{(\mu, s') \in \text{supp}(\underline{a})} \left(\sum_{n=0}^{x_\mu-1} \bar{u}_{(\mu, s'Y_\mu+n)} \right) \\ &= \sum_{\mu: (\mu, 0) \in \text{supp}(\underline{a})} \sum_{n=0}^{x_\mu Y_\mu} \bar{u}_{(\mu, n)} = \sum_{(\mu, n) \in \text{supp}(a)} \bar{u}_{(\mu, n)} = a \cdot \bar{u} \end{aligned}$$

By Proposition 2.9 applied to $A^{1,\dots,\ell}$, we have $a \cdot p^* = a \cdot \bar{u}$. Hence $\underline{a} \cdot q^* = \underline{a} \cdot \tilde{u}$, as needed.

Now we must argue that q^* lies in the variety of $I(\underline{A}^{1,\dots,\ell})$. Let $\underline{A}^{1,\dots,\ell}$ denote the matrix obtained from $A^{1,\dots,\ell}$ by removing all repeated columns. The columns of $\underline{A}^{1,\dots,\ell}$ are indexed by $\mu \in [\beta]$. Let J denote the inclusion of $I(\underline{A}^{1,\dots,\ell})$ into the polynomial ring,

$$\mathbb{C}[p_{(\mu, s)} \mid \mu \in [\beta], s \in \{0, \dots, x_\mu Y_\mu - 1\}]$$

obtained by mapping p_μ to $p_{(\mu, 0)}$. Then by definition of the toric ideal associated to a matrix, we have that

$$I(A^{1,\dots,\ell}) = \langle p_{(\mu, s)} - p_{(\mu, r)} \mid \mu \in [\beta], s, r \in \{0, \dots, x_\mu Y_\mu - 1\} \rangle + J.$$

Similarly, let \tilde{J} denote the inclusion of $I(\underline{A}^{1,\dots,\ell})$ into R_B obtained by mapping p_μ to $q_{(\mu,0)}$. Then we have that

$$I(\underline{A}^{1,\dots,\ell}) = \langle q_{(\mu,s)} - q_{(\mu,r)} \mid \mu \in [\beta], s, r \in \{0, \dots, x_\mu - 1\} \rangle + \tilde{J}.$$

Thus we must check that q^* lies in each of the summands of $I(\underline{A}^{1,\dots,\ell})$.

First, $q_{(\mu,s)}^* = q_{(\mu,r)}^*$ for all $\mu \in [\beta]$ and $s, r \in \{0, \dots, x_\mu - 1\}$ by construction. Let

$$\tilde{f} = \prod_{n=1}^N q_{(\mu_n,0)} - \prod_{n=1}^N q_{(\nu_n,0)}$$

lie in \tilde{J} and $f = \prod_{n=1}^N p_{(\mu_n,0)} - \prod_{n=1}^N p_{(\nu_n,0)}$ be the corresponding element of J where $N = \deg(\tilde{f}) = \deg(f)$. By Proposition 3.14, \tilde{J} is multihomogeneous with respect to the multigrading specified by D . Since the columns of D are linearly independent, we may assume that for each $n \in [N]$, $a_{\mu_n}^B$ and $a_{\nu_n}^B$ lie in the same floret of B with respect to $A^{\ell+1}$. Thus $Y_{\mu_n} = Y_{\nu_n}$ for all n . Evaluating \tilde{f} at q^* yields

$$\begin{aligned} \tilde{f}(q^*) &= \prod_{n=1}^N q_{(\mu_n,0)}^* - \prod_{n=1}^N q_{(\nu_n,0)}^* = \prod_{n=1}^N Y_{\mu_n} \cdot p_{(\mu_n,0)}^* - \prod_{n=1}^N Y_{\nu_n} \cdot p_{(\nu_n,0)}^* \\ &= \left(\prod_{n=1}^N Y_{\mu_n} \right) \left(\prod_{n=1}^N p_{(\mu_n,0)}^* - \prod_{n=1}^N p_{(\nu_n,0)}^* \right) = 0, \end{aligned}$$

since $p^* \in V(J)$. Thus q^* lies in the closure of $\mathcal{M}(\underline{A}^{1,\dots,\ell})$, and in fact is in the model itself, since the entries of q^* are strictly positive. So by Proposition 2.9, q^* is the MLE for \tilde{u} in $\mathcal{M}(\underline{A}^{1,\dots,\ell})$. \square

Let A be a multipartition matrix with k partitions that satisfies the GRIP. The ℓ th level **florets** of the matrix A are given by the florets of $A^{\ell+1}$ (as defined in Definition 3.6) in the pair of matrices $B_\ell = \Psi_{n=1}^\ell A^n$ and $A^{\ell+1}$, with the convention that the 1st level florets are given by the rows of A^1 . We denote the ℓ th level floret containing the row a_i^ℓ as $\mathcal{F}^\ell[i]$; thus the ℓ th level floret corresponding to the j th column of A is $\mathcal{F}^\ell[S(\ell, j)]$.

Remark 3.17. There is a clear relationship between the notation described above and that used for florets of rows up until this point. In particular for the pair of matrices B_ℓ and $A^{\ell+1}$ we have that $\mathcal{F}^{\ell+1}[S(\ell+1, j)] = \mathcal{F}^{A^{\ell+1}}_{t(\bullet, S(\ell+1, j))}$ for $1 \leq \ell \leq k-1$. We make this distinction in order to avoid dependence on the matrix B produced at each step.

In order to prove one-cycle convergence for models satisfying the GRIP in Theorem 3.22 we take three important steps. These are similar to the structure of the proof of one-cycle convergence in the case of the RIP in [26] Theorem 5.3. There the author first proves in Lemma 5.8 that the MLE is a normalized product of the MLE for two submodels. We use a generalization of this result from [3], given here in Theorem 2.14, and apply it in Corollary 3.18 to show the structure of the MLE. This corresponds to Theorem 5.1 in [26]. Finally, in Theorem 3.22 and Theorem 5.3 in [26], respectively, the now known structure of the MLE is used in the induction to show one-cycle convergence.

Corollary 3.18. Let A be a multipartition matrix with k partitions that satisfies the GRIP. Then the MLE p^* of \bar{u} has as its j th coordinate function:

$$p_j^* = \frac{1}{c_j} \left(\prod_{\ell=1}^k \frac{a_{S(\ell, j)}^\ell(\bar{u})}{\sum_{a_i^\ell \in \mathcal{F}^\ell[S(\ell, j)]} a_i^\ell(\bar{u})} \right). \quad (10)$$

Proof. We induct on k . First, let $k = 1$. Then $I(A^1) = \langle p_j - p_{j'} \mid \lambda(j) = \lambda(j') \rangle$. Thus, it is straightforward to check that $p_j^* = \frac{1}{c_j} \sum_{i \in I_{S(1, j)}^1} \bar{u}_i = \frac{1}{c_j} a_{S(1, j)}^1(\bar{u})$, as needed.

Now suppose that the result holds for all $\ell \leq K$ for some natural number $K \geq 1$. Consider a partition matrix $A^{1,\dots,K+1}$ that satisfies the GRIP; then by definition of the GRIP so does $A^{1,\dots,K}$. Now let $B = \Psi_{\ell=1}^K A^\ell$, $C = A^{K+1}$, and $D = B \bowtie C$. By the GRIP, BC satisfies the floret condition. By Lemma 3.13, for each $\mu \in [\beta]$ and $\nu \in [\gamma]$, there exist $x_\mu, y_\nu \in \mathbb{Z}_+$ such that

$$\omega_{\mu\nu} = \begin{cases} x_\mu y_\nu & \text{if } t(\mu, \bullet) = t(\bullet, \nu) \\ 0 & \text{otherwise.} \end{cases}$$

Let \underline{A}^{K+1} have columns indexed by (ν, s'') where $s'' \in \{0, \dots, y_\nu - 1\}$. The (ν, s'') column of \underline{A}^{K+1} is equal to the ν th standard basis vector in \mathbb{Z}^γ . For each $\nu \in [\gamma]$, let

$$X_\nu = \sum_{\mu: a_\mu^B \in \mathcal{F}^B_{t(\bullet, \nu)}} x_\mu$$

so that $\sum_{\mu: a_{\mu}^B \in F_{t(\mu, \nu)}^B} \omega_{\mu\nu} = X_{\nu} y_{\nu}$. Let \bar{u}^C be the vector of data with entries indexed by (ν, s'') such that

$$\bar{u}_{(\nu, s'')}^C = \sum_{\mu: a_{\mu}^B \in F_{t(\mu, \nu)}^B} \sum_{n=0}^{x_{\mu}-1} \bar{u}_{(\mu, \nu, s'' x_{\mu} + n)}.$$

We note that while the entries of \bar{u}^C are ordered pairs, we do not view it as a matrix but instead as a vector (or flattened matrix), since we wish to treat it as a data vector to perform IPS on it.

Finally, let \bar{D} be the $f \times f$ identity matrix where f is the number of distinct florets of BC . Let \bar{u}^D be the vector of data indexed by the distinct florets of BC with t th component

$$\bar{u}_t^D = \sum_{\mu, \nu: t(\mu, \nu)=t} \sum_{s=0}^{\omega_{\mu\nu}-1} \bar{u}_{(\mu, \nu, s)},$$

where the triple index on \bar{u} is as in the paragraph before Definition 3.9.

Let $(p^*)^B$ denote the MLE for \bar{u} in $\mathcal{M}(\underline{A}^{1, \dots, K})$, $(p^*)^C$ denote the MLE for \bar{u}^C in $\mathcal{M}(\underline{A}^{K+1})$ and $(p^*)^D$ denote the MLE for \bar{u}^D in $\mathcal{M}(\bar{D})$. By Theorem 3.15, $I(\underline{A}^{1, \dots, K+1}) = I(\underline{A}^{1, \dots, K}) \times_D I(\underline{A}^{K+1})$. So by Theorem 2.14, we have that

$$p_{(\mu, \nu, s)}^*(\bar{u}) = \frac{(p^*)_{(\mu, s')}^B (p^*)_{(\nu, s'')}^C}{(p^*)_{t(\mu, \nu)}^D},$$

where s' and s'' are such that $s = s' y_{\nu} + s''$. We compute $(p^*)^B$, $(p^*)^C$ and $(p^*)^D$ to prove the desired result. Fix $\mu \in [\beta]$ and $\nu \in [\gamma]$. By induction and Proposition 3.16, we have that

$$\begin{aligned} (p^*)_{(\mu, s')}^B &= Y_{\mu} \cdot p_{(\mu, \nu, 0)}^* \\ &= \frac{Y_{\mu}}{x_{\mu} Y_{\mu}} \left(\prod_{\ell=1}^K \frac{a_{S(\ell, (\mu, \nu, 0))}^{\ell}(\bar{u})}{\sum_{a_{\mu'}^{\ell} \in F^{\ell}[S(\ell, (\mu, \nu, 0))]} a_{\mu'}^{\ell}(\bar{u})} \right) = \frac{1}{x_{\mu}} \left(\prod_{\ell=1}^K \frac{a_{S(\ell, (\mu, \nu, s))}^{\ell}(\bar{u})}{\sum_{a_{\mu'}^{\ell} \in F^{\ell}[S(\ell, (\mu, \nu, s))]} a_{\mu'}^{\ell}(\bar{u})} \right), \end{aligned}$$

since $S(\ell, (\mu, \nu, s)) = S(\ell, (\mu, \nu, 0))$ for all s . Since \underline{A}^{K+1} is an identity matrix with repeated columns, we have that

$$\begin{aligned} (p^*)_{(\nu, s'')}^C &= \frac{1}{y_{\nu}} \sum_{n=0}^{y_{\nu}-1} \bar{u}_{(\nu, n)}^C = \frac{1}{y_{\nu}} \sum_{\mu: a_{\mu}^B \in F_{t(\mu, \nu)}^B} \sum_{n=0}^{\omega_{\mu\nu}-1} \bar{u}_{(\mu, \nu, n)} \\ &= \frac{1}{y_{\nu}} (a_{\nu}^{K+1}(\bar{u})) = \frac{1}{y_{\nu}} \left(a_{S(K+1, (\mu, \nu, s))}^{K+1}(\bar{u}) \right). \end{aligned}$$

The equality in the first line comes from the definition of the vector \bar{u}^C .

Let $t = t(\mu, \nu)$. Finally, since \bar{D} is an identity matrix, we have that

$$(p^*)_t^D = \bar{u}_t^D = \sum_{\mu', \nu': t(\mu', \nu')=t} \sum_{s=0}^{\omega_{\mu'\nu'}-1} \bar{u}_{(\mu', \nu', s)} = \sum_{a_{\nu'}^{K+1} \in F^{K+1}[S(\ell, (\mu, \nu, s))]} a_{\nu'}^{K+1}(\bar{u}).$$

Thus by Theorem 5.5 of [3], the (μ, ν, s) coordinate of the MLE for \bar{u} in $\mathcal{M}(\underline{A}^{1, \dots, K+1})$ is

$$\begin{aligned} p_{(\mu, \nu, s)}^*(\bar{u}) &= \frac{1}{x_{\mu}} \left(\prod_{\ell=1}^K \frac{a_{S(\ell, (\mu, \nu, s))}^{\ell}(\bar{u})}{\sum_{a_{\mu'}^{\ell} \in F^{\ell}[S(\ell, (\mu, \nu, s))]} a_{\mu'}^{\ell}(\bar{u})} \right) \\ &\quad \times \frac{1}{y_{\nu}} (a_{S(K+1, (\mu, \nu, s))}^{K+1}(\bar{u})) \times \frac{1}{\sum_{a_{\nu'}^{K+1} \in F^{K+1}[S(\ell, (\mu, \nu, s))]} a_{\nu'}^{K+1}(\bar{u})} \\ &= \frac{1}{x_{\mu} y_{\nu}} \left(\prod_{\ell=1}^{K+1} \frac{a_{S(\ell, (\mu, \nu, s))}^{\ell}(\bar{u})}{\sum_{a_{\mu'}^{\ell} \in F^{\ell}[S(\ell, (\mu, \nu, s))]} a_{\mu'}^{\ell}(\bar{u})} \right) = \frac{1}{\omega_{\mu\nu}} \left(\prod_{\ell=1}^{K+1} \frac{a_{S(\ell, (\mu, \nu, s))}^{\ell}(\bar{u})}{\sum_{a_{\mu'}^{\ell} \in F^{\ell}[S(\ell, (\mu, \nu, s))]} a_{\mu'}^{\ell}(\bar{u})} \right). \end{aligned}$$

This proves the result for $\underline{A}^{1, \dots, K+1}$ and we are done by induction. \square

We now define an iterated toric fiber product of linear ideals. These are defined so that if $I(\underline{A}^{1, \dots, k})$ is an iterated toric fiber product of models corresponding to partition matrices, then $\underline{A}^{1, \dots, k}$ satisfies the GRIP.

Definition 3.19. Let A^1, \dots, A^k be partition matrices such that $A^{1, \dots, \ell}$ and $A^{\ell+1}$ satisfy the floret condition for each $\ell \in [k-1]$. Let $D_\ell = (\Psi_{\eta=1}^\ell A^n) \mathbb{M} A^{\ell+1}$ for each $\ell \in [k-1]$. Suppose there exist compressions $\underline{A}^{1, \dots, \ell}$ and $\underline{A}^{\ell+1}$ of $A^{1, \dots, \ell}$ and $A^{\ell+1}$, respectively, such that $I(\underline{A}^{1, \dots, \ell}) \times_{D_\ell} I(\underline{A}^{\ell+1})$ is defined and equal to $I(A^{1, \dots, \ell+1})$ for each $\ell \in [k-1]$. Then the ideal $I(A^{1, \dots, k})$ is an **iterated toric fiber product of linear ideals**.

Proposition 3.20. Let $A^{1, \dots, k}$ be a multipartition matrix such that $I(A^{1, \dots, k})$ is an iterated toric fiber product of linear ideals. Then $A^{1, \dots, k}$ is well-connected and satisfies the GRIP.

Proof. Fix $\ell \in [k-1]$. Let $B = \Psi_{\eta=1}^\ell A^n$ and let $C = A^{\ell+1}$. By Definition 3.19, BC satisfies the floret condition. Thus, $D_\ell = B \mathbb{M} C$ is well-defined. Let $\underline{A}^{1, \dots, \ell}$ and $\underline{A}^{\ell+1}$ be such that $I(\underline{A}^{1, \dots, \ell}) \times_{D_\ell} I(\underline{A}^{\ell+1})$ is defined and equal to $I(A^{1, \dots, \ell+1})$. Then since the toric fiber product is defined, both $I(\underline{A}^{1, \dots, \ell})$ and $I(\underline{A}^{\ell+1})$ are multihomogeneous with respect to the multigrading given by D_ℓ . But $A^{1, \dots, \ell}$ is obtained from $\underline{A}^{1, \dots, \ell}$ by adding repeated columns. Thus, each row of D_ℓ is contained in the orthogonal complement of the kernel of $A^{1, \dots, \ell}$, which is equal to the rowspan of $A^{1, \dots, \ell}$.

Let $\underline{A}^{1, \dots, \ell}$ have β distinct columns with labels $1, \dots, \beta$. Let x_μ be the number of copies of column μ in $\underline{A}^{1, \dots, \ell}$. Let y_v be the number of copies of the v th standard basis vector in $\tilde{A}^{\ell+1}$. Then by the construction of the toric fiber product, there are $\omega_{\mu v} = x_\mu y_v$ columns of $A^{1, \dots, \ell+1}$ that consist of column μ of $\underline{A}^{1, \dots, \ell}$ concatenated with e_v .

Let $\mu, \mu' \in [\beta]$ be such that a_μ^B and $a_{\mu'}^B$ belong to the same floret of B . Let a_v^C belong to the corresponding floret of C and let \mathcal{F} denote the set of all rows of C in the same floret as a_v^C . Then the number of copies of column μ in $A^{1, \dots, \ell}$ is $x_\mu \sum_{v': a_{v'}^C \in \mathcal{F}} y_{v'}$. Similarly, the number of copies of column μ' in $A^{1, \dots, \ell}$ is $x_{\mu'} \sum_{v': a_{v'}^C \in \mathcal{F}} y_{v'}$. Thus we have

$$\omega_{\mu v} / (x_\mu \sum_{v': a_{v'}^C \in \mathcal{F}} y_{v'}) = y_v / (\sum_{v': a_{v'}^C \in \mathcal{F}} y_{v'}) = \omega_{\mu' v} / (x_{\mu'} \sum_{v': a_{v'}^C \in \mathcal{F}} y_{v'}).$$

Thus BC is well-connected. The above argument holds for all $\ell \in [k-1]$. Thus $A^{1, \dots, k}$ satisfies the GRIP. \square

We now show that for multipartition matrices that satisfy the GRIP, the IPS algorithm produces the MLE in exactly one cycle. In order to understand this MLE, we must first establish the following equality.

Lemma 3.21. If A is a multipartition matrix with k partitions that satisfies the GRIP, then the following equation holds for any row vector a_i^k of A^k :

$$\sum_{j \in I_i^k} \frac{1}{c_j^{k-1}} \left(\prod_{\ell=1}^{k-1} \frac{a_{S(\ell, j)}^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in \mathcal{F}^\ell[S(\ell, j)]} a_{i'}^\ell(\bar{u})} \right) = C_i^k \sum_{a_{i'}^k \in \mathcal{F}^k[i]} a_{i'}^k(\bar{u}).$$

Proof. By Corollary 3.18, the MLE of \mathcal{M}_A is given by $p = (p_1, \dots, p_m)^T$ where

$$p_j = \frac{1}{c_j^k} \left(\prod_{\ell=1}^k \frac{a_{S(\ell, j)}^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in \mathcal{F}^\ell[S(\ell, j)]} a_{i'}^\ell(\bar{u})} \right),$$

for $j \in \{1, \dots, m\}$. Since p is the MLE it also satisfies $Ap = A\bar{u}$ by Proposition 2.9. In particular for any row vector a_i^k we have that

$$a_i^k(p) = \sum_{j \in I_i^k} \frac{1}{c_j^k} \left(\prod_{\ell=1}^k \frac{a_{S(\ell, j)}^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in \mathcal{F}^\ell[S(\ell, j)]} a_{i'}^\ell(\bar{u})} \right) = a_i^k(\bar{u}).$$

Note that for any $j \in I_i^k$, we have $(1/c_j^k) = (c_j^{k-1}/c_j^{k-1}c_j^k) = (1/C_i^k c_j^{k-1})$ since A is well-connected. Thus we pull out a factor of $(1/C_i^k)$ out of each term. Additionally for any $j \in I_i^k$, we have that $S(k, j) = i$ by definition, and hence we rewrite the middle term of the above equation as

$$\frac{a_i^k(\bar{u})}{C_i^k \sum_{a_{i'}^k \in \mathcal{F}^k[i]} a_{i'}^k(\bar{u})} \left(\sum_{j \in I_i^k} \frac{1}{c_j^{k-1}} \left(\prod_{\ell=1}^{k-1} \frac{a_{S(\ell, j)}^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in \mathcal{F}^\ell[S(\ell, j)]} a_{i'}^\ell(\bar{u})} \right) \right).$$

The result immediately follows from this. \square

Theorem 3.22. *If A is a multipartition matrix with k partitions that satisfies the GRIP, then the IPS algorithm applied to A results in the MLE after one cycle.*

Proof. We proceed by induction on k . For $k = 1$, IPS produces the MLE of \mathcal{M}_A after just one step. Indeed the information projection onto the linear family defined by A in this case is simply the MLE of \mathcal{M}_A (see (4) and Proposition 2.9). Now assume that for any multipartition matrix A satisfying the GRIP where $k \leq K$ for some natural number $K \geq 1$, the IPS algorithm results in the MLE. By Corollary 3.18, the MLE p^k for $\mathcal{M}_{A^{1,\dots,k}}$ is of the form in (10) in this case.

Now let $A^{1,\dots,K+1}$ be a multipartition matrix with $K + 1$ partitions satisfying the GRIP. Note that based on the definition of the GRIP, $A^{1,\dots,K}$ also satisfies the GRIP.

Performing the first K steps of the IPS algorithm on A is equivalent to one cycle on $A^{1,\dots,K}$. Thus we have, by the induction hypothesis and Corollary 3.18, that p^K is the MLE of $\mathcal{M}_{A^{1,\dots,K}}$, i.e.

$$p_j^K = \frac{1}{c_j^K} \left(\prod_{\ell=1}^K \frac{a_{S(\ell,j)}^\ell(\bar{u})}{\sum_{a_i^\ell \in F^\ell[S(\ell,j)]} a_i^\ell(\bar{u})} \right).$$

Then the $(K + 1)$ st step of IPS scales each p_j^K by

$$\frac{a_{S(K+1,j)}^{K+1}(\bar{u})}{\sum_{j' \in I_{S(K+1,j)}^{K+1}} \frac{1}{c_{j'}^K} \left(\prod_{\ell=1}^K \frac{a_{S(\ell,j')}^\ell(\bar{u})}{\sum_{a_i^\ell \in F^\ell[S(\ell,j')] } a_i^\ell(\bar{u})} \right)} = \frac{a_{S(K+1,j)}^{K+1}(\bar{u})}{C_{S(K+1,j)}^K \sum_{a_i^{K+1} \in F^{K+1}[S(K+1,j)]} a_i^{K+1}(\bar{u})},$$

where equality follows from Lemma 3.21. Since $C_{S(K+1,j)}^{K+1} = \frac{c_j^{K+1}}{c_j^K}$ this implies that p^{K+1} is exactly the MLE of $\mathcal{M}_{A^{1,\dots,K+1}}$ by Proposition 3.18. By induction this proves the result. \square

Not only does the proof of Theorem 3.22 show that the IPS algorithm constructs the MLE in one cycle, but also that at the ℓ -th step the vector p^ℓ is exactly the MLE of the $A^{1,\dots,\ell}$. It is natural to ask if the property of producing the MLE in one cycle also characterizes multipartition matrices satisfying the GRIP. However, we immediately see that this is not the case in general. Consider a $n \times m$ multipartition matrix A with k partitions such that A^k is the identity matrix I_m . Then \mathcal{M}_A is the entire simplex Δ_{m-1} and IPS automatically produces vector \bar{u} (the MLE of \mathcal{M}_A in this case) regardless of the first $k - 1$ partitions. Another example is given below.

Example 3.23. Consider the multipartition matrix A below. This matrix does not satisfy the GRIP because it is not well-connected.

$$A = \begin{pmatrix} 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \\ \hline 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 \\ \hline 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{pmatrix}$$

The column weights of the columns 6 and 8 of the first two partitions are

$$c_6^1 = 3, c_8^1 = 1, c_6^2 = 1 \text{ and } c_8^2 = 1. \text{ Hence, } \frac{c_6^2}{c_6^1} \neq \frac{c_8^2}{c_8^1}.$$

However, the IPS converges to the MLE

$$p^*(\bar{u}) = ((\bar{u}_1 + \bar{u}_2)(\bar{u}_1 + \bar{u}_3), (\bar{u}_1 + \bar{u}_2)(\bar{u}_2 + \bar{u}_4), (\bar{u}_3 + \bar{u}_4)(\bar{u}_1 + \bar{u}_3), (\bar{u}_3 + \bar{u}_4)(\bar{u}_2 + \bar{u}_4), \frac{1}{2}(\bar{u}_5 + \bar{u}_7), \bar{u}_6, \frac{1}{2}(\bar{u}_5 + \bar{u}_7), \bar{u}_8)^T$$

in one cycle.

Remark 3.24. For a multipartition matrix A with $k = 2$, it is possible to show that if the IPS algorithm produces the MLE in one cycle, then A satisfies the GRIP. This relies on being able to explicitly write the coordinate functions of the MLE as a product of linear forms that must satisfy the Horn uniformization [18,31]. While the above counter-example implies that this is not the case for $k = 3$,

it would be interesting to investigate whether requiring IPS to produce the MLE of $A^{1,\dots,\ell}$ at the ℓ -th step is sufficient to guarantee that A satisfies the GRIP.

Example 3.25. Here we apply IPS to the matrix A from Example 2.5. The rows of A are colored blue in this Example. (For interpretation of the colors in the text, the reader is referred to the web version of this article.)

Let $\bar{u} = (\bar{u}_1, \dots, \bar{u}_{14})^T$ be the normalized data vector. Projecting to the first partition of A leads to $p_1^0 = \dots = p_7^0 = \frac{1}{7} a_1^1(\bar{u})$ and $p_8^0 = \dots = p_{14}^0 = \frac{1}{7} a_2^1(\bar{u})$. The second step of the algorithm results in four different types of indices. These are given by the different rows in the matrix $A^1 \uplus A^2$, indicated by the different dashed and dotted lines below the matrix.

$$A^1 \uplus A^2 = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 \end{pmatrix} \begin{matrix} a_1^{1\mathfrak{m}2} \\ a_1^{1\psi2} \\ a_2^{1\psi2} \\ a_3^{1\psi2} \\ a_4^{1\psi2} \end{matrix}$$

- - - - -

Note that since $A^1 \uplus A^2$ consists only of the ones vector and \bar{u} is normalized, we have $a_1^{1\mathfrak{m}2}(\bar{u}) = 1$. Therefore the second projection simplifies in the following way:

$$\begin{aligned} p_1^1 &= p_2^1 = p_3^1 = \frac{1}{C_1^1} a_1^1(\bar{u}) \frac{a_1^2(\bar{u})}{C_1^2 a_1^{1\mathfrak{m}2}(\bar{u})} = \frac{1}{3} a_1^1(\bar{u}) a_1^2(\bar{u}) \\ p_4^1 &= p_5^1 = p_6^1 = p_7^1 = \frac{1}{C_1^1} a_1^1(\bar{u}) \frac{a_2^2(\bar{u})}{C_2^2 a_1^{1\mathfrak{m}2}(\bar{u})} = \frac{1}{4} a_1^1(\bar{u}) a_2^2(\bar{u}) \\ p_8^1 &= p_9^1 = p_{10}^1 = \frac{1}{C_2^1} a_2^1(\bar{u}) \frac{a_1^2(\bar{u})}{C_1^2 a_1^{1\mathfrak{m}2}(\bar{u})} = \frac{1}{3} a_2^1(\bar{u}) a_1^2(\bar{u}) \\ p_{11}^1 &= p_{12}^1 = p_{13}^1 = p_{14}^1 = \frac{1}{C_2^1} a_2^1(\bar{u}) \frac{a_2^2(\bar{u})}{C_2^2 a_1^{1\mathfrak{m}2}(\bar{u})} = \frac{1}{4} a_2^1(\bar{u}) a_2^2(\bar{u}) \end{aligned}$$

For the last projection we only demonstrate four indices as an example.

$$(A^1 \uplus A^2) \mathfrak{m} A^3 = \begin{pmatrix} 1 & 1 & 1 & \cdot & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 \end{pmatrix} \begin{matrix} a_1^{(1\psi2)\mathfrak{m}3} \\ a_2^{(1\psi2)\mathfrak{m}3} \end{matrix}$$

In this case $A^2 = (A^1 \uplus A^2) \mathfrak{m} A^3$, hence $a_i^2(\bar{u}) = a_i^{(1\psi2)\mathfrak{m}3}(\bar{u})$. In order to visualize the structure created by IPS, we also display the general formula of p_j^2 without simplifications

$$\begin{aligned} p_1^2 &= \frac{1}{C_1^1} a_1^1(\bar{u}) \frac{a_1^2(\bar{u})}{C_1^2 a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_1^3(\bar{u})}{C_1^3 a_1^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{7} a_1^1(\bar{u}) \frac{a_1^2(\bar{u})}{\frac{3}{7} a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_1^3(\bar{u})}{\frac{1}{3} a_1^{(1\psi2)\mathfrak{m}3}(\bar{u})} = a_1^1(\bar{u}) a_1^3(\bar{u}) \\ p_4^2 &= \frac{1}{C_1^1} a_1^1(\bar{u}) \frac{a_2^2(\bar{u})}{C_2^2 a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_4^3(\bar{u})}{C_4^3 a_2^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{7} a_1^1(\bar{u}) \frac{a_2^2(\bar{u})}{\frac{4}{7} a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_4^3(\bar{u})}{\frac{2}{4} a_2^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{2} a_1^1(\bar{u}) a_4^3(\bar{u}) \\ p_8^2 &= \frac{1}{C_2^1} a_2^1(\bar{u}) \frac{a_1^2(\bar{u})}{C_1^2 a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_1^3(\bar{u})}{C_1^3 a_1^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{7} a_2^1(\bar{u}) \frac{a_1^2(\bar{u})}{\frac{3}{7} a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_1^3(\bar{u})}{\frac{1}{3} a_1^{(1\psi2)\mathfrak{m}3}(\bar{u})} = a_2^1(\bar{u}) a_1^3(\bar{u}) \\ p_{11}^2 &= \frac{1}{C_2^1} a_2^1(\bar{u}) \frac{a_2^2(\bar{u})}{C_2^2 a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_4^3(\bar{u})}{C_4^3 a_2^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{7} a_2^1(\bar{u}) \frac{a_2^2(\bar{u})}{\frac{4}{7} a_1^{1\mathfrak{m}2}(\bar{u})} \frac{a_4^3(\bar{u})}{\frac{2}{4} a_2^{(1\psi2)\mathfrak{m}3}(\bar{u})} = \frac{1}{2} a_2^1(\bar{u}) a_4^3(\bar{u}) \end{aligned}$$

There exist two columns identical to column 4 and two columns identical to column 11. Hence the corresponding connection ratios do not add up to one, but to $\frac{1}{2}$.

We end this section by noting that the formula for the MLE of a matrix satisfying the GRIP in (10) factors through the associated monomial map ϕ_A . Re-ordering if necessary, we denote \bar{u} in the ambient space of \mathcal{M}_A as $\bar{u} = (\bar{u}_1^1, \dots, \bar{u}_1^{c_1}, \bar{u}_2^1, \dots, \bar{u}_2^{c_2}, \dots, \bar{u}_m^1, \dots, \bar{u}_m^{c_m})^T$ where c_j is the column weight. We have that

$$p^*(\bar{u}) = \phi_A \left(\frac{s_1^1(\bar{u})}{C_1^1}, \dots, \frac{s_{n_1}^1(\bar{u})}{C_{n_1}^1}, \frac{s_1^2(\bar{u})}{C_1^2}, \dots, \frac{s_{n_k}^k(\bar{u})}{C_{n_k}^k} \right) \text{ where } s_i^\ell(\bar{u}) = \frac{a_i^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in F^\ell[i]} a_{i'}^\ell(\bar{u})}.$$

This factorization implies a nice correspondence between the MLE of the matrix \underline{A} obtained by removing all repeated columns of A and the MLE of A .

Proposition 3.26. *Let A be a multipartition matrix with k partitions that satisfies the GRIP and let \underline{A} be the $n \times m$ matrix obtained by removing all repeated columns of A . Let $\bar{u} = (\bar{u}_1, \dots, \bar{u}_m)$ be a normalized data vector of counts. Then the MLE of $\mathcal{M}_{\underline{A}}$ for a data vector \bar{u} is given by*

$$p_{\underline{A}}^*(\bar{u}) = \phi_{\underline{A}} \left(s_1^1(\bar{u}), \dots, s_{n_1}^1(\bar{u}), s_1^2(\bar{u}), \dots, s_{n_k}^k(\bar{u}) \right), \text{ where } s_i^\ell(\bar{u}) = \frac{a_i^\ell(\bar{u})}{\sum_{a_{i'}^\ell \in F^\ell[i]^\ell} a_{i'}^\ell(\bar{u})}$$

for row vectors a_i^ℓ of A' .

Proof. Clearly $p_{\underline{A}}^*(\bar{u})$ lies on the model $\mathcal{M}_{\underline{A}}$; it remains to show that $\underline{A}(p_{\underline{A}}^*(\bar{u})) = \underline{A}(\bar{u})$ which implies that $p_{\underline{A}}^*(\bar{u})$ is the MLE of $\mathcal{M}_{\underline{A}}$ by Proposition 2.9. Consider a data vector \bar{u} in the ambient space of \mathcal{M}_A where $\bar{u}_j^i = (\bar{u}_j / c_j)$ for all $j \in [m]$ and $i \in [c_j]$. Then we have that $\bar{u}_j = \bar{u}_j^1 + \bar{u}_j^2 + \dots + \bar{u}_j^{c_j}$.

Since A satisfies the GRIP, the MLE of \mathcal{M}_A with respect to \bar{u} , denoted $p_A^*(\bar{u})$, is given in (10). For any row a_i^ℓ of \underline{A} , it is easy to see that $a_i^\ell(\bar{u}) = a_i^\ell(\bar{u})$. Thus we write the coordinate functions of $p_A^*(\bar{u})$ in terms of the coordinate functions of $p_{\underline{A}}^*(\bar{u})$ as

$$p_A^*(\bar{u})_j^i = \frac{p_{\underline{A}}^*(\bar{u})_j}{c_j}$$

for all $i \in [c_j]$. So we have that

$$\sum_{i=1}^{c_j} p_A^*(\bar{u})_j^i = p_{\underline{A}}^*(\bar{u})_j,$$

which implies that

$$a_i^\ell(p_A^*(\bar{u})) = \sum_{j=1}^m \sum_{i=1}^{c_j} p_A^*(\bar{u})_j^i = \sum_{j=1}^m p_{\underline{A}}^*(\bar{u})_j = a_i^\ell(p_{\underline{A}}^*(\bar{u})).$$

Putting the above two equalities together, along with the fact that $p_A^*(\bar{u})$ is the MLE of \mathcal{M}_A , we have that $a_i^\ell(\bar{u}) = a_i^\ell(\bar{u}) = a_i^\ell(p_A^*(\bar{u})) = a_i^\ell(p_{\underline{A}}^*(\bar{u}))$, which in turn implies that $a_i^\ell(\bar{u}) = a_i^\ell(p_{\underline{A}}^*(\bar{u}))$. \square

In Proposition 3.26, we assume that the data vector is normalized only to be consistent with the notation of the previous results. Indeed, the assumption that the data is normalized is not necessary as the normalizing factor cancels in each function s_i^ℓ .

Example 3.27. In Proposition 3.26, if one removes all repeated columns from a multipartition matrix that satisfies the GRIP, the result also satisfies the GRIP. The reverse is, however, not true. Consider the independence model from Example 2.10. We see in that example that the IPS algorithm produces the MLE in one cycle. One can easily check that this model satisfies the GRIP as well, so this is not surprising.

However, if we simply repeat the last column to obtain a new multipartition matrix, shown here,

$$\begin{pmatrix} 1 & 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & 1 \\ 1 & \cdot & 1 & \cdot & \cdot \\ \cdot & 1 & \cdot & 1 & 1 \end{pmatrix}$$

then the IPS algorithm does not produce the MLE after one cycle. In fact, one can check (by computing the solution set to the polynomial equations defining the ML-degree) that the resulting multipartition model has non-rational MLE. This implies that the IPS algorithm *cannot* produce the MLE exactly.

4. Model families satisfying the GRIP

4.1. Hierarchical models

In this section, we justify our use of the name “generalized running intersection property”. As discussed in Section 2.4, we associate a hierarchical model with a simplicial complex Γ . Recall that if a simplicial complex is decomposable, then by Lemma 2.19,

there exists an ordering on the facets that satisfies RIP. Recall further the matrix that defines a binary hierarchical model, as given in Equation (7).

In the following proposition, we prove that a binary hierarchical model with the RIP also satisfies the GRIP. We note that an analogous argument proves the result for non-binary random variables as well; we do not include a proof here to maintain notational simplicity.

Proposition 4.1. *Let Γ be a simplicial complex and let A be a multipartition matrix representing the binary hierarchical model on Γ that satisfies the running intersection property. Then A satisfies the generalized running intersection property as well.*

Proof. Let Γ be a simplicial complex on $[p]$ with facets F_1, \dots, F_k such that the resulting multipartition matrix $A^{1, \dots, k}$ satisfies the running intersection property. Let $1 \leq \ell \leq k$, $B = \bigcup_{n=1}^{\ell} A^n$ and let $C = A^{\ell+1}$. Without loss of generality, let F_1, \dots, F_{ℓ} form a simplicial complex on $[q]$ for $q \leq p$. Since $A^{1, \dots, k}$ satisfies the RIP, there exists an $s \leq \ell$ such that $[q] \cap F_{\ell+1} = F_s \cap F_{\ell+1}$.

First, we wish to show that the matrix BC satisfies the floret condition. The matrix B has rows indexed by subsets of $[q]$. For each $S \subset [p]$ and $T \subset [q]$, the S entry of a_T^B is equal to 1 if $T = S \cap [q]$ and 0 otherwise. Let $U \subset F_{\ell+1}$. We claim that a_U^C is connected to a_T^B if and only if $U \cap [q] = T \cap F_{\ell+1}$.

Indeed, suppose that $U \cap [q] = T \cap F_{\ell+1}$. Let $S = U \cap T$. Then

$$S \cap [q] = (U \cap [q]) \cup (T \cap [q]) = (T \cap F_{\ell+1}) \cup T = T.$$

Similarly, $S \cap F_{\ell+1} = U$. Hence the S coordinates of a_T^B and a_U^C are both equal to one, and these rows are connected.

Conversely, suppose that a_T^B and a_U^C are connected. Then there exists an $S \subset [p]$ such that $U = S \cap F_{\ell+1}$ and $T = S \cap [q]$. Intersecting the first equation with $[q]$ and the second with $F_{\ell+1}$ yields that $U \cap [q] = T \cap F_{\ell+1}$, as needed.

Thus, for each $T, T' \subset [q]$, if $T \cap F_{\ell+1} = T' \cap F_{\ell+1}$, then we have that the sets of rows of C connected to a_T^B and $a_{T'}^B$ are equal. Otherwise these sets are disjoint. Hence BC satisfies the floret condition.

For each $S \subset [p]$, the number of columns of B identical to the column associated to S is 2^{p-q} . Indeed, the columns of B associated to S and S' are equal if and only if $S \cap [q] = S' \cap [q]$, and there are 2^{p-q} such subsets of $[p]$. Similarly, the number of columns of BC identical to the column associated to S is $2^{p-q'}$, where $q' = \#([q] \cup F_{\ell+1})$. Thus, since all column weights within B are equal to one another and all column weights within BC are equal to one another, BC is well-connected.

Finally, we must show that $B \bowtie C$ lies in the rowspan of $A^{1, \dots, \ell}$. Recall that there exists an $s \leq \ell$ such that $[q] \cap F_{\ell+1} = F_s \cap F_{\ell+1}$. We showed above that the rows of $D = B \bowtie C$ are indexed by subsets of $[q] \cap F_{\ell+1}$. For each $T \subset [q] \cap F_{\ell+1}$ and each $S \subset [p]$, the row a_T^D has S entry equal to 1 if $S \cap ([q] \cap F_{\ell+1}) = T$ and 0 otherwise.

We claim that for each $T \subset [q] \cap F_{\ell+1}$, a_T^D is the sum of all rows of A^s , a_U^s , such that $U \cap F_{\ell+1} = T$, i.e. a_T^D is equal to

$$\sum_{\substack{U \subset F_{\ell+1}: \\ U \cap F_{\ell+1} = T}} a_U^s. \quad (11)$$

Indeed, let $S \subset [p]$ satisfy $S \cap F_{\ell+1} \cap [q] = T$. Then we have that $(S \cap F_s) \cap F_{\ell+1} = T$. In particular, the term of Equation (11) corresponding to $U = S \cap F_s$ has S entry equal to 1, as needed.

Similarly, if $S \cap F_{\ell+1} \cap [q] = T' \neq T$, then we have $T' = (S \cap F_s) \cap F_{\ell+1}$. In particular, the row $a_{S \cap F_s}^s$ is the row of A^s supported on S . But it is not a term of the sum in Equation (11). Hence the sum in (11) has S entry equal to 0, as needed.

Since each row of D lies in the rowspan of A^s , it lies in the rowspan of $A^{1, \dots, \ell}$, as needed. Hence $A^{1, \dots, k}$ satisfies the generalized running intersection property. \square

4.2. Balanced staged tree models

In this section, we explore the implications of the results in Section 3 for staged tree models. An introduction to staged tree models can be found in Appendix A. We show that a staged tree model is balanced and stratified if and only if the associated matrix satisfies the GRIP, leading to a new characterization of balanced, stratified staged trees. In particular for the associated multipartition matrix, the IPS algorithm results in the MLE after one cycle.

To define the notion of a balanced staged tree, first we must describe the interpolating polynomial of a staged tree \mathcal{T} , first introduced in [24]. For any vertex $v \in V$, we consider the subtree \mathcal{T}_v rooted at v . We denote the set of root-to-leaf paths of \mathcal{T}_v (or equivalently the v -to-leaf paths in \mathcal{T}) as Λ_v .

Definition 4.2. For a staged tree (\mathcal{T}, θ) and a vertex $v \in V$, let $t(v) := \sum_{\lambda \in \Lambda_v} \prod_{s_i \in \theta(\lambda)} s_i$. If v_0 is the root of \mathcal{T} , then $t(v_0)$ is the interpolating polynomial of \mathcal{T} .

Definition 4.3. A staged tree (\mathcal{T}, θ) is **balanced** if for any two vertices $v, w \in V$ such that $F_v = F_w$ the following equation

$$t(v')t(w'') = t(w')t(v'') \quad (12)$$

is satisfied for all distinct pairs of vertices $v', v'' \in \text{ch}(v)$ and $w', w'' \in \text{ch}(w)$ where $\theta(v, v') = \theta(w, w')$ and $\theta(v, v'') = \theta(w, w'')$.

Example 4.4. Consider the staged tree (\mathcal{T}, θ) in Example A.4. To determine if the tree is balanced we must compare all pairs of vertices in the same stage. The pairs $v_{s_0 t_0}, v_{s_1 t_0}$ and $v_{s_0 t_1}, v_{s_1 t_1}$ trivially satisfy the condition in (12). Thus to determine whether \mathcal{T} is balanced it is enough to check whether (12) holds for v_{s_0} and v_{s_1} . We have that $\text{ch}(v_{s_0}) = \{v_{s_0 t_0}, v_{s_0 t_1}\}$ and $\text{ch}(v_{s_1}) = \{v_{s_1 t_0}, v_{s_1 t_1}\}$. So we have to check that $t(v_{s_0 t_0})t(v_{s_1 t_1}) = t(v_{s_1 t_0})t(v_{s_0 t_1})$. Writing out the corresponding polynomials,

$$t(v_{s_0 t_0}) = t(v_{s_1 t_0}) = r_0 + r_1 + r_2, \quad t(v_{s_0 t_1}) = t(v_{s_1 t_1}) = r_3 + r_4,$$

we see that this is satisfied. The staged tree (\mathcal{T}, θ) additionally satisfies the property that for any two vertices v, w in the same stage, $t(v) = t(w)$. In [4, Lemma 2.13] it was shown that this property implies balanced. To “unbalance” the staged tree, one could re-color as blue one of the red vertices, changing the resulting floret. For another example of an unbalanced staged tree see [4, Example 2.14].

One can determine statistical and algebraic properties of a staged tree model from its interpolating polynomial [23,24] and whether or not it is balanced [4,17]. In particular the following proposition follows from [17, Thm. 10] and applies to our definition of staged tree models due to [4, Lem. 5.13]

Proposition 4.5. A staged tree model $\mathcal{M}_{\mathcal{T}}$ is equal to the toric model $\widetilde{\mathcal{M}}_{\mathcal{T}}$ if and only if (\mathcal{T}, θ) is balanced.

Example 4.6. For the staged tree (\mathcal{T}, θ) introduced in Example A.4, we define the toric model $\widetilde{\mathcal{M}}_{\mathcal{T}}$ by the monomial map $\phi_{\mathcal{T}}$.

This is equivalent to the monomial map $\phi_{A_{\mathcal{T}}}$ arising from the multipartition matrix $A_{\mathcal{T}}$ below.

$$A_{\mathcal{T}} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & 1 & 1 \\ 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 \end{pmatrix} \begin{matrix} s_0 \\ s_1 \\ t_0 \\ t_1 \\ r_0 \\ r_1 \\ r_2 \\ r_3 \\ r_4 \end{matrix}$$

Thus $\widetilde{\mathcal{M}}_{\mathcal{T}}$ is the log-linear model $\mathcal{M}_{A_{\mathcal{T}}}$. In Example 4.4 we showed that (\mathcal{T}, θ) is a balanced staged tree, and hence Proposition 4.5 implies that $\mathcal{M}_{\mathcal{T}} = \widetilde{\mathcal{M}}_{\mathcal{T}} = \mathcal{M}_{A_{\mathcal{T}}}$.

In the case that (\mathcal{T}, θ) is stratified and balanced, the staged tree model $\mathcal{M}_{\mathcal{T}}$ is also a multipartition model. Since the tree is stratified the associated matrix $A_{\mathcal{T}}$ with the monomial map $\phi_{\mathcal{T}}$ (defined in (14)) is a multipartition matrix, and being balanced implies that $\mathcal{M}_{\mathcal{T}} = \widetilde{\mathcal{M}}_{\mathcal{T}} = \mathcal{M}_{A_{\mathcal{T}}}$.

Note that the set of models defined by multipartition matrices A whose associated tree \mathcal{T}_A is a balanced, stratified staged tree is larger than the set of staged tree models arising from balanced, stratified staged trees. When the matrix A has repeated columns, this information is lost when constructing the tree graph \mathcal{T}_A . Thus the staged tree model arising from $(\mathcal{T}_A, \theta_A)$ is the model obtained by removing repeated columns from a multipartition matrix A .

Example 4.7. Consider the matrix A from Example 3.4, which is used throughout Section 3. On the left of Fig. 7 the associated tree $(\mathcal{T}_A, \theta_A)$ is drawn and on the right we color portions of the matrix corresponding to the blue and red florets. We draw double circles around four of the leaves to indicate that the associated column is repeated in the matrix A . However, the underlying labeled tree graph is exactly the same as in Examples A.4 and 4.4, and hence is a stratified and balanced staged tree.

We show that the stratified and balanced conditions on a staged tree are related to the generalized running intersection property for multipartition matrices.

Lemma 4.8. The multipartition matrix associated to a stratified and balanced staged tree is well-connected.

Proof. For a balanced and stratified staged tree we have that $\mathcal{M}_A = \mathcal{M}_{\mathcal{T}_A}$ and each column of A represents a distinct root-to-leaf path λ_j in \mathcal{T}_A . Consider an edge (v, v') of \mathcal{T}_A with label s_i^ℓ , $1 \leq \ell \leq m$. Then there exists a subset of $I \subset I_i^\ell \subset \{1, \dots, m\}$ representing the set of root-to-leaf paths containing (v, v') , i.e. $\{\lambda \in \Lambda \mid (v, v') \in E(\lambda)\} = \{\lambda_j \mid j \in I \subset I_i^\ell\}$. The number of these paths, $|I|$ is exactly c_j^ℓ for any $j \in I$ since c_j^ℓ describes the number of repeated columns of the first ℓ partitions each of which is a distinct column of A . Thus the number of terms in the summation $t(v') = \sum_{\lambda \in \Lambda_v} \prod_{s \in \theta(\lambda)} s = \sum_{\lambda_j \mid j \in I} \prod_{s \in \theta(\lambda_j)} s$ is c_j^ℓ . In other words if we denote $\sigma(t(v'))$

Theorem 4.12. *For any multipartition matrix A satisfying the GRIP, the associated tree \mathcal{T}_A is a stratified and balanced staged tree.*

Proof. Consider the matrix \bar{A} obtained by removing the repeated columns of A . The associated tree $\mathcal{T}_{\bar{A}}$ is equivalent to \mathcal{T}_A , and hence is a stratified staged tree. Since there are no repeated columns, we have that the toric model for the staged tree \mathcal{T}_A is given by the multipartition model of \bar{A} , i.e. $\widetilde{\mathcal{M}}_{\mathcal{T}_A} = \mathcal{M}_{\bar{A}}$.

By Proposition 4.5, the staged tree model $\mathcal{M}_{\mathcal{T}_A}$ is equal to $\mathcal{M}_{\bar{A}}$ if and only if \mathcal{T}_A is a balanced staged tree. By Proposition 3.26 the MLE of $\mathcal{M}_{\bar{A}}$ is equal to the MLE of the staged tree model \mathcal{T}_A (Remark 4.11). Since the maximum likelihood estimator also parameterizes the model, this implies that $\widetilde{\mathcal{M}}_{\mathcal{T}_A} = \mathcal{M}_{\bar{A}} = \mathcal{M}_{\mathcal{T}_A}$. Hence by Proposition 4.5, \mathcal{T}_A is balanced. \square

Putting the previous results together, we see that a multipartition matrix without repeated columns satisfies the GRIP if and only if the associated tree is a stratified and balanced staged tree.

5. Discussion

In this work we explored the relationship between multipartition matrices, their associated multipartition model, and the IPS algorithm. In particular we showed equivalences between the GRIP and other conditions from disparate areas of algebraic statistics and examined their implications for exact convergence under the IPS algorithm. We conclude with a discussion of some natural questions that arise from these results.

Question 1. For a multipartition matrix A with k partitions, does the requirement that the IPS algorithm produces the MLE for $A^{1,\dots,\ell}$ at each step imply that A must satisfy the GRIP?

In Remark 3.24, we claim that IPS producing the MLE in 2 steps for a multipartition matrix with 2 partitions is enough to guarantee that the GRIP is satisfied. While the GRIP is sufficient for the algorithm to produce the MLE for $A^{1,\dots,\ell}$ at each step, it is possible that the reverse logical direction holds.

Question 2. Let \mathcal{M} be a log-linear multipartition model with rational MLE. Does there always exist a multipartition matrix A representing \mathcal{M} that satisfies the GRIP?

It is clear that the matrix representation of a particular model affects the IPS algorithm. In [11], the authors show sufficient conditions for 2-way quasi-independence models to have rational MLE. Since k -way quasi-independence models are just multipartition models without repeated columns, it would be interesting to see if these results can be used to show that every such matrix has a representation satisfying the GRIP. Indeed our preliminary computations indicate that this is true.

Furthermore, we only address one-cycle convergence in this work. It has been conjectured in [26] that all decomposable hierarchical models converge to the MLE in two cycles regardless of the order of the partition matrices. However, the author of [41] presents a counter-example to that conjecture. Whether a modified version of the RIP or GRIP can guarantee convergence in two or more cycles remains an open question.

Finally there are many results and questions related to the convergence of the generalized IPS algorithm on log-linear models and its connections to tools from algebraic statistics (see [2, Sec. 5] and [16, Sec. 7.3]). Our work falls adjacent to this line of inquiry; it would be worthwhile to investigate connections between the generalized running intersection property and recent advances in this area.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

We would like to thank Bernd Sturmfels, Nihat Ay, Jiří Vomlel, Eliana Duarte, Orlando Marigliano, Seth Sullivan and Carlos Améndola for helpful discussions and suggestions regarding this project. We also thank the anonymous reviewers for their careful reading of the manuscript and their very helpful feedback. Jane Coons was partially supported by the US National Science Foundation (DGE 1746939) and the L'Oreal-UNESCO For Women in Science UK and Ireland Rising Talent Programme. Carlotta Langer was supported by the priority programme “The Active Self” (SPP 2134) of the Deutsche Forschungsgemeinschaft. The authors thank the Max Planck Institute for Mathematics in the Sciences for facilitating a visit to the institute to conduct this research.

Appendix A. Staged tree models

We introduce staged tree models following the presentation in [4]. Staged tree models are a type of probability tree model, which are graphical models that encode conditional independence statements between events. For a detailed introduction see [10]. Every discrete Bayesian network and decomposable graphical model is a staged tree model [19]. In this section, we define the notion of balanced and stratified staged trees, and we describe the correspondence between balanced, stratified staged tree models and multipartition models with a rational maximum likelihood estimator.

Let $\mathcal{T} = (V, E)$ be a tree with vertex set V and edge set E . A directed tree is *rooted* if there exists a vertex v_0 , which we called the *root*, such that the unique path from v_0 to each leaf in \mathcal{T} is a directed path.

For a particular vertex $v \in V$, we define the set of *children* of v as the set $\text{ch}(v) = \{w \mid (v, w) \in E\}$. Then the set of outgoing edges from v will be denoted as the set $E(v) = \{(v, w) \mid w \in \text{ch}(v)\}$, and v is a leaf if $E(v) = \emptyset$. A *labeling* for a tree is a surjective mapping $\theta : E \rightarrow S$ to a label set S . From now on we assume that (\mathcal{T}, θ) is a rooted tree graph with labeling θ .

Definition A.1. For any vertex $v \in V$, the set of labels of its outgoing edges is the **floret** of v , denoted $F_v := \{\theta(e) \mid e \in E(v)\}$. The set of florets of (\mathcal{T}, θ) is denoted $F_{\mathcal{T}}$.

Definition A.2. The tree graph (\mathcal{T}, θ) is a **staged tree** if the following are satisfied:

- For each $v \in V$, no two edges in $E(v)$ have the same labeling.
- For any two vertices $v, w \in V$, the florets F_v and F_w are equal or disjoint. When $F_v = F_w$ we say that the vertices v and w are in the same *stage*.

A *staged tree model* is associated to a staged tree (\mathcal{T}, θ) as follows. Let Λ denote the set of root-to-leaf paths in \mathcal{T} and suppose $|\Lambda| = m$. For $\lambda \in \Lambda$ let $E(\lambda)$ be the set of edges in Λ and define the multiset $\theta(\lambda) := \{\{s_i \mid \theta(e) = s_i \text{ for } e \in E(\lambda)\}\}$. This is a multiset because several edges on the same path can have the same labeling; the multiplicity of s_i in this multiset is the number of edges in $E(\lambda)$ with $\theta(e) = s_i$. The paths λ parameterize the model $\mathcal{M}_{\mathcal{T}}$ defined below.

Definition A.3. For a staged tree (\mathcal{T}, θ) with label set $S = \{s_1, \dots, s_r\}$, the **staged tree model** $\mathcal{M}_{\mathcal{T}}$ is the image of the map $\phi_{\mathcal{T}} : \Theta \rightarrow \Delta_{m-1}$ defined by

$$\phi_{\mathcal{T}}(s_1, \dots, s_r) = \left(\prod_{s_i \in \theta(\lambda)} s_i \right)_{\lambda \in \Lambda} \quad (14)$$

where the parameter space is

$$\Theta := \left\{ (s_1, \dots, s_r) \in (0, 1]^r \mid \sum_{s_i \in F} s_i = 1 \text{ for all florets } F \in F_{\mathcal{T}} \right\}.$$

Example A.4. Consider the staged tree (\mathcal{T}, θ) in Fig. 8.

In this diagram, the vertices in the same stage have the same coloring. We immediately see that each floret has equal or disjoint label set. We refer to the vertices of the tree by the edge labels defining the path to the vertex, i.e. the blue vertices are $v_{s_0 t_0}$ and $v_{s_1 t_0}$. If v_0 is the root vertex then the florets of the tree are

$$\begin{aligned} F_{v_0}, \quad F_{v_{s_0}} &= F_{v_{s_1}}, \\ F_{v_{s_0 t_0}} = F_{v_{s_1 t_0}}, \quad F_{v_{s_0 t_1}} &= F_{v_{s_1 t_1}}. \end{aligned}$$


The associated staged tree model $\mathcal{M}_{\mathcal{T}}$ is image of the map $\phi_{\mathcal{T}} : \Theta \rightarrow \Delta_9$ where

$$\begin{aligned} \phi_{\mathcal{T}}(s_0, s_1, t_0, t_1, r_0, r_1, r_2, r_3, r_4) = & (s_0 t_0 r_0, s_0 t_0 r_1, s_0 t_0 r_2, s_0 t_1 r_3, s_0 t_1 r_4, \\ & s_1 t_0 r_0, s_1 t_0 r_1, s_1 t_0 r_2, s_1 t_1 r_3, s_1 t_1 r_4), \end{aligned}$$

and

$$\begin{aligned} \Theta = \{ (s_0, s_1, t_0, t_1, r_0, r_1, r_2, r_3, r_4) \in (0, 1]^9 \mid & s_0 + s_1 = 1, t_0 + t_1 = 1, \\ & r_0 + r_1 + r_2 = 1, r_3 + r_4 = 1 \}. \end{aligned}$$

Remark A.5. We use the slightly more general definitions of staged tree and staged tree models introduced in [4] that allow for staged trees with singleton florets. The authors show that the staged tree obtained by contracting these florets, which satisfies the definition used in previous literature (i.e. [10,24]), results in the same staged tree model [4, Lemma 5.11].

Fig. 8. Staged tree associated to the model $\mathcal{M}_{\mathcal{T}}$.

The map $\phi_{\mathcal{T}}$ in Definition A.3 also defines a toric model if one “forgets” the conditions imposed on the parameter space by the florets. We denote the map obtained by extending $\phi_{\mathcal{T}}$ to all of \mathbb{R}^r as $\tilde{\phi}_{\mathcal{T}} : \mathbb{R}^r \rightarrow \mathbb{R}^m$ and the corresponding toric model as $\tilde{\mathcal{M}}_{\mathcal{T}}$. Immediately it follows that $\mathcal{M}_{\mathcal{T}} \subset \tilde{\mathcal{M}}_{\mathcal{T}}$, but in general it is not true that $\mathcal{M}_{\mathcal{T}} = \tilde{\mathcal{M}}_{\mathcal{T}}$. In [17] the authors found necessary and sufficient graphical conditions on (\mathcal{T}, θ) such that $\mathcal{M}_{\mathcal{T}} = \tilde{\mathcal{M}}_{\mathcal{T}}$. We detail those conditions in Section 4.2.

Definition A.6. For any vertex in a staged tree (\mathcal{T}, θ) , we define the **level** $\ell(v)$ as the number of edges in the unique root-to- v path. A staged tree (\mathcal{T}, θ) is **stratified** if all its leaves have the same level and for any two vertices v, w such that $F_v = F_w$, we have $\ell(v) = \ell(w)$. We say that the **level** of a stratified staged tree is the level of its leaves.

There is a natural partition of the label set of a stratified staged tree by level. For a stratified staged tree of level ℓ , we write $S = \{s_1^1, s_2^1, \dots, s_{n_1}^1, s_1^2, \dots, s_{n_\ell}^\ell\}$ where n_k is the number of distinct edge labels for edges emanating from vertices with level k . Thus the monomial map $\phi_{\mathcal{T}}$ for a stratified staged tree of level ℓ is multilinear and homogeneous, which implies that the associated matrix is a multipartition matrix with ℓ partitions. We denote the matrix associated with $\phi_{\mathcal{T}}$ as $A_{\mathcal{T}}$.

Similarly with multipartition matrix A , we can associate a directed, rooted tree with a labeling induced by the rows of A . We define the label set associated with A as $S_A = \{s_1^1, \dots, s_{n_1}^1, \dots, s_1^\ell, \dots, s_{n_\ell}^\ell\}$ associated to the row vectors $a_1^1, a_2^1, \dots, a_{n_\ell}^\ell$.

Starting with a root vertex, we add a labeled edge s_i^1 for each associated row vector in A^1 . To each edge s_i^1 , we attach an edge labeled s_j^2 to the outgoing vertex for each s_j^2 is the label set S_A such that $I_i^1 \cap I_j^2 \neq \emptyset$. The resulting directed, rooted, and labeled tree graph is $\mathcal{T}_{A^{1,2}}$ where $A^{1,2}$ is the matrix with partitions A_1, A_2 . The paths of $\mathcal{T}_{A^{1,2}}$ exactly correspond to the distinct columns of $A^{1,2}$. We note that if $A^{1,2}$ has two columns that are equal to one another, these correspond to the same path in $\mathcal{T}_{A^{1,2}}$. We repeat this process for each A^k . In this way we define \mathcal{T}_A as the directed, rooted tree graph with labeling such that the paths in $\mathcal{T}_A^{(\ell)}$ correspond to the distinct columns of $A^{1, \dots, \ell}$. For a visual example of this process, consider Example 4.7 where we draw the staged tree associated with a multipartition matrix A .

We have seen that staged trees are useful in describing certain properties of multipartition matrices. In particular, the GRIP condition, under which multipartition matrices produce the MLE after one IPS cycle, requires that \mathcal{T}_A is a staged tree. We have also shown in Section 4.2 that the staged tree models for which $\mathcal{M}_{\mathcal{T}} = \tilde{\mathcal{M}}_{\mathcal{T}}$ also produce the MLE after one cycle of IPS.

Appendix B. Glossary

Multipartition models

ϕ_A	Polynomial map defined by the matrix A
Δ_{m-1}	$(m-1)$ -dimensional probability simplex
\mathcal{M}_A	Log-linear model associated to the integer matrix A
$A^{1, \dots, k}$	Multipartition matrix consisting of the stacked partition matrices A^1, \dots, A^k
a_i^ℓ	The i -th row of the ℓ -th partition of a multipartition matrix A
I_i^ℓ	The support of the row a_i^ℓ of a multipartition matrix A
$S(\ell, j)$	Index such that $j \in I_{S(\ell, j)}^\ell$
c_j	Column weight of the j -th column of a matrix A
$a_i^\ell(\bar{u})$	Row of a multipartition matrix multiplied with the normalized data
C_μ^ℓ	Connection ratio for a_μ^ℓ
β	Number of distinct columns of a multipartition matrix
λ	Labeling of the columns of a multipartition matrix marking the distinct columns
$\Psi_{n=1}^\ell A^n$	Matrix with the j th column $e_{\lambda(j)}$
\mathcal{F}_t^C	Floret of C with respect to the matrix B with BC satisfying the floret condition

$t(\mu, \nu)$	Index of the floret such that $a_\mu^B \in \mathcal{F}_{t(\mu, \nu)}^B$ and $a_\nu^C \in \mathcal{F}_{t(\mu, \nu)}^C$
$\omega_{\mu\nu}$	Number of columns of BC of the form $[e_\mu e_\nu]^T$
$B \bowtie C$	Matrix in which the columns are indicator vectors for the florets in BC
\underline{A}	Compression of the matrix A

Iterative proportional scaling

u	Vector of counts of the observed data
\bar{u}	Normalized data
$p^*(\bar{u})$	Maximum likelihood estimate of \bar{u}

Staged tree models

\mathcal{T}	Tree with vertex set V and edge set E
(\mathcal{T}, θ)	Labeled tree graph with labeling $\theta : E \rightarrow S$
Λ	Set of root-to-leaf paths, λ , in the tree \mathcal{T}
$\mathcal{M}_{\mathcal{T}}$	Staged tree model
$\phi_{\mathcal{T}}$	Polynomial map defined by the tree \mathcal{T}

References

- [1] C. Améndola, N. Bliss, I. Burke, C.R. Gibbons, M. Helmer, S. Hoşten, E.D. Nash, J.I. Rodriguez, D. Smolkin, The maximum likelihood degree of toric varieties, *J. Symb. Comput.* 92 (2019) 222–242.
- [2] C. Améndola, K. Kohn, P. Reichenbach, A. Seigal, Toric invariant theory for maximum likelihood estimation in log-linear models, *Algebraic Stat.* 12 (2021) 2.
- [3] C. Améndola, D. Kosta, K. Kubjas, Maximum likelihood estimation of toric Fano varieties, *Algebraic Stat.* 11 (1) (2020) 5–30.
- [4] L. Ananiadi, E. Duarte, Gröbner bases for staged trees, *Algebraic Stat.* 12 (1) (2021) 1–20.
- [5] N. Ay, J. Jost, H.V. Lê, L. Schwachhöfer, *Information Geometry*, vol. 64, Springer International Publishing, 2017.
- [6] R.J. Berman, The Sinkhorn algorithm, parabolic optimal transport and geometric Monge–Ampère equations, *Numer. Math.* 145 (2020) 771–836.
- [7] J.B. Brown, P.J. Chase, A.O. Pittenger, Order independence and factor convergence in iterative scaling, *Linear Algebra Appl.* 190 (1993) 1–38.
- [8] F.F. Catanese, S. Hoşten, A. Khetan, B. Sturmfels, The maximum likelihood degree, *Am. J. Math.* 128 (3) (2006) 671–697.
- [9] S. Chang, E. Pierson, P. Koh, J. Geradin, B. Redbird, D. Grusky, J. Leskovec, Mobility network models of COVID-19 explain inequities and inform reopening, *Nature* 589 (2021) 82–87.
- [10] R.A. Collazo, C. Görgen, J.Q. Smith, *Chain Event Graphs*, CRC Press, 2018.
- [11] J.I. Coons, S. Sullivant, Quasi-independence models with rational maximum likelihood estimator, *J. Symb. Comput.* 104 (2021) 917–941.
- [12] I. Csiszár, I-divergence geometry of probability distributions and minimization problems, *Ann. Probab.* 3 (1975).
- [13] I. Csiszár, P.C. Shields, *Information Theory and Statistics: A Tutorial*, Foundations and Trends in Communications and Information Theory, Now Publishers Inc, 2004.
- [14] J. Darroch, D. Ratcliff, Generalized iterative scaling for log-linear models, *Ann. Math. Stat.* (1972).
- [15] W.E. Deming, F.F. Stephan, On a least squares adjustment of a sampled frequency table when the expected marginal totals are known, *Ann. Math. Stat.* (1940).
- [16] M. Drton, B. Sturmfels, S. Sullivant, *Lectures on Algebraic Statistics*, vol. 39, Springer Science & Business Media, 2008.
- [17] E. Duarte, C. Görgen, Equations defining probability tree models, *J. Symb. Comput.* 99 (2020) 127–146.
- [18] E. Duarte, O. Marigliano, B. Sturmfels, Discrete statistical models with rational maximum likelihood estimator, *Bernoulli* 27 (2021) 1.
- [19] E. Duarte, L. Solus, A new characterization of discrete decomposable models, *arXiv:2105.05907*, 2021.
- [20] Y. Endo, A. Takemura, Iterative proportional scaling via decomposable submodels for contingency tables, *Comput. Stat. Data Anal.* 53 (4) (2009) 966–978.
- [21] A. Engström, T. Kahle, S. Sullivant, Multigraded commutative algebra of graph decompositions, *J. Algebraic Comb.* 39 (2) (2014) 335–372.
- [22] D. Geiger, C. Meek, B. Sturmfels, On the toric algebra of graphical models, *Ann. Stat.* 34 (3) (2006) 1463–1492.
- [23] C. Görgen, A. Bigatti, E. Riccomagno, J.Q. Smith, Discovery of statistical equivalence classes using computer algebra, *Int. J. Approx. Reason.* 95 (2018) 167–184.
- [24] C. Görgen, J.Q. Smith, Equivalence classes of staged trees, *Bernoulli* 24 (4A) (2018) 2676–2692.
- [25] C. Görgen, A. Maraj, L. Nicklasson, Staged tree models with toric structure, *J. Symb. Comput.* 113 (2022) 242–268.
- [26] S.J. Haberman, *The Analysis of Frequency Data*, University of Chicago Press, 1974.
- [27] J. Harris, H. Stoecker, *The Handbook of Mathematics and Computational Science*, Springer, 1998.
- [28] M.E. Hauer, R.D. Hardy, D.R. Mishra, J.S. Pippin, No landward movement: examining 80 years of population migration and shoreline change in Louisiana, *Popul. Environ.* 40 (2019) 369–387.
- [29] J. Huh, The maximum likelihood degree of a very affine variety, *Compos. Math.* 149 (8) (2013) 1245–1266.
- [30] J. Huh, Varieties with maximum likelihood degree one, *J. Algebraic Stat.* 5 (1) (2014) 1–17.
- [31] J. Huh, B. Sturmfels, Likelihood Geometry, In *Lecture Notes in Mathematics*, Springer International Publishing, 2014, pp. 63–117.
- [32] E.T. Jaynes, *Information Theory and Statistical Mechanics*, *Phys. Rev.* 106 (1957).
- [33] A. Klenke, *Probability Theory: A Comprehensive Course*, Springer, 2007.
- [34] R. Millar, Maximum likelihood estimation and inference: with examples in R, SAS and ADMB, Wiley, 2011.
- [35] L. Pachter, B. Sturmfels (Eds.), *Algebraic Statistics for Computational Biology*, Cambridge University Press, New York, 2005, ch. § 1.2.
- [36] G. Peyré, M. Cuturi, Computational Optimal Transport, *Found. Trends Mach. Learn.* 11 (5–6) (2019) 355–607.
- [37] J. Rauh, Optimally approximating exponential families, *Kybernetika* 49 (2011).
- [38] J.Q. Smith, P.E. Anderson, Conditional independence and chain event graphs, *Artif. Intell.* 172 (1) (2008) 42–68.
- [39] S. Sullivant, Toric fiber products, *J. Algebra* 316 (2) (2007) 560–577.
- [40] S. Sullivant, *Algebraic Statistics*, vol. 194, American Mathematical Soc., 2018.
- [41] J. Vomlel, *Methods of Probabilistic Knowledge Integration*, Dissertation, Czech Technical University, Faculty of Electrical Engineering, 1999.
- [42] P.-F. Xu, J. Sun, N. Shan, Local computations of the iterative proportional scaling procedure for hierarchical models, *Comput. Stat. Data Anal.* 95 (2016) 17–23.
- [43] C.T. Yiannoutsos, P.K. Halverson, N. Menachemi, Bayesian estimation of SARS-CoV-2 prevalence in Indiana by random testing, *Proc. Natl. Acad. Sci.* 118 (2021) 5.