

ITERATIVE SCALING IN CURVED EXPONENTIAL FAMILIES

Anna Klimova

Institute of Science and Technology (IST) Austria,
Klosterneuburg, Austria

Tamás Rudas

Eötvös Loránd University,
Budapest, Hungary

Abstract

The paper describes a generalized iterative proportional fitting procedure which can be used for maximum likelihood estimation in a special class of the general log-linear model. The models in this class, called relational, apply to multivariate discrete sample spaces which do not necessarily have a Cartesian product structure and may not contain an overall effect. When applied to the cell probabilities, the models without the overall effect are curved exponential families and the values of the sufficient statistics are reproduced by the MLE only up to a constant of proportionality. The paper shows that Iterative Proportional Fitting, Generalized Iterative Scaling and Improved Iterative Scaling, fail to work for such models. The algorithm proposed here is based on iterated Bregman projections. As a by-product, estimates of the multiplicative parameters are also obtained.

KEY WORDS: Bregman divergence, contingency tables, curved exponential family, generalized odds ratio, iterative proportional fitting, maximum likelihood estimate, overall effect, relational model

Introduction

This paper deals with variants of the general log-linear model

$$\log \boldsymbol{\delta} = \mathbf{A}'\boldsymbol{\beta}, \quad (1)$$

where $\boldsymbol{\delta}$ denotes a vector of probabilities (appropriate in the case of multinomial sampling) or a vector of intensities (appropriate in the case of Poisson sampling) and the model matrix \mathbf{A} has non-negative integer entries.

When the sample space is the Cartesian product of the ranges of categorical variables and the rows of \mathbf{A} are indicators of cylinder sets of the sample space, (1) specifies a conventional log-linear model [cf. Bishop, Fienberg, and Holland, 1975]. When the sample space does not have a Cartesian product structure and \mathbf{A} is the indicator matrix of arbitrary subsets of the cells in the sample space, (1) defines a relational model [Klimova, Rudas, and Dobra, 2012].

Models that associate parameters with subsets of cells appear, for instance, in areas of machine learning which deal with feature selection. Features are characteristics of objects, and a subset of cells comprises the objects that possess a particular feature. The goal of the analysis is to choose features so that a Markov field based on them approximates the observed distribution well, see, e.g., Lafferty, Della Pietra, and Della Pietra [1997], Lafferty, McCallum, and Pereira [2001], Malouf [2002], Huang, Hsieh, Chang, and Lin [2010]. Maximum entropy models used in machine learning, see, e.g., Lafferty et al. [2001], are special cases of relational models. Feature selection techniques are used, among others, in text processing [cf. McCallum, Freitag, and Pereira, 2000], in computer tomography [cf. O’Sullivan and Benac, 2007], and in the analysis of social mobility [cf. Klimova and Rudas, 2012].

Many problems in feature selection allow for the existence of unaffected cases, i.e., objects who do not possess any of the characteristics of interest. However, in some contexts of feature selection, such an assumption may not be feasible. For example, in market basket analysis [cf. Brin, Motwani, and Silverstein, 1997, Wu, Barbara, and Ye, 2003], where records of purchases are analyzed to reveal patterns of associations among the different goods bought, each purchase consists of one item, at least. Similarly, a registry of congenital abnormalities [cf. Källén et al., 1999, Boyd et al., 2011, Copeland and Kirby, 2007] lists only affected newborns. While in the case of birth defects, unaffected newborns exist and their total number may be known, in market basket analysis there is no purchase with nothing bought. In such problems, testing hypotheses of association between features, e.g., independence, cannot be performed using conventional log-linear models, but may be done within the relational model framework [Klimova et al., 2012].

In Section 1, the formal definition and the main properties of relational models are reviewed. The characteristics of model (1) are affected by the presence or absence of the vector of 1's, denoted in the sequel as $\mathbf{1}$, in the row space $R(\mathbf{A})$ of the model matrix. If $\mathbf{1} \in R(\mathbf{A})$, there exists a parameterization of the model in which one of the parameters appears in every cell; such a model is said to have the overall effect. Relational models for probabilities with the overall effect and relational models for intensities are regular exponential families, and standard results about the MLE in such families apply. Relational models for probabilities without the overall effect are curved exponential families, and some properties of the MLE are fundamentally different.

In Section 2, three iterative scaling algorithms used for models of type (1) are reviewed and their applicability for determining the MLE under relational models is investigated. The iterative proportional fitting (IPF) procedure [Deming and Stephan, 1940] is employed for conventional log-linear models [cf. Bishop et al., 1975] and can be easily modified to suit relational models with the overall effect. Generalized Iterative Scaling (GIS), proposed by Darroch and Ratcliff [1972], is used in feature selection and works under an assumption which is sometimes called “a constant sum of features”. It is shown, that this assumption is equivalent to the presence of the overall effect in the model, and thus GIS can only be applied to relational models with the overall effect. Improved Iterative Scaling (IIS) was proposed by Della Pietra, Della Pietra, and Lafferty [1997] as a generalization of GIS which does not rely on the assumption of constant sum of features. One might expect that IIS could be used for relational models without the overall effect, but it is shown that this is not the case.

The main contribution of the paper, described in Section 3, is an iterative proportional fitting procedure that generalizes the traditional IPF and can be used for all relational models, with or without the overall effect. The algorithm constructs a sequence of projections on convex sets, and, while for the traditional IPF procedure the projections minimize the Kullback-Leibler distance, the projections for the new algorithm minimize the Bregman distance [Bregman, 1967]. Iterative proportional fitting is typically used for estimating the cell parameters and, to the best knowledge of the authors, it is not used to estimate the model parameters. It is also shown here, that the estimates of the model parameters can be found as a by-product of the IPF procedure. A numerical variant of the generalized IPF is also described, and it is proven that its output approximates the true limit with any specified precision.

1 Preliminaries

Let Y_1, \dots, Y_K be random variables taking values in finite sets $\mathcal{Y}_1, \dots, \mathcal{Y}_K$, respectively. A combination of values $(y_1, y_2, \dots, y_K) \in \mathcal{Y}_1 \times \dots \times \mathcal{Y}_K$ is called a cell, and a sequence \mathcal{I} of cells is a table. The population distribution is parameterized by $\boldsymbol{\delta} = \{\delta(i), \text{ for } i \in \mathcal{I}\}$. Depending on whether the population is characterized by probabilities or by intensities, $\delta(i) \equiv p(i) \in (0, 1)$, with $\sum_{i \in \mathcal{I}} p(i) = 1$, or $\delta(i) \equiv \lambda(i) > 0$, for all $i \in \mathcal{I}$. The distinction between data collected through multinomial or Poisson sampling procedures is necessary, because maximum likelihood estimates under the two sampling schemes may have fundamentally different properties under the models considered here.

Let \mathcal{P} be the set of positive distributions, parameterized by $\boldsymbol{\delta}$, on \mathcal{I} . A relational model is generated by a class $\mathbf{S} = \{S_1, \dots, S_J\}$ of non-empty subsets of the table \mathcal{I} . A distribution $P_{\boldsymbol{\delta}} \in \mathcal{P}$ is in the model if and only if

$$\log \boldsymbol{\delta} = \mathbf{A}'\boldsymbol{\beta}, \text{ for some } \boldsymbol{\beta} \in \mathbb{R}^J. \quad (2)$$

Here, the components of $\boldsymbol{\beta}$ are the log-linear parameters of the model and the rows of the model matrix \mathbf{A} are indicators of the generating subsets.

A relational model is also an exponential (multiplicative) family

$$\mathcal{M}(\mathbf{A}) = \{P_{\boldsymbol{\delta}} \in \mathcal{P} : \boldsymbol{\delta} = \boldsymbol{\theta}^{\mathbf{A}'}, \text{ for some } \boldsymbol{\theta} \in \mathbb{R}_{>0}^J\}, \quad (3)$$

where, for every $i \in \mathcal{I}$, $\delta(i) = \prod_{j=1}^J \theta_j^{a_{ji}}$, and the components of $\boldsymbol{\theta}$ are the multiplicative parameters.

The overall effect, present in every cell, plays the role of the normalizing constant, say β_0 , which is often included in exponential family models:

$$\boldsymbol{\delta} = \frac{1}{\beta_0} \exp \{\mathbf{A}'\boldsymbol{\beta}\}.$$

Such models can be re-written in the form (2) by adding a row of 1's to \mathbf{A} . Therefore, relational models do not assume that there is no normalizing constant, they only allow for this possibility.

A dual representation of a relational model is obtained using an integer kernel basis matrix \mathbf{D} [Klimova et al., 2012], whose rows are a basis of $\text{Ker}(\mathbf{A})$. Thus, (3) can be re-written as

$$\mathcal{M}(\mathbf{A}) = \{P_{\boldsymbol{\delta}} \in \mathcal{P} : \mathbf{D} \log \boldsymbol{\delta} = \mathbf{0}\}. \quad (4)$$

The dual representation says that the generalized odds ratios,

$$\mathcal{OR}_d = \boldsymbol{\delta}^{d^+} / \boldsymbol{\delta}^{d^-}, \quad (5)$$

associated with the rows of \mathbf{D} , are all equal to 1, where \mathbf{d}^+ and \mathbf{d}^- denote, respectively, the positive and negative parts of row \mathbf{d} . If the degrees of the numerator and of the denominator in (5) are equal, the odds ratio $\mathcal{OR}_{\mathbf{d}}$ is called homogeneous and, otherwise it is called non-homogeneous. For a relational model with the overall effect, all generalized odds ratios are homogeneous, and it can be shown that any model without the overall effect has a dual representation with exactly one non-homogeneous odds ratio.

Some of the properties of maximum likelihood estimates under relational models are also affected by presence or absence of the overall effect. Let \mathbf{Y} be a random variable that has a distribution parameterized by $\boldsymbol{\delta}$ and let \mathbf{y} be a realization of \mathbf{Y} . Assume that the MLE $\hat{\boldsymbol{\delta}}$, under the relational model generated by \mathbf{A} , exists. The properties of the MLE in $\mathcal{M}(\mathbf{A})$ are best described by using the linear family

$$\mathcal{L}_{\boldsymbol{\delta}}(\mathbf{A}, \mathbf{q}, \gamma) = \{P_{\boldsymbol{\delta}} \in \mathcal{P} : \mathbf{A}\boldsymbol{\delta} = \gamma\mathbf{A}\mathbf{q}\},$$

where

$$\mathbf{q} = \begin{cases} \mathbf{y}, & \text{if } \boldsymbol{\delta} \equiv \boldsymbol{\lambda}, \\ \mathbf{y}/(\mathbf{1}\mathbf{y}), & \text{if } \boldsymbol{\delta} \equiv \mathbf{p}. \end{cases} \quad (6)$$

As was shown by Klimova et al. [2012], when the relational model is a regular exponential family, the distribution parameterized by $\hat{\boldsymbol{\delta}}$ is the unique point in

$$\mathcal{M}(\mathbf{A}) \cap \mathcal{L}_{\hat{\boldsymbol{\delta}}}(\mathbf{A}, \mathbf{q}, 1), \quad (7)$$

and, when the relational model is a curved exponential family, the distribution parameterized by $\hat{\boldsymbol{\delta}}$ is the unique common point in

$$\mathcal{M}(\mathbf{A}) \cap \mathcal{L}_{\hat{\boldsymbol{\delta}}}(\mathbf{A}, \mathbf{q}, \gamma), \quad (8)$$

for the unique $\gamma > 0$ for which $\mathbf{1}\hat{\boldsymbol{\delta}} = 1$. The components of $\mathbf{A}\boldsymbol{\delta}$ are called the subset sums of $P_{\boldsymbol{\delta}}$. The coefficient of proportionality γ is called the adjustment factor. The properties of maximum likelihood estimates under relational models are summarized in Table 1.

*** Table 1 here. ***

In general, there are no closed form expressions for the MLE. Maximum likelihood estimates for the cell parameters can be computed using the Newton-Raphson algorithm or algorithms for convex optimization [cf. Bertsekas, 1999, Aitchison and Silvey, 1960, Evans and Forcina, 2013]. This paper focuses on the iterative scaling approach to computing the MLE, as it is often used in feature selection procedures [cf. Huang et al., 2010].

2 Why IPF, GIS, IIS cannot be used for relational models?

The traditional iterative proportional fitting (IPF) procedure is used for computing maximum likelihood estimates of the cell frequencies or probabilities under log-linear models [cf. Bishop et al., 1975]. The IPF algorithm starts with a contingency table of the same structure as the sample space, with all cell frequencies equal to one: $\boldsymbol{\delta}^{(0)} = \mathbf{1}$. The cell frequencies are then adjusted until the marginal sums $A_j \boldsymbol{\delta}^{(d)}$, where $d + 1 \equiv j \bmod J$, become equal or close enough to the observed values:

$$\delta^{(d+1)}(i) = \delta^{(d)}(i) \left(\frac{A_j \mathbf{q}}{A_j \boldsymbol{\delta}^{(d)}} \right)^{a_{ji}}, \text{ for all } i \in \mathcal{I}. \quad (9)$$

The model structure, expressed in terms of the odds ratios, is preserved during iterations. Given that the MLEs of the cell parameters exist, the sequence $\boldsymbol{\delta}^{(d)}$ converges, as $d \rightarrow \infty$, to the maximum likelihood estimate $\hat{\boldsymbol{\delta}}$. The proof of convergence of IPF relies on a particular order of the subsets. In this order, the cylinder sets associated with the k -th marginal are listed in one block indexed from $j = [k]_s$ to $j = [k]_e$. Then the distribution parameterized by $\delta^{(tJ+[k]_e)}$ is the I-projection of the distribution parameterized by $\delta^{(tJ+[k-1]_e)}$ in $\cap_{j=[k]_s}^{[k]_e} \mathcal{L}_{\boldsymbol{\delta}}(A_j, \mathbf{q}, 1)$, for every non-negative integer t , see Csiszár [1975].

A straightforward modification of IPF, which updates, instead of marginal totals, the subset sums, may be considered for computing the MLEs of the cell parameters under relational models. This will work when the overall effect is present in the model, however, the proof needs modification, because in the actual parameterization, there may be no sets S_j that form a partitioning of the sample space, like $S_{[k]_s}, \dots, S_{[k]_e}$ do, and, therefore, the algorithm may not produce I-projections. When the overall effect is not present, (9) works for intensities, but as the total is not reproduced by the MLE, a new proof is needed, see Section 3.

The generalized iterative scaling (GIS) algorithm was proposed by Darroch and Ratcliff [1972] for maximum likelihood estimation in discrete exponential families of the form $\log \mathbf{p} = \mathbf{A}'\boldsymbol{\beta}$, where \mathbf{A} is a non-negative real matrix whose rows add to the vector of 1's:

$$A_1 + \dots + A_J = \left(\sum_{j=1}^J a_{j1}, \dots, \sum_{j=1}^J a_{j|\mathcal{I}|} \right) = \mathbf{1}. \quad (10)$$

The GIS procedure organizes updating cycles not according to subsets (or marginals), rather according to cells. It performs all adjustments that apply to a cell in one step, and recalculates the multipliers to be used in the next cycle of adjustment only after all cells were

updated:

$$p^{(n+1)}(i) = p^{(n)}(i) \prod_{j=1}^J \left(\frac{A_j \mathbf{q}}{A_j \mathbf{p}^{(n)}} \right)^{a_{ji}}, \quad \text{for all } i \in \mathcal{I}. \quad (11)$$

As $n \rightarrow \infty$, $\mathbf{p}^{(n)}$ converges to the maximum likelihood estimate $\hat{\mathbf{p}}$, see Theorem 1 in Darroch and Ratcliff [1972]. The condition (10) is used in the proof of convergence, but it can be shown that GIS also applies to relational models with a model matrix whose rows sum to a constant vector:

$$A_1 + \dots + A_J = c\mathbf{1}. \quad (12)$$

In fact, (12) and the presence of the overall effect are equivalent:

Proposition 2.1. *For a 0-1 matrix \mathbf{A} , the following two conditions are equivalent:*

- (i) $\mathbf{1} \in R(\mathbf{A})$.
- (ii) *There exists a non-negative matrix $\tilde{\mathbf{A}}$, with $R(\tilde{\mathbf{A}}) = R(\mathbf{A})$, whose rows add to a vector $c\mathbf{1}$, for some $c \in \mathbb{Z}_+$.*

Proof. Suppose (i) holds. Without loss of generality, assume that $A_1 = \mathbf{1}$. The matrix \mathbf{A}^* consisting of rows

$$A_1^* = A_1 - A_2, \quad A_2^* = A_2 - A_3, \quad \dots, \quad A_{J-1}^* = A_{J-1} - A_J, \quad A_J^* = A_J, \quad (13)$$

has the same row space as \mathbf{A} . The rows of \mathbf{A}^* sum to $\mathbf{1}$, but its entries may not be all positive. Let $a_j^* = \min_{i \in \mathcal{I}}(a_{ji}^*)$, for $j = 1, \dots, J$.

The matrix $\tilde{\mathbf{A}}$ with rows $\tilde{A}_j = A_j^* - a_j^*\mathbf{1}$ is nonnegative, has the same row space as \mathbf{A} , and the rows of $\tilde{\mathbf{A}}$ sum to $c\mathbf{1}$, where $c = 1 - a_1^* - a_2^* - \dots - a_J^* \in \mathbb{Z}_+$.

Suppose (ii) holds. Then $\mathbf{1}$ is a linear combination of some vectors in $R(\mathbf{A})$ and thus $\mathbf{1} \in R(\mathbf{A})$. The proof is now complete. \square

In feature selection procedures, before GIS is applied, a model matrix with a non-constant sum of rows is often converted to a matrix with a constant row sum by adding a "slack feature", which balances other rows, to the model [cf. Lafferty et al., 2001, McCallum et al., 2000]. By Proposition 2.1, a model with a slack feature always has the overall effect. Adding a slack feature to a model without the overall effect changes the model.

Another algorithm, Improved Iterative Scaling (IIS), proposed by Della Pietra et al. [1997], claims that it "is an improvement of the Generalized Iterative Scaling algorithm of Darroch and Ratcliff in that it does not require that the features sum to a constant". The IIS algorithm cyclically updates the cells of the table, and each iteration consists of two steps - updating the multipliers and updating the cell probabilities:

1. For each $j = 1, \dots, J$, solve

$$\sum_{i \in \mathcal{I}} a_{ji} p^{(n)}(i) (\zeta_j^{(n)})^{a(i)} = A_j \mathbf{q}, \text{ where } a(i) = \sum_{j=1}^J a_{ji}, \quad (14)$$

to compute $\boldsymbol{\zeta}^{(n)} = (\zeta_1^{(n)}, \dots, \zeta_J^{(n)})'$.

2. Compute

$$p^{(n+1)}(i) = p^{(n)}(i) \prod_{j=1}^J (\zeta_j^{(n)})^{a_{ji}}, \quad (15)$$

and set $n = n + 1$.

In the case when all of $a(i) = \sum_{j=1}^J a_{ji}$ are equal to 1, (15) reduces to the updating step (11) of the GIS procedure.

The literature on the IIS algorithm seems implicit about whether or not the model matrix has a row of 1's. For example, Della Pietra et al. [1997] and Bancarz and Osborne [2002] mention that a normalizing constant is included in the model, however Lafferty et al. [1997] do not explicitly say it, but prove that the sequence $\mathbf{p}^{(n)}$ converges to the MLE. In fact, if the model has the overall effect, it does not matter whether (10) or (12) holds in the actual parameterization. But if the model does not contain the overall effect, IIS may converge to a vector of probabilities that does not sum to 1 and thus can not be the MLE. Normalization of the limit vector would not help: a relational model without the overall effect is not scale invariant [Klimova et al., 2012], that is, a non-normalized vector of probabilities, that has the multiplicative structure prescribed by the model, loses this structure after being normalized. For example, the model of independence between three features A , B , or C is defined by Aitchison and Silvey [1960] as

$$\frac{p_{AB}}{p_A p_B} = 1, \quad \frac{p_{AC}}{p_A p_C} = 1, \quad \frac{p_{BC}}{p_B p_C} = 1, \quad \frac{p_{ABC}}{p_A p_B p_C} = 1. \quad (16)$$

Here p_A , p_B , p_C , p_{AB} , p_{AC} , p_{BC} , p_{ABC} denote positive probabilities of having the corresponding combination of features, and $p_A + p_B + p_C + p_{AB} + p_{AC} + p_{BC} + p_{ABC} = 1$. Let $\mathbf{p} = (p_A, p_B, p_C, p_{AB}, p_{AC}, p_{BC}, p_{ABC})'$. This is a relational model generated by the subsets S_1 (possessing the feature A), S_2 (possessing the feature B), and S_3 (possessing the feature C) and is a variant of independence when no unaffected cases exist. It is defined using non-homogeneous odds ratios in (16) and thus does not have the overall effect. Let $\mathbf{q} = (0.04, 0.04, 0.04, 0.04, 0.04, 0.24, 0.56)'$ be the parameter of the observed probability distribution. The vector sequence $\mathbf{p}^{(n)}$, produced by IIS, converges to $\mathbf{p}^* = (0.3202, 0.4574, 0.4574, 0.1464, 0.1464, 0.2092, 0.0670)'$, with the total $\mathbf{1p}^* = 1.804$. The normalized vector:

$$\mathbf{p}_1^* = \mathbf{p}^* / (\mathbf{1p}^*) = (0.1775, 0.2535, 0.2535, 0.0812, 0.0812, 0.1160, 0.0371)'$$

does not have the multiplicative structure implied by (16).

In order to illustrate that this problem is not isolated, rather quite common, the histogram of sums of the limit vectors obtained for 13352 observed distributions \mathbf{q} , generated with equal spacing on the parameter space, is given in Figure 3.

*** Figure 1 here. ***

In summary, a relational model with the overall effect can be parameterized using a model matrix whose rows do or do not sum to a constant vector. Depending on the parameterization, IPF, GIS or IIS can be used to compute the MLE under such a model. However, a relational model without the overall effect, under any parameterization, will have a model matrix whose rows do not sum to a constant vector, and, if the model is for probabilities, IPF, GIS, IIS do not converge to the MLE.

The next section presents an iterative fitting procedure that converges to the MLE, whether or not the overall effect is present in the relational model.

3 Iterative proportional fitting in relational models

Let \mathbf{Y} be a random variable that has a Poisson distribution parameterized by $\boldsymbol{\delta} \equiv \boldsymbol{\lambda}$ or a multinomial distribution parameterized by N and $\boldsymbol{\delta} \equiv \mathbf{p}$. Consider the relational model generated by a matrix \mathbf{A} , and let \mathbf{D} be a kernel basis matrix of it. Suppose \mathbf{y} is a realization of \mathbf{Y} . Consider any $\gamma > 0$, such that $\mathcal{L}_{\boldsymbol{\delta}}(\mathbf{A}, \mathbf{q}, \gamma)$, with \mathbf{q} given in (6), is not empty.

An iterative proportional fitting algorithm that computes the cell parameters, $\boldsymbol{\delta}_{\gamma}^*$, and the model parameters for the distribution that is the unique common point of the linear family $\mathcal{L}_{\boldsymbol{\delta}}(\mathbf{A}, \mathbf{q}, \gamma)$ and the multiplicative family $\mathcal{M}(\mathbf{A})$ is given next.

IPF(γ) Algorithm:

Set $d = 0$; $\delta_{\gamma}^{(0)}(i) = 1$ for all $i \in \mathcal{I}$; $\theta_{\gamma}^{(0)}(j) = 1$ for all $j = 1, \dots, J$, and proceed as follows.

Step 1: Find $j \in \{1, 2, \dots, J\}$, such that $d + 1 \equiv j \pmod{J}$;

Step 2: Compute

$$\delta_{\gamma}^{(d+1)}(i) = \delta_{\gamma}^{(d)}(i) \left(\gamma \frac{A_j \mathbf{q}}{A_j \boldsymbol{\delta}_{\gamma}^{(d)}} \right)^{a_{ji}} \text{ for all } i \in \mathcal{I}; \quad (17)$$

$$\theta_{\gamma}^{(d+1)}(j) = \theta_{\gamma}^{(d)}(j) \frac{\delta_{\gamma}^{(d+1)}(i^*)}{\delta_{\gamma}^{(d)}(i^*)}, \text{ where } i^* \in \mathcal{I} : a_{ji^*} > 0. \quad (18)$$

Step 3: While $\gamma A_j \mathbf{q} \neq A_j \boldsymbol{\delta}_\gamma^{(d+1)}$ for at least one j , set $d = d + 1$, go to **Step 1**.

Step 4: Set $\boldsymbol{\delta}_\gamma^* = \boldsymbol{\delta}_\gamma^{(d)}$, and finish. \square

The proof of convergence of $\text{IPF}(\gamma)$ is based on showing that $\text{IPF}(\gamma)$ is an instantiation of the algorithm proposed by Bregman [1967] to find a common point of convex sets.

Let $\mathbf{t}, \mathbf{u} \in \mathbb{R}_{>0}^{|\mathcal{I}|}$ and let $D(\mathbf{t}||\mathbf{u})$ denote the Bregman divergence associated with the function $F(\mathbf{x}) = \sum_{i \in \mathcal{I}} x(i) \log x(i)$:

$$D(\mathbf{t}||\mathbf{u}) = \sum_{i \in \mathcal{I}} t(i) \log (t(i)/u(i)) + \left(\sum_{i \in \mathcal{I}} u(i) - \sum_{i \in \mathcal{I}} t(i) \right). \quad (19)$$

If P and Q are probability distributions parameterized by \mathbf{p} and \mathbf{q} respectively, then $D(\mathbf{p}||\mathbf{q}) = I(P||Q)$ is the Kullback-Leibler divergence between P and Q .

Lemma 3.1. *For the sequence $\{\boldsymbol{\delta}_\gamma^{(d)}\}$ of vectors obtained with $\text{IPF}(\gamma)$,*

$$D(\boldsymbol{\delta}_\gamma^{(d+1)}||\boldsymbol{\delta}_\gamma^{(d)}) = \min_{\boldsymbol{\zeta}: P_{\boldsymbol{\zeta}} \in \mathcal{L}_{\boldsymbol{\delta}}(A_j, \mathbf{q}, \gamma)} D(\boldsymbol{\zeta}||\boldsymbol{\delta}_\gamma^{(d)}). \quad (20)$$

Proof. The function $D(\boldsymbol{\zeta}||\boldsymbol{\delta}_\gamma^{(d)})$ is convex with respect to $\boldsymbol{\zeta}$ and, therefore, its minimum, say $\boldsymbol{\zeta}^*$, on the set $\{\boldsymbol{\zeta} : P_{\boldsymbol{\zeta}} \in \mathcal{L}_{\boldsymbol{\delta}}(A_j, \mathbf{q}, \gamma)\}$ exists and is unique.

Setting the derivatives of the Lagrangian

$$\begin{aligned} L &= D(\boldsymbol{\zeta}||\boldsymbol{\delta}_\gamma^{(d)}) - \alpha(A_j \boldsymbol{\zeta} - \gamma A_j \mathbf{q}) \\ &= \sum_{i \in \mathcal{I}} \zeta(i) \log (\zeta(i)/\delta_\gamma^{(d)}(i)) + \left(\sum_{i \in \mathcal{I}} \delta_\gamma^{(d)}(i) - \sum_{i \in \mathcal{I}} \zeta(i) \right) - \alpha(A_j \boldsymbol{\zeta} - \gamma A_j \mathbf{q}) \end{aligned}$$

equal to zero, one obtains that $\boldsymbol{\zeta}^*$ is the unique solution to the equations

$$\begin{aligned} \log (\zeta^*(i)/\delta_\gamma^{(d)}(i)) &= \alpha a_{ji}, \quad i \in \mathcal{I}, \\ A_j \boldsymbol{\zeta}^* &= \gamma A_j \mathbf{q}. \end{aligned} \quad (21)$$

Since $\boldsymbol{\delta}_\gamma^{(d+1)}$, given in (17), satisfies (21) with $\alpha = \log \gamma \frac{A_j \mathbf{q}}{A_j \boldsymbol{\delta}_\gamma^{(d)}}$, it is equal to $\boldsymbol{\zeta}^*$. \square

Lemma 3.1 implies that, for each $d \geq 0$, the distribution parameterized by $\boldsymbol{\delta}_\gamma^{(d+1)}$ is the D-projection of the distribution parameterized by $\boldsymbol{\delta}_\gamma^{(d)}$ on the set $\mathcal{L}_{\boldsymbol{\delta}}(A_j, \mathbf{q}, \gamma)$, that is, $\{\boldsymbol{\delta}_\gamma^{(d)}\}$ is a relaxation sequence with respect to Bregman divergence. Note that the distribution parameterized by $\boldsymbol{\delta}_\gamma^{(d+1)}$ does not necessarily have the same total as the distribution parameterized by $\boldsymbol{\delta}_\gamma^{(d)}$ did and, even if $\boldsymbol{\delta} \equiv \mathbf{p}$, is not necessarily a probability distribution.

It is proved in the following theorem that $\text{IPF}(\gamma)$ converges to the parameter $\boldsymbol{\delta}_\gamma^*$ of the unique common point of the linear family $\mathcal{L}_{\boldsymbol{\delta}}(\mathbf{A}, \mathbf{q}, \gamma)$ and the multiplicative family $\mathcal{M}(\mathbf{A})$.

Theorem 3.2. *The sequence $\boldsymbol{\delta}_\gamma^{(d)}$, obtained from $\text{IPF}(\gamma)$, converges, as $d \rightarrow \infty$, and the limit $\boldsymbol{\delta}_\gamma^*$ satisfies*

- (i) $\mathbf{A}\boldsymbol{\delta}_\gamma^* = \gamma\mathbf{A}\mathbf{q}$,
- (ii) $\mathbf{D}\log \boldsymbol{\delta}_\gamma^* = \mathbf{0}$.

Proof. (i) Since $\{\boldsymbol{\delta}_\gamma^{(d)}\}$ is a relaxation sequence with respect to the function $D(\mathbf{t}||\mathbf{u})$, it converges to a $\boldsymbol{\delta}_\gamma^*$ [Bregman, 1967, Theorem 1] which belongs to the set

$$\cap_{j=1}^J \{\boldsymbol{\delta}_\gamma : A_j \boldsymbol{\delta}_\gamma = \gamma A_j \mathbf{q}\},$$

and thus

$$\mathbf{A}\boldsymbol{\delta}_\gamma^* = \gamma\mathbf{A}\mathbf{q}.$$

(ii) $\text{IPF}(\gamma)$ multiplies the current value in each cell that belongs to a subset S_j by the ratio of the desired subset sum to the actual one. This transformation leaves the values of the generalized odds ratios unchanged. The formal argument proceeds by induction. Since $\delta_\gamma^{(0)}(i) = 1$, for all $i \in \mathcal{I}$, $\mathbf{D}\log \boldsymbol{\delta}_\gamma^{(0)} = \mathbf{0}$, and the statement holds for $d = 0$. Assume that $\mathbf{D}\log \boldsymbol{\delta}_\gamma^{(d)} = \mathbf{0}$ for a positive integer d . Set $C_j = \frac{\gamma A_j \boldsymbol{\delta}}{A_j \boldsymbol{\delta}_\gamma^{(d)}}$. Then,

$$\begin{aligned} \mathbf{D}\log \boldsymbol{\delta}_\gamma^{(d+1)} &= \mathbf{D}\log \begin{bmatrix} \delta_\gamma^{(d)}(1) \cdot C_j^{a_{j1}} \\ \delta_\gamma^{(d)}(2) \cdot C_j^{a_{j2}} \\ \vdots \\ \delta_\gamma^{(d)}(|\mathcal{I}|) \cdot C_j^{a_{j|\mathcal{I}|}} \end{bmatrix} = \mathbf{D} \begin{bmatrix} \log \delta_\gamma^{(d)}(1) + a_{j1} \log C_j \\ \log \delta_\gamma^{(d)}(2) + a_{j2} \log C_j \\ \vdots \\ \log \delta_\gamma^{(d)}(|\mathcal{I}|) + a_{j|\mathcal{I}|} \log C_j \end{bmatrix} \\ &= \mathbf{D}\log \boldsymbol{\delta}_\gamma^{(d)} + \log C_j \mathbf{D}A'_j = \mathbf{0}, \end{aligned}$$

as \mathbf{D} is a kernel basis matrix and thus $\mathbf{D}A'_j = \mathbf{0}$. Therefore, $\mathbf{D}\log \boldsymbol{\delta}_\gamma^{(d)} = \mathbf{0}$ for all $d = 0, 1, 2, \dots$. Finally, by continuity of matrix multiplication and logarithm, $\mathbf{D}\log \boldsymbol{\delta}_\gamma^* = \mathbf{0}$. \square

The next statement specifies when $\text{IPF}(\gamma)$ converges to the maximum likelihood estimates under a relational model.

Corollary 3.3. *The following statements hold:*

1. Let \mathbf{Y} be a random variable that has a multinomial distribution with parameters N and \mathbf{p} . If the overall effect is present, then the sequence obtained from $\text{IPF}(1)$ converges to the maximum likelihood estimate $\hat{\mathbf{p}}$ of \mathbf{p} under the model. If the overall effect is not present, then, if the adjustment factor γ^* is known, the sequence obtained from $\text{IPF}(\gamma^*)$ converges to the maximum likelihood estimate $\hat{\mathbf{p}}$ of \mathbf{p} under the model.

2. Let \mathbf{Y} be a random variable that has a Poisson distribution with parameter $\boldsymbol{\lambda}$. Then, whether or not the overall effect is present, the sequence obtained from IPF(1) converges to the maximum likelihood estimate $\hat{\boldsymbol{\lambda}}$ of $\boldsymbol{\lambda}$ under the model. \square

The conventional IPF procedure computes maximum likelihood estimates for the cell parameters, and it is often pointed out that this algorithm is not suitable for estimating the model parameters [cf. Fienberg and Rinaldo, 2012, p.14]. In fact, this is not the case, and it will be proven next that the sequence $\boldsymbol{\theta}_\gamma^{(d)}$, obtained from IPF(γ), converges to estimates of the model parameters.

Theorem 3.4. *For all $j = 1, \dots, J$, the sequence $\theta_\gamma^{(d)}(j)$, obtained in (18), converges, as $d \rightarrow \infty$, to a $\theta_\gamma^*(j)$ and*

$$\delta_\gamma^*(i) = \prod_{j=1}^J (\theta_\gamma^*(j))^{a_{ji}}, \quad \text{for all } i \in \mathcal{I}.$$

Proof. By the choice of the initial values, $\delta_\gamma^{(0)}(i) = \prod_{j=1}^J (\theta_\gamma^{(0)}(j))^{a_{ji}}$, for all $i \in \mathcal{I}$. The further argument is by induction.

Let $d \geq 0$ and $d+1 \equiv j \pmod{J}$. By the induction hypothesis,

$$\delta_\gamma^{(d)}(i) = \prod_{j=1}^J (\theta_\gamma^{(d)}(j))^{a_{ji}}, \quad \text{for all } i \in \mathcal{I}.$$

During the $(d+1)$ -st iteration, only the parameters of the cells that belong to the subset S_j , and the model parameter corresponding to S_j are updated; $\delta_\gamma^{(d+1)}(i) = \delta_\gamma^{(d)}(i)$ for $i \notin S_j$, and $\theta_\gamma^{(d+1)}(l) = \theta_\gamma^{(d)}(l)$ for $l \neq j$. Therefore, for $i \notin S_j$

$$\delta_\gamma^{(d+1)}(i) = \prod_{l=1}^J (\theta_\gamma^{(d+1)}(l))^{a_{li}},$$

because for those parameters that changed during the last step of iteration, the exponent a_{li} is zero. Further, for $i \in S_j$

$$\delta_\gamma^{(d+1)}(i) = \delta_\gamma^{(d)}(i) \frac{\theta_\gamma^{(d+1)}(j)}{\theta_\gamma^{(d)}(j)} = \prod_{l=1}^J (\theta_\gamma^{(d)}(l))^{a_{li}} \frac{\theta_\gamma^{(d+1)}(j)}{\theta_\gamma^{(d)}(j)} = \prod_{l=1}^J (\theta_\gamma^{(d+1)}(l))^{a_{li}}.$$

By the principle of induction, for all $d \geq 0$,

$$\delta_\gamma^{(d+1)}(i) = \prod_{j=1}^J (\theta_\gamma^{(d+1)}(j))^{a_{ji}}, \quad \text{for all } i \in \mathcal{I}. \quad (22)$$

During the iterations, the value of the parameter associated with the subset S_j is updated at a step $j + kJ$, for some non-negative integer k , and then remains constant till it is updated again in the step $j + (k + 1)J$. The ratio of two subsequent values of this parameter equals

$$\frac{\theta_\gamma^{(j+(k+1)J)}(j)}{\theta_\gamma^{(j+kJ)}(j)} = \frac{\theta_\gamma^{(j+(k+1)J)}(j)}{\theta_\gamma^{(j-1+(k+1)J)}(j)} \cdot \frac{\theta_\gamma^{(j-1+(k+1)J)}(j)}{\theta_\gamma^{(j-2+(k+1)J)}(j)} \cdots \frac{\theta_\gamma^{(j-(J-1)+(k+1)J)}(j)}{\theta_\gamma^{(j+kJ)}(j)},$$

which, by repeated application of (18), may be written as

$$\frac{\delta_\gamma^{(j+(k+1)J)}(i^*)}{\delta_\gamma^{(j+kJ)}(i^*)} = \frac{\delta_\gamma^{(j+(k+1)J)}(i^*)}{\delta_\gamma^{(j-1+(k+1)J)}(i^*)} \cdot \frac{\delta_\gamma^{(j-1+(k+1)J)}(i^*)}{\delta_\gamma^{(j-2+(k+1)J)}(i^*)} \cdots \frac{\delta_\gamma^{(j-(J-1)+(k+1)J)}(i^*)}{\delta_\gamma^{(j+kJ)}(i^*)}, \quad (23)$$

where i^* is specified prior to (18).

By Theorem 3.2, as $d \rightarrow \infty$, $\delta_\gamma^{(d)}(i)$ converges to a $\delta_\gamma^*(i)$ for every $i \in \mathcal{I}$. Therefore, the right hand side of (23) converges to 1 and thus

$$\frac{\theta_\gamma^{(j+(k+1)J)}(j)}{\theta_\gamma^{(j+kJ)}(j)} \rightarrow 1,$$

when $d \rightarrow \infty$, i.e., for every fixed j , when $k \rightarrow \infty$.

This, by applying an elementary calculus argument, implies that for every $j = 1, \dots, J$, $\theta_\gamma^{(d)}(j)$, obtained in (18), converges to a $\theta_\gamma^*(j)$. From (22), by continuity,

$$\delta_\gamma^*(i) = \prod_{j=1}^J (\theta_\gamma^*(j))^{a_{ji}}, \quad \text{for all } i \in \mathcal{I}.$$

□

The G-IPF algorithm described next can be used for computing the maximum likelihood estimates under relational models both for probabilities and for intensities, with or without the overall effect. In the case when the model is for probabilities and does not include the overall effect, G-IPF complements IPF(γ) with a step that computes the adjustment factor.

G-IPF Algorithm:

If $\delta \equiv \lambda$, compute $\hat{\lambda}$ using IPF(1), and finish.

If $\delta \equiv \mathbf{p}$, compute \mathbf{p}^* using IPF(1).

If $\mathbf{1p}^* = 1$, set $\hat{\mathbf{p}} = \mathbf{p}^*$, and finish. Otherwise,

compute $\gamma_L = (\mathbf{1Aq})^{-1}$, $\gamma_R = \min \{1/A_1\mathbf{q}, \dots, 1/A_J\mathbf{q}\}$, and proceed as follows:

Step 1: Find $\delta_{(\gamma_L + \gamma_R)/2}^*$ using IPF(γ).

Step 2: While $\mathbf{1}\delta_{(\gamma_L+\gamma_R)/2}^* \neq 1$,
 if $\mathbf{1}\delta_{(\gamma_L+\gamma_R)/2}^* < 1$, set $\gamma_L = \frac{\gamma_L+\gamma_R}{2}$,
 else set $\gamma_R = \frac{\gamma_L+\gamma_R}{2}$;
 go to **Step 1**.

Step 3: Set $\hat{\mathbf{p}} = \delta_{(\gamma_L+\gamma_R)/2}^*$, and finish. □

The next lemma justifies the initial choice of γ_L and γ_R .

Lemma 3.5. *Let $\gamma > 0$ and let δ_γ be a solution to the system of equations*

$$\mathbf{A}\delta = \gamma\mathbf{A}\mathbf{q}, \quad \mathbf{D}\log \delta = 0. \quad (24)$$

Then, for $\gamma_L = (\mathbf{1Aq})^{-1}$ and $\gamma_R = \min \{1/A_1\mathbf{q}, \dots, 1/A_J\mathbf{q}\}$, $\mathbf{1}\delta_{\gamma_L} \leq 1$ and $\mathbf{1}\delta_{\gamma_R} \geq 1$.

Proof. Each column of \mathbf{A} contains at least one 1, and thus

$$\mathbf{1}\delta_\gamma \leq \mathbf{1A}\delta_\gamma = \gamma\mathbf{1Aq},$$

which implies the first inequality. Since \mathbf{A} is a 0-1 matrix,

$$\gamma \max \{A_1\mathbf{q}, \dots, A_J\mathbf{q}\} = \max \{A_1\delta_\gamma, \dots, A_J\delta_\gamma\} \leq \mathbf{1}\delta_\gamma.$$

Using that

$$(\max \{A_1\mathbf{q}, \dots, A_J\mathbf{q}\})^{-1} = \min \{1/A_1\mathbf{q}, \dots, 1/A_J\mathbf{q}\},$$

implies the second inequality. □

The next lemma is needed for the proof of convergence of G-IPF.

Lemma 3.6. *The solution, δ_γ , to the system of equations (24) and thus its total, $\mathbf{1}\delta_\gamma$, are continuous functions of γ .*

Proof. Fix an arbitrary $\gamma_0 > 0$. It will be shown that for a sequence $\{\gamma_n\} > 0$, which converges to γ_0 , as $n \rightarrow \infty$, $\delta_{\gamma_n} \rightarrow \delta_{\gamma_0}$.

Since γ_n converges, there exists a finite cover of $\{\gamma_n\}$ and so does a finite cover of the sequence of segments $\{[0, \gamma_n]\}$. Since $\mathbf{1q} = 1$, $\delta_{\gamma_n} \leq \gamma_n$, for every $n \geq 1$. Hence, there exists a finite cover of the sequence $\{\delta_{\gamma_n}\}$, implying that the set $\{\delta_{\gamma_n}\}$ is compact, and, therefore, there exists a subsequence $\{\delta_{\gamma_{n_k}}\}$ that converges to a δ^* . By continuity of the logarithm and matrix multiplication,

$$\mathbf{A}\delta^* = \gamma_0\mathbf{Aq}, \quad \mathbf{D}\log \delta^* = 0.$$

Since the solution of this system is unique, $\delta^* = \delta_{\gamma_0}$, and thus $\delta_{\gamma_{n_k}} \rightarrow \delta_{\gamma_0}$. □

The next theorem states that G-IPF produces maximum likelihood estimates of the cell parameters under a relational model, whether or not the overall effect is present.

Theorem 3.7. *Assume that the maximum likelihood estimate of $\boldsymbol{\delta}$ under the relational model generated by \mathbf{A} exists and is unique. Then $\hat{\boldsymbol{\delta}}$, obtained from G-IPF, is the maximum likelihood estimate of $\boldsymbol{\delta}$.*

Proof. In the case when $\boldsymbol{\delta} \equiv \boldsymbol{\lambda}$, G-IPF computes $\hat{\boldsymbol{\lambda}}$ and $\hat{\boldsymbol{\theta}}$ using IPF(1). By Corollary 3.3, $\hat{\boldsymbol{\lambda}}$ is the MLE of the cell intensities.

In the case when $\boldsymbol{\delta} \equiv \mathbf{p}$, G-IPF first calls IPF(1) and computes a vector of probabilities \mathbf{p}^* .

If $\mathbf{1p}^* = 1$, which happens if and only if $\mathbf{1} \in R(\mathbf{A})$, then, by Corollary 3.3, $\hat{\mathbf{p}} = \mathbf{p}^*$ is the MLE of the cell probabilities.

Suppose $\mathbf{1p}^* \neq 1$. By Lemma 3.5, the function $\mathbf{1p}_\gamma - 1$ has opposite signs at the endpoints of the segment $[(\mathbf{1Aq})^{-1}, \min \{1/A_1\mathbf{q}, \dots, 1/A_J\mathbf{q}\}]$. Since $\mathbf{1p}_\gamma - 1$ is a continuous function of γ , the segment contains a value

$$\gamma^*, \text{ such that } \mathbf{1p}_{\gamma^*} = 1. \quad (25)$$

In order to find γ^* the algorithm uses the bisection method [cf. Boyd and Vandenberghe, 2004]. By the uniqueness of the MLE, $\hat{\mathbf{p}} = \mathbf{p}_{\gamma^*}$. \square

A reviewer to an earlier version of the paper pointed out that the above algorithm cannot be programmed because any implementation of IPF(γ) will return the value of $\boldsymbol{\delta}_\gamma^*$ with some error. A programmable variant of the algorithm is obtained by choosing a desired precision $\epsilon > 0$. Define IPF(γ, ϵ) by replacing Steps 3 and 4 of IPF(γ) by the following steps:

Step 3: While $|\gamma A_j \mathbf{q} - A_j \boldsymbol{\delta}_\gamma^{(d+1)}| > \epsilon$ for at least one j , set $d = d + 1$, go to Step 1.

Step 4: Set $\boldsymbol{\delta}_{\gamma, \epsilon}^* = \boldsymbol{\delta}_\gamma^{(d)}$, and finish.

The sequence of vectors $\boldsymbol{\delta}_{\gamma, \epsilon}^*$, obtained using IPF(γ, ϵ), converges pointwise, as $\epsilon \rightarrow 0$, to $\boldsymbol{\delta}_\gamma^*$, the result of IPF(γ).

The G-IPF(ϵ) algorithm described next computes the maximum likelihood estimates of cell parameters under a relational model with a given precision.

G-IPF(ϵ) Algorithm:

If $\boldsymbol{\delta} \equiv \boldsymbol{\lambda}$, compute $\hat{\boldsymbol{\lambda}}_\epsilon$ using IPF(1, ϵ), and finish.

If $\delta \equiv \mathbf{p}$, compute \mathbf{p}_ϵ^* using IPF(1, ϵ).

If $|\mathbf{1p}_\epsilon^* - 1| < \epsilon$, set $\hat{\mathbf{p}}_\epsilon = \mathbf{p}_\epsilon^*$, and finish. Otherwise, set $T = 1$,

compute $\gamma_L = (\mathbf{1Aq})^{-1}$, $\gamma_R = \min \{1/A_1\mathbf{q}, \dots, 1/A_J\mathbf{q}\}$, and proceed as follows:

Step 1: for $t = 0, \dots, T$, let $\gamma_t = \gamma_L + \frac{t}{T}(\gamma_R - \gamma_L)$, and compute $\mathbf{p}_{\gamma_t, \epsilon}^*$ using IPF(γ_t , ϵ/T). Whenever $|\mathbf{1p}_{\gamma_t, \epsilon}^* - 1| < \epsilon$, go to **Step 3**.

Step 2: Set $T = T + 1$, go to **Step 1**.

Step 3: Set $\hat{\mathbf{p}}_\epsilon = \mathbf{p}_{\gamma_t, \epsilon}^*$, $\hat{\gamma}_\epsilon = \gamma_t$, and finish. \square

The next theorem states that G-IPF(ϵ) produces maximum likelihood estimates of the cell parameters under a relational model with precision ϵ .

Theorem 3.8. *Assume that the maximum likelihood estimate of δ under the relational model generated by \mathbf{A} exists and is unique. For every $\epsilon > 0$, G-IPF(ϵ) terminates and returns $\hat{\delta}_\epsilon$ and $\hat{\gamma}_\epsilon$, which satisfy:*

$$|\mathbf{A}\hat{\delta}_\epsilon - \hat{\gamma}_\epsilon\mathbf{Aq}| < \epsilon; \quad (26)$$

$$\mathbf{D} \log \hat{\delta}_\epsilon = 0; \quad (27)$$

$$|\mathbf{1}\hat{\delta}_\epsilon - 1| < \epsilon \quad (\text{for } \delta \equiv \mathbf{p}). \quad (28)$$

Proof. Fix an $\epsilon > 0$. When $\delta \equiv \mathbf{\lambda}$, G-IPF(ϵ) computes $\hat{\mathbf{\lambda}}_\epsilon$ using IPF(1, ϵ). Since IPF(1, ϵ) terminates, so does G-IPF(ϵ). The termination condition of IPF(1, ϵ) implies (26), with $\hat{\gamma}_\epsilon = 1$.

In the case when $\delta \equiv \mathbf{p}$, G-IPF(ϵ) first calls IPF(1, ϵ), which terminates and results in \mathbf{p}_ϵ^* , for which (26) holds with $\hat{\gamma}_\epsilon = 1$. If (28) also holds, G-IPF(ϵ) is terminated after the first iteration with $\hat{\mathbf{p}}_\epsilon = \mathbf{p}_\epsilon^*$.

Suppose $|\mathbf{1p}_\epsilon^* - 1| > \epsilon$, and thus G-IPF(ϵ) does not stop after the first iteration. By Lemma 3.5, the (unknown) adjustment factor corresponding to the MLE, say $\hat{\gamma}$, belongs to the segment $[\gamma_L, \gamma_R]$. By Lemma 3.6, $\mathbf{1p}_\gamma - 1$ is a continuous function of γ , and, since $\mathbf{1p}_\gamma = 1$, for every $T_0 \in \mathbb{Z}_{>0}$, there exists a $t_0 \in \{0, \dots, T_0\}$, such that $\hat{\gamma} \in [\gamma_{t_0}, \gamma_{t_0+1}]$, where $\gamma_{t_0} = \gamma_L + \frac{t_0}{T_0}(\gamma_R - \gamma_L)$. As $T_0 \rightarrow \infty$, the length of the interval $[\gamma_{t_0}, \gamma_{t_0+1}]$ approaches 0, and thus, for a large enough T_0 , both, $\mathbf{p}_{\gamma_{t_0}, \epsilon}^*$ and $\mathbf{p}_{\gamma_{t_0+1}, \epsilon}^*$, obtained from IPF(γ_{t_0} , ϵ/T_0) and IPF(γ_{t_0+1} , ϵ/T_0), respectively, satisfy (28). Thus, G-IPF(ϵ) will terminate at the $(T_0 + 1)$ -st iteration with $\hat{\mathbf{p}}_\epsilon$ that fulfills (26), with some $\hat{\gamma}_\epsilon$, and (28).

Since the multiplicative structure of the distributions does not change during the iterations, (27) holds. \square

The theorem implies that the sequence of vectors $\hat{\boldsymbol{\delta}}_\epsilon$, obtained using G-IPF(ϵ), converges pointwise, as $\epsilon \rightarrow 0$, to the maximum likelihood estimate of $\boldsymbol{\delta}$.

The G-IPF(ϵ) algorithm is implemented in the R-package `gIPFrm` [Klimova and Rudas, 2014]. In this implementation, there is also an option to update the adjustment factor using the bisection method. Alternatively, any numerical technique available for monotone functions [cf. Boyd and Vandenberghe, 2004] may be used.

Acknowledgments

Part of the material presented here was contained in the PhD thesis of the first author to which the second author and Thomas Richardson were advisers. The authors wish to thank him for several comments and suggestions. The proof of Proposition 2.1 uses the idea of Olga Klimova, to whom the authors are also indebted. The second author was supported in part by Grant K-106154 from the Hungarian National Scientific Research Fund (OTKA).

References

- J. Aitchison and S. D. Silvey. Maximum-likelihood estimation procedures and associated tests of significance. *J. Roy. Statist. Soc. Ser.B*, 22:154–171, 1960.
- I. Bancarz and M. Osborne. Improved iterative scaling can yield multiple globally optimal models with radically differing performance levels. In S. C. Tseng, T. E. Chen, and Y. F. Liu, editors, *Proceedings of the Nineteenth International Conference on Computational Linguistics (COLING 2002)*, Howard International House, Taipei, Taiwan, pages 1–7. Morgan Kaufmann, 2002.
- D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.
- Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete multivariate analysis: theory and practice*. MIT, 1975.
- P. A. Boyd, A. M. Tonks, J. Rankin, C. Rounding, D. Wellesley, and E. S. Draper. Monitoring the prenatal detection of structural fetal congenital anomalies in England and Wales: register-based study. *J. Med. Screen.*, 18:2–7, 2011.
- S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004.

- L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *U.S.S.R. Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.
- S. Brin, R. Motwani, and R. Silverstein. Beyond market basket: generalizing association rules to correlations. In J. M. Peckham, editor, *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data (SIGMOD 1997)*, Tucson, AZ, USA, pages 265–276, 1997.
- S. Classen, O. Shechtman, K. D. Awadzi, Y. Joo, and D. N. Lanford. Traffic violations versus driving errors of older adults: Informing clinical practice. *The American Journal of Occupational Therapy*, 64:233–241, 2010.
- G. E. Copeland and R. S. Kirby. Using birth defects registry data to evaluate infant and childhood mortality associated with birth defects: an alternative to traditional mortality assessment using underlying cause of death statistics. *Birth Defects Res. A Clin. Mol. Teratol.*, 79:792–797, 2007.
- I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Probab.*, 3:146–158, 1975.
- J. N. Darroch and D. Ratcliff. Generalized iterative scaling for log-linear models. *The Annals of Mathematical Statistics*, 43:1470–1480, 1972.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. Inducing features of random fields. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19:283–297, 1997.
- W. E. Deming and F.F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11: 427–444, 1940.
- R. J. Evans and A. Forcina. Two algorithms for fitting constrained marginal models. *Comput. Statist. Data Anal.*, 2013.
- S. E. Fienberg and A. Rinaldo. Maximum likelihood estimation in log-linear models - supplementary material. Technical report, Carnegie Mellon University, 2012.
- F. L. Huang, C. J. Hsieh, K. W. Chang, and C. J. Lin. Iterative scaling and coordinate descent methods for maximum entropy models. *J. Mach. Learn. Res.*, 11:815–848, 2010.

- R. Jin, R. Yan, J. Zhang, and A. G. Hauptmann. A faster iterative scaling algorithm for conditional exponential model. In T. Fawcett and N. Mishra, editors, *Proceedings of the Twentieth International Conference on Machine Learning (ICML 2003)*, Washington DC, pages 282–289. The AAAI Press, Menlo Park, CA, USA, 2003.
- K. B. M. Källén, E. E. Castilla, M. D. G. Dutra, P. Mastroiacovo, E. Robert, and B. A. J. Källén. A modified method for the epidemiological analysis of registry data on infants with multiple malformations. *International Journal of Epidemiology*, 28:701–710, 1999.
- A. Klimova and T. Rudas. Coordinate free analysis of trends in British social mobility. *J. Appl. Stat.*, 39:1681–1691, 2012.
- A. Klimova and T. Rudas. *gIPFrm: Generalized Iterative Proportional Fitting for Relational Models*, 2014. URL <http://CRAN.R-project.org/package=gIPFrm>. R package version 2.0.
- A. Klimova, T. Rudas, and A. Dobra. Relational models for contingency tables. *J. Multivariate Anal.*, 104:159–173, 2012.
- J. Lafferty, S. Della Pietra, and V. Della Pietra. Statistical learning algorithms based on Bregman distances. In *Proceedings of the 1997 Canadian Workshop on Information Theory*, Fields Institute, Toronto, Canada, pages 77–80, 1997.
- J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: probabilistic models for segmenting and labeling sequence data. In C. E. Brodley and A. Pohorecky Danyluk, editors, *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, Williams College, Williamstown, MA, USA, pages 282–289. Morgan Kaufman, 2001.
- R. Malouf. A comparison of algorithms for maximum entropy parameter estimation. In *Proceedings of the Sixth Conference on Natural Language Learning (CoNLL-2002)*, Stroudsburg, PA, USA, pages 49–55. Association for Computational Linguistics, 2002.
- A. McCallum, D. Freitag, and F. C. N. Pereira. Maximum entropy Markov models for information extraction and segmentation. In P. Langley, editor, *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford University, Stanford, CA, USA, pages 591–598. Morgan Kaufman, 2000.
- J. A. O’Sullivan and J. Benac. Alternating minimization algorithms for transmission tomography. *IEEE Trans. Med. Imag.*, 26:283–297, 2007.




Figure 1: Histogram of sums of the limit vectors obtained by IIS, under the model of Aitchison and Silvey [1960], for 13352 generated distributions.

X. Wu, D. Barbara, and Y. Ye. Screening and interpreting multi-item associations based on log-linear modeling. In *Proceedings of the Ninth ACM SIGKDD international conference on Knowledge discovery and data mining (KDD 2003)*, Washington, DC, USA, pages 276–285. ACM, New York, NY, USA, 2003.

Table 1: Relational models and properties of the MLE [Klimova et al., 2012].

	Models with the overall effect	Models without the overall effect	
	Probabilities & Intensities	Probabilities	Intensities
Exponential family	Regular	Curved	Regular
Subset sums of the MLE vs observed subset sums	Equal	Proportional	Equal
Adjustment for total	1	1	Depends on the data
Adjustment for subset sums	1	Depends on the data	1