

Published in final edited form as:

Comput Stat Data Anal. 2013 October 1; 66: 1–7. doi:10.1016/j.csda.2013.02.001.

Two algorithms for fitting constrained marginal models

R.J. Evans^{a,*} and A. Forcina^b

^aStatistical Laboratory, University of Cambridge, UK

^bDipartimento di Economia, Finanza e Statistica, University of Perugia, Italy

Abstract

The two main algorithms that have been considered for fitting constrained marginal models to discrete data, one based on Lagrange multipliers and the other on a regression model, are studied in detail. It is shown that the updates produced by the two methods are identical, but that the Lagrangian method is more efficient in the case of identically distributed observations. A generalization is given of the regression algorithm for modelling the effect of exogenous individual-level covariates, a context in which the use of the Lagrangian algorithm would be infeasible for even moderate sample sizes. An extension of the method to likelihood-based estimation under L_1 -penalties is also considered.

Keywords

categorical data; L_1 -penalty; marginal log-linear model; maximum likelihood; non-linear constraint

1. Introduction

The application of marginal constraints to multi-way contingency tables has been much investigated in the last 20 years; see, for example, McCullagh and Nelder (1989); Liang et al. (1992); Lang and Agresti (1994); Glonek and McCullagh (1995); Agresti (2002); Bergsma et al. (2009). Bergsma and Rudas (2002) introduced marginal log-linear parameters (MLLPs), which generalize other discrete parameterizations including ordinary log-linear parameters and Glonek and McCullagh's multivariate logistic parameters. The flexibility of this family of parameterizations enables their application to many popular classes of conditional independence models, and especially to graphical models (Forcina et al., 2010; Rudas et al., 2010; Evans and Richardson, 2013). Bergsma and Rudas (2002) show that, under certain conditions, models defined by linear constraints on MLLPs are curved exponential families. However, naïve algorithms for maximum likelihood estimation with MLLPs face several challenges: in general, there are no closed form equations for computing raw probabilities from MLLPs, so direct evaluation of the log-likelihood can be time consuming; in addition, MLLPs are not necessarily variation independent and, as noted by Bartolucci et al. (2007), ordinary Newton-Raphson or Fisher scoring methods may become stuck by producing updated estimates which are incompatible.

© 2013 Elsevier B.V. All rights reserved.

*Address for correspondence: R.J. Evans, Statistical Laboratory, Wilberforce Road, Cambridge, CB3 0WB, UK; rje42@cam.ac.uk; phone: +44 (0) 1223 337952; fax: +44 (0) 1223 337956.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Lang (1996) and Bergsma (1997), amongst others, have tried to adapt a general algorithm introduced by Aitchison and Silvey (1958) for constrained maximum likelihood estimation to the context of marginal models. In this paper we provide an explicit formulation of Aitchison and Silvey's algorithm, and show that an alternative method due to Colombi and Forcina (2001) is equivalent; we term this second approach the *regression algorithm*. Though the regression algorithm is less efficient, we show that it can be extended to deal with individual-level covariates, a context in which Aitchison and Silvey's approach is infeasible, unless the sample size is very small. A variation of these algorithms, which can be used to fit marginal log-linear models under L_1 -penalties, and therefore perform automatic model selection, is also given.

Section 2 reviews marginal log-linear models and their basic properties, while in Section 3 we formulate the two algorithms, show that they are equivalent and discuss their properties. In Section 4 we derive an extension of the regression algorithm which can incorporate the effect of individual-level covariates. Finally Section 5 considers similar methods for L_1 -constrained estimation.

2. Notations and preliminary results

Let $X_j, j = 1, \dots, d$ be categorical random variables taking values in $\{1, \dots, c_j\}$. The joint distribution of X_1, \dots, X_d is determined by the vector of joint probabilities $\boldsymbol{\pi}$ of dimension $t = \prod_{j=1}^d c_j$, whose entries correspond to cell probabilities, and are assumed to be strictly positive; we take the entries of $\boldsymbol{\pi}$ to be in lexicographic order. Further, let \mathbf{y} denote the vector of cell frequencies with entries arranged in the same order as $\boldsymbol{\pi}$. We write the multinomial log-likelihood in terms of the canonical parameters as

$$l(\boldsymbol{\theta}) = \mathbf{y}' \mathbf{G} \boldsymbol{\theta} - n \log[\mathbf{1}'_t \exp(\mathbf{G} \boldsymbol{\theta})]$$

(see, for example, Bartolucci et al., 2007, p. 699); here n is the sample size, $\mathbf{1}_t$ a vector of length t whose entries are all 1, and \mathbf{G} a $t \times (t-1)$ full rank design matrix which determines the log-linear parameterization. The mapping between the canonical parameters and the joint probabilities may be expressed as

$$\log(\boldsymbol{\pi}) = \mathbf{G} \boldsymbol{\theta} - \mathbf{1}_t \log[\mathbf{1}'_t \exp(\mathbf{G} \boldsymbol{\theta})] \iff \boldsymbol{\theta} = \mathbf{L} \log(\boldsymbol{\pi}),$$

where \mathbf{L} is a $(t-1) \times t$ matrix of row contrasts and $\mathbf{L} \mathbf{G} = \mathbf{I}_{t-1}$.

The score vector, \mathbf{s} , and the expected information matrix, \mathbf{F} , with respect to $\boldsymbol{\theta}$ take the form

$$\mathbf{s} = \mathbf{G}' (\mathbf{y} - n\boldsymbol{\pi}) \quad \text{and} \quad \mathbf{F} = n \mathbf{G}' \boldsymbol{\Omega} \mathbf{G};$$

here $\boldsymbol{\Omega} = \text{diag}(\boldsymbol{\pi}) - \boldsymbol{\pi} \boldsymbol{\pi}'$.

2.1. Marginal log-linear parameters

Marginal log-linear parameters (MLLPs) enable the simultaneous modelling of several marginal distributions (see, for example, Bergsma et al., 2009, Chapters 2 and 4) and the specification of suitable conditional independencies within marginal distributions of interest

(see Evans and Richardson, 2013). In the following let $\boldsymbol{\eta}$ denote an arbitrary vector of MLLPs; it is well known that this can be written as

$$\boldsymbol{\eta} = \mathbf{C} \log(\mathbf{M}\boldsymbol{\pi}),$$

where \mathbf{C} is a suitable matrix of row contrasts, and \mathbf{M} a matrix of 0's and 1's producing the appropriate margins (see, for example, Bergsma et al., 2009, Section 2.3.4).

Bergsma and Rudas (2002) have shown that if a vector of MLLPs $\boldsymbol{\eta}$ is *complete* and *hierarchical*, two properties defined below, models determined by linear restrictions on $\boldsymbol{\eta}$ are curved exponential families, and thus smooth. Like ordinary log-linear parameters, MLLPs may be grouped into interaction terms involving a particular subset of variables; each interaction term must be defined within a margin of which it is a subset.

Definition 1

A vector of MLLPs $\boldsymbol{\eta}$ is called complete if every possible interaction is defined in precisely one margin.

Definition 2

A vector of MLLPs $\boldsymbol{\eta}$ is called hierarchical if there is a nondecreasing ordering of the margins of interest M_1, \dots, M_s such that, for each $j = 1, \dots, s$, no interaction term which is a subset of M_j is defined within a later margin.

3. Two algorithms for fitting marginal log-linear models

Here we describe the two main algorithms used for fitting models of the kind described above.

3.1. An adaptation of Aitchison and Silvey's algorithm

Aitchison and Silvey (1958) study maximum likelihood estimation under non-linear constraints in a very general context, showing that, under certain conditions, the maximum likelihood estimates exist and are asymptotically normal; they also outline an algorithm for computing those estimates. Suppose we wish to maximize $l(\boldsymbol{\theta})$ subject to $\mathbf{h}(\boldsymbol{\theta}) = \mathbf{0}$, a set of r non-linear constraints, under the assumption that the second derivative of $\mathbf{h}(\boldsymbol{\theta})$ exists and is bounded. Aitchison and Silvey consider stationary points of the function $l(\boldsymbol{\theta}) + \mathbf{h}(\boldsymbol{\theta})' \boldsymbol{\lambda}$, where $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers; this leads to the system of equations

$$\begin{aligned} s(\hat{\boldsymbol{\theta}}) + \mathbf{H}(\hat{\boldsymbol{\theta}}) \hat{\boldsymbol{\lambda}} &= \mathbf{0} \\ \mathbf{h}(\hat{\boldsymbol{\theta}}) &= \mathbf{0}, \end{aligned} \tag{1}$$

where $\hat{\boldsymbol{\theta}}$ is the ML estimate and \mathbf{H} the derivative of \mathbf{h}' with respect to $\boldsymbol{\theta}$. Since these are non-linear equations, they suggest an iterative algorithm which proceeds as follows: suppose that at the current iteration we have $\boldsymbol{\theta}_0$, a value reasonably close to $\hat{\boldsymbol{\theta}}$. Replace s and \mathbf{h} with first order approximations around $\boldsymbol{\theta}_0$; in addition replace $\mathbf{H}(\hat{\boldsymbol{\theta}})$ with $\mathbf{H}(\boldsymbol{\theta}_0)$ and the second derivative of the log-likelihood with $-\mathbf{F}$, minus the expected information matrix. The resulting equations, after rearrangement, may be written in matrix form as

$$\begin{pmatrix} \widehat{\theta} - \theta_0 \\ \widehat{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{F}_0 & -\mathbf{H}_0 \\ -\mathbf{H}'_0 & \mathbf{0} \end{pmatrix}^{-1} \begin{pmatrix} s_0 \\ h_0 \end{pmatrix},$$

where $s_0, \mathbf{F}_0, \mathbf{H}_0$ and so on denote the corresponding quantities evaluated at θ_0 . To compute a solution, Aitchison and Silvey (1958) exploit the structure of the partitioned matrix, while Bergsma (1997) solves explicitly for $\widehat{\theta}$ by substitution; in both cases, if we are uninterested in the Lagrange multipliers, we get the updating equation

$$\widehat{\theta} = \theta_0 + \mathbf{F}_0^{-1} s_0 - \mathbf{F}_0^{-1} \mathbf{H}_0 (\mathbf{H}'_0 \mathbf{F}_0^{-1} \mathbf{H}_0)^{-1} (\mathbf{H}'_0 \mathbf{F}_0^{-1} s_0 + h_0). \quad (2)$$

As noted by Bergsma (1997), the algorithm does not always converge unless some sort of step length adjustment is introduced.

Linearly constrained marginal models are defined by $\mathbf{K}' \boldsymbol{\eta} = \mathbf{0}$, where \mathbf{K} is a matrix of full column rank $r - t - 1$. The multinomial likelihood is a regular exponential family, so these models may be fitted using the smooth constraint $\mathbf{h}(\boldsymbol{\theta}) = \mathbf{K}' \boldsymbol{\eta}(\boldsymbol{\theta}) = \mathbf{0}$, which implies that

$$\mathbf{H}' = \frac{\partial \mathbf{h}}{\partial \boldsymbol{\theta}} = \frac{\partial \mathbf{h}}{\partial \boldsymbol{\eta}'} \frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\theta}} = \mathbf{K}' \mathbf{C} \text{diag}(\mathbf{M}\boldsymbol{\pi})^{-1} \mathbf{M} \text{diag}(\boldsymbol{\pi}) \mathbf{G}.$$

Remark 1—In the equation above we have replaced $\boldsymbol{\Omega}$ with $\text{diag}(\boldsymbol{\pi})$ by exploiting the fact that $\boldsymbol{\eta}$ is a homogeneous function of $\boldsymbol{\pi}$ (see Bergsma et al., 2009, Section 2.3.4). If the constrained model were not smooth then at singular points the Jacobian matrix \mathbf{R} would not be invertible, implying that \mathbf{H} is not of full rank and thus violating a crucial assumption in Aitchison and Silvey (1958). It has been shown (Bergsma and Rudas, 2002, Theorem 3) that completeness is a necessary condition for smoothness.

Calculation of (2) may be simplified by noting that $\mathbf{K}' \mathbf{C}$ does not need to be updated; in addition, if we choose, for example, \mathbf{G} to be the identity matrix of size t with the first column removed, an explicit inverse of \mathbf{F} exists:

$$\mathbf{F}^{-1} = \left[n(\text{diag}(\dot{\boldsymbol{\pi}}) - \dot{\boldsymbol{\pi}} \dot{\boldsymbol{\pi}}') \right]^{-1} = n^{-1} \left[\text{diag}(\dot{\boldsymbol{\pi}})^{-1} + \mathbf{1}_{t-1} \mathbf{1}'_{t-1} / (1 - \mathbf{1}'_{t-1} \dot{\boldsymbol{\pi}}) \right],$$

where $\dot{\boldsymbol{\pi}}$ denotes the vector $\boldsymbol{\pi}$ with the first element removed; this expression may be exploited when computing $\mathbf{F}^{-1} \mathbf{H}$.

3.2. A regression algorithm

By noting that the Aitchison-Silvey algorithm is essentially based on a quadratic approximation of $\mathcal{K}(\boldsymbol{\theta})$ with a linear approximation of the constraints, Colombi and Forcina (2001) designed an algorithm which they believed to be equivalent to the original, though no formal argument was provided; this equivalence is proven in Proposition 1 below. Recall that, by elementary linear algebra, there exists a $(t-1) \times (t-r-1)$ design matrix \mathbf{X} of full column rank such that $\mathbf{K}' \mathbf{X} = \mathbf{0}$, from which it follows that $\boldsymbol{\eta} = \mathbf{X} \boldsymbol{\beta}$ for a vector of $t-r-1$ unknown parameters $\boldsymbol{\beta}$. Let

$$R = \frac{\partial \theta}{\partial \eta'} = [C \text{diag}(M\pi)^{-1} M \text{diag}(\pi) G]^{-1},$$

and $\bar{s} = R' s$, $\bar{F} = R' F R$ respectively denote the score and information relative to η ; then the regression algorithm consists of alternating the following steps:

1. update the estimate of β by

$$\widehat{\beta} - \beta_0 = (X' \bar{F}_0 X)^{-1} X' (\bar{F}_0 \gamma_0 + \bar{s}_0), \quad (3)$$

where $\gamma_0 = \eta_0 - X \beta_0$;

2. update θ by

$$\widehat{\theta} - \theta_0 = R_0 [X (\widehat{\beta} - \beta_0) - \gamma_0]. \quad (4)$$

Proposition 1—The updating equation in (2) is equivalent to the combined steps given in (3) and (4).

Proof: First, consider matrices X and K such that the columns of X span the orthogonal complement of the space spanned by the columns of K . Then we claim that for any symmetric and positive definite matrix W

$$W^{-1} - W^{-1} K (K' W^{-1} K)^{-1} K' W^{-1} = X (X' W X)^{-1} X'. \quad (5)$$

To see this, let $U = W^{-1/2} K$ and $V = W^{1/2} X$ and note that $U' V = K' X = 0$, then (5) follows from the identity $U (U' U)^{-1} U' + V (V' V)^{-1} V' = I$.

Now, recall $\bar{s} = R' s$ and $\bar{F} = R' F R$, and note that

$$H' F^{-1} H = K' R^{-1} F^{-1} (R^{-1})' K = K' \bar{F}^{-1} K;$$

using this in the updating equation (2) enables us to rewrite it as

$$R_0^{-1} (\theta - \theta_0) = [\bar{F}_0^{-1} - \bar{F}_0^{-1} K (K' \bar{F}_0^{-1} K)^{-1} K' \bar{F}_0^{-1}] \bar{s}_0 + \bar{F}_0^{-1} K (K' \bar{F}_0^{-1} K) K' \bar{F}_0^{-1} \bar{F}_0 \eta_0. \quad (6)$$

Set $W = \bar{F}_0$ and note that (5) may be substituted into the first component of (6) and that its equivalent formulation

$$\bar{F}_0^{-1} K (K' \bar{F}_0^{-1} K)^{-1} K' \bar{F}_0^{-1} = \bar{F}_0^{-1} - X (X' \bar{F}_0 X)^{-1} X'$$

may be substituted into the second component, giving

$$\mathbf{R}_0^{-1}(\theta - \theta_0) = \mathbf{X}(\mathbf{X}'\overline{\mathbf{F}}_0\mathbf{X})^{-1}\mathbf{X}'\overline{\mathbf{s}}_0 - \eta_0 + \mathbf{X}(\mathbf{X}'\overline{\mathbf{F}}_0\mathbf{X})^{-1}\mathbf{X}'\overline{\mathbf{F}}_0\eta_0.$$

This is easily seen to be the same as combining equations (3) and (4).

Remark 2—From the form of the updating equations (2), (3) and (4) it is clear that Proposition 1 remains true if identical step length adjustments are applied to the $\boldsymbol{\theta}$ updates. This does not hold, however, if adjustments are applied to the $\boldsymbol{\beta}$ updates of the regression algorithm.

3.2.1. Derivation of the regression algorithm—In a neighbourhood of $\boldsymbol{\theta}_0$, approximate $l(\boldsymbol{\theta})$ by a quadratic function Q having the same information matrix and the same score vector as l at $\boldsymbol{\theta}_0$,

$$l(\boldsymbol{\theta}) \cong Q(\boldsymbol{\theta}) = -\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{t}_0)' \mathbf{F}_0(\boldsymbol{\theta} - \boldsymbol{t}_0), \quad \text{where } \boldsymbol{t}_0 = \boldsymbol{\theta}_0 + \mathbf{F}_0^{-1}\mathbf{s}_0.$$

Now compute a linear approximation of $\boldsymbol{\theta}$ with respect to $\boldsymbol{\beta}$ in a neighbourhood of $\boldsymbol{\theta}_0$,

$$\boldsymbol{\theta} - \boldsymbol{\theta}_0 \cong \mathbf{R}_0(\mathbf{X}\boldsymbol{\beta} - \eta_0); \quad (7)$$

substituting into the expression for Q we obtain a quadratic function in $\boldsymbol{\beta}$. By adding and subtracting $\mathbf{R}_0\mathbf{X}\boldsymbol{\beta}_0$ and setting $\boldsymbol{\delta} = \boldsymbol{\beta} - \boldsymbol{\beta}_0$, we have

$$Q(\boldsymbol{\beta}) = -\frac{1}{2}[\mathbf{R}_0\mathbf{X}\boldsymbol{\delta} - \mathbf{R}_0\gamma_0 - \mathbf{F}_0^{-1}\mathbf{s}_0]' \mathbf{F}_0[\mathbf{R}_0\mathbf{X}\boldsymbol{\delta} - \mathbf{R}_0\gamma_0 - \mathbf{F}_0^{-1}\mathbf{s}_0].$$

A weighted least square solution of this local maximization problem gives (3); substitution into (7) gives (4).

Remark 3—The choice of \mathbf{X} is somewhat arbitrary because the design matrix $\mathbf{X}\mathbf{A}$, where \mathbf{A} is any non-singular matrix, implements the same set of constraints as \mathbf{X} . In many cases an obvious choice for \mathbf{X} is provided by the context; otherwise, if we are not interested in the interpretation of $\boldsymbol{\beta}$, any numerical complement of \mathbf{K} will do.

3.3. Comparison of the two algorithms

Since the matrices \mathbf{C} and \mathbf{M} have dimensions $(t-1) \times u$ and $u \times t$ respectively, where the value of $u-t$ depends upon the particular parametrization, the hardest step in the Aitchison-Silvey's algorithm is $(\mathbf{K}'\mathbf{C}) \text{diag}(\mathbf{M}\boldsymbol{\pi})^{-1}\mathbf{M}$ whose computational complexity is $O(rut)$. In contrast, the hardest step in the regression algorithm is the computation of \mathbf{R} , which has computational complexity $O(u^2 + t^3)$, making this procedure clearly less efficient. However, the regression algorithm can be extended to models with individual covariates, a context in which it is usually much faster than a straightforward extension of the ordinary algorithm; see Section 4.

Note that because step adjustments, if used, are not made on the same scale, each algorithm may take a slightly different number of steps to converge.

3.4. Properties of the algorithms

Detailed conditions for the asymptotic existence of the maximum likelihood estimates of constrained models are given by Aitchison and Silvey (1958); see also Bergsma and Rudas (2002), Theorem 8. Much less is known about existence for finite sample sizes where estimates might fail to exist because of observed zeros. In this case, some elements of $\hat{\boldsymbol{\pi}}$ may converge to $\mathbf{0}$, leading the Jacobian matrix \mathbf{R} to become ill-conditioned and making the algorithm unstable.

Concerning the convergence properties of their algorithm, Aitchison and Silvey (1958, p. 827) noted only that it could be seen as a modified Newton algorithm and that similar modifications had been used successfully elsewhere. However, it is clear from the form of the updating equations that, if the algorithms converge to some $\boldsymbol{\theta}^*$, then the constraints $\mathbf{h}(\boldsymbol{\theta}^*) = \mathbf{0}$ are satisfied, and $\boldsymbol{\theta}^*$ is a stationary point of the constrained likelihood. In addition, as a consequence of the Karush-Kuhn-Tucker conditions, if a local maximum of the constrained objective function exists, then it will be a saddle point of the Lagrangian (see, for example, Bertsekas, 1999).

To ensure that a stationary point $\hat{\boldsymbol{\beta}}$ reached by the algorithm is indeed a local maximum, one could check that the observed information with respect to $\boldsymbol{\beta}$ is positive definite. An efficient formula for computing the observed information matrix is given in Appendix A. Since the log-likelihood of constrained marginal models is not, in general, concave, it might be advisable to apply the algorithm to a range of starting values, in order to check that the achieved maximum is the global one.

3.5. Extension to more general constraints

Occasionally, one may wish to fit general constraints on marginal probabilities without the need to define a marginal log-linear parameterization; an interesting example is provided by the *relational models* of Klimova et al. (2011). They consider constrained models of the form $\mathbf{h}(\boldsymbol{\theta}) = \mathbf{A} \log(\mathbf{M} \boldsymbol{\pi}) = \mathbf{0}$, where \mathbf{A} is an arbitrary matrix of full row rank. Redefine

$$\mathbf{K}' = \frac{\partial \mathbf{h}}{\partial \boldsymbol{\theta}} = \mathbf{A} \text{diag}(\mathbf{M}\boldsymbol{\pi})^{-1} \mathbf{M} \boldsymbol{\Omega} \mathbf{G}$$

and note that, because \mathbf{A} is not a matrix of row contrasts, \mathbf{h} is not homogeneous in \mathbf{p} , and thus the simplification of $\boldsymbol{\Omega}$ mentioned in Remark 1 does not apply. If the resulting model is smooth, implying that \mathbf{K} is a matrix of full column rank r everywhere in the parameter space, it can be fitted with the ordinary Aitchison-Silvey algorithm. We now show how the same model can also be fitted by a slight extension of the regression algorithm.

Let $\boldsymbol{\theta}_0$ be a starting value and $\bar{\mathbf{K}}_0$ be a right inverse of \mathbf{K}' at $\boldsymbol{\theta}_0$; consider a first order expansion of the constraints

$$\mathbf{h} = \mathbf{h}_0 + \mathbf{K}'_0(\boldsymbol{\theta} - \boldsymbol{\theta}_0) = \mathbf{K}'_0(\bar{\mathbf{K}}_0 \mathbf{h}_0 + \boldsymbol{\theta} - \boldsymbol{\theta}_0) = \mathbf{0}$$

and let \mathbf{X}_0 be a matrix that spans the orthogonal complement of \mathbf{K}_0 . Then, with the same order of approximation,

$$\bar{K}_0 h_0 + \theta - \theta_0 = X_0 \beta;$$

by solving the above equation for $\theta - \theta_0$ and substituting into the quadratic approximation of the log-likelihood, we obtain an updating equation similar to (3):

$$\widehat{\beta} - \beta_0 = (X_0' F_0 X_0)^{-1} X_0' [s_0 + F_0 (\bar{K}_0 h_0 - X_0 \beta_0)].$$

4. Modelling the effect of individual-level covariates

When exogenous individual-level covariates are available, it may be of interest to allow the marginal log-linear parameters to depend upon them as in a linear model: $\eta_i = C \log(M \pi_i) = X_i \beta$; here the matrix X_i specifies how the non-zero marginal log-linear parameters depend on individual specific information, in addition to structural restrictions such as conditional independencies. Let y_i , $i = 1, \dots, n$, be a vector of length t with a 1 in the entry corresponding to the response pattern of the i th individual, and all other values 0; define y to be the vector obtained by stacking the vectors y_i , one below the other. Alternatively, if the sample size is large and the covariates can take only a limited number of distinct values, y_i may contain the frequency table of the response variables within the sub-sample of subjects with the i th configuration of the covariates; in this case n denotes the number of strata. This arrangement avoids the need to construct a joint contingency table of responses and covariates; in addition the covariate configurations with no observations are simply ignored.

In either case, to implement the Aitchison-Silvey approach, stack the X_i matrices one below the other into the matrix X , and let K span the orthogonal complement of X ; as before, we have to fit the set of constraints $K' \eta = 0$. However, whilst q , the size of β , does not depend on the number of subjects, H is now of size $[n(t-1) - q] \times n(t-1)$, and its computation has complexity $O(n^3 t^2 u)$, where $u \sim t$ as before; in addition, the inversion of the $[n(t-1) - q] \times [n(t-1) - q]$ -matrix $H' F^{-1} H$ has complexity $O(n^3 t^3)$. With n moderately large, this approach becomes almost infeasible.

For the regression algorithm, let θ_i denote the vector of canonical parameters for the i th individual and $l(\theta_i) = y_i' G \theta_i - \log[1_t' \exp(G \theta_i)]$ be the contribution to the log-likelihood. Note that X_i need not be of full column rank, a property which must instead hold for the matrix X ; for this reason our assumptions are much weaker than those used by Lang (1996), and allow for more flexible models. Both the quadratic and the linear approximations must be applied at the individual level; thus we set $\theta_i - \theta_0 = R_{i0}(X_i \beta - \eta_{i0})$, and the log-likelihood becomes

$$\sum_{i=1}^n l(\theta_i) \cong -\frac{1}{2} \sum_{i=1}^n [R_{i0}(X_i \delta - \gamma_{i0}) - F_{i0}^{-1} s_{i0}]' F_{i0} [R_{i0}(X_i \delta - \gamma_{i0}) - F_{i0}^{-1} s_{i0}],$$

where $\gamma_{i0} = \eta_{i0} - X_i \beta_0$, $s_i = G'(y_i - \pi_i)$ and $F_i = G' \Omega_i G$.

Direct calculations lead to the updating expression

$$\widehat{\beta} - \beta_0 = \left(\sum_i \mathbf{X}_i' \mathbf{W}_i \mathbf{X}_i \right)^{-1} \left[\sum_i \mathbf{X}_i' (\mathbf{W}_i \gamma_{i0} + \mathbf{R}'_{i0} s_{i0}) \right],$$

where $\mathbf{W}_i = \mathbf{R}'_{i0} \mathbf{G}' \boldsymbol{\Omega}_{i0} \mathbf{G} \mathbf{R}_{i0}$. Thus, the procedure depends upon n only in that we have to sum across subjects, and so the complexity is $\mathcal{O}(n(t^2 u + t^3))$.

As an example of the utility of the method described above, consider the application to social mobility tables in Dardanoni et al. (2012). Social mobility tables are cross classifications of subjects according to their social class (columns) and that of their fathers (rows). The hypothesis of equality of opportunity would imply that the social class of sons is independent of that of their fathers. Mediating covariates may induce positive dependence between the social classes of fathers and sons, leading to the appearance of limited social mobility; to assess this, Dardanoni et al. (2012) fitted a model in which the vector of marginal parameters for each father-son pair was allowed to depend on individual covariates, including the father's age, the results of cognitive and non-cognitive test scores taken by the son at school, and his academic qualifications. The analysis, based on the UK's National Child Development Survey, included 1,942 father-son pairs classified in a 3×3 table. All marginal log-linear parameters for the father were allowed to depend on father's age, the only available covariate for fathers; the parameters for the son and the interactions were allowed to depend on all 11 available covariates. The fitted model used 76 parameters.

5. L_1 -penalized parameters

Evans (2011) shows that, in the context of marginal log-linear parameters, consistent model selection can be performed using the so-called adaptive lasso. Since the adaptive lasso uses L_1 -penalties, we might therefore be interested in relaxing the equality constraints discussed above to a penalization framework, in which we maximize the penalized log-likelihood

$$\varphi(\theta) \equiv l(\theta) - \sum_{j=1}^{t-1} v_j |\eta_j(\theta)|,$$

for some vector of penalties $\mathbf{v} = (v_j) \geq \mathbf{0}$.

The advantage of penalties of this form is that one can obtain parameter estimates which are exactly zero (Tibshirani, 1996). Setting parameters of the form $\boldsymbol{\eta}$ to zero corresponds to many interesting submodels, such as those defined by conditional independences, (Forcina et al., 2010; Rudas et al., 2010), we can therefore perform model selection without the need to fit many models separately. For now, assume that no equality constraints hold for $\boldsymbol{\eta}$, so we can take \mathbf{X} to be the identity, and $\boldsymbol{\beta} = \boldsymbol{\eta}$. This gives the quadratic form

$$Q(\boldsymbol{\eta}) = -\frac{1}{2} [\mathbf{R}_0(\boldsymbol{\eta} - \boldsymbol{\eta}_0) - \mathbf{F}_0^{-1} \mathbf{s}_0]' \mathbf{F}_0 (\mathbf{R}_0(\boldsymbol{\eta} - \boldsymbol{\eta}_0) - \mathbf{F}_0^{-1} \mathbf{s}_0)$$

approximating $l(\boldsymbol{\theta})$ as before. Then $\boldsymbol{\varphi}$ is approximated by

$$\tilde{\varphi}(\eta) \equiv -\frac{1}{2}[\mathbf{R}_0(\eta-\eta_0)-\mathbf{F}_0^{-1}\mathbf{s}_0]' \mathbf{F}_0[\mathbf{R}_0(\eta-\eta_0)-\mathbf{F}_0^{-1}\mathbf{s}_0] - \sum_j v_j |\eta_j|,$$

and we can attempt to maximize φ by repeatedly solving the sub-problem of maximizing $\tilde{\varphi}$. Now, because the quadratic form $Q(\eta)$ is concave and differentiable, and the absolute value function $|\cdot|$ is concave, coordinate-wise ascent is guaranteed to find a local maximum of $\tilde{\varphi}$ (Tseng, 2001). Coordinate-wise ascent cycles through $j = 1, 2, \dots, t-1$, at each step minimizing

$$-\frac{1}{2}[\mathbf{R}_0(\eta-\eta_0)-\mathbf{F}_0^{-1}\mathbf{s}_0]' \mathbf{F}_0[\mathbf{R}_0(\eta-\eta_0)-\mathbf{F}_0^{-1}\mathbf{s}_0] - v_j |\eta_j|$$

with respect to η_j with $\eta_1, \dots, \eta_{j-1}, \eta_{j+1}, \dots, \eta_{t-1}$ held fixed. This is solved just by taking

$$\eta_j = \text{sign}(\check{\eta}_j)(|\check{\eta}_j| - v_j)_+,$$

where $a_+ = \max\{a, 0\}$, and $\check{\eta}_j$ minimizes Q with respect to η_j (Friedman et al., 2010). This approach to the sub-problem may require a large number of iterations, but it is extremely fast in practice because each step is so simple. If the overall algorithm converges, then by a similar argument to that of Section 3.4, together with the fact that $\tilde{\varphi}$ has the same supergradient as φ at $\eta = \eta_0$, we see that we must have reached a local maximum of φ .

Since penalty selection for the lasso and adaptive lasso is typically performed using computationally intensive procedures such as cross validation, its implementation makes fast algorithms such as the one outlined above essential.

Acknowledgments

We thank two anonymous referees, the associate editor, and editor for their suggestions, corrections, and patience.

References

- Agresti, A. Categorical data analysis. John Wiley and Sons; 2002.
- Aitchison J, Silvey SD. Maximum-likelihood estimation of parameters subject to restraints. *Ann Math Stat.* 1958; 29(3):813–828.
- Bartolucci F, Colombi R, Forcina A. An extended class of marginal link functions for modelling contingency tables by equality and inequality constraints. *Statist Sinica.* 2007; 17(2):691.
- Bergsma, W.; Croon, M.; Hagenaars, JA. *Marginal Models: For Dependent, Clustered, and Longitudinal Categorical Data.* Springer Verlag; 2009.
- Bergsma, WP. *Marginal models for categorical data.* Tilburg University Press; Tilburg: 1997.
- Bergsma WP, Rudas T. Marginal models for categorical data. *Ann Statist.* 2002; 30(1):140–159.
- Bertsekas, DP. *Nonlinear programming.* 2. Athena Scientific; 1999.
- Colombi R, Forcina A. Marginal regression models for the analysis of positive association of ordinal response variables. *Biometrika.* 2001; 88(4):1001–1019.
- Dardanoni V, Fiorini M, Forcina A. Stochastic monotonicity in intergen-erational mobility tables. *J Appl Econometrics.* 2012; 27:85–107.
- Evans, RJ. PhD thesis. University of Washington; 2011. Parametrizations of discrete graphical models.

- Evans RJ, Richardson TS. Marginal log-linear parameters for graphical markov models. *J Roy Statist Soc B*. To appear 2013
- Forcina A, Lupporelli M, Marchetti GM. Marginal parameterizations of discrete models defined by a set of conditional independencies. *Journ Mult Analysis*. 2010; 101(10):2519–2527.
- Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J Stat Soft*. 2010; 33(1)
- Glonek GFV, McCullagh P. Multivariate logistic models. *J R Statist Soc B*. 1995; 57(3):533–546.
- Klimova A, Rudas T, Dobra A. Relational models for contingency tables. 2011 arXiv:1102.5390.
- Lang JB. Maximum likelihood methods for a generalized class of log-linear models. *Ann Statist*. 1996; 24(2):726–752.
- Lang JB, Agresti A. Simultaneously modeling joint and marginal distributions of multivariate categorical responses. *J Amer Statist Assoc*. 1994; 89(426):625–632.
- Liang KY, Zeger SL, Qaqish B. Multivariate regression analyses for categorical data. *J R Statist Soc B*. 1992; 54(1):3–40.
- McCullagh, P.; Nelder, JA. *Generalized linear models*. Chapman & Hall/CRC; 1989.
- Rudas T, Bergsma WP, Németh R. Marginal log-linear parameterization of conditional independence models. *Biometrika*. 2010; 97(4):1006–1012.
- Tibshirani R. Regression shrinkage and selection via the lasso. *J Royal Statist Soc B*. 1996; 58(1):267–288.
- Tseng P. Convergence of a block coordinate descent method for nondifferentiable minimization. *J Optim Theory Appl*. 2001; 109(3):475–494.

Appendix A. Computation of the observed information matrix

Lemma 1

Suppose that \mathbf{A} is a $p \times q$ matrix, and that \mathbf{y} , \mathbf{b} , \mathbf{x} and \mathbf{u} are column vectors with respective lengths q , p , k and r . Then if \mathbf{A} and \mathbf{b} are constant,

$$\frac{\partial}{\partial \mathbf{x}'} \text{diag}(\mathbf{A}\mathbf{y})\mathbf{b} = \text{diag}(\mathbf{b})\mathbf{A} \frac{\partial \mathbf{y}}{\partial \mathbf{u}'} \frac{\partial \mathbf{u}}{\partial \mathbf{x}'}. \quad (\text{A.1})$$

Proof

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}'} \text{diag}(\mathbf{A}\mathbf{y})\mathbf{b} &= \frac{\partial}{\partial \mathbf{u}'} \text{diag}(\mathbf{A}\mathbf{y})\mathbf{b} \frac{\partial \mathbf{u}}{\partial \mathbf{x}'} \\ &= (\text{diag}(\mathbf{A}\mathbf{y}_{u_1})\mathbf{b}, \dots, \text{diag}(\mathbf{A}\mathbf{y}_{u_h})\mathbf{b}) \frac{\partial \mathbf{u}}{\partial \mathbf{x}'} \\ &= (\text{diag}(\mathbf{b})\mathbf{A}\mathbf{y}_{u_1}, \dots, \text{diag}(\mathbf{b})\mathbf{A}\mathbf{y}_{u_h}) \frac{\partial \mathbf{u}}{\partial \mathbf{x}'} \\ &= \text{diag}(\mathbf{b})\mathbf{A} \frac{\partial \mathbf{y}}{\partial \mathbf{u}'} \frac{\partial \mathbf{u}}{\partial \mathbf{x}'}. \end{aligned}$$

The observed information matrix is minus the second derivative of the log-likelihood with respect to $\boldsymbol{\beta}$, that is

$$\begin{aligned} -\frac{\partial}{\partial \boldsymbol{\beta}'} \left[\frac{\partial l(\boldsymbol{\theta})}{\partial \boldsymbol{\beta}} \right] &= -\frac{\partial}{\partial \boldsymbol{\beta}'} \mathbf{X}' \mathbf{R}' \mathbf{G}' (\mathbf{y} - n\boldsymbol{\pi}) = - \left[\frac{\partial}{\partial \boldsymbol{\theta}'} \mathbf{X}' \mathbf{R}' \mathbf{G}' (\mathbf{y} - n\boldsymbol{\pi}) \right] \mathbf{R}\mathbf{X} \\ &= n\mathbf{X}' \mathbf{R}' \mathbf{G}' \boldsymbol{\Omega} \mathbf{G} \mathbf{R} \mathbf{X} - \mathbf{X}' \frac{\partial \mathbf{R}}{\partial \boldsymbol{\theta}'} \mathbf{G}' (\mathbf{y} - n\boldsymbol{\pi}) \mathbf{R} \mathbf{X}. \end{aligned}$$

Since s depends on $\boldsymbol{\theta}$ through both $(\mathbf{y} - n\boldsymbol{\pi})$ and \mathbf{R} , the above derivative has two main components, where the one obtained by differentiating $(\mathbf{y} - n\boldsymbol{\pi})$ is minus the expected

information. Using the well known expression for the derivative of an inverse matrix, it only remains to compute

$$\begin{aligned} X' \frac{\partial R'}{\partial \theta'} G' (y - n\pi) R X \\ &= X' R' \frac{\partial R'^{-1}}{\partial \theta'} R' G' (y - n\pi) R X \\ &= A \frac{\partial R'^{-1}}{\partial \theta'} b R X \end{aligned}$$

where $A = X' R'$ and $b = R' G' (y - n\pi)$, giving

$$= A G' \frac{\partial [\text{diag}(\pi) M' \text{diag}(M\pi)^{-1}]}{\partial \theta'} C' b R X.$$

By two applications of (A.1), this is

$$A G' \left[\text{diag}(M' \text{diag}(M\pi)^{-1} C' b) - \text{diag}(\pi) M' \text{diag}(C' b) \text{diag}(M\pi)^{-2} M \right] \Omega G R X.$$