Subspace decompositions for association structure learning in multivariate categorical response regression

Hongru Zhao[©], ^{1,*} Aaron J. Molstad^{1,2} and Adam J. Rothman¹

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

Abstract

Modeling the complex relationships between multiple categorical response variables as a function of predictors is a fundamental task in the analysis of categorical data. However, existing methods can be difficult to interpret and may lack flexibility. To address these challenges, we introduce a penalized likelihood method for multivariate categorical response regression that relies on a novel subspace decomposition to parameterize interpretable association structures. Our approach models the relationships between categorical responses by identifying mutual, joint, and conditionally independent associations, which yields a linear problem within a tensor product space. We establish theoretical guarantees for our estimator, including error bounds in high-dimensional settings, and demonstrate the method's interpretability and prediction accuracy through comprehensive simulation studies.

Key words: multinomial logistic regression, log-linear models, categorical data analysis, Potts model, hypergraph model

1. Introduction

We consider a multivariate response regression where each of the response variables is categorical. Specifically, let $X \in \mathcal{X} \subseteq \mathbb{R}^p$ be the predictor vector and let $Z = (Z_1, \dots, Z_q)^{\top}$ be the multivariate categorical response. The kth component of the response, Z_k , has J_k numerically coded outcome categories with $J_k \geq 2$ for $k \in [q]$, where [m] is defined as $\{1, \dots, m\}$ for positive integer m. The essential problem is to model the conditional distribution Z|X=x whose joint probability mass function is given by

$$\pi_j(x) := \mathbb{P}(Z_1 = j_1, \dots, Z_q = j_q | X = x) \ge 0,$$
 (1)

for any $j=(j_1,\cdots,j_q)\in J:=[J_1]\times\cdots\times [J_q]$, where $j_l\in [J_l]$ for all $l\in [q]$. For a given $\boldsymbol{x},\,\boldsymbol{Z}$ has a multivariate version of the single-trial multinomial distribution. If, for a given \boldsymbol{x} , one were to observe $v\geq 1$ independent realizations of \boldsymbol{Z} , say $\boldsymbol{z}_1,\ldots,\boldsymbol{z}_v$, then the probability mass function corresponding to (1) would be given by

$$\frac{v!}{\prod_{j\in J} y_j!} \prod_{j\in J} \left\{ \pi_j(x) \right\}^{y_j},$$

¹School of Statistics, University of Minnesota, 224 Church St SE, Minneapolis, MN 55455, USA and ²Department of Statistics, University of Florida, 227 Hull Rd, Gainesville, 32611, USA

^{*}Address for correspondence: Hongru Zhao, School of Statistics, University of Minnesota, Minneapolis, MN 55455, USA. zhao1118@umn.edu

[©] The Author XXXX. Published by Oxford University Press. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

where $y_j := \sum_{i=1}^{v} \mathbf{1}(z_i = j)$ for each $j \in J$.

For a given x, if v is sufficiently large, one could model (1) using standard methods for the analysis of q-way contingency tables, a classical problem in categorical data analysis (McCullagh and Nelder, 1989; Christensen, 1997; Agresti, 2002). However, when one needs to model (1) for all $x \in \mathcal{X}$, methods for contingency tables cannot be applied. For example, in many applications, for every subject in the study we observe (or impose) a distinct x, and observe the outcome of only a single trial, v = 1. Instead, one could model (1) using existing methods for multinomial regression (or in statistical learning terminology, multiclass classification). Notice that (1) could be equivalently defined in terms of a "univariate" categorical response variable, Z^* , with the cardinality of J, |J|, many outcome categories: one corresponding to each distinct element of J. Letting $f: J \to [|J|]$ be any bijective function, it would thus be natural to model $\mathbb{P}(Z^* = f(j) \mid X = x) = \pi_j(x) = \mathbb{P}(Z_1 = j_1, \cdots, Z_q = j_q \mid X = x)$ using multinomial logistic regression (Agresti, 2002; Vincent and Hansen, 2014); linear or quadratic discriminant analysis (Hastie et al., 2009; Mai et al., 2019); or nonparametric methods. Modeling the conditional distribution $Z^* \mid X$ using one of these methods is appealing because they allow for arbitrary dependence among the q categorical response variables.

However, off-the-shelf application of methods designed for a univariate categorical response may be problematic. In particular, these methods would fail to exploit that Z^* is constructed from q distinct response variables. This negatively affects both estimation efficiency and interpretability of the fitted model. Moreover, for even moderate q, |J| will be large. As a consequence, with small sample sizes, many outcome category combinations j will not be observed in the training data. If, for example, one used a multinomial logistic regression in this situation, the maximum likelihood estimator would not exist. In this work, we propose a new method for fitting (1) that allows practitioners to discover parsimonious and interpretable dependence structures among responses, and exploits the multivariate nature of the response.

To motivate our approach, consider a multinomial logistic regression model for (1) with $x \in \mathbb{R}$ (i.e., p = 1),

$$\pi_{j}(\boldsymbol{x}) = \mathbb{P}(Z_{1} = j_{1}, \dots, Z_{q} = j_{q} | \boldsymbol{X} = \boldsymbol{x}) = \frac{\exp(\boldsymbol{x} \cdot \boldsymbol{\zeta}_{j})}{\sum_{j \in J} \exp(\boldsymbol{x} \cdot \boldsymbol{\zeta}_{j})}, \quad \boldsymbol{j} \in \boldsymbol{J}, \quad \sum_{j \in J} \boldsymbol{\zeta}_{j} = 0,$$
(2)

where $\zeta = \{\zeta_{\boldsymbol{j}}\}_{\boldsymbol{j} \in J}$ is an unknown tensor. In full generality, $\zeta \in \{v \in \mathbb{R}^{[J_1] \times \cdots \times [J_q]} : \sum_{\boldsymbol{j} \in J} v_{\boldsymbol{j}} = 0\}$, which implies no restrictions on the dependence amongst responses: their dependence can be arbitrarily complex. Restrictions on the dependence between responses under (1) can often be represented as constraints on the space of the coefficients ζ . For example, in the case that q = 2, $J_1 = J_2 = 2$, if $\zeta \in \mathfrak{C}^0$, where

$$\mathfrak{C}^0 = \{ \boldsymbol{\zeta} \in \mathbb{R}^{[2] \times [2]} : \boldsymbol{\zeta}_{(1,1)} + \boldsymbol{\zeta}_{(2,2)} - \boldsymbol{\zeta}_{(1,2)} - \boldsymbol{\zeta}_{(2,1)} = 0 \},$$

then $Z_1 \perp \!\!\! \perp Z_2 \mid X$. Intuitively, \mathfrak{C}^0 is the set of coefficients for which the log odds ratio between the two responses is zero for all x. This observation motivated Molstad and Rothman (2023) to propose a regularized maximum likelihood estimator of ζ that shrinks coefficients towards the set \mathfrak{C}^0 . For applications with $J_l \geq 2$ and $q \leq 3$, Molstad and Rothman (2023) generalized the set \mathfrak{C}^0 to correspond to coefficients with all local log odds ratios equal to zero. Their approach thus allowed practitioners to discover only whether responses are mutually independent ($\zeta \in \mathfrak{C}^0$) or are arbitrarily dependent ($\zeta \notin \mathfrak{C}^0$). When $q \geq 3$, however, there are many other parsimonious dependence structures that are "intermediate" to mutual independence and arbitrary dependence. In this work, we generalize the approach of Molstad and Rothman (2023), allowing practitioners to discover much more complex, interpretable dependence structures.

As we just described, to learn the association structure for (1), it is crucial to identify whether the regression coefficients reside within a specific subspace. Representing the linear subspace $\mathfrak L$ of $\mathbb R^k$ can be approached in two ways: externally and internally. For the external representation, consider $\mathfrak L = \ker(A) = \{v \in \mathbb R^k; Av = 0\}$ for some matrix A. Then regularization of v towards the subspace $\ker(A)$ can be achieved by penalizing the term $\|Av\|_2$. In this sense, Molstad and Rothman (2023) achieves structure learning via an external subspace representation. In contrast, for the internal representation, we can set $\mathfrak L = \operatorname{span}\left(\{e_1,\ldots,e_s\}\right)$ and $v = v_1e_1 + v_2e_2 + \cdots + v_ke_k$, where e_1,\ldots,e_k form an orthonormal basis for $\mathbb R^k$. Then regularization of v towards the subspace $\operatorname{span}\left(\{e_1,\ldots,e_s\}\right)$ can be achieved by penalizing the terms v_{s+1},\cdots,v_k . This is the approach we take in this paper by selecting an orthonormal basis and penalizing the coordinates to achieve association structure learning.

For example, if $J_1 = J_2 = 2$, we can define

$$\boldsymbol{\zeta}^* = v_{11}g_1g_1^\top + v_{12}g_1g_2^\top + v_{21}g_2g_1^\top + v_{22}g_2g_2^\top,$$

where $g_i = \frac{1}{\sqrt{2}} (1, (-1)^{i+1})^{\top}, i \in \{1, 2\}$. Here, $g_1 g_1^{\top}$ represents the *overall association*, $g_1 g_2^{\top}$ denotes the *main association* of the first response, $g_2 g_1^{\top}$ is for the main association for the second response, and $g_2 g_2^{\top}$ captures the *interaction association* between the two responses.

Because $\mathfrak{C}^0 = \operatorname{span}\{g_1g_1^\top, g_1g_2^\top, g_2g_1^\top\}$, if $v_{22} = 0$, then $\zeta^* \in \mathfrak{C}^0$. That is, by carefully constructing the internal subspace representation, sparsity in the corresponding coefficients can imply parsimonious association structures among responses. This observation is central to our methodological developments, and one of our main contributions is the explicit construction of a flexible, interpretable internal subspace representation.

Multivariate categorical response regression without predictors serves as an extension of contingency table analysis, allowing for a more comprehensive examination of categorical variable interrelations. The Poisson log-linear model is used for association structure modeling of multiple categorical responses without predictors, with the connection between log-linear models for frequencies and multinomial response models for proportions being extensively studied (McCullagh and Nelder, 1989; Christensen, 1997).

In this paper, we will study structure learning via an internal subspace representation. We present a reparameterization via subspace decomposition and obtain a unifying framework for both multinomial and Poisson categorical response regression models in high dimensions. Complex dependencies between response variables can be systematically modeled, encompassing all possible association structures, including mutual independence, joint independence, and conditional independence among response variables. To achieve structure learning, we utilize variations of the group lasso (Yuan and Lin, 2006) and overlapping group lasso (Zhao et al., 2009; Jenatton et al., 2011) penalties. We apply accelerated proximal gradient descents algorithm to solve the corresponding convex optimization problems. We prove an error bound that illustrates our estimator's performance in high-dimensional settings. A key theoretical development is our derivation of restricted strong convexity conditions specific to multivariate categorical response regression, which notably incorporates the Rademacher complexity associated with general norm penalties. Finally, simulation studies validate our method's effectiveness in terms of interpretability and prediction accuracy.

We conclude this section by introducing notation to be used for the remainder of the article. First, let $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ denote the maximum and minimum eigenvalues of the real symmetric matrix A. For any vector (resp. matrix) X, define the Euclidean (resp. Frobenius) norm $||X|| = \sqrt{\operatorname{tr}(X^{\top}X)}$. Let $\mathbf{1}_m$ denote a vector of ones of length m. For matrices X and Y of the same size, define the Frobenius inner product $\langle X,Y\rangle = \operatorname{tr}(X^{\top}Y)$ and define the operator norm $||X||_{\operatorname{op}} = \sqrt{\lambda_{\max}(X^{\top}X)}$. Define the maximum norm $||X||_{\infty} = \max_{i,j} |x_{ij}|, X = \{x_{i,j}\}_{1 \leq i \leq n, 1 \leq j \leq m}$. Let $||\beta||_{0} = \sum_{j,k} \mathbf{1}(\beta_{j,k} \neq 0)$ for matrix β . Let I be the identity matrix. When X and Y are matrices, let $X \otimes Y$ denote the Kronecker product between X and Y. When U and V are vector spaces, let $U \otimes V$ be the tensor product of U and V. Finally, let $\otimes_t(u,v)$ denote the tensor product between vector u and v.

2. Association structure learning via subspace decomposition

2.1. Overview

Assume the response has $q \geq 2$ categorical components with J_1, \cdots, J_q categories, respectively. Define the Cartesian product of an indexed family of sets $J = [J_1] \times [J_2] \times \cdots \times [J_q]$. The cardinality of set J is $|J| = \prod_{i=1}^q J_i$. Let \mathbb{R}^J and \mathbb{N}^J denote the spaces of J arrays with entries that are real numbers and whole numbers, respectively. That is, $\mathbf{y} = \{y_j\}_{j \in J} \in \mathbb{F}^J$ if and only if $y_j \in \mathbb{F}$ for any $j \in J$, where field \mathbb{F} can take \mathbb{R} and \mathbb{N} . For a q-way array of shape J, let $\mathbf{y}^J = \{y_j\}_{j \in J} \in \mathbb{R}^J$, where $\mathbf{j} = (j_1, \cdots, j_q), j_1 \in [J_1], \cdots, j_q \in [J_q]$. Define the J-vectorization of \mathbf{y}^J as

$$\operatorname{vec}_{\mathbf{J}}(\mathbf{y}^{\mathbf{J}}) := (y_{1,1,\dots,1}, y_{2,1,\dots,1}, \dots, y_{J_1,1,\dots,1}, y_{1,2,\dots,1}, \dots, y_{J_1,J_2,\dots,J_q})^{\top} \in \mathbb{R}^{|\mathbf{J}|}.$$
 (3)

Define the inverse J-vectorization for vector $\text{vec}_J(y^J)$ as $\text{vec}_J^{-1}(\text{vec}_J(y^J)) = y^J$.

Let the sample data from the *i*th observational unit be denoted (x_i, y_i^J) , where $x_i \in \mathbb{R}^p$ and $y_i^J \in \mathbb{N}^J$, and let $y_i = \text{vec}_J(y_i^J) \in \mathbb{N}^{|J|}$. We assume that y_1^J, \dots, y_n^J are independent. Similar to Molstad and Rothman

(2023), we use a multinomial logistic regression model for the q response variables. Specifically, we assume that y_i is a multinomial random vector with index $n_i \ge 1$ and category probabilities

$$\operatorname{vec}_{\mathbf{J}}\{\boldsymbol{\pi}^{\mathbf{J}}(\boldsymbol{x}_i)\} = \frac{e^{\boldsymbol{\theta}\boldsymbol{x}_i}}{\langle \mathbf{1}_{|\mathbf{J}|}, e^{\boldsymbol{\theta}\boldsymbol{x}_i} \rangle} \in \mathbb{R}^{|\mathbf{J}|}, \tag{4}$$

where $\theta \in \mathbb{R}^{|J| \times p}$. In certain settings, we may treat the elements of y_i as independent Poisson random variables with mean

$$\operatorname{vec}_{J}\{\mu^{J}(x_{i})\} = e^{\theta x_{i}} \in \mathbb{R}^{|J|}, \tag{5}$$

and category probabilities $\text{vec}_{J}\{\pi^{J}(x_{i})\}=e^{\theta x_{i}}\{\langle \mathbf{1}_{|J|},e^{\theta x_{i}}\rangle\}^{-1}$, where $\mu^{J}(x)=\{\mu_{j}(x)\}_{j\in J}$ and $\pi^{J}(x)=\{\pi_{j}(x)\}_{j\in J}$ are q-dimensional arrays.

The ith observational unit's contribution to the negative log-likelihood of multinomial and Poisson categorical response models are

$$\ell_{\text{Mult}}(\boldsymbol{\theta}\boldsymbol{x}_{i},\boldsymbol{y}_{i}) = -\langle \boldsymbol{y}_{i}, \boldsymbol{\theta}\boldsymbol{x}_{i} \rangle + n_{i} \cdot \log\left(\langle \mathbf{1}_{|\boldsymbol{J}|}, e^{\boldsymbol{\theta}\boldsymbol{x}_{i}} \rangle\right), \qquad n_{i} = \langle \mathbf{1}_{|\boldsymbol{J}|}, \boldsymbol{y}_{i} \rangle, \tag{6}$$

and

$$\ell_{\text{Pois}}(\boldsymbol{\theta}\boldsymbol{x}_{i}, \boldsymbol{y}_{i}) = -\langle \boldsymbol{y}_{i}, \boldsymbol{\theta}\boldsymbol{x}_{i} \rangle + \langle \mathbf{1}_{|\boldsymbol{J}|}, e^{\boldsymbol{\theta}\boldsymbol{x}_{i}} \rangle. \tag{7}$$

The function in (6) is sometimes referred to as cross-entropy loss.

When considering the multinomial categorical response model, we impose the constraint $\mathbf{1}_{|J|}^{\top} \theta = \mathbf{0}$ for identifability. Define the linear subspace $\mathcal{V} = \{\alpha \in \mathbb{R}^{|J|} : \mathbf{1}_{|J|}^{\top} \alpha = \mathbf{0}\}$, and define the orthogonal projection matrix $P_{\mathcal{V}} = (I - |J|^{-1}\mathbf{1}_{|J|}\mathbf{1}_{|J|}^{\top})$. Notice that $\ell_{\text{Mult}}(\theta x, y) = \ell_{\text{Mult}}(\theta' x, y)$ for any (x, y) if and only if $P_{\mathcal{V}}\theta = P_{\mathcal{V}}\theta'$ (see Lemma S1 in the online supplementary material).

2.2. Subspace decomposition

We now introduce the subspace decomposition that allows us to parsimoniously model the mass function of interest. Naturally, the dependence between response variables is arbitrarily complex when $\theta \in \mathbb{R}^{|J| \times p}$ without additional constraints. To discover parsimonious association structures, we decompose θ into a sum of components, each of which spans a particular subspace. Returning to an example from the introduction, when q=2, we can decompose $\theta=H_{\{0\}}\beta_{\{0\}}+H_{\{1\}}\beta_{\{1\}}+H_{\{2\}}\beta_{\{2\}}+H_{\{1,2\}}\beta_{\{1,2\}}$, where each H is a basis matrix and β are the corresponding coefficients. With appropriately constructed H, if $\beta_{\{1,2\}}=0$, then the two response variables are independent. We demonstrate how to construct such bases H in the following example.

Example 1 (Subspace decomposition of a $J_1 \times J_2$ contingency table) Consider the intercept only model (i.e., p=1 with $x_i=1$) with q=2, and the categorical responses having $J_1=2$ categories for the first response and $J_2=3$ categories for the second. We can write $\boldsymbol{\theta}=(a_{11},\ldots,a_{J_1J_2})^{\top}\in\mathbb{R}^{J_1J_2}$ as

$$\left(\operatorname{vec}_{J}^{-1}(\boldsymbol{\theta})\right)_{j_{1},j_{2}} = a_{j_{1},j_{2}},$$
 (8)

where $a_{j_1,j_2} \in \mathbb{R}$ for any $(j_1,j_2) \in J$. Accordingly, we can rewrite (8) as

$$\boldsymbol{\theta} = \sum_{j_2=1}^{J_2} \sum_{j_1=1}^{J_1} a_{j_1,j_2} \boldsymbol{e}_{j_2}^{J_2} \otimes \boldsymbol{e}_{j_1}^{J_1} \ \ and \ \ \text{vec}_{\boldsymbol{J}}^{-1} \left(\boldsymbol{\theta}\right) = \sum_{j_2=1}^{J_2} \sum_{j_1=1}^{J_1} a_{j_1,j_2} \boldsymbol{E}_{j_1j_2}^{\boldsymbol{J}},$$

where $e_{j_i}^{J_i}$ is the j_i th standard basis vector for \mathbb{R}^{J_i} (i.e., the j_i th column of I_{J_i}) for $i \in [2]$, and $E_{j_1j_2}^J$ denotes the standard basis 2-way array for $\mathbb{R}^{J_1 \times J_2}$, which is defined as the array whose (j_1, j_2) -th entry is 1 and all other entries are 0.

Define U_m as a matrix such that $[\frac{1}{\sqrt{m}}\mathbf{1}_m, U_m]$ is an orthogonal matrix of order m. Let $\mathcal{R}(U)$ denote the column space of the matrix U. Then, we can rewrite \mathbb{R}^{J_i} as the internal direct sum between $\mathcal{R}(U_{J_i})$ and

 $\mathcal{R}(\mathbf{1}_{J_i})$, denoted as $\mathbb{R}^{J_i} = \mathcal{R}(U_{J_i}) \oplus \mathcal{R}(\mathbf{1}_{J_i})$. The definitions of the internal direct sum and tensor product can be found in Section S1 of the online supplementary material.

Due to the bilinearity of the tensor product, $\mathbb{R}^{J_2} \otimes \mathbb{R}^{J_1}$ can be decomposed into an internal direct sum

$$\begin{split} \mathbb{R}^{J_2} \otimes \mathbb{R}^{J_1} = & \{ \mathcal{R}(\boldsymbol{U}_{J_2}) \oplus \mathcal{R}(\boldsymbol{1}_{J_2}) \} \otimes \{ \mathcal{R}(\boldsymbol{U}_{J_1}) \oplus \mathcal{R}(\boldsymbol{1}_{J_1}) \} \\ = & \{ \mathcal{R}(\boldsymbol{1}_{J_2}) \otimes \mathcal{R}(\boldsymbol{1}_{J_1}) \} \oplus \{ \mathcal{R}(\boldsymbol{1}_{J_2}) \otimes \mathcal{R}(\boldsymbol{U}_{J_1}) \} \oplus \{ \mathcal{R}(\boldsymbol{U}_{J_2}) \otimes \mathcal{R}(\boldsymbol{1}_{J_1}) \} \\ \oplus & \{ \mathcal{R}(\boldsymbol{U}_{J_2}) \otimes \mathcal{R}(\boldsymbol{U}_{J_1}) \}. \end{split}$$

Consider an isomorphism T from the tensor product $\mathbb{R}^{J_2} \otimes \mathbb{R}^{J_1}$ to $\mathbb{R}^{|J|}$. The isomorphism T is uniquely determined by the change of basis $T(\otimes_t(e_{j_2}^{J_2}, e_{j_1}^{J_1})) = e_{j_2}^{J_2} \otimes e_{j_1}^{J_1}$, where the tensor product $\otimes_t(u, v)$ denotes the bilinear map of (u, v) from the Cartesian product $\mathbb{R}^{J_2} \times \mathbb{R}^{J_1}$, whose basis can be chosen as $\{ \otimes_t (e_{j_2}^{J_2}, e_{j_1}^{J_1}); 1 \leq v \}$ $j_1 \leq J_1, 1 \leq j_2 \leq J_2$.

Applying the isomorphism T onto each subspace, we obtain that

$$T(\mathcal{R}(\boldsymbol{V}_2) \otimes \mathcal{R}(\boldsymbol{V}_1)) = \mathcal{R}(\boldsymbol{V}_2 \otimes \boldsymbol{V}_1), \ \ \text{where} \ \boldsymbol{V}_i \in \left\{\frac{1}{\sqrt{J_i}} \mathbf{1}_{J_i}, \boldsymbol{U}_{J_i}\right\} \ \text{for} \ i \in [2].$$

Hence, for $\theta \in \mathbb{R}^{|J|}$, we can write

$$heta = \sum_{m{V}_1 \in \left\{ rac{1}{\sqrt{J_1}} m{1}_{J_1}, m{U}_{J_1}
ight\} m{V}_2 \in \left\{ rac{1}{\sqrt{J_2}} m{1}_{J_2}, m{U}_{J_2}
ight\}} \Big(m{V}_2 \otimes m{V}_1 \Big) m{lpha}_{m{V}_2, m{V}_1}$$

for vectors α_{V_2,V_1} of appropriate size. More simply, we may write $\theta = \sum_{k \in \mathcal{K}} H_k \beta_k$, with

$$egin{aligned} &m{H}_{\{0\}} = rac{1}{\sqrt{|J|}} \mathbf{1}_{|J|}, \quad m{H}_{\{1\}} = rac{1}{\sqrt{J_2}} \mathbf{1}_{J_2} \otimes m{U}_{J_1}, \quad m{H}_{\{2\}} = m{U}_{J_2} \otimes rac{1}{\sqrt{J_1}} \mathbf{1}_{J_1} \\ &m{H}_{\{1,2\}} = m{U}_{J_2} \otimes m{U}_{J_1}, \quad \mathcal{K} = \{\{0\}, \{1\}, \{2\}, \{1, 2\}\}, \end{aligned}$$

where each β_k is simply the corresponding α_{V_2,V_1} . As we will formalize in Lemma 1, Lemma 2, and Theorem 1, $\operatorname{span}(H_{\{0\}})$ is the subspace for overall association, $\operatorname{span}(H_{\{1\}})$, $\operatorname{span}(H_{\{2\}})$ are the subspaces for marginal association on Z_1 and Z_2 , respectively, and $\operatorname{span}(H_{\{1,2\}})$ is the subspace for joint association on (Z_1, Z_2) . Because of this, sparsity in the coefficients corresponding to each subspace can imply an interpretable restriction on θ .

The discussion outlined above can be generalized to any J, with the corresponding isomorphism, denoted as T_{J} , being applicable to each subspace.

Lemma 1 (Isomorphism) Define an isomorphism T_J from the tensor product space $\mathbb{R}^{J_q} \otimes \cdots \otimes \mathbb{R}^{J_1}$ to $\mathbb{R}^{|J|}$, which is uniquely determined by the change of basis $T_{\boldsymbol{J}}(\otimes_t(\boldsymbol{e}_{j_q}^{J_q},\cdots,\boldsymbol{e}_{j_1}^{J_1})) = \boldsymbol{e}_{j_q}^{J_q} \otimes \cdots \otimes \boldsymbol{e}_{j_1}^{J_1}. \text{ Here, } \{\otimes_t(\boldsymbol{e}_{j_q}^{J_q},\cdots,\boldsymbol{e}_{j_1}^{J_1})\}_{(j_1,\cdots,j_q) \in \boldsymbol{J}} \text{ and }$

$$T_{\boldsymbol{J}}(\otimes_t(e_{j_q}^{\gamma_q},\cdots,e_{j_1}^{\gamma_1})) = e_{j_q}^{\gamma_q}\otimes\cdots\otimes e_{j_1}^{\gamma_1}.$$
 Here, $\{\otimes_t(e_{j_q}^{\gamma_q},\cdots,e_{j_1}^{\gamma_1})\}_{(j_1,\cdots,j_q)\in\boldsymbol{J}}$ and $\{e_{j_q}^{J_q}\otimes\cdots\otimes e_{j_1}^{J_1}\}_{(j_1,\cdots,j_q)\in\boldsymbol{J}}$ denote the basis of $\mathbb{R}^{J_q}\otimes\cdots\otimes\mathbb{R}^{J_1}$ and $\mathbb{R}^{|\boldsymbol{J}|}$, respectively. Then,

(i) for any vector $\mathbf{v}_i \in \mathbb{R}^{J_i}$, $i \in [q]$, we have

$$T_{\mathbf{J}}\Big(\otimes_t (\mathbf{v}_q, \cdots, \mathbf{v}_1)\Big) = \mathbf{v}_q \otimes \cdots \otimes \mathbf{v}_1,$$
 (9)

(ii) and for any $V_i \in \{\frac{1}{\sqrt{J_i}} \mathbf{1}_{J_i}, U_{J_i}\}, i \in [q], \text{ we have }$

$$T_{J}(\mathcal{R}(V_{q}) \otimes \cdots \otimes \mathcal{R}(V_{1})) = \mathcal{R}(V_{q} \otimes \cdots \otimes V_{1}). \tag{10}$$

In Lemma 1, the isomorphism T_J is not only a vector space isomorphism between $\mathbb{R}^{J_q} \otimes \cdots \otimes \mathbb{R}^{J_1}$ and $\mathbb{R}^{|J|}$, i.e.,

$$T_{J}(\boldsymbol{v}+\boldsymbol{w})=T_{J}(\boldsymbol{v})+T_{J}(\boldsymbol{w}), \quad T_{J}(a\boldsymbol{v})=a\cdot T_{J}(\boldsymbol{v}), \quad \boldsymbol{v},\boldsymbol{w}\in\mathbb{R}^{J_{q}}\otimes\cdots\otimes\mathbb{R}^{J_{1}}, \quad a\in\mathbb{R},$$

but also preserves the bilinear structure, i.e., equation (9) holds. Here, the Kronecker product serves as the bilinear function.

To summarize Example 1 and Lemma 1, we define U_m as any matrix such that $\left[\frac{1}{\sqrt{m}}\mathbf{1}_m, U_m\right]$ is an orthogonal matrix of order m. Without loss of generality, we take

$$\boldsymbol{U}_{m} = \left[\frac{(1, -1, 0, \cdots, 0)^{\top}}{\sqrt{2}}, \frac{(1, 1, -2, 0, \cdots, 0)^{\top}}{\sqrt{6}}, \cdots, \frac{(1, \cdots, 1, -(m-1))^{\top}}{\sqrt{(m-1)m}} \right].$$

Define the index space $\mathcal{K}_s = \{ \mathbf{k} = \{k_1, \cdots, k_s\} \subset [q] : 1 \leq k_1 < k_2 < \cdots < k_s \leq q \}$ for $s \in [q]$. Define the order of the \mathbf{k} -interaction as $\|\mathbf{k}\|_0 = s$ and its number of parameters as $\|\mathbf{k}\|_J = (J_{k_1} - 1) \cdots (J_{k_s} - 1)$, if $\mathbf{k} \in \mathcal{K}_s$. Thus, the space \mathcal{K}_s defines the set of all possible joint associations of order s. Let $\mathcal{K}_0 = \{\{0\}\}$, and define $\|\mathbf{k}\|_0 = 0$ and $\|\mathbf{k}\|_J = 1$ when $\mathbf{k} = \{0\}$. In this context, $\mathbf{k} = \{0\}$ is treated as an empty set. The number of parameters (per predictor) for all sth order associations is given by s0 by s1 be s2 be s3. Define

$$oldsymbol{H}_0 = rac{\mathbf{1}_{|oldsymbol{J}|}}{\sqrt{|oldsymbol{J}|}} = rac{\mathbf{1}_{J_q}}{\sqrt{J_q}} \otimes rac{\mathbf{1}_{J_{q-1}}}{\sqrt{J_{q-1}}} \otimes \cdots \otimes rac{\mathbf{1}_{J_2}}{\sqrt{J_2}} \otimes rac{\mathbf{1}_{J_1}}{\sqrt{J_1}}.$$

For any $\mathbf{k} = \{k_1, \dots, k_s\} \in \mathcal{K}_s, s \geq 1$, define

$$\boldsymbol{H}_{k} = \boldsymbol{V}_{q} \otimes \boldsymbol{V}_{q-1} \otimes \cdots \otimes \boldsymbol{V}_{2} \otimes \boldsymbol{V}_{1}, \qquad \boldsymbol{V}_{i} = \begin{cases} \boldsymbol{U}_{J_{i}} & i \in \boldsymbol{k} \\ \frac{\mathbf{1}_{J_{i}}}{\sqrt{J_{i}}} & i \in [q] \backslash \boldsymbol{k} \end{cases}$$
 (11)

Following from Example 1, we see that $\operatorname{span}(H_k)$ is the subspace corresponding to the $k = \{k_1, k_2, \dots, k_s\}$ -joint associations. Importantly, it can be verified that the columns of H_k are orthonormal.

Lemma 2 (Subspace decomposition) We can express $\mathbb{R}^{|J|}$ as the orthogonal direct sum of the family $\{\mathcal{R}(H_k)\}_{k\in \cup_{s=0}^q \mathcal{K}_s}$ of subspaces of $\mathbb{R}^{|J|}$, where $\mathcal{R}(H_k)$ denotes the column space of H_k and the orthogonal direct sum is defined in Section S1 of the online supplementary material. Furthermore, for any $k\in \cup_{s=0}^q \mathcal{K}_s$, the orthogonal projection matrix onto $\mathcal{R}(H_k)$ is given by $H_kH_k^{\perp}$.

Lemma 2 has two key implications. First, it ensures invariance of our proposed estimator (Subsection 4.2), and second, it allows us to identify the association encoded by the span of each H_k , as we discuss in the following subsection.

2.3. Reparameterization via subspace decomposition

When fitting (4), it is common to restrict the hypothesis space of models to include joint associations of at most order d for some $1 \le d \le q$. For example, if q = 3, we only want to consider models with all possible main associations and two-way interaction associations. In this case, we would take d = 2. As such, define the space of possible associations as $\mathcal{K} = \bigcup_{s=0}^{d} \mathcal{K}_{s}$. We call \mathcal{K} the association index space. Though \mathcal{K} is a function of d, the maximal order of association considered, we omit notation indicating this dependence for notational simplicity.

The following theorem elucidates how the reparameterization of θ through our subspace decomposition neatly characterizes relationships of mutual, joint, and conditional independence among categorical responses.

Theorem 1 (Sparsity and interpretable models) Let $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_m$ be a partition of [q]. For any $k \in \mathcal{K}_s$, let β_k be a matrix in $\mathbb{R}^{|k|_J \times p}$. For any $\mathcal{I} = \{i_1, \dots, i_s\} \subset [q]$ such that $1 \leq i_1 < \dots < i_s \leq q$, define $J^{\mathcal{I}} = [J_{i_1}] \times \dots \times [J_{i_s}]$.

(Mutual and joint independence) Let S_{joint} = {k ∈ K : ∃ i ∈ [m] such that k ⊂ I_i}. If β_k = 0 for all k ∉ S_{joint}, i.e., the parameter θ under either Poisson (5) categorical response model or multinomial (4) categorical response model is given by θ = ∑_{k∈K} H_kβ_k = ∑_{k∈S_{joint}} H_kβ_k, then

$$\pi_{j}(x) = \prod_{l=1}^{m} \pi_{j_{\mathcal{I}_{l}},+}(x),$$
(12)

where $\mathbf{j}_{\mathcal{I}_l} = \{j_{i_1^l}, \cdots, j_{i_s^l}\} \in \mathbf{J}^{\mathcal{I}^l}$, $\mathcal{I}^l = \{i_1^l, \cdots, i_s^l\}$ for some s (which may depend on l), and $\pi_{\mathbf{j}_{\mathcal{I}_l}, +}$ is the marginal probability mass function of the responses corresponding to \mathcal{I}_l .

2. (Conditional independence) Let $S_{\text{joint}|\mathcal{I}_m} = \{ \mathbf{k} \in \mathcal{K} : \exists \ i \in [m-1] \ \text{such that } \mathbf{k} \subset \mathcal{I}_i \cup \mathcal{I}_m \}$. If $\beta_{\mathbf{k}} = \mathbf{0}$ for all $\mathbf{k} \notin S_{\text{joint}|\mathcal{I}_m}$, i.e., the parameter $\boldsymbol{\theta}$ under either Poisson (5) categorical response model or multinomial (4) categorical response model is given by $\boldsymbol{\theta} = \sum_{\mathbf{k} \in \mathcal{K}} \mathbf{H}_{\mathbf{k}} \beta_{\mathbf{k}} = \sum_{\mathbf{k} \in S_{\text{loint}|\mathcal{I}_m}} \mathbf{H}_{\mathbf{k}} \beta_{\mathbf{k}}$, then

$$\pi_{j_{\mathcal{I}_{1}},j_{\mathcal{I}_{2}},\cdots,j_{\mathcal{I}_{m-1}}|j_{\mathcal{I}_{m}}}(x) = \prod_{l=1}^{m-1} \pi_{j_{\mathcal{I}_{l}},+|j_{\mathcal{I}_{m}}}(x),$$
(13)

where

$$\pi_{oldsymbol{j}_{\mathcal{I}_1},\cdots,oldsymbol{j}_{\mathcal{I}_{m-1}}|oldsymbol{j}_{\mathcal{I}_m}}(oldsymbol{x}) = rac{\pi_{oldsymbol{j}}(oldsymbol{x})}{\pi_{oldsymbol{j}_{\mathcal{I}_m},+}(oldsymbol{x})}, \quad and \quad \pi_{oldsymbol{j}_{\mathcal{I}_l},+|oldsymbol{j}_{\mathcal{I}_m}}(oldsymbol{x}) = rac{\pi_{oldsymbol{j}_{\mathcal{I}_l}\cup\mathcal{I}_m},+(oldsymbol{x})}{\pi_{oldsymbol{j}_{\mathcal{I}_m},+}(oldsymbol{x})}.$$

To illustrate the practical implications and applications of Theorem 1, we present the following example.

Example 2 Suppose q=4. The following types of dependence structures—akin to those in Chapter 6 of McCullagh and Nelder (1989)—are encoded in the sparsity of the β_k . Recall that the random multivariate categorical response is $(Z_1, \dots, Z_q) \in J$.

1. Mutual independence. If $\theta = H_{\{0\}}\beta_{\{0\}} + H_{\{1\}}\beta_{\{1\}} + H_{\{2\}}\beta_{\{2\}} + H_{\{3\}}\beta_{\{3\}} + H_{\{4\}}\beta_{\{4\}}$, then Z_1, Z_2, Z_3 and Z_4 are mutually independent for any given x, i.e., for all $x \in \mathcal{X}$

$$\pi_{j_1,j_2,j_3,j_4}(x) = \pi_{j_1,+,+,+}(x) \cdot \pi_{+,j_2,+,+}(x) \cdot \pi_{+,+,j_3+,}(x) \cdot \pi_{+,+,+,j_4}(x),$$

for all $(j_1, j_2, j_3, j_4) \in J$.

2. Joint independence. If

$$\theta = \sum_{i=0}^{4} \boldsymbol{H}_{\{i\}} \boldsymbol{\beta}_{\{i\}} + \Big(\boldsymbol{H}_{\{2,3\}} \boldsymbol{\beta}_{\{2,3\}} + \boldsymbol{H}_{\{2,4\}} \boldsymbol{\beta}_{\{2,4\}} + \boldsymbol{H}_{\{3,4\}} \boldsymbol{\beta}_{\{3,4\}} + \boldsymbol{H}_{\{2,3,4\}} \boldsymbol{\beta}_{\{2,3,4\}} \Big),$$

then the variable Z_1 is jointly independent of $\{Z_2, Z_3, Z_4\}$ for any given x, i.e., for all $x \in \mathcal{X}$

$$\pi_{j_1,j_2,j_3,j_4}(\mathbf{x}) = \pi_{j_1,+,+,+}(\mathbf{x}) \cdot \pi_{+,j_2,j_3,j_4}(\mathbf{x}), \text{ for all } (j_1,j_2,j_3,j_4) \in \mathbf{J}.$$

3. Conditional independence. If

$$\boldsymbol{\theta} = \sum_{i=0}^4 \boldsymbol{H}_{\{i\}} \boldsymbol{\beta}_{\{i\}} + \Big(\sum_{2 < i < j < 4} \boldsymbol{H}_{\{i,j\}} \boldsymbol{\beta}_{\{i,j\}} + \boldsymbol{H}_{\{2,3,4\}} \boldsymbol{\beta}_{\{2,3,4\}} \Big) + \boldsymbol{H}_{\{1,4\}} \boldsymbol{\beta}_{\{1,4\}},$$

then the variable Z_1 and $\{Z_2, Z_3\}$ are conditionally independent for any given x and Z_4 , i.e., for all $x \in \mathcal{X}$,

$$\pi_{j_1,j_2,j_3|j_4}(\boldsymbol{x}) = \pi_{j_1,+,+|j_4}(\boldsymbol{x}) \cdot \pi_{+,j_2,j_3|j_4}(\boldsymbol{x}), \text{ for all } (j_1,j_2,j_3,j_4) \in \boldsymbol{J}.$$

The neat interpretations in Example 2 rely partly on a hierarchical structure of the associations. That is, highorder associations are included only if all the corresponding low-order associations are included. Formally, if association k is included in the model, then all $k' \in \mathcal{K}$ such that $k' \subset k$ must also be included in the model. For example, with q = 3, if the joint association $\{1, 2, 3\}$ is included in the model, then for the hierarchy to be enforced, the associations $\{0\}, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \text{ and } \{2, 3\}$ must all be included in the model.

Formally, given an association index space \mathcal{K} , the corresponding class of hierarchical association index space is the collection of all sets $\mathcal{N} \subset \mathcal{K}$ such that if $\mathbf{k} \in \mathcal{N}$, then $\mathcal{P}(\mathbf{k}) \subset \mathcal{N}$, where $\mathcal{P}(\mathbf{k})$ denotes the powerset of \mathbf{k} (with the null set replaced with $\{0\}$).

To restrict attention only to models that respect such a hierarchy, it is natural to consider a class of hierarchical hypotheses spaces

$$\left\{ \boldsymbol{\theta} \in \mathbb{R}^{|\boldsymbol{J}| \times p} : \mathcal{R}(\boldsymbol{\theta}) = \sum_{\boldsymbol{k} \in \mathcal{N}} \mathcal{R}(\boldsymbol{H}_{\boldsymbol{k}}) \subset \mathbb{R}^{|\boldsymbol{J}|}, \quad \mathcal{N} \subset \mathcal{K} \quad \text{s.t.} \quad \boldsymbol{k} \in \mathcal{N} \implies \mathcal{P}(\boldsymbol{k}) \subset \mathcal{N} \right\}.$$
(14)

In the next subsection, we will propose a penalized maximum likelihood estimator that allows to explore models in \mathcal{K} or its corresponding hierarchical association index space.

3. Penalized likelihood-based association learning

3.1. Penalized maximum likelihood estimation

Define the negative log-likelihoods as L_n^{Mult} , and its reparametarized versions as $\mathcal{L}_n^{\text{Mult}}$ where $L_n^{\text{Mult}}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \ell_{\text{Mult}}(\boldsymbol{\theta} \boldsymbol{x}_i, \boldsymbol{y}_i)$ and $\mathcal{L}_n^{\text{Mult}}(\boldsymbol{\beta}) = L_n^{\text{Mult}}(\boldsymbol{H}\boldsymbol{\beta})$. Similarly define $\mathcal{L}_n^{\text{Pois}}(\boldsymbol{\beta}) = L_n^{\text{Pois}}(\boldsymbol{H}\boldsymbol{\beta})$. To simplify the notation and unify the statements and analysis, set

$$\mathcal{L}_n(oldsymbol{eta}) = \left\{ egin{array}{ll} \mathcal{L}_n^{ ext{Mult}}(oldsymbol{eta}) &: \ell = \ell_{ ext{Mult}} \ \mathcal{L}_n^{ ext{Pois}}(oldsymbol{eta}) &: \ell = \ell_{ ext{Pois}} \end{array}
ight.$$

As described in the previous section, due to our subspace decomposition, association structure learning is achieved by learning the sparsity pattern of $\beta \in \mathbb{R}^{\sum_{s=0}^{d} L_s \times p}$. For this, we will use penalized maximum likelihood estimators of the form

$$\widehat{\boldsymbol{\beta}} \in \arg\min_{\boldsymbol{\beta}} \left\{ \mathcal{L}_n(\boldsymbol{\beta}) + \lambda \Omega(\boldsymbol{\beta}) \right\},$$
 (15)

for convex penalties $\Omega: \mathbb{R}^{\sum_{s=0}^d L_s \times p} \to [0, \infty)$ to be discussed in the next subsection.

3.2. Global versus local association learning

Given $d \in [q]$ and association index space \mathcal{K} (determined by d), to take the advantages of the subspace decomposition in Section 2, we parameterize θ as

$$\theta = \sum_{\mathbf{k} \in \mathcal{K}} H_{\mathbf{k}} \beta_{\mathbf{k}} =: H \beta, \tag{16}$$

with $\boldsymbol{H} = \{\boldsymbol{H_k}\}_{k \in \mathcal{K}} \in \mathbb{R}^{|J| \times \sum_{s=0}^d L_s}$ and $\boldsymbol{\beta} = \{\boldsymbol{\beta_k}\}_{k \in \mathcal{K}} \in \mathbb{R}^{\sum_{s=0}^d L_s \times p}$. Let $\boldsymbol{x}_{i(j)} \in \mathbb{R}^{p_j}$ be the jth subvector of \boldsymbol{x}_i , $j \in [t]$, where $\sum_{j=1}^t p_j = p$. Without loss of generality, we partition the matrix $\boldsymbol{\beta_k} = [\boldsymbol{\beta_{k,1}}, \boldsymbol{\beta_{k,2}}, \dots, \boldsymbol{\beta_{k,t}}]$ and vector $\boldsymbol{x}_i = (\boldsymbol{x}_{i(1)}^\top, \dots, \boldsymbol{x}_{i(t)}^\top)^\top \in \mathbb{R}^p$ so that

$$oldsymbol{ heta} oldsymbol{x}_i = \sum_{oldsymbol{k} \in \mathcal{K}} \sum_{i=1}^t oldsymbol{H_k} oldsymbol{eta_{k,j}} oldsymbol{x}_{i(j)}, \qquad oldsymbol{eta_{k,j}} \in \mathbb{R}^{|oldsymbol{k}|_J imes p_j}.$$

As discussed in the previous section, if $\beta_{\mathbf{k}} = 0$, then the corresponding association defined by \mathbf{k} is not included in our model. Our predictor grouping structure allows us to perform association learning at distinct resolutions: global association learning or local association learning (i.e., predictor-wise association learning).

The goal of global association learning is to discover associations such that all predictors contribute to the association, or none contribute to the association. For global association learning, we take t = 1. To encourage sparsity in our fitted model so as to discover a small number of global associations, we use a group

lasso-type penalty (Yuan and Lin, 2006) with a positive set of (user-specified) weights $\{w_k\}_{k\in\mathcal{K}}$ for β and θ , respectively, as

$$\Omega_{\text{global}}(\boldsymbol{\beta}) = \sum_{\boldsymbol{k} \in \mathcal{K}} w_{\boldsymbol{k}} \|\boldsymbol{\beta}_{\boldsymbol{k}}\|, \tag{17}$$

$$\Phi_{\text{global}}(\boldsymbol{\theta}) = \inf_{\boldsymbol{\theta} = \boldsymbol{H}\boldsymbol{\beta}} \Omega_{\text{global}}(\boldsymbol{\beta}) = \Omega_{\text{global}}(\boldsymbol{H}^{\top}\boldsymbol{\theta}). \tag{18}$$

Given that $\boldsymbol{\theta} = \boldsymbol{H}\boldsymbol{\beta}$ uniquely determines a $\boldsymbol{\beta} \in \mathbb{R}^{\sum_{s=0}^{d} L_s \times p}$, the infimum in (18) can be omitted. Because the Frobenius norm is nondifferentiable at the matrix of zeros, using Ω_{global} as a penalty can encourage estimates of the $\boldsymbol{\beta}$, $\hat{\boldsymbol{\beta}}$ such that $\hat{\boldsymbol{\beta}}_{\boldsymbol{k}} = 0$ for many $\boldsymbol{k} \in \mathcal{K}$.

In local association learning, we relax the assumption that all predictors either contribute to an association, or no predictors contribute to an association. For example, when q=2, it is possible that for the majority of predictors (but not all), a change in the predictor's value does not lead to a change in any of the local odd-ratios between response variables (i.e., these predictors only affect the marginal distributions of the response). This was exactly the type of association learning performed by Molstad and Rothman (2023). Our local association learning is much more general: we can discover which predictors modify certain high-order associations, and which predictors (or groups of predictors) only affect lower-order associations.

To achieve this type of learning, define the set $\mathcal{G}_{local} = \{(k, j) : k \in \mathcal{K}, j \in [p]\}$, let $\{w_{k,j}\}_{(k,j) \in \mathcal{G}_{local}}$ be a positive sequence, and define the penalty function

$$\Omega_{\text{local}}(\boldsymbol{\beta}) = \sum_{(\boldsymbol{k}, j) \in G_{\text{local}}} w_{\boldsymbol{k}, j} \| \boldsymbol{\beta}_{\boldsymbol{k}, j} \|, \qquad (19)$$

and similarly for Φ_{local} . In contrast to Ω_{global} , Ω_{local} has nondifferentiabilities when $\beta_{k,j} = 0$ for any $(k,j) \in \mathcal{G}_{\text{local}}$. As such, this penalty can encourage estimates such that $\widehat{\beta}_{k,j} = 0$ for many $j \in [p]$, but if $\widehat{\beta}_{k,j'} \neq 0$ for any j', then the k-joint association is included in the model.

Defining the set $\mathcal{G} = \{(\mathbf{k}, j) : \mathbf{k} \in \mathcal{K}, j \in [t]\}$, and defining $\Omega_{\mathcal{G}}(\beta) = \sum_{(\mathbf{k}, j) \in \mathcal{G}} w_{\mathbf{k}, j} \|\beta_{\mathbf{k}, j}\|$, we generalize both global association learning (t = 1) and predictor-wide local association learning (t = p). More generally, we can perform a version of local association learning with predictors partitioned into t sets. This may be useful, for example, if predictors are categorical and encoded via multiple dummy variables.

3.3. Association learning with hierarchical constraints

As mentioned in Section 2.3, it is often desirable to enforce a hierarchical structure for the associations. To this end, we can modify both our global and local association structure learning penalties to enforce the hierarchy. Recall that for the hierarchy to be enforced, we must have that for every association k included in the model, all elements of $\mathcal{P}(k)$ must also be included in the model.

To achieve model fits of this type, we utilize the overlapping group lasso penalty. This penalty is defined by

$$\Omega_{\mathcal{G}}^{H}(\boldsymbol{\beta}) = \sum_{(\boldsymbol{k},j)\in\mathcal{G}} w_{\boldsymbol{k},j} \sqrt{\sum_{\boldsymbol{k}':\boldsymbol{k}\subset\boldsymbol{k}'} \|\boldsymbol{\beta}_{\boldsymbol{k}',j}\|^{2}}$$
(20)

and

$$\Phi_{\mathcal{G}}^{H}(\boldsymbol{\theta}) = \Omega_{\mathcal{G}}^{H}(\boldsymbol{H}^{\top}\boldsymbol{\theta}). \tag{21}$$

The term $\sqrt{\sum_{k':k\subset k'}\|\beta_{k',j}\|^2}$ is a group lasso penalty on the entire set of coefficients corresponding to associations that include k in their powerset. For example, if q=3 and $k=\{1\}$, then $\sqrt{\sum_{k':k\subset k'}\|\beta_{k',j}\|^2}=\sqrt{\|\beta_{\{1\},j}\|^2+\|\beta_{\{1,2\},j}\|^2+\|\beta_{\{1,3\},j}\|^2+\|\beta_{\{1,2,3\},j}\|^2}$. Consequently, this penalty essentially precludes the possibility that $\hat{\beta}_{\{1,2\},j}\neq 0$ but $\hat{\beta}_{\{1\},j}=0$, for example, because the penalty enforces $\hat{\beta}_{\{1\},j}=0$ (via nondifferentiability at the origin) only when all higher order associations $\hat{\beta}_{\{1,2\},j}=\hat{\beta}_{\{1,3\},j}=\hat{\beta}_{\{1,2,3\},j}=0$ as well. See Yan and Bien (2017) for a comprehensive review of how hierarchical structures can be enforced with the overlapping group lasso and related penalties.

4. Relation to existing work

4.1. Alternative parametric links

Multivariate categorical response regression is a classical problem in categorical data analysis (e.g., see Chapter 6 of McCullagh and Nelder, 1989). The majority of existing methods designed specifically for this task utilize parametric links between predictors and responses that can yield interpretable fitted models. To best describe these methods, we will first consider the case that p = 1 and $x_i = 1$ for all $i \in [n]$ (i.e., the analysis of a q-way contingency table).

One particularly neat parametric link is the multivariate logistic transform. This transform maps probabilities $\pi \in \mathbb{R}^{|J|}$ to a set of parameters η . These parameters represent the logarithms of the marginal odds, pairwise odds ratios, and higher-order odds ratios, which are derived from all possible joint marginals of subsets Z_1, \ldots, Z_q (McCullagh and Nelder, 1989; Glonek, 1996; Molenberghs and Lesaffre, 1999). For a given π , the transformation $\pi \to \eta$ can be expressed as a matrix equation:

$$\eta = C \log(M\pi),\tag{22}$$

where C is a contrast matrix, and M is a marginalizing matrix that computes the joint marginals from the cell probabilities. A more general class of log-linear models (where C and M are more general, and $\eta = Z\eta_z$ for design matrix Z), was proposed by Lang (1996). According to the definition of Bergsma and Rudas (2002), a numerical value assigned to η is considered strongly compatible if there exists a valid probability distribution π that corresponds to it. Palmgren (1989) showed that excluding the cases when q=2 with $J_1=J_2$, no explicit solution is available. Glonek and McCullagh (1995) pointed out the difficulty in solving (22) for the analysis of contingency tables, stating that "no readily computable criterion, for determining whether a particular η is valid, is available". If there are more than two categorical variables, it can happen that no solution exists because of incompatibility of the lower dimensional marginals. Evidently, it remains unclear how to determine whether a specific η is strongly compatible. For Bernoulli response Z_1, \ldots, Z_q , Qaqish and Ivanova (2006) can determine the strong compatibility of η , and compute π from a strongly compatible η using a noniterative algorithm. When any $J_l \geq 3$, however, their results cannot be applied.

Matters become even more challenging when we consider the more general log-linear regression model $f(x_i) = C \log\{M\pi(x_i)\}$ where $x_i \in \mathbb{R}^p$ for linear function f. The goal of our work is to provide an alternative to log-linear models that (i) has parameters that can be interpreted in the same way as log-linear models and (ii) can be easily computed. Desideratum (i) is addressed by Theorem 1, and as we will show in a later section, because our estimator is the solution to a convex optimization problem, we can readily employ modern first order methods for (ii).

4.2. Generalizing log-linear models for contingency tables

In this section, we will explain how our method generalizes log-linear models used for the analysis of contingency tables. The key is that our method has the interpretability of "standard" log-linear models, but our specific subspace decomposition leads to an invariance property that is essential for penalized maximum likelihood-based association learning.

Log-linear models are a class of statistical models used to describe the relationship between categorical variables by modeling the expected cell counts in a contingency table. These models express the logarithm of expected frequencies as a linear combination of parameters corresponding to main associations and interactions of the variables. Specifically, for a contingency table (i.e., the intercept only model with p = 1) with variables Z_1 and Z_2 , the model can be written as

$$\log(\mu_{j_1j_2}) = \mathbf{\Lambda}_{j_1,j_2} := \mu + \mu_{j_1}^{Z_1} + \mu_{j_2}^{Z_2} + \mu_{j_1j_2}^{Z_1Z_2}$$
(23)

where $\mu_{j_1j_2}$ denotes the expected count in cell (j_1,j_2) , μ is the overall mean, $\mu_{j_1}^{Z_1}$ and $\mu_{j_2}^{Z_2}$ represent the main associations of variables Z_1 and Z_2 , respectively, and $\mu_{j_1j_2}^{Z_1Z_2}$ denotes the interaction association between Z_1

and Z_2 . Under a multinomial sampling scheme, the model can be written as

$$\pi_{j_1 j_2} = \frac{\exp(\mathbf{\Lambda}_{j_1, j_2})}{\sum_{j_1, j_2} \exp(\mathbf{\Lambda}_{j_1, j_2})}.$$
 (24)

The log-linear model and the multinomial model share the same linear structure of Λ_{j_1,j_2} .

To ensure the parameters in a log-linear model are uniquely estimable, certain constraints must be imposed. Commonly, sum-to-zero constraints are used, where the sum of the main associations and interaction associations for each variable is set to zero. For example, for the main associations, the constraints are: $\sum_{j_1=1}^{J_1} \mu_{j_1}^{Z_1} = 0$ and $\sum_{j_2=1}^{J_2} \mu_{j_2}^{Z_2} = 0$. Similarly, for the interaction associations:

$$\sum_{j_1} \mu_{j_1 j_2}^{Z_1 Z_2} = 0 \quad \text{for each } j_2 \quad \text{and} \quad \sum_{j_2} \mu_{j_1 j_2}^{Z_1 Z_2} = 0 \quad \text{for each } j_1.$$

Alternatively, one could define $\mu_1^{Z_1} = 0$, $\mu_1^{Z_2} = 0$ and $\mu_{j_1j_2}^{Z_1Z_2} = 0$ if $j_1 = 1$ or $j_2 = 1$. For maximum likelihood estimation (without penalization), the choice of constraint does not matter due to the invariance property of the maximum likelihood estimator. If, on the other hand, one wanted to impose sparsity inducing penalties on the μ , the choice of constraint may affect the solution.

To see this, recall that $\boldsymbol{\mu}^{\boldsymbol{J}} = \{\mu_{j_1,j_2}\}_{(j_1,j_2)\in \boldsymbol{J}}$ for log-linear model. Let $\boldsymbol{U}_m' = [e_2^m,\ldots,e_m^m]$. Similar to \boldsymbol{H}_k defined in (11), for any $\boldsymbol{k} = \{k_1,\ldots,k_s\} \in \mathcal{K}_s, s \geq 1$, define

$$\boldsymbol{H}_{k}' = \boldsymbol{V}_{q} \otimes \boldsymbol{V}_{q-1} \otimes \cdots \otimes \boldsymbol{V}_{2} \otimes \boldsymbol{V}_{1}, \qquad \boldsymbol{V}_{i} = \begin{cases} \boldsymbol{U}_{J_{i}}' & i \in \boldsymbol{k} \\ \boldsymbol{1}_{J_{i}} & i \in [q] \backslash \boldsymbol{k} \end{cases}$$
 (25)

We can thus rewrite (23) in matrix form as

$$\text{vec}_{\boldsymbol{J}}\{\log(\boldsymbol{\mu}^{\boldsymbol{J}})\} = \boldsymbol{H}'_{\{0\}}\boldsymbol{\beta}'_{\{0\}} + \boldsymbol{H}'_{\{1\}}\boldsymbol{\beta}'_{\{1\}} + \boldsymbol{H}'_{\{2\}}\boldsymbol{\beta}'_{\{2\}} + \boldsymbol{H}'_{\{1,2\}}\boldsymbol{\beta}'_{\{1,2\}},$$

where $\{H'_{k}\}_{k}$ are defined in (25) with q=2. Similarly, recall that $\pi^{J}=\{\pi_{j_1,j_2}\}_{(j_1,j_2)\in J}$ for the multinomial log-linear model so that

$$\operatorname{vec}_{\boldsymbol{J}}\{\boldsymbol{\pi}^{\boldsymbol{J}}\} = \frac{\exp(\boldsymbol{\theta})}{\langle \exp(\boldsymbol{\theta}), 1_{J_1, J_2} \rangle}, \quad \boldsymbol{\theta} = \boldsymbol{H}'_{\{0\}}\boldsymbol{\beta}'_{\{0\}} + \boldsymbol{H}'_{\{1\}}\boldsymbol{\beta}'_{\{1\}} + \boldsymbol{H}'_{\{2\}}\boldsymbol{\beta}'_{\{2\}} + \boldsymbol{H}'_{\{1,2\}}\boldsymbol{\beta}'_{\{1,2\}}.$$

Here, $H'_{\{0\}}\beta'_{\{0\}}, H'_{\{1\}}\beta'_{\{1\}}, H'_{\{2\}}\beta'_{\{2\}}, H'_{\{1,2\}}\beta'_{\{1,2\}}$ are the matrix forms of μ , $\mu^{Z_1}_{j_1}$, $\mu^{Z_2}_{j_2}$, and $\mu^{Z_1Z_2}_{j_1j_2}$, respectively for both log-linear model and multinomial model. Evidently the log-linear model can be parameterized as $\theta = \sum_{k \in \mathcal{K}} H'_k \beta'_k$. If we wanted to impose sparsity on the β' , it would be tempting to use the same group lasso penalty as defined before,

However, when considering $\Phi'(\theta) = \sum_{k} \|\beta'_{k}\|$, we see that $\Phi'(\cdot)$ is not invariant to the choice of identifability constraints. To be more specific, if $\max_{i \in [q]} J_i > 2$, $U''_m = [e_1^m, \dots, e_{m-1}^m]$, and we define H'' accordingly, then

$$\boldsymbol{H_k'\beta_k'} \not\equiv \boldsymbol{H_k''\beta_k''}, \sum_{\boldsymbol{k}} \|\boldsymbol{\beta_k''}\| \not\equiv \sum_{\boldsymbol{k}} \|\boldsymbol{\beta_k''}\| \text{ , where } \boldsymbol{\theta} = \sum_{\boldsymbol{k} \in \mathcal{K}} \boldsymbol{H_k'\beta_k'} = \sum_{\boldsymbol{k} \in \mathcal{K}} \boldsymbol{H_k''\beta_k''}.$$

Choosing H'' instead of H' changes how the k-joint association influences the categorical response, leading to results that may depend on this arbitrary selection rather than reflecting an inherent property.

To address the invariance issue, one might consider using an overparameterized version of the log-linear model with penalization of the parameters. However, this leads to an explosion in the number of parameters, and the parameters are more difficult to interpret. Moreover, statistical analysis of such an estimator is fundamentally more difficult than the analysis of our estimator.

In our reparameterization $\theta = \sum_{k \in \mathcal{K}} H_k \beta_k$, the corresponding group lasso penalty $\Phi(\theta) = \sum_k \|\beta_k\|$ is invariant under different choices of U_m such that $[\frac{1}{\sqrt{m}} \mathbf{1}_m, U_m]$ is a real orthogonal matrix. To be more specific, if we let \mathcal{U}_m be another real matrix such that $[\frac{1}{\sqrt{m}} \mathbf{1}_m, \mathcal{U}_m]$ is a real orthogonal matrix, and define $H_k^{\mathcal{U}}$ by replacing U_{J_i} with \mathcal{U}_{J_i} in (11), then

$$\boldsymbol{H}_{\boldsymbol{k}}\boldsymbol{\beta}_{\boldsymbol{k}} \equiv \boldsymbol{H}_{\boldsymbol{k}}^{\mathcal{U}}\boldsymbol{\beta}_{\boldsymbol{k}}^{\mathcal{U}}, \sum_{\boldsymbol{k}} \|\boldsymbol{\beta}_{\boldsymbol{k}}\| \equiv \sum_{\boldsymbol{k}} \left\|\boldsymbol{\beta}_{\boldsymbol{k}}^{\mathcal{U}}\right\|, \ \text{ where } \boldsymbol{\theta} = \sum_{\boldsymbol{k} \in \mathcal{K}} \boldsymbol{H}_{\boldsymbol{k}}\boldsymbol{\beta}_{\boldsymbol{k}} = \sum_{\boldsymbol{k} \in \mathcal{K}} \boldsymbol{H}_{\boldsymbol{k}}^{\mathcal{U}}\boldsymbol{\beta}_{\boldsymbol{k}}^{\mathcal{U}}.$$

4.3. Modern approaches to multivariate categorical response regression in high dimensions

Existing methods for multivariate categorical response regression with a large number of predictors, responses, and/or a large number categories per response typically rely on latent variable models (e.g., the regularized latent class model of Molstad and Zhang, 2022), or classifier chains (Read et al., 2021).

The latent class model is able to capture complex relationships between responses by assuming that given a latent variable W, Z_m and $Z_{m'}$ are independent given X, i.e., $Z_m \perp Z_{m'} \mid X, W$. Thus, fitted model coefficients cannot be straightforwardly interpreted in terms of the distribution of interest $Z_1, \ldots, Z_q \mid X$, as can the coefficients from our fitted model. Moreover, the order of associations in the latent class method cannot, generally speaking, be easily identified unless the association is null.

Along similar lines, it is common to decompose the joint mass function of interest into simpler, estimable parts. Methods utilizing to this approach include those most popular in the machine learning literature on "multilabel classification" (Herrera et al., 2016), namely, classifier chains (Read et al., 2021). A classifier chain estimates $Z_1, \ldots, Z_q \mid X$ by fitting a model for $Z_1 \mid X$, then $Z_2 \mid X, Z_1$, then $Z_3 \mid X, Z_1, Z_2$, and so on, and using their product as an estimate of the mass function of interest. This approach requires many ad-hoc decisions that can have a significant impact on how the model performs (e.g., in what order to fit the chain and how to model each specific conditional distribution). Like the latent class model approach, classifier chains cannot be used to identify the order of associations in a straightforward way, which is the primary motivation for our work.

Recently, to handle large p, Deng et al. (2024) proposed to model multivariate categorical responses under a multivariate generalization of the normal linear discriminant analysis model. Though their method can handle large p, it is specifically designed for a bivariate response, and does not provide the interpretability of our method.

4.4. Connection with Ising, Potts, and Hypergraph Models

In the classical Ising model, each vertex of a graph is assigned a binary state $j_i \in \{0,1\}$, with interactions limited to pairs of vertices connected by an edge. Cheng et al. (2014) extend the Ising model allow interactions to depend on covariates. They propose a sparse pseudo-likelihood estimation framework tailored to the binary Ising model, which does not extend to multi-level categorical variables. In contrast, in the Potts model, each vertex is assigned a variable $j_i \in \{1, 2, ..., J_i\}$, with pairwise interactions specified by edges (Shimagaki and Weigt, 2019; Wipf et al., 2007; Tanaka et al., 2011; Hirakida et al., 2016). Such models capture the essential physics of systems whose energy depends exclusively on pairwise interactions. Broadly speaking, maximum pseudo-likelihood estimation is widely used under both Ising and generalized Potts models without predictors (Razaee and Amini, 2020; Ekeberg et al., 2013; Okabayashi et al., 2011; Cheng et al., 2014; Keetelaar et al., 2024; Levada et al., 2009). However, as illustrated in Keetelaar et al. (2024), there is a key trade-off to be considered: while pseudo-likelihood is computationally tractable and approximates MLE well with sufficiently large sample sizes, it becomes less reliable under with small or moderate n. In contrast, our approach relies on the complete likelihood.

To the best of our knowledge, no existing study has introduced covariates into the Potts model. As we will show below, the generalized Potts model—already a special case of a hypergraph model—also emerges from our approach when the maximum interaction order is set to 2. Our method thus advances Potts modeling by incorporating covariates, extending them beyond the purely pairwise framework.

Many real-world systems exhibit interactions among more than two components simultaneously. For example, in certain magnetic materials or biological networks, the collective behavior depends on multi-body interactions. To model such scenarios, it is natural to generalize the notion of a graph to that of a (undirected) hypergraph (Bretto, 2013; Klamt et al., 2009; Yu et al., 2024; Robiglio et al., 2024; Mukherjee et al., 2022; Bhowal and Mukherjee, 2024). In Section S3 of the online supplementary material, we demonstrate that both the generalized Potts model and the hypergraph model coincide with our proposed multivariate categorical response regression model—under the reparameterization $\theta = \sum_k H_k \beta_k$ —when there are no predictors.

5. Computation

In this section, we outline a proximal gradient descent algorithm—described in Chapter 4 of Parikh et al. (2014)—for computing the group lasso and the overlapping group lasso-penalized estimators.

The proximal gradient descent algorithm can be understood from the perspective of the majorize-minimize principle (Lange, 2016). If there exists some L > 0 such that for arbitrary kth iterate β^k ,

$$\mathcal{L}_n(\beta) + \lambda \Omega(\beta) \le \mathcal{L}_n(\beta^k) + \left\langle \nabla \mathcal{L}_n(\beta^k), \beta - \beta^k \right\rangle + \frac{L}{2} \left\| \beta - \beta^k \right\|^2 + \lambda \Omega(\beta)$$
 (26)

for all β , then, if we define the (k+1)th iterate as

$$\beta^{k+1} = \arg\min_{\beta} \left[\frac{1}{2} \left\| \beta - \left\{ \beta^k - \frac{1}{L} \nabla \mathcal{L}_n(\beta^k) \right\} \right\|^2 + \frac{\lambda}{L} \Omega(\beta) \right], \tag{27}$$

we are ensured that the objective function at β^{k+1} is no greater than the objective function at β^k (i.e., the sequence of iterates $\{\beta^k\}_{k=1}^{\infty}$ have the descent property). When Ω is the group lasso penalty (i.e., $\Omega_{\mathcal{G}}$), then the proximal operator (27) has closed form solution

$$\boldsymbol{\beta}_g^{k+1} = \max\left(1 - \frac{\lambda w_g}{\left\|L \cdot \boldsymbol{\beta}_g^k - \frac{\partial}{\partial \boldsymbol{\beta}_g} \mathcal{L}_n(\boldsymbol{\beta}^k)\right\|}, 0\right) \left(\boldsymbol{\beta}_g^k - \frac{1}{L} \frac{\partial}{\partial \boldsymbol{\beta}_g} \mathcal{L}_n(\boldsymbol{\beta}^k)\right), \quad g \in \mathcal{G}.$$

When Ω is the overlapping group lasso penalty (i.e., $\Omega_{\mathcal{G}}^{H}$), we solve the proximal operator (27) using a blockwise coordinate algorithm (Jenatton et al., 2011; Yan and Bien, 2017).

In Lemma S8 of the online supplementary material, we show that for all β and β' , $\|\nabla \mathcal{L}_n^{\text{Mult}}(\beta) - \nabla \mathcal{L}_n^{\text{Mult}}(\beta')\| \leq (2n)^{-1} \lambda_{\max}(X^\top X) \|\beta - \beta'\|$ with $X = [x_1, \dots, x_n]$, which implies that with $L \geq (2n)^{-1} \lambda_{\max}(X^\top X)$, (26) will hold. However, in the case of a Poisson categorical response model, the inequality (26) cannot hold globally for any L. Therefore, we use a proximal gradient descent algorithm with the step size determined adaptively at each iteration by a backtracking line search.

For the sake of space, more details about tuning parameter selection, as well as the formulation of an accelerated variation of the proximal gradient descent algorithm, can be found in Section S4 of the online supplementary material. More details about the accelerated proximal gradient descent algorithm can be found in Section 4.3 of Parikh et al. (2014) and Algorithm 2 of Tseng (2008). We present our algorithm and all needed sub-algorithms in Section S4 of the online supplementary material.

6. Statistical properties

In this section, we examine the statistical properties of the group lasso estimator, as defined in (15), considering variations in n, p, and J. Let $\theta^* = H_{\text{full}}\beta^*$ represent the data generation parameter, where $H_{\text{full}} = \{H_k\}_{k \in \cup_{s=0}^q \mathcal{K}_s} \in \mathbb{R}^{|J| \times |J|}$. To establish an error bound, it is necessary to define an identifiable estimand: the parameter β^{\dagger} . Let the set \mathcal{F}_{θ^*} denote the set of all β , which leads to the same probability distribution, that is for multinomial and Poisson categorical response models,

$$\mathcal{F}_{\boldsymbol{\theta}^*} = \left\{ \widetilde{\boldsymbol{\beta}} \in \mathbb{R}^{|\boldsymbol{J}| \times p}; \ell(\boldsymbol{\theta}^* \boldsymbol{x}, \boldsymbol{y}) = \ell(\boldsymbol{H}_{\text{full}} \widetilde{\boldsymbol{\beta}} \boldsymbol{x}, \boldsymbol{y}), \forall (\boldsymbol{x}, \boldsymbol{y}) \right\}$$

$$= \left\{ \widetilde{\boldsymbol{\beta}} \in \mathbb{R}^{|\boldsymbol{J}| \times p}; \boldsymbol{P}_{\mathcal{V}} \boldsymbol{\theta}^* = \boldsymbol{P}_{\mathcal{V}} \boldsymbol{H}_{\text{full}} \widetilde{\boldsymbol{\beta}} \right\},$$
(28)

where $P_{\mathcal{V}} = I - |J|^{-1} \mathbf{1}_{|J|} \mathbf{1}_{|J|}^{\top}$ for multinomial categorical response model, and $P_{\mathcal{V}} = I$ for Poisson categorical response model.

Define $\theta^{\dagger} = P_{\mathcal{V}}\theta^*$ and $\beta^{\dagger} = H_{\text{full}}^{\top}\theta^{\dagger}$. By Lemma S10 in the online supplementary material, we know that

$$\beta^{\dagger} \in \underset{\beta \in \mathcal{F}_{\theta^*}}{\operatorname{arg min}} \{ \mathcal{L}_n(\beta) + \lambda \Omega_{\mathcal{G}}(\beta) \} = \underset{\beta \in \mathcal{F}_{\theta^*}}{\operatorname{arg min}} \Omega_{\mathcal{G}}(\beta). \tag{29}$$

Now, we introduce our assumptions. The first is a standard scaling assumption on the predictors.

Assumption 1 (Predictor scaling) The predictors are scaled so that for any $i \in [n], j \in [t]$, and $\|\mathbf{k}\|_0 \leq d$, $\|\mathbf{x}_{i(j)}\| \leq w_{\mathbf{k},j}C$ for finite constant C.

The following assumption regards the data generating process.

Assumption 2 The responses $\mathbf{y}_i^J = \{y_j^i\}_{j \in J}, 1 \leq i \leq n \text{ are independent given } \{x_i\}_{i=1}^n \text{ and generated under } (i) \text{ the Poisson categorical response model or } (ii) \text{ the multinomial categorical response model with } (n_i = 1 \text{ for } i \in [n], \text{ without loss of generality}).$

Assumption 3 (Poisson categorical response model) Under (i), the Poisson categorical response model with $\mathcal{V} = \mathbb{R}^{|J|}$, there exists a finite constant C_1 such that $\Lambda := \max_{i \in [n]} \|e^{\theta^{\dagger} \mathbf{x}_i}\|_{\infty} \leq C_1$.

Note that under (ii), the multinomial categorical response model, $\mathcal{V} = \{\theta' \in \mathbb{R}^{|J|} : \mathbf{1}_{|J|}^{\top}\theta' = \mathbf{0}\}$. This is not an assumption, but rather a definition.

Next, we make an assumption on the curvature of the negative log-likelihood in certain directions: this is commonly known as restricted strong convexity (Wainwright, 2019, Definition 9.15 and Theorem 9.36). Let $\mathcal{E}_n(\Delta\theta) := L_n(\theta^{\dagger} + \Delta\theta) - L_n(\theta^{\dagger}) - \langle \nabla L_n(\theta^{\dagger}), \Delta\theta \rangle$.

Assumption 4 (Restricted strong convexity) Let $\Phi_{\mathcal{G}}(\theta) = \Omega_{\mathcal{G}}(H^{\top}\theta)$ be the reparameterized group lasso penalty for the association learning. The quantity $\mathcal{E}_n(\Delta\theta)$ satisfies restricted strong convexity (RSC) condition with radius R > 0, constants A and C, and curvature $\kappa > 0$, i.e., $\Delta\theta \in \mathcal{M}_H := \{\theta : \theta = \sum_{g \in \mathcal{G}} H_k \beta_g \}$,

$$\mathcal{E}_{n}(\Delta \boldsymbol{\theta}) \geq \frac{\kappa}{2} \| \boldsymbol{P}_{\mathcal{V}}(\Delta \boldsymbol{\theta}) \|^{2} - A \cdot C^{2} \left(\frac{\log |\mathcal{G}|}{n} + \frac{m}{n} \right) \cdot \inf_{\boldsymbol{P}_{\mathcal{V}} \Delta \boldsymbol{\theta} = \boldsymbol{P}_{\mathcal{V}} \Delta \boldsymbol{\theta}'} \Phi_{\mathcal{G}}^{2}(\Delta \boldsymbol{\theta}'), \quad \|\Delta \boldsymbol{\theta}\| \leq R, \tag{30}$$

where $|\mathcal{G}|$ is the cardinality of \mathcal{G} , and $m = \max_{(\mathbf{k},j) \in \mathcal{G}} |\mathbf{k}|_{\mathbf{J}} \cdot p_j$. Under (i), the Poisson categorical response model, $P_{\mathcal{V}} = I$, and denote $\kappa = \kappa_{\mathbf{J}}^{Pois}$, whereas under (ii), the multinomial categorical response model, $P_{\mathcal{V}} = I - |\mathbf{J}|^{-1} \mathbf{1}_{|\mathbf{J}|} \mathbf{1}_{|\mathbf{J}|}^{\top}$, and denote $\kappa = \kappa_{\mathbf{J}}^{Mult}$.

The restricted strong convexity condition is a well-understood condition in penalized regression (Negahban et al., 2012, Section 2.4). Effectively, this condition requires that in a neighborhood of the true parameter, the negative log-likelihood has sufficient curvature.

Remark 1 In Lemma S6 of the online supplementary material, we verify that under mild assumptions on the distribution of predictors, restricted strong convexity holds with high probability for (i) Poisson and (ii) multinomial categorical response models.

Define the support of β^{\dagger} as $\mathcal{S} = \{g \in \mathcal{G}; \beta_g^{\dagger} \neq 0\}$ and define $\Psi(\mathcal{S})^2 = \sum_{(k,j) \in \mathcal{S}} w_{k,j}^2$. Clearly, if $w_g = 1$ for all $g \in \mathcal{G}$, then $\Psi(\mathcal{S})^2 = |\mathcal{S}|$. Note that $\Psi(\mathcal{S})$ is essentially the subspace compatibility constant (Wainwright, 2019, Definition 9.18): a quantity that often appears in error bounds for regularized M-estimators.

Note that the dimensionality of $\widehat{\boldsymbol{\beta}}$ depends on the user-specified d, whereas $\boldsymbol{\beta}^{\dagger} \in \mathbb{R}^{|J| \times p}$. Thus, to simplify notation, let $\widehat{\boldsymbol{\beta}}_0$ denote the version of $\widehat{\boldsymbol{\beta}}$ where all associations of order higher than d have been set to zero (i.e., $\widehat{\boldsymbol{\beta}}_0^{\top} = [\widehat{\boldsymbol{\beta}}^{\top}, \mathbf{0}] \in \mathbb{R}^{p \times |J|}$). We are now prepared to present our error bound for $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\dagger}\| = \|\widehat{\boldsymbol{\beta}}_0 - \boldsymbol{\beta}^{\dagger}\|$. Recall that $\mathcal{K} = \bigcup_{s=0}^{d} \mathcal{K}_s$, where d denotes the maximal number of association between response variables. Define the true maximal number of association as $d^* = \{\|\boldsymbol{k}\|_0 : (\boldsymbol{k}, j) \in \mathcal{S}, \boldsymbol{\beta}_k^{\dagger} \neq \mathbf{0}\}$.

Theorem 2 Let B, B_1, B_2 and B' be positive absolute constants, and let $\xi \geq 1$ be fixed. Suppose that d is chosen so that $d^* \leq d$ and that Assumptions 1-4 hold.

(i) Under the Poisson categorical response model, if $\lambda = \xi BC(\sqrt{\Lambda m/n} + \sqrt{\Lambda \log |\mathcal{G}|/n})$ with $0 \le (\xi - 1)(\sqrt{m/n} + \sqrt{\log |\mathcal{G}|/n}) \le B_2$, $\lambda \le R\kappa_J^{Pois}\{6\sqrt{|\mathcal{S}|}\}^{-1}$, and $(m/n + \log |\mathcal{G}|/n) \le B_1$.

 $\min\{1, \kappa_I^{Pois}(AC^2|\mathcal{S}|)^{-1}\}, then$

$$\|\boldsymbol{\theta}^{\dagger} - \widehat{\boldsymbol{\theta}}\| = \|\boldsymbol{\beta}^{\dagger} - \widehat{\boldsymbol{\beta}}_{0}\| \leq \frac{6\xi BC\sqrt{|\mathcal{S}|}}{\kappa_{J}^{Pois}} \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{\log|\mathcal{G}|}{n}}\right),$$

with probability at least $1 - e^{-B'(\xi-1)^2(m+\log|\mathcal{G}|)}$.

(ii) Under the multinomial categorical response model, if $\lambda = \xi BC(\sqrt{m/n} + \sqrt{\log |\mathcal{G}|/n}), \lambda \leq R\kappa_I^{Mult} \{6\sqrt{|\mathcal{S}|}\}^{-1}$, and $(m/n + \log |\mathcal{G}|/n) \leq B_1 \cdot \min\{1, \kappa_I^{Mult} (AC^2|\mathcal{S}|)^{-1}\}$, then

$$\|\boldsymbol{\theta}^{\dagger} - \widehat{\boldsymbol{\theta}}\| = \|\boldsymbol{\beta}^{\dagger} - \widehat{\boldsymbol{\beta}}_{0}\| \leq \frac{6\xi BC\sqrt{|\mathcal{S}|}}{\kappa_{\boldsymbol{J}}^{Mult}} \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{\log|\mathcal{G}|}{n}}\right),$$

with probability at least $1 - e^{-B'(\xi-1)^2(m+\log|\mathcal{G}|)}$.

In Lemma 3, we show that, under certain regularity assumptions, $\kappa_J^{Mult} \times |J|^{-1} \kappa_J^{Pois}$ and $\kappa_J^{Pois} = O(1)$. However, we cannot conclude that the Poisson-likelihood based estimator is better its multinomial counterpart because the data generating models being assumed are fundamentally different.

The result of Theorem 2 indicates that under the Poisson or multinomial sampling scheme, assuming $\kappa_J^{Pois} = O(1)$ or $\kappa_J^{Mult} = (|J|^{-1})$, we can achieve a Frobenius norm error rate of $O(\sqrt{m/n} + \sqrt{\log |\mathcal{G}|/n})$. Call that $|\mathcal{G}|$ is the number of groups of parameters being penalized in (15) under general local association learning. This would seem to suggest that having fewer groups is beneficial, but this term is counterbalanced with m, which is the largest number of parameters per group. Hence, since a small number of groups would require a larger number of parameters per group, there is a clear tradeoff between the two. Importantly, both terms are multiplied by $|\mathcal{S}|$, so ideally, we will select a number of groups that leads to small \mathcal{S} without inflating $|\mathcal{G}|$ or m.

Though not made explicit in our bounds, the association of a well-specified d is apparent in our error bounds. If $d=q\gg d^*$, then both m and $|\mathcal{G}|$ will be larger than if d were specified closer to d^* . Of course, if $d< d^*$, we could not expect consistent estimation since this will force estimates of truly nonzero associations to be zero.

The following corollary is a special case of Theorem 2 for multinomial categorical response model, letting $\mathcal{G} = \mathcal{G}_{\text{global}}$ or $\mathcal{G}_{\text{local}}$. Here, we replace the quantities from Theorem 2 with more explicit versions.

Corollary 1 Under the conditions of Theorem 2 and assuming the multinomial categorical response model, if tuning parameters are chosen according to part (ii) of Theorem 2 and $w_g = 1$ for all $g \in \mathcal{G}$, then the following statements holds.

1. For global association learning, with probability as specified in Theorem 2,

$$\|\boldsymbol{\theta}^{\dagger} - \widehat{\boldsymbol{\theta}}\| \leq \frac{6\xi BC}{\kappa_{\boldsymbol{J}}^{Mult}} \sqrt{\sum_{\boldsymbol{k} \in \mathcal{K}} \mathbf{1}(\boldsymbol{\beta}_{\boldsymbol{k}}^{\dagger} \neq 0)} \left\{ \sqrt{\frac{p \prod_{\ell=1}^{d} (J_{(\ell)} - 1)}{n}} + \sqrt{\frac{\log \sum_{l=0}^{d} \binom{q}{l}}{n}} \right\},$$

where $J_{(1)}, \dots, J_{(q)}$ is a permutation of J_1, \dots, J_q such that $J_{(1)} \geq J_{(2)} \geq \dots \geq J_{(q)}$.

2. For local association learning $(t \ge 2)$, with probability as specified in Theorem 2,

$$\|\boldsymbol{\theta}^{\dagger} - \widehat{\boldsymbol{\theta}}\| \leq \frac{6\xi BC}{\kappa_{\boldsymbol{J}}^{Mult}} \|\boldsymbol{\beta}^{\dagger}\|_{0,\mathcal{G}} \left\{ \sqrt{\frac{\max_{(\boldsymbol{k},j)\in\mathcal{G}} |\boldsymbol{k}|_{\boldsymbol{J}} \cdot p_{j}}{n}} + \sqrt{\frac{\log t + \log \sum_{l=0}^{d} \binom{q}{l}}{n}} \right\},$$

where
$$\|\boldsymbol{\beta}^{\dagger}\|_{0,\mathcal{G}} = \sqrt{\sum_{(\boldsymbol{k},j)\in\mathcal{G}} \mathbf{1}(\boldsymbol{\beta}_{\boldsymbol{k},j}^{\dagger} \neq 0)}$$
.

For the multinomial sampling scheme, as |J| increases, the upper bound of the estimation error worsens. This suggests that increasing the dimension of the response will lead to poorer estimation. We continue by demonstrating the reasonableness of Assumption 4, particularly regarding its validity under the assumption of random predictors. In section 9 of Wainwright (2019), the restricted strong convexity condition has been derived under a GLM setting (See Theorem 9.36 in Wainwright (2019)). Here, we generalized their results to a multivariate GLM setting, and calibrate the Rademacher complexity term of the group lasso penalty according to multivariate GLM setting. We summarize the results in Section S9 of the online supplementary material and incorporate both the multinomial and Poisson categorical response settings into the following lemma.

Lemma 3 Under Assumptions 1-3, assume that x_1, \ldots, x_n are independent and identically distributed with zero mean. Additionally, assume that for some positive constants (α, β) , we have

$$\mathbb{E} \|\Delta \boldsymbol{\theta} \boldsymbol{x}_i\|^2 \ge \alpha \quad and \quad \mathbb{E} \|\Delta \boldsymbol{\theta} \boldsymbol{x}_i\|^4 \le \beta,$$

for all vectors $\Delta \theta \in \mathcal{M}_H$ such that $\|\Delta \theta\| = 1$. Then, the following results hold. For both multinomial and Poisson categorical response models with the reparameterized group lasso penalty $\Phi_{\mathcal{G}}$, the restricted strong convexity condition (30) in Assumption 4 holds with probability at least $1 - c_1 e^{-c_2 n}$. Furthermore, $\kappa_J^{Mult} \approx |J|^{-1} \kappa_J^{Pois}$ and $\kappa_J^{Pois} = O(1)$.

The above lemma justifies condition (30) and the typical behavior of the curvature κ in Assumption 4, showing that both will hold with high probability under mild assumptions.

7. Numerical studies

7.1. Data generating models and competitors

We present a series of simulations designed to evaluate the performance of the proposed methods and applicable variations of existing methods under various scenarios. We consider a range of parameters, including different sample sizes, dimensions, and three different model generation schemes. A detailed description of this study is in Section S5 of the online supplementary material.

For $N_{\rm rep}=100$ independent replications, we simulate data from the multivariate multinomial response regression model, with d=4, q=4, and $(J_1,J_2,J_3,J_4)=(2,2,2,3)$. With $n\in\{100,300,500,1000,2000\}$ training samples, each observation x_i (excluding its riest element, because the first row of x_i is always 1) is drawn from a multivariate normal distribution $N_p(\mathbf{0},\Sigma)$, where the covariance entries $\Sigma_{jk}=0.5^{|j-k|}$ are defined for all pairs $(j,k)\in[p-1]\times[p-1]$. Given a coefficient matrix $\boldsymbol{\beta}^*\in\mathbb{R}^{|J|\times p}$, the probability vector is given by

$$\operatorname{vec}_{\boldsymbol{J}}\left(\boldsymbol{\pi}^{\boldsymbol{J}}(\boldsymbol{x})\right) = \frac{\exp(\boldsymbol{H}\boldsymbol{\beta}^*\boldsymbol{x})}{\langle \mathbf{1}_{|\boldsymbol{J}|}, \exp(\boldsymbol{H}\boldsymbol{\beta}^*\boldsymbol{x})\rangle},$$

from which we generate the response $y_i \in \mathbb{R}^p$ as a realization of

Multinomial
$$\left(n_i, \text{vec}_J\left(\boldsymbol{\pi}^J(\boldsymbol{x}_i)\right)\right), \quad n_i = 1.$$
 (31)

This process is also extended to generate 1000 validation samples for model tuning and $N_{\text{test}} = 10000$ test samples to evaluate model performance. We conduct our simulations over a range of dimensions $p \in \{10, 50\}$ to assess scalability and robustness. Let $\mathcal{G}_{\text{global}} = \{(\mathbf{k}, j); \mathbf{k} \in \mathcal{K}, j \in \{1, 2\}\}$ with $p_1 = 1, p_2 = p - 1$ and $\mathcal{G}_{\text{local}} = \{(\mathbf{k}, j); \mathbf{k} \in \mathcal{K}, j \in \{1, \dots, p\}\}$ with $p_1 = p_2 = \dots = p_p = 1$.

We consider three distinct structures for β^* . The parameter generation methods for β^* are designated as Scheme 1, Scheme 2, and Scheme 3, and are discussed in detail in Section S5 of the online supplementary material. These correspond to the three interpretable models—mutual independence, joint independence, and conditional independence, respectively—as presented in Example 2. Moreover, in each scheme, the associations are local—only the intercept and two randomly selected predictors have nonzero associations. Hence, these schemes are most well suited for the versions of our methods using Ω^H_{local} .

In our simulation studies, we will examine the following eight estimators. The first six estimators—O-Mult, O-Pois, L-Mult, L-Pois, G-Mult and G-Pois—are derived using the reparameterization technique. Here O,

L, and G denote estimators using overlapping group lasso with hierarchical structure built on local group $\mathcal{G}_{\text{local}}$, group lasso with local group $\mathcal{G}_{\text{local}}$, and group lasso with global group $\mathcal{G}_{\text{global}}$ penalties, respectively. Additionally, Mult and Pois refer to multinomial and Poisson multivariate categorical response models, respectively. Recall that the data-generating model is based on a multinomial model. Thus, the O-Mult, L-Mult, and G-Mult are penalized maximum likelihood estimators for a correctly specified model. In contrast, O-Pois, L-Pois, and G-Pois can be thought of as M-estimators. The seventh estimator, G-Mult- θ , employs the classical parameterization approach in θ . The eighth estimator, Sep-Mult, is designed to individually address each category in the multinomial vector, providing estimates of each response's probability mass function separately. The method denoted Oracle represents the true parameter, and is included to serve as a baseline.

We employ a train-validation split to select the tuning parameter λ in our simulation study, with the weight parameters $w_{k,j}$ set to 1. Specifically, we choose the candidate tuning parameter that minimizes the cross-entropy loss on the validation set.

7.2. Results

The estimators' performances, evaluated based on Hellinger distance and (joint) misclassification rate on a test set, is displayed in Figures 1 and 2. The Sep-Mult estimator is correctly specified under Scheme 1, where the responses are mutually independent. Unsurprisingly, Sep-Mult outperforms all other estimators under this scheme. Under Schemes 2 and 3 where responses are dependent, we see Sep-Mult perform very poorly relative to the other methods.

The estimators O-Mult, O-Pois, L-Mult, L-Pois, G-Mult, and G-Pois are all based on our parameterization. Considering the overall performance based on the Hellinger distance and the misclassification rate, the O-Mult estimator is generally the most favorable. This is expected as this method is based on a correct specification of the model and can exploit the hierarchical structure of the local associations. The estimator L-Mult tends to perform second best when sample sizes are large. Notably, the estimator O-Pois performs reasonably well when n=100: only O-Mult is evidently better. As n increases, however, O-Pois tends to be outperformed by the methods assuming a multinomial data generating model.

We caution that these results do not imply that estimators based on the multinomial negative log-likelihood are uniformly preferable to those utilizing the Poisson negative log-likelihood. In this case, the multinomial estimators assume a correctly specified model, and thus, as the sample size increases, tend to outperform their Poisson counterparts.

7.3. Poisson data generating model

Simulation study results under the Poisson data generating model are more difficult to interpret than those based the multinomial data generating model. This is partly because when fixing n, the effective sample size for the multinomial estimators is a random variable. Specifically, for each of the n samples, we draw a (possibly large) number of Poisson counts from the conditional distribution in (5). The multinomial estimators treat each count as an independent realization from a single-trial multinomial. Thus, the number of "samples" input into the multinomial estimators can be extremely large and vary greatly from simulation replicate to simulation replicate. For this reason, we exclude results under the Poisson from this manuscript. Nonetheless, to briefly summarize the results we observed in the simulation scenarios we considered (specifically Scheme 2 and 3), we found that under the Poisson data generating models, both L-Pois and G-Pois significantly outperformed L-Mult and G-Mult.

8. Colon tissue data analysis

We apply our proposed method to a dataset obtained from the Gene Expression Omnibus (accession number GDS3268). This dataset consists of 202 colon tissue samples, with each sample containing expression levels for tens of thousands of genes. Additionally, each sample includes three categorical response labels: (i) patient state (normal or ulcerative colitis), (ii) tissue inflammation status (inflamed or uninflamed), and (iii)

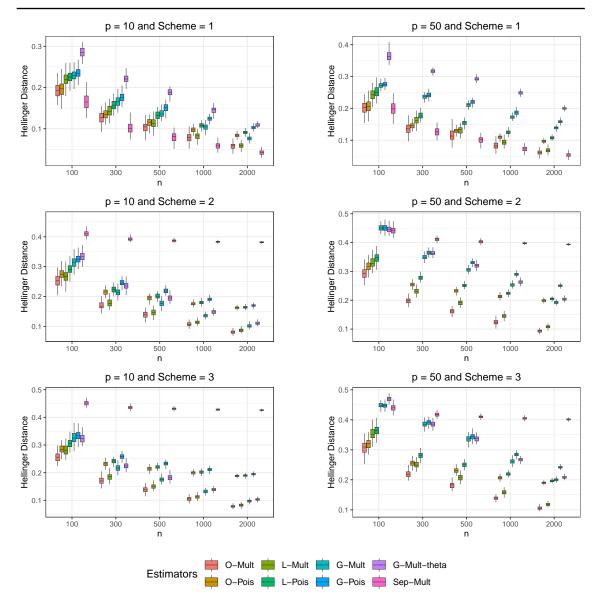


Fig. 1: Hellinger distances for the competing estimators with $p \in \{10, 50\}$ and Scheme $\in \{1, 2, 3\}$ as n varies.

anatomical location (ascending colon, descending colon, sigmoid colon, or terminal ileum). Thus, we have multiple response variables to jointly model.

First, the gene expression measurements are log-transformed and normalized. We then select a subset of the most variable genes, ranking them according to their median absolute deviation and retaining the top p genes. Highly correlated genes are pruned in order to avoid retaining genes with overlapping information. We consider several values for p; 100, 200, 300, 400, and 500.

We randomly partition the 202 samples into three subsets: 100 for training, 50 for validation (used for tuning parameter selection), and 52 for testing. On the training set, we fit the candidate estimators introduced in Section 7, modeling joint associations of at most order d = 1, 2, 3, and compare these to a

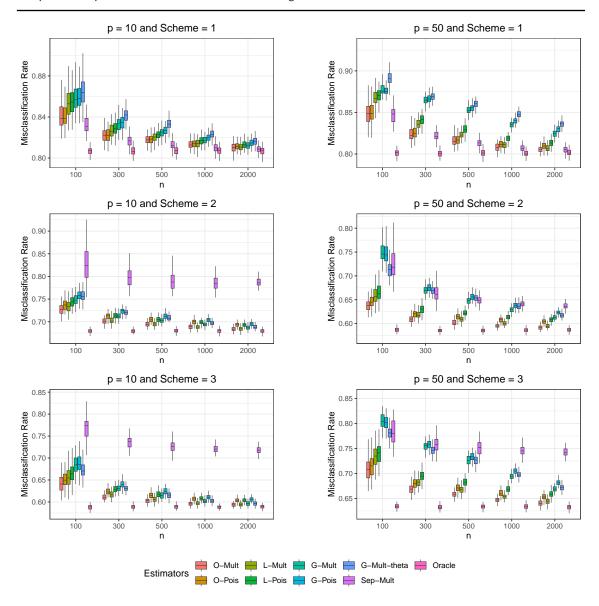


Fig. 2: Misclassification rates for the competing estimators with $p \in \{10, 50\}$ and Scheme $\in \{1, 2, 3\}$ as n varies.

separate logistic regression baseline. The tuning parameters are selected by minimizing cross-entropy loss on the validation set. Specifically, we tune (λ, r) using the weight $w_{k,j} = r^{\max\{\|k\|_0 - 1, 0\}}$, which penalizes higher-order associations more heavily. A standard fine grid is used for λ , whereas for r we choose from the smaller set $r \in \{1, 2, 4, 6, 10\}$ to improve computational efficiency.

The predictive performance of each fitted model is evaluated on the held-out test set using the misclassification rate, i.e., the proportion of samples in the test set whose predicted response combination does not match the observed response combination. We also evaluated predictive performance using empirical cross-entropy, estimating $-\mathbb{E}_{X,Z}\log[\hat{P}(Z\mid X)]$.

In Table 1, we report the misclassification rates and empirical cross-entropy averaged across 1000 random training-validation-testing splits. Overall, the estimators G-Mult with d=2 and d=3 consistently achieves the best performance on the misclassification rates, with most proposed methods outperforming the separate logistic regression baseline (Sep-Mult). Moreover, as p increases, our proposed methods continue to improve, the Sep-Mult method shows minimal change. The nearly identical results for d=2 and d=3 suggest that, although our model was designed to capture up to third-order associations, the higher-order interactions permitted by the model were effectively negligible in this dataset.

From the empirical cross-entropy perspective, the estimators G-Mult with d=2 and d=3, and O-Mult, perform equally well and outperform the others.

Table 1. Average misclassification rates and cross-entropy over 1000 independent training/validation/testing splits of the colon tissue dataset.

		Misclassification rate					Cross entropy				
		p = 100	p = 200	p = 300	p = 400	p = 500	p = 100	p = 200	p = 300	p = 400	p = 500
d=3	L-Mult	0.5033	0.4754	0.4712	0.4709	0.4662	1.4142	1.2983	1.2954	1.2942	1.2820
d=3	L-Pois	0.5129	0.4785	0.4732	0.4737	0.4687	2.6878	2.6415	2.6407	2.6411	2.6298
d=3	G-Mult	0.5017	0.4647	0.4520	0.4403	0.4332	1.3956	1.2917	1.2606	1.2384	1.2082
d=3	G-Pois	0.5072	0.4701	0.4637	0.4458	0.4342	2.7227	2.7053	2.7711	2.6384	2.6246
d=3	O-Mult	0.5070	0.4702	0.4586	0.4474	0.4416	1.4006	1.2961	1.2687	1.2481	1.2240
d=3	O-Pois	0.5111	0.4739	0.4682	0.4606	0.4533	2.7197	2.6270	2.5969	2.5932	2.5794
d=2	L-Mult	0.5022	0.4743	0.4700	0.4698	0.4652	1.4140	1.2981	1.2951	1.2937	1.2819
d=2	L-Pois	0.5117	0.4771	0.4721	0.4724	0.4676	2.6839	2.6388	2.6385	2.6396	2.6298
d=2	G-Mult	0.5015	0.4643	0.4515	0.4400	0.4327	1.3956	1.2918	1.2606	1.2385	1.2082
d=2	G-Pois	0.5067	0.4698	0.4634	0.4455	0.4341	2.7233	2.7100	2.7742	2.6387	2.6251
d=2	O-Mult	0.5034	0.4667	0.4542	0.4426	0.4358	1.3924	1.2908	1.2620	1.2403	1.2112
d=2	O-Pois	0.5070	0.4680	0.4581	0.4472	0.4388	2.7214	2.6324	2.6023	2.5974	2.5756
d=1	L-Mult	0.5227	0.4946	0.4924	0.4927	0.4873	1.4357	1.3238	1.3218	1.3200	1.3116
d=1	L-Pois	0.5346	0.4969	0.4903	0.4914	0.4843	3.0891	3.5829	3.3387	3.2668	3.5938
d=1	G-Mult	0.5184	0.4804	0.4677	0.4552	0.4469	1.4113	1.3087	1.2782	1.2530	1.2228
$_{d=1}$	G-Pois	0.5276	0.4876	0.4818	0.4622	0.4499	3.3714	3.9934	3.2810	3.0970	3.0893
d=1	O-Mult	0.5171	0.4793	0.4663	0.4545	0.4469	1.4048	1.3045	1.2759	1.2517	1.2210
d=1	O-Pois	0.5244	0.4813	0.4707	0.4570	0.4467	3.3548	3.6178	3.0861	3.0439	3.0182
	G-Mult-θ	0.6023	0.5769	0.5741	0.5680	0.5604	1.7416	1.6496	1.6420	1.6358	1.6162
	Sep-Mult	0.5117	0.5044	0.5102	0.5174	0.5152	1.4367	1.3849	1.4108	1.4290	1.4175
	Null model	0.9375	0.9375	0.9375	0.9375	0.9375					

9. Discussion

This article introduces an alternative approach to multivariate categorical response regression which relies on an explicit subspace decomposition. Our proposed decomposition allows practitioners to use standard regularization techniques to select the order of associations, and bypasses the issue of dependence on choice of identifability constraints. There are three key directions for future research.

9.1. More computationally efficient approaches to hierarchically-structured selection

The use of the overlapping group lasso penalty $\Omega_{\mathcal{G}}^{\mathcal{G}}$ to select s adhering to a hierarchy is especially appealing in practice, but leads to an estimator which is more computationally intensive than its counterpart excluding hierarchical constraints. In the future, it is important to consider alternative approaches to regularization that may be less computationally intensive, but encourage the desired hierarchy. One such approach may be to utilize the latent overlapping group lasso penalty (Obozinski et al., 2011), which allows the optimization problem to be separable across the (latent) parameters being penalized. This can afford more efficient computational algorithms and schemes to be developed. The estimator based on latent overlapping group lasso penalty is distinct from that based on the overlapping group lasso penalty in the sense that their solution paths are fundamentally distinct, but both can be used to enforce hierarchical constraints. Consequently, the

theoretical properties of the estimator based on the latent overlapping group lasso penalty are not immediate from the results we derived in Section 6, so this direction is nontrivial.

Another approach is to use a separable (non-overlapping) approximation to the overlapping group lasso penalty. Specifically, Qi and Li (2024) recently proposed a separable relaxation of the overlapping group lasso penalty, and showed that in terms of squared estimator error, the estimator using their relaxation is statistically equivalent to that using the overlapping group lasso penalty. Notably, because the relaxation is separable, the corresponding estimator can be computed much more efficiently—roughly at the same cost as estimators using nonoverlapping group lasso penalization schemes.

9.2. Other representations in predictors

Recall that in our model (6) and (7), $\theta x = H\beta x$ is linear in x. The subspace decomposition model can be extended to accommodate scenarios where the relationship with x is not necessarily linear, i.e., $\theta(x) = H\rho(x)$, where $\theta: \mathbb{R}^p \to \mathbb{R}^{|J|}$ and $\rho: \mathbb{R}^p \to \mathbb{R}^{\sum_{s=0}^d L_s}$. Here, ρ can be associated with both parametric models, such as polynomial regression, and non-parametric models, including splines, kernel-based models, additive models, and deep learning architectures.

9.3. Application to the analysis of large contingency tables

Finally, a direction not explored in this article is the use of our estimator for fitting traditional log-linear models for contingency tables. The traditional log-linear model is a special case of our model with the predictor consisting of the intercept only. Effect selection in standard log-linear models has been studied in the past (e.g, see Nardi and Rinaldo, 2012), but in the asymptotic regime with $n \to \infty$ and all other model dimensions fixed. Thus, it is of particular interest to study whether our finite sample error bounds can be applied, or even refined, in this context.

Data availability statement

Open-source R code implementing the proposed methods described in this paper is available via our GitHub repository. This repository provides clearly documented R scripts to reproduce all numerical results and figures presented. Additionally, it includes instructions for obtaining and preprocessing the publicly available gene expression dataset from the Gene Expression Omnibus (GEO) database (accession number: GDS3268).

Acknowledgments

A. J. Molstad's research was supported in part by a grant from the National Science Foundation (DMS-2415067).

References

- A. Agresti. Categorical Data Analysis. John Wiley & Sons, New York, 2 edition, 2002.
- W. P. Bergsma and T. Rudas. Marginal models for categorical data. The Annals of Statistics, 30(1):140–159, 2002.
- S. Bhowal and S. Mukherjee. Rates of convergence of the magnetization in the tensor curie-weiss potts model. Journal of Statistical Physics, 192(1):2, 2024.
- A. Bretto. Hypergraph theory. An introduction. Mathematical Engineering. Cham: Springer, 1, 2013.
- J. Cheng, E. Levina, P. Wang, and J. Zhu. A sparse ising model with covariates. Biometrics, 70(4):943–953, 2014.
- R. Christensen. Log-linear models and logistic regression. Springer Science & Business Media, 1997.
- K. Deng, X. Zhang, and A. J. Molstad. Multi-response linear discriminant analysis in high dimensions. *Journal of Machine Learning Research*, 25(413):1–66, 2024.
- M. Ekeberg, C. Lövkvist, Y. Lan, M. Weigt, and E. Aurell. Improved contact prediction in proteins: using pseudolikelihoods to infer potts models. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 87(1):012707, 2013.
- G. F. Glonek. A class of regression models for multivariate categorical responses. Biometrika, 83(1):15-28, 1996.
- G. F. Glonek and P. McCullagh. Multivariate logistic models. *Journal of the Royal Statistical Society: Series B* (Methodological), 57(3):533–546, 1995.

- T. Hastie, R. Tibshirani, and J. Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media, 2009.
- F. Herrera, F. Charte, A. J. Rivera, M. J. Del Jesus, F. Herrera, F. Charte, A. J. Rivera, and M. J. del Jesus. *Multilabel classification*. Springer, 2016.
- T. Hirakida, H. Kouno, J. Takahashi, and M. Yahiro. Interplay between sign problem and z 3 symmetry in three-dimensional potts models. *Physical Review D*, 94(1):014011, 2016.
- R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *The Journal of Machine Learning Research*, 12:2297–2334, 2011.
- S. Keetelaar, N. Sekulovski, D. Borsboom, and M. Marsman. Comparing maximum likelihood and maximum pseudolikelihood estimators for the ising model. *advances. in/psychology*, 2:e25745, 2024.
- S. Klamt, U.-U. Haus, and F. Theis. Hypergraphs and cellular networks. *PLoS computational biology*, 5(5): e1000385, 2009.
- J. B. Lang. Maximum likelihood methods for a generalized class of log-linear models. The Annals of Statistics, 24(2):726-752, 1996.
- K. Lange. MM Optimization Algorithms, volume 147. SIAM, 2016.
- A. L. M. Levada, N. D. A. Mascarenhas, and A. Tannús. Pseudo-likelihood equations for Potts model on higher-order neighborhood systems: A quantitative approach for parameter estimation in image analysis. Brazilian Journal of Probability and Statistics, 23(2):120 140, 2009. doi: 10.1214/08-BJPS018.
- Q. Mai, Y. Yang, and H. Zou. Multiclass sparse discriminant analysis. Statistica Sinica, 29(1):97–111, 2019.
- P. McCullagh and J. A. Nelder. Generalized Linear Models. Chapman and Hall, London, 2 edition, 1989.
- G. Molenberghs and E. Lesaffre. Marginal modelling of multivariate categorical data. Statistics in Medicine, 18 (17-18):2237–2255, 1999.
- A. J. Molstad and A. J. Rothman. A likelihood-based approach for multivariate categorical response regression in high dimensions. Journal of the American Statistical Association, 118(542):1402–1414, 2023.
- A. J. Molstad and X. Zhang. Conditional probability tensor decompositions for multivariate categorical response regression. arXiv preprint arXiv:2206.10676, 2022.
- S. Mukherjee, J. Son, and B. B. Bhattacharya. Estimation in tensor ising models. Information and Inference: A Journal of the IMA, 11(4):1457–1500, 2022.
- Y. Nardi and A. Rinaldo. The log-linear group-lasso estimator and its asymptotic properties. Bernoulli, 18(3): 945–974, 2012.
- S. N. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- G. Obozinski, L. Jacob, and J.-P. Vert. Group lasso with overlaps: the latent group lasso approach. $arXiv\ preprint\ arXiv:1110.0413,\ 2011.$
- S. Okabayashi, L. Johnson, and C. J. Geyer. Extending pseudo-likelihood for potts models. Statistica Sinica, pages 331–347, 2011.
- J. Palmgren. Regression models for bivariate binary responses. Technical Report 101, Department of Biostatistics, University of Washington, Seattle, 1989.
- N. Parikh, S. Boyd, et al. Proximal algorithms. Foundations and trends® in Optimization, 1(3):127-239, 2014.
- B. F. Qaqish and A. Ivanova. Multivariate logistic models. Biometrika, 93(4):1011-1017, 2006.
- M. Qi and T. Li. The non-overlapping statistical approximation to overlapping group lasso. *Journal of Machine Learning Research*, 25(115):1–70, 2024.
- Z. Razaee and A. Amini. The potts-ising model for discrete multivariate data. Advances in Neural Information Processing Systems, 33:13727–13737, 2020.
- J. Read, B. Pfahringer, G. Holmes, and E. Frank. Classifier chains: A review and perspectives. Journal of Artificial Intelligence Research, 70:683-718, 2021.
- T. Robiglio, L. Di Gaetano, A. Altieri, G. Petri, and F. Battiston. Higher-order ising model on hypergraphs. arXiv preprint arXiv:2411.19618, 2024.
- K. Shimagaki and M. Weigt. Selection of sequence motifs and generative hopfield-potts models for protein families. Physical Review E, 100(3):032128, 2019.
- S. Tanaka, R. Tamura, and N. Kawashima. Phase transition of generalized ferromagnetic potts model-effect of invisible states. In *Journal of Physics: Conference Series*, volume 297, page 012022. IOP Publishing, 2011.
- P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. submitted to SIAM Journal on Optimization, 2(3), 2008.
- M. Vincent and N. R. Hansen. Sparse group lasso and high dimensional multinomial classification. Computational Statistics & Data Analysis, 71:771–786, 2014.
- M. J. Wainwright. High-dimensional statistics: A non-asymptotic viewpoint, volume 48. Cambridge university press, 2019.

- A. Wipf, T. Heinzl, T. Kaestner, C. Wozar, et al. Generalized potts-models and their relevance for gauge theories. SIGMA. Symmetry, Integrability and Geometry: Methods and Applications, 3:006, 2007.
- X. Yan and J. Bien. Hierarchical sparse modeling: A choice of two group lasso formulations. *Statistical Science*, 32:531–560, 2017.
- H. Yu, Y. Zhong, L. Hong, C. Xu, W. Ren, X. Gong, and H. Xiang. Spin-dependent graph neural network potential for magnetic materials. *Physical Review B*, 109(14):144426, 2024.
- M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(1):49–67, 2006.
- P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *Annals of Statistics*, 37(6A):3468–3497, 2009.