Statistica Sinica Preprint No: SS-2021-0082

Title	Simplex-based Multinomial Logistic Regression with Diverging Numbers of Categories and Covariates
Manuscript ID	SS-2021-0082
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202021.0082
Complete List of Authors	Sheng Fu, Piao Chen, Yufeng Liu and Zhisheng Ye
Corresponding Author	Piao Chen
E-mail	p.chen-6@tudelft.nl
Notice: Accepted version subject to English editing.	

Simplex-based Multinomial Logistic Regression with Diverging Numbers of Categories and Covariates

Sheng Fu¹, Piao Chen², Yufeng Liu³ and Zhisheng Ye¹

¹*National University of Singapore*, ²*Delft University of Technology*

³*University of North Carolina at Chapel Hill*

Abstract: Multinomial logistic regression models are popular in multiclass classification analysis, but existing models suffer several intrinsic drawbacks. In particular, the parameters could not be uniquely determined because of the over-specification. Some additional constraints have been imposed to refine the model but such modifications can be inefficient and complicated. In this paper, we propose a novel and efficient simplex-based multinomial logistic regression technique, seamlessly connecting binomial and multinomial cases under a unified framework. Compared with the existing models, our model has less parameters and is free of any constraints, which can be efficiently solved via the Fisher scoring algorithm. In addition, the proposed model enjoys some theoretical advantages including Fisher consistency and sharp comparison inequality. Under mild conditions, we establish the asymptotical normality and convergence for the new model even when the numbers of categories and covariates increase with the sample size. The proposed framework is illustrated by extensive simulations and real applications.

Key words and phrases: Asymptotics; Classification; Fisher consistency; Kernel learning; MLR; Simplex coding scheme

Corresponding author: p.chen-6@tudelft.nl

1. Introduction

Logistic Regression (LR) is the most frequently used regression model for the analysis of categorical outcomes (Cramer, 2003; Yee, 2015; Fang and Yi, 2021; Mo and Liu, 2021). As a prevalent tool to practitioners, LR has been widely applied in epidemiology, biology, economics, social sciences, etc (Hosmer et al., 2013; Lemeshow and Hosmer, 2014). According to the number of categories, LR models could be divided into two types, i.e., binomial LR and multinomial LR (MLR). The MLR model is also known as the conditional maximum entropy model in natural language processing (Berger et al., 1996; Malouf, 2002) and the softmax regression in neural network (Ng et al., 2013; Goodfellow et al., 2016).

The statistical theory of binomial LR has been well established, whereas modeling and inference for MLR is more complicated. There have been numerous attempts to generalize the original binomial LR to the multinomial case; see McCullagh and Nelder (1989); Hastie et al. (2009); Tutz (2011). For a k -categorical regression problem, a natural approach is to estimate k regression functions, one for each category. The unknown parameters in these functions are typically jointly estimated by maximum likelihood, or through Bayesian updating. As shown in Bühlmann and Van De Geer (2011, §3.3.3), such an extension is over-specified with unidentifiable parameters. Generally, $k-1$ functions are sufficient to determine a k -categorical LR model. For instance, the binomial LR ($k = 2$) is defined by a single regression function. Therefore, additional restrictions are needed on the regression functions to make the model identifiable. Roughly speaking, there are two common schemes. The first is to prespecify a reference category and set its regression function to be zero, and the other is to use the sum-to-zero constraint on the k functions (Hastie et al., 2015, §3.3). Both extensions subsume the

binomial LR as a special example, and they will be introduced in Section 3. However, the extra constraints complicate the parameter estimation and theoretical analysis. Specifically, the reference-based MLR does not treat all categories equally, and the explicit sum-to-zero constraint increases the computational cost. In addition, the relationship between these two constrained MLR (CMLR) models is not entirely clear.

A powerful method to circumvent the explicit constraints is to use the simplex coding. The basic idea is to construct a k -vertex simplex structure in the $(k - 1)$ -dimensional Euclidean space, where each vertex represents one category. The covariate vector of each instance is then mapped to a point in this $(k - 1)$ -dimensional space. In other words, only $k - 1$ regression functions need to be estimated under the simplex coding scheme, and the non-identifiability issue is automatically resolved without further constraint on these $k - 1$ functions. Therefore, the simplex coding for multiclass learning is expected to have lower computational complexity in model training (Hill and Doucet, 2007; Lange and Wu, 2008; Mroueh et al., 2012). Such a scheme is also called angle-based classification by Zhang and Liu (2014), Zhang et al. (2018) and Fu et al. (2018) because the predicted label of a new observation corresponds to the vertex that has the smallest angle with the mapped $(k - 1)$ -dimensional point of the covariates. This geometric interpretation makes the coding scheme easy to understand.

The first objective of our study is to remove the cumbersome constraints in the existing MLR models, and propose a novel MLR model using the delicate simplex structure, called Simplex-based MLR (SMLR). Inheriting from the simplex coding, the redundancy of the categorical space is removed and the representation of the regression functions is identifiable. Hence, the resulting SMLR model has clear geometric explanation and it is computation-

ally efficient in the parameter estimation. Compared with the regular CMLRs, the SMLR model enjoys more parsimonious parameter specification. In specific, SMLR avoids the subjective selection of a reference category, gets rid of the sum-to-zero constraint, and provides a symmetric insight on all categories by treating them equally. With fewer parameters, the likelihood estimation of the SMLR solves an unconstrained optimization problem, which can be efficiently implemented via the Fisher scoring algorithm. The proposed SMLR can be treated as a unified framework recovering CMLRs, but the involved parameters have different interpretations.

The second objective of the study is to establish the asymptotic properties of the MLEs of the SMLR with a diverging number of categories, which is peculiar to multcategory classification applications. In practical problems, the granularity of the classification, in terms of the number of categories, is usually determined based on the size of the available training data (Dekel and Shamir, 2010). For example, photo sharing websites allow users to annotate their photos with keywords. The key task is to recommend keywords whenever new photos are uploaded. Assuming there is no restrictions on the keywords that may be used, the set of distinct keywords is likely to grow as more and more photos are uploaded to the site. Similarly, web directory classification with Yahoo! taxonomies incurred some rare categories, which are ignored under the small sample size, but considered for large sample problems (Liu et al., 2005). Other examples include textual document categorization (Dekel and Shamir, 2010), and identification of flowers, plants and products using image (Nilsback and Zisserman, 2008; Deng et al., 2010). The phenomenon of diverging number of categories has attracted some attention in the literature, and most existing studies are concerned with

algorithmic development, e.g., distributed computing, hierarchical classification, and penalization techniques. For instance, Deng et al. (2009, 2010) exploited the semantic hierarchy of categories to obtain more informative image classifiers. Based on the hierarchical structure of categories, Price et al. (2019) proposed a group-fused MLR model to automatically combine the categories. However, there is scant literature on the asymptotic studies on the MLR model when the class size increases with the sample size, probably because the constraint on the existing CMLRs makes the asymptotic properties hard to establish. On the contrary, when the dimensions of the covariates and the categories are fixed, asymptotic properties of the MLRs have been well-established (Fahrmeir and Kaufmann, 1985; Van der Vaart, 1998; Tutz, 2011). In this study, we focus on the asymptotics of the SMLR model with varying category sizes. Meanwhile, an important byproduct is that asymptotics under a diverging number of covariates are also established, which has received scant attention in the literature of MLR models.

The third objective of this study is to show the theoretical advantages of the SMLR model under certain settings. In particular, we explore the kernel learning for SMLR, which enjoys faster convergence rate than the existing MLR models. The convergence analysis for kernel MLR under a diverging number of categories is rarely investigated in the literature, and this study fills the gap by establishing the consistency of kernel SMLR while letting the number of categories k go to infinity at the order of $o(n)$. In addition, we show that the proposed SMLR enjoys some desirable statistical properties including Fisher consistency and the comparison inequality, which are fundamental in understanding the nature of the SMLR model.

The rest of the paper is organized as follows. In Section 2, we introduce the notations

and briefly review the regular MLR models. In Section 3, we propose the simplex-based MLR method, and explore its connections with the regular MLR models. We establish the asymptotical results for the SMLR model with diverging number of parameters in Section 4. Simulation studies and real applications are used to demonstrate the performances of the proposed approach in Section 5. Section 6 provides some discussions on the MLR models and related problems. The main proofs are given in the supplementary materials.

Throughout the paper, $\mathbf{0}$ and $\mathbf{1}$ are respectively vectors of 0's and 1's, and $\mathbf{e}_j$ is the j -th column vector of an identity matrix $\mathbf{I}$, whose dimensions can be contextually inferred. $\mathbf{diag}(\mathbf{u})$ is a diagonal matrix with entries determined by a vector $\mathbf{u}$. The vectorization of a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{m \times n}$ is defined as

$$\vec{\mathbf{A}} := \text{vec}(\mathbf{A}) = (a_{11}, \dots, a_{m1}, a_{12}, \dots, a_{m2}, \dots, a_{1n}, \dots, a_{mn})^\top \in \mathbb{R}^{mn}.$$

Let $\|\mathbf{A}\|_2$ be the spectral norm defined as the largest singular value of $\mathbf{A}$. Generally, $\|\mathbf{A}\| := \sqrt{\sum_{i,j} a_{ij}^2}$ is the Frobenius norm, including the Euclidean norm of a vector. For a square matrix $\mathbf{A}$, $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ denote the maximum and minimum eigenvalues, respectively. For any matrices $\mathbf{A}$ and $\mathbf{B}$ of the same dimensions, $\mathbf{A} \succeq \mathbf{B}$ or $\mathbf{B} \preceq \mathbf{A}$ denotes $\mathbf{A} - \mathbf{B}$ is positive semi-definite. $\mathbf{A} \succ \mathbf{0}$ implies that $\mathbf{A}$ is positive definite, where $\mathbf{0}$ is a matrix of 0's. For $\mathbf{A} \succ \mathbf{0}$, $\mathbf{A}^{\frac{1}{2}}$ denotes its symmetric square root with $\mathbf{A} = (\mathbf{A}^{\frac{1}{2}})^2$, and $\mathbf{A}^{-\frac{1}{2}} = (\mathbf{A}^{\frac{1}{2}})^{-1}$. The Kronecker product of two matrices $\mathbf{A}$ and $\mathbf{B}$ is denoted by $\mathbf{A} \otimes \mathbf{B}$.

2. Review of Regular MLR Models

Consider a multicategory classification problem with k possible categories $\mathcal{Y} \triangleq \{1, 2, \dots, k\}$.

Suppose that we are given a set of observations $\mathcal{D} = \{(\mathbf{x}_i, y_i), i = 1, \dots, n\}$, where $\mathbf{x}_i =$

$(x_{i1}, \dots, x_{id})^\top \in \mathcal{X} \subseteq \mathbb{R}^d$ is the covariate vector and $y_i \in \mathcal{Y}$ is the corresponding response category. We define $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times d}$ as the design matrix.

Let $\pi_y(\mathbf{x})$ be the conditional probability that the response category is y for the given covariate $\mathbf{x}$. Note that y must take one and only one class label from $\mathcal{Y}$. We need the sum-to-one condition $\sum_{y=1}^k \pi_y(\mathbf{x}) = 1$ to reflect such implicit nature of multinomial regression. Hence, under the sum-to-one condition, only $k - 1$ free probabilities are informative, and the rest is redundant. For the observations $\mathcal{D}$, the joint distribution is $\prod_{i=1}^n \pi_{y_i}(\mathbf{x}_i)$. To link the probability $\{\pi_1, \dots, \pi_k\}$ to the covariate $\mathbf{x}$, a generic MLR model considers the multinomial-Poisson transformation (Baker, 1994; Lang, 1996)

$$\pi_y(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \boldsymbol{\theta}_y}}{\sum_{j=1}^k e^{\mathbf{x}^\top \boldsymbol{\theta}_j}}, \quad \mathbf{x} \in \mathcal{X}, \quad y = 1, \dots, k, \quad (2.1)$$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_k) \in \mathbb{R}^{d \times k}$ is the regression coefficient matrix and $\boldsymbol{\theta}_j \in \mathbb{R}^d$ is the j -th column vector. Model (2.1) can be interpreted as a neural network (Ripley, 1996). Based on the observations $\mathcal{D}$, the log-likelihood function for (2.1) is

$$\ell_n(\boldsymbol{\theta}) = \log \left\{ \prod_{i=1}^n \pi_{y_i}(\mathbf{x}_i) \right\} = \sum_{i=1}^n \mathbf{x}_i^\top \boldsymbol{\theta}_{y_i} - \sum_{i=1}^n \log \left(\sum_{j=1}^k e^{\mathbf{x}_i^\top \boldsymbol{\theta}_j} \right). \quad (2.2)$$

Note that if we add a common vector $\mathbf{b} \in \mathbb{R}^d$ to each $\boldsymbol{\theta}_j$, the probabilities in (2.1) remain unchanged and $\ell_n(\boldsymbol{\theta} + \mathbf{b}\mathbf{1}_k^\top) = \ell_n(\boldsymbol{\theta})$. Therefore, the MLR model (2.1) is not identifiable with over-specified $\boldsymbol{\theta}$. In fact, $\boldsymbol{\theta}$ has $k - 1$ free columns, since there are $k - 1$ free informative conditional category probabilities. This problem may be readily resolved by adopting some restrictions on $\boldsymbol{\theta}_j$'s, leading to the constrained MLR (CMLR).

There are two customary constraints to refine the parameters in the literature. The first is to choose a reference category, denoted by $r \in \mathcal{Y}$, and set $\boldsymbol{\theta}_r$ as the zero vector $\mathbf{0}$,

leading to the reference-based MLR (CMLR1); see Anderson (1972), Anderson and Blair (1982), Albert and Anderson (1984), Böhning (1992), Krishnapuram et al. (2005), Hastie et al. (2009), among others. As a result, the CMLR1 estimator is given by

$$\hat{\boldsymbol{\theta}}^{\text{rb}} = \arg \max_{\boldsymbol{\theta}} \ell_n(\boldsymbol{\theta}), \quad \text{s.t. } \boldsymbol{\theta}_r = \mathbf{0}. \quad (2.3)$$

The other is a sum-to-zero constraint $\sum_{j=1}^k \boldsymbol{\theta}_j = \mathbf{0}$, and the symmetric constrained MLR (CMLR2) has been studied in Friedman et al. (2000), Zhu and Hastie (2004, 2005), Friedman et al. (2010), Zahid and Tutz (2013b), Hastie et al. (2015), etc. Consequently, the estimator for the CMLR2 model is defined as

$$\hat{\boldsymbol{\theta}}^{\text{sc}} = \arg \max_{\boldsymbol{\theta}} \ell_n(\boldsymbol{\theta}), \quad \text{s.t. } \sum_{j=1}^k \boldsymbol{\theta}_j = \mathbf{0}. \quad (2.4)$$

Note that the total number of parameters in (2.3) or (2.4) is dk , while the exact number of free parameters is $d(k-1)$ owing to the extra constraints. Comprehensive discussions on MLRs can be found in Tutz (2011, §8), Hosmer et al. (2013, §8), Yee (2015) and the references therein.

The above-mentioned CMLRs have some inherent deficiencies. On the one hand, when the reference category for CMLR1 alters, the estimated parameters and the corresponding interpretation change as well. Thus, CMLR1 lacks some systematic insights on the categories, and the choice of reference category is subjective and confusing in practice, especially with the increasing number of categories. On the other hand, as shown in Zhang and Liu (2014), the explicit sum-to-zero constraint in CMLR2 can be theoretically inefficient and computationally expensive. We next propose a competent MLR formulation with appealing properties, which is free of constraints on the parameters.

3. Simplex-based MLR

In this section, we formulate a novel and efficient simplex-based MLR (SMLR) in Section 3.1, develop the estimation procedure for the SMLR model in Section 3.2 and reveal the connections among the proposed SMLR and the existing MLRs in Section 3.3.

3.1 Methodology

To address limitations in the existing CMLRs, we borrow an idea in the recent multiclassification studies (Zhang and Liu, 2014; Zhang et al., 2018; Fu et al., 2018), and propose an attractive simplex-based MLR model. A well-designed simplex in a $(k - 1)$ -dimensional Euclidean space plays a central role in reducing the parameter redundancy. To begin with, consider k vertices $\{\mathbf{w}_j \in \mathbb{R}^{k-1}, j = 1, \dots, k\}$

$$\mathbf{w}_j = \begin{cases} (k - 1)^{-1/2} \mathbf{1}_{k-1}, & \text{if } j = 1 \\ -(1 + k^{1/2})/\{(k - 1)^{3/2}\} \mathbf{1}_{k-1} + \{k/(k - 1)\}^{1/2} \mathbf{e}_{j-1}, & \text{if } 2 \leq j \leq k \end{cases}, \quad (3.5)$$

where $\mathbf{1}$ and $\mathbf{e}_j$'s are vectors in $\mathbb{R}^{k-1}$. One can verify that each $\mathbf{w}_j$ has Euclidean norm 1 and $\mathbf{W} \mathbf{1}_k = \sum_{j=1}^k \mathbf{w}_j = \mathbf{0}$. Denote the matrix of vertex vectors as $\mathbf{W} \triangleq (\mathbf{w}_1, \dots, \mathbf{w}_k) \in \mathbb{R}^{(k-1) \times k}$. Clearly, $\mathbf{W}$ is of full row rank $k - 1$. Note that there exist some alternative constructions of the simplex, such as that of Hill and Doucet (2007); Mroueh et al. (2012). However, we can always connect $\mathbf{W}$ with those simplices via some proper linear transformations in $\mathbb{R}^{k-1}$, and show that they are equivalent.

The simplex $\mathbf{W}$ can be used to reduce the dimension of categorical space to $k - 1$. Consider a k -categorical distribution with $P(Y = y) = p_y > 0$ ($y \in \mathcal{Y}$), with the sum-to-one condition $\sum_{j=1}^k p_j = 1$. The conventional multinomial distribution encodes Y as a one-hot vector in

$\mathbb{R}^k$, i.e., $\mathbf{Y} = \mathbf{e}_Y$. Here, $\mathbf{Y}$ is redundant due to $\mathbf{1}_k^\top \mathbf{Y} \equiv 1$. Let $\mathbf{p} = (p_1, \dots, p_k)^\top$ be the probability vector, then the covariance matrix of $\mathbf{Y}$ is $\mathbf{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top$ (Forbes et al., 2011). Without loss of generality, we can encode category j to the j -th vertex $\mathbf{w}_j$ to obtain another multinomial random vector $\mathbf{Z} = \mathbf{W}\mathbf{Y} \in \mathbb{R}^{k-1}$ with $P(\mathbf{Z} = \mathbf{w}_j) = P(\mathbf{Y} = \mathbf{e}_j) = p_j$. The following Proposition states a useful result on the covariance matrix of $\mathbf{Z}$.

Proposition 1. *The covariance matrix of $\mathbf{Z} = \mathbf{W}\mathbf{Y}$ is positive definite.*

Note that the original $\mathbf{Y}$ is redundant with the covariance matrix $\mathbf{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top \succeq \mathbf{0}$. On the other hand, by Proposition 1, the refined category vector $\mathbf{Z} = \mathbf{W}\mathbf{Y}$ has the covariance matrix $\mathbf{W}(\mathbf{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top)\mathbf{W}^\top \succ \mathbf{0}$. This result indicates that the simplex $\mathbf{W}$ in $\mathbb{R}^{k-1}$ leads to a refined categorical space without redundancy. We can use the vertex $\mathbf{w}_j$ to represent the j -th category, and such a strategy is called simplex coding.

Next, we use $\mathbf{W}$ to remove the explicit sum-to-zero constraint in CMLR2. For the coefficients $\boldsymbol{\theta} \in \mathbb{R}^{d \times k}$ with $\sum_{j=1}^k \boldsymbol{\theta}_j = \mathbf{0}$, we can find a matrix $\boldsymbol{\beta} \in \mathbb{R}^{d \times (k-1)}$ such that $\boldsymbol{\theta}_j = \boldsymbol{\beta}\mathbf{w}_j$ for $j = 1, \dots, k$. It is obvious that $\sum_{j=1}^k \boldsymbol{\theta}_j = \sum_{j=1}^k \boldsymbol{\beta}\mathbf{w}_j = \boldsymbol{\beta}(\sum_{j=1}^k \mathbf{w}_j) = \mathbf{0}$ using the fact that $\sum_{j=1}^k \mathbf{w}_j = \mathbf{0}$. For example, a possible choice of $\boldsymbol{\beta}$ is $(1 - \frac{1}{k})\boldsymbol{\theta}\mathbf{W}^\top$. The equivalence between $\boldsymbol{\theta}$ and $\boldsymbol{\beta}\mathbf{W}$ is shown in the following Proposition.

Proposition 2. $\{\boldsymbol{\theta} \in \mathbb{R}^{d \times k} \mid \sum_{j=1}^k \boldsymbol{\theta}_j = \mathbf{0}\}$ is equivalent to $\{\boldsymbol{\beta}\mathbf{W} \mid \boldsymbol{\beta} \in \mathbb{R}^{d \times (k-1)}\}$.

By Proposition 2, the parameters $\boldsymbol{\theta}$ with a sum-to-zero constraint can be reformulated as its equivalent counterpart $\boldsymbol{\beta}\mathbf{W}$ without loss of information. By replacing $\boldsymbol{\theta}_j$'s in the classical MLR (2.1) with $\boldsymbol{\beta}\mathbf{w}_j$'s, the simplex-based MLR model takes the form

$$\tilde{\pi}_j(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \boldsymbol{\beta}\mathbf{w}_j}}{\sum_{s=1}^k e^{\mathbf{x}^\top \boldsymbol{\beta}\mathbf{w}_s}}, \quad \mathbf{x} \in \mathcal{X}, \quad j = 1, \dots, k. \quad (3.6)$$

We can interpret the quantity $\mathbf{x}^\top \beta \mathbf{w}_j$ as the inner product $\langle \beta^\top \mathbf{x}, \mathbf{w}_j \rangle$ in $\mathbb{R}^{k-1}$. In other words, a covariate $\mathbf{x}$ is first mapped to a point $\beta^\top \mathbf{x}$ in $\mathbb{R}^{k-1}$ by the coefficient matrix β , and then taken the inner product with the encoded vertex $\mathbf{w}_j$. Since $\tilde{\pi}_j(\mathbf{x})$ is increasing in $\mathbf{x}^\top \beta \mathbf{w}_j$, the predicted rule is $\hat{y}(\mathbf{x}) = \arg \max_j \tilde{\pi}_j(\mathbf{x}) = \arg \max_j \{\mathbf{x}^\top \beta \mathbf{w}_j\}$, which is computationally more efficient than that of the reference-based MLR. In particular, SMLR simply requires computing the k individual inner products, while the reference-based MLR involves the calculation of k probabilities via (2.1). As shown in Zhang and Liu (2014), the largest inner product rule is equivalent to the least angle rule, and hence the SMLR may be treated as an alternative multicategory angle-based classifier.

The proposed SMLR model enjoys parsimonious model specification without constraints. Each of the CMLRs involves dk parameters under a linear constraint, while SMLR only requires $d(k-1)$ parameters. Note that $\beta \mathbf{W}$ plays the same role as θ in the existing MLRs. Under the factorization $\beta \mathbf{W}$, the $k-1$ rows of $\mathbf{W}$ can be viewed as latent outcome variables, and each row has a loading on each of the k categories. The $k-1$ columns of β specify parameter vectors for these latent outcome variables.

To further interpret the parameters of SMLR and investigate their relationships with the CMLRs, we look into the log-odds forms. For CMLR1 model with a reference category r (i.e., $\theta_r \equiv \mathbf{0}$), we have

$$\log \left(\frac{\pi_j(\mathbf{x})}{\pi_r(\mathbf{x})} \right) = \mathbf{x}^\top \theta_j, \quad j = 1, \dots, k.$$

Clearly, θ_j depends on the the selection of the reference category and different selections yield different interpretations of the parameters of CMLR1. On the other hand, the parameter interpretation for CMLR2 is related to the median response (Tutz, 2011), which is defined

as the geometric mean

$$\text{GM}(\mathbf{x}) = \prod_{j=1}^k \{\pi_j(\mathbf{x})\}^{1/k}.$$

Then for CMLR2 with the constraint $\sum_{j=1}^k \boldsymbol{\theta}_j = \mathbf{0}$, we have

$$\log \left(\frac{\pi_j(\mathbf{x})}{\text{GM}(\mathbf{x})} \right) = \mathbf{x}^\top \boldsymbol{\theta}_j, \quad j = 1, \dots, k.$$

Therefore, $\boldsymbol{\theta}_j$ reflects the effects of $\mathbf{x}$ on the comparison of $Y = j$ to the geometric mean response $\text{GM}(\mathbf{x})$. Similarly, we can define the simplex-based geometric mean

$$\text{SGM}(\mathbf{x}) = \prod_{j=1}^k \{\tilde{\pi}_j(\mathbf{x})\}^{1/k}.$$

After some algebra, we can show that

$$\log \left(\frac{\tilde{\pi}_j(\mathbf{x})}{\text{SGM}(\mathbf{x})} \right) = \mathbf{x}^\top \boldsymbol{\beta} \mathbf{w}_j, \quad j = 1, \dots, k.$$

Thus, SMLR is a generalization of CMLR2 under the simplex-based coding scheme. In a nutshell, the modeling of CMLR1 is based on the asymmetric comparisons of the categories, while CMLR2 and SMLR use the symmetric comparisons. However, since SMLR removes the sum-to-zero constraint in CMLR2 and involves less parameters, SMLR is more desirable and computationally efficient.

Note that Zhang and Liu (2014) proposed a flexible multicategory classification framework under the simplex structure, which considered a general large-margin loss function $\ell(\cdot)$. According to Theorem 3 in Zhang and Liu (2014), if we consider the exponential loss of the form $\ell(z) = e^{-z}$ and $\ell'(z) = -e^{-z}$, the relationship between the conditional class probability and the theoretical minimizer $\mathbf{f}^*$ can be expressed as

$$P_j(\mathbf{x}) = \frac{\ell'(\langle \mathbf{f}^*(\mathbf{x}), \mathbf{w}_j \rangle)^{-1}}{\sum_{i=1}^k \ell'(\langle \mathbf{f}^*(\mathbf{x}), \mathbf{w}_i \rangle)^{-1}} = \frac{\exp(\langle \mathbf{f}^*(\mathbf{x}), \mathbf{w}_j \rangle)}{\sum_{s=1}^k \exp(\langle \mathbf{f}^*(\mathbf{x}), \mathbf{w}_s \rangle)}, \quad \mathbf{x} \in \mathcal{X}, \quad j = 1, \dots, k,$$

which recovers the probabilistic assumption (3.6) of the proposed SMLR. This interesting connection sheds some light on SMLR from the perspective of large-margin classification.

3.2 Maximum Likelihood Estimation

The log-likelihood function for the SMLR model (3.6) based on the dataset $\mathcal{D}$ is

$$\mathcal{L}_n(\boldsymbol{\beta}) = \ell_n(\boldsymbol{\beta}\mathbf{W}) = \sum_{i=1}^n \mathbf{x}_i^\top \boldsymbol{\beta} \mathbf{w}_{y_i} - \sum_{i=1}^n \log \left(\sum_{j=1}^k e^{\mathbf{x}_i^\top \boldsymbol{\beta} \mathbf{w}_j} \right), \quad (3.7)$$

and the corresponding estimator by maximizing likelihood is defined as

$$\hat{\boldsymbol{\beta}} = \arg \max_{\boldsymbol{\beta}} \mathcal{L}_n(\boldsymbol{\beta}). \quad (3.8)$$

Compared with CMLRs, our SMLR solves an unconstrained problem with less parameters.

The Fisher scoring algorithm, which is equivalent to Newton's method for the SMLR model, can be used to solve (3.8). For notational simplicity, let $\tilde{\boldsymbol{\pi}}_i \triangleq \tilde{\boldsymbol{\pi}}(\mathbf{x}_i)$ be the probability vector for the i th observation, where $\tilde{\boldsymbol{\pi}}_i$ depends on $\boldsymbol{\beta}$ as determined by (3.6). Let $\boldsymbol{\Lambda}(\mathbf{u}) = \mathbf{diag}(\mathbf{u}) - \mathbf{u}\mathbf{u}^\top$ be a matrix associated with $\mathbf{u}$. After some tedious computation, the score vector is

$$U_n(\boldsymbol{\beta}) = \frac{\partial \mathcal{L}_n(\boldsymbol{\beta})}{\partial \vec{\boldsymbol{\beta}}} = \sum_{i=1}^n (\mathbf{W} \otimes \mathbf{x}_i)(\mathbf{e}_{y_i} - \tilde{\boldsymbol{\pi}}_i), \quad (3.9)$$

and the negative Hessian matrix is

$$\mathbf{Q}_n(\boldsymbol{\beta}) = -\frac{\partial^2 \mathcal{L}_n(\boldsymbol{\beta})}{\partial \vec{\boldsymbol{\beta}} \partial \vec{\boldsymbol{\beta}}^\top} = \sum_{i=1}^n [\mathbf{W} \boldsymbol{\Lambda}(\tilde{\boldsymbol{\pi}}_i) \mathbf{W}^\top] \otimes (\mathbf{x}_i \mathbf{x}_i^\top). \quad (3.10)$$

Observe that $\mathbf{Q}_n(\boldsymbol{\beta})$ only involves the design matrix $\mathbf{X}$, and does not depend on the response Y . In what follows, we consider the setting of a fixed design. Then, $\mathbf{Q}_n(\boldsymbol{\beta})$ is also known as the Fisher information. The following Proposition presents some basic results for $\mathbf{Q}_n(\boldsymbol{\beta})$.

Proposition 3. *For any design $\mathbf{X}$, $\mathbf{Q}_n(\boldsymbol{\beta}) \succeq \mathbf{0}$ and $\mathcal{L}_n(\boldsymbol{\beta})$ is concave in $\boldsymbol{\beta}$. In addition, $\mathbf{Q}_n(\boldsymbol{\beta})$ is positive definite if and only if the design matrix $\mathbf{X}$ is of full column rank.*

Due to the concavity of $\mathcal{L}_n(\boldsymbol{\beta})$, the MLE $\hat{\boldsymbol{\beta}}$ is a solution to the likelihood equation $U_n(\boldsymbol{\beta}) = \mathbf{0}$, which can be solved by the standard convex programming (Boyd and Vandenberghe, 2004). In addition, the full column rank of the design matrix $\mathbf{X}$ implies that the number of observations n could be as low as d , i.e., $n \geq d$. If $\mathbf{Q}_n(\boldsymbol{\beta}) \succ \mathbf{0}$, then $\mathcal{L}_n(\boldsymbol{\beta})$ is strictly concave. In this case, as long as the estimate $\hat{\boldsymbol{\beta}}$ exists, it must be unique.

The Fisher scoring procedure can be further reformulated as an iteratively reweighted least squares algorithm, and the corresponding updating scheme is

$$\begin{aligned} \vec{\boldsymbol{\beta}}^{\text{new}} &= \vec{\boldsymbol{\beta}}^{\text{old}} + \mathbf{Q}_n^{-1}(\boldsymbol{\beta}^{\text{old}}) U_n(\boldsymbol{\beta}^{\text{old}}) \\ &= \vec{\boldsymbol{\beta}}^{\text{old}} + \left(\sum_{i=1}^n [\mathbf{W} \boldsymbol{\Lambda}(\tilde{\boldsymbol{\pi}}_i^{\text{old}}) \mathbf{W}^\top] \otimes (\mathbf{x}_i \mathbf{x}_i^\top) \right)^{-1} \left(\sum_{i=1}^n (\mathbf{W} \otimes \mathbf{x}_i) (\mathbf{e}_{y_i} - \tilde{\boldsymbol{\pi}}_i^{\text{old}}) \right), \end{aligned}$$

where $\tilde{\boldsymbol{\pi}}_i^{\text{old}}$ has the j -th entry $\tilde{\pi}_{ij}^{\text{old}} = \frac{e^{\mathbf{x}_i^\top \boldsymbol{\beta}^{\text{old}} \mathbf{w}_j}}{\sum_{s=1}^k e^{\mathbf{x}_i^\top \boldsymbol{\beta}^{\text{old}} \mathbf{w}_s}}$ for $j = 1, \dots, k$.

3.3 Connections with Regular CMLRs

As a new member of the MLR family, the SMLR has intimate relation with the existing MLRs. The main results are stated in the following Theorem.

Theorem 1. *For the three estimators (2.3), (2.4) and (3.8) defined for the same observations $\mathcal{D}$, they achieve the same log-likelihood values, i.e., $\ell_n(\hat{\boldsymbol{\theta}}^{\text{rb}}) = \ell_n(\hat{\boldsymbol{\theta}}^{\text{sc}}) = \mathcal{L}_n(\hat{\boldsymbol{\beta}})$. Moreover, if $\hat{\boldsymbol{\beta}}$ is unique, then the regular estimators can be recovered by $\hat{\boldsymbol{\theta}}^{\text{rb}} = \hat{\boldsymbol{\beta}}(\mathbf{W} - \mathbf{w}_r \mathbf{1}_k^\top)$ and $\hat{\boldsymbol{\theta}}^{\text{sc}} = \hat{\boldsymbol{\beta}} \mathbf{W}$.*

Theorem 1 says that, the two CMLR estimators $\hat{\boldsymbol{\theta}}^{\text{rb}}$ and $\hat{\boldsymbol{\theta}}^{\text{sc}}$ can be uniquely determined

by the SMLR estimator $\hat{\beta}$, under some linear transformations. For any observation $\mathbf{x}$, we have the following prediction results

$$\hat{\pi}_j^{\text{rb}}(\mathbf{x}) = \hat{\pi}_j^{\text{sc}}(\mathbf{x}) = \hat{\pi}_j(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \hat{\beta} \mathbf{w}_j}}{\sum_{s=1}^k e^{\mathbf{x}^\top \hat{\beta} \mathbf{w}_s}}.$$

Therefore, these three MLR models are equivalent in terms of probability estimation and label prediction. If an alternative simplex is applied, the estimated coefficients of SMLR may be different, but these three MLR models still share the same prediction outputs.

Theorem 1 also implies the connections between the two CMLRs, i.e., $\hat{\theta}^{\text{rb}} = \hat{\theta}^{\text{sc}}(\mathbf{I}_k - \mathbf{e}_r \mathbf{1}_k^\top)$ and $\hat{\theta}^{\text{sc}} = \hat{\theta}^{\text{rb}}(\mathbf{I}_k - \frac{\mathbf{1}_k \mathbf{1}_k^\top}{k})$, where $\mathbf{I}_k$ is a $k \times k$ identity matrix. To depict the connections of the three MLRs, we consider a corner case where there are two distinct categories, i.e., $k = 2$. Let $P(\mathbf{x}) = P(Y = 1|\mathbf{x})$ be the probability of the first category happens conditioning on the given covariates $\mathbf{x}$. Clearly, $P(Y = 2|\mathbf{x}) = 1 - P(\mathbf{x})$. Then, the CMLR1 model with parameter $(\theta, \mathbf{0})$ is

$$P(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \theta}}{e^{\mathbf{x}^\top \theta} + 1} = \frac{1}{1 + e^{-\mathbf{x}^\top \theta}};$$

where the second category is regarded as a reference. For the CMLR2, due to the sum-to-zero constraint $\theta_1 + \theta_2 = \mathbf{0}$, we can simplify (θ_1, θ_2) as $(\theta_1, -\theta_1)$, and the resulting model is

$$P(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \theta_1}}{e^{\mathbf{x}^\top \theta_1} + e^{\mathbf{x}^\top \theta_2}} = \frac{1}{1 + e^{\mathbf{x}^\top (\theta_2 - \theta_1)}} = \frac{1}{1 + e^{-2\mathbf{x}^\top \theta_1}}.$$

The SMLR model (3.6) with parameter β becomes

$$P(\mathbf{x}) = \frac{e^{\mathbf{x}^\top \beta}}{e^{\mathbf{x}^\top \beta} + e^{-\mathbf{x}^\top \beta}} = \frac{1}{1 + e^{-2\mathbf{x}^\top \beta}}.$$

Denote the three corresponding estimates as $(\hat{\theta}, \mathbf{0})$, $(\hat{\theta}_1, -\hat{\theta}_1)$ and $\hat{\beta} \in \mathbb{R}^d$, respectively. Based on Theorem 1, we have $\hat{\theta} = 2\hat{\beta}$ and $\hat{\theta}_1 = \hat{\beta}$. Then, we conclude the following Corollary.

Corollary 1. *If $k = 2$, the three MLR models are equivalent up to a scaling factor.*

Theorem 1 can also be used to reveal some interesting properties of the regular CMLRs. For example, consider the relation between the CMLR1 estimators under two different reference categories r and s . By Theorem 1, we have $\hat{\boldsymbol{\theta}}^{(r)} = \hat{\boldsymbol{\beta}}(\mathbf{W} - \mathbf{w}_r \mathbf{1}_k^\top)$ and $\hat{\boldsymbol{\theta}}^{(s)} = \hat{\boldsymbol{\beta}}(\mathbf{W} - \mathbf{w}_s \mathbf{1}_k^\top)$. The relation between $\hat{\boldsymbol{\theta}}^{(r)}$ and $\hat{\boldsymbol{\theta}}^{(s)}$ is shown in the following Corollary.

Corollary 2. *For the CMLR1 model with two different reference categories r and s , $\hat{\boldsymbol{\theta}}^{(r)}$ and $\hat{\boldsymbol{\theta}}^{(s)}$ have the following relation*

$$\hat{\boldsymbol{\theta}}^{(r)} = \hat{\boldsymbol{\theta}}^{(s)} - \hat{\boldsymbol{\theta}}_r^{(s)} \mathbf{1}_k^\top.$$

For illustration, consider $k = 2$ and an available estimate $(\hat{\boldsymbol{\theta}}, \mathbf{0})$ for the CMLR1 model with the second reference. By Corollary 2, the estimate for CMLR1 with the first reference is $(\mathbf{0}, -\hat{\boldsymbol{\theta}})$.

According to Theorem 1, it is possible to transfer the properties for regular CMLRs to the SMLR framework. For instance, Albert and Anderson (1984) showed that the MLE for the CMLR1 model exists only when the datasets are overlapped, where the existence means finiteness of an estimate. We can extend the concept of overlapping for SMLR model as follows.

Definition 1 (Overlapping). We say that the observations $\mathcal{D}$ are overlapping, if for every nonzero matrix $\boldsymbol{\beta} \in \mathbb{R}^{d \times (k-1)}$, there exists a duplet (i, t) with $i \in \{1, \dots, n\}$ and $t \in \mathcal{Y} \setminus y_i$ such that $\mathbf{x}_i^\top \boldsymbol{\beta}(\mathbf{w}_{y_i} - \mathbf{w}_t) < 0$.

Applying Theorem 1, we have the following Corollary for the existence of the SMLR estimator.

Corollary 3. *The MLE for the SMLR model exists if and only if the observations $\mathcal{D}$ are overlapped.*

Although the three MLRs have intimate connections, the SMLR model treats all categories equally, and provides systematic insights in a concise framework without further constraints. Hence, the SMLR model serves as a unified framework which covers the binomial LR and two regular CMLRs. In what follows, we focus on the SMLR model for succinct technical presentation. All conclusions for the SMLR model can be adapted to the regular CMLRs through the link functions in Theorem 1.

4. Theoretical Properties

In this section, we establish the asymptotical results and the statistical learning theory for the proposed SMLR model, including the existence and uniqueness of MLE, Fisher consistency and the comparison inequalities as a classifier and the convergence results for the kernel SMLR.

4.1 Asymptotical Results

We are interested in the asymptotic behavior of the SMLR model under a complicated diverging setting, where both the number of covariates and the number of categories can increase with the sample size. In the literature, the diverging number of covariates has been well studied while the diverging number of categories is rarely investigated, although such a setting is not uncommon in practice. In the general “large n , diverging d and diverging k ” setup, we denote $d = d_n$ and $k = k_n$ to emphasize the impact of sample size n . Then, the number of

parameters for the SMLR is $(k_n - 1)d_n$. Hereinafter, we replace $(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$ by $(\boldsymbol{\beta}_n, \hat{\boldsymbol{\beta}}_n)$ to stress their dependencies on the sample size n .

For the matrix $\boldsymbol{\Lambda}(\tilde{\boldsymbol{\pi}}_i)$ in $\mathbf{Q}_n(\boldsymbol{\beta}_n)$, since $\tilde{\boldsymbol{\pi}}_i$ depends on $\boldsymbol{\beta}_n$, we rewrite $\boldsymbol{\Lambda}_i(\boldsymbol{\beta}_n) = \boldsymbol{\Lambda}(\tilde{\boldsymbol{\pi}}_i)$. Hence, the Fisher information matrix $\mathbf{Q}_n(\boldsymbol{\beta}_n) = \sum_{i=1}^n [\mathbf{W}\boldsymbol{\Lambda}_i(\boldsymbol{\beta}_n)\mathbf{W}^\top] \otimes (\mathbf{x}_i\mathbf{x}_i^\top)$ is a function of $\boldsymbol{\beta}_n$. Assume that the true coefficient matrix is $\boldsymbol{\beta}_{n0}$. Denote $\mathbf{G}_{n0} = \mathbf{Q}_n(\boldsymbol{\beta}_{n0})$. When the underlying model is correctly specified,

$$P(Y = j|X = \mathbf{x}) = \frac{e^{\mathbf{x}^\top \boldsymbol{\beta}_{n0} \mathbf{w}_j}}{\sum_{s=1}^k e^{\mathbf{x}^\top \boldsymbol{\beta}_{n0} \mathbf{w}_s}}.$$

For the score vector $U_n(\cdot)$ as defined in (3.9), it follows that $E[U_n(\boldsymbol{\beta}_{n0})] = \mathbf{0}$. Define $\mathbf{S}_n = \mathbf{X}^\top \mathbf{X} = \sum_{i=1}^n \mathbf{x}_i\mathbf{x}_i^\top$. Since the covariates $\mathbf{x}_i$'s are assumed deterministic, $\mathbf{Q}_n(\boldsymbol{\beta}_n)$ and $\mathbf{S}_n$ are not random. Given a fixed $\delta > 0$, we consider a region of interest for $\boldsymbol{\beta}_n$, i.e., $N_n(\delta) = \{\boldsymbol{\beta} : \|\mathbf{G}_{n0}^{1/2}(\vec{\boldsymbol{\beta}} - \vec{\boldsymbol{\beta}}_{n0})\| \leq \delta\{d_n(k_n - 1)\}^{1/2}\}$.

In order to obtain the asymptotic results, we need the following assumptions.

Assumption 1. The true parameter $\boldsymbol{\beta}_{n0}$ is contained in the interior of a compact subset $\mathcal{B}_n$ in $\mathbb{R}^{d_n \times (k_n - 1)}$.

Assumption 2. Each covariate of X is uniformly bounded by a constant $C > 0$.

Assumption 3. There exist two positive constants c_1 and c_2 such that $c_1 \leq \lambda_{\min}(\mathbf{S}_n/n) \leq \lambda_{\max}(\mathbf{S}_n/n) \leq c_2$.

Assumption 4. There exists a constant $L_0 > 0$ such that $\|\boldsymbol{\Lambda}(\boldsymbol{\beta}) - \boldsymbol{\Lambda}(\boldsymbol{\beta}')\|_2 \leq L_0 \|\boldsymbol{\beta} - \boldsymbol{\beta}'\|$ for any $\boldsymbol{\beta}, \boldsymbol{\beta}' \in \mathcal{B}_n$.

The above assumptions are motivated from the literature of M -estimation with diverging dimension of covariates, see Portnoy (1985, 1988); He and Shao (2000); Wang (2011); Liang

and Du (2012); Gao et al. (2018). Assumption 1 makes a restriction on the true parameter β_{n0} under the doubly diverging setting. Assumption 2 ensures that each observed covariate vector is bounded, that is, $\|\mathbf{x}_i\| \leq C\sqrt{d_n}$ for any $i = 1, \dots, n$. Assumption 3 gives some regular conditions on the design matrix $\mathbf{X}$. Assumption 4 is specially tailored for the SMLR model, which can be viewed as a generalized Lipschitz condition.

Remark 1. Assumption 3 is also used by Wang (2011); Liang and Du (2012) to establish asymptotics for binomial LR with diverging covariates. Based on Assumption 3, we can find a positive constant c_0 such that $c_0 \leq \max_{i=1, \dots, n} \min_{j=1, \dots, k} \tilde{\pi}_{ij}(\beta_{n0})$. Furthermore, using Lemma 4 in the Supplementary Material, for Fisher information matrix $\mathbf{G}_{n0}$, we have

$$\frac{\mathbf{G}_{n0}}{n} = \frac{1}{n} \sum_{i=1}^n [\mathbf{W} \Lambda_i(\beta_n) \mathbf{W}^\top] \otimes (\mathbf{x}_i \mathbf{x}_i^\top) \succeq \frac{k_n c_0}{k_n - 1} \frac{\mathbf{S}_n}{n} \succeq c \mathbf{I}_{(k_n-1)d_n}, \quad (4.11)$$

with a constant $c \in (0, c_0 c_1]$. If $k = 2$, the result (4.11) is reduced to that of Liang and Du (2012). In summary, there exists a positive constant c such that $\lambda_{\min}(\mathbf{G}_{n0}/n) \geq c$, and (4.11) is required in establishing the convergence and asymptotic normality of estimated coefficients under certain diverging settings.

Remark 2. Under Assumption 2, we have $\|\mathbf{S}_n/n\|_2 \leq n^{-1} \sum_{i=1}^n \mathbf{x}_i^\top \mathbf{x}_i \leq d_n C^2$. Therefore, Assumption 3 naturally holds in the “fixed d ” case.

Remark 3. Due to the complicated nature of MLR, Assumption 4 is necessary to control the behavior of different coefficients on the component of Fisher information matrix, which depicts the continuousness of a matrix function in another manner. The technical proofs on consistency and asymptotical normality can be simplified by using Assumption 4.

Next, when the number of parameters $d_n(k_n - 1)$ increases, the existence and consistency of the MLE is guaranteed by the following Theorem.

Theorem 2. *Suppose Assumptions 1–4 hold. If $\sqrt{d_n k_n / n} \rightarrow 0$, then there exists a sequence $\{\hat{\beta}_n\}$ such that as $n \rightarrow \infty$,*

$$P\left(\hat{\beta}_n \in N_n(\delta) \text{ and } U_n(\hat{\beta}_n) = \mathbf{0}\right) \rightarrow 1 \text{ and } \|\hat{\beta}_n - \beta_{n0}\| = O_p(\sqrt{d_n k_n / n}).$$

For the classical setting with fixed d and k , the conditions become $1/n \rightarrow 0$. Consider the setting with fixed dimension d and varying categories k_n , we require that $k_n = o(n)$. When the number of categories k is fixed, d_n should satisfy $d_n = o(n)$.

The following Theorem ensures the asymptotic normality of $\hat{\beta}_n$.

Theorem 3. *Suppose Assumptions 1–4 hold. If $d_n k_n / \sqrt{n} \rightarrow 0$, then for any $d_n(k_n - 1) \times l$ matrix $\mathbf{V}_n$ with l fixed and such that $\mathbf{V}_n^\top \mathbf{V}_n = \mathbf{I}_l$,*

$$\mathbf{V}_n^\top \mathbf{G}_{n0}^{1/2} \left(\vec{\hat{\beta}}_n - \vec{\beta}_{n0} \right) \longrightarrow \mathcal{N}(\mathbf{0}, \mathbf{I}_l), \text{ in distribution,}$$

where $\mathcal{N}(\cdot, \cdot)$ is a multivariate normal distribution.

In particular, let $\mathbf{U}_n = \mathbf{G}_{n0}^{-1/2} \mathbf{V}_n (\mathbf{V}_n^\top \mathbf{G}_{n0}^{-1} \mathbf{V}_n)^{-1/2}$, then $\mathbf{U}_n^\top \mathbf{U}_n = \mathbf{I}_l$. Based on Theorem 3, we have the following Corollary which gives the asymptotic distribution of $\mathbf{V}_n^\top (\vec{\hat{\beta}}_n - \vec{\beta}_{n0})$.

Corollary 4. *Under the same conditions as in Theorem 3, as $n \rightarrow \infty$, we have*

$$(\mathbf{V}_n^\top \mathbf{G}_{n0}^{-1} \mathbf{V}_n)^{-1/2} \mathbf{V}_n^\top \left(\vec{\hat{\beta}}_n - \vec{\beta}_{n0} \right) \longrightarrow \mathcal{N}(\mathbf{0}, \mathbf{I}_l), \text{ in distribution.}$$

If d and k are fixed, we only need $n^{-1/2} \rightarrow 0$, which is trivially true. When the number of categories k_n diverges and d is fixed, the conditions in Theorem 3 reduce into $k_n / \sqrt{n} \rightarrow 0$,

which implies a sufficient condition $k_n = o(\sqrt{n})$. Consider the diverging number of covariates d_n and fixed k , we require $d_n/\sqrt{n} \rightarrow 0$, i.e., $d_n = o(\sqrt{n})$. In particular, for the binomial LR model, we recover the result in Portnoy (1985, 1988), and our result is stronger than that in He and Shao (2000); Wang (2011). Hence, our conditions in Theorem 2 and 3 are general for the MLR models with a diverging number of parameters, even with a diverging number of categories.

The following Theorem and its Corollary suggest that one can approximate $\mathbf{G}_{n0}$ via $\mathbf{Q}_n(\hat{\beta}_n)$, when applying Theorem 3 and Corollary 4 for interval estimation.

Theorem 4. *Under the same conditions as in Theorem 3, as $n \rightarrow \infty$, we have*

$$\|\mathbf{V}_n^\top \mathbf{Q}_n^{-1}(\hat{\beta}_n) \mathbf{V}_n - \mathbf{V}_n^\top \mathbf{G}_{n0}^{-1} \mathbf{V}_n\|_2 \rightarrow 0, \text{ in probability.}$$

Corollary 5. *Under the same conditions as in Theorem 3, as $n \rightarrow \infty$, we have*

$$[\mathbf{V}_n^\top \mathbf{Q}_n^{-1}(\hat{\beta}_n) \mathbf{V}_n]^{-\frac{1}{2}} \mathbf{V}_n^\top (\vec{\beta}_n - \vec{\beta}_{n0}) \rightarrow \mathcal{N}(\mathbf{0}, \mathbf{I}_l), \text{ in distribution,}$$

and

$$(\vec{\beta}_n - \vec{\beta}_{n0})^\top \mathbf{V}_n [\mathbf{V}_n^\top \mathbf{Q}_n^{-1}(\hat{\beta}_n) \mathbf{V}_n]^{-1} \mathbf{V}_n^\top (\vec{\beta}_n - \vec{\beta}_{n0}) \rightarrow \chi_l^2, \text{ in distribution,}$$

where χ_l^2 is the χ^2 distribution with l degrees of freedom.

According to Corollary 5 and the MLE, an asymptotic $1 - \alpha$ confidence interval ($0 < \alpha < 1$) for β_{ij} is

$$\hat{\beta}_{ij} \pm z_{\alpha/2} \mathbf{v}_{ij}^\top \mathbf{Q}_n^{-1} \mathbf{v}_{ij},$$

where $z_{\alpha/2}$ denotes the upper $(\alpha/2)$ -quantile of the standard normal distribution and $\mathbf{v}_{ij}$ is the unit vector of length $d_n(k_n - 1)$ with the $[(j - 1)d_n + i]$ -th element equal to 1 and all the

other elements equal to 0. We can also apply these results in testing for the following linear hypothesis:

$$H_0 : \mathbf{V}_n^\top \overrightarrow{\beta_{n0}} = \mathbf{a} \longleftrightarrow H_1 : \mathbf{V}_n^\top \overrightarrow{\beta_{n0}} \neq \mathbf{a},$$

where the vector $\mathbf{a} \in \mathbb{R}^l$ is known and $\mathbf{V}_n$ is a $\{(k_n - 1)d_n\} \times l$ matrix such that $\mathbf{V}_n^\top \mathbf{V}_n = \mathbf{I}_l$.

The large sample Wald test statistic is defined as

$$T_n = \left(\overrightarrow{\hat{\beta}_n}^\top \mathbf{V}_n - \mathbf{a}^\top \right) \left[\mathbf{V}_n^\top \mathbf{Q}_n^{-1}(\hat{\beta}_n) \mathbf{V}_n \right]^{-1} \left(\mathbf{V}_n^\top \overrightarrow{\hat{\beta}_n} - \mathbf{a} \right).$$

Corollary 5 shows that the Wald test remains valid, i.e., $T_n \longrightarrow \chi_l^2$ in distribution under the null hypothesis H_0 , even when the numbers of covariates and categories diverge with the sample size.

4.2 Fisher Consistency and Error Analysis

Fisher consistency is a fundamental property for various classifiers, which is also called infinite-sample consistency by Zhang (2004) and classification calibration by Bartlett et al. (2006); Tewari and Bartlett (2007). In this section, we aim to show these desired properties for the SMLR model.

First, let $P_j(\mathbf{x}) = P(Y = j | X = \mathbf{x})$ be the underlying class conditional probability for any $\mathbf{x} \in \mathcal{X}$, and we define a vector $\mathbf{p}(\mathbf{x}) = (P_1(\mathbf{x}), \dots, P_k(\mathbf{x}))^\top$. Consider a classifier $\mathcal{C} : \mathcal{X} \mapsto \mathcal{Y}$, and the expected misclassification error is given by

$$\mathcal{R}(\mathcal{C}) = E[\mathbb{1}(Y \neq \mathcal{C}(X))] = 1 - E_X[P(Y = \mathcal{C}(\mathbf{x}) | X = \mathbf{x})].$$

One can verify that the optimal classifier minimizing $\mathcal{R}(\mathcal{C})$, often called the Bayes rule, is denoted as $\mathcal{C}_B(\mathbf{x}) = \arg \max_j P_j(\mathbf{x})$. Denote $\mathcal{R}^* = \mathcal{R}(\mathcal{C}_B) = 1 - E_X[\max_j P_j(X)]$.

Under the simplex coding scheme, it is sufficient to directly use $k - 1$ functions for multicategory classification. Let $\mathbf{f} = (f_1, \dots, f_{k-1})^\top : X \mapsto \mathbb{R}^{k-1}$ be a generic classification function, and then the prediction rule induced by $\mathbf{f}$ is $\mathcal{C}_{\mathbf{f}}(\mathbf{x}) = \arg \max_j \langle \mathbf{f}(\mathbf{x}), \mathbf{w}_j \rangle$. Consider the SMLR log-likelihood (3.7), we can define the SMLR loss function for an observation $(\mathbf{x}, y)$,

$$V(\mathbf{f}(\mathbf{x}), y) = \log \left(\sum_{j=1}^k e^{\langle \mathbf{f}(\mathbf{x}), \mathbf{w}_j \rangle} \right) - \langle \mathbf{f}(\mathbf{x}), \mathbf{w}_y \rangle. \quad (4.12)$$

We are interested in the expected V -risk

$$\mathcal{E}(\mathbf{f}) = E[V(\mathbf{f}(X), Y)] = E_X\{E[V(\mathbf{f}(X), Y)|X]\}.$$

Consider the hypothesis space

$$\mathfrak{F} = \{\mathbf{f} : \mathcal{X} \mapsto \mathbb{R}^{k-1} \mid E_X[\|\mathbf{f}(X)\|] < \infty\},$$

where $\|\cdot\|$ is the standard Euclidean norm in $\mathbb{R}^{k-1}$. Note that $\mathcal{E}(\mathbf{f})$ is a functional of $\mathbf{f}$, and we define its minimizer over $\mathfrak{F}$, $\mathbf{f}^* = \arg \min_{\mathbf{f} \in \mathfrak{F}} \mathcal{E}(\mathbf{f})$. Fisher consistency requires that $\mathcal{C}_B(\mathbf{x}) = \mathcal{C}_{\mathbf{f}^*}(\mathbf{x})$ for any $\mathbf{x} \in \mathcal{X}$. We reveal the Fisher consistency for SMLR in the following Theorem.

Theorem 5. *Assume that $P_j(\mathbf{x}) > 0$ for $j = 1, \dots, k$. The expected V -risk $\mathcal{E} : \mathfrak{F} \mapsto \mathbb{R}_+$ is a convex and continuous functional, with the minimizer $\mathbf{f}^*(\mathbf{x}) = (1 - 1/k) \sum_{i=1}^k [\log P_i(\mathbf{x})] \mathbf{w}_i$. Moreover, the SMLR loss function (4.12) is Fisher consistent.*

Theorem 5 states a one-to-one correspondence between $\mathbf{f}^*$ and $\mathbf{p}$, where the argument $\mathbf{x}$ is suppressed for brevity. In specific, the explicit form of $\mathbf{f}^*$ is a linear expression of $\mathbf{w}_i$ s with coefficients uniquely determined by P_j 's. Since $\langle \mathbf{f}^*(\mathbf{x}), \mathbf{w}_j \rangle = \log P_j(\mathbf{x}) - \frac{1}{k} \sum_{i=1}^k \log P_i(\mathbf{x})$

for $j = 1, \dots, k$, we conclude that if $P_i > P_j$, then $\langle \mathbf{f}^*, \mathbf{w}_i \rangle > \langle \mathbf{f}^*, \mathbf{w}_j \rangle$, and if $P_i = P_j$, then $\langle \mathbf{f}^*, \mathbf{w}_i \rangle = \langle \mathbf{f}^*, \mathbf{w}_j \rangle$. However, if $P_j \rightarrow 0$, $\mathbf{f}^*$ becomes unbounded and meaningless.

For a classification function $\mathbf{f}$, $\mathcal{R}(\mathcal{C}_{\mathbf{f}}) - \mathcal{R}^*$ and $\mathcal{E}(\mathbf{f}) - \mathcal{E}(\mathbf{f}^*)$ are called the excess misclassification risk and the excess V -risk in SMLR, respectively. Then, the following theorem provides an essential comparison inequality.

Theorem 6. *For any $\mathbf{f} \in \mathfrak{F}$, we have $\mathcal{R}(\mathcal{C}_{\mathbf{f}}) - \mathcal{R}^* \leq \sqrt{2}\{\mathcal{E}(\mathbf{f}) - \mathcal{E}(\mathbf{f}^*)\}^{1/2}$.*

Theorem 6 covers the results for binary LR in Bartlett et al. (2006). The upper bound can be improved under some regularity conditions. To this end, we introduce the generalized Tsybakov's low-noise assumption (Tsybakov, 2004) as follows.

Assumption 5. Let $P_{(1)}(\mathbf{x})$ and $P_{(2)}(\mathbf{x})$ be the largest conditional probability and the second largest conditional probability, respectively. Assume that there exist $C > 0$ and $\alpha \geq 0$ such that for all $0 < h \leq 1$,

$$P_X(\{\mathbf{x} \in \mathcal{X} | P_{(1)}(\mathbf{x}) - P_{(2)}(\mathbf{x}) \leq h\}) \leq Ch^\alpha. \quad (4.13)$$

Intuitively, it is clear that misclassification error is particularly large when it is difficult to separate the class with the highest probability from others. In many multicategory classification problems, it is reasonable to assume that the $P_{(1)}(\mathbf{x})$ is unlikely to be very close to $P_{(2)}(\mathbf{x})$ for $\mathbf{x} \in \mathcal{X}$. Hence, Assumption 5 is a meaningful low-noise condition, which depends on the parameter α . We consider two extreme values of α . If $\alpha = 0$, this imposes the case without any assumption on the noise, discussed in Theorem 6. If $\alpha = \infty$, it reduces to the noiseless case.

We can improve the bound for the excess misclassification risk under Assumption 5. Note that similar results have been established in Theorem 2 of Mroueh et al. (2012) for

multicategory support vector machines equipped with hinge loss or quadratic loss. The Theorem 7 below further establishes the results for the MLR models.

Theorem 7. *For each $\mathbf{f} \in \mathfrak{F}$, if Assumption 5 holds, then we have*

$$\mathcal{R}(\mathcal{C}_{\mathbf{f}}) - \mathcal{R}^* \leq (8C_{\alpha})^{\frac{\alpha+1}{\alpha+2}} \{\mathcal{E}(\mathbf{f}) - \mathcal{E}(\mathbf{f}^*)\}^{\frac{\alpha+1}{\alpha+2}},$$

where $C_{\alpha} = (\alpha + 1)C^{\frac{1}{\alpha+1}}\alpha^{-\frac{\alpha}{\alpha+1}} > 0$ is a constant.

Remarkably, Theorem 6 is a particular case of Theorem 7 with $\alpha = 0$. Furthermore, Theorem 7 is a refined version of Theorem 6, since $\frac{\alpha+1}{\alpha+2} > 1/2$ when $\alpha > 0$.

4.3 Convergence Analysis of Kernel SMLR

Motivated by the kernel MLR in Zhu and Hastie (2005), we investigate the consistency of kernel SMLR and conduct the convergence analysis under a diverging number of categories, which is rarely studied in the literature.

To start with, we consider a Mercer kernel K defined over $\mathcal{X} \times \mathcal{X}$ and the reproducing kernel Hilbert space H_K induced by K with the inner product $\langle \cdot, \cdot \rangle$, stratifying $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle = K(\mathbf{x}, \mathbf{x}')$ with a feature map $\phi : \mathcal{X} \mapsto H_K$. Then, we introduce the following notations:

$$\begin{aligned} \mathcal{F} &= \left\{ \mathbf{f} : \mathcal{X} \mapsto \mathbb{R}^{k-1}, f_j(\mathbf{x}) = \langle \mathbf{u}_j, \phi(\mathbf{x}) \rangle \mid E_X[\|\mathbf{f}(X)\|] < \infty, \mathbf{u}_j \in H_K, \forall j = 1, \dots, k-1 \right\}, \\ \mathcal{F}_A &= \left\{ \mathbf{f} \in \mathcal{F} \mid \left(\sum_{j=1}^{k-1} \langle \mathbf{u}_j, \mathbf{u}_j \rangle \right)^{1/2} \leq A \right\} \subseteq \mathcal{F}, \end{aligned}$$

where $A > 0$ is a constant to bound the hypothesis in $\mathcal{F}$. Note that if $\phi(\mathbf{x}) = \mathbf{x}$, the classical setting of linear learning is recovered. We also need some technical assumptions about the boundedness of the kernel function.

Assumption 6. There exists a constant $C > 0$ such that $\sqrt{K(\mathbf{x}, \mathbf{x})} \leq C$ for any $\mathbf{x} \in \mathcal{X}$.

Under Assumption 6, we know $E_X[\|\mathbf{f}(X)\|] \leq AC$ for any $\mathbf{f} \in \mathcal{F}_A$, which ensures $\mathcal{F}_A \subseteq \mathcal{F}$.

Theoretically, the expected risk minimizer for the kernel SMLR model is denoted as

$$\mathbf{f}^* = \arg \min_{\mathbf{f} \in \mathcal{F}} \mathcal{E}(\mathbf{f}) = E[V(\mathbf{f}(X), Y)].$$

To ensure the uniqueness of $\mathbf{f}^*$, we further define $\mathbf{f}^* = \arg \min_{\mathbf{f} \in \mathcal{B}} (\sum_{j=1}^{k-1} \langle \mathbf{u}_j, \mathbf{u}_j \rangle)^{1/2}$, where $\mathcal{B} = \{\mathbf{f} \in \mathcal{F} | \mathbf{f} = \arg \min_{\mathbf{f} \in \mathcal{F}} \mathcal{E}(\mathbf{f})\}$. In fact, our aim is to learn the empirical risk minimizer from $\mathcal{F}_A$, defined by

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f} \in \mathcal{F}_A} \hat{\mathcal{E}}_{\mathcal{D}}(\mathbf{f}) = \frac{1}{n} \sum_{i=1}^n V(\mathbf{f}(\mathbf{x}_i), y_i).$$

Proposition 4. Suppose Assumption 6 is met. For the function class $\mathcal{F}_A$, we have

$$E_{\mathcal{D}} \left[\sup_{\mathbf{f} \in \mathcal{F}_A} \left| \frac{1}{n} \sum_{i=1}^n V(\mathbf{f}(\mathbf{x}_i), y_i) - \mathcal{E}(\mathbf{f}) \right| \right] \leq 4\sqrt{2}AC(k-1)^{1/2}n^{-1/2}.$$

The concentration inequality in Proposition 4 involves the expectation of the supremum of the empirical process, and it is a main tool to establish the consistency of $\hat{\mathbf{f}}$ under the diverging number of categories, as shown in the following Theorem.

Theorem 8. Assume Assumption 6 holds and there exists a proper A such that the theoretical minimizer $\mathbf{f}^* \in \mathcal{F}_A$. If $\sqrt{(k_n - 1)/n} \rightarrow 0$, then we have

$$\lim_{n \rightarrow \infty} E_{\mathcal{D}}[\mathcal{E}(\hat{\mathbf{f}})] = \min_{\mathbf{f} \in \mathcal{F}_A} \mathcal{E}(\mathbf{f}) = \mathcal{E}(\mathbf{f}^*).$$

Theorem 8 implies that the consistency of $\hat{\mathbf{f}}$ requires that k should increase at the order $o(n)$. In particular, this order is identical to that of Theorem 2 for the linear SMLR model. The main difference is that Theorem 2 is established for the MLE of the linear SMLR, while

Theorem 8 is based on the negative log-likelihood loss function for the kernel SMLR. For the kernel CMLR2 with fixed k , Zhang (2004) showed a convergence result under the condition $\sqrt{k/n} \ln^{3/2} n \rightarrow 0$. On the other hand, our Theorem 8 requires $\sqrt{(k-1)/n} \rightarrow 0$, which is faster than that of Zhang (2004).

5. Numerical Studies

In this section, we conduct several experiments to demonstrate the numerical performance of the proposed SMLR model. In particular, we study three simulated examples in Section 5.1, and consider two real world applications in Section 5.2 and 5.3, respectively.

5.1 Simulated Examples

Consider the following SMLR model

$$\tilde{\pi}_y(\mathbf{x}_i) = \frac{e^{\mathbf{x}_i^\top \boldsymbol{\beta}_{n0} \mathbf{w}_y}}{\sum_{j=1}^k e^{\mathbf{x}_i^\top \boldsymbol{\beta}_{n0} \mathbf{w}_j}}, \quad i = 1, \dots, n; \quad y \in \mathcal{Y}, \quad (5.14)$$

where $\boldsymbol{\beta}_{n0}$ is a $d_n \times (k_n - 1)$ matrix of parameters. In specific, $\mathbf{X}_i = (X_{i1}, \dots, X_{id_n})^\top$, $i = 1, \dots, n$ are independently generated from a multivariate normal distribution with mean zero and marginal standard deviation 0.2. In the simulation, we concentrate on the model (5.14) with diverging number of parameters, and the explicit settings are shown as follows. Let $\lfloor a \rfloor$ be the largest integer not greater than the number a .

Example 1 (diverging k and fixed d). Consider $k_n = \lfloor \sqrt{n} \rfloor$ and $d = 3$. Let $\mathbf{a}_n = \{(k_n - 1)d\}^{-1/2} \mathbf{1}_{(k_n-1)d}$. The true parameter matrix is

$$\boldsymbol{\beta}_{n0} = \begin{pmatrix} \frac{k_n-1}{k_n} \cdot \mathbf{1}_d & -\frac{k_n-2}{k_n} \cdot \mathbf{1}_d & \dots & (-1)^{k_n-1} \frac{2}{k_n} \cdot \mathbf{1}_d & (-1)^{k_n} \frac{1}{k_n} \cdot \mathbf{1}_d \end{pmatrix}.$$

Example 2 (diverging d and fixed k). Consider $d_n = \lfloor 2\sqrt{n} \rfloor$ and $k = 3$. Let $d_0 = \lfloor d_n/4 \rfloor$ and $\mathbf{V}_n = \begin{pmatrix} \mathbf{v}_{n1} & \mathbf{v}_{n2} \end{pmatrix} = d_n^{-1/2} \mathbf{I}_2 \otimes \mathbf{1}_{d_n}$. The true parameter matrix is

$$\beta_{n0} = \begin{pmatrix} 0.4 \cdot \mathbf{1}_{d_0} & -0.1 \cdot \mathbf{1}_{d_0} \\ -0.3 \cdot \mathbf{1}_{d_0} & 0.2 \cdot \mathbf{1}_{d_0} \\ 0.2 \cdot \mathbf{1}_{d_0} & -0.3 \cdot \mathbf{1}_{d_0} \\ -0.1 \cdot \mathbf{1}_{d_n-3d_0} & 0.4 \cdot \mathbf{1}_{d_n-3d_0} \end{pmatrix}.$$

Example 3 (diverging k and d simultaneously). Consider $k_n = \lfloor 3n^{1/4} \rfloor$ and $d_n = \lfloor 2n^{1/4} \rfloor$. Let $d_0 = \lfloor d_n/2 \rfloor$ and $\mathbf{b}_n = \{(k_n - 1)d_n\}^{-1/2} \mathbf{1}_{(k_n-1)d_n}$. The true parameter matrix is

$$\beta_{n0} = \begin{pmatrix} \frac{k_n-1}{k_n} \cdot \mathbf{1}_{d_0} & -\frac{k_n-2}{k_n} \cdot \mathbf{1}_{d_0} & \cdots & (-1)^{k_n-1} \frac{2}{k_n} \cdot \mathbf{1}_{d_0} & (-1)^{k_n} \frac{1}{k_n} \cdot \mathbf{1}_{d_0} \\ -\frac{k_n-1}{k_n} \cdot \mathbf{1}_{d_n-d_0} & \frac{k_n-2}{k_n} \cdot \mathbf{1}_{d_n-d_0} & \cdots & (-1)^{k_n} \frac{2}{k_n} \cdot \mathbf{1}_{d_n-d_0} & (-1)^{k_n+1} \frac{1}{k_n} \cdot \mathbf{1}_{d_n-d_0} \end{pmatrix}.$$

The sample size n varies from $\{100, 500, 1000, 5000, 10000\}$ for each example, and we conduct 10000 replications for each simulation. Since the dimension of the estimated coefficient matrix changes as n increases, we measure the accuracy of estimation by the simulated average mean square error (AMSE), which is obtained by averaging $\frac{\|\hat{\beta}_n - \beta_{n0}\|^2}{d_n(k_n-1)}$ over all simulated samples. We also report the asymptotic behaviors of some linear combinations of $\hat{\beta}_n$, including the average estimation bias (Bias), the standard error of estimates (SE_emp), the average estimated standard error (SE_est) and the coverage probability (CP) of a 95% confidence interval over 10000 replications. All simulations are implemented in R (R Core Team, 2021).

According to the established theoretical results in Section 4.1, we can design certain diverging settings as shown in Examples 1–3. Tables 1–3 demonstrate the simulation results for the three considered examples, wherein #para stands for the total number of unknown parameters. As seen, the biases and standard errors generally decrease as the sample size

n increases, and the coverage probabilities are close to the nominal value when n is large. These results suggest that the performance of the MLE is satisfactory under the considered diverging settings, and the asymptotic properties are valid.

Table 1: Results for simulated Example 1.

n	k	#para	AMSE	Bias	SE_emp	SE_est	CP
100	10	27	2.9282	0.0110	1.7587	1.6932	0.9278
500	22	63	1.0520	0.0114	1.1039	1.0905	0.9326
1000	31	90	0.7218	0.0023	0.8439	0.8458	0.9422
5000	70	207	0.3473	0.0080	0.6015	0.5997	0.9426
10000	100	297	0.2472	-0.0011	0.5077	0.5092	0.9493

Table 2: Results for simulated Example 2.

n	d	#para	AMSE	Estimate: $\mathbf{v}_{n1}^\top \vec{\beta}_n$				Estimate: $\mathbf{v}_{n2}^\top \vec{\beta}_n$			
				Bias	SE_emp	SE_est	CP	Bias	SE_emp	SE_est	CP
100	20	40	1.2580	0.1010	1.0868	0.9118	0.9137	0.0646	1.0786	0.9129	0.9173
500	44	88	0.1496	0.0323	0.4050	0.3813	0.9342	0.0305	0.4141	0.3906	0.9355
1000	63	126	0.0687	0.0252	0.2627	0.2517	0.9394	0.0445	0.2656	0.2538	0.9370
5000	141	282	0.0133	0.0223	0.1141	0.1133	0.9452	0.0198	0.1141	0.1134	0.9462
10000	200	400	0.0068	0.0169	0.0818	0.0817	0.9467	0.0171	0.0818	0.0821	0.9473

At last, we explore the hypothesis testing based on the large sample Wald test. Consider the model (5.14) with $n = 10000$, and the other settings remain the same as before. For each scenario with a null hypothesis H_0 , we are interested in the comparison between its

Table 3: Results for simulated Example 3.

n	d	k	#para	AMSE	Bias	SE_emp	SE_est	CP
100	6	9	48	2.9184	0.0229	1.7543	1.6355	0.9233
500	9	14	117	0.6827	0.0070	0.8175	0.8085	0.9401
1000	11	16	165	0.3995	0.0020	0.6576	0.6482	0.9418
5000	16	25	384	0.1233	-0.0030	0.3675	0.3656	0.9466
10000	20	30	580	0.0744	0.0024	0.2841	0.2819	0.9479

Figure 1: Asymptotic results for testing $H_0 : \mathbf{a}_n^\top \vec{\beta}_n = \mathbf{a}_n^\top \vec{\beta}_{n0}$ in Example 1. The left panel gives the estimated null density of large-sample Wald test (circle points) and the density of the chi-square distribution χ^2_1 under H_0 (red curve). The right panel gives the Q-Q plot for the Wald test statistic under H_0 .

estimated density curve and the density curve of its corresponding χ^2 distribution, and the Q-Q plots for the Wald test statistic under H_0 . The related results are shown in Figures 1–3. As seen, the χ^2 approximation for the null distribution is reasonably accurate, and the

Figure 2: Asymptotic results for testing $H_0 : \mathbf{V}_n^\top \vec{\beta}_n = \mathbf{V}_n^\top \vec{\beta}_{n0}$ in Example 2. The left panel gives the estimated null density of large-sample Wald test (circle points) and the density of the chi-square distribution χ^2_2 under H_0 (red curve). The right panel gives the Q-Q plot for the Wald test statistic under H_0 .

theoretical quantiles are approximated very well by the sample quantiles.

5.2 Application I

In this section, we apply the asymptotical results of the SMLR model to a real dataset for statistical inference. We consider the 1996 American National Election Study (NES96) dataset, which can be found in the CRAN package `faraway`. The data contains 944 respondents and each respondent is consisted of 10 related variables. In accordance with the studies in Faraway (2016) and Price et al. (2019), we take 3 explanatory variables into account, including education level (categorical with 7 levels), income (categorical with 24 levels), and age (numerical). Each covariate is standardized to have mean 0 and standard deviation 1.

Figure 3: Asymptotic results for testing $H_0 : \mathbf{b}_n^\top \vec{\beta}_n = \mathbf{b}_n^\top \vec{\beta}_{n0}$ in Example 3. The left panel gives the estimated null density of large sample Wald test (circle points) and the density of the chi-square distribution χ_1^2 under H_0 (red curve). The right panel gives the Q-Q plot for the Wald test statistic under H_0 .

The intercepts are considered as well. In addition, the response variable is the self-identified political affiliation of voters. As seen from Table 4, the response variable originally has 7 categories (Original Categories) and later the number is reduced to 3 (Grouped Categories) by Faraway (2016). The objective here is to check whether the subjective grouping of the response categories are recommended by using the proposed SMLR model.

The fitted coefficient matrix $\hat{\beta} = (\hat{\beta}_{ij}) \in \mathbb{R}^{4 \times (k-1)}$ for the original and grouped categories are shown in Table 5, where the j -th row vector of β is denoted by $\beta_{j\bullet}$. We observe that $\hat{\beta}_{1\bullet}$ is very close to 0 under the grouped categories, which implies the effect of the age factor could be weak. However, such effect of age on the original categories seems strong. To verify

Table 4: Summary of NES96 categories.

Original Categories (size)	Grouped Categories (size)
Strong Democrat (200)	Democrat (380)
Weak Democrat (180)	
Independent Democrat (108)	Independent (239)
Independent (37)	
Independent Republican (94)	
Weak Republican (150)	Republican (325)
Strong Republican (175)	

Table 5: Fitted coefficients of the SMLR model based on the NES96 data.

	Original Categories						Grouped Categories	
	$\hat{\beta}_{\bullet 1}$	$\hat{\beta}_{\bullet 2}$	$\hat{\beta}_{\bullet 3}$	$\hat{\beta}_{\bullet 4}$	$\hat{\beta}_{\bullet 5}$	$\hat{\beta}_{\bullet 6}$	$\hat{\beta}_{\bullet 1}$	$\hat{\beta}_{\bullet 2}$
Intercept ($\hat{\beta}_{0\bullet}$)	0.6304	0.1824	-0.8353	0.0667	0.5098	0.6198	0.0094	0.2519
Age ($\hat{\beta}_{1\bullet}$)	-0.1222	-0.0794	0.0815	0.2118	0.0728	0.1836	-0.0474	0.0103
Education Level ($\hat{\beta}_{2\bullet}$)	0.0391	0.1050	-0.2889	0.0010	0.0175	0.0925	-0.0365	0.0120
Income ($\hat{\beta}_{3\bullet}$)	-0.5525	-0.1564	0.0168	-0.1116	-0.1451	-0.0377	-0.2570	-0.2312

this, we conduct the following hypothesis test to investigate the effects of the age factor,

$$H_0 : \beta_{1\bullet} = \mathbf{0} \iff H_1 : \beta_{1\bullet} \neq \mathbf{0}. \tag{5.15}$$

By Corollary 5 with $\mathbf{V} = \mathbf{I}_{k-1} \otimes (0, 1, 0, 0)^\top$, the Wald test statistic $\hat{\beta}_{1\bullet}^\top [\mathbf{V}^\top \mathbf{Q}_n^{-1}(\hat{\beta}) \mathbf{V}]^{-1} \hat{\beta}_{1\bullet}$ follows χ^2_{k-1} under H_0 . For the grouped categories with $k = 3$, the Wald test statistic is 1.0572 and the p -value is 0.5894. As a result, the age factor is of no significance to identify

the political affiliation of voters under the grouped categories. On the other hand, the Wald test statistic is 18.3178 and the p -value is 0.0055 for the original categories, meaning the effects of the age factor are non-negligible. These results imply that the grouping of the response categories is not supported by the data, which is consistent with the conclusion in Price et al. (2019).

5.3 Application II

As discussed, the main advantages of the proposed SMLR over CMLR1 and CMLR2 are from the computational and asymptotical perspectives, while these three MLRs have been proved to be equivalent in Section 3.3. Therefore, these advantages may not be evident in a finite-sample real application. On the other hand, the penalized counterparts of these three MLRs are no longer equivalent. The penalty terms of CMLR1 and CMLR2 depend on the pre-specified constraints in order to make the models identifiable, while the regularized SMLR solves a unconstrained optimization problem without the reference category. The simplex-based structure is expected to be more convenient and efficient in terms of statistical analysis and algorithmic design (Zhang and Liu, 2014).

In this section, we apply the three ridge-penalized MLR models to three real datasets from the UCI machine learning repository to further illustrate the usefulness of the SMLR-based models. The summary of the datasets are provided in Table 6, where $n/d/k$ stand for the sample size of the training set, the dimension of covariates and the number of categories, respectively. In addition, $n_{\min}/n_{\max}$ represent the size of the minority/majority categories, respectively. Note that the ridge-penalized CMLR1 may have several versions depending on the choice of the reference category, while the regularized CMLR2 and SMLR are unique.

Given a fixed penalty factor λ , we train all the three ridge-penalized MLR models for each dataset, and compare their performances in terms of the label prediction accuracy.

Table 6: Summary of real datasets and results of training ridge penalized MLRs.

Dataset	Information					Accuracy			
	n	d	k	$n_{\min}$	$n_{\max}$	λ	CMLR1	CMLR2	SMLR
Breast	106	9	6	14	22	1.5	0.7830; 0.8019; 0.7925; 0.7547; 0.7830; 0.8113	0.8019	0.8113
Segmentation	210	19	7	30	30	1.1	0.9524; 0.9524; 0.9429; 0.9571 0.9619; 0.9571; 0.9524	0.9571	0.9619
Glass	214	9	6	9	76	0.5	0.7196; 0.7196; 0.7150; 0.7243; 0.7243; 0.7196	0.7336	0.7336

As seen from Table 6, the selection of reference category has a significant effect on the prediction accuracies, which is expected and was widely accepted in the literature (e.g., Tutz et al., 2015). This finding suggests that the use of penalized CMLR1 has to be cautious and an appropriate reference category needs to be selected beforehand. In practice, an extra tuning process on all possible reference categories may be needed to guarantee a stable performance of the penalized CMLR1. On the other hand, the penalized CMLR2 and SMLR do not depend on the reference categories. Such kind of MLR-based models offers a symmetric and systematic insight into all categories, and becomes more appealing in the literature; see Friedman et al. (2010); Zahid and Tutz (2013a,b); Hastie et al. (2015); Powers et al. (2018); de Jong et al. (2019). Between these two MLRs, the penalized SMLR is preferred in terms of computational convenience and establishment of the statistical properties. The similar

arguments were also made in Zhang and Liu (2014). Nevertheless, substantial efforts are needed to investigate the theoretical properties of the penalized SMLR, which is an on-going work of the authors. In summary, the advantages of the SMLR could be better exploited in a regularization setting and this study will serve as a building block in further exploration of the SMLR-based models.

6. Conclusion

In this paper, we have proposed a novel simplex-based MLR model which has several distinct features. Compared to the regular MLRs, the proposed SMLR circumvents the restrictions such as reference category selection and sum-to-zero constraint, and hence it is computationally efficient. In addition, the SMLR can be treated as an unified framework which connects the binomial and multinomial logistic regressions. Moreover, the asymptotic results, statistical learning theory and the properties under a general kernel of the SMLR have been well established, even when the number of covariates and the number of categories increase with the sample size. It is worth mentioning that these statistical properties for the MLR under a diverging number of categories are rarely investigated in the literature, and this study fills the gap due to the intimate relation between the SMLR and the regular MLRs. At last, numerical simulations and real examples have shown the excellent performance of the SMLR under a variety of scenarios.

One possible future research direction is to extend the SMLR to more complicated datasets such as the clustered multinomial data and the count data, which are ubiquitous in application areas including genomics, sports, imaging analysis and text mining. For example, Wang

(2011) studied the clustered binary LR model with a diverging number of covariates, while its extension to the clustered multinomial data remains challenging. In terms of the categorical counts data, Zhang et al. (2017) proposed the reference-based MLR model. However, the selection of the reference category is subjective and it is generally not computationally efficient. The proposed SMLR model could be a promising tool to deal with such multicategory data. Another potential research topic is to adapt the SMLR to the high-dimensional settings. As discussed in Section 5.3, the penalized SMLR is useful in such cases and the investigation is on-going.

Supplementary Materials

The online Supplementary Materials provide (i). Preliminary lemmas used for establishing the theoretical results in the manuscript and (ii). Technical proofs for all the propositions and theorems in the manuscript.

Acknowledgement

We are grateful to the editors and three referees for their insightful comments that have led to a substantial improvement to an earlier version of the paper.

References

- Albert, A. and J. A. Anderson (1984). On the existence of maximum likelihood estimates in logistic regression models. *Biometrika* 71(1), 1–10.

- Anderson, J. and V. Blair (1982). Penalized maximum likelihood estimation in logistic regression and discrimination. *Biometrika* 69(1), 123–136.
- Anderson, J. A. (1972). Separate sample logistic discrimination. *Biometrika* 59(1), 19–35.
- Baker, S. G. (1994). The multinomial-poisson transformation. *Journal of the Royal Statistical Society: Series D* 43(4), 495–504.
- Bartlett, P. L., M. I. Jordan, and J. D. McAuliffe (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association* 101(473), 138–156.
- Berger, A. L., V. J. D. Pietra, and S. A. D. Pietra (1996). A maximum entropy approach to natural language processing. *Computational Linguistics* 22(1), 39–71.
- Böhning, D. (1992). Multinomial logistic regression algorithm. *Annals of the Institute of Statistical Mathematics* 44(1), 197–200.
- Boyd, S. and L. Vandenberghe (2004). *Convex Optimization*. Cambridge University Press.
- Bühlmann, P. and S. Van De Geer (2011). *Statistics for High-dimensional Data: methods, theory and applications*. Springer.
- Cramer, J. S. (2003). *Logit Models from Economics and Other Fields*. Cambridge University Press.
- de Jong, V. M., M. J. Eijkemans, B. van Calster, D. Timmerman, K. G. Moons, E. W. Steyerberg, and M. van Smeden (2019). Sample size considerations and predictive performance of multinomial logistic prediction models. *Statistics in Medicine* 38(9), 1601–1619.
- Dekel, O. and O. Shamir (2010). Multiclass-multilabel classification with more classes than examples. In *International Conference on Artificial Intelligence and Statistics*, pp. 137–144.
- Deng, J., A. C. Berg, K. Li, and F.-F. Li (2010). What does classifying more than 10,000 image categories tell us? In *European Conference on Computer Vision*, pp. 71–84. Springer.

- hr/>
- Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and F.-F. Li (2009). Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255.
- Fahrmeir, L. and H. Kaufmann (1985). Consistency and asymptotic normality of the maximum likelihood estimator in generalized linear models. *Annals of Statistics* 13(1), 342–368.
- Fang, J. and G. Y. Yi (2021). Matrix-variate logistic regression with measurement error. *Biometrika* 108(1), 83–97.
- Faraway, J. J. (2016). *Extending The Linear Model with R: generalized linear, mixed effects and nonparametric regression models* (2nd ed.). Chapman and Hall/CRC.
- Forbes, C., M. Evans, N. Hastings, and B. Peacock (2011). *Statistical Distributions* (4th ed.). John Wiley & Sons.
- Friedman, J., T. Hastie, and R. Tibshirani (2000). Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *Annals of Statistics* 28(2), 337–407.
- Friedman, J., T. Hastie, and R. Tibshirani (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software* 33(1), 1–22.
- Fu, S., S. Zhang, and Y. Liu (2018). Adaptively weighted large-margin angle-based classifiers. *Journal of Multivariate Analysis* 166, 282–299.
- Gao, Q., X. Du, X. Zhou, and F. Xie (2018). Asymptotic properties of maximum quasi-likelihood estimators in generalized linear models with diverging number of covariates. *Journal of Systems Science and Complexity* 31(5), 1362–1376.
- Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press.
- Hastie, T., R. Tibshirani, and J. H. Friedman (2009). *The Elements of Statistical Learning: data mining, inference and prediction* (2nd ed.). Springer.
- Hastie, T., R. Tibshirani, and M. Wainwright (2015). *Statistical Learning with Sparsity: the lasso and generalizations*. CRC Press.

- He, X. and Q.-M. Shao (2000). On parameters of increasing dimensions. *Journal of Multivariate Analysis* 73(1), 120–135.
- Hill, S. I. and A. Doucet (2007). A framework for kernel-based multi-category classification. *Journal of Artificial Intelligence Research* 30, 525–564.
- Hosmer, D. W., S. Lemeshow, and R. X. Sturdivant (2013). *Applied Logistic Regression* (3rd ed.). John Wiley & Sons.
- Krishnapuram, B., L. Carin, M. A. Figueiredo, and A. J. Hartemink (2005). Sparse multinomial logistic regression: Fast algorithms and generalization bounds. *IEEE Transactions on Pattern Analysis & Machine Intelligence* 27(6), 957–968.
- Lang, J. B. (1996). On the comparison of multinomial and poisson log-linear models. *Journal of the Royal Statistical Society: Series B* 58(1), 253–266.
- Lange, K. and T. T. Wu (2008). An MM algorithm for multicategory vertex discriminant analysis. *Journal of Computational and Graphical Statistics* 17(3), 527–544.
- Lemeshow, S. and D. W. Hosmer (2014). Logistic regression in practice. In *Wiley StatsRef: Statistics Reference Online*, pp. 1–15. Wiley Online Library.
- Liang, H. and P. Du (2012). Maximum likelihood estimation in logistic regression models with a diverging number of covariates. *Electronic Journal of Statistics* 6, 1838–1846.
- Liu, T.-Y., Y. Yang, H. Wan, Q. Zhou, B. Gao, H.-J. Zeng, Z. Chen, and W.-Y. Ma (2005). An experimental study on large-scale web categorization. In *International Conference on World Wide Web*, pp. 1106–1107. Association for Computing Machinery.
- Malouf, R. (2002). A comparison of algorithms for maximum entropy parameter estimation. In *Conference on Natural Language Learning*, Volume 20, pp. 1–7. Association for Computational Linguistics.

- McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models*. CRC Press.
- Mo, W. and Y. Liu (2021). Supervised learning. In *Wiley StatsRef: Statistics Reference Online*, pp. 1–20. Wiley Online Library.
- Mroueh, Y., T. Poggio, L. Rosasco, and J.-J. Slotine (2012). Multiclass learning with simplex coding. In *Advances in Neural Information Processing Systems*, pp. 2789–2797.
- Ng, A., J. Ngiam, C. Y. Foo, Y. Mai, and C. Suen (2013). *Softmax Regression*. http://ufldl.stanford.edu/wiki/index.php/Softmax_Regression, UFLDL Tutorial.
- Nilsback, M.-E. and A. Zisserman (2008). Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729.
- Portnoy, S. (1985). Asymptotic behavior of M -estimators of p regression parameters when p^2/n is large; ii. normal approximation. *Annals of Statistics* 13(4), 1403–1417.
- Portnoy, S. (1988). Asymptotic behavior of likelihood methods for exponential families when the number of parameters tends to infinity. *Annals of Statistics* 16(1), 356–366.
- Powers, S., T. Hastie, and R. Tibshirani (2018). Nuclear penalized multinomial regression with an application to predicting at bat outcomes in baseball. *Statistical Modelling* 18(5-6), 388–410.
- Price, B. S., C. J. Geyer, and A. J. Rothman (2019). Automatic response category combination in multinomial logistic regression. *Journal of Computational and Graphical Statistics* 28(3), 758–766.
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge University Press.
- Tewari, A. and P. L. Bartlett (2007). On the consistency of multiclass classification methods. *Journal of Machine Learning Research* 8(May), 1007–1025.

- Tsybakov, A. B. (2004). Optimal aggregation of classifiers in statistical learning. *Annals of Statistics* 32(1), 135–166.
- Tutz, G. (2011). *Regression for Categorical Data*. Cambridge University Press.
- Tutz, G., W. Pöbnecker, and L. Uhlmann (2015). Variable selection in general multinomial logit models. *Computational Statistics & Data Analysis* 82, 207–222.
- Van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.
- Wang, L. (2011). GEE analysis of clustered binary data with diverging number of covariates. *Annals of Statistics* 39(1), 389–417.
- Yee, T. W. (2015). *Vector Generalized Linear and Additive Models: with an implementation in R*. Springer.
- Zahid, F. M. and G. Tutz (2013a). Multinomial logit models with implicit variable selection. *Advances in Data Analysis and Classification* 7(4), 393–416.
- Zahid, F. M. and G. Tutz (2013b). Ridge estimation for multinomial logit models with symmetric side constraints. *Computational Statistics* 28(3), 1017–1034.
- Zhang, C. and Y. Liu (2014). Multicategory angle-based large-margin classification. *Biometrika* 101(3), 625–640.
- Zhang, C., M. Pham, S. Fu, and Y. Liu (2018). Robust multicategory support vector machines using difference convex algorithm. *Mathematical Programming* 169(1), 277–305.
- Zhang, T. (2004). Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research* 5(Oct), 1225–1251.
- Zhang, Y., H. Zhou, J. Zhou, and W. Sun (2017). Regression models for multivariate count data. *Journal of Computational and Graphical Statistics* 26(1), 1–13.
- Zhu, J. and T. Hastie (2004). Classification of gene microarrays by penalized logistic regression. *Biostatistics* 5(3), 427–443.

Zhu, J. and T. Hastie (2005). Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics* 14(1), 185–205.

Sheng Fu

Department of Industrial Systems Engineering & Management, National University of Singapore

E-mail: fusheng1007@gmail.com

Piao Chen

Delft Institute of Applied Mathematics, Delft University of Technology, The Netherlands

E-mail: p.chen-6@tudelft.nl

Yufeng Liu

Department of Statistics and Operations Research, University of North Carolina at Chapel Hill

E-mail: yfliu@email.unc.edu

Zhisheng Ye

Department of Industrial Systems Engineering & Management, National University of Singapore

E-mail: yez@nus.edu.sg