

Reduced rank vector generalized linear models for feature extraction

YIYUAN SHE*

Supervised linear feature extraction can be achieved by fitting a reduced rank multivariate model. This paper studies rank penalized and rank constrained vector generalized linear models. From the perspective of thresholding rules, we build a framework for fitting singular value penalized models and use it for feature extraction. Through solving the rank constraint form of the problem, we propose progressive feature space reduction for fast computation in high dimensions with little performance loss. A novel projective cross-validation is proposed for parameter tuning in such nonconvex setups. Real data applications are given to show the power of the methodology in supervised dimension reduction and feature extraction.

AMS 2000 SUBJECT CLASSIFICATIONS: Primary 62H30, 62J12; secondary 62H12.

KEYWORDS AND PHRASES: Reduced rank regression, Generalized linear models, Feature extraction, Projective cross-validation, Progressive dimension reduction.

1. BACKGROUND

Recently, high dimensional data analysis attracts a great deal of interest from statisticians. The availability of a large pool of variables (relative to the sample size) poses challenges in statistical modeling because in this high-dimensional setup, both estimation accuracy and model interpretability can be seriously hurt. Dimension reduction is a natural and effective means to reduce the number of unknowns. One can remove nuisance and/or redundant variables, referred to as variable/feature selection; alternatively, one can find low dimensional linear or nonlinear projections of the input data, referred to as feature extraction. In this paper, we focus on linear feature extraction for dimension reduction purposes.

The most popular approach for linear feature extraction is perhaps the principle component analysis (PCA). Given $\mathbf{X} \in \mathbb{R}^{n \times p}$ with n observations and p features, perform the Singular Value Decomposition (SVD) on the data $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T$. Given any $1 \leq r \leq \text{rank}(\mathbf{X})$, denote by \mathbf{V}_r the submatrix of \mathbf{V} consisting of its first r columns. Then

$$\mathbf{Z}_r = \mathbf{X}\mathbf{V}_r$$

constructs r new features as linear combinations of the original features. The extraction is optimal in the sense that $\hat{\mathbf{B}} = \mathbf{X}\mathcal{P}_{\mathbf{V}_r} = \mathbf{X}\mathbf{V}_r\mathbf{V}_r^T$, with $\mathcal{P}_{\mathbf{V}_r}$ being the projection matrix onto the column space of \mathbf{V}_r , gives the best rank- r approximation to \mathbf{X} :

$$\hat{\mathbf{B}} = \arg \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq r} \|\mathbf{X} - \mathbf{B}\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm. A by-product is that the gram matrix of $\mathbf{X}\mathbf{V}_r$ is diagonal, which means all new features are uncorrelated with each other.

On the other hand, PCA is unsupervised. In many statistical learning problems, we want to construct new features that best predict the responses. Suppose $\mathbf{Y} \in \mathbb{R}^{n \times m}$ is the response matrix, n being the sample size and m being the number of response variables. Supervised feature extraction can be given by the *reduced rank regression* (RRR) (Anderson, 1951), with the RRR estimator $\hat{\mathbf{B}}$ defined by

$$\hat{\mathbf{B}} = \arg \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq r} \|\mathbf{Y} - \mathbf{X}\mathbf{B}\|_F^2.$$

Assume \mathbf{X} has full column rank and define $\mathbf{H} = \mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T$. Then $\hat{\mathbf{B}} = \hat{\mathbf{B}}_{ols}\mathbf{V}_r\mathbf{V}_r^T$, where $\hat{\mathbf{B}}_{ols} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y}$. \mathbf{V}_r is formed by the first r columns of \mathbf{V} from the spectral decomposition $\mathbf{Y}^T\mathbf{H}\mathbf{Y} = \mathbf{V}\mathbf{D}\mathbf{V}^T$. See Reinsel and Velu (1998) for more details. Therefore,

$$\mathbf{Z}_r = \mathbf{X}(\hat{\mathbf{B}}_{ols}\mathbf{V}_r)$$

constructs r new features that best approximate \mathbf{Y} in Frobenius norm, and these new features are, again, uncorrelated. The RRR framework includes the PCA as a special case, by setting $\mathbf{Y} = \mathbf{X}$ (Izenman, 2008). The above RRR solves a *nonconvex* optimization problem in the classical setup $n > p$. Recently, Bunea, She and Wegkamp (2011) studied the problem under $p > n$ and developed finite-sample theories as well as a computational algorithm.

On the other hand, the squared error loss may not always be appropriate. For vector generalized linear models (GLMs), such as discrete responses arising in classification problems, deviance loss is much more reasonable.

Although there is a large body of literature on the RRR—Robinson (1974), Rao (1979), and Brillinger (1981), to name a few, to the best of our knowledge, there is very little work beyond the Gaussian model. Yee and Hastie (2003)

studied the reduced-rank vector GLM problem and used an iterative approximate estimation by fitting RRR repeatedly. Yet this only provides an approximation solution to the original problem and there is no guarantee of converge. Heinen and Rengifo (2008) resorted to a continuation technique to deal with discrete responses.

This paper aims to tackle the penalized and constrained vector GLMs

$$(1.1) \quad \min_{\mathbf{B}} -\log\text{-likelihood}(\mathbf{B}; \mathbf{Y}, \mathbf{X}) + \frac{\lambda^2}{2} \text{rank}(\mathbf{B}), \quad \text{and}$$

$$(1.2) \quad \min_{\mathbf{B}} -\log\text{-likelihood}(\mathbf{B}; \mathbf{Y}, \mathbf{X}) \quad \text{s.t. } \text{rank}(\mathbf{B}) \leq r.$$

The imposed reduced rank structure is based on the belief that the features' relevant directions, in response to \mathbf{Y} , define a lower dimensional subspace in \mathbb{R}^p . The rank of such an estimator determines the number of new features to construct. These two problems are *not* equivalent to each other due to their nonconvexity. In fact, the rank function is nonconvex and discrete (and thus nondifferentiable), thereby posing a challenge in optimization. Our algorithms boil down to an iterative version, which is not surprising in the GLM setup.

The rest of the paper is organized as follows. Section 2 starts by studying a matrix approximation problem, and then builds a framework for fitting singular-value penalized multivariate GLMs. Supervised feature extraction can be attained for non-Gaussian models not necessarily using the squared error loss. The framework covers a wide family of penalty functions. A new parameter tuning strategy is proposed. Section 3 tackles the rank constrained GLM problem and comes up with a feature space reduction technique. Through this, (1.1) and (1.2) can be combined to achieve better estimation accuracy and computational efficiency. Section 4 illustrates real applications of the proposed methodology. We conclude in Section 5. All technical details are left to the Appendix.

2. PENALIZED VECTOR GLMS FOR FEATURE EXTRACTION

In this section, we study the *penalized* form reduced rank GLMs (1.1). Our algorithm and analysis apply to $p > n$ situations and cover a large family of singular-value penalties, including nuclear norm, Frobenius norm, Schatten p -penalties ($0 < p < 1$), and rank penalty. To achieve such generality, we start by studying a simpler matrix approximation problem.

2.1 Singular-value penalized matrix approximation

We consider the problem of matrix approximation with a singular value penalty

$$(2.1) \quad \min_{\mathbf{B}} \frac{1}{2} \|\mathbf{Y} - \mathbf{B}\|_F^2 + \sum_i P(\sigma_i^{(\mathbf{B})}; \lambda),$$

where $\sigma_i^{(\mathbf{B})}$ denote the singular values of \mathbf{B} . The choice of the penalty function P is flexible. For example, $P(t; \lambda) = \lambda|t|$ gives a multiple of the sum of singular values corresponding to the trace norm or nuclear norm penalty. When $P(t; \lambda) = \lambda^2 1_{t \neq 0}/2$, we get the rank penalty which is discrete and nonconvex. For a general P , the closed-form solution to (2.1) is not known (to the best of our knowledge). We address the problem from the standpoint of threshold functions.

Definition 2.1 (Threshold function). A threshold function is a real valued function $\Theta(t; \lambda)$ defined for $-\infty < t < \infty$ and $0 \leq \lambda < \infty$ such that

1. $\Theta(-t; \lambda) = -\Theta(t; \lambda)$,
2. $\Theta(t; \lambda) \leq \Theta(t'; \lambda)$ for $t \leq t'$,
3. $\lim_{t \rightarrow \infty} \Theta(t; \lambda) = \infty$, and
4. $0 \leq \Theta(t; \lambda) \leq t$ for $0 \leq t < \infty$.

Remarks. (i) A vector version of Θ (still denoted by Θ) is defined componentwise if either t or λ is replaced by a vector. (ii) There may be some ambiguity in defining a threshold function. For example, the hard-thresholding can be defined as $\Theta_H(t; \lambda) = t 1_{|t| > \lambda}$ or $\Theta_H(t; \lambda) = t 1_{|t| \geq \lambda}$. Fortunately, commonly used thresholding rules have at most finitely many discontinuity points and such discontinuities rarely occur in real data. When applying Θ to a quantity (say t), we always make the implicit assumption that Θ is continuous at t . (iii) By definition, $\Theta^{-1}(u; \lambda) \triangleq \sup\{t : \Theta(t; \lambda) \leq u\}$, $\forall u > 0$ must be monotonically increasing on $(0, \infty)$ and bounded between the identity line and $u = 0$; its derivative is defined almost everywhere on $(0, \infty)$. We assume that

$$d\Theta^{-1}(u; \lambda)/du \geq 1 - \mathcal{L}_\Theta \quad \text{a.e. on } (0, \infty)$$

for some constant $\mathcal{L}_\Theta \in [0, 1]$ independent of λ . (In fact, for all *convex* penalties constructed through (2.5), \mathcal{L}_Θ can be set to be 0.)

Next we introduce the matrix thresholding.

Definition 2.2 (Matrix threshold function). Given any threshold function $\Theta(\cdot; \lambda)$, its matrix version Θ^σ is defined as follows

$$(2.2) \quad \Theta^\sigma(\mathbf{B}; \lambda) \triangleq \mathbf{U} \text{diag}\{\Theta(\sigma_i^{(\mathbf{B})}; \lambda)\} \mathbf{V}^T, \quad \forall \mathbf{B} \in \mathbb{R}^{n \times m},$$

where \mathbf{U} , \mathbf{V} , and $\text{diag}\{\sigma_i^{(\mathbf{B})}\}$ are obtained from the SVD of \mathbf{B} : $\mathbf{B} = \mathbf{U}^T \text{diag}\{\sigma_i^{(\mathbf{B})}\} \mathbf{V}$.

Note that $\Theta(0; \lambda) = 0$ by definition, and $\Theta^\sigma(\mathbf{B}; \lambda)$ is not affected by the ambiguity of the SVD form.

Proposition 2.1. *Given an arbitrary thresholding rule Θ , let P be any function satisfying*

$$(2.3) \quad P(\theta; \lambda) - P(0; \lambda) = \int_0^{|\theta|} (\sup\{s : \Theta(s; \lambda_k) \leq u\} - u) du + q(\theta; \lambda),$$

Table 1. Some basic singular-value penalties and their coupled thresholding functions

	Nuclear norm	Frobenius	Rank	Schatten-p, $p \in (0, 1)$
Penalty function	$\lambda \ \mathbf{B}\ _* = \lambda \sum \sigma_i^{(\mathbf{B})}$	$\lambda \ \mathbf{B}\ _F^2$	$\frac{\lambda^2}{2} \text{rank}(\mathbf{B})$	$\lambda \sum (\sigma_i^{(\mathbf{B})})^p$
Thresholding rule	$(t - \text{sgn}(t)\lambda)1_{ t >\lambda}$ (soft)	$\frac{t}{1+\lambda}$ (ridge)	$t1_{ t >\lambda}$ (hard)	Ex 2.7 in She (2012)

where $q(\cdot; \lambda)$ is nonnegative and $q(\Theta(t; \lambda); \lambda) = 0$ for all $t \in \mathbb{R}$. Then, the singular-value penalized minimization

$$(2.4) \quad \min_{\mathbf{B}} F(\mathbf{B}) = \|\mathbf{Y} - \mathbf{B}\|_F^2/2 + \sum P(\sigma_i^{(\mathbf{B})}; \lambda)$$

has a unique optimal solution $\hat{\mathbf{B}} = \Theta^\sigma(\mathbf{Y}; \lambda)$ for every \mathbf{Y} , provided $\Theta(\cdot; \lambda)$ is continuous at any singular value of \mathbf{Y} .

See Appendix A for its proof.

The function q is often just zero, but can be nonzero in certain cases. In fact, we can use a nontrivial q to attain the exact rank penalty; see (2.7). The proposition implies that multiple (infinitely many, as a matter of fact) penalties can result in the same solution, which justifies our thresholding launching point (rather than a penalty one). Some examples of the penalty P and the coupled Θ are listed in Table 1.

We point out two special cases of Proposition 2.1 as follows.

A fusion between nuclear norm and Frobenius norm Define a continuous thresholding rule

$$(2.5) \quad \Theta_B(t; \lambda, M) = \begin{cases} 0, & \text{if } |t| \leq \lambda, \\ t - \lambda \text{sgn}(t), & \text{if } \lambda < |t| < \lambda + M, \\ \frac{t}{1 + \frac{\lambda}{M}}, & \text{if } |t| > \lambda + M. \end{cases}$$

When $M \rightarrow \infty$, Θ_B becomes the soft-thresholding. When $M = 0$, Θ_B reduces to the ridge thresholding. The penalty constructed from (2.5) is given by

$$P(\theta; \lambda, M) = \begin{cases} \lambda|\theta|, & \text{if } |\theta| \leq M, \\ \lambda \frac{\theta^2 + M^2}{2M}, & \text{if } |\theta| > M, \end{cases}$$

which is exactly the ‘‘Berhu’’ penalty (Owen, 2007) whose composition reverses that of Huber’s robust loss function. The Berhu penalty on the singular values provides a convex fusion of the nuclear norm penalty and the Frobenius norm (squared) penalty in the problem of (2.4). Unlike the elastic net (Zou and Hastie, 2005), this fusion is nonlinear and fully preserves the nondifferentiable behavior (around zero) of the nuclear norm.

A fusion between rank and Frobenius norm A direct thresholding rule that fuses the hard-thresholding and the ridge-thresholding is the hard-ridge thresholding (She, 2009)

$$(2.6) \quad \Theta_{HR}(t; \lambda, \eta) = \begin{cases} 0, & \text{if } |t| < \lambda, \\ \frac{t}{1+\eta}, & \text{if } |t| \geq \lambda. \end{cases}$$

Setting $q \equiv 0$ in Proposition 2.1, we obtain one associated penalty

$$P(\theta; \lambda, \eta) = \begin{cases} -\frac{1}{2}\theta^2 + \lambda|\theta|, & \text{if } |\theta| < \frac{\lambda}{1+\eta}, \\ \frac{1}{2}\eta\theta^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}, & \text{if } |\theta| \geq \frac{\lambda}{1+\eta}. \end{cases}$$

Interestingly, noticing that Θ_{HR} is discontinuous at λ , we can choose

$$(2.7) \quad q(\theta; \lambda, \eta) = \begin{cases} \frac{(1+\eta)(\lambda-|\theta|)^2}{2}, & \text{if } 0 < |\theta| < \lambda, \\ 0, & \text{if } \theta = 0 \text{ or } |\theta| > \lambda, \end{cases}$$

which leads to $P(\theta) = \frac{1}{2}\eta\theta^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}1_{\theta \neq 0}$. Therefore, $\Theta_{HR}(\cdot; \lambda, \eta)$ can deal with the following rank-Frobenius penalty in (2.4)

$$(2.8) \quad \frac{1}{2}\eta\|\mathbf{B}\|_F^2 + \frac{1}{2}\frac{\lambda^2}{1+\eta}\text{rank}(\mathbf{B}).$$

This penalty may be of interest in statistical learning tasks that have joint concerns of accuracy and parsimony: the rank portion enforces high rank deficiency, while the ridge (Frobenius) portion shrinks \mathbf{B} to compensate for large noise and decorrelates the input variables in large- p applications.

At the end of this subsection, we present a perturbation result which will be used to establish the main result in the next subsection.

Proposition 2.2. Given $\mathbf{Y} \in \mathbb{R}^{n \times m}$, let $Q(\mathbf{B}) = \|\mathbf{Y} - \mathbf{B}\|_F^2/2 + \sum P_\Theta(\sigma_i^{(\mathbf{B})}; \lambda)$, where P_Θ is the penalty obtained from (2.5). Denote by $\hat{\mathbf{B}}$ the minimizer of $Q(\mathbf{B})$. Then for any matrix $\Delta \in \mathbb{R}^{n \times m}$

$$Q(\hat{\mathbf{B}} + \Delta) - Q(\hat{\mathbf{B}}) \geq \frac{C_1}{2}\|\Delta\|_F^2,$$

where $C_1 = 1 - \mathcal{L}_\Theta \geq 0$.

See Appendix B for its proof.

2.2 Singular-value penalized vector GLMs

In this subsection, we generalize the results obtained for matrix approximation to vector GLMs.

Let $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_m] \in \mathbb{R}^{n \times m}$ be the response matrix with m response variables and n samples for each. Assume $y_{i,k}$ are independent and each follows a distribution in the natural exponential family $f(y_{i,k}; \theta_{i,k}) = \exp(y_{i,k}\theta_{i,k} - b(\theta_{i,k}) + c(y_{i,k}))$, where $\theta_{i,k}$ is the natural parameter. Let $L_{i,k} = \log f(y_{i,k}, \theta_{i,k})$, $L = \sum_k \sum_i L_{i,k}$. The canonical link

function $g = (b')^{-1}$ is applied throughout the paper. Let the model matrix and the corresponding coefficient matrix be

$$(2.9) \quad \begin{aligned} \mathbf{X} &= [\mathbf{x}_1, \dots, \mathbf{x}_n]^T = [\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_p] \in \mathbb{R}^{n \times (p+1)} \quad \text{and} \\ \mathbf{B} &= [\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_m] = [\tilde{\boldsymbol{\beta}}_0, \tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_p]^T \in \mathbb{R}^{(p+1) \times m}, \end{aligned}$$

respectively. If $\tilde{\mathbf{x}}_0 = \mathbf{1}$, $\tilde{\boldsymbol{\beta}}_0$ represents the intercept vector. For convenience, we use $\mathbf{B}^\circ = [\tilde{\boldsymbol{\beta}}_1, \dots, \tilde{\boldsymbol{\beta}}_p]^T$ to denote the submatrix of \mathbf{B} obtained by deleting the first row $\tilde{\boldsymbol{\beta}}_0^T$, and use \mathbf{X}° to denote the submatrix of \mathbf{X} obtained by deleting the first column $\tilde{\mathbf{x}}_0$. Given any GLM with coefficients $\boldsymbol{\beta}$, we introduce

$$\begin{aligned} \boldsymbol{\mu}(\boldsymbol{\beta}) &\triangleq [g^{-1}(\mathbf{x}_i^T \boldsymbol{\beta})]_{n \times 1} \quad \text{and} \\ \mathcal{I}(\boldsymbol{\beta}) &\triangleq \mathbf{X}^T \mathbf{W} \mathbf{X} = \mathbf{X}^T \text{diag} \{b''(\mathbf{x}_i^T \boldsymbol{\beta})\}_{i=1}^n \mathbf{X} \end{aligned}$$

to denote the mean vector and the information matrix at $\boldsymbol{\beta}$. For the m -response vector GLM, the mean matrix $\boldsymbol{\mu}(\mathbf{B}) = [\boldsymbol{\mu}(\boldsymbol{\beta}_1), \dots, \boldsymbol{\mu}(\boldsymbol{\beta}_m)]$.

Remarks. (i) Having $\tilde{\mathbf{x}}_0$ and $\tilde{\boldsymbol{\beta}}_0$ is necessary. For non-Gaussian GLMs, one cannot center the response variables because this may violate the distributional assumption. (ii) For clarity, the above setup does not include any dispersion parameter. But all discussions in this subsection can be trivially extended to the exponential dispersion family $f(y_{i,k}; \theta_{i,k}, \phi) = \exp\{(y_{i,k}\theta_{i,k} - b(\theta_{i,k}))/a(\phi) + c(y_{i,k}, \phi)\}$ which covers the vector Gaussian regression.

Our goal is to minimize (1.1) or more generally, the following objective function

$$(2.10) \quad F(\mathbf{B}) \triangleq - \sum_{k=1}^m \sum_{i=1}^n L_{i,k}(\boldsymbol{\beta}_k; \mathbf{x}_i, y_{ik}) + \sum_{s=1}^{p \wedge m} P(\sigma_s^{(\mathbf{B}^\circ)}; \lambda)$$

for a large family of penalty functions (possibly nonconvex). The penalty is *not* imposed on $\tilde{\boldsymbol{\beta}}_0$.

We construct the following sequence of iterates for solving the problem: given $\mathbf{B}^{(j)}$, perform the update

$$(2.11) \quad \begin{cases} \mathbf{B}^{\circ(j+1)} &= \Theta^\sigma(\mathbf{B}^{\circ(j)} + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\mathbf{B}^{(j)}); \lambda), \\ \tilde{\boldsymbol{\beta}}_0^{(j+1)} &= \tilde{\boldsymbol{\beta}}_0^{(j)} + (\mathbf{Y} - \boldsymbol{\mu}(\mathbf{B}^{(j)}))^T \tilde{\mathbf{x}}_0. \end{cases}$$

Theorem 2.1. *Given an arbitrary thresholding rule Θ , let $P(\cdot)$ be any function satisfying (2.5). Starting with any $\mathbf{B}^{(0)} \in \mathbb{R}^{(p+1) \times m}$, run (2.11) to obtain a sequence $\{\mathbf{B}^{(j)}\}$. Denote by A_k the set of $\{t\boldsymbol{\beta}_k^{(j)} + (1-t)\boldsymbol{\beta}_k^{(j+1)} : t \in (0, 1), j = 1, 2, \dots\}$, $1 \leq k \leq K$, and define*

$$\rho = \max_{1 \leq k \leq m} \sup_{\boldsymbol{\xi}_k \in A_k} \|\mathcal{I}(\boldsymbol{\xi}_k)\|_2.$$

Suppose $\rho < 2 - \mathcal{L}_\Theta$. Then $F(\mathbf{B}^{(j)})$ is decreasing and satisfies

$$(2.12) \quad \begin{aligned} F(\mathbf{B}^{(j)}) - F(\mathbf{B}^{(j+1)}) &\geq C \|\mathbf{B}^{(j)} - \mathbf{B}^{(j+1)}\|_F^2 / 2, \\ j &= 1, 2, \dots \end{aligned}$$

where $C = 2 - \mathcal{L}_\Theta - \rho$. Any limit point \mathbf{B}^* of the sequence $\mathbf{B}^{(j)}$, referred to as a Θ^σ -estimator, is a solution to the following equation:

$$(2.13) \quad \begin{cases} \mathbf{B}^\circ = \Theta^\sigma(\mathbf{B}^\circ + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\mathbf{B}); \lambda), \\ \mathbf{0} = (\mathbf{Y} - \boldsymbol{\mu}(\mathbf{B}))^T \tilde{\mathbf{x}}_0, \end{cases}$$

under the assumption that Θ is continuous at all singular values of $\mathbf{B}^{\circ*} + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\mathbf{B}^*)$.

The proof details are given in Appendix C.

Recall that $\mathcal{L}_\Theta \leq 1$. In implementation, we can scale the model matrix by \mathbf{X}/k_0 for any $k_0 \geq \sqrt{\rho}$ regardless of Θ , and then perform (2.11). The Θ^σ -estimate, obtained on the scaled data, can be scaled back to give an estimate on the original \mathbf{X} . The procedure has a theoretical guarantee of convergence and (2.12) yields a good stopping criterion based on the change in $\mathbf{B}^{(j)}$. Empirically, we always observe $\mathbf{B}^{(j)}$ has a unique limit point. Similar to She (2012), when it is possible to explicitly calculate the curvature parameter \mathcal{L}_Θ , say for SCAD or soft-thresholding, we recommend using the smallest possible value of $k_0 = \sqrt{\rho/(2 - \mathcal{L}_\Theta)}$, which significantly speeds the convergence of the algorithm based on extensive experience. (For example, with a convex penalty we can set $k_0 = \sqrt{\rho/2}$.) We give two typical situations to find an upper bound for ρ in theory.

Example 2.1 (Penalized Vector Gaussian GLM). For Gaussian regression, we can ignore the intercept term (after centering both responses and predictors beforehand), and the objective function (2.10) becomes

$$(2.14) \quad \|\mathbf{Y} - \mathbf{X}\mathbf{B}\|_F^2 / 2 + \sum_{s=1}^{p \wedge m} P(\sigma_s^{(\mathbf{B})}; \lambda).$$

(2.11) reduces to

$$(2.15) \quad \mathbf{B}^{(j+1)} = \Theta^\sigma(\mathbf{B}^{(j)} + \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \mathbf{X} \mathbf{B}; \lambda).$$

Here, $\mathcal{I} = \mathbf{X}^T \mathbf{W} \mathbf{X} = \mathbf{X}^T \mathbf{X}$. According to the theorem, k_0 can be chosen to be $\|\mathbf{X}\|_2$ regardless of the thresholding rule and the penalty, where $\|\cdot\|_2$ denotes the spectral norm.

In the special case of imposing a direct rank penalty, where $\sum_{s=1}^{p \wedge m} P(\sigma_s^{(\mathbf{B})}; \lambda) = \frac{\lambda^2}{2} \text{rank}(\mathbf{B})$, another computational procedure based on the classical RRR algorithm can be used. In fact, RRR studies a relevant but different problem, with no penalty but subject to a low rank constraint. But we can adapt the procedure to minimizing $\|\mathbf{Y} - \mathbf{X}\mathbf{B}\|_F^2 / 2 + \lambda^2 / 2 \cdot \text{rank}(\mathbf{B})$ as follows (cf.

Bunea, She and Wegkamp (2011)). Suppose $\mathbf{X}^T \mathbf{X}$ is non-singular and \mathbf{H} is the hat matrix $\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$. (i) Apply spectral decomposition to $\mathbf{Y}^T \mathbf{H} \mathbf{Y}$: $\mathbf{Y}^T \mathbf{H} \mathbf{Y} = \mathbf{V} \mathbf{D}^2 \mathbf{V}^T$ where $\mathbf{D} = \text{diag}\{d_1, \dots, d_m\}$ with $d_1 \geq d_2 \geq \dots \geq d_m \geq 0$. (ii) Given any value of λ , define $r = \max\{i : d_i \geq \lambda\}$ and $\mathbf{V}_r = \mathbf{V}[:, 1:r]$, by taking the first r columns in \mathbf{V} . (iii) Then the (globally) optimal solution is given by

$$\hat{\mathbf{B}}(\lambda) = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y} \mathcal{P}_{\mathbf{V}_r} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y} \mathbf{V}_r \mathbf{V}_r^T,$$

where $\mathcal{P}_{\mathbf{V}_r}$ is the orthogonal projection onto the column space of \mathbf{V}_r . We can show the Θ^σ -estimate defined by (2.15) reduces to the RRR estimate in this case, the proof details given in Appendix D.

Proposition 2.3. *Suppose $\mathbf{X} \in \mathbb{R}^{n \times p}$ ($n \geq p$) has full column rank and $\|\mathbf{X}\|_2 \leq 1$. Then the RRR estimate $\hat{\mathbf{B}}(\lambda)$ constructed above must satisfy the Θ^σ -equation $\hat{\mathbf{B}} = \Theta_H^\sigma(\hat{\mathbf{B}} + \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \mathbf{X} \hat{\mathbf{B}}; \lambda)$ for matrix hard-thresholding Θ_H^σ .*

Unlike RRR, our algorithm and convergent analysis do not require \mathbf{X} to have full rank or $n > p$. In comparison with the large- p RSC (Bunea, She and Wegkamp, 2011), (2.11) applies to any Θ (and covers all vector GLMs).

Example 2.2 (Penalized Vector Logistic GLM). Assume a classification setup where y_{ik} are all binary. The singular-value penalized vector logistic regression minimizes

$$(2.16) \quad - \sum_{k=1}^m \sum_{i=1}^n (y_{i,k} \mathbf{x}_i^T \boldsymbol{\beta}_k - \log(1 + \exp(\mathbf{x}_i^T \boldsymbol{\beta}_k))) + \sum_{s=1}^{p \wedge m} P(\sigma_s(\mathbf{B}^\circ); \lambda).$$

The first iteration step in (2.11) becomes

$$(2.17) \quad \mathbf{B}^{\circ(j+1)} = \Theta^\sigma \left(\mathbf{B}^{\circ(j)} + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \left[1 / (1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta}_k^{(j)})) \right]_{n \times m}; \lambda \right).$$

In R, the matrix $\boldsymbol{\mu}(\mathbf{B})$ can be simply constructed by $1 / (1 + \exp(-\mathbf{X} \% * \% \mathbf{B}))$. Since $\mathbf{W}(\boldsymbol{\beta}) = \text{diag}\{b''(\mathbf{x}_i^T \boldsymbol{\beta})\} = \text{diag}\{\pi_i(1 - \pi_i)\} \preceq \mathbf{I}/4$ with $\pi_i = 1 / (1 + \exp(-\mathbf{x}_i^T \boldsymbol{\beta}))$, a crude but general choice of the scaling constant is $k_0 \geq \|\mathbf{X}\|_2/2$, again, regardless of Θ and λ . Yet in applying a convex penalty such as the nuclear norm penalty, we can use $k_0 = \|\mathbf{X}\|_2 / (2\sqrt{2})$ to speed the convergence.

Some related works There has been a surge of interest in nuclear norm penalization recently, in which case the penalty in (2.10) simplifies to a multiple of the sum of all singular values of \mathbf{B}° or $\lambda \|\mathbf{B}^\circ\|_*$. This gives a *convex* optimization problem. In the statistics community, Yuan et al.

(2007) seem to be the first to study the nuclear norm penalized least squares estimator. A popular equivalent formulation of the nuclear norm minimization in optimization is through semidefinite programming (SDP) Fazel (2002). See, e.g., Candès and Recht (2009), Candès and Tao (2010), and Ma et al. (2009) for some recent theoretical and computational achievements.

Although the nuclear norm provides a convex relaxation to the rank penalty, this approximation works only under certain regularity conditions (e.g., Candès and Plan (2011)). Bunea, She and Wegkamp (2011) show that direct rank penalization achieves the same oracle rate in a much less restrictive manner. Yet in addition to the reduced rank regression studies (see Example 2.1), there have been very few attempts to extend the rank penalization beyond the Gaussian framework. Two commonly cited works are Yee and Hastie (2003) and Heinen and Rengifo (2008). See Section 1 for their limitations. In comparison with these works, our matrix thresholding algorithm has a theoretical guarantee of convergence, is simple to implement, and covers a wide family of penalty functions as well as loss functions.

Finally, we point out a major difference between the thresholding-based iterative selection procedures (TISP) (She, 2009) and the proposed algorithm which can be referred to as *matrix-TISP*. TISP aims for variable selection in a single-response model, while here we discuss singular value regularization in vector GLMs. The singular-value sparsity or low rankness, different than coefficient sparsity, offers a new type of parsimony that can be used for supervised feature extraction. It brings a true multivariate flavor into our analysis.

Feature extraction In many high-dimensional problems, *feature extraction*, by transforming the input variables and creating a reduced set of new features, is a useful technique for dimension-reduction. For example, PCA considers linear projections of correlated variables to construct new orthogonal features ordered by decreasing variances. For singular-value penalized models, once a low-rank estimate $\hat{\mathbf{B}}$ is obtained, one can attain the same goal. Suppose the rank of $\hat{\mathbf{B}}$ is r . A direct way is to apply the reduced form SVD on $\hat{\mathbf{B}}$, getting $\hat{\mathbf{B}} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ with \mathbf{D} an $r \times r$ diagonal matrix. Next, construct a new model matrix

$$(2.18) \quad \text{Type-I: } \mathbf{Z} \triangleq \mathbf{X} \mathbf{U} \text{ (or } \mathbf{X} \mathbf{U} \mathbf{D}),$$

which has only r new predictors. We refer to this as *Type-I extraction*. It can be used for parameter tuning later.

On the other hand, it may be preferred to work on $\mathbf{X} \hat{\mathbf{B}}$ in some situations. Perform the spectral decomposition $\hat{\mathbf{B}}^T \mathbf{X}^T \mathbf{X} \hat{\mathbf{B}} = \mathbf{V} \mathbf{D} \mathbf{V}^T$, where \mathbf{V} is an $m \times r$ orthogonal matrix. It follows that $\mathbf{X} \hat{\mathbf{B}} = \mathbf{X} \hat{\mathbf{B}} \mathbf{V} \mathbf{V}^T$. Therefore, for the new design matrix \mathbf{Z} defined by

$$(2.19) \quad \text{Type-II (Post-Decorrelation): } \mathbf{Z} \triangleq \mathbf{X}(\hat{\mathbf{B}} \mathbf{V}) \in \mathbb{R}^{n \times r},$$

each column (z -predictor) can be represented as a linear combination of the columns of \mathbf{X} , and the r newly obtained z -predictors are uncorrelated with each other, i.e., $\mathbf{Z}^T \mathbf{Z}$ is diagonal. We refer to this as *Type-II extraction* or *post-decorrelation*. Type-I and Type-II are not equivalent in general (but coincide for the RRR estimate). The (linear) feature extraction is supervised and the corresponding dimension reduction can be dramatic when r is much smaller than p .

Initial point When the problem (2.10) is convex, we can further show (based on Theorem 2.1) that any Θ^σ -estimate is a global minimum point. In this case, the choice of the initial estimate $\mathbf{B}^{(0)}$ is not essential and a pathwise algorithm with warm starts can be used in computing the solution path $\hat{\mathbf{B}}(\lambda)$ for a series of values of λ . However, for nonconvex problems, we do not have such global optimality given any initial point \mathbf{B}^0 . Although one can try multiple random starts, we found that empirically, simply setting $\mathbf{B}^{(0)}$ to be the zero matrix leads to a solution with good statistical performance. Intuitively, such initialization looks for a local optimum that is close to zero. Of course, other initialization choices are possible.

Parameter tuning The challenge still comes from nonconvexity. Take the rank penalty as an example. The solution path $\hat{\mathbf{B}}(\lambda)$ is *discontinuous*, while the optimal penalty parameter λ (as a surrogate for the Lagrange multiplier in convex programmings) is a function of both the data (\mathbf{X}, \mathbf{Y}) and the true coefficient \mathbf{B} . Therefore, plain cross-validation with respect to λ does not seem to be appropriate, as slightly perturbed data may result in serious regularization parameter mismatches. We propose to cross-validate the *range space* of the low rank estimator (as a function of λ) and refer to it as the projective cross-validation (**PCV**). In the following, we focus on the rank-Frobenius penalty (2.8) and the associated hard-ridge thresholding rule (2.6) to describe the idea. Let $\hat{\mathbf{B}}$ be a Θ_{HR}^σ estimator obtained from (2.11) and write it as $\begin{bmatrix} \hat{\beta}_0^T \\ \hat{\mathbf{B}}^\circ \end{bmatrix}$ following our previous notation (cf. (2.9)). For the penalized part $\hat{\mathbf{B}}^\circ$, denote its rank by r and its SVD by $\hat{\mathbf{B}}^\circ = \mathbf{U} \mathbf{D} \mathbf{V}^T$ with $\mathbf{D} \in \mathbb{R}^{r \times r}$. Construct r new features (Type-I) $\mathbf{Z}^\circ = \mathbf{X}^\circ \mathbf{U}$ and set $\mathbf{Z} = [\tilde{\mathbf{x}}_0, \mathbf{Z}^\circ] = [\mathbf{z}_1, \dots, \mathbf{z}_n]^T$. Let $\hat{\mathbf{C}}^\circ = \mathbf{D} \mathbf{V}^T$ and $\hat{\mathbf{C}} = [\hat{\beta}_0, \hat{\mathbf{C}}^\circ]^T = [\mathbf{c}_1, \dots, \mathbf{c}_m] \in \mathbb{R}^{(r+1) \times m}$.

Proposition 2.4. *Under the condition on ρ given in Theorem 2.1, for any Θ_{HR}^σ -estimator $\hat{\mathbf{B}}, \hat{\mathbf{C}}$ defined above is a Frobenius penalized estimator associated with new model matrix \mathbf{Z} , i.e.,*

$$(2.20) \quad \hat{\mathbf{C}} \in \arg \min_{\mathbf{C} \in \mathbb{R}^{(r+1) \times m}} - \sum_{k=1}^m \sum_{i=1}^n L_{i,k}(\mathbf{c}_k; \mathbf{z}_i, y_{ik}) + \frac{\eta}{2} \|\mathbf{C}^\circ\|_F^2.$$

See its proof in Appendix E. When $\eta > 0$ or \mathbf{Z} has full column rank, the optimization problem (2.20) is strictly convex and so $\hat{\mathbf{C}}$ is the unique optimal solution.

The proposition implies that once \mathbf{U} is extracted, we can simply use maximum likelihood estimation on the projected data to obtain the rank penalized estimator, or ridge penalized estimation to obtain the rank-Frobenius penalized estimator. There is no need to run the more expensive reduced rank fitting algorithms.

We now state the K -fold PCV procedure for tuning the rank penalty parameter in (1.1).

1. Run Algorithm (2.11) on the *whole* dataset for a grid of values for λ . The solution path is denoted by $\hat{\mathbf{B}}(\lambda_l)$, $l = 1, \dots, L$.
2. Obtain L candidate models via (2.18), each with a new model matrix $\mathbf{Z}(l) = \mathbf{X} \mathbf{U}(l)$, $1 \leq l \leq L$.
3. Compute the cross-validation error for each model. Concretely, partition the sample index set into K (roughly) even subsets $\mathcal{T}_1, \dots, \mathcal{T}_K$. Given $\mathbf{Z}(l)$, fit a vector GLM on the data without the subset indexed by \mathcal{T}_k , and evaluate its validation error (measured by deviance) on the left-out subset. In all, K maximum likelihood estimates are obtained and their validation errors are summed up to yield the CV error of the candidate model $\mathbf{Z}(l)$. Repeat this for all $l : 1 \leq l \leq L$.
4. Find the optimal model that minimizes the CV error.

In the pursuit of a parsimonious model with very low rank, a BIC penalty term can be added to the CV error (She, 2012). This is necessary in the large- p setup (Chen and Chen, 2008).

PCV is much more efficient than CV because the more involved reduced rank fitting algorithm runs only once beforehand, rather than K times in the CV trainings. The ML fitting in Step 3, justified by Proposition 2.4, involves very few predictors. Another benefit of PCV is that the parameter mismatch issue is eliminated and all K trainings are regarding the same model and feature space.

When there is an additional ridge parameter (cf. (2.8)), the procedure still applies, but a two-dimensional grid for (λ, η) has to be used. Fortunately, according to our experience, the statistical performance is not very sensitive to small changes in the ridge parameter and we can choose a sparse grid for it. Step 3 now fits a series of l_2 -penalized GLMs. But again, this type of problems is smooth and convex; Newton-based algorithms are reasonably fast. Finally, we mention that PCV shares some similarities to the selective cross-validation (SCV) proposed for variable selection (She, 2012).

3. RANK CONSTRAINED VECTOR GLMS FOR FEATURE SPACE REDUCTION

In this section, we study the reduced rank GLMs in constraint form (cf. (1.2)). For any $r \geq 1$ and $\eta \geq 0$, the problem

of interest is

$$(3.1) \quad \min_{\mathbf{B} \in \mathbb{R}^{(p+1) \times m}} - \sum_{k=1}^m \sum_{i=1}^n L_{i,k}(\boldsymbol{\beta}_k; \mathbf{x}_i, y_{ik}) + \frac{\eta}{2} \|\mathbf{B}^\circ\|_F^2$$

s.t. $\text{rank}(\mathbf{B}^\circ) \leq r.$

The additional Frobenius norm penalty is to handle collinearity. Again, neither the penalty nor the constraint is imposed on the first row of \mathbf{B} .

We introduce a *quantile thresholding rule* $\Theta^\#(\cdot; r, \eta)$ as a variant of the hard-ridge thresholding. Given $1 \leq r \leq p$ and $\eta \geq 0$, $\Theta^\#(\mathbf{a}; r, \eta) : \mathbb{R}^p \rightarrow \mathbb{R}^p$ is defined for any $\mathbf{a} \in \mathbb{R}^p$ such that the r largest components of \mathbf{a} (in absolute value) are shrunk by a factor of $(1 + \eta)$ and the remaining components are all set to be zero. In the case of ties, a random tie breaking rule is used. The matrix version of $\Theta^\#$ is defined as

$$\Theta^{\#\sigma}(\mathbf{B}; \lambda) \triangleq \mathbf{U} \text{diag}\{\Theta^\#([\sigma_i^{(\mathbf{B})}]; r, \eta)\} \mathbf{V}^T, \quad \forall \mathbf{B} \in \mathbb{R}^{p \times m}$$

where \mathbf{U} , \mathbf{V} , and $\text{diag}\{\sigma_i^{(\mathbf{B})}\}$ are obtained from the SVD of \mathbf{B} : $\mathbf{B} = \mathbf{U}^T \text{diag}\{\sigma_i^{(\mathbf{B})}\} \mathbf{V}$.

Then, a simple procedure similar to (2.11) can be used to solve (3.1): given $\mathbf{B}^{(j)}$, perform the update

$$(3.2) \quad \begin{cases} \mathbf{B}^{\circ(j+1)} = \Theta^{\#\sigma}(\mathbf{B}^{\circ(j)} + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\mathbf{B}^{(j)}); r, \eta), \\ \tilde{\boldsymbol{\beta}}_0^{(j+1)} = \tilde{\boldsymbol{\beta}}_0^{(j)} + (\mathbf{Y} - \boldsymbol{\mu}(\mathbf{B}^{(j)}))^T \tilde{\mathbf{x}}_0. \end{cases}$$

Starting with any $\mathbf{B}^{(0)} \in \mathbb{R}^{(p+1) \times m}$, denote the sequence obtained from (3.2) by $\{\mathbf{B}^{(j)}\}$. Let F be the objective function (3.1). Define ρ as in Theorem 2.1.

Theorem 3.1. *If $\rho \leq 1$, $F(\mathbf{B}^{(j)})$ is decreasing and satisfies*

$$F(\mathbf{B}^{(j)}) - F(\mathbf{B}^{(j+1)}) \geq (1 - \rho) \|\mathbf{B}^{(j)} - \mathbf{B}^{(j+1)}\|_F^2 / 2,$$

and $\text{rank}(\mathbf{B}^{\circ(j)}) \leq r, \forall j \geq 1.$

See Appendix F for its proof. The preliminary scaling of \mathbf{X}/k_0 for any $k_0 \geq \sqrt{p}$ guarantees the convergence. Still, PCV can be used for parameter tuning and model selection although the obtained estimate may not be globally optimal due to nonconvexity.

Rank penalty vs. rank constraint We have developed algorithms (2.11) and (3.2) for solving (1.1) (or (2.10)) and (1.2) (or (3.1)), respectively. The obtained estimates are usually local optimizers of the corresponding objective functions. However, in non-Gaussian GLM setups, we found that the nonconvexity of either problem can be very strong. For instance, there may exist many local optima all having the same rank but spanning different subspaces in \mathbb{R}^p . In consideration of this, the penalized solution path $\hat{\mathbf{B}}(\lambda)$ ($0 \leq \lambda < +\infty$) may provide more candidate models of certain rank (if existing) than the constrained solution path $\hat{\mathbf{B}}(r)$ ($r = 1, 2, \dots, p \wedge m$), which is advantageous in the

stage of parameter tuning. This phenomenon is often observed in datasets where p is comparable to or larger than n . Note that typically the direct rank penalized $\hat{\mathbf{B}}(\lambda)$ has no rank monotonicity.

On the other hand, computing the solution path for the penalty form is often more time-consuming in large- p applications. The path $\hat{\mathbf{B}}(\lambda)$ has jumps. Assuming no prior knowledge of the appropriate interval for λ , one has to specify a large search grid fine enough to cover a reasonable number of candidate models. By contrast, for the problem of constraint form, we can set a small upper bound for r in pursuing a low rank model (say, $r \leq 0.5n \wedge p \wedge m$ could be good), and the natural grid spacing is 1. With the grid focusing on small values of r (which amounts to applying large thresholds in the iteration steps), Algorithm (3.2) runs efficiently.

Feature space reduction To combine the virtues of both approaches, we propose to solve the rank constrained problem to perform feature space reduction, and then run the rank penalized algorithm in the reduced feature space. This is very helpful in large- p applications. A crude sketch is as follows. First, we set $r = \alpha n \wedge p \wedge m$ with $\alpha < 1$ (e.g., $\alpha = 0.5$) and solve (1.2). Using the estimate $\hat{\mathbf{B}}(r)$, we execute Type-I feature extraction (2.18) to construct a new model matrix $\mathbf{Z} = \mathbf{X} \mathbf{U}_1(r)$ with only r factors (in addition to the intercept). Next, we turn to the penalized problem (2.10) on (\mathbf{Y}, \mathbf{Z}) : Get the solution path from running Algorithm (2.11), and tune the regularization parameters to find the optimal estimate, denoted by $\hat{\mathbf{B}}'(\lambda_o)$. Our final coefficient matrix estimate is given by $\mathbf{U}_1(r) \hat{\mathbf{B}}'(\lambda_o)$. A small number of new predictive features can be constructed (and decorrelated) based on (2.19).

According to this scheme, the sample size of the reduced problem on \mathbf{Z} is large relative to the reduced dimension. It is not difficult to show that for $n > p$, the update in (2.11) is essentially a contraction, and so Algorithm (2.11) converges fast.

A crucial assumption here is that the rank of the true model, denoted by r^* , is very small, compared with the sample size n . This makes it possible to choose a safe rank constraint value r in (1.2), which, though possibly much less than p , is still much larger than the true r^* . Hence the computational cost of obtaining a solution path according to Algorithm (2.11) can be effectively reduced with little performance loss. This idea shares similarity with the variable screening (Fan and Lv, 2008) proposed in the context of sparse variable selection. In the process of screening, all relevant variables should be kept, while in feature space reduction, only the necessary factors, being linear combinations of the original predictors and typically as few as a handful, are required to lie in reduced feature space we project \mathbf{X} onto.

In implementation, we further adopt a path-following (annealing) idea to reduce computational load and avoid greedy reduction. Define a cooling schedule $r(t)$ ($0 \leq t \leq T$) with $r(0) = p$ and $r(T) = r$, where r is an upper bound of

Table 2. Rank constraint vs. rank penalty. Misclassification rates of the constrained and penalized reduced rank logistic regressions ($RRL^{(c)}$ and $RRL^{(p)}$) are shown for the zipcode (sub)dataset where $p = 257, n = 300$. The rank r controls the # of newly constructed features

r	1	2	3	4
$RRL^{(c)}$	66.52%	55.06%	38.47%	33.83%
$RRL^{(p)}$	66.52%	55.06%	38.32%, 38.37%	33.83%
$RRL^{(c)}$ +SVM	59.24%	47.48%	33.58%	30.79%
$RRL^{(p)}$ +SVM	59.24%	47.48%	33.63%, 33.58%	30.79%
5	6	7	8	9
24.86%	22.27%	21.33%	21.33%	20.43%
24.81%, 24.86%	22.42%	21.47%, 21.33%	21.33%	20.43%, 19.13%
23.02%	20.38%	20.43%	20.28%	18.53%
23.02%	21.08%	20.33%, 20.43%	20.33%, 20.28%	18.53%, 15.84%

the target rank. We conduct *progressive* feature space reduction as follows. (As aforementioned, \mathbf{Z}° refers to \mathbf{Z} without the first column, \mathbf{B}° refers to \mathbf{B} without the first row, and \mathbf{U}° refers to \mathbf{U} without the first row and the first column.)

1. Let $t \leftarrow 0, \mathbf{Z} \leftarrow \mathbf{X}, \mathbf{U} \leftarrow \mathbf{I}$.
2. Iterate until $r(t) \leq r$:
 - (a) Set the rank constraint value to be $r(t)$ and perform the update (3.2) on (\mathbf{Y}, \mathbf{Z}) for at most M times (with M pre-specified);
 - (b) Obtain the left singular vectors of the current slope estimate \mathbf{B}° , denoted by $\mathbf{U}_1(r(t))$;
 - (c) Let $\mathbf{Z}^\circ \leftarrow \mathbf{Z}^\circ \mathbf{U}_1(r(t)), \mathbf{U}^\circ \leftarrow \mathbf{U}^\circ \mathbf{U}_1(r(t))$;
 - (d) $t \leftarrow t + 1$.

At the end, \mathbf{Z} is delivered as the new design, and the orthogonal matrix \mathbf{U} gives the accumulated transformation matrix.

The previously described prototype reduction scheme corresponds to $r(t) = r$ for any t . With an annealing algorithm design, the dimensionality of the feature space keeps dropping; the \mathbf{B} involved in (3.2) has only $r(t)$ columns. A slow cooling schedule with a small number of M is recommended. Based on our experiments, it is not too greedy and is usually computationally affordable for large- p problems.

4. DATA EXAMPLES

We use two real data examples to illustrate the proposed methodology for dimension reduction and supervised feature extraction.

Example 4.1. First, we make a practical comparison of the rank penalized estimators from solving (1.1) and the rank constrained estimators from solving (1.2) by use of a zip-code dataset. The whole dataset (available at the website of [Hastie et al. \(2001\)](#)) contains normalized handwritten digits in 16×16 grayscale images. The digits were originally scanned from envelopes by the U.S. Postal Service and have

been deslanted and size normalized. The space of pixel predictors is of dimension 256. We standardized all such predictors. The intercept term is included in the model and is always unpenalized. We introduced $m = 9$ indicator response variables for digits 0–8, using 9 as the reference class.

The training set is large in comparison with p and m (7,291 images). We chose a subset of $n = 300$ at random in this experiment to compare the penalized solution path and the constrained solution path. No additional Frobenius-norm penalty was enforced. The prediction results of the estimates are shown in Table 2, evaluated on 2,007 test observations.

From the table, at certain values of r , the rank penalty offered more candidate models along its solution path than the rank constraint. Note that these rank- r estimators may behave differently in prediction and feature extraction. For $p \sim n$ or $p > n$, this phenomenon is commonly seen. With an appropriate parameter tuning strategy, the penalty form gives better chances to achieve a low error rate.

Of course, this comes with a price in computation. In our experiment, the time for obtaining the $RRL^{(c)}$ path was less than one minute, while computing the $RRL^{(p)}$ path, with a 50-point grid for λ , took about four minutes.

There is no obligation to predict through the obtained estimator; perhaps more useful is the much lower dimensional feature space yielded from such an estimator. Fancier classifiers such as SVM can be applied with the new features automatically extracted and decorrelated via (2.19), and result in lower error rates as shown in the table.

Finally, we add a comment that in some situations there may exist no penalized solution at certain rank values. Yet with a large λ -grid chosen, the performance of the penalized estimator (after parameter tuning) does not seem to be worse than that of the constrained estimator.

Example 4.2. The *Computer Audition Lab 500* (CAL500) dataset is collected by [Turnbull et al. \(2008\)](#) and involves 502 Western popular songs by different artists selected from

the past 50 years. Digital audio files were played to students to annotate these songs with $m = 174$ words representing emotion, genre, instrument, vocals, etc. The concepts characterized by the words are not mutually exclusive and one song can be annotated with multiple labels. This is called *multi-label* data in machine learning. The predictors are MFCC-Delta audio features from analyzing a short-time segment of the audio signal. [Turnbull et al. \(2008\)](#) used 68 such feature vectors. To allow for interactions between these audio features and to make a more challenging problem, we consider a full quadratic model including all main effects, quadratic effects, and pairwise interactions. Hence $p = 68 + 68(69)/2 + 1 = 2,415$. We split the data into two halves and used $n = 251$ songs for training and the other 251 for testing.

For this small-sample-size-high-dimensional problem, the SVM using all 2,415 predictors gave a total misclassification rate of 21.2%, which is not all bad. On the other hand, the proposed reduced rank methodology can be applied to automatically construct new predictive audio features, possibly much fewer than 2,415. The supervised nature is important because only the audio features helpful in annotation (classification) are truly meaningful in this learning task.

First, we conducted the progressive feature space reduction introduced in Section 3, with the upper bound of the target rank set to be 20. Then we ran Algorithm (2.11) to fit a penalized reduced rank vector logistic regression with the 20 extracted features. The rank-Frobenius penalty was chosen due to serious collinearity arising from the high-dimensional quadratic model. The parameters were tuned by 5-fold PCV with BIC correction.

Surprisingly, our final estimate $\hat{\mathbf{B}}$ has $\text{rank}(\hat{\mathbf{B}}^\circ) = 2$, which gives a dramatic dimension reduction from 2,514 to 2. But the SVM trained based on just the two new features yielded an improved error rate of 14.13%. In fact, even using the vanilla reduced rank estimator, we can achieve an error rate of 14.36%. The per-word precision and recall (cf. [Turnbull et al. \(2008\)](#) for the detailed definitions) are, respectively, 35.6% and 8.7% on the test dataset, comparable to the rates of the two advocated approaches in [Turnbull et al. \(2008\)](#). But our model is more parsimonious and creates two concise audio summary indexes for semantic annotation.

5. CONCLUSION

Supervised linear feature extraction can be obtained from a reduced rank vector model. We studied rank penalized and rank constrained generalized linear models and discussed how to adapt them to feature extraction and feature space reduction. The latter technique helps to reduce the computational cost significantly in high dimensions. We also noticed the strong nonconvexity of such problems raises some serious issues in data-resampling based parameter tunings, but the proposed projective cross-validation works decently

in general and is efficient. Through reduced rank modeling, dimension reduction can be attained if the rank of the model is small relative to the number of predictors. The work can be viewed as a supervised and parametric generalization of the principle component analysis.

APPENDIX A. PROOF OF PROPOSITION 2.1

To prove Proposition 2.1, we first introduce two lemmas.

Lemma A.1 ([von Neumann \(1937\)](#)). *Let \mathbf{A}, \mathbf{B} be two $n \times n$ matrices. Then*

$$(A.1) \quad |\text{Tr}(\mathbf{AB})| \leq \sum \sigma_i(\mathbf{A})\sigma_i(\mathbf{B}),$$

where $\sigma_1(\mathbf{A}) \geq \sigma_2(\mathbf{A}) \geq \dots \geq \sigma_n(\mathbf{A})$ and $\sigma_1(\mathbf{B}) \geq \sigma_2(\mathbf{B}) \geq \dots \geq \sigma_n(\mathbf{B})$ are ordered singular values of \mathbf{A} and \mathbf{B} respectively.

We refer to [Grigorieff \(1991\)](#) for an elementary proof.

Lemma A.2. *Given a thresholding rule Θ , let P be any penalty satisfying condition (2.5) in Proposition 2.1. Then, the univariate minimization problem $\min_{\theta} (t - \theta)^2 / 2 + P(\theta; \lambda)$ has a unique optimal solution $\hat{\theta} = \Theta(t; \lambda)$ for every t at which $\Theta(\cdot; \lambda)$ is continuous.*

Proof. Apply Lemma 1 in [She \(2012\)](#). \square

Proof of the optimality part of Proposition 2.1. Let $\mathbf{Y} \in \mathbb{R}^{n \times m}$ and assume $n \geq m$ without any loss of generality. Let $\mathbf{Y} = \mathbf{U}_0 \mathbf{D}_0 \mathbf{V}_0^T$ and $\mathbf{B} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ be the SVDs where $\mathbf{D}_0 = \text{diag}(d_{0,i})$ and $\mathbf{D} = \text{diag}(d_i)$ with $d_{0,1} \geq d_{0,2} \geq \dots \geq d_{0,m}$ and $d_1 \geq d_2 \geq \dots \geq d_m$. Clearly,

$$\begin{aligned} \|\mathbf{Y} - \mathbf{B}\|_F^2 &= \|\mathbf{Y}\|_F^2 + \|\mathbf{B}\|_F^2 - 2\text{Tr}(\mathbf{Y}^T \mathbf{B}) \\ &= \|\mathbf{Y}\|_F^2 + \|\mathbf{B}\|_F^2 - 2\text{Tr}([\mathbf{Y} \ \mathbf{0}]^T [\mathbf{B} \ \mathbf{0}]), \end{aligned}$$

where $[\mathbf{B} \ \mathbf{0}] \in \mathbb{R}^{n \times n}$ and $[\mathbf{Y} \ \mathbf{0}] \in \mathbb{R}^{n \times n}$. It follows from Lemma A.1 that $\text{Tr}(\mathbf{Y}^T \mathbf{B}) \leq \sum d_{0,i} d_i$. Hence

$$(A.2) \quad \begin{aligned} F(\mathbf{B}) &\geq (\|\mathbf{D}_0\|_F^2 + \|\mathbf{D}\|_F^2 - 2\text{Tr}(\mathbf{D}_0 \mathbf{D})) / 2 + \sum P(d_i; \lambda) \\ &= \sum (d_{0,i} - d_i)^2 / 2 + \sum P(d_i; \lambda). \end{aligned}$$

Now the problem reduces to

$$\min_{d_i} \sum (d_{0,i} - d_i)^2 / 2 + \sum P(d_i; \lambda).$$

The optimal solution $\hat{\mathbf{B}}$ then follows from Lemma A.2. \square

The argument above only implies the singular values of $\hat{\mathbf{B}}$ are unique (up to permutation). Although one can possibly argue the uniqueness of $\hat{\mathbf{B}}$ by studying the condition under which equality is achieved in (A.2), another formal proof of the uniqueness is deferred to Appendix B.

APPENDIX B. PROOF OF PROPOSITION 2.2

Let $\mathbf{B} = \hat{\mathbf{B}} + \mathbf{\Delta}$. Suppose $\mathbf{Y} = \mathbf{U}_0 \mathbf{D}_0 \mathbf{V}_0^T$, $\hat{\mathbf{B}} = \mathbf{U}_0 \hat{\mathbf{D}} \mathbf{V}_0^T$, and $\mathbf{B} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ are the SVDs. We have

$$\begin{aligned} & \|\mathbf{Y} - \mathbf{B}\|_F^2/2 - \|\mathbf{D}_0 - \mathbf{D}\|_F^2/2 \\ &= -\text{Tr}(\mathbf{B}^T \mathbf{Y}) + \text{Tr}(\mathbf{D}_0 \mathbf{D}) \\ &= -\text{Tr}(\mathbf{B}^T (\mathbf{Y} - \hat{\mathbf{B}})) + \text{Tr}((\mathbf{D}_0 - \hat{\mathbf{D}}) \mathbf{D}) \\ &\quad + \text{Tr}(\mathbf{D} \hat{\mathbf{D}}) - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}}) \\ &= -\text{Tr}(\mathbf{V} \mathbf{D} \mathbf{U}^T \mathbf{U}_0 (\mathbf{D}_0 - \hat{\mathbf{D}}) \mathbf{V}_0^T) + \text{Tr}((\mathbf{D}_0 - \hat{\mathbf{D}}) \mathbf{D}) \\ &\quad + \text{Tr}(\mathbf{D} \hat{\mathbf{D}}) - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}}). \end{aligned}$$

By Proposition 2.1 $\hat{\mathbf{D}} \preceq \mathbf{D}_0$, i.e., $\mathbf{D}_0 - \hat{\mathbf{D}}$ is positive semi-definite. By augmenting $\mathbf{Y} - \mathbf{B}$ and \mathbf{B} and applying Lemma A.1, we can prove

$$\text{Tr}(\mathbf{D}(\mathbf{D}_0 - \hat{\mathbf{D}})) \geq \text{Tr}(\mathbf{V} \mathbf{D} \mathbf{U}^T \mathbf{U}_0 (\mathbf{D}_0 - \hat{\mathbf{D}}) \mathbf{V}_0^T),$$

from which it follows that

$$\begin{aligned} & \|\mathbf{Y} - \mathbf{B}\|_F^2/2 - \|\mathbf{D}_0 - \mathbf{D}\|_F^2/2 \\ & \geq \text{Tr}(\mathbf{D} \hat{\mathbf{D}}) - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}}) \geq C_1 (\text{Tr}(\mathbf{D} \hat{\mathbf{D}}) - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}})). \end{aligned}$$

Now we have

$$\begin{aligned} & Q(\mathbf{B}) - Q(\hat{\mathbf{B}}) \\ &= \|\mathbf{Y} - \mathbf{B}\|_F^2/2 - \|\mathbf{Y} - \hat{\mathbf{B}}\|_F^2/2 \\ &\quad + \sum P_\Theta(d_i; \lambda) - \sum P_\Theta(\hat{d}_i; \lambda) \\ &\geq \|\mathbf{D}_0 - \mathbf{D}\|_F^2/2 - \|\mathbf{D}_0 - \hat{\mathbf{D}}\|_F^2/2 \\ &\quad + \sum P_\Theta(d_i; \lambda) - \sum P_\Theta(\hat{d}_i; \lambda) + C_1 \text{Tr}(\mathbf{D} \hat{\mathbf{D}} - \mathbf{B}^T \hat{\mathbf{B}}) \\ &= \sum ((d_{0,i} - d_i)^2/2 + P_\Theta(d_i; \lambda)) \\ &\quad - ((d_{0,i} - \hat{d}_i)^2/2 + P_\Theta(\hat{d}_i; \lambda)) + C_1 \text{Tr}(\mathbf{D} \hat{\mathbf{D}} - \mathbf{B}^T \hat{\mathbf{B}}) \\ &\geq C_1 \sum (d_i - \hat{d}_i)^2/2 + C_1 \text{Tr}(\mathbf{D} \hat{\mathbf{D}} - \mathbf{B}^T \hat{\mathbf{B}}) \\ &= C_1 (\|\mathbf{D} - \hat{\mathbf{D}}\|_F^2/2 + \text{Tr}(\mathbf{D} \hat{\mathbf{D}}) - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}})) \\ &= C_1 (\|\mathbf{D}\|_F^2/2 + \|\hat{\mathbf{D}}\|_F^2/2 - \text{Tr}(\mathbf{B}^T \hat{\mathbf{B}})) \\ &= C_1 \|\mathbf{B} - \hat{\mathbf{B}}\|_F^2/2. \end{aligned}$$

The second inequality is due to Lemma 2 in She (2012).

Proof of the optimality part of Proposition 2.1. From the comment in Appendix A, any optimal solution \mathbf{B} must have the same nonzero singular values (up to permutation) as $\hat{\mathbf{B}}$, i.e., $d_i = \hat{d}_i$, seen from the proof of Proposition 2.1. A more careful examination of the proof of Proposition 2.2 shows $Q(\mathbf{B}) - Q(\hat{\mathbf{B}}) \geq \text{Tr}(\mathbf{D} \hat{\mathbf{D}} - \mathbf{B}^T \hat{\mathbf{B}}) = \|\mathbf{B} - \hat{\mathbf{B}}\|_F^2/2$. Therefore, the globally optimal solution $\hat{\mathbf{B}}$ in Proposition 2.1 must be unique. \square

APPENDIX C. PROOF OF THEOREM 2.1

The proof is similar to that of Theorem 2.1 in She (2012). Define a surrogate function G for any $\mathbf{A} = [\alpha_1, \dots, \alpha_m] = [\tilde{\alpha}_0, \tilde{\alpha}_1, \dots, \tilde{\alpha}_p]^T$ and $\mathbf{B} = [\beta_1, \dots, \beta_m] \in \mathbb{R}^{(p+1) \times m}$

$$\begin{aligned} & G(\mathbf{B}, \mathbf{A}) \\ &= -\sum_{k=1}^m \sum_{i=1}^n L_{i,k}(\alpha_k) + \sum_{s=1}^{p \wedge m} P(\sigma_s^{(\mathbf{A}^\circ)}; \lambda) \\ &\quad + \frac{1}{2} \|\mathbf{A} - \mathbf{B}\|_F^2 - \sum_{k=1}^m \sum_{i=1}^n (b(\mathbf{x}_i^T \alpha_k) - b(\mathbf{x}_i^T \beta_k)) \\ &\quad + \sum_{k=1}^m \sum_{i=1}^n \mu_{i,k} (\mathbf{x}_i^T \alpha_k - \mathbf{x}_i^T \beta_k), \end{aligned}$$

where $\mu_{i,k} = g^{-1}(\mathbf{x}_i^T \beta_k) = b'(\mathbf{x}_i^T \beta_k)$. It can be shown that given \mathbf{B} , minimizing G over \mathbf{A} is equivalent to

$$\begin{aligned} & \arg \min_{\mathbf{A}} \frac{1}{2} \left\| \mathbf{A} - \left[\mathbf{B} + \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \boldsymbol{\mu}(\mathbf{B}) \right] \right\|_F^2 \\ &\quad + \sum_{s=1}^{p \wedge m} P(\sigma_s^{(\mathbf{A}^\circ)}; \lambda). \end{aligned}$$

By Proposition 2.1, $\mathbf{B}^{(j+1)}$ in (2.11) can be characterized by $\arg \min_{\mathbf{A}} G(\mathbf{B}^{(j)}, \mathbf{A})$. Furthermore, we have for any $\mathbf{\Delta} \in \mathbb{R}^{(p+1) \times m}$

(C.1)

$$G(\mathbf{B}^{(j)}, \mathbf{B}^{(j+1)} + \mathbf{\Delta}) - G(\mathbf{B}^{(j)}, \mathbf{B}^{(j+1)}) \geq \frac{C_1}{2} \|\mathbf{\Delta}\|_F^2$$

with $C_1 = \max(0, 1 - \mathcal{L}_\Theta)$, by applying Proposition 2.2 and Lemma 1 in She (2012), and noting that $q_s(\sigma^{(\mathbf{B}^{(j+1)})}) = 0$, for $\mathbf{B}^{(j+1)}$ obtained by Θ^σ -thresholding.

Next, Taylor series expansion gives

$$\begin{aligned} & F(\mathbf{B}^{(j+1)}) \\ &= \sum_k \frac{1}{2} (\beta_k^{(j+1)} - \beta_k^{(j)})^T (\mathbf{I} - \mathcal{I}(\boldsymbol{\xi}_k^{(j)})) (\beta_k^{(j+1)} - \beta_k^{(j)}) \\ &= G(\mathbf{B}^{(j)}, \mathbf{B}^{(j+1)}) \\ &\leq G(\mathbf{B}^{(j)}, \mathbf{B}^{(j)}) \\ &\quad - \sum_k \frac{C_1}{2} (\beta_k^{(j+1)} - \beta_k^{(j)})^T (\beta_k^{(j+1)} - \beta_k^{(j)}) \\ &= F(\mathbf{B}^{(j)}) - \frac{C_1}{2} \|\mathbf{B}^{(j+1)} - \mathbf{B}^{(j)}\|_F^2 \end{aligned}$$

(2.12) can be obtained. In fact, this decreasing property holds for any $\rho \leq 2 - \mathcal{L}_\Theta$.

Let $\mathbf{B}^{(j_l)} \rightarrow \mathbf{B}^*$ as $l \rightarrow \infty$. Under the condition $\rho \leq 2 - \mathcal{L}_\Theta$, C is strictly positive and $\|\mathbf{B}^{(j_{l+1})} - \mathbf{B}^{(j_l)}\|_F^2/2 \leq (F(\mathbf{B}^{(j_l)}) - F(\mathbf{B}^{(j_{l+1})}))/C \leq (F(\mathbf{B}^{(j_l)}) - F(\mathbf{B}^{(j_{k+1})}))/C \rightarrow 0$. That is, $\Theta^\sigma(\mathbf{B}^{(j_l)} + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\mathbf{B}^{(j_l)}); \lambda) - \mathbf{B}^{(j_l)} \rightarrow 0$. Therefore, \mathbf{B}^* is a solution to (2.13) due to the continuity assumption.

APPENDIX D. PROOF OF PROPOSITION 2.3

Let $M = (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{X}^T \mathbf{Y}$ and $r_0 = \text{rank}(M)$. Obviously, $r_0 \leq p \wedge m$. Note that $M^T M = \mathbf{Y}^T \mathbf{H} \mathbf{Y}$. Assume $M = \mathbf{U} \mathbf{D} \mathbf{V}^T$ is the SVD of M with $\mathbf{U} \in \mathbb{R}^{p \times r_0}$, $\mathbf{V} \in \mathbb{R}^{m \times r_0}$, and $\mathbf{D} \in \mathbb{R}^{r_0 \times r_0}$. Suppose all (positive) diagonal entries of \mathbf{D} are arranged in decreasing order. Let $\mathbf{A} \triangleq \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \mathbf{X} \hat{\mathbf{B}}$. To prove $\hat{\mathbf{B}}$ obeys the Θ^σ -equation (2.13) for hard-thresholding, it suffices to show that (i) there exists a $p \times r_0$ orthogonal matrix \mathbf{U}_* satisfying $\mathbf{U}_*^T \mathbf{U}_* = \mathbf{I}$ such that $\mathbf{U}_*^T (\hat{\mathbf{B}} \hat{\mathbf{B}}^T) \mathbf{U}_*$ and $\mathbf{U}_*^T (\mathbf{A} \mathbf{A}^T) \mathbf{U}_*$ are both diagonal; (ii) there exists an $m \times r_0$ orthogonal matrix \mathbf{V}_* such that $\mathbf{V}_*^T (\hat{\mathbf{B}}^T \hat{\mathbf{B}}) \mathbf{V}_*$ and $\mathbf{V}_*^T (\mathbf{A}^T \mathbf{A}) \mathbf{V}_*$ are both diagonal; (iii) $\text{Tr}(\hat{\mathbf{B}}^T \mathbf{A}) = 0$; (iv) the singular values of \mathbf{A} are all bounded by λ .

Recall that $r = \max\{i : d_i \geq \lambda\}$ and $\mathbf{V}_r = \mathbf{V}[:, 1:r]$. Introduce $\mathbf{V}_{-r} = \mathbf{V}[:, (r+1):r_0]$, formed by deleting the first r columns in \mathbf{V} . Then we have

$$\begin{aligned} \text{(D.1)} \quad \mathbf{A} &= \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \mathbf{X} \hat{\mathbf{B}} = \mathbf{X}^T \mathbf{Y} - \mathbf{X}^T \mathbf{H} \mathbf{Y} \mathcal{P}_{\mathbf{V}_r} \\ &= \mathbf{X}^T \mathbf{Y} (\mathbf{I} - \mathcal{P}_{\mathbf{V}_r}) = \mathbf{X}^T \mathbf{Y} \mathcal{P}_{\mathbf{V}_{-r}} \\ &= \mathbf{X}^T \mathbf{Y} \mathbf{V}_{-r} \mathbf{V}_{-r}^T = (\mathbf{X}^T \mathbf{X})^{1/2} \mathbf{M} \mathbf{V}_{-r} \mathbf{V}_{-r}^T. \end{aligned}$$

Obviously, $\text{Tr}(\hat{\mathbf{B}}^T \mathbf{A}) = 0$. (iii) is true. On the other hand, we can rewrite $\hat{\mathbf{B}}$ as

$$\begin{aligned} \text{(D.2)} \quad \hat{\mathbf{B}} &= (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{M} \mathbf{V}_r \mathbf{V}_r^T \\ &= (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{M} \mathbf{V} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{V}^T \\ &= (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{V}^T \\ \text{(D.3)} \quad &= (\mathbf{X}^T \mathbf{X})^{-1/2} \Theta_H^\sigma(\mathbf{M}; \lambda). \end{aligned}$$

Now we obtain

$$\text{(D.4)} \quad \hat{\mathbf{B}}^T \hat{\mathbf{B}} = \mathbf{V} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{D} \mathbf{U}^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{V}^T,$$

$$\text{(D.5)} \quad \mathbf{A}^T \mathbf{A} = \mathbf{V} \begin{bmatrix} \mathbf{0}_{r \times r} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{D} \mathbf{U}^T (\mathbf{X}^T \mathbf{X}) \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0}_{r \times r} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{V}^T.$$

(iv) is straightforward from (D.5):

$$\begin{aligned} \|\mathbf{A}^T \mathbf{A}\|_2 &\leq \|\mathbf{X}\|_2^2 \cdot \left\| \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0}_{r \times r} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{V}^T \right\|_2^2 \\ &\leq 1 \cdot d_{r+1}^2 \leq \lambda^2. \end{aligned}$$

(D.4) + (D.5) also implies (ii). In fact, introducing $\mathbf{G} = \mathbf{D} \mathbf{U}^T (\mathbf{X} \mathbf{X})^{-1} \mathbf{U} \mathbf{D}$, $\mathbf{H} = \mathbf{D} \mathbf{U}^T (\mathbf{X} \mathbf{X}) \mathbf{U} \mathbf{D}$, $\mathbf{G}_{11} = \mathbf{G}[1:r, 1:r]$, $\mathbf{H}_{22} = \mathbf{H}[(r+1):r_0, (r+1):r_0]$, and assuming the spectral decompositions of the two submatrices are given by $\mathbf{G}_{11} = \mathbf{U}_{11}^G \mathbf{D}_{11}^G (\mathbf{U}_{11}^G)^T$ and $\mathbf{H}_{22} = \mathbf{U}_{22}^H \mathbf{D}_{22}^H (\mathbf{U}_{22}^H)^T$, respectively, then,

$$\mathbf{V}_* = \mathbf{V} \begin{bmatrix} \mathbf{U}_{11}^G & \\ & \mathbf{U}_{22}^H \end{bmatrix}$$

simultaneously diagonalizes $\hat{\mathbf{B}}^T \hat{\mathbf{B}}$ and $\mathbf{A}^T \mathbf{A}$ and satisfies $\mathbf{V}_*^T \mathbf{V}_* = \mathbf{I}$.

Finally, we construct \mathbf{U}_* to prove (i). From (D.3) and (D.2),

$$\text{(D.6)} \quad \hat{\mathbf{B}} \mathbf{V}_* = (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{11}^G & \\ & \mathbf{U}_{22}^H \end{bmatrix},$$

$$\text{(D.7)} \quad \mathbf{A} \mathbf{V}_* = (\mathbf{X}^T \mathbf{X})^{1/2} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0}_{r \times r} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{11}^G & \\ & \mathbf{U}_{22}^H \end{bmatrix}.$$

Let $\tilde{\mathbf{G}} = (\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{U} \mathbf{D}$ and $\tilde{\mathbf{H}} = (\mathbf{X}^T \mathbf{X})^{1/2} \mathbf{U} \mathbf{D}$. Then $\tilde{\mathbf{G}}^T \tilde{\mathbf{G}} = \mathbf{G}$, $\tilde{\mathbf{H}}^T \tilde{\mathbf{H}} = \mathbf{H}$. By construction, \mathbf{U}_{11}^G and \mathbf{U}_{22}^H must be the right-singular vectors of $\tilde{\mathbf{G}}_1 = \tilde{\mathbf{G}}[:, 1:r]$ and $\tilde{\mathbf{H}}_2 = \tilde{\mathbf{H}}[:, (r+1):r_0]$ respectively. Denoting by $\mathbf{U}_1^{\tilde{\mathbf{G}}}$ and $\mathbf{U}_2^{\tilde{\mathbf{H}}}$ their associated left-singular vectors respectively, we get

$$\mathbf{U}_* = \begin{bmatrix} \mathbf{U}_1^{\tilde{\mathbf{G}}} & \mathbf{U}_2^{\tilde{\mathbf{H}}} \end{bmatrix},$$

which makes both $\mathbf{U}_*^T \hat{\mathbf{B}} \hat{\mathbf{B}}^T \mathbf{U}_*$ and $\mathbf{U}_*^T \mathbf{A} \mathbf{A}^T \mathbf{U}_*$ diagonal. To prove \mathbf{U}_* is the desired matrix in (i), it remains to show the orthogonality of \mathbf{U}_* . It follows from (D.6) and (D.7) that

$$\begin{aligned} &(\tilde{\mathbf{G}}_1 \mathbf{U}_{11}^G)^T \tilde{\mathbf{H}}_2 \mathbf{U}_{22}^H \\ &= (\mathbf{U}_{11}^G)^T \left((\mathbf{X}^T \mathbf{X})^{-1/2} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{I}_{r \times r} & \\ & \mathbf{0} \end{bmatrix} \right)^T \\ &\quad \cdot (\mathbf{X}^T \mathbf{X})^{1/2} \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{U}_{22}^H \\ &= (\mathbf{U}_{11}^G)^T \begin{bmatrix} \mathbf{I}_{r \times r} & \mathbf{0} \end{bmatrix} \mathbf{D} \mathbf{U}^T (\mathbf{X}^T \mathbf{X})^{-1/2} (\mathbf{X}^T \mathbf{X})^{1/2} \\ &\quad \cdot \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{U}_{22}^H \\ &= (\mathbf{U}_{11}^G)^T \begin{bmatrix} \mathbf{I}_{r \times r} & \mathbf{0} \end{bmatrix} \mathbf{D} \mathbf{U}^T \mathbf{U} \mathbf{D} \begin{bmatrix} \mathbf{0} & \\ & \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{U}_{22}^H \end{aligned}$$

$$\begin{aligned}
&= (\mathbf{U}_{11}^G)^T \begin{bmatrix} \mathbf{I}_{r \times r} & \mathbf{0} \end{bmatrix} \mathbf{D}^2 \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_{(r_0-r) \times (r_0-r)} \end{bmatrix} \mathbf{U}_{22}^H \\
&= \mathbf{0}.
\end{aligned}$$

Since \mathbf{G}_{11} and \mathbf{H}_{22} are positive definite (noting that $\mathbf{D} \in \mathbb{R}^{r_0 \times r_0}$ is nonsingular), we further obtain $(\mathbf{U}_1^{\hat{G}})^T \mathbf{U}_2^{\hat{H}} = \mathbf{0}$. Hence $\mathbf{U}_*^T \mathbf{U}_* = \mathbf{I}$. The proof is now complete.

APPENDIX E. PROOF OF PROPOSITION 2.4

From Theorem 2.1, $\hat{\mathbf{B}}$ satisfies

$$(E.1) \quad \begin{cases} \hat{\mathbf{B}}^\circ = \Theta_{HR}^\sigma(\hat{\mathbf{B}}^\circ + \mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X}); \lambda, \eta), \\ \mathbf{0} = (\mathbf{Y} - \boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X}))^T \hat{\mathbf{x}}_0. \end{cases}$$

Here, we write $\boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X})$ to emphasize the dependence of the mean matrix on the design. In this proof, we use the same submatrix notation as in Appendix D.

Given the SVD $\mathbf{B}^\circ = \mathbf{U} \mathbf{D} \mathbf{V}^T$, by Definition 2.2, there exist orthogonal matrices $\bar{\mathbf{U}}$ and $\bar{\mathbf{V}}$, as augmented versions of \mathbf{U} and \mathbf{V} , respectively, i.e., $\mathbf{U} = \bar{\mathbf{U}}[I,]$ and $\mathbf{V} = \bar{\mathbf{V}}[I,]$ for some index set I , such that $\hat{\mathbf{B}}^\circ = \bar{\mathbf{U}} \boldsymbol{\Sigma} \bar{\mathbf{V}}^T$ and $\mathbf{X}^{\circ T} \mathbf{Y} - \mathbf{X}^{\circ T} \boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X}) = \bar{\mathbf{U}} \mathbf{W} \bar{\mathbf{V}}^T$ are both the SVDs. Clearly, $\boldsymbol{\Sigma}[I, I] = \mathbf{D}$, $\boldsymbol{\Sigma}[I^c, I^c] = \mathbf{0}$. Using the hard-ridge thresholding (2.6), we rewrite the first equation in (E.1) as

$$(E.2) \quad (1 + \eta) \mathbf{B}^\circ + \lambda(1 + \eta) \bar{\mathbf{U}} \boldsymbol{\Sigma} \bar{\mathbf{V}}^T = \mathbf{B}^\circ + \mathbf{X}^{\circ T} (\mathbf{Y} - \boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X})),$$

where $\boldsymbol{\Sigma}$ is diagonal and satisfies $\boldsymbol{\Sigma}[I, I] = \mathbf{0}$ and $\boldsymbol{\Sigma}[i, i] \leq 1$ for any $i \in I^c$. Left-multiplying both sides of (E.2) by \mathbf{U}^T yields

$$\eta \mathbf{D} \mathbf{V}^T = \mathbf{U}^T \mathbf{X}^{\circ T} (\mathbf{Y} - \boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X})).$$

On the other hand, from the construction of $\hat{\mathbf{C}}$ and \mathbf{Z} , it is easy to verify $\mathbf{x}_i^T \hat{\mathbf{B}} = \mathbf{z}_i^T \hat{\mathbf{C}}$, from which it follows that $\boldsymbol{\mu}(\hat{\mathbf{B}}; \mathbf{X}) = \boldsymbol{\mu}(\hat{\mathbf{C}}; \mathbf{Z})$. Therefore, $\hat{\mathbf{C}}$ satisfies

$$(E.3) \quad \begin{cases} \eta \mathbf{C}^\circ &= \mathbf{Z}^{\circ T} (\mathbf{Y} - \boldsymbol{\mu}(\hat{\mathbf{C}}; \mathbf{Z})), \\ \mathbf{0} &= (\mathbf{Y} - \boldsymbol{\mu}(\hat{\mathbf{C}}; \mathbf{Z}))^T \hat{\mathbf{z}}_0. \end{cases}$$

Noticing that the optimization problem in (2.20) is convex and (E.3) gives its KKT equation, the conclusion follows.

APPENDIX F. PROOF OF THEOREM 3.1

Lemma F.1. *Given any $\mathbf{Y} \in \mathbb{R}^{n \times r}$, $\hat{\mathbf{B}} = \Theta^{\#\sigma}(\mathbf{Y}; r, \eta)$ is a globally optimal solution to*

$$(F.1) \quad \min_{\mathbf{B}} \frac{1}{2} \|\mathbf{Y} - \mathbf{B}\|_F^2 + \frac{\eta}{2} \|\mathbf{B}\|_F^2 \quad \text{s.t.} \quad \text{rank}(\mathbf{B}) \leq r.$$

Proof. The problem is equivalent to minimizing

$$\frac{1 + \eta}{2} \sum (\sigma_i^{(\mathbf{B})})^2 - \langle \mathbf{Y}, \mathbf{B} \rangle$$

subject to $\text{rank}(\mathbf{B}) \leq r$. Applying Lemma A.1 yields the result. \square

The remainder of the proof follows the same lines to the proof of Theorem 2.1. See Appendix C for details.

ACKNOWLEDGEMENTS

The author is grateful to the anonymous referee and the associate editor for their careful comments and useful suggestions. This work was partially supported by NSF grant CCF-1116447.

Received 19 December 2011

REFERENCES

- ANDERSON, T. W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *Annals of Mathematical Statistics* **22** 327–351. [MR0042664](#)
- BRILLINGER, D. R. (1981). *Time Series: Data Analysis and Theory*, expanded edn, San Francisco, CA: Holden-Day. [MR0595684](#)
- BUNEA, F., SHE, Y. and WEGKAMP, M. (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *Annals of Statistics* **39**(2) 1282–1309. [MR2816355](#)
- CANDÈS, E. J. and PLAN, Y. (2011). Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory* **57**(4) 2342–2359. [MR2809094](#)
- CANDÈS, E. and RECHT, B. (2009). Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* **9** 717–772. [MR2565240](#)
- CANDÈS, E. and TAO, T. (2010). The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory* **56**(5) 2053–2080. [MR2723472](#)
- CHEN, J. and CHEN, Z. (2008). Extended Bayesian information criterion for model selection with large model space. *Biometrika* **95** 759–771. [MR2443189](#)
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B* **70**(5) 849–911. [MR2530322](#)
- FAZEL, M. (2002). Matrix rank minimization with applications, PhD thesis, Stanford University.
- GRIGORIEFF, R. D. (1991). A note on von Neumann’s trace inequality. *Math. Nachr.* **151** 327–328.
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2001). *The Elements of Statistical Learning*, New York: Springer-Verlag. [MR1121214](#)
- HEINEN, A. and RENGIFO, E. (2008). Multivariate reduced rank regression in non-Gaussian contexts, using copulas. *Computational Statistics & Data Analysis* **52**(6) 2931–2944. [MR2424769](#)
- IZENMAN, A. (2008). *Modern Multivariate. Statistical Techniques: Regression, Classification and Manifold Learning*, New York: Springer.
- MA, S., GOLDFARB, D. and CHEN, L. (2009). Fixed point and Bregman iterative methods for matrix rank minimization. arXiv:0905.1643 [math.OA]. [MR2445017](#)
- OWEN, A. B. (2007). A robust hybrid of lasso and ridge regression. In: *Prediction and Discovery*, Contemporary Mathematics, vol. 443, pp. 59–71. [MR2433285](#)
- RAO, C. R. (1979). Separation theorems for singular values of matrices and their applications in multivariate analysis. *Journal of Multivariate Analysis* **9**(3) 362–377. [MR0548787](#)

- REINSEL, G. and VELU, R. (1998). *Multivariate Reduced-Rank Regression: Theory and Applications*, New York: Springer. [MR1719704](#)
- ROBINSON, P. M. (1974). Identification, estimation and large sample theory for regression containing unobservable variables. *International Economic Review* **15** 680–692. [MR0356376](#)
- SHE, Y. (2009). Thresholding-based iterative selection procedures for model selection and shrinkage. *Electronic Journal of Statistics* **3** 384–415. [MR2501318](#)
- SHE, Y. (2012). An iterative algorithm for fitting nonconvex penalized generalized linear models with grouped predictors. *Computational Statistics & Data Analysis* **56** 2976–2990.
- TURNBULL, D., BARRINGTON, L., TORRES, D. and LANCKRIET, G. (2008). Semantic annotation and retrieval of music and sound effects. *IEEE Transactions on Audio, Speech and Language Processing* **16**(2) 467–476.
- VON NEUMANN, J. (1937). Some matrix-inequalities and metrization of matrix-space. *Tomskii Univ. Rev.* **1** 286–300.
- YEE, T. W. and HASTIE, T. J. (2003). Reduced-rank vector generalized linear models. *Statistical Modelling* **3** 15–41.
- YUAN, M., EKICI, A., LU, Z. and MONTEIRO, R. (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B* **69**(3) 329–346. [MR1977163](#)
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *JRSSB* **67**(2) 301–320. [MR2137327](#)

Yiyuan She
Department of Statistics
Florida State University
Tallahassee, FL 32306-4330
USA
E-mail address: yshe@stat.fsu.edu