

This article was downloaded by: [Academia Sinica - Taiwan]

On: 02 July 2012, At: 13:23

Publisher: Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



## Journal of the American Statistical Association

Publication details, including instructions for authors and subscription information:  
<http://www.tandfonline.com/loi/uasa20>

### Likelihood-Based Selection and Sharp Parameter Estimation

Xiaotong Shen<sup>a</sup>, Wei Pan<sup>b</sup> & Yunzhang Zhu<sup>a</sup>

<sup>a</sup> School of Statistics, University of Minnesota, Minneapolis, MN, 55455

<sup>b</sup> Division of Biostatistics, University of Minnesota, Minneapolis, MN, 55455

Accepted author version posted online: 31 Jan 2012. Version of record first published: 11 Jun 2012

To cite this article: Xiaotong Shen, Wei Pan & Yunzhang Zhu (2012): Likelihood-Based Selection and Sharp Parameter Estimation, Journal of the American Statistical Association, 107:497, 223-232

To link to this article: <http://dx.doi.org/10.1080/01621459.2011.645783>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

# Likelihood-Based Selection and Sharp Parameter Estimation

Xiaotong SHEN, Wei PAN, and Yunzhang ZHU

In high-dimensional data analysis, feature selection becomes one effective means for dimension reduction, which proceeds with parameter estimation. Concerning accuracy of selection and estimation, we study nonconvex constrained and regularized likelihoods in the presence of nuisance parameters. Theoretically, we show that constrained  $L_0$  likelihood and its computational surrogate are optimal in that they achieve feature selection consistency and sharp parameter estimation, under one necessary condition required for any method to be selection consistent and to achieve sharp parameter estimation. It permits up to exponentially many candidate features. Computationally, we develop difference convex methods to implement the computational surrogate through prime and dual subproblems. These results establish a central role of  $L_0$  constrained and regularized likelihoods in feature selection and parameter estimation involving selection. As applications of the general method and theory, we perform feature selection in linear regression and logistic regression, and estimate a precision matrix in Gaussian graphical models. In these situations, we gain a new theoretical insight and obtain favorable numerical results. Finally, we discuss an application to predict the metastasis status of breast cancer patients with their gene expression profiles. This article has online supplementary material.

KEY WORDS: Continuous but nonsmooth minimization; Coordinate descent; General likelihood; Graphical models; Nonconvex;  $(p, n)$ -Asymptotics.

## 1. INTRODUCTION

Feature selection is essential to battle the inherited “curse of dimensionality” in high-dimensional analysis. It removes non-informative features to derive simpler models for interpretability, prediction, and inference. In cancer studies, for instance, a patient’s gene expression is linked to her metastasis status of breast cancer, for identifying cancer genes. In a situation as such, our ability of identifying cancer genes is as critical as a model’s predictive accuracy, where selection accuracy becomes extremely important to reproducible findings and generalizable conclusions. Toward accuracy of selection and parameter estimation, we address several core issues in high-dimensional likelihood-based selection.

Consider a selection problem with nuisance parameters, based on a random sample  $\mathbf{Y} = (Y_1, \dots, Y_n)$  with each  $Y_i$  following probability density  $g(\boldsymbol{\theta}^0, y)$ , where  $\boldsymbol{\theta}^0 = (\boldsymbol{\beta}^0, \boldsymbol{\eta}^0)$  is a true parameter vector,  $\boldsymbol{\beta}^0 \equiv (\beta_1^0, \dots, \beta_p^0) = (\boldsymbol{\beta}_{A_0}^0, \mathbf{0}_{A_0^c})$  and  $\boldsymbol{\eta}^0 \equiv (\eta_1^0, \dots, \eta_q^0)$  are the parameters of interest and nuisance parameters, respectively,  $A_0 = \{j : \beta_j^0 \neq 0\}$  is a set of nonzero coefficients of  $\boldsymbol{\beta}^0$  with size  $|A_0| = p_0$ , and  $\mathbf{0}_{A_0^c}$  is a vector of 0’s with  $^c$  denoting the set complement. Here, we estimate  $(\boldsymbol{\beta}^0, A_0)$ , where  $p$  may greatly exceed  $n$ , and  $q = 0$  is permitted.

For estimation and selection, a likelihood is regularized with regard to  $\boldsymbol{\beta}$ , particularly when  $p > n$ . This leads to an information criterion,

$$-L(\boldsymbol{\theta}) + \lambda \sum_{j=1}^p I(\beta_j \neq 0), \quad (1)$$

where  $L(\boldsymbol{\theta}) = \sum_{i=1}^n \log g(\boldsymbol{\theta}, Y_i)$  is the log likelihood based on  $\mathbf{Y}$ ,  $\lambda > 0$  is a regularization parameter, and  $\sum_{j=1}^p I(\beta_j \neq 0)$  is

the  $L_0$  function penalizing an increase in a model’s size. In (1), when  $\boldsymbol{\theta} = \boldsymbol{\beta}$  without nuisance parameters,  $\lambda = 1$  is Akaike’s information criterion,  $\lambda = \frac{\log n}{2}$  is Bayesian information criterion (Schwarz 1978), among others. In fact, essentially all selection rules can be cast into the framework of (1).

Regularization (1) has been of considerable interest for its interpretability and computational merits. Yet its constrained counterpart (2) has not received much attention, which is

$$-L(\boldsymbol{\theta}), \quad \text{subject to} \quad \sum_{j=1}^p I(\beta_j \neq 0) \leq K, \quad (2)$$

where  $K \geq 0$  is a tuning parameter corresponding to  $\lambda$  in (1). Minimizing (1) or (2) in  $\boldsymbol{\theta}$  gives a global minimizer leading to an estimate  $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}_{\hat{A}}, \mathbf{0}_{\hat{A}^c})^T$ , with  $\hat{A}$  the estimated  $A_0$ , where  $\boldsymbol{\eta}$  is unregularized and possibly profiled out. Note that (1) and (2) may not be equivalent in their global minimizers, which is unlike a convex problem.

This article systematically investigates constrained and regularized likelihoods involving nuisance parameters, for estimating zero components of  $\boldsymbol{\beta}^0$  as well as nonzero ones of  $\boldsymbol{\theta}^0$ . This includes, but is not limited to, estimating nonzero entries of a precision matrix in graphical models.

There is a huge body of literature on parameter estimation through  $L_1$  regularization in linear regression (see, e.g., Negahban et al. 2010 for a comprehensive review). For feature selection, consistency of the Lasso (Tibshirani 1996) has been extensively studied under the irrepresentable assumption (cf., Meinshausen and Bühlmann 2006; Zhao and Yu 2006). Other methods, such as the smoothly clipped absolute deviation (SCAD) (Fan and Li 2001), have been studied. Yet  $L_0$  constrained or regularized likelihoods remain largely unexplored. Despite progress, many open issues remain. First, what is the maximum number of candidate features allowed

Xiaotong Shen and Yunzhang Zhu, School of Statistics, University of Minnesota, Minneapolis, MN 55455. Wei Pan, Division of Biostatistics, University of Minnesota, Minneapolis, MN 55455. Research supported in part by NSF grant DMS-0906616, NIH grants 1R01GM081535, 2R01GM081535, HL65462, and R01HL105397. The authors thank the editor, the associate editor, and anonymous referees for helpful comments and suggestions.

for a likelihood method to reconstruct informative features? Results, such as Kim, Choi, and Oh (2008), seem to suggest that the capacity of handling exponentially many features may be attributed primarily to the exponential tail of a Gaussian distribution, which we show is not necessary. Second, can parameter estimation be enhanced through removal of zero components of  $\beta$ ? Third, can a selection method continue to perform well for parameters of interest in the presence of a large number of nuisance parameters, as in covariance selection for off-diagonal entries of a precision matrix?

This article intends to address the foregoing three issues. First, we establish finite-sample misselection error bounds for constrained  $L_0$  likelihood as well as its computational surrogate, given  $(n, p_0, p)$ , where the surrogate—a likelihood based on a truncated  $L_1$  function (TLP) approximating the  $L_0$  function—permits efficient computation (see Section 2.1 for a definition). On this basis, we establish feature selection consistency for them as  $n, p \rightarrow \infty$ , under one key condition that is necessary for any method to be selection consistent:

$$C_{\min}(\theta^0) \geq d_0 \frac{\log p}{n}, \quad (3)$$

where  $C_{\min}(\theta^0) \equiv \inf_{\{\theta_A = ((\beta_A, \theta_{A^c}), \eta) : A \neq A_0, |A| \leq p_0\}} \frac{-\log(1 - h^2(\theta_A, \theta^0))}{\max(|A_0 \setminus A|, 1)}$ ,  $d_0 > 0$  is a constant,  $|\cdot|$  and  $\setminus$  denote the size of a set and that of set difference, respectively,  $h(\theta, \theta^0) = \frac{1}{2}(\int (g^{1/2}(\theta, y) - g^{1/2}(\theta^0, y))^2 d\mu(y))^{1/2}$  is the Hellinger distance with respect to a dominating measure  $\mu$ , and  $g(\theta, y)$  is a probability density for  $Y_1$ . As one consequence, exponentially many candidate features  $p = \exp(n \frac{C_{\min}(\theta^0)}{d_0})$  are permitted for selection consistency with a broad class of constrained likelihoods. This challenges the well-established result that the maximum number of candidate features permitted for selection consistency depends highly on a likelihood's tail behavior (c.f., Chen and Chen 2008). In fact, selection consistency continues to hold even if the error distribution does not have an exponential tail (see Proposition 1 for linear regression). Second, sharper parameter estimation results from accurate selection by  $L_0$  likelihood and its surrogate as compared to that without such selection. For feature selection in linear regression and logistic regression, the optimal Hellinger risk of the oracle estimator, the maximum likelihood estimate (MLE) based on  $A_0$  as if the true  $A_0$  were known a priori, is recovered by these methods, which is of order of  $\sqrt{\frac{p_0}{n}}$  and is uniform over a certain  $L_0$  band of  $\theta^0$  excluding the origin. This is in contrast to the minimax rate  $\sqrt{\frac{u \log(p/u)}{n}}$  with  $u \geq p_0$  for estimation without feature selection in linear regression (Raskutti, Wainwright, and Yu 2009). In other words, accurate selection by  $L_0$  likelihood and its surrogate over the  $L_0$  band improves accuracy of estimation after noninformative features are removed, without introducing additional bias to estimation. Moreover, in estimating a precision matrix in Gaussian graphical models, the foregoing conclusions extend but with a different rate at  $\sqrt{\frac{p_0 \log p}{n}}$ , where a  $\log p$  factor is due to estimation of  $2p$  nuisance parameters as compared to logistic regression. Third, two difference of convex (DC) methods are employed for computation of (1) and (2), which relax nonconvex minimization through a sequence of convex problems.

Two disparate applications are considered, namely, feature selection in generalized linear models (GLMs) and estimation

of a precision matrix in Gaussian graphical models. In GLMs, feature selection in nonlinear regression appears more challenging than linear regression for a high-dimensional problem. In statistical modeling of a precision matrix in Gaussian graphical models, two major approaches have emerged to exploit matrix sparsity by likelihood selection and neighborhood selection. Articles based on these two approaches include Meinshausen and Bühlmann (2006), Li and Gui (2006), Yuan and Lin (2007), Benerjee, Ghmoui, and dAspremont (2008), Rothman, Bickel, Levina, and Zhu (2008, 2009), Rocha, Zhao, and Yu (2008), among others. As suggested by Rothman et al. (2008), existing methods may not perform well when the dimension of a matrix exceeds the sample size  $n$ , although they give estimates better than the sample covariance matrix. In addition, theoretical aspects for a likelihood approach remain to be under studied. In these situations, the proposed method compares favorably against its competitors in simulations, and novel theoretical results provide an insight into a selection process.

This article is organized as follows. Section 2 develops the proposed method for  $L_0$  regularized and constrained likelihoods. Section 3 presents main theoretical results for selection consistency and parameter estimation involving selection, followed by a necessary condition for selection consistency. Section 4 applies the general method and theory to feature selection in GLMs. Section 5 is devoted to estimation of a precision matrix in Gaussian graphical models. Section 6 presents an application to predict the metastasis status of breast cancer patients with their gene expression profiles. Section 7 contains technical proofs.

## 2. METHOD AND COMPUTATION

### 2.1 Method

In a high-dimensional situation, it is computationally infeasible to minimize a discontinuous cost function involving the  $L_0$  function in (1) and (2). As a surrogate, we seek a good approximation of the  $L_0$  function by the TLP, defined as  $J(|z|) = \min(\frac{|z|}{\tau}, 1)$ , with  $\tau > 0$  a tuning parameter controlling the degree of approximation (see Figure 1). This  $\tau$  decides which individual coefficients to be shrunk toward zero. The advantages of  $J(|z|)$  are fourfold, although  $J(z)$  has been considered in other contexts (Gasso, Raotomamonjy, and Canu 2009):

- (1) (Surrogate) It performs the model selection task of the  $L_0$  function, while providing a computationally efficient means. Note that the approximation error of the TLP function to the  $L_0$  function becomes zero when  $\tau$  is tuned such that  $\tau < \min\{|\beta_k^0| : k \in A_0\}$ , seeking the *sparsest* solution by minimizing the number of nonzero coefficients.
- (2) (Adaptive model selection through adaptive shrinkage) It performs adaptive model selection (Shen and Ye 2002) through a computationally efficient means when  $\lambda$  is tuned. Moreover, it corrects the Lasso bias through adaptive shrinkage combining shrinkage with thresholding.
- (3) (Piecewise linearity) It is piecewise linear, gaining computational advantages.
- (4) (Low resolutions) It discriminates small from large coefficients through thresholding. Consequently, it is capable of handling many low-resolution coefficients, through tuning  $\tau$ .




Figure 1. Truncated  $L_1$  function  $J(|\beta_j|)$  with  $\tau = 1$  in (a), and its DC decomposition into a difference of two convex functions  $J_1$  and  $J_2$  in (b).

To treat nonconvex minimization, we replace the  $L_0$  function by its surrogate  $J(\cdot)$  to construct an approximation of (2) and that of (1):

$$-L(\theta), \quad \text{subject to} \quad \sum_{j=1}^p J(|\beta_j|) \leq K, \quad (4)$$

$$S(\theta) = -L(\theta) + \lambda \sum_{j=1}^p J(|\beta_j|), \quad (5)$$

where (5) is a dual problem of (4). To solve (5) and (4), we develop difference convex methods for the primal and dual problems, for efficient computation.

### 2.2 Unconstrained Dual and Constrained Primal Problems

Our DC method for the dual problem (5) begins with a DC decomposition of  $S(\theta)$ :  $S(\theta) = S_1(\theta) - S_2(\beta)$ , where  $S_1(\theta) = -L(\theta) + \lambda \sum_{j=1}^p J_1(|\beta_j|)$ ,  $S_2(\beta) = \lambda \sum_{j=1}^p J_2(|\beta_j|)$ ,  $J_1(|\beta_j|) = \frac{|\beta_j|}{\tau}$ , and  $J_2(|\beta_j|) = \frac{|\beta_j|}{\tau} - \max(\frac{|\beta_j|}{\tau} - 1, 0)$ . Without loss of generality, assume that  $-L$  is convex in  $\theta$ ; otherwise, a DC decomposition of  $-L$  is required and can be treated similarly. Given this DC decomposition, a sequence of upper approximations of  $S(\theta)$  is constructed iteratively, say, at iteration  $m$ , with  $\nabla S_2$  a subgradient of  $S_2$  in  $|\beta|$ :  $S^{(m)}(\theta) = S_1(\theta) - (S_2(\hat{\beta}^{(m-1)}) + (|\beta| - |\hat{\beta}^{(m-1)}|)^T \nabla S_2(|\hat{\beta}^{(m-1)}|))$ , by successively replacing  $S_2(\beta)$  by its minorization, where  $|\cdot|$  for a vector takes the absolute value in each component. After ignoring  $S_2(\hat{\beta}^{(m-1)}) - \frac{\lambda}{\tau} \sum_{j=1}^p |\hat{\beta}_j^{(m-1)}| I(|\hat{\beta}_j^{(m-1)}| > \tau)$  that is independent of  $\theta$ , the problem reduces to

$$S^{(m)}(\theta) = -L(\theta) + \frac{\lambda}{\tau} \sum_{j=1}^p |\beta_j| I(|\hat{\beta}_j^{(m-1)}| \leq \tau). \quad (6)$$

Minimizing (6) in  $\theta$  yields its minimizer  $\hat{\theta}^{(m)}$ . The process continues in  $m$  until termination occurs. Our unconstrained DC method is summarized as follows.

*Algorithm 1:*

- Step 1.* (Initialization) Supply a good initial estimate  $\hat{\theta}^{(0)}$ , such as the minimizer of  $S_1(\theta)$ .
- Step 2.* (Iteration) At iteration  $m$ , compute  $\hat{\theta}^{(m)}$  by solving (6).
- Step 3.* (Termination) Terminate when  $S(\hat{\theta}^{(m-1)}) - S(\hat{\theta}^{(m)}) \leq \varepsilon$ , and no components of  $\hat{\beta}^{(m)}$  is at  $\pm\tau$ . Otherwise, add  $\varepsilon$  to that components whose absolute value is  $\tau$ , and go to Step 2, where  $\varepsilon$  is the square root of the machine precision. Then the estimate  $\hat{\theta} = \hat{\theta}^{(m^*)}$ , where  $m^*$  is the smallest index at the termination criterion.

In Algorithm 1, (6) reduces to a general weighted Lasso problem:  $-L(\theta) + \sum_{j=1}^p \lambda_j |\beta_j|$ , with  $\lambda_j = \frac{\lambda}{\tau} I(|\hat{\beta}_j^{(m-1)}| \leq \tau)$ . Therefore, any efficient software is applicable.

For (4), we decompose the nonconvex constraint into a difference of two convex functions to construct a sequence of approximating convex constraints. This amounts to solving the  $m$ th subproblem in a parallel fashion as in (6):

$$\begin{aligned} \min_{\beta} -L(\theta), \quad \text{subject to} \quad & \frac{1}{\tau} \sum_{j=1}^p |\beta_j| I(|\hat{\beta}_j^{(m-1)}| \leq \tau) \\ & \leq K - \sum_{j=1}^p I(|\hat{\beta}_j^{(m-1)}| > \tau). \end{aligned} \quad (7)$$

This leads to a constrained DC algorithm—Algorithm 2—for solving (4) by replacing (5) in Algorithm 1 by (4).

Algorithms 1 and 2 are a generalization of those in Shen et al. (2010) for a general likelihood, where all the computational properties there extend to the present situation, including equivalence of the DC solutions of the two algorithms and their convergence. Next, we shall work with (5) due to its computational advantage. For instance, a coordinate descent method that works well with (5) breaks down for (4) (cf., Shen et al. 2010).

### 3. THEORY

This section presents a general theory for accuracy of reconstruction of the oracle estimator  $\hat{\theta}^{ml} = (\hat{\beta}^{ml}, \hat{\eta}^{ml})$  with  $\hat{\beta}^{ml} = (\hat{\beta}_{A_0}^{ml}, \mathbf{0}_{A_0^c})$  given  $A_0$ , which is the MLE provided that the knowledge about  $A_0$  were known a priori. As direct consequences, feature selection consistency is studied as well as optimal parameter estimation defined by the oracle estimator. In addition, a necessary condition for feature selection will be established as well. A parallel theory for regularized likelihood is similar and thus is omitted.

#### 3.1 Constrained $L_0$ Likelihood

In (2), assume that a global minimizer exists, denoted by  $\hat{\theta}^{L_0} = (\hat{\beta}^{L_0}, \hat{\eta}^{L_0})$  with  $\hat{\beta}^{L_0} = (\hat{\beta}_{\hat{A}^{L_0}}^{L_0}, \mathbf{0}_{(\hat{A}^{L_0})^c})$ . Write  $\beta$  as  $(\beta_A, \mathbf{0}_{|A^c|})$ , with  $\beta_A$  being  $(\beta_1, \dots, \beta_{|A|})^T$  for any subset  $A \subset \{1, \dots, p\}$  of nonzero coefficients.

Before proceeding, we define a complexity measure for the size of a space  $\mathcal{F}$ . The bracketing Hellinger metric entropy of  $\mathcal{F}$ , denoted by the function  $H(\cdot, \mathcal{F})$ , is defined by logarithm of the cardinality of the  $u$ -bracketing (of  $\mathcal{F}$ ) of the smallest size. That is, for a bracket covering  $S(\varepsilon, m) = \{f_1^l, f_1^u, \dots, f_m^l, f_m^u\} \subset \mathcal{L}_2$  satisfying  $\max_{1 \leq j \leq m} \|f_j^u - f_j^l\|_2 \leq \varepsilon$  and for any  $f \in \mathcal{F}$ , there exists a  $j$  such that  $f_j^l \leq f \leq f_j^u$ , a.e.  $P$ , then  $H(u, \mathcal{F})$  is  $\log(\min\{m : S(u, m)\})$ , where  $\|f\|_2 = \int f^2(z) d\mu$ . For more discussions about metric entropy of this type, see Kolmogorov and Tihomirov (1959).

*Assumption A.* (Size of parameter space) For some constant  $c_0 > 0$  and any  $\frac{\varepsilon}{2t} < t < \varepsilon \leq 1$ ,  $H(t, \mathcal{B}_A) \leq c_0(\log p)^2 |A| \log(2\varepsilon/t)$ , with  $|A| \leq p_0$ , where  $\mathcal{B}_A = \mathcal{F}_A \cap \{h(\theta, \theta^0) \leq 2\varepsilon\}$  is a local parameter space, and  $\mathcal{F}_A = \{g^{1/2}(\theta, y) : \theta = (\beta, \eta), \beta = (\beta_A, 0)\}$  be a collection of square root densities.

*Theorem 1.* (Error bound and oracle properties) Under Assumption A, if  $K = p_0$ , then, there exists a constant  $c_2 > 0$ , say  $c_2 = \frac{2}{27} \frac{1}{963}$ , such that for  $(n, p_0, p)$ ,

$$P(\hat{\theta}^{L_0} \neq \hat{\theta}^{ml}) \leq \exp(-c_2 n C_{\min}(\theta^0) + 2 \log(p + 1) + 3). \quad (8)$$

Moreover, under (3) with  $d_0 > \max(\frac{2}{c_2}, (2c_0)^{1/2} c_4^{-1} \log(2^{1/2}/c_3))$ ,  $\hat{\theta}^{L_0}$  reconstructs the oracle estimator  $\hat{\theta}^{ml}$  with probability tending to one as  $n, p \rightarrow \infty$ . Three oracle properties hold as  $n, p \rightarrow \infty$ :

- (A) (Selection consistency) Estimator  $\hat{A}^{L_0}$  is selection consistent, that is,  $P(\hat{A}^{L_0} \neq A_0) \rightarrow 0$ .
- (B) (Optimal parameter estimation) For  $\theta^0$ ,  $Eh^2(\hat{\theta}^{L_0}, \theta^0) = (1 + o(1))Eh^2(\hat{\theta}^{ml}, \theta^0) = O(\varepsilon_{n,p}^2)$  and  $h^2(\hat{\theta}^{L_0}, \theta^0) = O_p(\varepsilon_{n,p_0,p}^2)$ , provided that  $Eh^2(\hat{\theta}^{ml}, \theta^0)$  does not tend to zero too fast in that  $\frac{c_2}{2} n C_{\min}(\theta^0) + \log Eh^2(\hat{\theta}^{ml}, \theta^0) \rightarrow \infty$ , where  $\varepsilon_{n,p_0,p}$  is any solution for  $\varepsilon$ :

$$\int_{2^{-8\varepsilon^2}}^{2^{1/2\varepsilon}} H^{1/2}(t/c_3, \mathcal{B}_{A_0}) dt \leq c_4 n^{1/2} \varepsilon^2. \quad (9)$$

- (C) (Uniformity over a  $L_0$  band) The reconstruction holds uniformly over  $B_0(u, l)$ , namely,  $\sup_{\theta^0 \in B_0(u,l)} P(\hat{\theta}^{L_0} \neq \hat{\theta}^{ml}) \rightarrow 0$ , where  $B_0(u, l)$  is a

$L_0$  band, defined as  $\{(\beta, \eta) : p_0 = \sum_{j=1}^p I(\beta_j \neq 0) \leq u, C_{\min}(\theta) \geq l\}$  with  $0 < u \leq \min(n, p)$ ,  $l = d_0 \sigma^2 \frac{\log p}{n}$ , and  $u < \min(n, p)$ . This implies feature selection consistency  $\sup_{\theta^0 \in B_0(u,l)} P(\hat{A}^{L_0} \neq A_0) \rightarrow 0$ , and optimal parameter estimation  $\frac{\sup_{\theta^0 \in B_0(u,l)} Eh^2(\hat{\theta}^{L_0}, \theta^0)}{\sup_{\theta^0 \in B_0(u,l)} Eh^2(\hat{\theta}^{ml}, \theta^0)} \rightarrow 1$ , with  $\sup_{\theta^0 \in B_0(u,l)} Eh^2(\hat{\theta}^{ml}, \theta^0) = O(\varepsilon_{n,u,p}^2)$ , provided that  $Eh^2(\hat{\theta}^{ml}, \theta^0)$  does not tend to zero too fast in that  $\frac{c_2}{2} n C_{\min}(\theta^0) + \log \sup_{\theta^0 \in B_0(u,l)} Eh^2(\hat{\theta}^{ml}, \theta^0) \rightarrow \infty$ .

The  $L_0$  method consistently reconstructs the oracle estimator when the degree of separation exceeds the minimal level, precisely under (3). As a result, selection consistency is established for the  $L_0$  method. This, combined with that in Theorem 3, suggests that the  $L_0$  method is optimal in feature selection against any method, matching up with the lower bound requirement under the degree of separation with respect to  $(p, p_0, n)$  except a constant factor  $d_0 > 0$  in Theorem 3. Moreover, the optimality extends further to parameter estimation, where sharper parameter estimation is obtained from accurate  $L_0$  selection, achieving the optimal Hellinger risk of the oracle estimator asymptotically. By comparison, such a result is not expected for  $L_1$  regularization. As suggested by Raskutti et al. (2009), selection consistency of Lasso does not give sharper parameter estimation, where the rate of convergence of a  $L_1$  method in the  $L_2$  risk remains to be  $\sqrt{\frac{p_0 \log(p/p_0)}{n}}$  in linear regression. This is because a  $L_1$  method is nonadaptive and overpenalizes large coefficients as a result of shrinking small coefficients toward zero. Similarly, in feature selection in logistic regression, the  $L_0$  method is expected to give better estimation precision than a  $L_1$  method, although a parallel result for a  $L_1$  method has not been available. Finally, the uniform result in (C) is over a  $L_0$  band  $B_0(u, l)$ , which is not expected over a  $L_0$  ball  $B_0(u, 0)$  in view of the result of Theorem 3.

#### 3.2 Constrained Truncated $L_1$ Likelihood

For constrained truncated  $L_1$  likelihood, one additional regularity condition—Assumption B—is assumed, which is generally met with a smooth likelihood (see Section 4 for an example). It requires the Hellinger distance to be smooth so that the TLP approximation to the  $L_0$  function becomes adequate through tuning  $\tau$ .

*Assumption B.* For some constants  $d_1 - d_3 > 0$ ,  $-\log(1 - h^2(\theta, \theta^0)) \geq -d_1 \log(1 - h^2(\theta_{\tau^+}, \theta^0)) - d_3 p \tau^{d_2}$ ,  $A^{\tau^+} \equiv \{j : |\beta_j| \geq \tau\}$  (10)

where  $\theta_{\tau^+} = (\beta_1 I(|\beta_1| \geq \tau), \dots, \beta_p I(|\beta_p| \geq \tau), \eta_1, \dots, \eta_q)$ .

*Theorem 2.* (Error bound and oracle properties) Under Assumption A with  $\mathcal{F}_A$  replaced by  $\{g^{1/2}(\theta, y) : \theta = (\beta, \eta) : \beta = (\beta_A, \beta_{A^c}), \|\beta_{A^c}\|_{\ell_\infty} \leq \tau\}$ , say  $0 \leq \tau \leq c'\varepsilon$  for some constant  $c'$ , and Assumption B, if  $K = p_0$  and  $\tau \leq \max(c', (d_1 C_{\min}(\theta^0)/2pd_3)^{1/d_2})$ , then there exists a constant  $c_2 > 0$ , such that for any  $(n, p_0, p)$ ,

$$P(\hat{\theta}^T \neq \hat{\theta}^{ml}) \leq \exp(-c_2 n C_{\min}(\theta^0) + 2 \log(p + 1) + 3). \quad (11)$$

Moreover, under (3) with sufficiently large constant  $d_0 > 0$ ,  $\hat{\theta}^T$  has the three oracle properties (A)–(C) of  $\hat{\theta}^{L_0}$ ,

provided that  $\frac{c_2 d_l}{4} n C_{\min}(\theta^0) + \log Eh^2(\hat{\theta}^{ml}, \theta^0) \rightarrow \infty$ . For (C),  $\tau \leq (\frac{d_l l}{2 p d_3})^{1/d_2}$  is required as well as  $\frac{c_2 d_l}{4} n C_{\min}(\theta^0) + \log \sup_{\theta^0 \in B_0(u,l)} Eh^2(\hat{\theta}^{ml}, \theta^0) \rightarrow \infty$ .

*Remark.* Constants in Theorem 1 can be made precise. For instance,  $c_2 = \frac{4}{27} \frac{1}{1926}$  and  $d_0 > \max(\frac{4}{c_2 d_1}, (2c_0)^{1/2} c_4^{-1} \log(2^{1/2}/c_3))$ .

Theorem 2 says that the oracle properties of the  $L_0$  function are attained by its computational surrogate when  $\tau$  is sufficiently small.

### 3.3 Necessary Condition for Selection Consistency

This section establishes the necessary condition (3) by estimating the minimal value  $d_0$  in (3), required for feature selection consistency.

Let  $K(\theta_1, \theta_2) = E \log(g(\theta_1, Y)/g(\theta_2, Y))$  be the Kullback–Leibler loss for  $\theta_1$  versus  $\theta_2$ , where  $E$  is taken with regard to  $g(\theta_1, Y)$ . Let  $\gamma_{\min}(\theta^0) \equiv \min\{|\beta_k^0| : k \in A_0\} > 0$ .

*Assumption C.* For a constant  $r > 0$ ,  $K(\theta_j, \theta_k) \leq r \gamma_{\min}^2(\beta)$ . Here  $\{\theta_j = (\beta_j, \eta^0), j = 1, \dots, p\}$  is a set of parameters, where  $\beta_j = \sum_{k=1}^{p_0} \gamma_{\min} e_k - \gamma_{\min} e_j; j = 1, \dots, p_0$ , and  $\beta_j = \sum_{k=1}^{p_0} \gamma_{\min} e_k + \gamma_{\min} e_j; j = p_0 + 1, \dots, p$ , and  $e_j = (\underbrace{0, \dots, 0}_{j-1}, \underbrace{1, 0, \dots, 0}_{p-j})^T$ . Assume that  $s \equiv \inf_{\theta^0} \frac{C_{\min}(\theta^0)}{\gamma_{\min}^2(\theta^0)} > 0$ .

*Theorem 3.* (Necessary condition for feature selection consistency) Under Assumption C, for any constant  $c_* \in (0, 1)$ , any  $(n, p_0, p)$  with  $p_0 \leq p/2$ , and any  $\eta^0$ , we have

$$\inf_{\hat{A}} \sup_{\{\beta^0: C_{\min}(\theta^0)=R^*\}} P(\hat{A} \neq A_0) \geq c_*, \tag{12}$$

with  $R^* = \frac{s(1-c_*) \log p}{4rn}$ . Moreover,

$$\inf_{\hat{A}} \sup_{\theta^0 \in B_0(u,l)} P(\hat{A} \neq A_0) \geq c_*, \text{ as } n, p \rightarrow \infty, \tag{13}$$

where  $u \leq \min(p/2, n)$ ,  $l = d_0 \frac{\log p}{n}$ , and  $d_0 = \frac{(1-c_*)s}{4r}$ .

Theorem 3 says that feature selection inconsistency occurs when  $d_0 < \frac{s}{4r}$  in (3). There the minimal value  $d_0 = \frac{s}{4r}$  yields a requirement for feature selection consistency in (3).

## 4. GENERALIZED LINEAR MODELS

For GLMs, observations  $Y_i = (Z_i, X_i)$  are paired, response  $Z_i$  is assumed to follow an exponential family with density function  $g(z_i; \theta_i, \phi) = \exp\{[z_i \theta_i - b(\theta_i)]/a(\phi) + c(z_i, \phi)\}$ , where  $\theta_i$  is the natural parameter that is related to the mean  $\mu_i = E(z_i) = b'(\theta_i)$ , and  $\phi$  is a dispersion parameter. With a link function  $g$ , a regression model becomes  $\eta_i = g(\mu_i) = \beta^T x_i$ . The penalized likelihood for estimating regression coefficient vectors  $\beta$  is  $-L(\beta) + \lambda \sum_{j=1}^p J(|\beta_j|)$ , where  $L(\beta) = \sum_{i=1}^n [z_i \mu_i - b(\mu_i)]/a(\phi) + c(z_i, \phi)$  is the log-likelihood, and  $J(|\beta_j|) = \frac{\lambda}{\tau} \min(|\beta_j|, \tau)$  is the TLP penalty.

For parameter estimation and feature selection, we apply Algorithm 1, where (6) becomes a series of weighted lasso for GLMs, for which some existing routines are applicable, for simplicity. In implementation, we use the function `wtlassoglm()` in R package SIS.

Next, we examine effectiveness of the proposed method through simulated examples in feature selection. In linear regression and logistic regression, the Lasso, SCAD (Fan and Li 2001), SCAD-OS, TLP, and TLP-OS are compared in terms of predictive accuracy and identification of the true model, where SCAD-OS and TLP-OS are SCAD and TLP with only one iteration step in the DC iterative process, and SCAD-OS is proposed in Zou and Li (2008). The latter four methods use the Lasso as an initial estimate.

### 4.1 Simulations

For simulations predictors  $X_i$ 's are independent and identically distributed from  $N(0, V)$ , where  $V$  is a  $p \times p$  matrix whose  $ij$ th element is  $0.5^{|i-j|}$ . In linear regression,  $Z_i = \beta^T X_i + \epsilon_i, \epsilon_i \sim N(0, \sigma^2); i = 1, \dots, n$ , and random error  $\epsilon_i$  is independent of  $X_i$ ; in logistic regression, a binary response is generated from  $\text{logit Pr}(Z_i = 1) = \beta^T X_i$ . In both cases,  $\beta = (\beta_1, \dots, \beta_p)^T$  with  $\beta_1 = 1, \beta_2 = 0.5$  and  $\beta_5 = 0.75; \beta_j = 0$  for  $j \neq 1, 2, 5$ . This setup was similar to that considered by Zou and Li (2008); here, we examine various situations with respect to  $p, n$ . Each simulation is based on 1000 independent replications.

For any given tuning parameter  $\lambda$ , all other methods use the Lasso estimate as an initial estimate. For each method, we choose its tuning parameter values by maximizing the log-likelihood based on a common tuning dataset with an equal sample size of the training data and independent of the training data. This is achieved through a grid search over 21  $\lambda$  values returned by `glmnet()` for all the methods, and additionally over a grid of 10  $\tau$  values that are the 9th, 19th, 29th, ..., 99th percentiles of the final Lasso estimate for the TLP.

The model error (ME) is used to evaluate predictive performance of  $\hat{\beta}$ , defined as  $\text{ME}(\hat{\beta}) = (\hat{\beta} - \beta^0)^T V (\hat{\beta} - \beta^0)$ , which is the prediction error minus  $\sigma^2$  in linear regression, corresponding to the test error over an independent test sample of size  $T = \infty$ . In our context, the median MEs are reported over 1000 simulation replications, due to possible skewness of the distribution of ME. In addition, the mean parameter estimates of the nonzero elements of  $\beta$  will be reported, together with the mean true positive (TP) and mean false positive (FP) numbers:  $\#\text{TP} = \sum_{j=1}^p I(\beta_j \neq 0, \hat{\beta}_j \neq 0)$  and  $\#\text{FP} = \sum_{j=1}^p I(\beta_j = 0, \hat{\beta}_j \neq 0)$ .

For linear regression, simulation results are reported for the cases of  $p = 12, 500, 1000, n = 50, 100$ , and  $\sigma^2 = 1$  in Table 1. As suggested by Table 1, the TLP performs best: it gives the smallest estimation and prediction error as measured by the ME, the smallest mean FP number while maintaining a comparable mean number of TPs around 3. Most critically, as  $p$  increases, the TLP's performance remains much more stable than its competitors. On a relative basis, the TLP outperforms its competitors more in more difficult situations.

For logistic regression, simulation results are summarized for the cases of  $p = 12, 200, 500$  and  $n = 100, 200$  in Table 2. As expected, the TLP continues to outperform other methods with the smallest median MEs. It gives less biased estimates than the Lasso estimates. The TLP's superior performance remains strong over other methods, as  $p$  increases.

### 4.2 Theory for Feature Selection

This section establishes some theoretical results to gain an insight into performance of the proposed method in feature

Table 1. Median MEs, means (SD in parentheses) nonzero coefficients ( $\beta_1, \beta_2, \beta_5$ ), and TP and FP numbers of nonzero estimates, for linear regression, based on 1000 simulation replications

$n$	$p$	Method	ME	$\beta_1 = 1$	$\beta_2 = 0.5$	$\beta_5 = 0.75$	#TP	#FP
50	12	Lasso	0.129	0.91(0.17)	0.41(0.18)	0.60(0.16)	2.98(0.14)	3.82(2.39)
		SCAD-OS	0.109	1.02(0.19)	0.40(0.22)	0.68(0.18)	2.92(0.27)	2.50(1.97)
		SCAD	0.118	1.04(0.20)	0.39(0.24)	0.71(0.18)	2.88(0.32)	2.30(1.90)
		TLP-OS	0.088	1.01(0.18)	0.41(0.20)	0.68(0.17)	2.94(0.25)	1.65(2.04)
		TLP	0.090	1.01(0.19)	0.41(0.21)	0.69(0.17)	2.92(0.27)	1.57(1.98)
50	500	Lasso	0.431	0.76(0.19)	0.29(0.18)	0.41(0.17)	2.90(0.30)	14.7(10.48)
		SCAD-OS	0.327	1.01(0.24)	0.25(0.25)	0.52(0.24)	2.70(0.47)	14.81(8.69)
		SCAD	0.301	1.09(0.26)	0.21(0.27)	0.59(0.26)	2.53(0.53)	12.25(7.63)
		TLP-OS	0.150	1.02(0.21)	0.39(0.26)	0.67(0.22)	2.75(0.45)	4.27(6.86)
		TLP	0.143	1.02(0.21)	0.39(0.26)	0.68(0.22)	2.75(0.45)	4.10(6.89)
50	1000	Lasso	0.501	0.72(0.19)	0.28(0.18)	0.37(0.18)	2.88(0.33)	17.20(11.49)
		SCAD-OS	0.370	0.99(0.25)	0.26(0.25)	0.51(0.26)	2.67(0.49)	18.76(9.60)
		SCAD	0.327	1.08(0.26)	0.20(0.28)	0.57(0.29)	2.49(0.54)	15.19(8.41)
		TLP-OS	0.182	1.01(0.20)	0.40(0.27)	0.66(0.25)	2.72(0.47)	5.43(8.69)
		TLP	0.175	1.02(0.20)	0.40(0.27)	0.66(0.25)	2.72(0.47)	5.06(8.30)
100	12	Lasso	0.063	0.94(0.12)	0.44(0.13)	0.65(0.11)	3.00(0.00)	3.94(2.42)
		SCAD-OS	0.042	1.01(0.12)	0.45(0.14)	0.72(0.11)	2.99(0.08)	2.17(2.04)
		SCAD	0.042	1.02(0.13)	0.45(0.15)	0.74(0.11)	2.99(0.09)	2.06(2.01)
		TLP-OS	0.037	1.01(0.12)	0.45(0.13)	0.71(0.11)	3.00(0.06)	1.51(1.99)
		TLP	0.036	1.01(0.12)	0.45(0.13)	0.72(0.11)	3.00(0.07)	1.46(1.95)
100	500	Lasso	0.186	0.84(0.12)	0.36(0.12)	0.52(0.11)	3.00(0.04)	15.61(11.04)
		SCAD-OS	0.118	1.06(0.14)	0.32(0.18)	0.66(0.14)	2.94(0.24)	14.92(10.45)
		SCAD	0.121	1.10(0.15)	0.30(0.21)	0.71(0.12)	2.89(0.31)	14.20(10.00)
		TLP-OS	0.036	1.01(0.12)	0.47(0.14)	0.72(0.11)	2.99(0.12)	3.64(6.57)
		TLP	0.035	1.01(0.12)	0.46(0.14)	0.72(0.11)	2.99(0.12)	3.49(6.69)
100	1000	Lasso	0.211	0.83(0.13)	0.34(0.13)	0.51(0.12)	3.00(0.06)	18.10(12.50)
		SCAD-OS	0.142	1.06(0.15)	0.30(0.19)	0.66(0.15)	2.91(0.29)	19.70(13.35)
		SCAD	0.147	1.10(0.15)	0.27(0.22)	0.72(0.14)	2.83(0.38)	18.80(12.74)
		TLP-OS	0.037	1.01(0.13)	0.46(0.15)	0.74(0.12)	2.97(0.18)	3.93(7.04)
		TLP	0.037	1.01(0.13)	0.46(0.15)	0.74(0.12)	2.97(0.18)	3.80(6.94)

selection. Let  $Y = (Z, \mathbf{X})$ , and  $g(\boldsymbol{\beta}, Z) = \frac{1}{2\sqrt{\pi}\sigma} \exp(-\frac{1}{2\sigma^2}(Z - \boldsymbol{\beta}^T \mathbf{X})^2)$  and  $g(\boldsymbol{\beta}, Z) = p^Z(1-p)^{1-Z}$  in linear and logistic regression. Assume that  $\boldsymbol{\beta}^T \mathbf{x} = \boldsymbol{\beta}_A^T \mathbf{x}_A$  belongs to a compact parameter space for any model size  $|A| \leq p_0$ . In this case, selection does not involve nuisance parameters, where  $\boldsymbol{\theta} = \boldsymbol{\beta}$ . Under (14), we establish feature selection consistency as well as optimal parameter estimation for the TLP:

$$\begin{aligned} & \min_{A:|A| \leq p_0, A \neq A_0} \frac{1}{\max(|A_0 \setminus A|, 1)} (\boldsymbol{\beta}_{A_0 \setminus A}^0)^T \\ & \times (\boldsymbol{\Sigma}_{A_0 \setminus A} - \boldsymbol{\Sigma}_{A_0 \setminus A, A} \boldsymbol{\Sigma}_A^{-1} \boldsymbol{\Sigma}_{A, A_0 \setminus A}) \boldsymbol{\beta}_{A_0 \setminus A}^0 \\ & \geq d_0 \frac{\log p}{n}, \end{aligned} \tag{14}$$

where  $d_0 > 0$  is a constant independent of  $(n, p, p_0)$ , and  $\boldsymbol{\Sigma}_B$  is a submatrix given a subset  $B$  of predictors, of covariance matrix  $\boldsymbol{\Sigma}$  with the  $jk$ th element  $\text{cov}(X_j, X_k)$ , independent of  $\boldsymbol{\beta}^0$ . A simpler but stronger condition can be used for verification of (14):

$$\begin{aligned} & \gamma_{\min}^2 \min_{A:|A| \leq p_0, A \neq A_0} c_{\min}(\boldsymbol{\Sigma}_{A_0 \setminus A} - \boldsymbol{\Sigma}_{A_0 \setminus A, A} \boldsymbol{\Sigma}_A^{-1} \boldsymbol{\Sigma}_{A, A_0 \setminus A}) \\ & \geq d_0 \frac{\log p}{n}, \end{aligned} \tag{15}$$

where  $\gamma_{\min} = \gamma_{\min}(\boldsymbol{\beta}^0) \equiv \min\{|\beta_k^0| : \beta_k^0 \neq 0\}$  is the resolution level of the true regression coefficients,  $\min_{A:|A| \leq p_0, A \neq A_0} c_{\min}$

$(\boldsymbol{\Sigma}_{A_0 \setminus A} - \boldsymbol{\Sigma}_{A_0 \setminus A, A} \boldsymbol{\Sigma}_A^{-1} \boldsymbol{\Sigma}_{A, A_0 \setminus A}) \geq \min_{B \supset A_0: |B| \leq 2p_0} c_{\min}(\boldsymbol{\Sigma}_B)$ , and  $c_{\min}$  denotes the smallest eigenvalue. Note that (14) is necessary for any method to be selection consistent except constant  $d_0$  if  $\min_{A:|A| \leq p_0, A \neq A_0} c_{\min}(\boldsymbol{\Sigma}_{A_0 \setminus A} - \boldsymbol{\Sigma}_{A_0 \setminus A, A} \boldsymbol{\Sigma}_A^{-1} \boldsymbol{\Sigma}_{A, A_0 \setminus A}) > 0$ .

*Proposition 1.* Under (14), the constrained MLE  $\hat{\boldsymbol{\beta}}^T$  of (4) consistently reconstructs the oracle estimate  $\hat{\boldsymbol{\beta}}^{ml}$ . As  $n, p \rightarrow \infty$ , feature selection consistency is established for the TLP as well as optimal parameter estimation  $Eh^2(\hat{\boldsymbol{\beta}}^T, \boldsymbol{\beta}^0) = Eh^2(\hat{\boldsymbol{\beta}}^{ml}, \boldsymbol{\beta}^0) = O(\frac{p_0}{n})$  under the Hellinger distance  $h(\cdot, \cdot)$ . Moreover, the results hold uniformly over a  $L_0$ -band  $B_0(u, l) = \{\boldsymbol{\beta}_0 : \sum_{j=1}^p I(\beta_j^0 \neq 0) \leq u, \gamma_{\min}^2(\boldsymbol{\beta}^0) \min_{B \supset A_0: |B| \leq 2p_0} c_{\min}(\boldsymbol{\Sigma}_B) \geq l\}$ , with  $0 < u \leq \min(n, p), l = d_0 \sigma^2 \frac{\log p}{n}$ , that is, as  $n, p \rightarrow \infty$ ,

$$\begin{aligned} & \sup_{\boldsymbol{\beta}^0 \in B_0(u, l)} P(\hat{\boldsymbol{\beta}}^T \neq \hat{\boldsymbol{\beta}}^{ml}) \rightarrow 0, \\ & \times \frac{\sup_{\boldsymbol{\beta}^0 \in B_0(u, l)} Eh^2(\hat{\boldsymbol{\beta}}^T, \boldsymbol{\beta}^0)}{\sup_{\boldsymbol{\beta}^0 \in B_0(u, l)} Eh^2(\hat{\boldsymbol{\beta}}^{ml}, \boldsymbol{\beta}^0)} \rightarrow 1, \end{aligned}$$

with  $\sup_{\boldsymbol{\beta}^0 \in B_0(u, l)} Eh^2(\hat{\boldsymbol{\beta}}^{ml}, \boldsymbol{\beta}^0) = d^* \frac{u}{n}$  for some  $d^*$ .

Various conditions have been proposed for studying feature selection consistency in linear regression. In particular, a condition on  $\gamma_{\min}$  is usually imposed, in addition to assumptions on

Table 2. Median MEs, means (SD in parentheses) nonzero coefficients ( $\beta_1, \beta_2, \beta_5$ ), and TP and FP numbers of nonzero estimates, for logistic regression, based on 1000 simulation replications

$n$	$p$	Method	ME	$\beta_1 = 1$	$\beta_2 = 0.5$	$\beta_5 = 0.75$	#TP	#FP
100	12	Lasso	0.388	0.80(.27)	0.35(0.25)	0.49(0.27)	2.8(0.4)	3.8(2.2)
		SCAD-OS	0.416	1.03(0.37)	0.39(0.36)	0.61(0.40)	2.5(0.7)	1.6(1.9)
		SCAD	0.472	1.10(0.41)	0.38(0.39)	0.67(0.44)	2.4(0.7)	1.1(1.9)
		TLP-OS	0.350	0.98(0.35)	0.36(0.32)	0.59(0.35)	2.7(0.5)	1.8(2.0)
		TLP	0.355	0.98(0.35)	0.35(0.32)	0.58(0.35)	2.6(0.6)	1.8(2.0)
100	200	Lasso	0.947	0.57(0.25)	0.20(0.19)	0.26(0.21)	2.6(0.6)	11.7(7.1)
		SCAD-OS	0.733	0.96(0.45)	0.23(0.36)	0.40(0.41)	2.0(0.7)	3.1(2.9)
		SCAD	0.827	1.08(0.53)	0.23(0.42)	0.46(0.53)	1.7(0.6)	1.1(1.4)
		TLP-OS	0.649	0.99(0.42)	0.31(0.36)	0.49(0.46)	2.2(0.7)	3.8(5.2)
		TLP	0.664	0.99(0.43)	0.30(0.37)	0.48(0.47)	2.2(0.7)	3.6(5.2)
100	500	Lasso	1.166	0.48(0.24)	0.18(0.19)	0.19(0.19)	2.4(0.7)	13.6(9.1)
		SCAD-OS	0.867	0.84(0.48)	0.23(0.35)	0.29(0.37)	1.8(0.7)	3.9(3.6)
		SCAD	0.847	1.00(0.57)	0.25(0.46)	0.34(0.50)	1.6(0.6)	1.3(1.5)
		TLP-OS	0.791	0.93(0.45)	0.30(0.39)	0.38(0.45)	2.0(0.7)	4.4(6.5)
		TLP	0.811	0.94(0.46)	0.29(0.40)	0.38(0.45)	2.0(0.7)	4.1(6.4)
200	12	Lasso	0.203	0.87(0.20)	0.39(0.20)	0.57(0.20)	3.0(0.2)	4.3(2.4)
		SCAD-OS	0.173	1.06(0.25)	0.44(0.28)	0.72(0.26)	2.8(0.4)	1.6(2.2)
		SCAD	0.202	1.08(0.25)	0.45(0.30)	0.77(0.25)	2.8(0.5)	1.2(2.1)
		TLP-OS	0.155	1.00(0.24)	0.40(0.24)	0.67(0.24)	2.9(0.3)	1.8(2.1)
		TLP	0.157	1.00(0.24)	0.40(0.24)	0.67(0.24)	2.9(0.3)	1.8(2.1)
200	200	Lasso	0.540	0.68(0.18)	0.27(0.17)	0.38(0.17)	2.9(0.3)	14.1(8.9)
		SCAD-OS	0.271	1.07(0.26)	0.34(0.32)	0.64(0.32)	2.6(0.6)	3.2(3.3)
		SCAD	0.262	1.12(0.29)	0.30(0.37)	0.68(0.36)	2.3(0.6)	0.8(1.4)
		TLP-OS	0.204	1.04(0.25)	0.40(0.29)	0.68(0.27)	2.7(0.5)	3.3(5.5)
		TLP	0.204	1.04(0.26)	0.39(0.30)	0.68(0.28)	2.7(0.5)	3.2(5.8)
200	500	Lasso	0.651	0.64(0.17)	0.24(0.16)	0.33(0.15)	2.9(0.3)	18.0(10.5)
		SCAD-OS	0.289	1.07(0.27)	0.31(0.32)	0.57(0.31)	2.5(0.5)	4.1(4.0)
		SCAD	0.262	1.13(0.28)	0.29(0.37)	0.66(0.36)	2.3(0.6)	1.4(1.7)
		TLP-OS	0.231	1.04(0.27)	0.39(0.30)	0.65(0.29)	2.7(0.5)	4.1(6.8)
		TLP	0.231	1.04(0.27)	0.38(0.30)	0.65(0.30)	2.7(0.5)	3.8(6.8)

the design matrix  $X$  such as the sparse Riesz condition in Zhang (2010). To compare (14) with existing assumptions for consistent selection, note that these assumptions imply a fixed design version of (14) by necessity of consistent feature selection. For instance, as shown in Zhang (2010), the sparse Riesz condition with dimension restriction and  $\gamma_{\min}^2 \geq c' \frac{\log(p-u)}{n}$ , required for the minimum concavity penalty to be consistent, imply (15) with  $p$  replaced by  $p - u$  thus (14) when  $p/u$  bounded away from 1, where  $u \geq p_0$ . Moreover, the number of overselected variables is proved to be bounded but may not tend to zero for thresholding Lasso in Theorem 1.1 of Zhou (2010), under a restrictive eigenvalue condition (Bickel et al. 2008) and a requirement on  $\gamma_{\min}$ . Finally, in linear regression, only finite variance  $\sigma^2$  is required for the proposed method, which is in contrast to a commonly used assumption on sub-Gaussian distribution of  $\epsilon_i$ .

In conclusion, the computational surrogate—the TLP—method indeed shares desirable oracle properties of the  $L_0$  method, which is optimal against any selection method, for feature selection and parameter estimation.

### 5. ESTIMATION OF A PRECISION MATRIX

Given  $n$  random samples from a  $p$ -dimensional normal distribution  $\mathbf{Y}_1, \dots, \mathbf{Y}_n \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , we estimate the inverse covariance matrix  $\boldsymbol{\Omega} = \boldsymbol{\Sigma}^{-1}$  that is  $p \times p$  positive definite, denoted

by  $\boldsymbol{\Omega} \succ 0$ . For estimation of  $(\boldsymbol{\mu}, \boldsymbol{\Omega})$ , the log-likelihood is proportional to

$$\frac{n}{2} \log \det(\boldsymbol{\Omega}) - \frac{1}{2} \sum_{i=1}^n (\mathbf{Y}_i - \boldsymbol{\mu})^T \boldsymbol{\Omega} (\mathbf{Y}_i - \boldsymbol{\mu}). \quad (16)$$

The profile log-likelihood for  $\boldsymbol{\Omega}$ , after  $\boldsymbol{\mu}$  is maximized out, is proportional to  $\frac{n}{2} \log \det(\boldsymbol{\Omega}) - \frac{1}{2} \text{tr}(\mathbf{S}\boldsymbol{\Omega})$ , where  $\bar{\mathbf{Y}} = n^{-1} \sum_{i=1}^n \mathbf{Y}_i$  and  $\mathbf{S} = n^{-1} \sum_{i=1}^n (\mathbf{Y}_i - \bar{\mathbf{Y}})(\mathbf{Y}_i - \bar{\mathbf{Y}})^T$  are the corresponding sample mean and covariance matrix,  $\det$  and  $\text{tr}$  denote the determinant and trace. In (16), the number of unknown parameters  $p^2$  in  $\boldsymbol{\Omega}$  can greatly exceed the sample size  $n$  in the presence of  $2p$  nuisance parameters  $(\boldsymbol{\mu}, \{\Omega_{jj} : j = 1, \dots, p\})$ , where  $\Omega_{jk}$  denotes the  $jk$ th elements of  $\boldsymbol{\Omega}$ . To avoid nonidentifiability in estimation, we regularize off-diagonal elements of  $\boldsymbol{\Omega}$  in (16) through a nonnegative penalty function  $J(\cdot)$  for the  $\frac{p(p-1)}{2}$  parameters of interest:

$$S(\boldsymbol{\Omega}) = \log \det(\boldsymbol{\Omega}) - \frac{1}{2} \text{tr}(\mathbf{S}\boldsymbol{\Omega}) - \lambda \sum_{j,k=1, j \neq k}^p J(\{\Omega_{jk}, j \neq k\}). \quad (17)$$

In estimation, the TLP function  $J(\{\Omega_{jk}, j \neq k\}) = \frac{1}{\tau} \min(|\Omega_{jk}|, \tau)$  is employed for both parameter estimation and



covariance selection in (17). Toward this end, we apply Algorithm 1 to solve (6) sequentially, which reduces to a series of weighted graphical lasso problems, and is solved by taking advantage of existing software. In implementation, we use R package *glasso* (Friedman, Hastie, and Tibshirani 2008) for (6).

### 5.1 Simulations

Simulations are performed, where a tridiagonal precision matrix is used as in Fan, Feng, and Wu (2009). In particular,  $\Sigma$  is AR(1)-structured with its  $ij$ -element being  $\sigma_{ij} = \exp(-a|s_i - s_j|)$ , and  $s_1 < s_2 < \dots < s_p$  are randomly chosen:  $s_i - s_{i-1} \sim \text{Unif}(0.5, 1)$ , for some  $a > 0$ ;  $i = 2, \dots, p$ . The following situations are considered:  $(n, p) = (120, 30)$  or  $(n, p) = (120, 200)$ , and  $a = 0.9$  or  $a = 0.6$ , based on 100 replications.

Five competing methods are compared, including Lasso, adaptive Lasso (ALasso), SCAD-OS and SCAD, and TLP-OS and TLP. ALasso uses weight  $\lambda/|\hat{\beta}_j^{(0)}|^\gamma$ , where  $\hat{\beta}^{(0)}$  is an initial estimate and  $\gamma = 1/2$  as in Fan et al. (2009).

To measure performance of estimator  $\hat{\Omega}$ , we use the entropy loss and quadratic loss:  $\text{loss}_1(\Omega, \hat{\Omega}) = \text{tr}(\Omega^{-1}\hat{\Omega}) - \log|\Omega^{-1}\hat{\Omega}| - p$ , and  $\text{loss}_2(\Omega, \hat{\Omega}) = \text{tr}(\Omega^{-1}\hat{\Omega} - I)^2$ , as well as the TP and FP numbers:  $\#\text{TP} = \sum_{i,j} I(\Omega_{ij} \neq 0, \hat{\Omega}_{ij} \neq 0)$ ;  $\#\text{FP} = \sum_{i,j} I(\Omega_{ij} = 0, \hat{\Omega}_{ij} \neq 0)$ .

For small  $p = 30$ , TLP and TLP-OS are always among the winners. It is also confirmed that the one-step approximation to SCAD or TLP gives similar performance to that of the fully iterated SCAD or TLP, respectively. For large  $p$ , to save computing time, as advocated in Fan et al. (2009), we only run SCAD-OS and TLP-OS. In such a situation, an improvement of TLP-OS over other methods is more substantial for large  $p = 200$  than for small  $p = 30$ . Overall, the proposed method delivers higher performance in low-dimensional and high-dimensional situations, respectively.

### 5.2 Theory for Precision Matrix

To perform theoretical analysis, we specify a parameter space  $\Theta$  in which  $\Omega > 0$  with  $0 < \max_{1 \leq j \leq p} |\Omega_{jj}| \leq M_2$ ,  $c_{\min}(\Omega) \geq M_1 > 0$ , for some constants  $M_1, M_2 > 0$ , independent of  $(n, p, p_0)$ . Let  $A = \{(j, k) : j \neq k, \Omega_{jk} \neq 0\}$  be the set of nonzero off-diagonal elements of  $\Omega$ , where  $|A| = p_0$  is an even number by symmetry of  $\Omega$ , and  $\Omega$  depends on  $A$ . Results in Theorem 1 imply that the constrained MLE yields covariance selection consistency under one assumption:

$$C_{\min}(\Omega^0) \geq d_0 \frac{\log p}{n}, \tag{18}$$

which is necessary for covariance selection consistency indeed for any method, up to constant  $d_0$  when  $c_{\min}(H) > 0$ , where  $d_0 > 0$  is a constant independent of  $(n, p, p_0)$ , and  $H = (\frac{\partial^2(-\log \det(\Omega))}{\partial^2 \Omega})|_{\Omega=\Omega_0}$  is the  $p^2 \times p^2$  Hessian matrix of  $-\log \det(\Omega)$ , whose  $(\Omega_{jk}, \Omega_{j'k'})$  element is  $\text{tr}(\Sigma^0 \Delta_{jk} \Sigma^0 \Delta_{j'k'})$  (cf., Boyd and Vandenberghe 2004),  $\Delta_{jk}$  is a  $p \times p$  with the  $jk$ -element being 1 and 0 otherwise. Sufficiently, (18) can be verified using

$$C_{\min}(\Omega^0) \geq \gamma_{\min}^2 c_{\min}(H), \tag{19}$$

with  $\gamma_{\min}(\Omega^0) \equiv \gamma_{\min} = \min\{|\Omega_{jk}^0| : \Omega_{jk}^0 \neq 0, j \neq k\}$ .

*Proposition 2.* Under (18), the constrained MLE  $\hat{\Omega}^T$  of (4) consistently reconstructs the oracle estimator  $\hat{\Omega}^{ml}$ ,

$$c^* = \frac{1}{8} \min_{\Omega: \#\{\omega_{jk} \neq 0, j \neq k\} \leq p_0} \frac{c_{\min}(H(\Omega))}{c_{\min}(H_0)}.$$

As  $n, p \rightarrow \infty$ , covariance selection consistency is established for the TLP as well as optimal parameter estimation

$$Eh^2(\hat{\Omega}^T, \Omega^0) = (1 + o(1))Eh^2(\hat{\Omega}^{ml}, \Omega^0) = O\left(\frac{p_0 \log p}{n}\right),$$

where

$$h^2(\Omega, \Omega^0) = 1 - \sqrt{\frac{\det(\Omega)\det(\Omega^0)}{\det(\frac{\Omega+\Omega^0}{2})}}^{1/2}$$

is the squared Hellinger distance for  $\Omega$  versus  $\Omega^0$ . Moreover, the above results hold uniformly over a  $L_0$  band  $B_0(u, l) = \{\Omega^0 : \sum_{j,k=1, j \neq k}^p I(\Omega_{jk}^0 \neq 0) \leq u, \gamma_{\min}^2(\Omega^0)c_{\min}(H) \geq l\}$ , with  $0 < u \leq \min(n, p)$  and  $l = d_0 \sigma^2 \frac{\log p}{n}$ , that is, as  $n, p \rightarrow \infty$ ,

$$\sup_{\Omega^0 \in B_0(u, l)} P(\hat{\Omega}^T \neq \hat{\Omega}^{ml}) \rightarrow 0, \quad \frac{\sup_{\Omega^0 \in B_0(u, l)} Eh^2(\hat{\Omega}^T, \Omega^0)}{\sup_{\Omega^0 \in B_0(u, l)} Eh^2(\hat{\Omega}^{ml}, \Omega^0)} \rightarrow 1,$$

with  $\sup_{\Omega^0 \in B_0(u, l)} Eh^2(\hat{\Omega}^{ml}, \Omega^0) = d^* \frac{u \log p}{n}$  for some  $d^* > 0$ .

In short, the TLP method is optimal against any method in covariance selection, permitting  $p$  up to exponentially large in the sample size, or  $p^2 \leq p_0 \exp(n \frac{\gamma_{\min}^2 c_{\min}(H)}{d_0})$ . Moreover, as a result of accurate selection of this method, parameter estimation can be sharply enhanced at an order of  $\sqrt{\frac{p_0 \log p}{n}}$ , as measured by the Hellinger distance, after zero off-diagonal elements are removed. Note that the  $\log p$  factor is due to estimation of  $2p$  nuisance parameters as compared to the rate of  $\sqrt{\frac{p_0}{n}}$  in logistic regression. In view of the result in Lemma 1, this result seems to be consistent with the minimax rate  $\sqrt{\frac{\log p}{n}}$  under the  $L_\infty$  matrix norm (Rothman et al. 2009).

## 6. METASTASIS STATUS OF BREAST CANCER PATIENTS

We apply the penalized logistic regression methods to analyze a microarray gene expression dataset of Wang et al. (2005), where our objectives are (1) to develop a model predicting the metastasis status, and (2) to identify cancer genes, for breast cancer patients. Among the 286 patients, metastasis was detected in 106 patients during follow-ups within five years after surgery. Their expression profiles were obtained from primary breast tumors with Affymetrix HG-133a GeneChips.

In Wang et al. (2005), a 76-gene signature was developed based on a training set of 115 patients, which yielded a misclassification error rate of  $64/171 = 37.4\%$  when applied to the remaining samples. Wei and Li (2007) compared the performance of a variety of classifiers using a subset of 245 genes drawn from 33 cancer-related pathways based on a 10-fold cross-validation (CV). Their nonparametric pathway-based regression method yielded the smallest error rate at 29%, while random forest, bagging, and support vector machine (SVM) had error rates of 33%, 35%, and 42%, respectively.

Table 3. Averaged (with SD in parentheses) entropy loss ( $loss_1$ ), quadratic loss ( $loss_2$ ), TP and FP numbers of nonzero parameters based on 100 simulations, for estimating a precision matrix in Gaussian graphical models in Section 4

Setup	Method	$loss_1$	$loss_2$	#TP	#FP
$p = 30, a = 0.9$	Lasso	1.55(0.15)	2.96(0.42)	88.0(0.0)	314.0(41.6)
	ALasso	1.02(0.15)	1.99(0.37)	88.0(0.0)	95.5(30.2)
	SCAD-OS	0.93(0.16)	1.99(0.44)	88.0(0.0)	126.0(39.4)
	SCAD	0.74(0.16)	1.60(0.42)	87.9(0.5)	85.5(18.0)
	TLP-OS	0.66(0.18)	1.47(0.47)	87.9(0.5)	28.1(21.4)
	TLP	0.63(0.18)	1.39(0.48)	87.8(0.7)	22.4(17.0)
$p = 30, a = 0.6$	Lasso	1.69(0.16)	3.28(0.46)	88.0(0.0)	342.5(35.5)
	ALasso	1.01(0.15)	1.97(0.37)	88.0(0.0)	103.9(17.4)
	SCAD-OS	0.75(0.14)	1.61(0.36)	88.0(0.0)	83.8(29.0)
	SCAD	0.56(0.12)	1.20(0.30)	88.0(0.2)	26.1(15.0)
	TLP-OS	0.57(0.14)	1.26(0.37)	88.0(0.0)	14.7(13.0)
	TLP	0.54(0.14)	1.18(0.36)	88.0(0.0)	7.3(10.6)
$p = 200, a = 0.9$	Lasso	20.16(0.50)	34.50(1.85)	597.9(0.4)	4847.8(614.7)
	ALasso	10.62(0.53)	19.64(1.20)	597.3(1.2)	936.8(37.9)
	SCAD-OS	11.46(0.60)	24.03(1.67)	597.7(0.8)	2453.6(251.2)
	TLP-OS	6.16(0.77)	13.99(2.10)	593.6(3.0)	284.8(158.0)
$p = 200, a = 0.6$	Lasso	24.86(0.54)	46.18(3.72)	598.0(0.0)	6161.7(863.0)
	ALasso	11.06(0.48)	21.53(1.18)	598.0(0.0)	1526.1(118.6)
	SCAD-OS	9.43(0.49)	20.87(1.77)	598.0(0.0)	2754.8(523.6)
	TLP-OS	4.45(0.48)	9.89(1.29)	597.7(0.7)	185.5(76.3)

In our analysis, we first performed a preliminary screening of the genes using a marginal  $t$  test to select the top  $p$  genes with most significant  $p$ -values, based on the training data for each fold of a 10-fold CV. Then the training data were split into two parts to fit penalized logistic models and to select tuning parameters, respectively. The results were summarized in Table 4, including the total misclassification errors and average model sizes (i.e., nonzero estimates) based on 10-fold CV. A final model is obtained by fitting the best model selected from a 10-fold CV to the entire dataset.

With regard to prediction, no large difference is seen among various methods, with the error rates ranging from  $102/286 =$

$35.7\%$  (of TLP and TLP-OS with  $p = 200$ ) to  $118/286 = 41.3\%$  (of ALasso with  $p = 200$ ). The TLP performed similar to the TLP-OS, both were among the winners. In addition, the Lasso gave the least sparse models while the SCAD gave the most sparse models.

With regard to identifying cancer genes, the Lasso, TLP-OS, and TLP yield the same model, identifying the largest number of cancer genes, whereas the SCAD and SCAD-OS give the most sparse models with only at most two cancer genes, and ALasso only yields 10 cancer genes. Here cancer genes are defined according to the Cancer Gene Database (Higgins, Claremont, Major, Sander, and Lash 2007).

In summary, the TLP and TLP-OS identify a good proportion of cancer genes and lead to a model giving a reasonably good predictive accuracy of the metastasis status. In this sense, they perform well with regard to the foregoing two objectives.

Table 4. Analysis results with various numbers ( $p$ ) of predictors for the breast cancer data. The numbers of total classification errors (#Err), including FPs (#FP), and mean numbers of nonzero estimates (#Nonzero) from 10-fold CV, and the total numbers of nonzero estimates and cancer genes in the final models are shown

$p$	Method	10-fold CV			Final model	
		#Err	#FP	#Nonzero	#Nonzero	#Cancer genes
200	Lasso	107	17	40.1	62	13
	ALasso	118	27	18.8	39	9
	SCAD-OS	107	4	9.5	15	2
	SCAD	107	1	4.7	2	0
	TLP-OS	102	8	33.5	62	13
	TLP	102	8	33.2	62	13
400	Lasso	107	19	46.9	95	26
	ALasso	112	19	14.4	32	10
	SCAD-OS	108	8	11.1	15	2
	SCAD	106	0	4.1	2	0
	TLP-OS	106	15	40.1	95	26
	TLP	106	14	38.2	95	26

## 7. SUPPLEMENTARY MATERIAL

### Appendix

[Received February 2011. Revised August 2011.]

### REFERENCES

Banerjee, O., El Ghaoui, L., and d'Aspremont, A. (2008), "Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data," *Journal of Machine Learning Research*, 9, 485–516. [224]

Bickel, P., Ritov, Y., and Tsybakov, A. (2008), "Simultaneous Analysis of Lasso and Dantzig Selector," *The Annals of Statistics*, 37, 1705–1732. [228]

Boyd, S., and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge UK: Cambridge University Press. [230]

Chen, J., and Chen, Z. (2008), "Extended Bayesian Information Criterion for Model Selection with Large Model Space," *Biometrika*, 95, 759–771. [224]

- Fan, J., Feng, Y., and Wu, Y. (2009), "Network Exploration Via the Adaptive Lasso and SCAD Penalties," *Annals of Applied Statistics*, 3, 521–541. [230]
- Fan, J., and Li, R. (2001), "Variable Selection Via Nonconcave Penalized Likelihood and its Oracle Properties," *Journal of the American Statistical Association*, 96, 1348–1360. [223,227]
- Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation with the Graphical Lasso," *Biostatistics*, 9, 432–441. [229]
- Gasso, G., Raotomamonjy, A., and Canu, S. (2009), "Recovering Sparse Signals with Nonconvex Penalties and DC Programming," *Signal Processing*, IEEE, 57, 4686–4698. [224]
- Higgins, M. E., Claremont, M., Major, J. E., Sander, C., and Lash, A. E. (2007), "CancerGenes: A Gene Selection Resource for Cancer Genome Projects," *Nucleic Acids Research*, 35(suppl 1), D721–D726. [231]
- Kim, Y., Choi, H., and Oh, H.-S. (2008), "Smoothly Clipped Absolute Deviation of High Dimensions," *Journal of the American Statistical Association*, 103, 1665–1673. [223]
- Kolmogorov, A. N., and Tihomirov, V. M. (1959), " $\epsilon$ -entropy and  $\epsilon$ -capacity of Sets in Function Spaces," *Uspekhi Matematicheskikh Nauk*, 14, 3–86. [In Russian. English translation, *American Mathematical Society*, 2, 17, 277–364. (1961).] [226]
- Li, H., and Gui, J. (2006), "Gradient Directed Regularization for Sparse Gaussian Concentration Graphs, with Applications to Inference of Genetic Networks," *Biostatistics*, 7, 302–317. [224]
- Meinshausen, N., and Bühlmann, P. (2006), "High Dimensional Graphs and Variable Selection with the Lasso," *The Annals of Statistics*, 34, 1436–1462. [223,224]
- Negahban, S., Wainwright, M., Ravikumar, P., and Yu, B. (2010), "A Unified Framework for High-Dimensional Analysis of M-estimators with Decomposable Regularizers," Technical Report, Department of Statistics, University of California, Berkeley. [223]
- Raskutti, G., Wainwright, M., and Yu, B. (2009), "Minimax Rates of Estimation for High-Dimensional Linear Regression over  $l_q$  Balls," University California, Berkeley: Tech Report. [224,226]
- Rocha, G., Zhao, P., and Yu, B. (2008), "A Path Following Algorithm for Sparse Pseudo-Likelihood Inverse Covariance Estimation," Technical Report 759, University of California, Berkeley. [224]
- Rothman, A., Bickel, P., Levina, E., and Zhu, J. (2008), "Sparse Permutation Invariant Covariance Estimation," *Electronic Journal of Statistics*, 2, 494–515. [224]
- (2009), "A New Approach to Cholesky-based Covariance Regularization in High Dimensions," *Biometrika*, 97, 539–550. [224,230]
- Schwarz, G. (1978), "Estimating the Dimension of a Model," *The Annals of Statistics*, 6, 461–64. [223]
- Shen, X., and Ye, J. (2002), "Adaptive Model Selection," *Journal of the American Statistical Association*, 97, 210–221. [224]
- Shen, X., Pan, W., Zhu, Y., and Zhou, H. (2010), "On Constrained and Regularization for High-Dimensional Regression," Technical Report, Division of Biostatistics, University of Minnesota. [225]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection Via the LASSO," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [223]
- Wang, Y., Klijn, J. G., E-Wang, Y., Klijn, J. G., Zhang, Y., Sieuwerts, A., Look, M., Yang, F., Talantov, D., Timmermans, M., Meijer-van Gelder, M., Yu, J., Jatke, T., Berns, E., Atkins, D., and Foekens, J. (2005), "Gene-Expression Profiles to Predict Distant Metastasis of Lymph-Node-Negative Primary Breast Cancer," *Lancet*, 365, 671–679. [230]
- Wei, Z., and Li, H. (2007), "Nonparametric Pathway-based Regression Models for Analysis of Genomic Data," *Biostatistics*, 8, 265–284. [230]
- Yuan, M., and Lin, Y. (2007), "Model Selection and Estimation in the Gaussian Graphical Model," *Biometrika*, 94, 19–35. [224]
- Zhang, C.-H. (2010), "Nearly Unbiased Variable Selection Under Minimax Concave Penalty," *The Annals of Statistics*, 38, 894–942. [228]
- Zou, H., and Li, R. (2008), "One-Step Sparse Estimates in Nonconcave Penalized Likelihood Models" (with discussion), *Annals of Statistics*, 36, 1509–1566. [227]
- Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," *Journal of Machine Learning Research*, 7, 2541–2563. [223]
- Zhou, S. (2010), "Thresholded Lasso for High-Dimensional Variable Selection and Statistical Estimation," Technical Report, Department of Statistics, University of Michigan. [228]