

This article was downloaded by: [Brown University]

On: 17 July 2012, At: 22:53

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://amstat.tandfonline.com/loi/uasa20>

Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties

Jianqing Fan and Runze Li

Jianqing Fan is Professor of Statistics, Department of Statistics, Chinese University of Hong Kong and Professor, Department of Statistics, University of California, Los Angeles, CA 90095. Runze Li is Assistant Professor, Department of Statistics, Pennsylvania State University, University Park, PA 16802-2111. Fan's research was partially supported by National Science Foundation (NSF) grants DMS-0196041 and DMS-9977096 and a grant from University of California at Los Angeles. Li's research was supported by a NSF grant DMS-0102505. The authors thank the associate editor and the referees for constructive comments that substantially improved an earlier draft.

Version of record first published: 31 Dec 2011

To cite this article: Jianqing Fan and Runze Li (2001): Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties, Journal of the American Statistical Association, 96:456, 1348-1360

To link to this article: <http://dx.doi.org/10.1198/016214501753382273>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://amstat.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.

Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties

Jianqing FAN and Runze LI

Variable selection is fundamental to high-dimensional statistical modeling, including nonparametric regression. Many approaches in use are stepwise selection procedures, which can be computationally expensive and ignore stochastic errors in the variable selection process. In this article, penalized likelihood approaches are proposed to handle these kinds of problems. The proposed methods select variables and estimate coefficients simultaneously. Hence they enable us to construct confidence intervals for estimated parameters. The proposed approaches are distinguished from others in that the penalty functions are symmetric, nonconcave on $(0, \infty)$, and have singularities at the origin to produce sparse solutions. Furthermore, the penalty functions should be bounded by a constant to reduce bias and satisfy certain conditions to yield continuous solutions. A new algorithm is proposed for optimizing penalized likelihood functions. The proposed ideas are widely applicable. They are readily applied to a variety of parametric models such as generalized linear models and robust regression models. They can also be applied easily to nonparametric modeling by using wavelets and splines. Rates of convergence of the proposed penalized likelihood estimators are established. Furthermore, with proper choice of regularization parameters, we show that the proposed estimators perform as well as the oracle procedure in variable selection; namely, they work as well as if the correct submodel were known. Our simulation shows that the newly proposed methods compare favorably with other variable selection techniques. Furthermore, the standard error formulas are tested to be accurate enough for practical applications.

KEY WORDS: Hard thresholding; LASSO; Nonnegative garrote; Penalized likelihood; Oracle estimator; SCAD; Soft thresholding.

1. INTRODUCTION

Variable selection is an important topic in linear regression analysis. In practice, a large number of predictors usually are introduced at the initial stage of modeling to attenuate possible modeling biases. On the other hand, to enhance predictability and to select significant variables, statisticians usually use stepwise deletion and subset selection. Although they are practically useful, these selection procedures ignore stochastic errors inherited in the stages of variable selections. Hence, their theoretical properties are somewhat hard to understand. Furthermore, the best subset variable selection suffers from several drawbacks, the most severe of which is its lack of stability as analyzed, for instance, by Breiman (1996). In an attempt to automatically and simultaneously select variables, we propose a unified approach via penalized least squares, retaining good features of both subset selection and ridge regression. The penalty functions have to be singular at the origin to produce sparse solutions (many estimated coefficients are zero), to satisfy certain conditions to produce continuous models (for stability of model selection), and to be bounded by a constant to produce nearly unbiased estimates for large coefficients. The bridge regression proposed in Frank and Friedman (1993) and the least absolute shrinkage and selection operator (LASSO) proposed by Tibshirani (1996, 1997) are members of the penalized least squares, although their associated L_q penalty functions do not satisfy all of the preceding three required properties.

The penalized least squares idea can be extended naturally to likelihood-based models in various statistical contexts. Our approaches are distinguished from traditional methods (usually quadratic penalty) in that the penalty functions are sym-

metric, convex on $(0, \infty)$ (rather than concave for the negative quadratic penalty in the penalized likelihood situation), and possess singularities at the origin. A few penalty functions are discussed. They allow statisticians to select a penalty function to enhance the predictive power of a model and engineers to sharpen noisy images. Optimizing a penalized likelihood is challenging, because the target function is a high-dimensional nonconcave function with singularities. A new and generic algorithm is proposed that yields a unified variable selection procedure. A standard error formula for estimated coefficients is obtained by using a sandwich formula. The formula is tested accurately enough for practical purposes, even when the sample size is very moderate. The proposed procedures are compared with various other variable selection approaches. The results indicate the favorable performance of the newly proposed procedures.

Unlike the traditional variable selection procedures, the sampling properties on the penalized likelihood can be established. We demonstrate how the rates of convergence for the penalized likelihood estimators depend on the regularization parameter. We further show that the penalized likelihood estimators perform as well as the oracle procedure in terms of selecting the correct model, when the regularization parameter is appropriately chosen. In other words, when the true parameters have some zero components, they are estimated as 0 with probability tending to 1, and the nonzero components are estimated as well as when the correct submodel is known. This improves the accuracy for estimating not only the null components, but also the nonnull components. In short, the penalized likelihood estimators work as well as if the correct submodel were known in advance. The significance of this is that the proposed procedures outperform the maximum likelihood estimator and perform as well as we hope. This is very analogous to the superefficiency phenomenon in the Hodges example (see Lehmann 1983, p. 405).

Jianqing Fan is Professor of Statistics, Department of Statistics, Chinese University of Hong Kong and Professor, Department of Statistics, University of California, Los Angeles, CA 90095. Runze Li is Assistant Professor, Department of Statistics, Pennsylvania State University, University Park, PA 16802-2111. Fan's research was partially supported by National Science Foundation (NSF) grants DMS-0196041 and DMS-9977096 and a grant from University of California at Los Angeles. Li's research was supported by a NSF grant DMS-0102505. The authors thank the associate editor and the referees for constructive comments that substantially improved an earlier draft.

The proposed penalized likelihood method can be applied readily to high-dimensional nonparametric modeling. After approximating regression functions using splines or wavelets, it remains very critical to select significant variables (terms in the expansion) to efficiently represent unknown functions. In a series of work by Stone and his collaborators (see Stone, Hansen, Kooperberg, and Truong 1997), the traditional variable selection approaches were modified to select useful spline subbases. It remains very challenging to understand the sampling properties of these data-driven variable selection techniques. Penalized likelihood approaches, outlined in Wahba (1990), and Green and Silverman (1994) and references therein, are based on a quadratic penalty. They reduce the variability of estimators via the ridge regression. In wavelet approximations, Donoho and Johnstone (1994a) selected significant subbases (terms in the wavelet expansion) via thresholding. Our penalized likelihood approach can be applied directly to these problems (see Antoniadis and Fan, in press). Because we select variables and estimate parameters simultaneously, the sampling properties of such a data-driven variable selection method can be established.

In Section 2, we discuss the relationship between the penalized least squares and the subset selection when design matrices are orthonormal. In Section 3, we extend the penalized likelihood approach discussed in Section 2 to various parametric regression models, including traditional linear regression models, robust linear regression models, and generalized linear models. The asymptotic properties of the penalized likelihood estimators are established in Section 3.2. Based on local quadratic approximations, a unified iterative algorithm for finding penalized likelihood estimators is proposed in Section 3.3. The formulas for covariance matrices of the estimated coefficients are also derived in this section. Two data-driven methods for finding unknown thresholding parameters are discussed in Section 4. Numerical comparisons and simulation studies also are given in this section. Some discussion is given in Section 5. Technical proofs are relegated to the Appendix.

2. PENALIZED LEAST SQUARES AND VARIABLE SELECTION

Consider the linear regression model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (2.1)$$

where \mathbf{y} is an $n \times 1$ vector and \mathbf{X} is an $n \times d$ matrix. As in the traditional linear regression model, we assume that y_i 's are conditionally independent given the design matrix. There are strong connections between the penalized least squares and the variable selection in the linear regression model. To gain more insights about various variable selection procedures, in this section we assume that the columns of \mathbf{X} in (2.1) are orthonormal. The least squares estimate is obtained via minimizing $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$, which is equivalent to $\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|^2$, where $\boldsymbol{\beta}^* = \mathbf{X}^T \mathbf{y}$ is the ordinary least squares estimate.

Denote $\mathbf{z} = \mathbf{X}^T \mathbf{y}$ and let $\hat{\mathbf{y}} = \mathbf{X}\mathbf{X}^T \mathbf{y}$. A form of the penalized least squares is

$$\begin{aligned} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \lambda \sum_{j=1}^d p_j(|\beta_j|) &= \frac{1}{2} \|\mathbf{y} - \hat{\mathbf{y}}\|^2 + \frac{1}{2} \sum_{j=1}^d (z_j - \beta_j)^2 \\ &\quad + \lambda \sum_{j=1}^d p_j(|\beta_j|). \end{aligned} \quad (2.2)$$

The penalty functions $p_j(\cdot)$ in (2.2) are not necessarily the same for all j . For example, we may wish to keep important predictors in a parametric model and hence not be willing to penalize their corresponding parameters. For simplicity of presentation, we assume that the penalty functions for all coefficients are the same, denoted by $p(\cdot)$. Furthermore, we denote $\lambda p(\cdot)$ by $p_\lambda(\cdot)$, so $p(\cdot)$ may be allowed to depend on λ . Extensions to the case with different thresholding functions do not involve any extra difficulties.

The minimization problem of (2.2) is equivalent to minimizing componentwise. This leads us to consider the penalized least squares problem

$$\frac{1}{2}(z - \theta)^2 + p_\lambda(|\theta|). \quad (2.3)$$

By taking the hard thresholding penalty function [see Fig. 1(a)]

$$p_\lambda(|\theta|) = \lambda^2 - (|\theta| - \lambda)^2 I(|\theta| < \lambda), \quad (2.4)$$

we obtain the hard thresholding rule (see Antoniadis 1997 and Fan 1997)

$$\hat{\theta} = z I(|z| > \lambda); \quad (2.5)$$

see Figure 2(a). In other words, the solution to (2.2) is simply $z_j I(|z_j| > \lambda)$, which coincides with the best subset selection, and stepwise deletion and addition for orthonormal designs. Note that the hard thresholding penalty function is a smoother penalty function than the entropy penalty $p_\lambda(|\theta|) = (\lambda^2/2) I(|\theta| \neq 0)$, which also results in (2.5). The former facilitates computational expedience in other settings.

A good penalty function should result in an estimator with three properties.

1. *Unbiasedness*: The resulting estimator is nearly unbiased when the true unknown parameter is large to avoid unnecessary modeling bias.
2. *Sparsity*: The resulting estimator is a thresholding rule, which automatically sets small estimated coefficients to zero to reduce model complexity.
3. *Continuity*: The resulting estimator is continuous in data z to avoid instability in model prediction.

We now provide some insights on these requirements.

The first order derivative of (2.3) with respect to θ is $\text{sgn}(\theta)\{|\theta| + p'_\lambda(|\theta|)\} - z$. It is easy to see that when $p'_\lambda(|\theta|) = 0$ for large $|\theta|$, the resulting estimator is z when $|z|$ is sufficiently large. Thus, when the true parameter $|\theta|$ is large, the observed value $|z|$ is large with high probability. Hence the penalized least squares simply is $\hat{\theta} = z$, which is approximately unbiased. Thus, the condition that $p'_\lambda(|\theta|) = 0$




Figure 1. Three Penalty Functions $p_\lambda(\theta)$ and Their Quadratic Approximations. The values of λ are the same as those in Figure 5(c).

for large $|\theta|$ is a sufficient condition for unbiasedness for a large true parameter. It corresponds to an improper prior distribution in the Bayesian model selection setting. A sufficient condition for the resulting estimator to be a thresholding rule is that the minimum of the function $|\theta| + p'_\lambda(|\theta|)$ is positive. Figure 3 provides further insights into this statement. When $|z| < \min_{\theta \neq 0} \{|\theta| + p'_\lambda(|\theta|)\}$, the derivative of (2.3) is positive for all positive θ 's (and is negative for all negative θ 's). Therefore, the penalized least squares estimator is 0 in this situation, namely $\hat{\theta} = 0$ for $|z| < \min_{\theta \neq 0} \{|\theta| + p'_\lambda(|\theta|)\}$. When $|z| > \min_{\theta \neq 0} \{|\theta| + p'_\lambda(|\theta|)\}$, two crossings may exist as shown in Figure 1; the larger one is a penalized least squares estimator. This implies that a sufficient and necessary condition for continuity is that the minimum of the function $|\theta| + p'_\lambda(|\theta|)$ is attained at 0. From the foregoing discussion, we conclude that a penalty function satisfying the conditions of sparsity and continuity must be singular at the origin.

It is well known that the L_2 penalty $p_\lambda(|\theta|) = \lambda|\theta|^2$ results in a ridge regression. The L_1 penalty $p_\lambda(|\theta|) = \lambda|\theta|$ yields a soft thresholding rule

$$\hat{\theta}_j = \text{sgn}(z_j)(|z_j| - \lambda)_+, \quad (2.6)$$

which was proposed by Donoho and Johnstone (1994a). LASSO, proposed by Tibshirani (1996, 1997), is the penalized least squares estimate with the L_1 penalty in the general least squares and likelihood settings. The L_q penalty $p_\lambda(|\theta|) = \lambda|\theta|^q$ leads to a bridge regression (Frank and Friedman 1993 and Fu 1998). The solution is continuous only when $q \geq 1$. However, when $q > 1$, the minimum of $|\theta| + p'_\lambda(|\theta|)$ is zero and hence it does not produce a sparse solution [see Fig. 4(a)]. The only continuous solution with a thresholding rule in this family is the L_1 penalty, but this comes at the price of shifting the resulting estimator by a constant λ [see Fig. 2(b)].

2.1 Smoothly Clipped Absolute Deviation Penalty

The L_q and the hard thresholding penalty functions do not simultaneously satisfy the mathematical conditions for unbiasedness, sparsity, and continuity. The continuous differentiable penalty function defined by

$$p'_\lambda(\theta) = \lambda \left\{ I(\theta \leq \lambda) + \frac{(a\lambda - \theta)_+}{(a-1)\lambda} I(\theta > \lambda) \right\} \quad \text{for some } a > 2 \text{ and } \theta > 0, \quad (2.7)$$




Figure 2. Plot of Thresholding Functions for (a) the Hard, (b) the Soft, and (c) the SCAD Thresholding Functions With $\lambda = 2$ and $a = 3.7$ for SCAD.




Figure 3. A Plot of $\theta + p'_\lambda(\theta)$ Against $\theta(\theta > 0)$.

improves the properties of the L_1 penalty and the hard thresholding penalty function given by (2.4) [see Fig. 1(c) and subsequent discussion]. We call this penalty function the smoothly clipped absolute deviation (SCAD) penalty. It corresponds to a quadratic spline function with knots at λ and $a\lambda$. This penalty function leaves large values of θ not excessively penalized and makes the solution continuous. The resulting solution is given by

$$\hat{\theta} = \begin{cases} \text{sgn}(z)(|z| - \lambda)_+, & \text{when } |z| \leq 2\lambda, \\ \{(a-1)z - \text{sgn}(z)a\lambda\}/(a-2), & \text{when } 2\lambda < |z| \leq a\lambda, \\ z, & \text{when } |z| > a\lambda \end{cases} \quad (2.8)$$

[see Fig. 2(c)]. This solution is owing to Fan (1997), who gave a brief discussion in the settings of wavelets. In this article, we use it to develop an effective variable selection procedure for a broad class of models, including linear regression models and generalized linear models. For simplicity of presentation, we use the acronym SCAD for all procedures using the SCAD penalty. The performance of SCAD is similar to that of firm shrinkage proposed by Gao and Bruce (1997) when design matrices are orthonormal.

The thresholding rule in (2.8) involves two unknown parameters λ and a . In practice, we could search the best pair (λ, a) over the two-dimensional grids using some criteria, such as cross-validation and generalized cross-validation (Craven and Wahba 1979). Such an implementation can be computationally expensive. To implement tools in Bayesian risk analysis, we assume that for given a and λ , the prior distribution for θ is a normal distribution with zero mean and variance $a\lambda$. We compute the Bayes risk via numerical integration. Figure 5(a) depicts the Bayes risk as a function of a under the squared loss, for the universal thresholding $\lambda = \sqrt{2 \log(d)}$ (see Donoho and Johnstone, 1994a) with $d = 20, 40, 60$, and 100; and Figure 5(b) is for $d = 512, 1024, 2048$, and 4096. From Figure 5, (a) and (b), it can be seen that the Bayesian risks are not very sensitive to the values of a . It can be seen from Figure 5(a) that the Bayes risks achieve their minimums at $a \approx 3.7$ when the value of d is less than 100. This choice gives pretty good practical performance for various variable selection problems. Indeed, based on the simulations in Section 4.3, the choice of $a = 3.7$ works similarly to that chosen by the generalized cross-validation (GCV) method.

2.2 Performance of Thresholding Rules

We now compare the performance of the four previously stated thresholding rules. Marron, Adak, Johnstone, Neumann, and Patil (1998) applied the tool of risk analysis to understand the small sample behavior of the hard and soft thresholding rules. The closed forms for the L_2 risk functions $R(\hat{\theta}, \theta) = E(\hat{\theta} - \theta)^2$ were derived under the Gaussian model $Z \sim N(\theta, \sigma^2)$ for hard and soft thresholding rules by Donoho and Johnstone (1994b). The risk function of the SCAD thresholding rule can be found in Li (2000). To gauge the performance of the four thresholding rules, Figure 5(c) depicts their L_2 risk functions under the Gaussian model $Z \sim N(\theta, 1)$. To make the scale of the thresholding parameters roughly comparable, we took $\lambda = 2$ for the hard thresholding rule and adjusted the values of λ for the other thresholding rules so that their estimated values are the same when $\theta = 3$. The SCAD




Figure 4. Plot of $p'_\lambda(\theta)$ Functions Over $\theta > 0$ (a) for L_q Penalties, (b) the Hard Thresholding Penalty, and (c) the SCAD Penalty. In (a), the heavy line corresponds to L_1 , the dash-dot line corresponds to L_s , and the thin line corresponds to L_2 penalties.




Figure 5. Risk Functions of Proposed Procedures Under the Quadratic Loss. (a) Posterior risk functions of the SCAD under the prior $\theta \sim N(0, a\lambda)$ using the universal thresholding $\lambda = \sqrt{2 \log(d)}$ for four different values d : heavy line, $d = 20$; dashed line, $d = 40$; medium line, $d = 60$; thin line, $d = 100$. (b) Risk functions similar to those for (a): heavy line, $d = 572$; dashed line, $d = 1,024$; medium line, $d = 2,048$; thin line, $d = 4,096$. (c) Risk functions of the four different thresholding rules. The heavy, dashed, and solid lines denote minimum SCAD, hard, and soft thresholding rules, respectively.

performs favorably compared with the other two thresholding rules. This also can be understood via their corresponding penalty functions plotted in Figure 1. It is clear that the SCAD retains the good mathematical properties of the other two thresholding penalty functions. Hence, it is expected to perform the best. For general σ^2 , the picture is the same, except it is scaled vertically by σ^2 , and the θ axis should be replaced by θ/σ .

3. VARIABLE SELECTION VIA PENALIZED LIKELIHOOD

The methodology in the previous section can be applied directly to many other statistical contexts. In this section, we consider linear regression models, robust linear models, and likelihood-based generalized linear models. From now on, we assume that the design matrix $\mathbf{X} = (x_{ij})$ is standardized so that each column has mean 0 and variance 1.

3.1 Penalized Least Squares and Likelihood

In the classical linear regression model, the least squares estimate is obtained via minimizing the sum of squared residual errors. Therefore, (2.2) can be extended naturally to the situation in which design matrices are not orthonormal. Similar to (2.2), a form of penalized least squares is

$$\frac{1}{2}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^T(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) + n \sum_{j=1}^d p_{\lambda}(|\beta_j|). \quad (3.1)$$

Minimizing (3.1) with respect to $\boldsymbol{\beta}$ leads to a penalized least squares estimator of $\boldsymbol{\beta}$.

It is well known that the least squares estimate is not robust. We can consider the outlier-resistant loss functions such as the L_1 loss or, more generally, Huber's ψ function (see Huber 1981). Therefore, instead of minimizing (3.1), we minimize

$$\sum_{i=1}^n \psi(|y_i - \mathbf{x}_i \boldsymbol{\beta}|) + n \sum_{j=1}^d p_{\lambda}(|\beta_j|) \quad (3.2)$$

with respect to $\boldsymbol{\beta}$. This results in a penalized robust estimator for $\boldsymbol{\beta}$.

For generalized linear models, statistical inferences are based on underlying likelihood functions. The penalized maximum likelihood estimator can be used to select significant variables. Assume that the data $\{(\mathbf{x}_i, Y_i)\}$ are collected independently. Conditioning on \mathbf{x}_i , Y_i has a density $f_i(g(\mathbf{x}_i^T \boldsymbol{\beta}), y_i)$, where g is a known link function. Let $\ell_i = \log f_i$ denote the conditional log-likelihood of Y_i . A form of the penalized likelihood is

$$\sum_{i=1}^n \ell_i(g(\mathbf{x}_i^T \boldsymbol{\beta}), y_i) - n \sum_{j=1}^d p_{\lambda}(|\beta_j|).$$

Maximizing the penalized likelihood function is equivalent to minimizing

$$- \sum_{i=1}^n \ell_i(g(\mathbf{x}_i^T \boldsymbol{\beta}), y_i) + n \sum_{j=1}^d p_{\lambda}(|\beta_j|) \quad (3.3)$$

with respect to $\boldsymbol{\beta}$. To obtain a penalized maximum likelihood estimator of $\boldsymbol{\beta}$, we minimize (3.3) with respect to $\boldsymbol{\beta}$ for some thresholding parameter λ .

3.2 Sampling Properties and Oracle Properties

In this section, we establish the asymptotic theory for our nonconcave penalized likelihood estimator. Let

$$\boldsymbol{\beta}_0 = (\beta_{10}, \dots, \beta_{d0})^T = (\boldsymbol{\beta}_{10}^T, \boldsymbol{\beta}_{20}^T)^T.$$

Without loss of generality, assume that $\boldsymbol{\beta}_{20} = \mathbf{0}$. Let $I(\boldsymbol{\beta}_0)$ be the Fisher information matrix and let $I_1(\boldsymbol{\beta}_{10}, \mathbf{0})$ be the Fisher information knowing $\boldsymbol{\beta}_{20} = \mathbf{0}$. We first show that there exists a penalized likelihood estimator that converges at the rate

$$O_p(n^{-1/2} + a_n), \quad (3.4)$$

where $a_n = \max\{p'_{\lambda_n}(|\beta_{j0}|) : \beta_{j0} \neq 0\}$. This implies that for the hard thresholding and SCAD penalty functions, the penalized likelihood estimator is root- n consistent if $\lambda_n \rightarrow 0$. Furthermore, we demonstrate that such a root- n consistent estimator

must satisfy $\hat{\beta}_2 = \mathbf{0}$ and $\hat{\beta}_1$ is asymptotic normal with covariance matrix I_1^{-1} , if $n^{1/2}\lambda_n \rightarrow \infty$. This implies that the penalized likelihood estimator performs as well as if $\beta_{20} = \mathbf{0}$ were known. In language similar to Donoho and Johnstone (1994a), the resulting estimator performs as well as the oracle estimator, which knows in advance that $\beta_{20} = \mathbf{0}$.

The preceding oracle performance is closely related to the superefficiency phenomenon. Consider the simplest linear regression model $\mathbf{y} = \mathbf{1}_n\mu + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, I_n)$. A super-efficient estimate for μ is

$$\delta_n = \begin{cases} \bar{Y}, & \text{if } |\bar{Y}| \geq n^{-1/4}, \\ c\bar{Y}, & \text{if } |\bar{Y}| < n^{-1/4}, \end{cases}$$

owing to Hodges (see Lehmann 1983, p. 405). If we set c to 0, then δ_n coincides with the hard thresholding estimator with the thresholding parameter $\lambda_n = n^{-1/4}$. This estimator correctly estimates the parameter at point 0 without paying any price for estimating the parameter elsewhere.

We now state the result in a fairly general setting. To facilitate the presentation, we assume that the penalization is applied to every component of β . However, there is no extra difficulty to extend it to the case where some components (e.g., variance in the linear models) are not penalized.

Set $\mathbf{V}_i = (\mathbf{X}_i, Y_i)$, $i = 1, \dots, n$. Let $L(\beta)$ be the log-likelihood function of the observations $\mathbf{V}_1, \dots, \mathbf{V}_n$ and let $Q(\beta)$ be the penalized likelihood function $L(\beta) - n \sum_{j=1}^d p_{\lambda_n}(|\beta_j|)$. We state our theorems here, but their proofs are relegated to the Appendix, where the conditions for the theorems also can be found.

Theorem 1. Let $\mathbf{V}_1, \dots, \mathbf{V}_n$ be independent and identically distributed, each with a density $f(\mathbf{V}, \beta)$ (with respect to a measure μ) that satisfies conditions (A)–(C) in the Appendix. If $\max\{p''_{\lambda_n}(|\beta_{j0}|) : \beta_{j0} \neq 0\} \rightarrow 0$, then there exists a local maximizer $\hat{\beta}$ of $Q(\beta)$ such that $\|\hat{\beta} - \beta_0\| = O_p(n^{-1/2} + a_n)$, where a_n is given by (3.4).

It is clear from Theorem 1 that by choosing a proper λ_n , there exists a root- n consistent penalized likelihood estimator. We now show that this estimator must possess the sparsity property $\beta_2 = \mathbf{0}$, which is stated as follows.

Lemma 1. Let $\mathbf{V}_1, \dots, \mathbf{V}_n$ be independent and identically distributed, each with a density $f(\mathbf{V}, \beta)$ that satisfies conditions (A)–(C) in the Appendix. Assume that

$$\liminf_{n \rightarrow \infty} \liminf_{\theta \rightarrow 0^+} p'_{\lambda_n}(\theta)/\lambda_n > 0. \quad (3.5)$$

If $\lambda_n \rightarrow 0$ and $\sqrt{n}\lambda_n \rightarrow \infty$ as $n \rightarrow \infty$, then with probability tending to 1, for any given β_1 satisfying $\|\beta_1 - \beta_{10}\| = O_p(n^{-1/2})$ and any constant C ,

$$Q\left\{\begin{pmatrix} \beta_1 \\ \mathbf{0} \end{pmatrix}\right\} = \max_{\|\beta_2\| \leq Cn^{-1/2}} Q\left\{\begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix}\right\}.$$

Denote

$$\Sigma = \text{diag}\{p''_{\lambda_n}(|\beta_{10}|), \dots, p''_{\lambda_n}(|\beta_{s0}|)\}$$

and

$$\mathbf{b} = (p'_{\lambda_n}(|\beta_{10}|)\text{sgn}(\beta_{10}), \dots, p'_{\lambda_n}(|\beta_{s0}|)\text{sgn}(\beta_{s0}))^T,$$

where s is the number of components of β_{10} .

Theorem 2 (Oracle Property). Let $\mathbf{V}_1, \dots, \mathbf{V}_n$ be independent and identically distributed, each with a density $f(\mathbf{V}, \beta)$ satisfying conditions (A)–(C) in Appendix. Assume that the penalty function $p_{\lambda_n}(|\theta|)$ satisfies condition (3.5). If $\lambda_n \rightarrow 0$ and $\sqrt{n}\lambda_n \rightarrow \infty$ as $n \rightarrow \infty$, then with probability tending to 1, the root- n consistent local maximizers $\hat{\beta} = \begin{pmatrix} \hat{\beta}_1 \\ \hat{\beta}_2 \end{pmatrix}$ in Theorem 1 must satisfy:

- (a) Sparsity: $\hat{\beta}_2 = \mathbf{0}$.
- (b) Asymptotic normality:

$$\sqrt{n}(I_1(\beta_{10}) + \Sigma)\{\hat{\beta}_1 - \beta_{10} + (I_1(\beta_{10}) + \Sigma)^{-1}\mathbf{b}\} \rightarrow N\{\mathbf{0}, I_1(\beta_{10})\}$$

in distribution, where $I_1(\beta_{10}) = I_1(\beta_{10}, \mathbf{0})$, the Fisher information knowing $\beta_2 = \mathbf{0}$.

As a consequence, the asymptotic covariance matrix of $\hat{\beta}_1$ is

$$\frac{1}{n}\{I_1(\beta_{10}) + \Sigma\}^{-1}I_1(\beta_{10})\{I_1(\beta_{10}) + \Sigma\}^{-1},$$

which approximately equals $(1/n)I_1^{-1}(\beta_{10})$ for the thresholding penalties discussed in Section 2 if λ_n tends to 0.

Remark 1. For the hard and SCAD thresholding penalty functions, if $\lambda_n \rightarrow 0$, $a_n = 0$. Hence, by Theorem 2, when $\sqrt{n}\lambda_n \rightarrow \infty$, their corresponding penalized likelihood estimators possess the oracle property and perform as well as the maximum likelihood estimates for estimating β_1 knowing $\beta_2 = \mathbf{0}$. However, for the L_1 penalty, $a_n = \lambda_n$. Hence, the root- n consistency requires that $\lambda_n = O_p(n^{-1/2})$. On the other hand, the oracle property in Theorem 2 requires that $\sqrt{n}\lambda_n \rightarrow \infty$. These two conditions for LASSO cannot be satisfied simultaneously. Indeed, for the L_1 penalty, we conjecture that the oracle property does not hold. However, for L_q penalty with $q < 1$, the oracle property continues to hold with suitable choice of λ_n .

Now we briefly discuss the regularity conditions (A)–(C) for the generalized linear models (see McCullagh and Nelder 1989). With a canonical link, the condition distribution of Y given $\mathbf{X} = \mathbf{x}$ belongs to the canonical exponential family, that is, with a density function

$$f(y; \mathbf{x}, \beta) = c(y) \exp\left\{\frac{y\mathbf{x}^T\beta - b(\mathbf{x}^T\beta)}{a(\phi)}\right\}.$$

Clearly, the regularity conditions (A) are satisfied. The Fisher information matrix is

$$I(\beta) = E\{b''(\mathbf{x}^T\beta)\mathbf{xx}^T\}/a(\phi).$$

Therefore, if $E\{b''(\mathbf{x}^T\beta)\mathbf{xx}^T\}$ is finite and positive definite, then condition (B) holds. If for all β in some neighborhood of β_0 , $|b^{(3)}(\mathbf{x}^T\beta)| \leq M_0(\mathbf{x})$ for some function $M_0(\mathbf{x})$ satisfying $E_{\beta_0}\{M_0(\mathbf{x})X_jX_kX_l\} < \infty$ for all j, k, l , then condition (C) holds. For general link functions, similar conditions need to guarantee conditions (B) and (C). The mathematical derivation of those conditions does not involve any extra difficulty except more tedious notation. Results in Theorems 1 and 2 also can be established for the penalized least squares (3.1) and the penalized robust linear regression (3.2) under some mild regularity conditions. See Li (2000) for details.

3.3 A New Unified Algorithm

Tibshirani (1996) proposed an algorithm for solving constrained least squares problems of LASSO, whereas Fu (1998) provided a “shooting algorithm” for LASSO. See also LASSO2 submitted by Berwin Turlach at Statlib (<http://lib.stat.cmu.edu/S/>). In this section, we propose a new unified algorithm for the minimization problems (3.1), (3.2), and (3.3) via local quadratic approximations. The first term in (3.1), (3.2), and (3.3) may be regarded as a loss function of β . Denote it by $\ell(\beta)$. Then the expressions (3.1), (3.2), and (3.3) can be written in a unified form as

$$\ell(\beta) + n \sum_{j=1}^d p_{\lambda}(|\beta_j|). \quad (3.6)$$

The L_1 , hard thresholding, and SCAD penalty functions are singular at the origin, and they do not have continuous second order derivatives. However, they can be locally approximated by a quadratic function as follows. Suppose that we are given an initial value β_0 that is close to the minimizer of (3.6). If β_{j0} is very close to 0, then set $\hat{\beta}_j = 0$. Otherwise they can be locally approximated by a quadratic function as

$$[p_{\lambda}(|\beta_j|)]' = p'_{\lambda}(|\beta_j|) \text{sgn}(\beta_j) \approx \{p'_{\lambda}(|\beta_{j0}|)/|\beta_{j0}|\} \beta_j,$$

when $\beta_j \neq 0$. In other words,

$$p_{\lambda}(|\beta_j|) \approx p_{\lambda}(|\beta_{j0}|) + \frac{1}{2} \{p'_{\lambda}(|\beta_{j0}|)/|\beta_{j0}|\} (\beta_j^2 - \beta_{j0}^2),$$

$$\text{for } \beta_j \approx \beta_{j0}. \quad (3.7)$$

Figure 1 shows the L_1 , hard thresholding, and SCAD penalty functions, and their approximations on the right-hand side of (3.7) at two different values of β_{j0} . A drawback of this approximation is that once a coefficient is shrunk to zero, it will stay at zero. However, this method significantly reduces the computational burden.

If $\ell(\beta)$ is the L_1 loss as in (3.2), then it does not have continuous second order partial derivatives with respect to β . However, $\psi(|y - \mathbf{x}^T \beta|)$ in (3.2) can be analogously approximated by $\{\psi(y - \mathbf{x}^T \beta_0)/(y - \mathbf{x}^T \beta_0)^2\} (y - \mathbf{x}^T \beta)^2$, as long as the initial value β_0 of β is close to the minimizer. When some of the residuals $|y - \mathbf{x}^T \beta_0|$ are small, this approximation is not very good. See Section 3.4 for some slight modifications of this approximation.

Now assume that the log-likelihood function is smooth with respect to β so that its first two partial derivatives are continuous. Thus the first term in (3.6) can be locally approximated by a quadratic function. Therefore, the minimization problem (3.6) can be reduced to a quadratic minimization problem and the Newton–Raphson algorithm can be used. Indeed, (3.6) can be locally approximated (except for a constant term) by

$$\ell(\beta_0) + \nabla \ell(\beta_0)^T (\beta - \beta_0) + \frac{1}{2} (\beta - \beta_0)^T \nabla^2 \ell(\beta_0) (\beta - \beta_0) + \frac{1}{2} n \beta^T \Sigma_{\lambda}(\beta_0) \beta, \quad (3.8)$$

where

$$\nabla \ell(\beta_0) = \frac{\partial \ell(\beta_0)}{\partial \beta}, \quad \nabla^2 \ell(\beta_0) = \frac{\partial^2 \ell(\beta_0)}{\partial \beta \partial \beta^T},$$

$$\Sigma_{\lambda}(\beta_0) = \text{diag}\{p'_{\lambda}(|\beta_{10}|)/|\beta_{10}|, \dots, p'_{\lambda}(|\beta_{d0}|)/|\beta_{d0}|\}.$$

The quadratic minimization problem (3.8) yields the solution

$$\beta_1 = \beta_0 - \{\nabla^2 \ell(\beta_0) + n \Sigma_{\lambda}(\beta_0)\}^{-1} \{\nabla \ell(\beta_0) + n \mathbf{U}_{\lambda}(\beta_0)\}, \quad (3.9)$$

where $\mathbf{U}_{\lambda}(\beta_0) = \Sigma_{\lambda}(\beta_0) \beta_0$. When the algorithm converges, the estimator satisfies the condition

$$\frac{\partial \ell(\hat{\beta}_0)}{\partial \beta_j} + n p'_{\lambda}(|\hat{\beta}_{j0}|) \text{sgn}(\hat{\beta}_{j0}) = 0,$$

the penalized likelihood equation, for nonzero elements of $\hat{\beta}_0$. Specifically, for the penalized least squares problem (3.1), the solution can be found by iteratively computing the ridge regression

$$\beta_1 = \{\mathbf{X}^T \mathbf{X} + n \Sigma_{\lambda}(\beta_0)\}^{-1} \mathbf{X}^T \mathbf{y}.$$

Similarly we obtain the solution for (3.2) by iterating

$$\beta_1 = \{\mathbf{X}^T \mathbf{W} \mathbf{X} + \frac{1}{2} n \Sigma_{\lambda}(\beta_0)\}^{-1} \mathbf{X}^T \mathbf{W} \mathbf{y},$$

where $\mathbf{W} = \text{diag}\{\psi(|y_1 - \mathbf{x}_1^T \beta_0|)/(y_1 - \mathbf{x}_1^T \beta_0)^2, \dots, \psi(|y_n - \mathbf{x}_n^T \beta_0|)/(y_n - \mathbf{x}_n^T \beta_0)^2\}$.

As in the maximum likelihood estimation (MLE) setting, with the good initial value β_0 , the one-step procedure can be as efficient as the fully iterative procedure, namely, the penalized maximum likelihood estimator, when the Newton–Raphson algorithm is used (see Bickel 1975). Now regarding $\beta^{(k-1)}$ as a good initial value at the k th step, the next iteration also can be regarded as a one-step procedure and hence the resulting estimator still can be as efficient as the fully iterative method (see Robinson 1988 for the theory on the difference between the MLE and the k -step estimators). Therefore, estimators obtained by the aforementioned algorithm with a few iterations always can be regarded as a one-step estimator, which is as efficient as the fully iterative method. In this sense, we do not have to iterate the foregoing algorithm until convergence as long as the initial estimators are good enough. The estimators from the full models can be used as initial estimators, as long as they are not overly parameterized.

3.4 Standard Error Formula

The standard errors for the estimated parameters can be obtained directly because we are estimating parameters and selecting variables at the same time. Following the conventional technique in the likelihood setting, the corresponding sandwich formula can be used as an estimator for the covariance of the estimates β_1 , the nonvanishing component of β . That is,

$$\widehat{\text{cov}}(\hat{\beta}_1) = \{\nabla^2 \ell(\hat{\beta}_1) + n \Sigma_{\lambda}(\hat{\beta}_1)\}^{-1} \widehat{\text{cov}}\{\nabla \ell(\hat{\beta}_1)\} \times \{\nabla^2 \ell(\hat{\beta}_1) + n \Sigma_{\lambda}(\hat{\beta}_1)\}^{-1}. \quad (3.10)$$

Compare with Theorem 2(b). This formula is shown to have good accuracy for moderate sample sizes.

When the L_1 loss is used in the robust regression, some slight modifications are needed in the aforementioned algorithm and its corresponding sandwich formula. For $\psi(x) = |x|$, the diagonal elements of \mathbf{W} are $\{|r_i|^{-1}\}$ with $r_i = y_i - \mathbf{x}_i^T \boldsymbol{\beta}_0$. Thus, for a given current value of $\boldsymbol{\beta}_0$, when some of the residuals $\{r_i\}$ are close to 0, these points receive too much weight. Hence, we replace the weight by $(a_n + |r_i|)^{-1}$. In our implementations, we took a_n as the $2n^{-1/2}$ quantile of the absolute residuals $\{|r_i|, i = 1, \dots, n\}$. Thus, the constant a_n changes from iteration to iteration.

3.5 Testing Convergence of the Algorithm

We now demonstrate that our algorithm converges to the right solution. To this end, we took a 100-dimensional vector $\boldsymbol{\beta}$ consisting of 50 zeros and other nonzero elements generated from $N(0, 5^2)$, and used a 100×100 orthonormal design matrix \mathbf{X} . We then generated a response vector \mathbf{y} from the linear model (2.1). We chose an orthonormal design matrix for our testing case, because the penalized least squares has a closed form mathematical solution so that we can compare our output with the mathematical solution. Our experiment did show that the proposed algorithm converged to the right solution. It took MATLAB 0.27, 0.39, and 0.16 s to converge for the penalized least squares with the SCAD, L_1 , and hard thresholding penalties. The numbers of iterations were 30, 30, and 5, respectively for the penalized least squares with the SCAD, L_1 , and the hard thresholding penalty. In fact, after 10 iterations, the penalized least squares estimators are already very close to the true value.

4. NUMERICAL COMPARISONS

The purpose of this section is to compare the performance of the proposed approaches with existing ones and to test the accuracy of the standard error formula. We also illustrate our penalized likelihood approaches by a real data example. In all examples in this section, we computed the penalized likelihood estimate with the L_1 penalty, referred as to LASSO, by our algorithm rather than those of Tibshirani (1996) and Fu (1998).

4.1 Prediction and Model Error

The prediction error is defined as the average error in the prediction of Y given \mathbf{x} for future cases not used in the construction of a prediction equation. There are two regression situations, *X random* and *X controlled*. In the case that X is random, both Y and \mathbf{x} are randomly selected. In the controlled situation, design matrices are selected by experimenters and only y is random. For ease of presentation, we consider only the *X-random* case.

In *X-random* situations, the data (\mathbf{x}_i, Y_i) are assumed to be a random sample from their parent distribution (\mathbf{x}, Y) . Then, if $\hat{\mu}(\mathbf{x})$ is a prediction procedure constructed using the present data, the prediction error is defined as

$$\text{PE}(\hat{\mu}) = E\{Y - \hat{\mu}(\mathbf{x})\}^2,$$

where the expectation is taken only with respect to the new observation (\mathbf{x}, Y) . The prediction error can be decomposed as

$$\text{PE}(\hat{\mu}) = E\{Y - E(Y|\mathbf{x})\}^2 + E\{E(Y|\mathbf{x}) - \hat{\mu}(\mathbf{x})\}^2.$$

The first component is the inherent prediction error due to the noise. The second component is due to lack of fit to an underlying model. This component is called *model error* and is denoted $\text{ME}(\hat{\mu})$. The size of the model error reflects performances of different model selection procedures. If $Y = \mathbf{x}^T \boldsymbol{\beta} + \varepsilon$, where $E(\varepsilon|\mathbf{x}) = 0$, then $\text{ME}(\hat{\mu}) = (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})^T E(\mathbf{X}\mathbf{X}^T)(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})$.

4.2 Selection of Thresholding Parameters

To implement the methods described in Sections 2 and 3, we need to estimate the thresholding parameters λ and a (for the SCAD). Denote by $\boldsymbol{\theta}$ the tuning parameters to be estimated, that is, $\boldsymbol{\theta} = (\lambda, a)$ for the SCAD and $\boldsymbol{\theta} = \lambda$ for the other penalty functions. Here we discuss two methods of estimating $\boldsymbol{\theta}$: fivefold cross-validation and generalized cross-validation, as suggested by Breiman (1995), Tibshirani (1996), and Fu (1998).

For completeness, we now describe the details of the cross-validation and the generalized cross-validation procedures. Here we discuss only these two procedures for linear regression models. Extensions to robust linear models and likelihood-based linear models do not involve extra difficulties. The fivefold cross-validation procedure is as follows: Denote the full dataset by T , and denote cross-validation training and test set by $T - T^\nu$ and T^ν , respectively, for $\nu = 1, \dots, 5$. For each $\boldsymbol{\theta}$ and ν , we find the estimator $\hat{\boldsymbol{\beta}}^{(\nu)}(\boldsymbol{\theta})$ of $\boldsymbol{\beta}$ using the training set $T - T^\nu$. Form the cross-validation criterion as

$$\text{CV}(\boldsymbol{\theta}) = \sum_{\nu=1}^5 \sum_{(y_k, \mathbf{x}_k) \in T^\nu} \{y_k - \mathbf{x}_k^T \hat{\boldsymbol{\beta}}^{(\nu)}(\boldsymbol{\theta})\}^2.$$

We find a $\hat{\boldsymbol{\theta}}$ that minimizes $\text{CV}(\boldsymbol{\theta})$.

The second method is the generalized cross-validation. For linear regression models, we update the solution by

$$\boldsymbol{\beta}_1(\boldsymbol{\theta}) = \{\mathbf{X}^T \mathbf{X} + n\Sigma_\lambda(\boldsymbol{\beta}_0)\}^{-1} \mathbf{X}^T \mathbf{y}.$$

Thus the fitted value $\hat{\mathbf{y}}$ of \mathbf{y} is $\mathbf{X}\{\mathbf{X}^T \mathbf{X} + n\Sigma_\lambda(\boldsymbol{\beta}_0)\}^{-1} \mathbf{X}^T \mathbf{y}$, and

$$\mathbf{P}_X\{\hat{\boldsymbol{\beta}}(\boldsymbol{\theta})\} = \mathbf{X}\{\mathbf{X}^T \mathbf{X} + n\Sigma_\lambda(\hat{\boldsymbol{\beta}})\}^{-1} \mathbf{X}^T$$

can be regarded as a projection matrix. Define the number of effective parameters in the penalized least squares fit as $e(\boldsymbol{\theta}) = \text{tr}[\mathbf{P}_X\{\hat{\boldsymbol{\beta}}(\boldsymbol{\theta})\}]$. Therefore, the generalized cross-validation statistic is

$$\text{GCV}(\boldsymbol{\theta}) = \frac{1}{n} \frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}(\boldsymbol{\theta})\|^2}{\{1 - e(\boldsymbol{\theta})/n\}^2}$$

and $\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \{\text{GCV}(\boldsymbol{\theta})\}$.

4.3 Simulation Study

In the following examples, we numerically compare the proposed variable selection methods with the ordinary least squares, ridge regression, best subset selection, and nonnegative garrote (see Breiman 1995). All simulations are conducted using MATLAB codes. We directly used the constraint least squares module in MATLAB to find the nonnegative garrote estimate. As recommended in Breiman (1995), a fivefold cross-validation was used to estimate the tuning parameter for the nonnegative garrote. For the other model selection procedures, both fivefold cross-validation and generalized cross-validation were used to estimate thresholding parameters. However, their performance was similar. Therefore, we present only the results based on the generalized cross-validation.

Example 4.1 (Linear Regression). In this example we simulated 100 datasets consisting of n observations from the model

$$Y = \mathbf{x}^T \boldsymbol{\beta} + \sigma \varepsilon,$$

where $\boldsymbol{\beta} = (3, 1.5, 0, 0, 2, 0, 0, 0)^T$, and the components of \mathbf{x} and ε are standard normal. The correlation between x_i and x_j is $\rho^{|i-j|}$ with $\rho = .5$. This is a model used in Tibshirani (1996). First, we chose $n = 40$ and $\sigma = 3$. Then we reduced σ to 1 and finally increased the sample size to 60. The model error of the proposed procedures is compared to that of the least squares estimator. The median of relative model errors (MRME) over 100 simulated datasets are summarized in Table 1. The average of 0 coefficients is also reported in Table 1, in which the column labeled "Correct" presents the average restricted only to the true zero coefficients, and the column labeled "Incorrect" depicts the average of coefficients erroneously set to 0.

From Table 1, it can be seen that when the noise level is high and the sample size is small, LASSO performs the best and it significantly reduces both model error and model complexity, whereas ridge regression reduces only model error. The other variable selection procedures also reduce model error and model complexity. However, when the noise level is reduced, the SCAD outperforms the LASSO and the other penalized least squares. Ridge regression performs very poorly. The best subset selection method performs quite similarly to the SCAD. The nonnegative garrote performs quite well in various situations. Comparing the first two rows in Table 1, we can see that the choice of $a = 3.7$ is very reasonable. Therefore, we used it for other examples in this article. Table 1 also depicts the performance of an oracle estimator. From Table 1, it also can be seen that the performance of

Table 1. Simulation Results for the Linear Regression Model

		Avg. No. of 0 Coefficients	
Method	MRME (%)	Correct	Incorrect
$n = 40, \sigma = 3$			
SCAD ¹	72.90	4.20	.21
SCAD ²	69.03	4.31	.27
LASSO	63.19	3.53	.07
Hard	73.82	4.09	.19
Ridge	83.28	0	0
Best subset	68.26	4.50	.35
Garrote	76.90	2.80	.09
Oracle	33.31	5	0
$n = 40, \sigma = 1$			
SCAD ¹	54.81	4.29	0
SCAD ²	47.25	4.34	0
LASSO	63.19	3.51	0
Hard	69.72	3.93	0
Ridge	95.21	0	0
Best subset	53.60	4.54	0
Garrote	56.55	3.35	0
Oracle	33.31	5	0
$n = 60, \sigma = 1$			
SCAD ¹	47.54	4.37	0
SCAD ²	43.79	4.42	0
LASSO	65.22	3.56	0
Hard	71.11	4.02	0
Ridge	97.36	0	0
Best subset	46.11	4.73	0
Garrote	55.90	3.38	0
Oracle	29.82	5	0

NOTE: The value of a in SCAD¹ is obtained by generalized cross-validation, whereas the value of a in SCAD² is 3.7.

SCAD is expected to be as good as that of the oracle estimator as the sample size n increases (see Tables 5 and 6 for more details).

We now test the accuracy of our standard error formula (3.10). The median absolute deviation divided by .6745, denoted by SD in Table 2, of 100 estimated coefficients in the 100 simulations can be regarded as the true standard error. The median of the 100 estimated SD's, denoted by SD_m , and the median absolute deviation error of the 100 estimated standard errors divided by .6745, denoted by SD_{mad} , gauge the overall performance of the standard error formula (3.10). Table 2 presents the results for nonzero coefficients when the sample size $n = 60$. The results for the other two cases with $n = 40$ are similar. Table 2 suggests that the sandwich formula performs surprisingly well.

Table 2. Standard Deviations of Estimators for the Linear Regression Model ($n = 60$)

Method	$\hat{\beta}_1$		$\hat{\beta}_2$		$\hat{\beta}_5$	
	SD	SD_m (SD_{mad})	SD	SD_m (SD_{mad})	SD	SD_m (SD_{mad})
SCAD ¹	.166	.161 (.021)	.170	.160 (.024)	.148	.145 (.022)
SCAD ²	.161	.161 (.021)	.164	.161 (.024)	.151	.143 (.023)
LASSO	.164	.154 (.019)	.173	.150 (.022)	.153	.142 (.021)
Hard	.169	.161 (.022)	.174	.162 (.025)	.178	.148 (.021)
Best subset	.163	.155 (.020)	.152	.154 (.026)	.152	.139 (.020)
Oracle	.155	.154 (.020)	.147	.153 (.024)	.146	.137 (.019)

Table 3. Simulation Results for the Robust Linear Model

Method	MRME (%)	Avg. No. of 0 Coefficients	
		Correct	Incorrect
SCAD ($a = 3.7$)	35.52	4.71	0
LASSO	52.80	4.29	0
Hard	47.22	4.70	0
Best subset	41.53	4.85	.18
Oracle	23.33	5	0

Example 4.2 (Robust Regression). In this example, we simulated 100 datasets consisting of 60 observations from the model

$$Y = \mathbf{x}^T \boldsymbol{\beta} + \varepsilon,$$

where $\boldsymbol{\beta}$ and \mathbf{x} are the same as those in Example 1. The ε is drawn from the standard normal distribution with 10% outliers from the standard Cauchy distribution. The simulation results are summarized in Table 3. From Table 3, it can be seen that the SCAD somewhat outperforms the other procedures. The true and estimated standard deviations of estimators via sandwich formula (3.7) are shown in Table 4, which indicates that the performance of the sandwich formula is very good.

Example 4.3 (Logistic Regression). In this example, we simulated 100 datasets consisting of 200 observations from the model $Y \sim \text{Bernoulli}\{p(\mathbf{x}^T \boldsymbol{\beta})\}$, where $p(u) = \exp(u)/(1 + \exp(u))$, and the first six components of \mathbf{x} and $\boldsymbol{\beta}$ are the same as those in Example 1. The last two components of \mathbf{x} are independently identically distributed as a Bernoulli distribution with probability of success .5. All covariates are standardized. Model errors are computed via 1000 Monte Carlo simulations. The summary of simulation results is depicted in Tables 5 and 6. From Table 5, it can be seen that the performance of the SCAD is much better than the other two penalized likelihood estimates. Results in Table 6 show that our standard error estimator works well. From Tables 5 and 6, SCAD works as well as the oracle estimator in terms of the MRME and the accuracies of estimated standard errors.

We remark that the estimated SDs for the L_1 penalized likelihood estimator (LASSO) are consistently smaller than the SCAD, but its overall MRME is larger than that of the SCAD. This implies that the biases in the L_1 penalized likelihood estimators are large. This remark applies to all of our examples. Indeed, in Table 7, all coefficients were noticeably shrunk by LASSO.

Table 5. Simulation Results for the Logistic Regression

Method	MRME (%)	Avg. No. of 0 Coefficients	
		Correct	Incorrect
SCAD ($a = 3.7$)	26.48	4.98	.04
LASSO	53.14	3.76	0
Hard	59.06	4.27	0
Best subset	31.63	4.84	.01
Oracle	25.71	5	0

Example 4.4. In this example, we apply the proposed penalized likelihood methodology to the burns data, collected by the General Hospital Burn Center at the University of Southern California. The dataset consists of 981 observations. The binary response variable Y is 1 for those victims who survived their burns and 0 otherwise. Covariates $X_1 = \text{age}$, $X_2 = \text{sex}$, $X_3 = \log(\text{burn area} + 1)$, and binary variable $X_4 = \text{oxygen}$ (0 normal, 1 abnormal) were considered. Quadratic terms of X_1 and X_3 , and all interaction terms were included. The intercept term was added and the logistic regression model was fitted. The best subset variable selection with the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) was applied to this dataset. The unknown parameter λ was chosen by generalized cross-validation: it is .6932, .0015, and .8062 for the penalized likelihood estimates with the SCAD, L_1 , and hard thresholding penalties, respectively. The constant a in the SCAD was taken as 3.7. With the selected λ , the penalized likelihood estimator was obtained at the 6th, 28th, and 5th step iterations for the penalized likelihood with the SCAD, L_1 , and hard thresholding penalties, respectively. We also computed 10-step estimators, which took us less than 50 s for each penalized likelihood estimator, and the differences between the full iteration estimators and the 10-step estimators were less than 1%. The estimated coefficients and standard errors for the transformed data, based on the penalized likelihood estimators, are reported in Table 7.

From Table 7, the best subset procedure via minimizing the BIC scores chooses 5 out of 13 covariates, whereas the SCAD chooses 4 covariates. The difference between them is that the best subset keeps X_4 . Both SCAD and the best subset variable selection (BIC) do not include X_1^2 and X_3^2 in the selected subset, but both LASSO and the best subset variable selection (AIC) do. LASSO chooses the quadratic term of X_1 and X_3 rather than their linear terms. It also selects an interaction term X_2X_3 , which may not be statistically significant. LASSO shrinks noticeably large coefficients. In this example,

Table 4. Standard Deviations of Estimators for the Robust Regression Model

Method	$\hat{\beta}_1$		$\hat{\beta}_2$		$\hat{\beta}_5$	
	SD	SD_m (SD_{mad})	SD	SD_m (SD_{mad})	SD	SD_m (SD_{mad})
SCAD	.167	.171 (.018)	.185	.176 (.022)	.165	.155 (.020)
LASSO	.158	.165 (.022)	.159	.167 (.020)	.182	.154 (.019)
Hard	.179	.168 (.018)	.176	.176 (.025)	.157	.154 (.020)
Best subset	.198	.172 (.023)	.185	.175 (.024)	.199	.152 (.023)
Oracle	.163	.199 (.040)	.156	.202 (.043)	.166	.177 (.037)

Table 6. Standard Deviations of Estimators for the Logistic Regression

Method	$\hat{\beta}_1$		$\hat{\beta}_2$		$\hat{\beta}_5$	
	SD	SD _m (SD _{mad})	SD	SD _m (SD _{mad})	SD	SD _m (SD _{mad})
SCAD ($a = 3.7$)	.571	.538 (.107)	.383	.372 (.061)	.432	.398 (.065)
LASSO	.310	.379 (.037)	.285	.284 (.019)	.244	.287 (.019)
Hard	.675	.561 (.126)	.428	.400 (.062)	.467	.421 (.079)
Best subset	.624	.547 (.121)	.398	.383 (.067)	.468	.412 (.077)
Oracle	.553	.538 (.103)	.374	.373 (.060)	.432	.398 (.064)

the penalized likelihood with the hard thresholding penalty retains too many predictors. Particularly, it selects variables X_2 and X_2X_3 .

5. CONCLUSION

We proposed a variable selection method via penalized likelihood approaches. A family of penalty functions was introduced. Rates of convergence of the proposed penalized likelihood estimators were established. With proper choice of regularization parameters, we have shown that the proposed estimators perform as well as the oracle procedure for variable selection. The methods were shown to be effective and the standard errors were estimated with good accuracy. A unified algorithm was proposed for minimizing penalized likelihood function, which is usually a sum of convex and concave functions. Our algorithm is backed up by statistical theory and hence gives estimators with good statistical properties. Compared with the best subset method, which is very time consuming, the newly proposed methods are much faster, more effective, and have strong theoretical backup. They select variables simultaneously via optimizing a penalized likelihood, and hence the standard errors of estimated parameters can be estimated accurately. The LASSO proposed by Tibshirani (1996) is a member of this penalized likelihood family with L_1 penalty. It has good performance when the noise to signal ratio is large, but the bias created by this approach is noticeably large. See also the remarks in Example 4.3. The penalized likelihood with the SCAD penalty function gives the best performance in selecting significant variables without creating

excessive biases. The approach proposed here can be applied to other statistical contexts without any extra difficulties.

APPENDIX: PROOFS

Before we present the proofs of the theorems, we first state some regularity conditions. Denote by Ω the parameter space for β .

Regularity Conditions

(A) The observations V_i are independent and identically distributed with probability density $f(V, \beta)$ with respect to some measure μ . $f(V, \beta)$ has a common support and the model is identifiable. Furthermore, the first and second logarithmic derivatives of f satisfying the equations

$$E_{\beta} \left[\frac{\partial \log f(V, \beta)}{\partial \beta_j} \right] = 0 \quad \text{for } j = 1, \dots, d$$

and

$$\begin{aligned} I_{jk}(\beta) &= E_{\beta} \left[\frac{\partial}{\partial \beta_j} \log f(V, \beta) \frac{\partial}{\partial \beta_k} \log f(V, \beta) \right] \\ &= E_{\beta} \left[- \frac{\partial^2}{\partial \beta_j \partial \beta_k} \log f(V, \beta) \right]. \end{aligned}$$

(B) The Fisher information matrix

$$I(\beta) = E \left\{ \left[\frac{\partial}{\partial \beta} \log f(V, \beta) \right] \left[\frac{\partial}{\partial \beta} \log f(V, \beta) \right]^T \right\}$$

is finite and positive definite at $\beta = \beta_0$.

(C) There exists an open subset ω of Ω that contains the true parameter point β_0 such that for almost all V the density $f(V, \beta)$

Table 7. Estimated Coefficients and Standard Errors for Example 4.4

Method	MLE	Best Subset (AIC)	Best Subset (BIC)	SCAD	LASSO	Hard
Intercept	5.51 (.75)	4.81 (.45)	6.12 (.57)	6.09 (.29)	3.70 (.25)	5.88 (.41)
X_1	-8.83 (2.97)	-6.49 (1.75)	-12.15 (1.81)	-12.24 (.08)	0 (—)	-11.32 (1.1)
X_2	2.30 (2.00)	0 (—)	0 (—)	0 (—)	0 (—)	2.21 (1.41)
X_3	-2.77 (3.43)	0 (—)	-6.93 (.79)	-7.00 (.21)	0 (—)	-4.23 (.64)
X_4	-1.74 (1.41)	.30 (.11)	-.29 (.11)	0 (—)	-.28 (.09)	-1.16 (1.04)
X_1^2	-.75 (.61)	-1.04 (.54)	0 (—)	0 (—)	-1.71 (.24)	0 (—)
X_3^2	-2.70 (2.45)	-4.55 (.55)	0 (—)	0 (—)	-2.67 (.22)	-1.92 (.95)
X_1X_2	.03 (.34)	0 (—)	0 (—)	0 (—)	0 (—)	0 (—)
X_1X_3	7.46 (2.34)	5.69 (1.29)	9.83 (1.63)	9.84 (.14)	.36 (.22)	9.06 (.96)
X_1X_4	.24 (.32)	0 (—)	0 (—)	0 (—)	0 (—)	0 (—)
X_2X_3	-2.15 (1.61)	0 (—)	0 (—)	0 (—)	-0.10 (.10)	-2.13 (1.27)
X_2X_4	-.12 (.16)	0 (—)	0 (—)	0 (—)	0 (—)	0 (—)
X_3X_4	1.23 (1.21)	0 (—)	0 (—)	0 (—)	0 (—)	.82 (1.01)

admits all third derivatives $(\partial f(\mathbf{V}, \boldsymbol{\beta})) / (\partial \beta_j \partial \beta_k \partial \beta_l)$ for all $\boldsymbol{\beta} \in \omega$. Furthermore, there exist functions M_{jkl} such that

$$\left| \frac{\partial^3}{\partial \beta_j \partial \beta_k \partial \beta_l} \log f(\mathbf{V}, \boldsymbol{\beta}) \right| \leq M_{jkl}(\mathbf{V}) \quad \text{for all } \boldsymbol{\beta} \in \omega,$$

where $m_{jkl} = E_{\boldsymbol{\beta}_0}[M_{jkl}(\mathbf{V})] < \infty$ for j, k, l .

These regularity conditions guarantee asymptotic normality of the ordinary maximum likelihood estimates. See, for example, Lehmann (1983).

Proof of Theorem 1

Let $\alpha_n = n^{-1/2} + a_n$. We want to show that for any given $\varepsilon > 0$, there exists a large constant C such that

$$P \left\{ \sup_{\|\mathbf{u}\|=C} Q(\boldsymbol{\beta}_0 + \alpha_n \mathbf{u}) < Q(\boldsymbol{\beta}_0) \right\} \geq 1 - \varepsilon. \quad (\text{A.1})$$

This implies with probability at least $1 - \varepsilon$ that there exists a local maximum in the ball $\{\boldsymbol{\beta}_0 + \alpha_n \mathbf{u} : \|\mathbf{u}\| \leq C\}$. Hence, there exists a local maximizer such that $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\| = O_p(\alpha_n)$.

Using $p_{\lambda_n}(0) = 0$, we have

$$\begin{aligned} D_n(\mathbf{u}) &\equiv Q(\boldsymbol{\beta}_0 + \alpha_n \mathbf{u}) - Q(\boldsymbol{\beta}_0) \\ &\leq L(\boldsymbol{\beta}_0 + \alpha_n \mathbf{u}) - L(\boldsymbol{\beta}_0) - n \sum_{j=1}^s \{p_{\lambda_n}(|\beta_{j0} + \alpha_n u_j|) - p_{\lambda_n}(|\beta_{j0}|)\}, \end{aligned}$$

where s is the number of components of $\boldsymbol{\beta}_{10}$. Let $L'(\boldsymbol{\beta}_0)$ be the gradient vector of L . By the standard argument on the Taylor expansion of the likelihood function, we have

$$\begin{aligned} D_n(\mathbf{u}) &\leq \alpha_n L'(\boldsymbol{\beta}_0)^T \mathbf{u} - \frac{1}{2} \mathbf{u}^T I(\boldsymbol{\beta}_0) \mathbf{u} \alpha_n^2 \{1 + o_p(1)\} \\ &\quad - \sum_{j=1}^s [n \alpha_n p'_{\lambda_n}(|\beta_{j0}|) \text{sgn}(\beta_{j0}) u_j \\ &\quad + n \alpha_n^2 p''_{\lambda_n}(|\beta_{j0}|) u_j^2 \{1 + o(1)\}]. \end{aligned} \quad (\text{A.2})$$

Note that $n^{-1/2} L'(\boldsymbol{\beta}_0) = O_p(1)$. Thus, the first term on the right-hand side of (A.2) is on the order $O_p(n^{1/2} \alpha_n) = O_p(n \alpha_n^2)$. By choosing a sufficiently large C , the second term dominates the first term uniformly in $\|\mathbf{u}\| = C$. Note that the third term in (A.2) is bounded by

$$\sqrt{s} n \alpha_n a_n \|\mathbf{u}\| + n \alpha_n^2 \max\{|p''_{\lambda_n}(|\beta_{j0}|)| : \beta_{j0} \neq 0\} \|\mathbf{u}\|^2.$$

This is also dominated by the second term of (A.2). Hence, by choosing a sufficiently large C , (A.1) holds. This completes the proof of the theorem.

Proof of Lemma 1

It is sufficient to show that with probability tending to 1 as $n \rightarrow \infty$, for any $\boldsymbol{\beta}_1$ satisfying $\boldsymbol{\beta}_1 - \boldsymbol{\beta}_{10} = O_p(n^{-1/2})$ and for some small $\varepsilon_n = Cn^{-1/2}$ and $j = s+1, \dots, d$,

$$\frac{\partial Q(\boldsymbol{\beta})}{\partial \beta_j} < 0 \quad \text{for } 0 < \beta_j < \varepsilon_n \quad (\text{A.3})$$

$$> 0 \quad \text{for } -\varepsilon_n < \beta_j < 0. \quad (\text{A.4})$$

To show (A.3), by Taylor's expansion, we have

$$\begin{aligned} \frac{\partial Q(\boldsymbol{\beta})}{\partial \beta_j} &= \frac{\partial L(\boldsymbol{\beta})}{\partial \beta_j} - n p'_{\lambda_n}(|\beta_j|) \text{sgn}(\beta_j) \\ &= \frac{\partial L(\boldsymbol{\beta}_0)}{\partial \beta_j} + \sum_{l=1}^d \frac{\partial^2 L(\boldsymbol{\beta}_0)}{\partial \beta_j \partial \beta_l} (\beta_l - \beta_{l0}) + \sum_{l=1}^d \sum_{k=1}^d \frac{\partial^3 L(\boldsymbol{\beta}^*)}{\partial \beta_j \partial \beta_l \partial \beta_k} \\ &\quad \times (\beta_l - \beta_{l0})(\beta_k - \beta_{k0}) - n p'_{\lambda_n}(|\beta_j|) \text{sgn}(\beta_j), \end{aligned}$$

where $\boldsymbol{\beta}^*$ lies between $\boldsymbol{\beta}$ and $\boldsymbol{\beta}_0$. Note that by the standard arguments,

$$n^{-1} \frac{\partial L(\boldsymbol{\beta}_0)}{\partial \beta_j} = O_p(n^{-1/2})$$

and

$$\frac{1}{n} \frac{\partial^2 L(\boldsymbol{\beta}_0)}{\partial \beta_j \partial \beta_l} = E \left\{ \frac{\partial^2 L(\boldsymbol{\beta}_0)}{\partial \beta_j \partial \beta_l} \right\} + o_p(1).$$

By the assumption that $\boldsymbol{\beta} - \boldsymbol{\beta}_0 = O_p(n^{-1/2})$, we have

$$\frac{\partial Q(\boldsymbol{\beta})}{\partial \beta_j} = n \lambda_n \{-\lambda_n^{-1} p'_{\lambda_n}(|\beta_j|) \text{sgn}(\beta_j) + O_p(n^{-1/2}/\lambda_n)\}.$$

Whereas $\liminf_{n \rightarrow \infty} \liminf_{\theta \rightarrow 0^+} \lambda_n^{-1} p'_{\lambda_n}(\theta) > 0$ and $n^{-1/2}/\lambda_n \rightarrow 0$, the sign of the derivative is completely determined by that of β_j . Hence, (A.3) and (A.4) follow. This completes the proof.

Proof of Theorem 2

It follows by Lemma 1 that part (a) holds. Now we prove part (b). It can be shown easily that there exists a $\hat{\boldsymbol{\beta}}_1$ in Theorem 1 that is a root- n consistent local maximizer of $Q(\cdot)$, which is regarded as a function of $\boldsymbol{\beta}_1$, and that satisfies the likelihood equations

$$\frac{\partial Q(\boldsymbol{\beta})}{\partial \beta_j} \bigg|_{\boldsymbol{\beta} = (\hat{\boldsymbol{\beta}}_1)} = 0 \quad \text{for } j = 1, \dots, s. \quad (\text{A.5})$$

Note that $\hat{\boldsymbol{\beta}}_1$ is a consistent estimator,

$$\begin{aligned} \frac{\partial L(\boldsymbol{\beta})}{\partial \beta_j} \bigg|_{\boldsymbol{\beta} = (\hat{\boldsymbol{\beta}}_1)} &- n p'_{\lambda_n}(|\hat{\beta}_j|) \text{sgn}(\hat{\beta}_j) \\ &= \frac{\partial L(\boldsymbol{\beta}_0)}{\partial \beta_j} + \sum_{l=1}^s \left\{ \frac{\partial^2 L(\boldsymbol{\beta}_0)}{\partial \beta_j \partial \beta_l} + o_p(1) \right\} (\hat{\beta}_l - \beta_{l0}) \\ &\quad - n (p'_{\lambda_n}(|\beta_{j0}|) \text{sgn}(\beta_{j0}) + \{p''_{\lambda_n}(|\beta_{j0}|) + o_p(1)\} (\hat{\beta}_j - \beta_{j0})). \end{aligned}$$

It follows by Slutsky's theorem and the central limit theorem that

$$\sqrt{n} (I_1(\boldsymbol{\beta}_{10}) + \Sigma) \{\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_{10} + (I_1(\boldsymbol{\beta}_{10}) + \Sigma)^{-1} \mathbf{b}\} \rightarrow N\{\mathbf{0}, I_1(\boldsymbol{\beta}_{10})\}$$

in distribution.

[Received March 2000. Revised December 2000.]

REFERENCES

- Antoniadis, A. (1997), "Wavelets in Statistics: A Review" (with discussion), *Journal of the Italian Statistical Association*, 6, 97–144.
- Antoniadis, A., and Fan, J. (2001), "Regularization of Wavelets Approximations," *Journal of the American Statistical Association*, 96, 939–967.
- Bickel, P. J. (1975), "One-Step Huber Estimates in Linear Models," *Journal of the American Statistical Association*, 70, 428–433.
- Breiman, L. (1995), "Better Subset Regression Using the Nonnegative Garrote," *Technometrics*, 37, 373–384.
- Breiman, L. (1996), "Heuristics of Instability and Stabilization in Model Selection," *The Annals of Statistics*, 24, 2350–2383.
- Craven, P., and Wahba, G. (1979), "Smoothing Noisy Data With Spline Functions: Estimating the Correct Degree of Smoothing by the Method of Generalized Cross-Validation," *Numerische Mathematik*, 31, 377–403.
- Donoho, D. L., and Johnstone, I. M. (1994a), "Ideal Spatial Adaptation by Wavelet Shrinkage," *Biometrika*, 81, 425–455.
- (1994b), "Minimax Risk Over ℓ_p Balls for ℓ_q Error," *Probability Theory and Related Fields*, 99, 277–303.
- Fan, J. (1997), "Comments on 'Wavelets in Statistics: A Review' by A. Antoniadis," *Journal of the Italian Statistical Association*, 6, 131–138.
- Frank, I. E., and Friedman, J. H. (1993), "A Statistical View of Some Chemometrics Regression Tools," *Technometrics*, 35, 109–148.
- Fu, W. J. (1998), "Penalized Regression: The Bridge Versus the LASSO," *Journal of Computational and Graphical Statistics*, 7, 397–416.

- Gao, H. Y., and Bruce, A. G. (1997), "WaveShrink With Firm Shrinkage," *Statistica Sinica*, 7, 855–874.
- Green, P. J., and Silverman, B. W. (1994), *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*, London: Chapman and Hall.
- Huber, P. (1981), *Robust Estimation*, New York: Wiley.
- Lehmann, E. L. (1983), *Theory of Point Estimation*. Pacific Grove, CA: Wadsworth and Brooks/Cole.
- Li, R. (2000), "High-Dimensional Modeling via Nonconcave Penalized Likelihood and Local Likelihood," unpublished Ph.D. dissertation, University of North Carolina at Chapel Hill, Dept. of Statistics.
- Marron, J. S., Adak, S., Johnstone, I. M., Neumann, M. H., and Patil, P. (1998), "Exact Risk Analysis of Wavelet Regression," *Journal Computational and Graphical Statistics*, 7, 278–309.
- McCullagh, P., and Nelder, J. A. (1989), *Generalized Linear Models*. 2nd ed., London: Chapman and Hall.
- Robinson, P. M. (1988), "The Stochastic Difference Between Econometrics and Statistics," *Econometrica*, 56, 531–547.
- Stone, C. J., Hansen, M., Kooperberg, C., and Truong, Y. K. (1997), "Polynomial Splines and Their Tensor Products in Extended Linear Modeling" (with discussion), *The Annals of Statistics* 25, 1371–1470.
- Tibshirani, R. J. (1996), "Regression Shrinkage and Selection via the LASSO," *Journal of the Royal Statistical Society, Ser. B*, 58, 267–288.
- Tibshirani, R. J. (1997), "The LASSO Method for Variable Selection in the Cox Model," *Statistics in Medicine*, 16, 385–395.
- Wahba, G. (1990), *Spline Models for Observational Data*, Philadelphia: SIAM.