

APPLE: approximate path for penalized likelihood estimators

Yi Yu · Yang Feng

Received: 19 September 2012 / Accepted: 2 May 2013 / Published online: 1 June 2013
© Springer Science+Business Media New York 2013

Abstract In high-dimensional data analysis, penalized likelihood estimators are shown to provide superior results in both variable selection and parameter estimation. A new algorithm, APPLE, is proposed for calculating the Approximate Path for Penalized Likelihood Estimators. Both convex penalties (such as LASSO) and folded concave penalties (such as MCP) are considered. APPLE efficiently computes the solution path for the penalized likelihood estimator using a hybrid of the modified predictor-corrector method and the coordinate-descent algorithm. APPLE is compared with several well-known packages via simulation and analysis of two gene expression data sets.

Keywords APPLE · LASSO · MCP · Penalized likelihood estimator · Solution path

1 Introduction

Variable selection is a vital tool in statistical analysis of high-dimensional data. Typically, a large number of potential predictors are included during the first stage of modeling, in order to avoid missing important links between a predictor and the outcome. This practice has become more popular in recent years for two primary reasons. First, in

many recently promising fields, such as bioinformatics, genetics and finance, more and more high-throughput and high-dimensional data are being generated. Secondly, low cost and easy implementation for data collection and storage have made problems for which the number of variables is large, in comparison to the sample size, possible to be handled.

In order to provide more representative and reasonable applications of models in a mathematical framework, we often seek a smaller subset of important variables. The first attempt to variable selection was the ℓ_0 -type regularization methods, including AIC (Akaike 1973), C_p (Mallows 1973) and BIC (Schwarz 1978), which work well in low-dimensional cases. In addition, they also exhibit good sampling properties (Barron et al. 1999). However, searching all the possible subsets can be unstable (Breiman and Gao 1996), and in high-dimensional settings, the combinatorial problem has NP-complexity, which is computationally prohibitive. As a result, numerous attempts have been made to modify the ℓ_0 -type regularization to reduce the computational burden. The most popular penalized regression method is LASSO (Tibshirani 1996) or equivalently Basis Pursuit (Chen and Donoho 1994). Being a convex penalty, it is computationally convenient, but lacks the oracle property and shrinks estimators regardless of importance. Hence, some folded concave penalties have been proposed in order to yield better performance, such as SCAD (Fan and Li 2001) and MCP (Zhang 2010). We refer to Fan et al. (2012) for the detailed definition of folded concave penalties. Also, for generalized linear models (GLM), penalized likelihood methods have been studied for high-dimensional variable selection, for example in Fan and Li (2001) and van de Geer (2008). We refer to Fan and Lv (2010) for a review of variable selection in high-dimensionality.

Y. Yu
School of Mathematical Sciences, Fudan University, Shanghai
200433, China
e-mail: yuyi@fudan.edu.cn

Y. Feng (✉)
Department of Statistics, Columbia University, New York,
10027 NY, USA
e-mail: yangfeng@stat.columbia.edu

Throughout the paper, we assume we have i.i.d. observations $(x_i, y_i), i = 1, \dots, n$, where x_i is a p -dimensional predictor and y_i is the response. We further assume the conditional distribution of y given x belongs to an exponential family with canonical link, that is, it has the following density function

$$f(y; \mathbf{x}, \boldsymbol{\beta}) = c(y) \exp \left[\frac{y\theta - b(\theta)}{a(\phi)} \right], \quad (1)$$

where $\theta = \mathbf{x}'\boldsymbol{\beta}$ and $\phi \in (0, \infty)$ is the dispersion parameter.

In view of (1), the log-likelihood of the sample is given, up to an affine transformation, by

$$\ell(y; \boldsymbol{\beta}) = n^{-1} \sum_{i=1}^n [y_i \theta_i - b(\theta_i)].$$

Here, we are interested in estimating the p -dimensional vector $\boldsymbol{\beta}$, and the penalized likelihood estimator is defined as

$$\widehat{\boldsymbol{\beta}}(\lambda) = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \{-\ell(y; \boldsymbol{\beta}) + p_\lambda(\boldsymbol{\beta})\},$$

where $p_\lambda(\cdot)$ is the LASSO or folded concave penalty function and $\lambda > 0$ is the regularization parameter.

Developing an efficient algorithm for calculating the solution path of the coefficient vector $\widehat{\boldsymbol{\beta}}(\lambda)$, as λ varies along a possible set of values, is very desirable. There is a vast literature on calculating such a path for penalized linear regression. For the convex penalty LASSO, least angle regression (LARS) (Efron et al. 2004), or homotopy (Osborne et al. 2000) are efficient methods for computing the entire path of LASSO solutions in the linear regression case. For folded concave penalties including SCAD and MCP, Fan and Li (2001) used the local quadratic approximation (LQA); Zou and Li (2008) proposed the local linear approximation (LLA), which makes a local linear approximation to the penalty function, thereby yielding an objective function that can be optimized by using the LARS algorithm. Zhang (2010) proposed the penalized linear unbiased selection (PLUS), which is designed for the linear regression penalized by quadratic spline penalties, including LASSO, SCAD and MCP. More recently, coordinate descent methods have received considerable attention in high-dimensional settings, including Fu (1998), Shevade and Keerthi (2003), Krishnapuram et al. (2005), Genkin et al. (2007), Friedman et al. (2007), Wu and Lange (2008), among others. Other work on penalized linear regression includes Hastie et al. (2004), Daubechies et al. (2004), Kim et al. (2007) and Wei and Zhu (2012).

Different from linear regression, derivatives of the log-likelihood in GLM are changing with respect to the regularization parameter λ . There has been major research on calculating the solution path for penalized likelihood estimators in the GLM setting. Park and Hastie (2007) proposed

the *glm*path algorithm. They considered the solution $\widehat{\boldsymbol{\beta}}$ as a function of λ , and used a linear approximation of this function to update the estimator $\widehat{\boldsymbol{\beta}}$. They selected the step length in decreasing λ by using an approximate smallest length that will change the active set of variables. Yuan and Zou (2009) approximated the loss function by a quadratic spline and showed a generalized LARS algorithm is suitable for solution path computation. It is worth to point out that their method can also be extended to more general regularization framework, including a generalization of the elastic net and a new method that effectively exploits the so-called “support vectors” in kernel logistic regression. Friedman et al. (2010) proposed a coordinate descent algorithm for penalized GLM, in which they quadratically approximate the log-likelihood function and sequentially solve the resulting penalized weighted least squares problem on a grid of λ values. Wu (2011) proposed an ordinary differential equation-based solution path algorithm, which used quasi-likelihood instead of likelihood models, in order to use LARS more straightforwardly. However, all the numerical results in Wu (2011) are based on the small p large n setting. Breheny and Huang (2011) adopted a coordinate descent algorithm for MCP and SCAD penalized GLM. Like Friedman et al. (2010), they used a quadratic approximation to the log-likelihood part and then used coordinate descent to update the regression parameter estimator. For the MCP penalty (Zhang 2010), the tuning parameter γ is used to adjust the concavity of the penalty. The smaller γ is, the more concave the penalty is, which means finding a global minimizer is more difficult; but on the other hand, it results in less biased estimators. The tuning parameter γ can be changed freely from $1+$ to ∞ . In the GLM case, Breheny and Huang (2011) proposed the adaptive rescaling, which allows the range of the parameter γ to be as wide as it can be for the linear regression case. Other related papers include Zhu and Hastie (2004), Lee et al. (2006), Rosset and Zhu (2007) and Meier et al. (2008).

In this work, we propose a new path algorithm, the *approximate path for penalized likelihood estimators* (APPLE), under the setting of high-dimensional GLM. Different from linear regression, it is often difficult to get explicit solutions in GLM. Taking accuracy and feasibility into account, instead of linear approximation of the corresponding change in $\boldsymbol{\beta}$ with the decrease in λ , which is used by most of the previous work, we use quadratic approximation to get a warm-start in updating. Then targeting on the KKT conditions, we perform a correction by optimizing a convex problem. Inspired by the adaptive rescaling in Breheny and Huang (2011), we develop a modified concavity adaptation method for MCP when updating the solution, which is shown to have better performance when γ is small. In this paper, not only path algorithms for LASSO penalized GLM are derived, but also path algorithms for folded concave penalized GLM, which have appeared in few of the

previous work. Here we mainly focus on MCP as an example of folded concave penalty, but it can be easily extended to other quadratic spline penalty functions.

For LASSO, we detect the active set through the KKT conditions like most of the previous work. However, for some folded concave penalties, such as MCP, by fixing λ and the concavity parameter γ , the value of the derivative of the penalty function decreases towards zero as the absolute value of the estimator increases towards $\lambda\gamma$. We introduce a modified active set detection method, which has not appeared in any of the previous work.

The rest of the paper is organized as follows. In Sects. 2 and 3, we introduce the path algorithm APPLE for the LASSO and MCP penalties, respectively. We conduct simulation studies in Sect. 4 and two real data examples are presented in Sect. 5. A short summary is given in Sect. 6, while the technical details for logistic regression and Poisson regression are presented in the Appendix.

2 APPLE with LASSO penalty

LASSO (Tibshirani 1996) is a popular method for regression that uses an ℓ_1 -penalty to achieve simultaneous variable selection and parameter estimation. The idea has been broadly applied in GLM, where the problem is to minimize a convex function. In this section, we describe the details of the APPLE algorithm for LASSO penalized GLM.

2.1 Problem setup

Let $\{(x_i, y_i), i = 1, \dots, n\}$ be n i.i.d. pairs of p predictors and a response as described in the introduction. By adding an additional column of $\mathbf{1}$'s to the design matrix \mathbf{X} , the intercept β_0 is absorbed into the coefficient vector $\boldsymbol{\beta}$. We are interested in finding the maximum likelihood solution for $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_p)'$, with a penalization on the size of the ℓ_1 -norm of the coefficients excluding the intercept. With a little abuse of notation, we denote $\|\boldsymbol{\beta}\|_1 = \sum_{i=1}^p |\beta_j|$. Therefore, the optimization problem for a given λ is reduced to finding $\widehat{\boldsymbol{\beta}}$, which minimizes the following:

$$L_\lambda(\boldsymbol{\beta}) = -\ell(\mathbf{y}; \boldsymbol{\beta}) + \lambda \|\boldsymbol{\beta}\|_1$$

$$= -\frac{1}{n} \sum_{i=1}^n \{y_i \theta(\boldsymbol{\beta})_i - b(\theta(\boldsymbol{\beta})_i)\} + \lambda \|\boldsymbol{\beta}\|_1. \quad (2)$$

As is common in GLM, the function $b(\theta)$ is implicitly assumed to be twice continuously differentiable with $b''(\theta)$ always positive. It is straightforward to check that $L_\lambda(\cdot)$ is a convex function. Therefore, for a given λ , the unique minimizer $\widehat{\boldsymbol{\beta}}(\lambda)$ is the solution to the KKT conditions, which are

given as follows.

$$\begin{cases} \frac{\partial \ell}{\partial \beta_0} |_{\beta_0 = \hat{\beta}_0} = 0, \\ \frac{\partial \ell}{\partial \beta_j} |_{\beta_j = \hat{\beta}_j} = \lambda \text{sgn}(\hat{\beta}_j) & \text{for } j = 1, \dots, p, \text{ s.t. } \hat{\beta}_j \neq 0, \\ \left| \frac{\partial \ell}{\partial \beta_j} |_{\beta_j = \hat{\beta}_j} \right| \leq \lambda & \text{for } j = 1, \dots, p, \text{ s.t. } \hat{\beta}_j = 0. \end{cases} \quad (3)$$

2.2 Grid of penalty parameter

It is easy to notice from the KKT conditions that when

$$\lambda \geq \lambda_{\max} = \max_{1 \leq j \leq p} |\partial \ell / \partial \beta_j |_{\beta_j = 0},$$

$\hat{\beta}_j = 0$ for $1 \leq j \leq p$. As λ decreases from λ_{\max} to 0, $\widehat{\boldsymbol{\beta}} = \widehat{\boldsymbol{\beta}}(\lambda)$ changes from 0 (except for the intercept $\hat{\beta}_0$) to the MLE solution. However, the full MLE solution has poor predictive performance and lacks the sparsity property because of the high-dimensionality. Here, following Friedman et al. (2010) and Breheny and Huang (2011), we set the minimum value of λ to be $\lambda_{\min} = \delta \lambda_{\max}$ and construct a sequence of K values of λ decreasing from λ_{\max} to λ_{\min} on the logarithm scale. We denote the sequence of λ as λ_k , where $k = 1, \dots, K$. Typical values are $\delta = 0.01$ and $K = 100$.

2.3 Update

From the KKT conditions, we can see the relationship between $|\partial \ell / \partial \beta_j|$ and λ determines whether the variable β_j is activated or not. For λ_k , we define the active set A_k as follows,

$$A_k = \left\{ 1 \leq j \leq p : \left| \frac{\partial \ell}{\partial \beta_j} \right| \geq \lambda_k \right\} \cup \{0\}, \quad (4)$$

and the step size as $\Delta_k = \lambda_{k+1} - \lambda_k$. For a given λ_k and active set A_k , we update the active coordinates together using the quadratic approximation,

$$\widehat{\boldsymbol{\beta}}_{A_k}^{(k+1,0)} = \widehat{\boldsymbol{\beta}}_{A_k}^{(k)} + \mathbf{s}^{(k)} \cdot \Delta_k + \frac{1}{2} \mathbf{d}^{(k)} \cdot \Delta_k^2, \quad (5)$$

where $\mathbf{s}^{(k)}$ and $\mathbf{d}^{(k)}$ are the first and second derivatives of $\widehat{\boldsymbol{\beta}}^{(k)}$ with respect to λ , respectively, which are derived using the chain rule, i.e.

$$\begin{aligned} \left(\frac{\partial \ell}{\partial \boldsymbol{\beta}_{A_k}} \right) \Big|_{\boldsymbol{\beta} = \widehat{\boldsymbol{\beta}}^{(k)}} &= \lambda \text{sgn}(\widehat{\boldsymbol{\beta}}_{A_k}^{(k)}) \\ \Rightarrow \left(\frac{\partial^2 \ell}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T} \cdot \frac{\partial \boldsymbol{\beta}_{A_k}}{\partial \lambda} \right) \Big|_{\boldsymbol{\beta} = \widehat{\boldsymbol{\beta}}^{(k)}} &= \text{sgn}(\widehat{\boldsymbol{\beta}}_{A_k}^{(k)}) \\ \Rightarrow \mathbf{s}^{(k)} &= \left(\frac{\partial^2 \ell}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T} \right)^{-1} \Big|_{\boldsymbol{\beta} = \widehat{\boldsymbol{\beta}}^{(k)}} \cdot \text{sgn}(\widehat{\boldsymbol{\beta}}_{A_k}^{(k)}), \end{aligned}$$

and,

$$\mathbf{d}^{(k)} = \partial \left[\left(\frac{\partial^2 \ell}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T} \right)^{-1} \cdot \text{sgn}(\widehat{\boldsymbol{\beta}}_{A_k}^{(k)}) \right] / \left. \frac{\partial \lambda}{\partial \boldsymbol{\beta}} \right|_{\boldsymbol{\beta} = \widehat{\boldsymbol{\beta}}^{(k)}}.$$

The explicit formula for calculating $s^{(k)}$ and $\mathbf{d}^{(k)}$ are presented in the [Appendix](#) for both logistic regression and Poisson regression. Since the intercept is not penalized, the first coordinates of $s^{(k)}$ and $\mathbf{d}^{(k)}$ are both 0. Different from Park and Hastie (2007), where linear approximation was used, the quadratic approximation (5) is more accurate and computationally efficient as a warm-start. Keep in mind, that here, the coefficients for the variables outside of A_k are set to be 0 in $\widehat{\boldsymbol{\beta}}^{(k+1,0)}$. Additionally, the approximation (5) will typically cause a small deviation from the KKT conditions, which makes the following correction step necessary, in order to get the exact solution $\widehat{\boldsymbol{\beta}}^{(k+1)}$ at the current λ_{k+1} .

Here, we adopt two different correction methods depending on the current model size. To be more precise, at step k , we check the following inequality,

$$\#\{j : \widehat{\boldsymbol{\beta}}_j^{(k+1,0)} \neq 0\} \leq c\sqrt{n}, \quad (6)$$

where c is a user-specified constant. We set $c = 1$ in all our numerical examples. If (6) holds (i.e., the current solution is relatively sparse compared with the sample size), we use a Newton-Raphson correction, otherwise, a coordinate descent correction is applied. When the correction method is stopped by a convergence check, the last $\widehat{\boldsymbol{\beta}}^{(k+1,j)}$ is denoted as $\widehat{\boldsymbol{\beta}}^{(k+1)}$.

2.3.1 Newton-Raphson correction

Given the current solution $\widehat{\boldsymbol{\beta}}^{(k+1,0)}$, we use the following Newton-Raphson method to correct the estimate until convergence,

$$\widehat{\boldsymbol{\beta}}_{A_k}^{(k,j+1)} = \widehat{\boldsymbol{\beta}}_{A_k}^{(k,j)} - \left(\frac{\partial^2 L^{(k)}}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T} \right)^{-1} \left(\frac{\partial L^{(k)}}{\partial \boldsymbol{\beta}_{A_k}} \right).$$

Here, all the active variables are corrected together, which is different from coordinate descent method used in Friedman et al. (2010) and Breheny and Huang (2011). We notice in our simulation studies that when (6) holds, the Newton-Raphson type correction tends to be much faster than the coordinate descent correction method.

2.3.2 Coordinate descent correction

When (6) does not hold (i.e., the number of active variables is relatively large), the Newton-Raphson method involves inverting a big matrix $(\partial^2 L / \partial \boldsymbol{\beta}^2)$, which may be

come ill-conditioned and cause stability issues in the iteration. Therefore, under this scenario, the more stable coordinate descent method is applied. In the coordinate descent algorithm, we fix all coefficients except β_j , and minimize (2) for the current λ by updating β_j . The process is repeated for $j = 0, \dots, p$. After sweeping through all coordinates, we compare the new solution with $\widehat{\boldsymbol{\beta}}^{(k+1,0)}$. If they are sufficiently close, we have reached convergence; otherwise, the sweeping process is repeated until the two most recent estimators are close enough. What makes the coordinate descent algorithm particularly attractive is that there is an explicit formula for each coordinate update. The details for the coordinate descent algorithm may be found in Friedman et al. (2010) and Breheny and Huang (2011).

2.4 Stopping rules

Following the updating process, we will obtain a solution path $\widehat{\boldsymbol{\beta}}^{(k)}$ for $k = 1, \dots, K$. However, from our simulation results, we notice that in most cases, the solutions near the end of the path involve too many spurious variables. Therefore, the following two stopping rules are proposed to further speed up the path calculation process.

- (a) (Model saturation detection). The first rule is designed to terminate the path algorithm if the fitting value is too extreme. For example, in logistic regression, if the current estimated probability $\hat{p}_i = \exp(\mathbf{x}_i' \widehat{\boldsymbol{\beta}}^{(k)}) / (1 + \exp(\mathbf{x}_i' \widehat{\boldsymbol{\beta}}^{(k)}))$ satisfies

$$\max_{i=1, \dots, n} \hat{p}_i > 1 - \epsilon \quad \text{or} \quad \min_{i=1, \dots, n} \hat{p}_i < \epsilon,$$

where ϵ is a predefined positive constant, we terminate the algorithm.

- (b) (A pre-specified maximum size of the model). In some real applications, the practitioner has an upper bound on the size of the model for various reasons. For example, in the optimal portfolio allocation problem, one common restriction is the control of the transaction costs, which in turn puts a restriction on the maximum number of selected stocks. In order to avoid missing the important variables, we usually set the upper limit significantly larger than the model size we need.

Although early stopping is performed following these two rules, the optimal solution always occurs before the stopping point in our numerical experience.

2.5 Summary of the algorithm

- S1. Define the grid of penalty parameters λ as $\{\lambda_1, \dots, \lambda_K\}$, where $\lambda_1 = \lambda_{\max}$, $\lambda_K = \lambda_{\min} = \delta \lambda_{\max}$, and the remaining ones decrease on the logarithm scale. Set $k = 1$ and the initial estimate $\widehat{\boldsymbol{\beta}}^{(1)} = \mathbf{0}$.

S2. Calculate the active set by $A_k = \{j : |\partial\ell/\partial\beta_j| \geq \lambda_k\} \cup \{0\}$. Denote $\Delta_k = \lambda_{k+1} - \lambda_k$. The approximate solution is given by

$$\widehat{\beta}_{A_k}^{(k+1,0)} = \widehat{\beta}_{A_k}^{(k)} + s^{(k)} \cdot \Delta_k + \frac{1}{2} \mathbf{d}^{(k)} \cdot \Delta_k^2.$$

S3. Correct the current solution towards the KKT conditions. If (6) holds, we use the Newton-Raphson procedure; otherwise, coordinate descent method is adopted.
 S4. Check the two stopping rules, if at least one is satisfied, stop the algorithm; otherwise, set $k = k + 1$ and repeat S2–S4 until $k = K$.

2.6 Selection of tuning parameter

The performance of penalized likelihood estimators depends heavily on the choice of tuning parameters, that is λ in LASSO and (λ, γ) in MCP. This is usually accomplished through cross-validation or by using some information criterion such as AIC, BIC.

Information criteria derived using asymptotic arguments for the classical regression models are usually problematic when applied to penalized regression problems where $p \gg n$. For high-dimensional GLM, in Chen and Chen (2008), Extended BIC (EBIC) was proposed by adding an extra penalty term on top of BIC. It is defined as, for $0 \leq \gamma \leq 1$,

$$EBIC_\gamma(s) = -2 \log L_n\{\hat{\theta}(s)\} + v(s) \log n + 2\gamma \log \binom{p}{j},$$

where s is a subset of $\{1, \dots, p\}$, $\theta(s)$ is the parameter θ with those components outside s being set to 0 or some pre-specified values, $\hat{\theta}(s)$ is the maximum likelihood estimator of $\theta(s)$, and $v(s)$ is the number of components in s . In this paper, we investigate both cross-validation and EBIC.

3 APPLE with MCP penalty

Different from LASSO, MCP is a folded concave penalty which was proposed by Zhang (2010). The penalty is a quadratic spline defined on $[0, \infty)$ by

$$p_{\lambda, \gamma}(t) = \lambda \int_0^t (1 - x/(\gamma\lambda))_+ dx, \tag{7}$$

where the parameter $\gamma > 0$ measures the concavity of the penalty, and λ is the regularization parameter. The APPLE algorithm for the MCP penalized GLM is slightly different from what we proposed in Sect. 2 for LASSO. Due to the non-convexity of MCP penalty, in this section, we will only focus on the main differences from the LASSO case.

3.1 Problem setup

For MCP penalized GLM, the corresponding target function is

$$\begin{aligned} L_\lambda(\beta) &= -\ell(\mathbf{y}; \beta) + \lambda \sum_{j=1}^p \int_0^{|\beta_j|} \left(1 - \frac{x}{\lambda\gamma}\right)_+ dx \\ &= -\frac{1}{n} \sum_{i=1}^n \{y_i \theta(\beta)_i - b(\theta(\beta)_i)\} \\ &\quad + \lambda \sum_{j=1}^p \int_0^{|\beta_j|} \left(1 - \frac{x}{\lambda\gamma}\right)_+ dx. \end{aligned} \tag{8}$$

As introduced in Zhang (2010) and Zhang and Huang (2008), the sparse Riesz condition (SRC) (c_*, c^*, q) holds under some mild regularity conditions. As a result, in the low-dimensional manifolds with dimension smaller than q , the convexity of $-\ell(\mathbf{y}; \beta)$ can dominate the concavity of the penalty, which will lead to the convexity of the target function (8) even with the choice of folded concave penalty. Therefore, under the SRC, for estimator with sparsity smaller than q , the KKT conditions are still valid to obtain a global minimizer. The KKT conditions are given as follows,

$$\begin{cases} \frac{\partial \ell}{\partial \beta_0} \Big|_{\beta_0 = \hat{\beta}_0} = 0, \\ \frac{\partial \ell}{\partial \beta_j} \Big|_{\beta_j = \hat{\beta}_j} = \lambda \left(1 - \frac{|\hat{\beta}_j|}{\lambda\gamma}\right) \text{sgn}(\hat{\beta}_j) & \text{for } 0 < |\hat{\beta}_j| < \lambda\gamma, \\ \frac{\partial \ell}{\partial \beta_j} \Big|_{\beta_j = \hat{\beta}_j} = 0 & \text{for } |\hat{\beta}_j| \geq \lambda\gamma, \\ \left| \frac{\partial \ell}{\partial \beta_j} \Big|_{\beta_j = \hat{\beta}_j} \right| \leq \lambda & \text{for } \hat{\beta}_j = 0. \end{cases} \tag{9}$$

3.2 Grid of penalty parameter

The grid $\{\lambda_1, \dots, \lambda_K\}$ of penalty parameters is identical to that in the LASSO case.

3.3 Update

In the LASSO case, the effective penalty level is λ for all variables. Therefore, from the KKT condition (3), as long as a variable is activated, it stays in the active set as λ decreases. But in the MCP case, for the same λ , the effective penalty level on each variable is different depending on the magnitude of the estimate, as shown in (9). In all the existing work, such as NCVREG package, this specific property of MCP was not fully exploited and the same active set detection method was used as that for LASSO penalty (see (4) for details). Here we introduce a new active set detection method using the KKT conditions, that, to our best knowledge, has not appeared before for folded concave penalties in the literature. As will be shown later, the new detection

method is more suitable for MCP with a more efficient calculation. For a given λ_k , we define the active set A_k as

$$A_k = \{A_{k-1} \cup N_k\} \setminus D_k,$$

where

$$N_k = \{j \in \{1, \dots, p\} \setminus A_{k-1} : |\partial \ell / \partial \beta_j| \geq \lambda_k\},$$

and

$$D_k = \{j \in A_{k-1} \cap A_{k-2} : \text{sgn}(\hat{\beta}_j^{(k-1)})\text{sgn}(\hat{\beta}_j^{(k-2)}) < 0\}.$$

This means, with decreasing threshold λ_k , a particular variable becomes active when it satisfies the KKT condition (9). Then the variable will stay activated until it crosses 0 (i.e., the index lies in D_k), which means the covariates of the estimators have different signs in two consequent steps. From our experience, variables which cross 0 at some point in the path are usually noise variables. If this deleted variable satisfies the KKT condition (9) along the path later, we reactivate it. With decreasing λ , the optimization problem (8) will no longer be convex at some point. Therefore, the proposed treatment of deleting variables which cross 0 at some point will make the path more stable.

In accordance with Sect. 2.3, the step size is defined as $\Delta_k = \lambda_{k+1} - \lambda_k$. For a given λ_k and active set A_k , we update the active covariates altogether by using the quadratic approximation,

$$\hat{\beta}_{A_k}^{(k+1,0)} = \hat{\beta}_{A_k}^{(k)} + s^{(k)} \cdot \Delta_k + \frac{1}{2} d^{(k)} \cdot \Delta_k^2,$$

where $s^{(k)}$ and $d^{(k)}$ are the first and second derivatives of $\hat{\beta}^{(k)}$ with respect to λ , respectively. Since the intercept is not penalized, the first coordinates of $s^{(k)}$ and $d^{(k)}$ are both 0.

Now, we have $s^{(k)} = (0, s_{-0}^{(k)})'$, where $s_{-0}^{(k)}$ is defined as

$$s_{-0}^{(k)} = \left[\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - \Gamma \right]^{-1} (-\text{sgn}(\hat{\beta}_{A_k \setminus \{0\}}^{(k)})), \tag{10}$$

where $V^{(k)}$ is given by (11) in the Appendix, $\Gamma = \text{diag}\{1/\gamma, \dots, 1/\gamma\}$, and $\text{sgn}(\cdot)$ is the sign function of a vector.

In MCP (Zhang 2010), the tuning parameter γ is free to vary from 1+ to ∞ . For the derivative $s^{(k)}$ defined in (10), particularly in logistic regression, γ has to be large enough in order to make the matrix

$$n^{-1} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - \Gamma$$

invertible. However, if γ is too large, the MCP penalty (7) is approximately equal to $\lambda|t|$, which is the same as the LASSO penalty. In that case, it is hard to find the advantages which MCP enjoys over LASSO. In Breheny and Huang

(2011), adaptive rescaling was proposed to solve a similar issue. They replaced $p_{\lambda,\gamma}(|\beta_j|)$ with $p_{\lambda,\gamma}(|v_j \beta_j|)$, where v_j is the j -th diagonal element of the Hessian matrix. Since they used a coordinate descent algorithm, updating coordinates one-at-a-time, the rescaled updates are straightforward after this change. But in our algorithm, all the active variables are updated together. Therefore, a new adaptation method is needed.

The adaptation we use is to replace $p_{\lambda,\gamma}(|\beta_j|)$ by $p_{\lambda,\gamma u_{\min}^{(k)}}(|u_{\min}^{(k)} \beta_j|)$, where $u_{\min}^{(k)}$ is the smallest eigenvalue of the matrix $\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}}$. Then,

$$s_{-0}^{(k)} = \left(\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - u_{\min}^{(k)} \Gamma \right)^{-1} (-\text{sgn}(\hat{\beta}_{A_k \setminus \{0\}}^{(k)})).$$

Now, for any $\gamma > 1$,

$$\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - u_{\min}^{(k)} \Gamma > 0.$$

Therefore, the singularity problem in (10) is avoided for all $\gamma > 1$.

The correction method we introduced in the LASSO case is based on the fact that the problem (2) is convex. Here in MCP, although the original problem is not convex, in each step after the adaptation, our problem is still convex in a low-dimensional manifold as long as $\gamma > 1$. One important issue is that when calculating the first and second order derivatives in the Newton-Raphson correction, $u_{\min}^{(k)}$ is also a function of $\hat{\beta}^{(k)}$ (see Appendix). To avoid computing implicit derivatives, we use the popular quadratic approximation method (e.g., McCullagh and Nelder 1989) to the negative log-likelihood, which turns out to be very effective. Our new target function is

$$L(\lambda) = \frac{1}{2n} (\tilde{y} - X\beta)' V (\tilde{y} - X\beta) + \lambda \sum_{j=1}^p \int_0^{|u_{\min} \beta_j|} \left(1 - \frac{x}{\lambda \gamma u_{\min}} \right)_+ dx,$$

where $\tilde{y} = X\beta + V^{-1}(y - \pi)$. The detailed formulations are presented in the Appendix, (16)–(18) for logistic case and (19)–(21) for Poisson case.

The same sparsity criterion (6) is used, and the corresponding Newton-Raphson or coordinate descent method is applied.

3.4 Stopping rules

Stopping rules are the same as the ones discussed in Sect. 2.4. But different from the LASSO case, if a variable is activated in a certain step of the MCP procedure (9), it may turn inactive later, and even be activated again later on. So

for the same data, we can consider making the upper limit (Sect. 2.4, (b)) in MCP larger than that would be for the LASSO case.

3.5 Summary of the algorithm

The algorithm is the same as in LASSO, except S2 is replaced with S2' described by the following.

S2' Calculate the active set by $A_k = \{A_{k-1} \cup N_k\} \setminus D_k$, where

$$N_k = \{j \in \{1, \dots, p\} \setminus A_{k-1} : |\partial \ell / \partial \beta_j| \geq \lambda_k\},$$

and

$$D_k = \{j \in A_{k-1} \cap A_{k-2} : \text{sgn}(\hat{\beta}_j^{(k-1)}) \text{sgn}(\hat{\beta}_j^{(k-2)}) < 0\}.$$

3.6 Selection of tuning parameter

The selection methods are mainly the same as the ones discussed in Sect. 2.6. But as an advantage MCP is shown to possess in our numerical results, the estimators will stay in their optimal value for a certain interval of regularization parameter λ . This bears some advantages in selecting the tuning parameters.

4 Simulation results

In this section, we conduct simulation studies for comparing APPLE with the GLMNET package (Friedman et al. 2010) and NCVREG package (Breheny and Huang 2011) for LASSO and MCP penalties, respectively. Now we highlight the differences of APPLE, GLMNET, and NCVREG. First, APPLE uses vectorized update when the estimator is sparse enough, which is faster than the coordinate descent method used in both GLMNET and NCVREG packages. Second, GLMNET is only available for convex penalties, while APPLE can handle both convex and non-convex penalties. Third, to deal with non-convex penalties, APPLE uses a different γ adaptation and active set detection methods from those of NCVREG.

Logistic and Poisson regression models are two popular generalized linear models. For each model, we present results of LASSO/MCP penalized methods. For the LASSO penalty, we compare APPLE LASSO with the GLMNET package (Friedman et al. 2010). For MCP, we compare APPLE MCP with the NCVREG package (Breheny and Huang 2011). Since NCVREG only applies to Gaussian and logistic models, no comparable results are presented for MCP penalized Poisson model. For each setting, we report results for different tuning parameter selection methods, including EBIC and K -fold CV. The criteria include *false positives* (FP), *true positives* (TP), ℓ_1 loss = $\|\hat{\beta} - \beta^o\|_1$, ℓ_2 loss = $\|\hat{\beta} - \beta^o\|_2$. We also compare the computational cost with NCVREG for the MCP penalty case.

4.1 Logistic regression

Example 1 We consider a logistic regression model with different dimension, sparsity level and correlation settings. (i) Covariate dimension $p = 1000$, sample size $n = 500$ and $d = 3$, where d is the number of nonzero elements in β^o . The first 5 dimensions of β^o are (3, 1.5, 0, 0, 2), while the rest are all zeros and $\beta_0 = 0$. The vector x follows a multivariate normal distribution with zero mean and covariance between the i -th and j -th elements being $\rho^{|i-j|}$ with $\rho = 0, 0.2, 0.5$ and 0.7 in four different settings. The results are summarized in Table 1. (ii) Different dimension and different sparsity levels are considered. In particular, $(p, n, d) = (3000, 500, 3)$ and $(1000, 500, 24)$. And for both settings, we consider $\rho = 0$ and $\rho = 0.5$ with the results reported in Table 2. When $d = 24$, the first 56 dimensions are 8 repetitions of (3, 1.5, 0, 0, 2, 0, 0). In each setting, 100 repetitions are performed. Part of the setup is borrowed from Fan and Li (2001).

In Fig. 1(a), we compare the solution paths of the APPLE algorithm for MCP and LASSO. We can see that the MCP path is less smooth than the LASSO path, but has intervals at which the estimators stay constant, and yields a sparser model. With the convergence stopping rules adopted in all simulations here, the corresponding solutions of the MCP path are sparse even near the end of the path. Actually, the size of active set does not exceed the square root of the sample size, which means that the Newton-Raphson correction is used throughout the whole path. Notice that there is a jump on the LASSO path, which is caused by the change of correction method. See Feng et al. (2012) for the stability comparison of various penalty functions. In Fig. 1(b), the APPLE and GLMNET LASSO paths are illustrated. Before the change point, using the Newton-Raphson correction method, the APPLE path exhibits better estimation with a sparser model. After the change point, coordinate descent correction is employed, which makes the two paths identical. In Fig. 1(c), the APPLE and NCVREG MCP paths are compared given the same concavity parameter $\gamma = 1.3$. APPLE paths are significantly smoother than NCVREG paths. Although both paths stay at the “optimal” level, APPLE paths have a longer period, which makes the model selection task easier and leads to more stable estimation. In Fig. 1(d), APPLE MCP paths with different concavity parameters ($\gamma = 1.3, 3, 100$) are presented. This shows that as γ gets larger, the “flat” period of constant optimal magnitude gets shorter, and the APPLE MCP path eventually approaches the LASSO path when γ becomes sufficiently large. In Fig. 1(e), we show APPLE LASSO paths with different correction methods throughout the whole path. We can see that at the beginning of the path when λ is large, the differences between these two correction methods are negligible. However, as λ decreases, more variables are recruited

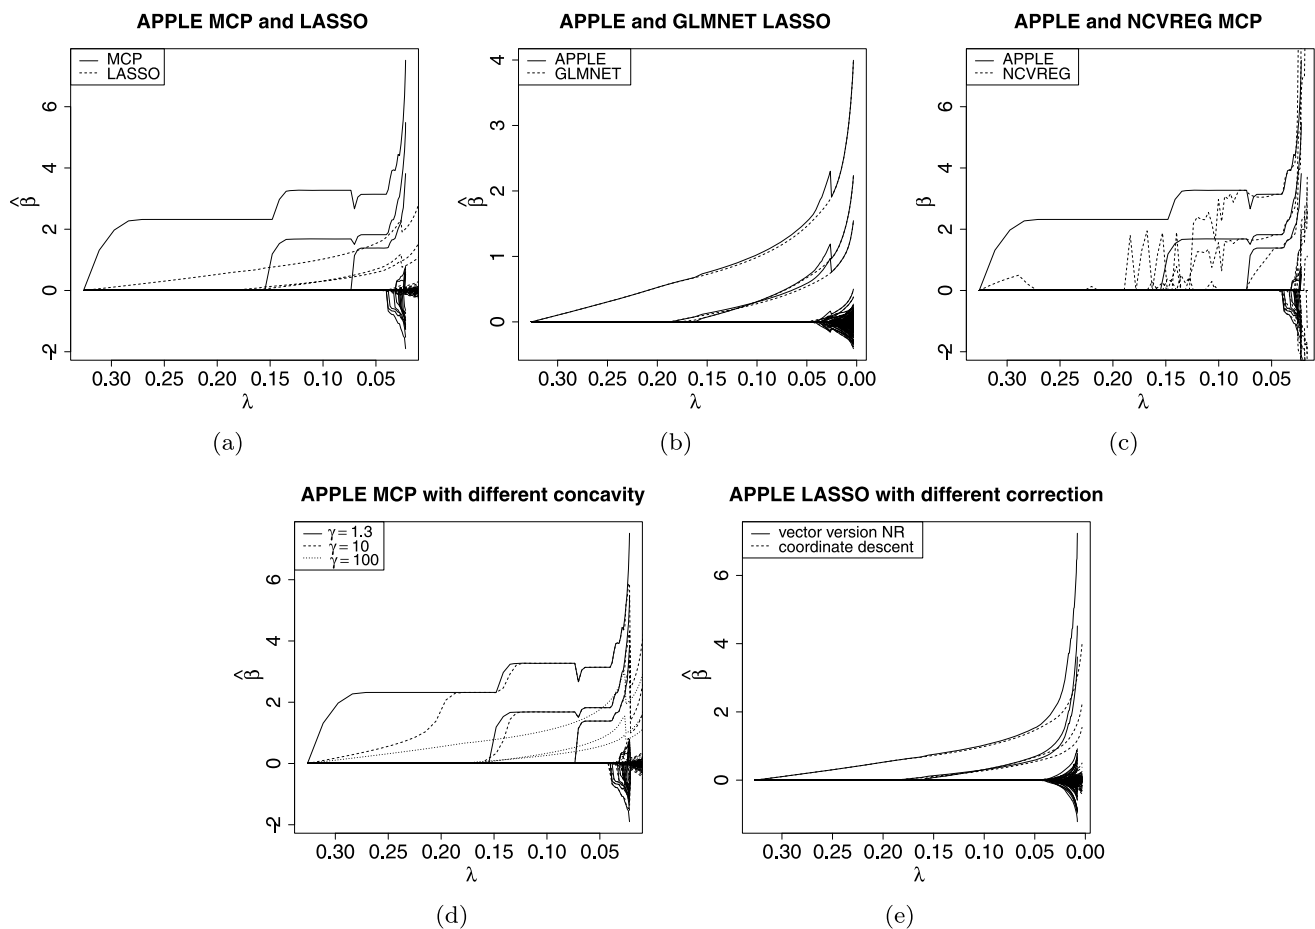


Fig. 1 Solution paths for logistic regression in Example 1, where $p = 1000$, $n = 500$, $d = 3$ and $\rho = 0$. In (a), the *solid lines* and *dotted lines* are the solution paths for APPLE MCP and APPLE LASSO, respectively. In (b), the *solid lines* and *dotted lines* are the solution paths for APPLE LASSO and GLMNET LASSO, respectively. In (c), the *solid lines* and *dotted lines* are for APPLE and NCVREG MCP, respectively. In (d), the *solid lines*, *dashed lines* and *dotted lines* are

solution paths of APPLE MCP with different γ values. In (e), the *solid lines* and *dotted lines* are the solution paths for APPLE LASSO with different correction methods. For each panel and each type of lines, the important variables are selected in the same order. As λ becomes smaller, variables with index 1, 5 and 2 are selected one by one. When λ gets smaller than 0.05, noise variables are selected

and the Newton-Raphson method becomes unstable, making the coefficient estimates “take-off” more quickly compared with the coordinate descent method. Therefore, we recommend the hybrid approach of using the Newton-Raphson in the first part of the path, and later switch to coordinate descent when the number of active variables becomes large enough.

The FP, TP, ℓ_1 loss and ℓ_2 loss results for Example 1(i) are summarized in Table 1. When $\rho = 0$, comparing the APPLE LASSO and GLMNET, we see the results from EBIC are similar for these two methods. However, for the CV, APPLE tends to provide a model with smaller FP values while keeping TP the same. When MCP is applied, similar observations can be found. Overall, comparing with the existing methods, APPLE does a better job than GLMNET and NCVREG in the LASSO and MCP cases, respectively. In

addition, for the MCP penalty, APPLE provides a smoother path than NCVREG.

The corresponding results for Example 1(ii) can be found in Table 2. When the dimension is increased to 3000 from 1000, the behaviors of APPLE applied to both LASSO and MCP cases are similar to those analyzed in Example 1(i). Recall that in APPLE, we use two different correction methods when the size of the current active set changes. To study the robustness of this dynamic correction method, we consider the case $d = 24$, which implies $d > \sqrt{n}$, i.e., the true model size exceeds the square root of sample size. In this scenario, for each setting and each method, some important variables are missing, particularly for EBIC based LASSO estimators. We conjecture the reason to be the overpenalization of EBIC. Recall that the modification in EBIC is a prior imposed on all the possible subsets of a given sparsity level. As a result, when the size of active set is large, the

Table 1 Comparisons for APPLE with GLMNET and NCVREG for LASSO and MCP penalties, respectively, in Example 1, where $p = 1000$, $n = 500$ and $d = 3$. Design matrices with different correlation, and different selection criteria are presented. The medians of false positive (FP), true positive (TP), ℓ_1 loss, and ℓ_2 loss are reported over 100 repetitions, enclosed in parentheses are the corresponding standard errors

Model	Package	Method	FP	TP	ℓ_1 loss	ℓ_2 loss
LASSO $\rho = 0$	APPLE	EBIC	0.30(0.54)	3.00(0.00)	3.64(0.28)	4.44(0.70)
		CV	23.51(7.68)	3.00(0.00)	4.93(1.27)	2.52(0.74)
	GLMNET	EBIC	0.25(0.48)	3.00(0.00)	3.83(0.23)	4.94(0.58)
		CV	48.44(26.02)	3.00(0.00)	4.93(1.27)	2.52(0.74)
MCP $\rho = 0, \gamma = 1.3$	APPLE	EBIC	0.03(0.17)	3.00(0.00)	0.68(0.38)	0.21(0.25)
		CV	0.26(0.03)	3.00(0.00)	0.86(0.96)	0.34(0.67)
	NCVREG	EBIC	0.65(0.45)	3.00(0.00)	2.50(0.31)	1.51(0.42)
		CV	1.86(3.44)	3.00(0.00)	2.05(0.48)	1.56(0.45)
MCP $\rho = 0, \gamma = 3$	APPLE	EBIC	0.03(0.17)	3.00(0.00)	0.68(0.38)	0.21(0.25)
		CV	0.26(1.03)	3.00(0.00)	0.87(0.96)	0.34(0.67)
	NCVREG	EBIC	0.03(0.17)	3.00(0.00)	0.85(0.47)	0.32(0.36)
		CV	2.65(5.37)	3.00(0.00)	1.04(0.78)	0.30(0.32)
LASSO $\rho = 0.2$	APPLE	EBIC	0.29(0.55)	3.00(0.00)	3.48(0.30)	4.03(0.70)
		CV	19.96(11.60)	3.00(0.00)	3.70(1.53)	2.31(0.61)
	GLMNET	EBIC	0.23(0.48)	3.00(0.00)	3.74(0.26)	4.71(0.67)
		CV	83.19(11.60)	3.00(0.00)	4.73(1.53)	2.31(0.61)
MCP $\rho = 0.2, \gamma = 1.3$	APPLE	EBIC	0.02(0.15)	3.00(0.00)	0.72(0.36)	0.22(0.25)
		CV	0.14(0.59)	3.00(0.00)	0.80(0.63)	0.29(0.47)
	NCVREG	EBIC	0.03(0.15)	3.00(0.00)	0.72(0.36)	0.22(0.25)
		CV	0.22(0.73)	3.00(0.00)	0.82(0.55)	0.27(0.34)
MCP $\rho = 0.2, \gamma = 3$	APPLE	EBIC	0.02(0.15)	3.00(0.00)	0.72(0.36)	0.22(0.25)
		CV	0.14(0.59)	3.00(0.00)	0.80(0.63)	0.29(0.47)
	NCVREG	EBIC	0.02(0.14)	3.00(0.00)	0.72(0.37)	0.25(0.22)
		CV	0.42(0.73)	3.00(0.00)	0.80(0.64)	0.30(0.34)
LASSO $\rho = 0.5$	APPLE	EBIC	0.10(0.30)	3.00(0.00)	3.63(0.27)	4.44(0.70)
		CV	19.20(0.45)	3.00(0.00)	3.72(0.26)	2.24(0.64)
	GLMNET	EBIC	0.10(0.30)	3.00(0.00)	3.78(0.22)	4.84(0.58)
		CV	43.88(26.02)	3.00(0.00)	4.77(1.27)	2.48(0.74)
MCP $\rho = 0.5, \gamma = 1.3$	APPLE	EBIC	0.02(0.15)	3.00(0.00)	0.66(0.32)	0.18(0.20)
		CV	0.12(0.25)	3.00(0.00)	0.79(0.49)	0.30(0.26)
	NCVREG	EBIC	0.07(0.26)	2.98(0.15)	0.73(0.44)	0.25(0.44)
		CV	3.33(2.67)	3.00(0.00)	2.84(0.42)	2.77(0.85)
MCP $\rho = 0.5, \gamma = 3$	APPLE	EBIC	0.01(0.10)	3.00(0.00)	0.72(0.34)	0.22(0.22)
		CV	0.14(0.41)	3.00(0.10)	0.87(0.60)	0.34(0.47)
	NCVREG	EBIC	0.01(0.11)	3.00(0.00)	1.62(0.54)	0.96(0.55)
		CV	4.93(5.60)	3.00(0.00)	1.55(0.97)	0.58(0.52)
LASSO $\rho = 0.7$	APPLE	EBIC	0.44(0.61)	3.00(0.00)	3.52(0.24)	4.15(0.60)
		CV	17.22(18.54)	3.00(0.00)	3.64(1.10)	3.48(1.57)
	GLMNET	EBIC	0.51(0.63)	3.00(0.00)	3.72(0.26)	5.65(0.64)
		CV	45.77(11.26)	3.00(0.00)	4.76(1.62)	3.88(1.73)
MCP $\rho = 0.7, \gamma = 1.3$	APPLE	EBIC	0.16(0.49)	2.87(0.34)	1.19(0.76)	0.72(0.95)
		CV	0.65(1.02)	2.79(0.41)	1.68(1.07)	1.18(1.20)
	NCVREG	EBIC	0.26(0.62)	2.81(0.39)	1.56(0.78)	1.09(1.01)
		CV	0.72(0.90)	2.85(0.36)	1.59(0.68)	1.07(0.91)
MCP $\rho = 0.7, \gamma = 3$	APPLE	EBIC	0.16(0.53)	2.88(0.32)	1.20(0.69)	0.72(0.97)
		CV	0.65(1.02)	2.79(0.41)	1.68(1.07)	1.18(1.20)
	NCVREG	EBIC	0.20(0.59)	2.80(0.39)	1.56(0.78)	1.09(1.01)
		CV	0.70(0.95)	2.89(0.36)	1.56(0.77)	1.17(0.95)

Table 2 Comparisons for APPLE with GLMNET and NCVREG for LASSO and MCP penalties, respectively, in Example 1, where $(p, n, d) = (3000, 500, 3)$ and $(1000, 500, 24)$. Design matrices with different correlation, and different selection criteria are presented. The medians of false positive (FP), true positive (TP), ℓ_1 loss, and ℓ_2 loss are reported over 100 repetitions, enclosed in parentheses are the corresponding standard errors

Model	Package	Method	FP	TP	ℓ_1 loss	ℓ_2 loss
LASSO $p = 3000$ $d = 3$ $\rho = 0$	APPLE	EBIC	0.06(0.09)	3.00(0.00)	3.73(0.67)	4.70(0.84)
		CV	53.93(15.90)	3.00(0.00)	5.27(1.52)	2.50(0.54)
	GLMNET	EBIC	0.07(0.05)	3.00(0.00)	4.00(1.09)	5.43(1.10)
		CV	118.03(20.83)	3.00(0.00)	5.79(1.46)	3.94(0.39f)
MCP $p = 3000$ $d = 3$ $\rho = 0$	APPLE	EBIC	0.03(0.18)	3.00(0.00)	0.77(0.41)	0.24(0.26)
		CV	0.23(0.25)	3.00(0.00)	0.78(0.39)	0.23(0.22)
	NCVREG	EBIC	0.03(0.10)	3.00(0.00)	0.78(0.40)	0.25(0.22)
		CV	0.27(0.64)	3.00(0.00)	1.02(0.88)	0.44(0.67)
LASSO $p = 3000$ $d = 3$ $\rho = 0.5$	APPLE	EBIC	0.19(0.43)	3.00(0.00)	3.61(0.26)	4.40(0.62)
		CV	32.93(42.06)	3.00(0.00)	4.28(1.89)	2.58(0.72)
	GLMNET	EBIC	0.21(0.44)	3.00(0.00)	3.83(0.26)	4.98(0.65)
		CV	190.17(13.96)	3.00(0.00)	3.83(0.25)	4.99(0.65)
MCP $p = 3000$ $d = 3$ $\rho = 0.5$	APPLE	EBIC	0.16(0.54)	3.00(0.00)	1.03(0.68)	0.50(0.64)
		CV	0.01(0.10)	3.00(0.00)	0.89(0.52)	0.38(0.50)
	NCVREG	EBIC	0.30(0.66)	3.00(0.00)	1.03(0.54)	0.53(0.49)
		CV	0.02(0.15)	3.00(0.00)	0.94(0.53)	0.41(0.51)
LASSO $p = 1000$ $d = 24$ $\rho = 0$	APPLE	EBIC	0.24(0.52)	11.17(7.67)	7.01(2.37)	10.35(3.36)
		CV	124.51(18.00)	23.90(0.29)	6.85(0.99)	8.52(2.36)
	GLMNET	EBIC	0.00(0.00)	0.12(0.54)	7.21(0.10)	11.03(0.58)
		CV	200.41(12.39)	23.91(0.28)	7.20(1.49)	6.73(1.82)
MCP $p = 1000$ $d = 24$ $\rho = 0$	APPLE	EBIC	0.12(0.38)	20.63(1.99)	4.48(2.50)	4.68(3.35)
		CV	0.09(0.29)	23.25(1.42)	4.16(2.32)	4.17(2.86)
	NCVREG	EBIC	0.14(0.36)	21.95(2.31)	4.38(3.68)	5.21(6.45)
		CV	4.59(1.70)	23.54(0.81)	17.82(11.51)	70.61(74.69)
LASSO $p = 1000$ $d = 24$ $\rho = 0.5$	APPLE	EBIC	0.90(0.98)	14.67(5.69)	6.94(1.80)	10.16(2.91)
		CV	110.84(14.05)	23.82(0.43)	6.82(1.01)	8.45(2.32)
	GLMNET	EBIC	0.00(0.00)	0.88(3.31)	7.20(0.61)	11.01(3.03)
		CV	167.15(10.24)	23.80(0.44)	7.06(2.12)	7.00(2.94)
MCP $p = 1000$ $d = 24$ $\rho = 0.5$	APPLE	EBIC	0.90(0.97)	15.78(2.12)	5.51(2.17)	6.54(3.51)
		CV	1.14(1.09)	16.93(1.88)	5.16(2.20)	5.88(3.21)
	NCVREG	EBIC	0.69(0.83)	15.61(2.34)	5.37(2.70)	6.41(4.02)
		CV	6.95(3.35)	18.29(1.99)	15.87(11.20)	58.76(53.40)

Table 3 Comparison of the computational cost for the APPLE and NCVREG packages in different simulation settings. The medians of the computation time (in seconds) are reported, enclosed in parentheses are the corresponding standard errors. CPU: Intel(R) Xeon(R) L5420 @ 2.50 GHz

	γ	$\rho = 0$		$\rho = 0.5$	
		NCVREG	APPLE	NCVREG	APPLE
$p = 2^7$	1.3	1.33(0.41)	0.13(0.07)	1.42(0.52)	0.14(0.06)
$n = 50$	3	0.22(0.28)	0.06(0.04)	0.15(0.32)	0.07(0.02)
$p = 2^8$	1.3	4.09(1.17)	0.37(0.21)	5.58(1.29)	0.42(0.11)
$n = 100$	3	0.35(0.53)	0.19(0.06)	0.32(0.40)	0.27(0.10)
$p = 2^9$	1.3	18.11(4.47)	1.15(0.18)	27.47(6.25)	1.21(0.20)
$n = 200$	3	1.58(0.77)	0.86(0.13)	1.32(0.55)	1.08(0.14)
$p = 2^{10}$	1.3	123.53(22.87)	6.29(0.67)	186.95(33.00)	6.55(0.80)
$n = 500$	3	6.55(1.49)	5.55(0.66)	9.01(4.42)	6.03(0.65)

penalty is too stringent for the variable selection purpose. Nevertheless, with the same method for choosing the tuning parameter, APPLE does a better job than GLMNET and NCVREG in the LASSO and MCP cases, respectively.

APPLE is an efficient algorithm for computing the solution path for penalized likelihood estimators, particularly for folded concave penalties. Table 3 illustrates the median time required to fit the entire path and the corresponding standard errors of the NCVREG and APPLE algorithms. Here, we use the same setting as Example 1 except for different p and n . It is clear that APPLE takes less time than NCVREG for the current example.

4.2 Poisson regression

Example 2 We consider a Poisson regression model with different dimension, sparsity level and correlation settings. (i) As Example 1(i), we set $(p, n, d) = (1000, 500, 3)$. The

Table 4 Comparison for APPLE with GLMNET for LASSO penalty and presenting APPLE MCP results, in Example 2, where $p = 1000$, $n = 500$ and $d = 3$. Design matrices with different correlation, and different selection criteria are presented. The medians of false positive (FP), true positive (TP), ℓ_1 loss, and ℓ_2 loss are reported over 100 repetitions, enclosed in parentheses are the corresponding standard errors

Model	Package	Method	FP	TP	ℓ_1 loss	ℓ_2 loss
LASSO $\rho = 0$	APPLE	EBIC	0.70(0.93)	3.00(0.00)	0.42(0.16)	0.06(0.04)
		CV	9.41(3.75)	3.00(0.00)	0.33(0.11)	0.02(0.01)
	GLMNET	EBIC	0.66(1.20)	3.00(0.00)	0.89(0.15)	0.24(0.08)
		CV	28.33(22.61)	3.00(0.00)	0.74(0.33)	0.06(0.03)
MCP $\rho = 0, \gamma = 1.3$	APPLE	EBIC	0.50(0.07)	2.99(0.08)	0.42(0.10)	0.06(0.03)
		CV	2.89(2.76)	2.99(0.26)	0.49(0.27)	0.07(0.05)
MCP $\rho = 0, \gamma = 3$	APPLE	EBIC	0.43(0.18)	2.97(0.62)	0.51(0.17)	0.08(0.05)
		CV	3.72(5.82)	2.97(0.43)	0.49(0.17)	0.05(0.03)
LASSO $\rho = 0.2$	APPLE	EBIC	0.56(0.80)	3.00(0.00)	0.37(0.13)	0.05(0.03)
		CV	9.68(5.41)	3.00(0.00)	0.32(0.12)	0.02(0.01)
	GLMNET	EBIC	1.63(1.43)	3.00(0.00)	0.63(0.13)	0.12(0.05)
		CV	55.56(25.21)	3.00(0.00)	2.61(0.67)	0.08(0.03)
MCP $\rho = 0.2, \gamma = 1.3$	APPLE	EBIC	0.81(2.39)	2.97(0.16)	0.35(1.19)	0.19(1.16)
		CV	6.60(15.19)	2.97(0.16)	1.82(12.20)	4.85(14.29)
MCP $\rho = 0.2, \gamma = 3$	APPLE	EBIC	0.79(2.41)	2.99(0.10)	0.30(1.21)	0.17(1.20)
		CV	6.98(13.90)	2.99(0.10)	1.89(13.02)	4.99(15.48)
LASSO $\rho = 0.5$	APPLE	EBIC	0.86(0.95)	3.00(0.00)	0.29(0.11)	0.03(0.03)
		CV	9.57(5.27)	3.00(0.00)	0.26(0.09)	0.01(0.01)
	GLMNET	EBIC	0.79(1.37)	3.00(0.00)	1.19(0.26)	0.51(0.23)
		CV	25.64(13.92)	3.00(0.00)	0.54(0.17)	0.04(0.02)
MCP $\rho = 0.5, \gamma = 1.3$	APPLE	EBIC	0.91(0.53)	2.91(0.22)	0.16(0.54)	0.00(0.45)
		CV	6.12(2.07)	2.97(0.28)	0.28(0.78)	0.03(0.75)
MCP $\rho = 0.5, \gamma = 3$	APPLE	EBIC	0.89(0.16)	2.98(0.12)	0.16(0.34)	0.00(0.42)
		CV	7.22(3.31)	2.99(0.19)	0.25(0.56)	0.02(0.64)
LASSO $\rho = 0.7$	APPLE	EBIC	1.01(0.93)	3.00(0.00)	0.18(0.13)	0.01(0.01)
		CV	11.08(10.03)	3.00(0.00)	0.19(0.09)	0.01(0.02)
	GLMNET	EBIC	2.58(1.49)	3.00(0.00)	0.41(0.39)	0.02(0.02)
		CV	27.39(18.42)	3.00(0.00)	0.43(0.20)	0.06(0.05)
MCP $\rho = 0.7, \gamma = 1.3$	APPLE	EBIC	1.02(0.19)	2.98(0.12)	0.18(0.10)	0.01(0.01)
		CV	16.10(6.10)	3.00(0.00)	0.15(0.11)	0.00(0.01)
MCP $\rho = 0.7, \gamma = 3$	APPLE	EBIC	0.99(0.21)	2.98(0.11)	0.18(0.21)	0.01(0.03)
		CV	15.39(4.91)	3.00(0.00)	0.18(0.10)	0.01(0.02)

first 5 dimensions of β^o are (1.2, 0.6, 0, 0, 0.8), and the rest are all zeros and $\beta_0 = 0$. The vector \mathbf{x} follows a multivariate normal distribution with zero mean and covariance between the i -th and j -th elements being $\rho^{|i-j|}$ with $\rho = 0, 0.2, 0.5$ and 0.7 in four different settings. The results are in Table 4. (ii) As Example 1(ii), different dimension and different sparsity levels are considered. We consider both $(p, n, d) = (3000, 500, 3)$ and $(1000, 500, 24)$ with $\rho = 0$ and $\rho = 0.5$. All the results are summarized in Table 5. When $d = 24$, the first 56 dimensions are 8 repetitions of (1.2, 0.6, 0, 0, 0.8). In each setting, 100 repetitions are performed. Part of the setup is borrowed from Zou and Li (2008).

In Fig. 2(a), solution paths for APPLE MCP and LASSO are presented. Similar to the logistic regression case, LASSO yields a smoother path, while MCP results in better estimation with a nearly “flat” region of optimal level. In addition, at the end of the MCP path, the solution is still sparse in terms of our “square root of sample size” criterion. In Fig. 2(b), we compare APPLE and GLMNET LASSO paths. As in the logistic regression model, there is a small jump in the APPLE LASSO path, which is caused by the change of correction method. After the correction method switches to the coordinate descent, the APPLE path coincides with the GLMNET LASSO path. In Fig. 2(c), APPLE MCP paths

Table 5 Comparison for APPLE with GLMNET for LASSO penalty and presenting APPLE MCP results, in Example 2, where $(p, n, d) = (3000, 500, 3)$ and $(1000, 500, 24)$. Design matrices with different correlation, and different selection criteria are presented. The medians of false positive (FP), true positive (TP), ℓ_1 loss, and ℓ_2 loss are reported over 100 repetitions, enclosed in parentheses are the corresponding standard errors

Model	Package	Method	FP	TP	ℓ_1 loss	ℓ_2 loss
LASSO $p = 3000$ $d = 3$ $\rho = 0$	APPLE	EBIC	0.49(0.25)	3.00(0.00)	0.33(0.10)	0.02(0.02)
		CV	5.75(0.18)	3.00(0.00)	0.30(0.04)	0.02(0.01)
	GLMNET	EBIC	1.10(0.13)	3.00(0.00)	0.56(0.23)	0.06(0.02)
		CV	98.28(12.17)	3.00(0.00)	1.36(0.32)	0.04(0.20)
MCP $p = 3000, d = 3, \rho = 0$	APPLE	EBIC	1.38(0.47)	3.00(0.00)	0.29(0.04)	0.02(0.01)
		CV	6.29(0.84)	3.00(0.00)	0.27(0.07)	0.01(0.01)
LASSO $p = 3000$ $d = 3$ $\rho = 0$	APPLE	EBIC	0.67(0.18)	3.00(0.00)	0.29(0.04)	0.03(0.01)
		CV	6.12(0.65)	3.00(0.00)	0.27(0.02)	0.01(0.02)
	GLMNET	EBIC	1.92(0.54)	3.00(0.00)	0.64(0.01)	0.12(0.01)
		CV	168.53(6.76)	3.00(0.00)	2.13(0.02)	0.12(0.02)
MCP $p = 3000, d = 3, \rho = 0.5$	APPLE	EBIC	2.39(0.55)	3.00(0.00)	0.26(0.02)	0.02(0.01)
		CV	10.67(1.27)	3.00(0.00)	0.37(0.04)	0.02(0.01)
LASSO $p = 1000$ $d = 24$ $\rho = 0$	APPLE	EBIC	0.73(0.18)	15.29(0.47)	3.92(0.53)	2.88(0.23)
		CV	57.29(12.58)	23.98(0.03)	3.21(0.42)	3.89(0.35)
	GLMNET	EBIC	0.65(0.10)	0.49(0.04)	4.28(1.20)	3.26(0.63)
		CV	142.48(22.58)	23.97(0.10)	4.29(1.02)	5.32(0.37)
MCP $p = 1000, d = 24, \rho = 0$	APPLE	EBIC	0.27(0.01)	22.19(0.27)	1.02(0.02)	0.98(0.02)
		CV	0.10(0.20)	22.30(0.39)	0.98(0.03)	0.74(0.04)
LASSO $p = 1000$ $d = 24$ $\rho = 0.5$	APPLE	EBIC	0.70(0.23)	12.14(2.49)	4.29(1.02)	5.20(0.43)
		CV	49.33(0.37)	22.96(1.18)	4.19(0.12)	4.32(0.48)
	GLMNET	EBIC	0.22(0.20)	1.02(0.04)	5.21(0.53)	4.94(0.39)
		CV	155.32(13.28)	23.01(0.20)	6.32(0.48)	5.23(0.47)
MCP $p = 1000, d = 24, \rho = 0.5$	APPLE	EBIC	0.25(0.10)	22.47(1.03)	1.20(0.07)	0.99(0.17)
		CV	0.24(0.08)	23.19(1.02)	1.04(0.03)	1.01(0.32)

with different concavity parameters are presented. The continuous gradual change with respect to γ is clear, with paths getting smoother and tending to the LASSO path as γ becomes sufficiently large. In Fig. 2(d), just as the logistic regression case, Newton-Raphson correction yields more aggressive solution. From the simulation results presented in Tables 4 and 5, APPLE LASSO performs much better than GLMNET when CV is applied in all different ρ cases. Also, it is obvious that MCP does a better job than LASSO in terms of FP and TP. Similar behaviors as in logistic regression are observed when $d > \sqrt{n}$, as what we conjectured there, we think the main reason is the over-penalization of EBIC for the larger models.

Another interesting observation is the behavior when using different values of γ for MCP in Fig. 2(c), from the simulation results presented in Table 4, neither the selection nor the estimation seems to be sensitive to the choice of γ . This shows the stability of MCP in terms of the concavity parameter γ .

4.3 Linear v.s. quadratic approximation

Different from most previous work where a linear approximation is used as the warm start in each update, APPLE

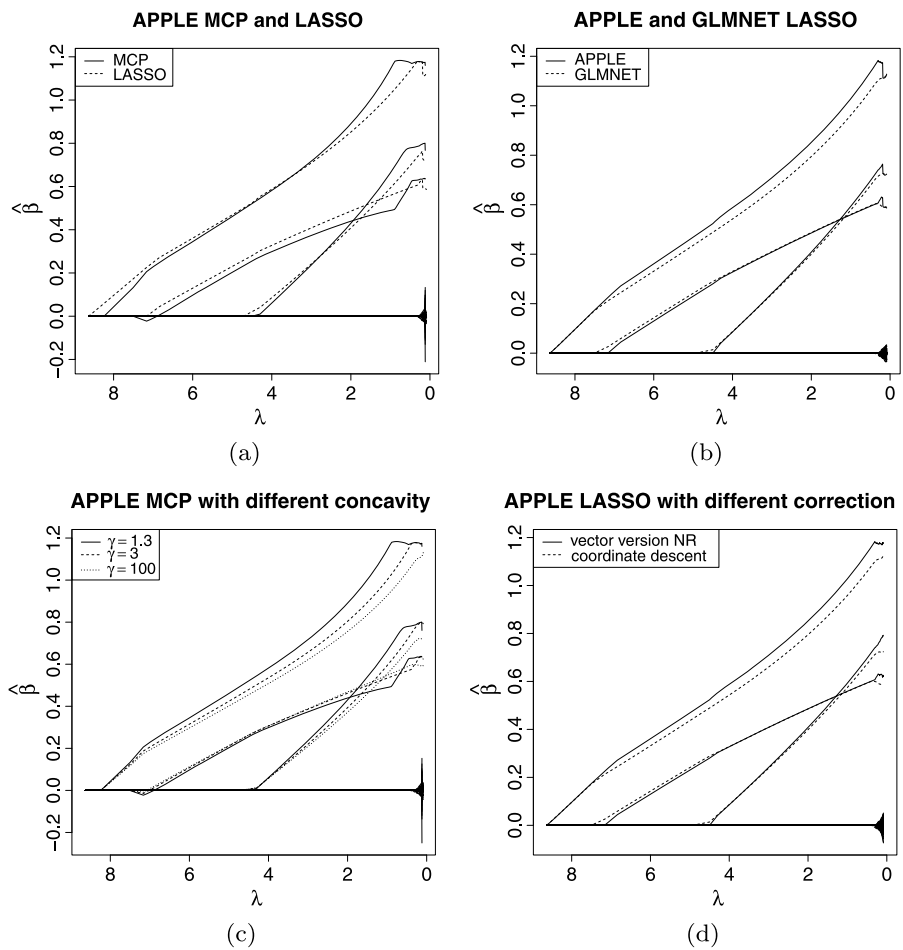
Table 6 Comparison for linear and quadratic approximation in Sect. 4.3. The median time (in seconds) are reported over 100 repetitions, enclosed in parentheses are the corresponding standard errors. CPU: Intel(R) Xeon(R) L5420 @ 2.50 GHz

	LASSO		MCP	
	$\rho = 0$	$\rho = 0.5$	$\rho = 0$	$\rho = 0.5$
Linear	3.22(0.10)	1.77(0.11)	2.26(0.20)	2.20(0.19)
Quadratic	1.50(0.12)	1.56(0.14)	1.90(0.20)	2.02(0.20)

uses a quadratic approximation. The technical details can be found in Sects. 2.3, 3.3 and Appendix. Here, we perform a simulation study to compare the solution path and the computation time under the settings $\rho = 0$ and $\rho = 0.5$ in Example 1(i).

Due to the correction step used after the covariates are updated when λ changes along the path, the linear and quadratic approximation yield essentially an identical solution path (not shown). The difference of these two methods lies in the quality of the warm starts in each step, which affects the computation cost. As expected, Table 6 shows quadratic approximation saves time over linear approximation.

Fig. 2 Solution paths for Poisson regression in Example 2, where $p = 1000$, $n = 500$, $d = 3$ and $\rho = 0$. In (a), the solid lines and dotted lines are the solution paths for APPLE MCP and APPLE LASSO, respectively. In (b), the solid lines and dotted lines are the solution paths for APPLE LASSO and GLMNET LASSO, respectively. In (c), the solid lines, dashed lines and dotted lines are solution paths of APPLE MCP with different γ values. In (d), the solid lines and dotted lines are the solution paths for APPLE LASSO with different correction methods. For each panel and each type of lines, the important variables are selected in the same order. As λ becomes smaller, variables with index 1, 5 and 2 are selected one by one. When λ gets near to zero, noise variables are selected



5 Applications

In this section, we present the analysis for two gene expression datasets with large dimension p and small sample size n .

Example 3 (i) We consider the leukemia dataset previously analyzed in Golub et al. (1999). There are $p = 7,129$ genes and $n = 72$ samples coming from two classes: 47 in class ALL (acute lymphocytic leukemia) and 25 in class AML (acute myelogenous leukemia). (ii) The Neuroblastoma data set, obtained via the MicroArray Quality Control phase-II (MAQC-II) project (Consortium 2010), consists of gene expression profiles for $p = 10,707$ genes from 251 patients of the German Neuroblastoma Trials NB90-NB2004, diagnosed between 1989 and 2004. We analyzed the gene expression data with the 3-year event-free survival (3-year EFS), which indicates whether a patient survived 3 years after the diagnosis of neuroblastoma. There are $n = 239$ subjects with the 3-year EFS information available (49 positives and 190 negatives).

Potentially, a large number of genes are affected by the two types of leukemia in (i) or negative/positive informa-

tion about 3-year EFS in (ii). In addition, the sample size n is much smaller than the dimension p for both problems. Therefore, a regularized logistic regression model is suitable. We impose LASSO and MCP penalties to these data sets, and compare the prediction accuracy yielded by the APPLE, GLMNET and NCVREG packages, respectively.

To check the stability of the results, we randomly split the data into training and testing sets 5 times for each example, and report the median prediction accuracy on the testing data and the median model size. For simplicity, EBIC was used to select the tuning parameter. For the MCP case, we fix $\gamma = 1.3$ in both examples, which turned out to have better performance than larger γ values in our simulation results. Notice that in some other high-dimensional variable selection literature, a larger γ was chosen to present the results. But when γ is too large, the MCP solution path has little difference from the LASSO path, as shown in the figures of our simulation examples.

From the results in Table 7, where test error is the number of misclassified subjects out of the size of the test dataset, we notice that for LASSO, APPLE leads to a smaller model size while having the same test error in both examples when compared with GLMNET. For MCP, in the

Table 7 Comparison for APPLE with GLMNET and NCVREG in LASSO and MCP ($\gamma = 1.3$), respectively. The medians of model size and test error (the ratio of number of wrongly classified subjects to size of test dataset) for two real data sets are reported, enclosed in the enclosed in parentheses are the corresponding standard errors

Data	Criteria	LASSO		MCP ($\gamma = 1.3$)	
		APPLE	GLMNET	APPLE	NCVREG
Leukemia	Model size	11	13	1	3
	Test error	4/36	4/36	4/36	5/36
Neuroblastoma	Model size	37	44	5	4
	Test error	22/123	22/123	22/123	23/123

leukemia example, APPLE only needs 1 variable to achieve the same test error as LASSO, and a better test error than the NCVREG. For the neuroblastoma example, MCP performs very well for both APPLE and NCVREG as compared with the LASSO.

6 Summary

In this paper, we propose a new algorithm, APPLE, for calculating the Approximate Path for Penalized Likelihood Estimators. The results from the simulation studies and real data examples provide compelling evidence that the APPLE algorithm is a worthwhile alternative to the existing methods.

APPLE takes significantly less time than NCVREG, and the same order of time as GLMNET. In each step, APPLE only updates the variables in the active set when the current model is sparse enough. When the model involves too many noise variables, APPLE switches to a coordinate descent correction.

The γ adaptation method we adopt here is different from the one originally introduced by Breheny and Huang (2011). It is due to the vector update performed in APPLE. Here, the minimum eigenvalue adaptation preserves the minimization of the maximum concavity of the MCP penalty while maintaining the stability in the Newton-Raphson update.

A public domain R language package *apple* is available from the CRAN website. <http://cran.r-project.org/web/packages/apple/>

Acknowledgements The authors thank the editor, the associate editor, and referees for their constructive comments. The authors thank Diego Franco Saldaña for proofreading.

Appendix A: Logistic regression

A.1 LASSO

In logistic regression, we assume $(x_i, y_i), i = 1, \dots, n$ are i.i.d. with $\mathbb{P}(y_i = 1|x_i) = p_i = \exp(\beta'x_i)/(1 + \exp(\beta'x_i))$.

Then the target function for the LASSO penalized logistic regression is defined as

$$L(\beta) = -\frac{1}{n} \sum_{i=1}^n \{y_i(\beta'x_i) - \log(1 + \exp(\beta'x_i))\} + \lambda \sum_{j=1}^p |\beta_j|.$$

The KKT conditions are given as follows.

$$\begin{cases} \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\beta}^{(k)' } x_i)}{1 + \exp(\hat{\beta}^{(k)' } x_i)} - y_i \right\} x_{ij} = \lambda \operatorname{sgn}(\hat{\beta}_j), & \hat{\beta}_j \neq 0; \\ \left| \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\beta}^{(k)' } x_i)}{1 + \exp(\hat{\beta}^{(k)' } x_i)} - y_i \right\} x_{ij} \right| \leq \lambda, & \hat{\beta}_j = 0; \\ \sum_{i=1}^n \frac{\exp(\hat{\beta}^{(k)' } x_i)}{1 + \exp(\hat{\beta}^{(k)' } x_i)} = \sum_{i=1}^n y_i. \end{cases}$$

We define active set A_k as

$$A_k = \left\{ j : \left| \frac{1}{n} \sum_{i=1}^n \left\{ y_i - \frac{\exp(\hat{\beta}^{(k)' } x_i)}{1 + \exp(\hat{\beta}^{(k)' } x_i)} \right\} x_{ij} \right| \geq \lambda_k \right\} \cup \{0\}.$$

To update, we define

$$\begin{aligned} \pi_i^{(k)} &= \frac{\exp(\hat{\beta}^{(k)' } x_i)}{1 + \exp(\hat{\beta}^{(k)' } x_i)}, \\ \mathbf{V}^{(k)} &= \operatorname{diag}\{\pi_1^{(k)}(1 - \pi_1^{(k)}), \dots, \pi_n^{(k)}(1 - \pi_n^{(k)})\}, \\ \mathbf{T}^{(k)} &= \operatorname{diag}\left\{ \pi_1^{(k)}(1 - \pi_1^{(k)}) \frac{1 - \exp(\hat{\beta}^{(k)' } x_1)}{1 + \exp(\hat{\beta}^{(k)' } x_1)}, \dots, \right. \\ &\quad \left. \pi_n^{(k)}(1 - \pi_n^{(k)}) \frac{1 - \exp(\hat{\beta}^{(k)' } x_n)}{1 + \exp(\hat{\beta}^{(k)' } x_n)} \right\}, \end{aligned} \tag{11}$$

then $\mathbf{s}^{(k)} = (0, \mathbf{s}_{-0}^{(k)'})', \mathbf{d}^{(k)} = (0, \mathbf{d}_{-0}^{(k)'})'$, where

$$\begin{aligned} \mathbf{s}_{-0}^{(k)} &= -[\mathbf{X}'_{A_k \setminus \{0\}} \mathbf{V}^{(k)} \mathbf{X}_{A_k \setminus \{0\}}]^{-1} \operatorname{sgn}(\hat{\beta}_{A_k \setminus \{0\}}^{(k)}), \\ \xi^{(k)} &= \operatorname{diag}(\mathbf{T}^{(k)} \mathbf{X}_{A_k \setminus \{0\}} \mathbf{s}_{-0}^{(k)}), \\ \mathbf{d}_{-0}^{(k)} &= -[\mathbf{X}'_{A_k \setminus \{0\}} \mathbf{V}^{(k)} \mathbf{X}_{A_k \setminus \{0\}}]^{-1} \mathbf{X}'_{A_k \setminus \{0\}} \xi^{(k)} \mathbf{X}_{A_k \setminus \{0\}} \mathbf{s}_{-0}^{(k)}. \end{aligned}$$

To correct,

$$\begin{aligned} \frac{\partial L^{(k)}}{\partial \beta_{A_k}} &= \frac{1}{n} \mathbf{X}'_{A_k} \left(\frac{\exp(\hat{\beta}^{(k)' } \mathbf{X})}{1 + \exp(\hat{\beta}^{(k)' } \mathbf{X})} - \mathbf{y} \right) \\ &\quad + \lambda_k \operatorname{sgn}(0, \hat{\beta}_{A_k \setminus \{0\}}^{(k)'})', \end{aligned}$$

$$\frac{\partial^2 L^{(k)}}{\partial \beta_{A_k} \partial \beta_{A_k}^T} = \frac{1}{n} \mathbf{X}'_{A_k} \mathbf{V}^{(k)} \mathbf{X}_{A_k}.$$

A.2 MCP

For MCP penalized logistic regression, we define the target function as

$$L(\boldsymbol{\beta}) = -\frac{1}{n} \sum_{i=1}^n \{y_i(\boldsymbol{\beta}'\mathbf{x}_i) - \log(1 + \exp(\boldsymbol{\beta}'\mathbf{x}_i))\} + \lambda \sum_{j=1}^p \int_0^{|\hat{\beta}_j|} \left(1 - \frac{t}{\lambda\gamma}\right)_+ dt.$$

The KKT conditions are given as follows.

$$\begin{cases} \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)} - y_i \right\} x_{ij} = \lambda \left(1 - \frac{|\hat{\beta}_j|}{\lambda\gamma}\right) \text{sgn}(\hat{\beta}_j), \\ 0 < |\hat{\beta}_j| < \lambda\gamma, \\ \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)} - y_i \right\} x_{ij} = 0, \\ |\hat{\beta}_j| \geq \lambda\gamma, \\ \left| \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)} - y_i \right\} x_{ij} \right| \leq \lambda, \\ \hat{\beta}_j = 0, \\ \sum_{i=1}^n \frac{\exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)} = \sum_{i=1}^n y_i. \end{cases}$$

For a given λ_k , define the active set A_k as

$$A_k = \{A_{k-1} \cup N_k\} \setminus D_k,$$

where

$$N_k = \left\{ j \in \{1, \dots, p\} \setminus A_{k-1} : \left| \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)} - y_i \right\} x_{ij} \right| \geq \lambda_k \right\},$$

and

$$D_k = \left\{ j \in A_{k-1} \cap A_{k-2} : \text{sgn}(\hat{\beta}_j^{(k-1)})\text{sgn}(\hat{\beta}_j^{(k-2)}) < 0 \right\}.$$

To perform adaptive rescaling on γ , define

$$\boldsymbol{\Gamma} = \text{diag}\{1/\gamma, \dots, 1/\gamma\}.$$

To update, the derivatives are defined as follows,

$$\mathbf{s}_{-0}^{(k)} = -\left(\frac{1}{n} \mathbf{X}'_{A_k \setminus \{0\}} \mathbf{V}^{(k)} \mathbf{X}_{A_k \setminus \{0\}} - u_{\min} \boldsymbol{\Gamma}\right)^{-1} \quad (12)$$

$$\times \text{sgn}(\hat{\boldsymbol{\beta}}_{A_k \setminus \{0\}}^{(k)}), \quad (13)$$

$$\mathbf{d}_{-0}^{(k)} = -\left[\frac{1}{n} \mathbf{X}'_{A_k \setminus \{0\}} \mathbf{V}^{(k)} \mathbf{X}_{A_k \setminus \{0\}} - u_{\min} \boldsymbol{\Gamma}\right]^{-1} \quad (14)$$

$$\times \mathbf{X}'_{A_k \setminus \{0\}} \boldsymbol{\xi}^{(k)} \mathbf{X}_{A_k \setminus \{0\}} \mathbf{s}_{-0}^{(k)}, \quad (15)$$

and

$$\text{sgn}(\hat{\boldsymbol{\beta}}_{-0}^{(k)}) = \begin{pmatrix} \text{sgn}(\hat{\beta}_{A_{k,1}}^{(k)}) I\{|\hat{\beta}_{A_{k,1}}^{(k)}| < \lambda^{(k)}\gamma\} \\ \vdots \\ \text{sgn}(\hat{\beta}_{A_{k,n_k}}^{(k)}) I\{|\hat{\beta}_{A_{k,n_k}}^{(k)}| < \lambda^{(k)}\gamma\} \end{pmatrix},$$

$$\text{sgn}(\hat{\beta}_j^{(k)}) = -\frac{1}{n} \sum_{i=1}^n \left\{ \frac{\exp(\hat{\boldsymbol{\beta}}^{(k)'}\mathbf{x}_i)}{1 + \exp(\hat{\boldsymbol{\beta}}^{(k)'}\mathbf{x}_i)} - y_i \right\} x_{ij}.$$

To correct we use

$$\hat{\boldsymbol{\beta}}_{A_k}^{(k,j+1)} = \hat{\boldsymbol{\beta}}_{A_k}^{(k,j)} - \left(\frac{\partial^2 L^{(k)}}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T}\right)^{-1} \left(\frac{\partial L^{(k)}}{\partial \boldsymbol{\beta}_{A_k}}\right),$$

where

$$\frac{\partial L^{(k)}}{\partial \boldsymbol{\beta}_{A_k}} = -\frac{1}{n} \mathbf{X}'_{A_k} \mathbf{V}^{(k)} (\tilde{\mathbf{y}} - \mathbf{X}_{A_k} \hat{\boldsymbol{\beta}}_{A_k}^{(k)}) + \lambda_k \text{sgn}(0, \hat{\boldsymbol{\beta}}_{A_k \setminus \{0\}}^{(k)})' \left(1 - \frac{|\hat{\boldsymbol{\beta}}_{A_k \setminus \{0\}}^{(k)}|}{\lambda_k \gamma}\right)_+, \quad (16)$$

$$\frac{\partial^2 L^{(k)}}{\partial \boldsymbol{\beta}_{A_k} \partial \boldsymbol{\beta}_{A_k}^T} = \frac{1}{n} \mathbf{X}'_{A_k} \mathbf{V}^{(k)} \mathbf{X}_{A_k} - u_{\min} \boldsymbol{\Gamma}, \quad (17)$$

and

$$\tilde{\mathbf{y}} = (\mathbf{V}^{(k)})^{-1} \left\{ \mathbf{y} - \frac{\exp(\hat{\boldsymbol{\beta}}^{(k)'} \mathbf{X})}{1 + \exp(\hat{\boldsymbol{\beta}}^{(k)'} \mathbf{X})} \right\}. \quad (18)$$

Appendix B: Poisson regression

B.3 LASSO

In Poisson regression, we assume (\mathbf{x}_i, y_i) , $i = 1, \dots, n$ are iid with $\mathbb{P}(Y = y_i) = e^{-\lambda_i} \lambda_i^{y_i} / (y_i)!$, where $\log \lambda_i = \boldsymbol{\beta}'\mathbf{x}_i$. Then criterion for the LASSO penalized Poisson regression is defined as

$$L(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n \{ \exp(\boldsymbol{\beta}'\mathbf{x}_i) - y_i(\boldsymbol{\beta}'\mathbf{x}_i) \} + \lambda \sum_{j=1}^p |\beta_j|.$$

The KKT conditions are given as follows.

$$\begin{cases} \frac{1}{n} \sum_{i=1}^n \{ \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i) - y_i \} x_{ij} = \lambda \text{sgn}(\hat{\beta}_j), & \hat{\beta}_j \neq 0, \\ \left| \frac{1}{n} \sum_{i=1}^n \{ \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i) - y_i \} x_{ij} \right| \leq \lambda, & \hat{\beta}_j = 0, \\ \hat{\beta}_0 = \log \frac{\sum_{i=1}^n y_i}{\sum_{i=1}^n \exp(\hat{\boldsymbol{\beta}}'\mathbf{x}_i)}. \end{cases}$$

For a given λ_k , we define the active set A_k as follows.

$$A_k = \left\{ j : \left| \frac{1}{n} \sum_{i=1}^n \{ y_i - \exp(\hat{\boldsymbol{\beta}}^{(k)'}\mathbf{x}_i) x_{ij} \} \right| \geq \lambda_k \right\} \cup \{0\}.$$

To update, we define

$$V^{(k)} = \text{diag}\{\exp(\widehat{\beta}^{(k)'} \mathbf{x}_1), \dots, \exp(\widehat{\beta}^{(k)'} \mathbf{x}_n)\},$$

then

$$s_{-0}^{(k)} = -[X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}}]^{-1} \text{sgn}(\widehat{\beta}_{A_k \setminus \{0\}}^{(k)}),$$

$$\xi^{(k)} = \text{diag}(V^{(k)} X_{A_k \setminus \{0\}} s_{-0}^{(k)}),$$

and

$$d_{-0}^{(k)} = -[X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}}]^{-1} \times X'_{A_k \setminus \{0\}} \xi^{(k)} X_{A_k \setminus \{0\}} s_{-0}^{(k)}.$$

To correct,

$$\frac{\partial L^{(k)}}{\partial \beta_{A_k}} = \frac{1}{n} X'_{A_k} (\exp(\widehat{\beta}^{(k)'} X) - y) + \lambda_k \text{sgn}(0, \widehat{\beta}_{A_k \setminus \{0\}}^{(k)'}),$$

$$\frac{\partial^2 L^{(k)}}{\partial \beta_{A_k} \partial \beta_{A_k}^T} = \frac{1}{n} X'_{A_k} V^{(k)} X_{A_k}.$$

B.4 MCP

For MCP penalized Poisson regression, we define the target function as

$$L(\beta) = \frac{1}{n} \sum_{i=1}^n \{\exp(\beta' \mathbf{x}_i) - y_i(\beta' \mathbf{x}_i)\} + \lambda \sum_{j=1}^p \int_0^{|\beta_j|} \left(1 - \frac{t}{\lambda \gamma}\right)_+ dt.$$

The KKT conditions are,

$$\begin{cases} \frac{1}{n} \sum_{i=1}^n \{\exp(\widehat{\beta}' \mathbf{x}_i) - y_i\} \mathbf{x}_{ij} = \lambda \left(1 - \frac{|\widehat{\beta}_j|}{\lambda \gamma}\right) \text{sgn}(\widehat{\beta}_j), \\ 0 < |\widehat{\beta}_j| < \lambda \gamma, \\ \frac{1}{n} \sum_{i=1}^n \{\exp(\widehat{\beta}' \mathbf{x}_i) - y_i\} \mathbf{x}_{ij} = 0, \\ |\widehat{\beta}_j| \geq \lambda \gamma, \\ \left| \frac{1}{n} \sum_{i=1}^n \{\exp(\widehat{\beta}' \mathbf{x}_i) - y_i\} \mathbf{x}_{ij} \right| \leq \lambda, \\ \widehat{\beta}_j = 0, \\ \widehat{\beta}_0 = \log \frac{\sum_{i=1}^n y_i}{\sum_{i=1}^n \exp(\widehat{\beta}' \mathbf{x}_i)}. \end{cases}$$

For a given λ_k , the active set is defined as

$$A_k = \{A_{k-1} \cup N_k\} \setminus D_k,$$

where

$$N_k = \left\{ j \in \{1, \dots, p\} \setminus A_{k-1} : \left| \frac{1}{n} \sum_{i=1}^n \{\exp(\widehat{\beta}' \mathbf{x}_i) - y_i\} \mathbf{x}_{ij} \right| \geq \lambda_k \right\},$$

and

$$D_k = \{j \in A_{k-1} \cap A_{k-2} : \text{sgn}(\widehat{\beta}_j^{(k-1)}) \text{sgn}(\widehat{\beta}_j^{(k-2)}) < 0\}.$$

To update, the derivatives are defined as follows,

$$s_{-0}^{(k)} = \left(\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - u_{\min} \Gamma \right)^{-1} \times \text{sgn}(\widehat{\beta}_{A_k \setminus \{0\}}^{(k)}),$$

$$d_{-0}^{(k)} = - \left[\frac{1}{n} X'_{A_k \setminus \{0\}} V^{(k)} X_{A_k \setminus \{0\}} - u_{\min} \Gamma \right]^{-1} \times X'_{A_k \setminus \{0\}} \xi^{(k)} X_{A_k \setminus \{0\}} s_{-0}^{(k)},$$

and

$$\text{sgn}(\widehat{\beta}_{-0}^{(k)}) = \begin{pmatrix} \text{sgn}(\widehat{\beta}_{A_{k,1}}^{(k)}) I_{\{|\widehat{\beta}_{A_{k,1}}^{(k)}| < \lambda^{(k)} \gamma\}} \\ \vdots \\ \text{sgn}(\widehat{\beta}_{A_{k,n_k}}^{(k)}) I_{\{|\widehat{\beta}_{A_{k,n_k}}^{(k)}| < \lambda^{(k)} \gamma\}} \end{pmatrix},$$

$$\text{sgn}(\widehat{\beta}_j^{(k)}) = -\frac{1}{n} \sum_{i=1}^n \{\exp(\widehat{\beta}^{(k)'} \mathbf{x}_i) - y_i\} \mathbf{x}_{ij}.$$

To correct we use

$$\widehat{\beta}_{A_k}^{(k,j+1)} = \widehat{\beta}_{A_k}^{(k,j)} - \left(\frac{\partial^2 L^{(k)}}{\partial \beta_{A_k} \partial \beta_{A_k}^T} \right)^{-1} \left(\frac{\partial L^{(k)}}{\partial \beta_{A_k}} \right),$$

where

$$\frac{\partial L^{(k)}}{\partial \beta_{A_k}} = -\frac{1}{n} X'_{A_k} V^{(k)} (\tilde{\mathbf{y}} - X_{A_k} \widehat{\beta}_{A_k}^{(k)}) + \lambda_k \text{sgn}(0, \widehat{\beta}_{A_k \setminus \{0\}}^{(k)'} \left(1 - \frac{|\widehat{\beta}_{A_k \setminus \{0\}}^{(k)}|}{\lambda_k \gamma}\right)_+, \tag{19}$$

$$\frac{\partial^2 L^{(k)}}{\partial \beta_{A_k} \partial \beta_{A_k}^T} = \frac{1}{n} X'_{A_k} V^{(k)} X_{A_k} - u_{\min} \Gamma, \tag{20}$$

and

$$\tilde{\mathbf{y}} = (V^{(k)})^{-1} \{y - \exp(\widehat{\beta}^{(k)'} X)\}. \tag{21}$$

References

Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: Proc. 2nd International Symposium on Information Theory (1973)

- Barron, A., Birge, L., Massart, P.: Risk bounds for model selection via penalization. *Probab. Theory Relat. Fields* **113**, 301–413 (1999)
- Breheny, P., Huang, J.: Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.* **5**(1), 232–253 (2011)
- Breiman, A.G., Gao, H.Y.: Understanding waveshrink: variance and bias estimation. *Biometrika* **83**, 727–745 (1996)
- Chen, J., Chen, Z.: Extended Bayesian information criteria for model selection with large model spaces. *Biometrika* **95**(3), 759–771 (2008)
- Chen, S., Donoho, D.L.: On basis pursuit. Tech. Rep., Dept. Statistics, Stanford Univ. (1994)
- Consortium, M.: The microarray quality control (maqc)-ii study of common practices for the development and validation of microarray-based predictive models. *Nat. Biotechnol.* **28**, 827–841 (2010)
- Daubechies, I., Defrise, M., Mol, C.D.: An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Commun. Pure Appl. Math.* **57**, 1413–1457 (2004)
- Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression (with discussion). *Ann. Stat.* **32**, 407–499 (2004)
- Fan, J., Li, R.: Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Am. Stat. Assoc.* **96**, 1348–1360 (2001)
- Fan, J., Lv, J.: A selective overview of variable selection in high dimensional feature space. *Stat. Sin.* **20**, 101–148 (2010)
- Fan, J., Xue, L., Zou, H.: Strong oracle optimality of folded concave penalized estimation (2012). [arXiv:1210.5992](https://arxiv.org/abs/1210.5992)
- Feng, Y., Li, T., Ying, Z.: Likelihood adaptively modified penalties. Manuscript (2012)
- Friedman, J., Hastie, T., Tibshirani, R.: Pathwise coordinate optimization. *Ann. Stat.* **1**, 302–332 (2007)
- Friedman, J., et al.: Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**, 1–22 (2010)
- Fu, W.: Penalized regressions: the bridge versus the lasso. *J. Comput. Graph. Stat.* **7**(3), 397–416 (1998)
- van de Geer, S.: High-dimensional generalized linear models and the lasso. *Ann. Stat.* **36**, 614–645 (2008)
- Genkin, A., Lewis, D.D., Madigan, D.: Large-scale Bayesian logistic regression for text categorization. *Technometrics* **49**(3), 291–304 (2007)
- Golub, T.R., Slonim, D.K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J.P., Coller, H.: Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* **286**(5439), 531–537 (1999)
- Hastie, T., Tibshirani, R., Rosset, S., Zhu, J.: The entire regularization path for the support vector machine. *J. Mach. Learn. Res.* **5**, 1391–1415 (2004)
- Kim, S.J., Koh, K., Lustig, M., Boyd, S., Gorinevsky, D.: An interior-point method for large-scale l_1 regularized least squares. *J. Mach. Learn. Res.* **8**, 1519–1555 (2007)
- Krishnapuram, B., Carin, L., Figueiredo, M., Hartemink, A.: Sparse multinomial logistic regression: fast algorithms and generalization bounds. *IEEE Trans. Pattern Anal. Mach. Intell.* **27**, 957–968 (2005)
- Lee, S.I., Lee, H., Abbeel, P., Ng, A.Y.: Efficient l_1 regularized logistic regression. In: *Proceedings of the National Conference on Artificial Intelligence*, vol. 21, pp. 401–408 (2006)
- Mallows, C.L.: Some comments on c_p . *Technometrics* **12**, 661–675 (1973)
- McCullagh, P., Nelder, J.A.: *Generalized Linear Model*, 2nd edn. Chapman and Hall, New York (1989)
- Meier, L., Geer, S.V.D., Bühlmann, P.: The group lasso for logistic regression. *J. R. Stat. Soc. B* **70**(1), 53–71 (2008)
- Osborne, M., Presnell, B., Turlach, B.: A new approach to variable selection in least squares problems. *IMA J. Numer. Anal.* **20**(3), 389–404 (2000)
- Park, M.Y., Hastie, T.: An l_1 regularization-path algorithm for generalized linear models. *J. R. Stat. Soc. B* **69**, 659–677 (2007)
- Rosset, S., Zhu, J.: Piecewise linear regularized solution paths. *Ann. Stat.* **35**(3), 1012–1030 (2007)
- Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978)
- Shevade, K., Keerthi, S.: A simple and efficient algorithm for gene selection using sparse logistic regression. *Bioinformatics* **19**, 2246–2253 (2003)
- Tibshirani, R.: Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **9**, 1135–1151 (1996)
- Wei, F., Zhu, H.: Group coordinate descent algorithms for nonconvex penalized regression. *Comput. Stat. Data Anal.* **56**(2), 316–326 (2012)
- Wu, T.T., Lange, K.: Coordinate descent method for lasso penalized regression. *Ann. Appl. Stat.* **2**, 224–244 (2008)
- Wu, Y.: An ordinary differential equation-based solution path algorithm. *J. Nonparametr. Stat.* **23**(1), 185–199 (2011)
- Yuan, M., Zou, H.: Efficient global approximation of generalized nonlinear ℓ_1 -regularized solution paths and its applications. *J. Am. Stat. Assoc.* **104**(488), 1562–1574 (2009)
- Zhang, C.H.: Nearly unbiased variable selection under minimax concave penalty. *Ann. Stat.* **38**(2), 894–942 (2010)
- Zhang, C.H., Huang, J.: The sparsity and bias of the lasso selection in high-dimensional regression. *Ann. Stat.* **36**, 1567–1594 (2008)
- Zhu, J., Hastie, T.: Classification of gene microarrays by penalized logistic regression. *Biostatistics* **5**(3), 427–443 (2004)
- Zou, H., Li, R.: One-step sparse estimates in nonconcave penalized likelihood models. *Ann. Stat.* **38**, 1509–1533 (2008)