


This article was downloaded by: [Seoul National University]

On: 28 April 2015, At: 01:12

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 Computational and Graphical Statistics

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/ucgs20>

High-Dimensional Fused Lasso Regression Using Majorization-Minimization and Parallel Processing

Donghyeon Yu, Joong-Ho Won, Taehoon Lee, Johan Lim & Sungroh Yoon

Accepted author version posted online: 24 Jan 2014. Published online: 31 Mar 2015.



CrossMark

[Click for updates](#)

To cite this article: Donghyeon Yu, Joong-Ho Won, Taehoon Lee, Johan Lim & Sungroh Yoon (2015) High-Dimensional Fused Lasso Regression Using Majorization-Minimization and Parallel Processing, Journal of Computational and Graphical Statistics, 24:1, 121-153, DOI: [10.1080/10618600.2013.878662](https://doi.org/10.1080/10618600.2013.878662)

To link to this article: <http://dx.doi.org/10.1080/10618600.2013.878662>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

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. Terms &

Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

High-Dimensional Fused Lasso Regression Using Majorization–Minimization and Parallel Processing

Donghyeon YU, Joong-Ho WON, Taehoon LEE, Johan LIM,
and Sungroh YOON

In this article, we propose a majorization–minimization (MM) algorithm for high-dimensional fused lasso regression (FLR) suitable for parallelization using graphics processing units (GPUs). The MM algorithm is stable and flexible as it can solve the FLR problems with various types of design matrices and penalty structures within a few tens of iterations. We also show that the convergence of the proposed algorithm is guaranteed. We conduct numerical studies to compare our algorithm with other existing algorithms, demonstrating that the proposed MM algorithm is competitive in many settings including the two-dimensional FLR with arbitrary design matrices. The merit of GPU parallelization is also exhibited. Supplementary materials are available online.

Key Words: Graphics processing unit; Majorization-minimization algorithm; Parallel computation.

1. INTRODUCTION

Regression methods using ℓ_1 -regularization are commonly employed to analyze high-dimensional data, as they tend to yield sparse estimates of regression coefficients. In this article, we consider the fused lasso regression (FLR), an important special case of ℓ_1 -regularized regression. The FLR minimizes

$$f(\beta) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \sum_{j=1}^p |\beta_j| + \lambda_2 \sum_{(j,k) \in E} |\beta_j - \beta_k|, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^n$ is the response vector, and $\mathbf{X} \in \mathbb{R}^{n \times p}$ is the design matrix; n is the sample size, and p is the number of variables; λ_1 and λ_2 are nonnegative regularization parameters, and E is the unordered index set of the adjacent pairs of variables specified in the model.

Donghyeon Yu, Department of Statistics, Keimyung University, Daegu, Korea (E-mail: dyu3@kmu.ac.kr).
Joong-Ho Won and Johan Lim, Department of Statistics, Seoul National University, Seoul, Korea (E-mail: wonj@stats.snu.ac.kr; johanlim@snu.ac.kr). Taehoon Lee and Sungroh Yoon, Department of Electrical and Computer Engineering, Seoul National University, Seoul, Korea (E-mail: taehoonlee@snu.ac.kr; sryoon@snu.ac.kr).

© 2015 American Statistical Association, Institute of Mathematical Statistics,
and Interface Foundation of North America

Journal of Computational and Graphical Statistics, Volume 24, Number 1, Pages 121–153

DOI: 10.1080/10618600.2013.878662

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jcgs.

Thus, the FLR not only promotes sparsity among the coefficients, but also encourages adjacent coefficients to have the same values (“fusion”), where adjacency is determined by the application.

The above specification of the FLR (1) generalizes the original proposal by Tibshirani et al. (2005), which considers adjacency on a one-dimensional chain graph, that is, $E = \{(j - 1, j) | j = 2, \dots, p\}$. We refer to this particular case as the *standard FLR*. Problem (1) also covers a wider class of the index set E , as considered by Chen et al. (2012) (see Section 2). Thus, we call the penalty terms associated with λ_2 the *generalized fusion penalty*. For instance, the *two-dimensional FLR* uses a two-dimensional lattice $E = \{(i, j - 1), (i, j) | 1 \leq i \leq q, j = 2, 3, \dots, q\} \cup \{(i - 1, j), (i, j) | i = 2, 3, \dots, q, 1 \leq j \leq q\}$ for coefficients $(\beta_{ij})_{1 \leq i, j \leq q}$; the *clustered lasso* (She 2010) and the *pairwise fused lasso* (Petry et al. 2011) use all pairs of the variables as the index set. Closely related but not exactly specified by (1) includes the *generalized lasso* (Tibshirani and Taylor 2011). Regardless of the penalty structure, or the choice of E , when $\mathbf{X} = \mathbf{I}$ the FLR is called the *fused lasso signal approximator* (FLSA).

The major challenge in the FLR as compared to the classical lasso regression (Tibshirani 1996) that corresponds to setting $\lambda_2 = 0$ in (1) is that the generalized fusion penalty terms are nonseparable as well as nonsmooth. The FLR is basically a quadratic programming (QP) and can be solved in principle using general purpose QP solvers (Gill et al. 1997; Grant et al. 2008). However, the standard QP formulation introduces a large number of additional variables and constraints, and is not adequate for general purpose QP solvers if the dimension p increases.

To resolve this difficulty, many special algorithms for solving the FLR problem have been proposed. Some are based on LARS (Efron et al. 2004)-flavored solution path methods (Friedman et al. 2007; Hoefling 2010; Tibshirani and Taylor 2011), and others on a fixed set of the regularization parameters (λ_1, λ_2) (Liu et al. 2010; Lin et al. 2011; Ye and Xie 2011; Chen et al. 2012). Although these algorithms are adequate for high-dimensional settings, restrictions may apply: some are only applicable for special design matrices (Friedman et al. 2007; Hoefling 2010; Tibshirani and Taylor 2011) or the standard FLR (Liu et al. 2010). In addition, computational issues still pertain for very high dimensional settings ($p \gg 1000$).

In this article, we apply the majorization–minimization (MM) algorithm (Lange et al. 2000) to solve the FLR problem. The MM algorithm iteratively finds and minimizes a surrogate function called the *majorizing function* that is generally convex and differentiable. It is well known that the efficiency of the MM algorithm depends largely on the choice of the majorizing function. Hence, we begin the main body of the article by proposing a majorizing function suitable for the FLR problem.

Motivated by Hunter and Li (2005) and their work on the classical lasso, we first introduce a perturbed version of the objective function (1) and propose a majorizing function for this perturbed objective. The MM algorithm we propose is based on these perturbed objective function and its majorizing function. The MM algorithm has several advantageous points over other existing algorithms. First, it has flexibility in the choice of both design matrix and penalty structure. We show that the proposed MM algorithm converges to the optimal objective (1) regardless of the rank of the design matrix. Furthermore, it can be applied to general penalty structures, while having comparable performance with state-of-the-art

FLR solvers, some of which require a special penalty structure. Second, as numerically demonstrated in Section 5, our MM algorithm is stable in the number of iterations to converge regardless of the choice of design matrix and the penalty structure. Finally, an additional benefit of our MM formulation is that it opens up the possibility of solving the FLR problem in a massively parallel fashion using graphics processing units (GPUs).

This article is organized as follows. In Section 2, we review existing algorithms for the FLR. In Section 3, we propose an MM algorithm for the FLR with the generalized fusion penalty and prove the convergence of the algorithm. We also introduce the preconditioned conjugate gradient (PCG) method to accelerate the proposed MM algorithm. In Section 4, we explain how to parallelize the MM algorithm using GPUs. In Section 5, we conduct numerical studies to compare the proposed MM algorithm with the other existing algorithms. In Section 6, we conclude the article with a brief summary.

2. REVIEW OF THE EXISTING ALGORITHMS

Existing algorithms for solving the FLR problem can be classified into two categories. One is based on solution path methods, whose goal is to find the whole solutions for all values of the regularization parameters. The other is based on first-order methods, which attempt to solve a large-scale problem given a fixed set of regularization parameters using first-order approximation.

2.1 SOLUTION PATH METHODS

Solution path methods aim to provide all the solutions of interest with a small computational cost after solving the initial problems to find all change points while varying the regularization parameters λ_1 and λ_2 . A change point refers to the value of λ_1 or λ_2 at which some coefficients fuse, or coefficients that used to be fused split.

Pathwise optimization algorithm. Friedman et al. (2007) proposed the pathwise optimization algorithm for the “standard” FLSA,

$$\min_{\beta} f(\beta) \equiv \frac{1}{2} \|\mathbf{y} - \beta\|_2^2 + \lambda_1 \sum_{j=1}^n |\beta_j| + \lambda_2 \sum_{j=2}^n |\beta_j - \beta_{j-1}|.$$

In this case, it can be shown that $\widehat{\beta}_j(\lambda_1, \lambda_2)$, the optimal solution for the standard FLSA for regularization parameters (λ_1, λ_2) , can be obtained from $\widehat{\beta}_j(0, \lambda_2)$ by soft-thresholding:

$$\widehat{\beta}_j(\lambda_1, \lambda_2) = \text{sign}(\widehat{\beta}_j(0, \lambda_2)) \cdot \max(|\widehat{\beta}_j(0, \lambda_2)| - \lambda_1, 0) \quad \text{for } j = 1, 2, \dots, p. \quad (2)$$

Hence, we can set $\lambda_1 = 0$ without loss of generality and focus on the regularization path by varying only λ_2 .

Because of the nonseparable nature of the generalized fusion penalty terms, typical coordinate descent (CD) algorithms may fail to reach the optimal solution, despite the convexity of the objective function (Tseng 2001). This is in contrast to the classical lasso, where the penalty terms (associated with λ_1 in the case of the FLR) are separable. Hence, the pathwise optimization algorithm uses a modified version of the CD algorithm so that

two coefficients move together if the coordinatewise move fails to reduce the objective function.

An iteration of this modified CD algorithm comprises of three cycles: descent, fusion, and smoothing. The descent cycle is to update the i th coordinate $\hat{\beta}_i$ of the current solution $\hat{\beta}$ in a usual coordinate descent fashion: minimize $f(\beta)$ with respect to β_i holding all the other coefficients fixed. Since the subdifferential of f with respect to β_i

$$\frac{\partial f(\beta)}{\partial \beta_i} = -(y_i - \beta_i) - \lambda_2 \text{sgn}(\tilde{\beta}_{i+1} - \beta_i) + \lambda_2 \text{sgn}(\beta_i - \tilde{\beta}_{i-1}),$$

where $\tilde{\beta}$ is the solution from the previous iteration; $\text{sgn}(x)$ is the sign of x if $x \neq 0$, and any number in $[-1, 1]$ if $x = 0$, is piecewise linear with breaks at $\tilde{\beta}_{i-1}$ and $\tilde{\beta}_{i+1}$, the desired coordinatewise minimization is simple. As discussed above, the descent cycle may not reduce the objective function $f(\beta)$, hence the fusion cycle minimizes $f(\beta)$ for two adjacent variables β_{i-1} and β_i , under the constraint $\beta_{i-1} = \beta_i$, that is, the two variables are fused. This cycle in effect reduces the number of the variables by one. Still, fusion of neighboring two variables is not enough to reduce the objective function. To overcome this drawback, the smoothing cycle increases λ_2 by a small amount and update the solution $\hat{\beta}(0, \lambda_2)$ by repeatedly applying the descent and the fusion cycles, and keeping track of the variable reduction. This strategy guarantees convergence to the exact solution if it starts from $\lambda_2 = 0$ where $\hat{\beta} = \mathbf{y}$, because in the standard FLSA, a fusion occurs at most two neighboring variables if the increment of λ_2 is sufficiently small, and the fused variables do not split for the larger values of λ_2 .

The resulting algorithm produces a forward stagewise-type path over a fine grid of λ_2 , which can be obtained for all values of λ_1 due to (2). This algorithm has also been applied to the two-dimensional FLSA in Friedman et al. (2007). However, its convergence is not guaranteed since the conditions described at the end of the last paragraph do not hold for this case. For the same reason, its extension to the FLR with general design matrices is limited.

Path algorithm for the FLSA. Hoefling (2010) proposed a path algorithm for the FLSA with the generalized fusion penalty,

$$\min_{\beta} f(\beta) \equiv \frac{1}{2} \|\mathbf{y} - \beta\|_2^2 + \lambda_1 \sum_{i=1}^n |\beta_i| + \lambda_2 \sum_{(j,k) \in E} |\beta_j - \beta_k|, \quad (3)$$

which gives an exact regularization path, in contrast to the approximate one for the standard FLSA by Friedman et al. (2007). The main idea of the path algorithm arises from the observation that the sets of fused coefficients do not change with λ_2 except for finitely many points. (Again (2) holds and we can assume $\lambda_1 = 0$ without loss of generality.) For each interval of λ_2 constructed by these finite change points, the objective function $f(\beta)$ in (3) can be written as

$$f^*(\beta) = \frac{1}{2} \sum_{i=1}^{n_F} \left(\sum_{j \in F_i(\lambda_2)} (y_j - \beta_{F_i}(\lambda_2))^2 \right) + \lambda_2 \sum_{1 \leq i < j \leq n_F} n_{ij} |\beta_{F_i}(\lambda_2) - \beta_{F_j}(\lambda_2)|, \quad (4)$$

where $F_1(\lambda_2), F_2(\lambda_2), \dots, F_{n_F}(\lambda_2)$ are the sets of fused coefficients, $n_F = n_F(\lambda_2)$ is the number of the sets, and $n_{ij} = |\{(k, l) \in E | k \in F_i(\lambda_2), l \in F_j(\lambda_2)\}|$ is the number of edges connecting the sets $F_i(\lambda_2)$ and $F_j(\lambda_2)$. The minimizer of (4), hence of (3), varies linearly with λ_2 within this interval:

$$\frac{\partial \widehat{\beta}_{F_i}(\lambda_2)}{\partial \lambda_2} = - \frac{\sum_{j \neq i} n_{ij} \operatorname{sgn}(\widehat{\beta}_{F_i}(\lambda_2) - \widehat{\beta}_{F_j}(\lambda_2))}{|F_i(\lambda_2)|}, \tag{5}$$

that is the solution path is piecewise linear. Since at $\lambda_2 = 0$, $\widehat{\beta}_{F_i}(0) = y_i$ and $F_i = \{i\}$ for $i = 1, 2, \dots, n$, by keeping track of the change points at which the solution path changes its slope (5) starting from $\lambda_2 = 0$, we can determine the entire solution path.

Regarding the change points, it can be shown that the optimality condition for (3)

$$\frac{\partial f}{\partial \beta_k}(\widehat{\beta}) = \widehat{\beta}_k - y_k + \lambda_2 \sum_{l:(k,l) \in E} \widehat{t}_{kl}(\lambda_2) = 0,$$

where $\widehat{t}_{kl}(\lambda_2) = \operatorname{sgn}(\widehat{\beta}_k(\lambda_2) - \widehat{\beta}_l(\lambda_2))$, is satisfied by affine functions $\widehat{\beta}_k(\lambda_2)$ and $\widehat{t}_{kl}(\lambda_2) = \lambda_2 \widehat{t}_{kl}(\lambda_2)$ for an interval $[\lambda_2^0, \lambda_2^0 + \Delta]$, with λ_2^0 being a change point. In this case, we have

$$\frac{\partial \widehat{\beta}_{F_i}(\lambda_2^0)}{\partial \lambda_2} + \sum_{l \in F_j(\lambda_2^0), j \neq i, (k,l) \in E} \widehat{t}_{kl}(\lambda_2^0) + \sum_{l \in F_i(\lambda_2^0), (k,l) \in E} \frac{\partial \widehat{t}_{kl}(\lambda_2^0)}{\partial \lambda_2} = 0, \tag{6}$$

for $k \in F_i(\lambda_2^0), i = 1, 2, \dots, n_F(\lambda_2^0)$. At the beginning of the interval, $\partial \widehat{\beta}_{F_i} / \partial \lambda_2$ is given by (5), and $\widehat{t}_{kl} = \pm 1$ for $(k, l) \in E, l \in F_j(\lambda_2^0)$ with $i \neq j$. The length of the interval Δ , hence the next change point, can be determined by examining where the sets of fused coefficients merge or split. Similar to the standard FLSA, at most two sets can merge, and a set splits into at most two sets for a sufficiently small increase in λ_2 . The first merge after λ_2^0 may occur when the paths for any two sets of fused coefficients meet:

$$h(\lambda_2^0) = \min_{i=1, \dots, n_F(\lambda_2^0)} \min_{j: h_{ij}(\lambda_2^0) > \lambda_2^0} h_{ij}(\lambda_2^0), \quad \text{where } h_{ij}(\lambda_2^0) = \frac{\widehat{\beta}_{F_i}(\lambda_2^0) - \widehat{\beta}_{F_j}(\lambda_2^0)}{\frac{\partial \widehat{\beta}_{F_i}(\lambda_2^0)}{\partial \lambda_2} - \frac{\partial \widehat{\beta}_{F_j}(\lambda_2^0)}{\partial \lambda_2}} + \lambda_2^0,$$

unless a split occurs before this point. The first split may occur at the point where the condition (6) is violated, that is,

$$v(\lambda_2^0) = \min_{i=1, \dots, n_F(\lambda_2^0)} \delta_i + \lambda_2^0,$$

where δ_i can be found by solving the linear program (LP)

$$\begin{aligned} &\text{minimize } 1/\delta \\ &\text{subject to } \sum_{l \in F_i(\lambda_2^0), (k,l) \in E} f_{kl} = p_k, \quad k \in F_i(\lambda_2^0), \\ &\quad -1 - (1/\delta)(\lambda_2^0 + \widehat{t}_{kl}(\lambda_2^0)) \leq f_{kl} \leq 1 + (1/\delta)(\lambda_2^0 - \widehat{t}_{kl}(\lambda_2^0)), \\ &\quad k, l \in F_i(\lambda_2^0), (k, l) \in E, \end{aligned} \tag{7}$$

in $1/\delta$ and f_{kl} , where $p_k = -\frac{\partial \widehat{\beta}_{F_i}(\lambda_2^0)}{\partial \lambda_2} - \sum_{l \in F_j(\lambda_2^0), j \neq i, (k,l) \in E} \widehat{t}_{kl}(\lambda_2^0)$, for each $i = 1, 2, \dots, n_F(\lambda_2^0)$. This LP can be solved by iteratively applying a maximum flow algorithm, for example, that of Ford and Fulkerson (1956). The maximum flow algorithm solves (7)

for fixed δ , and the final solution \hat{f}_{kl} to (7) gives the value of $\frac{\partial \hat{f}_{kl}}{\partial \lambda_2}(\lambda_2^0)$ in (6). Thus, $\Delta = \min\{h(\lambda_2^0), v(\lambda_2^0)\} - \lambda_2^0$. At split, (7) also determines how the set F_i is partitioned.

The computational complexity of this path algorithm depends almost on that of solving the maximum flow problem corresponding to (7). It becomes inefficient as the dimension p increases because a large-scale maximum flow problem is difficult to solve. However, when restricted to the standard FLSA, this path algorithm is very efficient, since the solution path of the standard FLSA has only fusions and not splits, hence does not require solving the maximum flow problem.

Path algorithm for the generalized lasso. Tibshirani and Taylor (2011) proposed a path algorithm for the generalized lasso problem, which replaces the penalty terms in (1) with a generalized ℓ_1 -norm penalty $\lambda \|\mathbf{D}\beta\|_1$:

$$\min_{\beta} f(\beta) \equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\mathbf{D}\beta\|_1, \quad (8)$$

where λ is a nonnegative regularization parameter and \mathbf{D} is a $m \times p$ dimensional matrix corresponding to the dependent structure of coefficients. The generalized lasso has the FLR with the generalized fusion penalty as a special case. The dual problem of (8) is given by

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|\tilde{\mathbf{y}} - \tilde{\mathbf{D}}^T \mathbf{u}\|_2^2 \\ & \text{subject to} && \|\mathbf{u}\|_{\infty} \leq \lambda, \mathbf{D}^T \mathbf{u} \in \text{row}(\mathbf{X}) \end{aligned} \quad (9)$$

on $\mathbf{u} \in \mathbb{R}^m$, where $\|\mathbf{u}\|_{\infty} = \max_j |u_j|$, $\tilde{\mathbf{y}} = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{\dagger} \mathbf{X}^T \mathbf{y}$, $\tilde{\mathbf{D}} = \mathbf{D}(\mathbf{X}^T \mathbf{X})^{\dagger} \mathbf{X}^T$, and $\text{row}(\mathbf{X})$ is a row space of \mathbf{X} ; A^{\dagger} denotes the Moore–Penrose pseudo-inverse of the matrix A . This dual problem is difficult to solve due to the constraint $\mathbf{D}^T \mathbf{u} \in \text{row}(\mathbf{X})$, but this row space constraint can be removed when $\text{rank}(\mathbf{X}) = p$. Thus, the path algorithm focuses on the case when the design matrix \mathbf{X} has a full rank.

The algorithm starts from $\lambda = \infty$ and decreases λ . Suppose λ_k is the current value of the regularization parameter λ . The current dual solution $\hat{\mathbf{u}}(\lambda_k)$ is given by

$$\begin{aligned} \hat{\mathbf{u}}_{\mathcal{B}}(\lambda_k) &= \lambda_k \mathbf{s}, \\ \hat{\mathbf{u}}_{-\mathcal{B}}(\lambda_k) &= (\mathbf{D}_{-\mathcal{B}}(\mathbf{D}_{-\mathcal{B}})^T)^{\dagger} \mathbf{D}_{-\mathcal{B}}(\mathbf{y} - \lambda_k (\mathbf{D}_{\mathcal{B}})^T \mathbf{s}), \end{aligned}$$

where $\mathcal{B} = \mathcal{B}(\lambda) = \{j \mid |u_j| = \lambda, j = 1, 2, \dots, m\}$ is the active set of the constraint $\|\mathbf{u}\|_{\infty} \leq \lambda$; $-\mathcal{B} = \mathcal{B}^c$; for an index set $A \subset \{1, 2, \dots, m\}$ and for a vector $\mathbf{v} = (v_i)_{1 \leq i \leq m} \in \mathbb{R}^m$, \mathbf{v}_A denotes the $|A|$ -dimensional vector such that $\mathbf{v}_A = (v_i)_{i \in A}$, and for a matrix $\mathbf{D} = (\mathbf{d}_i)_{1 \leq i \leq m} \in \mathbb{R}^{p \times m}$ with $\mathbf{d}_i \in \mathbb{R}^p$, \mathbf{D}_A denotes the $p \times |A|$ dimensional matrix such that $\mathbf{D}_A = (\mathbf{d}_i)_{i \in A}$; and $\mathbf{s} = (s_1, \dots, s_{|\mathcal{B}|})^T$ with $s_k = \text{sign}(u_{n_k}(\lambda_k))$, $n_k \in \mathcal{B}$. The primal solution is obtained from the dual solution by the following primal-dual relationship:

$$\hat{\beta}(\lambda) = \tilde{\mathbf{y}} - \tilde{\mathbf{D}}^T \hat{\mathbf{u}}(\lambda), \quad (10)$$

hence the solution path is piecewise linear. The intervals in which all $\hat{\beta}_j(\lambda)$ have constant slopes can be found by keeping track of the active set \mathcal{B} . An inactive coordinate $j \in -\mathcal{B}$ at $\lambda = \lambda_k$, possibly joins \mathcal{B} at the hitting time

$$t_j^{(\text{hit})} = \frac{[(\mathbf{D}_{-\mathcal{B}}(\mathbf{D}_{-\mathcal{B}})^T)^{\dagger} \mathbf{D}_{-\mathcal{B}} \mathbf{y}]_j}{[(\mathbf{D}_{-\mathcal{B}}(\mathbf{D}_{-\mathcal{B}})^T)^{\dagger} \mathbf{D}_{-\mathcal{B}} (\mathbf{D}_{\mathcal{B}})^T \mathbf{s}]_j \pm 1},$$

for which only one is less than λ_k , and an active coordinate $j' \in \mathcal{B}$, possibly leaves \mathcal{B} at the leaving time

$$t_{j'}^{(\text{leave})} = \begin{cases} \frac{s_{j'}[\mathbf{D}_{\mathcal{B}}[\mathbf{I} - \mathbf{D}_{-\mathcal{B}}^T(\mathbf{D}_{-\mathcal{B}}\mathbf{D}_{-\mathcal{B}}^T)^\dagger\mathbf{D}_{-\mathcal{B}}]\mathbf{y}]_{j'}}{s_{j'}[\mathbf{D}_{\mathcal{B}}[\mathbf{I} - \mathbf{D}_{-\mathcal{B}}^T(\mathbf{D}_{-\mathcal{B}}\mathbf{D}_{-\mathcal{B}}^T)^\dagger\mathbf{D}_{-\mathcal{B}}]\mathbf{D}_{\mathcal{B}}^T\mathbf{s}]_{j'}} \equiv \frac{c_{j'}}{d_{j'}}, & \text{if } c_{j'}, d_{j'} < 0, \\ 0, & \text{otherwise.} \end{cases}$$

Thus \mathcal{B} changes at $\lambda = \lambda_{k+1}$ where

$$\lambda_{k+1} = \max \left\{ \max_{j \in -\mathcal{B}} t_j^{(\text{hit})}, \max_{j \in \mathcal{B}} t_j^{(\text{leave})} \right\}.$$

The path algorithm for the generalized lasso problem can solve various ℓ_1 -regularization problems and obtains the exact solution path. After sequentially solving the dual, the solution path is obtained by the primal-dual relationship (10). It is also efficient for solving the standard FLSA. However, if the rank of the design matrix \mathbf{X} is not full rank, the additional constraint $\mathbf{D}^T \mathbf{u} \in \text{row}(\mathbf{X})$ does not disappear, making the dual (9) difficult to solve. To resolve this problem, Tibshirani and Taylor (2011) suggested adding an additional penalty $\epsilon \|\beta\|_2^2$ to the primal (8), which in effect substitutes \mathbf{y} and \mathbf{X} with $\mathbf{y}^* = (\mathbf{y}^T, \mathbf{0}^T)^T$ and $\mathbf{X}^* = (\mathbf{X}^T, \sqrt{\epsilon}\mathbf{I})^T$, respectively. This modification makes the new design matrix \mathbf{X}^* full rank, but a small value of ϵ may lead an ill-conditioned matrix $(\mathbf{X}^*)^T \mathbf{X}^*$. A further difficulty is that as the number of rows m of the matrix \mathbf{D} , that is, the number of penalty terms on β , increases, the path algorithm becomes less efficient since m is the number of dual variables.

2.2 FIRST-ORDER METHODS

To avoid the restrictions in the solution path methods, several optimization algorithms, which target to solve the FLR problem with a fixed set of regularization parameters, have been developed. These algorithms can solve the FLR problems with the general design matrix \mathbf{X} regardless of its rank. For scalability with the dimension p , they employ gradient descent-flavored first-order optimization methods.

Efficient fused lasso algorithm. Liu et al. (2010) proposed the efficient fused lasso algorithm (EFLA) that solves the standard FLR

$$\min_{\beta} f(\beta) \equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \sum_{j=1}^p |\beta_j| + \lambda_2 \sum_{j=2}^p |\beta_j - \beta_{j-1}|. \tag{11}$$

At the $(r + 1)$ th iteration, the EFLA solves (11) by minimizing an approximation of $f(\beta)$ in which the squared error term is replaced by its first-order Taylor expansion at the approximate solution $\widehat{\beta}^{(r)}$ obtained from the r th iteration, followed by an additional quadratic regularization term $\frac{L_r}{2} \|\beta - \widehat{\beta}^{(r)}\|_2^2$, $L_r > 0$. Minimization of this approximate objective can be written as

$$\min_{\beta} \frac{1}{2} \|\mathbf{v}^{(r)} - \beta\|_2^2 + \frac{\lambda_1}{L_r} \sum_{j=1}^p |\beta_j| + \frac{\lambda_2}{L_r} \sum_{j=2}^p |\beta_j - \beta_{j-1}|, \tag{12}$$

where $\mathbf{v}^{(r)} = \widehat{\beta}^{(r)} - (\mathbf{X}^T \mathbf{X} \widehat{\beta}^{(r)} + \mathbf{X}^T \mathbf{y})/L_r$. The minimizer of (12) is denoted by $\widehat{\beta}^{(r+1)}$. Note that this problem is the standard FLSA.

Although (12) can be solved by the path algorithms reviewed in Section 2, for fixed λ_1 and λ_2 , more efficient methods can be employed. Liu et al. (2010) advocated the use of a gradient descent method on the dual of (12), given by

$$\min_{\|\mathbf{u}\|_\infty \leq \lambda_2} \frac{1}{2} \mathbf{u}^T \mathbf{D} \mathbf{D}^T \mathbf{u} - \mathbf{u}^T \mathbf{D} \mathbf{v}^{(r)}, \quad (13)$$

where $\mathbf{D} \in \mathbb{R}^{(p-1) \times p}$ is the finite difference matrix on the one-dimensional grid of size p , a special case of the $m \times p$ penalty matrix of the generalized lasso (8). Note here that we set $\lambda_1 = 0$ without loss of generality, due to relation (2). Solution to (12) is obtained from the primal-dual relationship $\widehat{\beta}^{(r+1)} = \mathbf{v}^{(r)} - \mathbf{D}^T \mathbf{u}^I$, where \mathbf{u}^I is the solution to (13). Since (13) is a box-constrained quadratic program, it is solved efficiently by the following projected gradient method:

$$\widehat{\mathbf{u}}^{(k+1)} = P_{\lambda_2}(\widehat{\mathbf{u}}^{(k)} - \alpha(\mathbf{D} \mathbf{D}^T \widehat{\mathbf{u}}^{(k)} - \mathbf{D} \mathbf{v}^{(r)})),$$

where $P_{\lambda_2}(\cdot)$ is the projection onto the l_∞ -ball $\mathcal{B} = \{\mathbf{u} : \|\mathbf{u}\|_\infty \leq \lambda_2\}$ with radius λ_2 , and α is the reciprocal of the largest eigenvalue of the matrix $\mathbf{D} \mathbf{D}^T$. Further acceleration is achieved by using a restart technique that keeps track of the coordinates of \mathbf{u} for which the box constraints are active, and by replacing $\widehat{\beta}^{(r)}$ in (11) with $\tilde{\beta}^{(r)} = \widehat{\beta}^{(r)} + \eta_r(\widehat{\beta}^{(r)} - \widehat{\beta}^{(r-1)})$. The acceleration constant η_r and the additional regularization constant L_r in (12) are chosen using Nesterov's method (Nesterov 2004, 2007).

The efficiency of the EFLA strongly depends on the special structure that the finite difference matrix \mathbf{D} takes. Hence, its generalizability to the generalized fusion penalty other than the standard one is limited. One may want to use one of the path algorithms for FLSA shown in Section 2 instead of solving the dual (13). However, this modification does not assure the efficiency of the algorithm, because a path algorithm always starts from $\lambda_2 = 0$ (Hoeffling 2010) or $\lambda_2 = \infty$ (Tibshirani and Taylor 2011) while for EFLA a fixed value of λ_2 suffices.

Smoothing proximal gradient method. Chen et al. (2012) proposed the smoothing proximal gradient (SPG) method that solves the regression problems with structured penalties including the generalized fusion penalty. Recall the FLR problem (1) can be written as

$$\begin{aligned} \min_{\beta} f(\beta) &\equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_2 \sum_{(j,k) \in E} |\beta_j - \beta_k| + \lambda_1 \sum_{j=1}^p |\beta_j| \\ &\equiv \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_2 \|\mathbf{D}\beta\|_1 + \lambda_1 \|\beta\|_1, \end{aligned} \quad (14)$$

where \mathbf{D} is an $m \times p$ dimensional matrix such that $\|\mathbf{D}\beta\|_1 = \sum_{(j,k) \in E} |\beta_j - \beta_k|$. The main idea of the SPG method is to approximate the nonsmooth and nonseparable penalty term $\|\mathbf{D}\beta\|_1$ by a smooth function and to solve the resulting smooth surrogate problem iteratively in a computationally efficient fashion.

For the smooth approximation, recall that the penalty term $\|\mathbf{D}\beta\|_1$ can be reformulated as a maximization problem with an auxiliary vector $\alpha \in \mathbb{R}^m$,

$$\|\mathbf{D}\beta\|_1 \equiv \max_{\|\alpha\|_\infty \leq 1} \alpha^T \mathbf{D}\beta \equiv \Omega(\beta).$$

Now consider a smooth function

$$\tilde{\Omega}_\mu(\beta) \equiv \max_{\|\alpha\|_\infty \leq 1} \left(\alpha^T \mathbf{D}\beta - \frac{\mu}{2} \|\alpha\|_2^2 \right), \tag{15}$$

that approximates $\Omega(\beta)$ with a positive smoothing parameter μ , chosen as ϵ/m to obtain the desired accuracy $\epsilon > 0$ of the approximation. The function $\tilde{\Omega}_\mu(\beta)$ is convex and continuously differentiable with respect to β (Nesterov 2005).

For the smooth surrogate problem, the SPG method solves the following:

$$\min_{\beta} \tilde{f}(\beta) \equiv \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_2 \tilde{\Omega}(\beta, \mu)}_{h(\beta)} + \lambda_1 \|\beta\|_1. \tag{16}$$

This is the sum of a smooth convex function $h(\beta)$ and a nonsmooth but *separable* function $\lambda_1 \|\beta\|_1$, hence can be solved efficiently using a proximal gradient method, for example, the fast iterative shrinkage-thresholding algorithm (FISTA; Beck and Teboulle 2009). The FISTA iteratively approximates $h(\beta)$ by its first-order Taylor expansion with an additional quadratic regularization term of the form $\frac{L}{2} \|\beta - \hat{\beta}^{(r)}\|_2^2$, with essentially the same manner as the EFLA that yields (12). As a result, the FISTA solves at each iteration a subproblem that minimizes the sum of a quadratic function of β whose Hessian is the identity, and the classical lasso penalty $\lambda_1 \|\beta\|_1$. The solution of this subproblem is readily given by soft-thresholding.

The computational complexity of the SPG method depends on that of determining the constant L for the additional quadratic regularization term, which the FISTA chooses as the Lipschitz constant of the gradient of $h(\beta)$, that is, $L = \|\mathbf{X}\|_2^2 + \frac{\lambda_2}{\mu} \|\mathbf{D}\|_2^2$, where $\|\mathbf{A}\|_2$ is the spectral norm of a matrix \mathbf{A} . The computational complexity for computing L is $O(\min(m^2 p, mp^2))$, which is costly if either m or p increases. To reduce this computational cost, Chen et al. (2012) suggested using a line search on L , only to ensure that at the minimum of objective of the FISTA subproblem is not less than the surrogate function $\tilde{f}(\beta)$ evaluated at the minimizer of the same problem. However, since such an evaluation of $\tilde{f}(\beta)$ takes $O(\max(mp, np))$ time, this scheme does not save much cost if the surrogate function is evaluated frequently, especially when either m or p is large.

Split Bregman algorithm. Ye and Xie (2011) proposed the split Bregman (SB) algorithm for the standard FLR and it can directly extend to the general FLR (1) and (14). Note that (14) can be reformulated as a constrained form:

$$\begin{aligned} \min_{\beta, \mathbf{a}, \mathbf{b}} \quad & \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \|\mathbf{a}\|_1 + \lambda_2 \|\mathbf{b}\|_1 \\ \text{subject to} \quad & \mathbf{a} = \beta \\ & \mathbf{b} = \mathbf{D}\beta, \end{aligned} \tag{17}$$

with auxiliary variables $\mathbf{a} \in \mathbb{R}^p$ and $\mathbf{b} \in \mathbb{R}^m$. The SB algorithm is derived from the augmented Lagrangian (Hestenes 1969; Rockafellar 1973) of (17) given by

$$L(\beta, \mathbf{a}, \mathbf{b}, \mathbf{u}, \mathbf{v}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \|\mathbf{a}\|_1 + \lambda_2 \|\mathbf{b}\|_1 + \mathbf{u}^T(\beta - \mathbf{a}) + \mathbf{v}^T(\mathbf{D}\beta - \mathbf{b}) + \frac{\mu_1}{2} \|\beta - \mathbf{a}\|_2^2 + \frac{\mu_2}{2} \|\mathbf{D}\beta - \mathbf{b}\|_2^2, \quad (18)$$

where $\mathbf{u} \in \mathbb{R}^p$, $\mathbf{v} \in \mathbb{R}^m$ are dual variables, and μ_1, μ_2 are positive constants. This is the usual Lagrangian of (17) augmented by adding the quadratic penalty terms $\frac{\mu_1}{2} \|\beta - \mathbf{a}\|_2^2$ and $\frac{\mu_2}{2} \|\mathbf{D}\beta - \mathbf{b}\|_2^2$, for violating the equality constraints in (17). The primal problem associated with the augmented Lagrangian (18) is to minimize $\sup_{\mathbf{u}, \mathbf{v}} L(\beta, \mathbf{a}, \mathbf{b}, \mathbf{u}, \mathbf{v})$ over $(\beta, \mathbf{a}, \mathbf{b})$, and the dual is to maximize $\inf_{\beta, \mathbf{a}, \mathbf{b}} L(\beta, \mathbf{a}, \mathbf{b}, \mathbf{u}, \mathbf{v})$ over (\mathbf{u}, \mathbf{v}) .

The SB algorithm solves the primal and the dual in an alternating fashion, in which the primal is solved separately for each of β , \mathbf{a} , and \mathbf{b} :

$$\begin{aligned} \widehat{\beta}^{(r)} = \arg \min_{\beta} & \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + (\mathbf{u}^{(r-1)})^T (\beta - \mathbf{a}^{(r-1)}) + (\mathbf{v}^{(r-1)})^T (\mathbf{D}\beta - \mathbf{b}^{(r-1)}) \\ & + \frac{\mu_1}{2} \|\beta - \mathbf{a}^{(r-1)}\|_2^2 + \frac{\mu_2}{2} \|\mathbf{D}\beta - \mathbf{b}^{(r-1)}\|_2^2, \end{aligned} \quad (19)$$

$$\mathbf{a}^{(r)} = \arg \min_{\mathbf{a}} \lambda_1 \|\mathbf{a}\|_1 + (\widehat{\beta}^{(r)} - \mathbf{a})^T \mathbf{u}^{(r-1)} + \frac{\mu_1}{2} \|\widehat{\beta}^{(r)} - \mathbf{a}\|_2^2, \quad (20)$$

$$\mathbf{b}^{(r)} = \arg \min_{\mathbf{b}} \lambda_1 \|\mathbf{b}\|_1 + (\mathbf{D}\widehat{\beta}^{(r)} - \mathbf{b})^T \mathbf{v}^{(r-1)} + \frac{\mu_2}{2} \|\mathbf{D}\widehat{\beta}^{(r)} - \mathbf{b}\|_2^2, \quad (21)$$

while the dual is solved by gradient ascent:

$$\mathbf{u}^{(r)} = \mathbf{u}^{(r-1)} + \delta_1 (\widehat{\beta}^{(r)} - \mathbf{a}^{(r)}), \quad (22)$$

$$\mathbf{v}^{(r)} = \mathbf{v}^{(r-1)} + \delta_2 (\mathbf{D}\widehat{\beta}^{(r)} - \mathbf{b}^{(r)}). \quad (23)$$

where δ_1 and δ_2 are the step sizes.

The dual updates (22) and (23) are simple; the primal updates for the auxiliary variables (20) and (21) are readily given by soft-thresholding, in a similar manner as the SPG case (16). The primal update for the coefficient β (19) is equivalent to solving the linear system

$$(\mathbf{X}^T \mathbf{X} + \mu_1 \mathbf{I} + \mu_2 \mathbf{D}^T \mathbf{D})\beta = \mathbf{X}^T \mathbf{y} + (\mu_1 \mathbf{a}^{(r-1)} - \mathbf{u}^{(r-1)}) + \mathbf{D}^T (\mu_2 \mathbf{b}^{(r-1)} - \mathbf{v}^{(r-1)}) \quad (24)$$

whose computational complexity is $O(p^3)$ for general \mathbf{D} . However, for standard FLR in which \mathbf{D} is the finite difference matrix on a one-dimensional grid, (24) can be solved quickly by using the PCG algorithm (Demmel 1997). In this case, $\mathbf{M} = \mu_1 \mathbf{I} + \mu_2 \mathbf{D}^T \mathbf{D}$ is used as a preconditioner because \mathbf{M} is tridiagonal and easy to invert.

Note that the separation of updating Equations (19), (20), and (21) for the primal is possible because the nondifferentiable and nonseparable ℓ_1 -norm penalties on the coefficients are transferred to the auxiliary variables that are completely decoupled and separable. Also note that the updating equations for the auxiliary variables (20) and (21) reduce to soft-thresholding due to the augmented quadratic terms and the separable nondifferentiable terms.

The SB algorithm needs to choose the augmentation constants μ_1 and μ_2 and the step sizes δ_1 and δ_2 . In the implementation, Ye and Xie (2011) used a common value μ for all

of the four quantities, and consider a pretrial procedure to obtain highest convergence rate. Although the choice of μ does not affect the convergence of the algorithm, it is known that the rate of convergence is very sensitive to this choice (Lin et al. 2011; Ghadimi et al. 2012). This makes the SB algorithm to stall or fail to reach the optimal solution; see Appendix B for further details.

Alternating linearization. Lin et al. (2011) proposed the alternating linearization (ALIN) algorithm for the generalized lasso (8). Rewrite the problem (8) as a sum of two convex functions

$$\min_{\beta} f(\beta) \equiv \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2}_{g(\beta)} + \underbrace{\lambda \|\mathbf{D}\beta\|_1}_{h(\beta)}. \quad (25)$$

The main idea of the ALIN algorithm is to alternately solve two subproblems, in which $g(\beta)$ is linearized (h -subproblem) and added by a quadratic regularization term, and so is $h(\beta)$ (g -subproblem), respectively. This algorithm maintains three sequences of solutions: $\{\widehat{\beta}^{(r)}\}$ is the sequence of solutions of the problem (25); $\{\tilde{\beta}_h^{(r)}\}$ and $\{\tilde{\beta}_g^{(r)}\}$ are the sequences of solutions of the h -subproblem and the g -subproblem, respectively.

For r th iteration of subproblems, the h -subproblem is given by

$$\tilde{\beta}_h^{(r)} = \arg \min_{\beta} \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{X}\tilde{\beta}_g^{(r-1)}\|_2^2 + \mathbf{s}_g^T (\beta - \tilde{\beta}_g^{(r-1)})}_{\tilde{g}(\beta)} + \frac{1}{2} \|\beta - \widehat{\beta}^{(r-1)}\|_{\mathbf{A}}^2 + \lambda \|\mathbf{D}\beta\|_1, \quad (26)$$

where $\tilde{g}(\beta)$ is a linearization of $g(\beta)$ at $\tilde{\beta}_g^{(r-1)}$, $\mathbf{s}_g = \nabla g(\tilde{\beta}_g^{(r-1)}) = \mathbf{X}^T (\mathbf{X}\tilde{\beta}_g^{(r-1)} - \mathbf{y})$, and $\|\beta - \widehat{\beta}\|_{\mathbf{A}}^2 = (\beta - \widehat{\beta})^T \mathbf{A} (\beta - \widehat{\beta})$ with a diagonal matrix $\mathbf{A} = \text{diag}(\mathbf{X}^T \mathbf{X})$. Problem (26) is equivalent to

$$\min_{\beta, \mathbf{z}} \mathbf{s}_g^T \beta + \frac{1}{2} \|\beta - \widehat{\beta}^{(r-1)}\|_{\mathbf{A}}^2 + \lambda \|\mathbf{z}\|_1 \quad \text{subject to} \quad \mathbf{D}\beta = \mathbf{z}. \quad (27)$$

Similar to (9), the dual of (27) is given by

$$\min_{\mathbf{u}} \frac{1}{2} \mathbf{u}^T \mathbf{D} \mathbf{A}^{-1} \mathbf{D}^T \mathbf{u} - \mathbf{u}^T \mathbf{D} (\widehat{\beta}^{(r-1)} - \mathbf{A}^{-1} \mathbf{s}_g) \quad \text{subject to} \quad \|\mathbf{u}\|_{\infty} \leq \lambda. \quad (28)$$

The solution of h -subproblem is obtained by the primal–dual relationship $\tilde{\beta}_h^{(r)} = \widehat{\beta}^{(r-1)} - \mathbf{A}^{-1} (\mathbf{s}_g + \mathbf{D}^T \mathbf{u}^*)$, where \mathbf{u}^* is the optimal solution of (28). The dual (28) is efficiently solved by an active-set box-constrained PCG algorithm with $\text{diag}(\mathbf{D} \mathbf{A}^{-1} \mathbf{D}^T)$ as a preconditioner.

The g -subproblem is also given by

$$\begin{aligned} \tilde{\beta}_g^{(r)} &= \arg \min_{\beta} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \underbrace{\lambda \|\mathbf{D}\tilde{\beta}_h^{(r)}\|_1 + \mathbf{s}_h^T (\beta - \tilde{\beta}_h^{(r)})}_{\tilde{h}(\beta)} + \frac{1}{2} \|\beta - \widehat{\beta}^{(r-1)}\|_{\mathbf{A}}^2 \\ &= \arg \min_{\beta} \mathbf{s}_h^T \beta + \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \frac{1}{2} \|\beta - \widehat{\beta}^{(r-1)}\|_{\mathbf{A}}^2, \end{aligned} \quad (29)$$

where \mathbf{s}_h is the subgradient of $h(\beta)$ at $\tilde{\beta}_h^{(r)}$ that can be calculated as $\mathbf{s}_h = -\mathbf{s}_g - \mathbf{A} (\tilde{\beta}_h^{(r)} - \widehat{\beta}^{(r-1)})$. Problem (29) is equivalent to solving the following linear system:

$$(\mathbf{X}^T \mathbf{X} + \mathbf{A}) (\beta - \widehat{\beta}^{(r-1)}) = \mathbf{X}^T (\mathbf{y} - \mathbf{X}\widehat{\beta}^{(r-1)}) - \mathbf{s}_h, \quad (30)$$

which can be solved by using the PCG algorithm with $\mathbf{A} = \text{diag}(\mathbf{X}^T \mathbf{X})$ as a preconditioner.

After solving each subproblem, the ALIN algorithm checks its stopping and updating criteria. The stopping criteria for the h -subproblem and the g -subproblem are defined as

$$\begin{aligned} h\text{-sub} &: \tilde{g}(\tilde{\beta}_h^{(r)}) + h(\tilde{\beta}_h^{(r)}) \geq g(\widehat{\beta}^{(r-1)}) + h(\widehat{\beta}^{(r-1)}) - \epsilon, \\ g\text{-sub} &: g(\tilde{\beta}_g^{(r)}) + \tilde{h}(\tilde{\beta}_g^{(r)}) \geq g(\widehat{\beta}^{(r-1)}) + h(\widehat{\beta}^{(r-1)}) - \epsilon, \end{aligned}$$

where ϵ is a tolerance of the algorithm. If one of the stopping criteria is met, then the algorithm terminates; otherwise it proceeds to check the updating criteria:

$$\begin{aligned} h\text{-sub} &: g(\tilde{\beta}_h^{(r)}) + h(\tilde{\beta}_h^{(r)}) \leq (1 - \gamma)[g(\widehat{\beta}^{(r-1)}) + h(\widehat{\beta}^{(r-1)})] + \gamma[\tilde{g}(\tilde{\beta}_h^{(r)}) + h(\tilde{\beta}_h^{(r)})], \\ g\text{-sub} &: g(\tilde{\beta}_g^{(r)}) + h(\tilde{\beta}_g^{(r)}) \leq (1 - \gamma)[g(\widehat{\beta}^{(r-1)}) + h(\widehat{\beta}^{(r-1)})] + \gamma[g(\tilde{\beta}_g^{(r)}) + \tilde{h}(\tilde{\beta}_g^{(r)})], \end{aligned}$$

where $\gamma \in (0, 1)$. If one of the updating criterion is satisfied, then $\widehat{\beta}^{(r)}$ is updated as the solution of the corresponding subproblem, that is, $\tilde{\beta}_g^{(r)}$ or $\tilde{\beta}_h^{(r)}$. Otherwise $\widehat{\beta}^{(r)}$ remains unchanged: $\widehat{\beta}^{(r)} = \widehat{\beta}^{(r-1)}$.

As it solves the dual problem (28), the ALIN algorithm has a similar drawback as the path algorithm for the generalized lasso: if the number of rows m of \mathbf{D} increases, solving the dual is less efficient than solving the primal. Moreover, the matrix $\mathbf{D}\mathbf{A}^{-1}\mathbf{D}^T$ in (28) is not positive definite when m is greater than p . This violates the assumption of the PCG.

2.3 SUMMARY OF THE REVIEWED ALGORITHMS

Table 1 summarizes the existing algorithms explained in Sections 2.1 and 2.2 and the MM algorithm proposed in Section 3 according to types of the design matrix \mathbf{X} and penalty structure. The items marked with a filled circle (\bullet) denote the availability of the algorithm.

Table 1. Summary of algorithms for solving the FLSA and the FLR problems

Method	FLSA ($\mathbf{X} = \mathbf{I}$)		FLR (general \mathbf{X})	
	STD	GEN	STD	GEN
Pathwise optimization (pathwise)	\bullet			
Path algorithm for the FLSA (pathFLSA)	\bullet	\bullet		
Path algorithm for the generalized lasso (genlasso)	\bullet	Δ	\blacktriangle	\blacktriangle
Efficient Fused Lasso Algorithm (EFLA)	\bullet		\bullet	
Smooth Proximal Gradient (SPG)	\bullet	\circ	\bullet	\circ
Split Bregman (SB)	\bullet	\bullet	\bullet	\bullet
Alternating Linearization (ALIN)	\bullet	\bullet	\bullet	\bullet
Majorization–Minimization (MM)	\bullet	\bullet	\bullet	\bullet
MM with GPU parallelization (MMGPU)	\bullet		\bullet	

Abbreviations of the names of the algorithms are given within the parentheses for future references.

STD denotes the standard fusion penalty $\sum_{j=2}^p |\beta_j - \beta_{j-1}|$.

GEN denotes the generalized fusion penalty $\sum_{(j,k) \in E} |\beta_j - \beta_k|$ for a given index set E .

\bullet denotes that the method is applicable.

Δ denotes that the dual solution path is computationally infeasible in large scale image denoising.

\blacktriangle denotes the method is only applicable when \mathbf{X} has full rank.

\circ denotes the method is not adequate since computation of spectral matrix norm when p is large.

3. MM ALGORITHM FOR FUSED LASSO PROBLEM

3.1 MM ALGORITHM

In this section, we propose an MM algorithm to solve the FLR (1). The MM algorithm iterates two steps, the majorization step and the minimization step. Given the current estimate $\widehat{\beta}^0$ of the optimal solution to (1), the majorization step finds a majorizing function $g(\beta|\widehat{\beta}^0)$ such that $f(\widehat{\beta}^0) = g(\widehat{\beta}^0|\widehat{\beta}^0)$ and $f(\beta) \leq g(\beta|\widehat{\beta}^0)$ for all $\beta \neq \widehat{\beta}^0$. The minimization step updates the estimate with

$$\widehat{\beta} = \arg \min_{\beta} g(\beta|\widehat{\beta}^0).$$

It is known that the MM algorithm has a descent property in the sense that $f(\widehat{\beta}) \leq f(\widehat{\beta}^0)$. Furthermore, if $f(\beta)$ and $g(\beta|\widehat{\beta}^0)$ satisfy some regularity conditions, for example, those in Vaida (2005), the MM algorithm converges to the optimal solution of $f(\beta)$.

The MM algorithm was applied to the classical lasso problem by Hunter and Li (2005). They proposed to minimize a perturbed version of the objective function of the classical lasso

$$f_{\lambda,\epsilon}(\beta) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \sum_{j=1}^p \left(|\beta_j| - \epsilon \log \left(1 + \frac{|\beta_j|}{\epsilon} \right) \right)$$

instead of

$$f_{\lambda}(\beta) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_1$$

to avoid division by zero in the algorithm. The majorizing function for $f_{\lambda,\epsilon}(\beta)$ at $\widehat{\beta}^0$ is given by

$$g_{\lambda,\epsilon}(\beta|\widehat{\beta}^0) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \sum_{j=1}^p \left\{ |\widehat{\beta}_j^0| - \epsilon \log \left(1 + \frac{|\widehat{\beta}_j^0|}{\epsilon} \right) + \frac{\beta_j^2 - (\widehat{\beta}_j^0)^2}{2(|\widehat{\beta}_j^0| + \epsilon)} \right\}.$$

Hunter and Li (2005) showed that the sequence $\{\widehat{\beta}^{(r)}\}_{r \geq 0}$ with $\widehat{\beta}^{(r+1)} = \arg \min_{\beta} g_{\lambda,\epsilon}(\beta|\widehat{\beta}^{(r)})$ converges to the minimizer of $f_{\lambda,\epsilon}(\beta)$, and that $f_{\lambda,\epsilon}(\beta)$ converges to $f_{\lambda}(\beta)$ uniformly as ϵ approaches 0.

Motivated by the above development, we introduce a perturbed version $f_{\epsilon}(\beta)$ of the objective (1) and the majorizing function $g_{\epsilon}(\beta|\widehat{\beta}^0)$ for the FLR problem (1):

$$\begin{aligned} f_{\epsilon}(\beta) &= \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \sum_{j=1}^p \left\{ |\beta_j| - \epsilon \log \left(1 + \frac{|\beta_j|}{\epsilon} \right) \right\} \\ &\quad + \lambda_2 \sum_{(j,k) \in E} \left\{ |\beta_j - \beta_k| - \epsilon \log \left(1 + \frac{|\beta_j - \beta_k|}{\epsilon} \right) \right\}, \end{aligned} \tag{31}$$

and

$$\begin{aligned}
 g_\epsilon(\beta|\widehat{\beta}^0) &= \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \sum_{j=1}^p \left\{ |\widehat{\beta}_j^0| - \epsilon \log \left(1 + \frac{|\widehat{\beta}_j^0|}{\epsilon} \right) + \frac{\beta_j^2 - (\widehat{\beta}_j^0)^2}{2(|\widehat{\beta}_j^0| + \epsilon)} \right\} \\
 &+ \lambda_2 \sum_{(j,k) \in E} \left\{ |\widehat{\beta}_j^0 - \widehat{\beta}_k^0| - \epsilon \log \left(1 + \frac{|\widehat{\beta}_j^0 - \widehat{\beta}_k^0|}{\epsilon} \right) + \frac{(\beta_j - \beta_k)^2 - (\widehat{\beta}_j^0 - \widehat{\beta}_k^0)^2}{2(|\widehat{\beta}_j^0 - \widehat{\beta}_k^0| + \epsilon)} \right\}.
 \end{aligned} \tag{32}$$

Note that $g_\epsilon(\beta|\widehat{\beta}^0)$ is a smooth convex function in β . Below we see that $g_\epsilon(\beta|\widehat{\beta}^0)$ indeed majorizes $f_\epsilon(\beta)$ at $\widehat{\beta}^0$, and has a unique global minimum regardless of the rank of \mathbf{X} , together with several properties of the perturbed objective $f_\epsilon(\beta)$:

Proposition 1. For $\epsilon > 0$ and $\lambda_1 > 0$, (i) $f_\epsilon(\beta)$ is continuous, convex, and satisfies $\lim_{\|\beta\|_2 \rightarrow \infty} f_\epsilon(\beta) = \infty$, (ii) $f_\epsilon(\beta)$ converges to $f(\beta)$ uniformly on any compact set \mathbf{C} as ϵ approaches to zero. (iii) $g_\epsilon(\beta|\beta')$ majorizes $f_\epsilon(\beta)$ at β' , and $g_\epsilon(\beta|\beta')$ has a unique global minimum for all β' .

Proof. See Appendix A.1. □

Now let $\widehat{\beta}^{(r)}$ be the current estimate in the r th iteration. The majorizing function $g_\epsilon(\beta|\widehat{\beta}^{(r)})$ is minimized over β when $\partial g_\epsilon(\beta|\widehat{\beta}^{(r)})/\partial\beta = 0$, which can be written as a linear system of equations

$$(\mathbf{X}^T \mathbf{X} + \lambda_1 \mathbf{A}^{(r)} + \lambda_2 \mathbf{B}^{(r)})\beta = \mathbf{X}^T \mathbf{y}, \tag{33}$$

where $\mathbf{A}^{(r)} = \text{diag}(a_1^{(r)}, a_2^{(r)}, \dots, a_p^{(r)})$ with $a_j^{(r)} = 1/(|\widehat{\beta}_j^{(r)}| + \epsilon)$ for $1 \leq j \leq p$, and $\mathbf{B}^{(r)} = (b_{jk}^{(r)})_{1 \leq j, k \leq p}$ is a symmetric and positive semidefinite matrix with

$$\begin{aligned}
 b_{jj}^{(r)} &= \sum_{k:(j,k) \in E} \frac{1}{|\widehat{\beta}_j^{(r)} - \widehat{\beta}_k^{(r)}| + \epsilon}, \quad j = 1, \dots, p, \\
 b_{jk}^{(r)} &= b_{kj}^{(r)} = -\frac{1}{|\widehat{\beta}_j^{(r)} - \widehat{\beta}_k^{(r)}| + \epsilon}, \quad \forall (j, k) \in E.
 \end{aligned}$$

The procedure described so far is summarized as Algorithm 1. Note that once the matrix $\mathbf{X}^T \mathbf{X} + \lambda_1 \mathbf{A}^{(r)} + \lambda_2 \mathbf{B}^{(r)}$ is constructed, the linear system (33) can be efficiently solved. Since this matrix is symmetric and positive definite, hence Cholesky decomposition and back substitution can be applied, for example, using `dpotrf` and `dpotri` functions in LAPACK (Anderson et al. 1995).

Furthermore, the construction of the matrix $\lambda_1 \mathbf{A}^{(r)} + \lambda_2 \mathbf{B}^{(r)}$ is efficient in the MM algorithm. Recall from (14) that the generalized fusion penalty in (1) can be written as $\|\mathbf{D}\beta\|_1$ for a $m \times p$ dimensional matrix \mathbf{D} , where m is the number of penalty terms. In Algorithm 1, the matrix $\lambda_1 \mathbf{A}^{(r)} + \lambda_2 \mathbf{B}^{(r)}$ is constructed in $O(m)$ time, that is, independent of the problem dimension p . This is a great advantage of the MM algorithm over other algorithms reviewed in Section 2: SPG requires to check a specific condition to guarantee its convergence, which

takes at least $O(np)$ time for every iteration; SB needs to construct the matrix $\mu\mathbf{I} + \mu\mathbf{D}^T\mathbf{D}$, where μ is the parameter discussed in Section 2.2, and it takes $O(mp^2)$ time; ALIN and genlasso need to compute the matrix $\mathbf{D}\mathbf{D}^T$ for solving the dual, hence requires $O(m^2p)$ time. Although in the MM algorithm the matrix $\lambda_1\mathbf{A}^{(r)} + \lambda_2\mathbf{B}^{(r)}$ needs to be constructed for each iteration, our experience is that the MM algorithm tends to converge within a few tens of iterations; see Section 5. Thus, if p is more than a few tens, which is of our genuine interest, the MM algorithm is expected to run faster than the other algorithms.

In case the matrix $\mathbf{M}^{(r)} = \lambda_1\mathbf{A}^{(r)} + \lambda_2\mathbf{B}^{(r)}$ is tridiagonal as in the standard FLR, or block tridiagonal as in the two-dimensional FLR, we can employ the PCG method by using the matrix $\mathbf{M}^{(r)}$ as the preconditioner to efficiently solve the linear system (33). The PCG method is an iterative method for solving a linear system $\mathbf{Q}\mathbf{x} = \mathbf{c}$, where \mathbf{Q} is a symmetric positive definite matrix, with a preconditioner matrix \mathbf{M} (Demmel 1997). For example, we can use $\mathbf{M} = \text{diag}(\mathbf{X}^T\mathbf{X})$ for $\mathbf{Q} = \mathbf{X}^T\mathbf{X}$. The PCG method achieves the solution by iteratively updating four sequences $\{\mathbf{x}_j\}_{j \geq 0}$, $\{\mathbf{r}_j\}_{j \geq 0}$, $\{\mathbf{z}_j\}_{j \geq 0}$, and $\{\mathbf{p}_j\}_{j \geq 1}$ with

$$\begin{aligned} \mathbf{x}_j &= \mathbf{x}_{j-1} + v_j\mathbf{p}_j, \\ \mathbf{r}_j &= \mathbf{r}_{j-1} + v_j\mathbf{Q}\mathbf{p}_j, \\ \mathbf{z}_j &= \mathbf{M}^{-1}\mathbf{r}_j, \\ \mathbf{p}_{j+1} &= \mathbf{z}_j + \gamma_j\mathbf{p}_j, \end{aligned}$$

where $v_j = (\mathbf{z}_j^T\mathbf{r}_{j-1})/(\mathbf{p}_j^T\mathbf{Q}\mathbf{p}_j)$, $\gamma_j = (\mathbf{z}_j^T\mathbf{r}_j)/(\mathbf{z}_{j-1}^T\mathbf{r}_{j-1})$, $\mathbf{x}_0 = 0$, $\mathbf{r}_0 = \mathbf{c}$, $\mathbf{z}_0 = \mathbf{M}^{-1}\mathbf{c}$, and $\mathbf{p}_1 = \mathbf{z}_0$. For the standard FLR, the matrix $\mathbf{M}^{(r)}$ is tridiagonal and its Cholesky decomposition can be evaluated within $O(p)$ time, much faster than that of general positive-definite matrices. In this case, we can solve the linear system $\mathbf{M}^{(r)}\mathbf{z}_j = \mathbf{r}_j$ using LAPACK functions `dpbtrf` and `dpbtrs` that perform Cholesky decomposition and linear system solving for symmetric positive definite band matrix, respectively. Algorithm 2 describes the proposed MM algorithm using the PCG method for the standard FLR.

Algorithm 1 MM algorithm for the FLR

Input: \mathbf{y} , \mathbf{X} , λ_1 , λ_2 , convergence tolerance δ , perturbation constant ϵ .

- 1: **for** $j = 1, \dots, p$ **do**
- 2: $\widehat{\beta}_j^{(0)} \leftarrow \frac{X_j^T \mathbf{y}}{X_j^T X_j}$ ▷ initialization
- 3: **end for**
- 4: **repeat** $r = 0, 1, 2, \dots$
- 5: $\mathbf{M}^{(r)} \leftarrow \mathbf{X}^T\mathbf{X} + \lambda_1\mathbf{A}^{(r)} + \lambda_2\mathbf{B}^{(r)}$
- 6: $\widehat{\beta}^{(r+1)} \leftarrow (\mathbf{M}^{(r)})^{-1}\mathbf{X}^T\mathbf{y}$ ▷ using `dpotrf` and `dpotri` in LAPACK
- 7: **until** $\frac{|f(\widehat{\beta}^{(r+1)}) - f(\widehat{\beta}^{(r)})|}{|f(\widehat{\beta}^{(r)})|} \leq \delta$

Output: $\widehat{\beta}_\epsilon = \widehat{\beta}^{(r+1)}$

Algorithm 2 MM algorithm using the PCG method for the standard FLR and the two-dimensional FLR

Input: \mathbf{y} , \mathbf{X} , λ_1 , λ_2 , convergence tolerance δ , perturbation constant ϵ .

```

1: for  $j = 1, \dots, p$  do
2:    $\widehat{\beta}_j^{(0)} \leftarrow \frac{\mathbf{X}_j^T \mathbf{y}}{\mathbf{X}_j^T \mathbf{X}_j}$  ▷ initialization
3: end for
4: repeat  $r = 0, 1, 2, \dots$ 
5:    $\mathbf{M}^{(r)} \leftarrow \lambda_1 \mathbf{A}^{(r)} + \lambda_2 \mathbf{B}^{(r)}$  ▷ preconditioner
6:    $\mathbf{x}_0 \leftarrow \widehat{\beta}^{(r)}$  ▷ begin the PCG method
7:    $\mathbf{r}_0 \leftarrow \mathbf{X}^T \mathbf{y} - (\mathbf{X}^T \mathbf{X} + \mathbf{M}^{(r)}) \mathbf{x}_0$ 
8:    $\mathbf{z}_0 \leftarrow (\mathbf{M}^{(r)})^{-1} \mathbf{r}_0$ 
9:    $\mathbf{p}_1 \leftarrow \mathbf{z}_0$ 
10:  repeat  $j = 1, 2, \dots$ 
11:     $v_j \leftarrow \frac{\mathbf{r}_{j-1}^T \mathbf{z}_{j-1}}{\mathbf{p}_j^T (\mathbf{X}^T \mathbf{X} + \mathbf{M}^{(r)}) \mathbf{p}_j}$ 
12:     $\mathbf{x}_j \leftarrow \mathbf{x}_{j-1} + v_j \mathbf{p}_j$ 
13:     $\mathbf{r}_j \leftarrow \mathbf{r}_{j-1} - v_j (\mathbf{X}^T \mathbf{X} + \mathbf{M}^{(r)}) \mathbf{p}_j$ 
14:    Solve  $\mathbf{M}^{(r)} \mathbf{z}_j = \mathbf{r}_j$  ▷ using dgbtrf and dgbtrs in LAPACK
15:     $\gamma_j \leftarrow \frac{\mathbf{r}_j^T \mathbf{z}_j}{\mathbf{r}_{j-1}^T \mathbf{z}_{j-1}}$ 
16:     $\mathbf{p}_{j+1} \leftarrow \mathbf{z}_j + \gamma_j \mathbf{p}_j$  ▷ update conjugate gradient
17:  until  $\frac{\|\mathbf{r}_j\|_2}{\|\mathbf{X}^T \mathbf{y}\|_2} \leq \delta$  ▷ end the PCG method
18:   $\widehat{\beta}^{(r+1)} \leftarrow \mathbf{x}_j$  ▷ update solution
19: until  $\frac{|f(\widehat{\beta}^{(r+1)}) - f(\widehat{\beta}^{(r)})|}{|f(\widehat{\beta}^{(r)})|} \leq \delta$ 

```

Output: $\widehat{\beta}_\epsilon = \widehat{\beta}^{(r+1)}$

3.2 CONVERGENCE

In this section, we provide some results on the convergence of the proposed MM algorithm for the FLR. We first show that a minimizer $\widehat{\beta}_\epsilon$ of the perturbed version (31) of the objective (1) exhibits a minimal difference from a minimizer $\widehat{\beta}$ of the true objective (1) for sufficiently small ϵ . We then see that the sequence of the solutions $\{\widehat{\beta}_\epsilon^{(r)}\}_{r \geq 0}$ of the MM algorithm converges to $\widehat{\beta}_\epsilon$.

The key to the proof of the following lemma is to see that the level sets of both the perturbed objective $f_\epsilon(\beta)$ and the true objective $f(\beta)$ are compact, which is a consequence of Proposition 1.

Lemma 1. Consider an arbitrary decreasing sequence $\{\epsilon_n, n = 1, 2, \dots, \infty\}$ that converges to 0. Then, any limit point of $\widehat{\beta}_{\epsilon_n}$ is a minimizer of $f(\beta)$, provided that $\{\beta \mid f(\beta) = f(\widehat{\beta})\}$ is nonempty.

Proof. See Appendix A.2. □

Lemma 2 states that $f_\epsilon(\beta)$ serves as a Lyapunov function for the nonlinear recursive updating rule, implicitly defined by the proposed MM algorithm, that generates the sequence $\{\widehat{\beta}_\epsilon^{(r)}\}$. Monotone convergence immediately follows from the descent property of the MM algorithm.

Lemma 2. Let \mathbf{S} be a set of stationary points of $f_\epsilon(\beta)$ and $\{\widehat{\beta}_\epsilon^{(r)}\}_{r \geq 0}$ be a sequence of solutions of the proposed MM algorithm given by $\widehat{\beta}_\epsilon^{(r+1)} = M(\widehat{\beta}_\epsilon^{(r)}) = \arg \min_\beta g_\epsilon(\beta \mid \widehat{\beta}_\epsilon^{(r)})$. Then, the limit points of $\{\widehat{\beta}_\epsilon^{(r)}\}_{r \geq 0}$ are stationary points of $f_\epsilon(\beta)$, that is, $\partial f_\epsilon(\beta) / \partial \beta = 0$ at these limit points. Moreover, $f_\epsilon(\widehat{\beta}_\epsilon^{(r)})$ converges monotonically to $f_\epsilon(\beta_\epsilon^*)$ for some $\beta_\epsilon^* \in \mathbf{S}$.

Proof. See Appendix A.3. □

Combining Lemmas 1 and 2, it is straightforward to see that the MM algorithm generates solutions that converge to an optimal solution of the FLR problem (1) as ϵ decreases to zero.

Theorem 1. The sequence of the solutions $\{\widehat{\beta}_\epsilon^{(r)}\}_{r \geq 0}$ generated by the proposed MM algorithm converges to a minimizer of $f(\beta)$. Moreover, the sequence of functionals $\{f_\epsilon(\widehat{\beta}_\epsilon^{(r)})\}_{r \geq 0}$ converges to the minimum value of $f(\beta)$.

Proof. Note for the first part

$$\|\widehat{\beta}_\epsilon^{(r)} - \beta^{**}\| \leq \|\widehat{\beta}_\epsilon^{(r)} - \beta_\epsilon^*\| + \|\beta_\epsilon^* - \beta^{**}\|,$$

where β_ϵ^* is a stationary point of, hence minimizes, $f_\epsilon(\beta)$, and β^{**} is a limit point of $\{\beta_\epsilon^*\}$ as $\epsilon \downarrow 0$. The first term in the right-hand side becomes arbitrarily small for sufficiently large r by lemma 2, whereas the second term does for sufficiently small ϵ by Lemma 1. The limit point β^{**} is a minimizer of $f(\beta)$ also by Lemma 1. Now we see

$$\|f_\epsilon(\widehat{\beta}_\epsilon^{(r)}) - f(\beta^{**})\| \leq \|f_\epsilon(\widehat{\beta}_\epsilon^{(r)}) - f_\epsilon(\beta_\epsilon^*)\| + \|f_\epsilon(\widehat{\beta}_\epsilon) - f(\beta_\epsilon^*)\| + \|f(\beta_\epsilon^*) - f(\beta^{**})\|.$$

The first in the right-hand side vanishes by the continuity of $f_\epsilon(\beta)$; the second term by the uniform convergence of $f_\epsilon(\beta)$ to $f(\beta)$, as shown in the proof of Lemma 1; and the third term by the continuity of $f(\beta)$. □

4. PARALLELIZATION OF THE MM ALGORITHM WITH GPU

The GPU was originally developed as a co-processor to reduce the workload of the central processing unit (CPU) for computationally expensive graphics operations. A GPU has a collection of thousands of light-weight computing cores, which makes it suitable for massively parallel applications (Owens et al. 2008). While early GPUs could only perform simple bitmap processing for generating pixels on screen, as the applications require high-level three-dimensional graphics performance it has become possible that users can program so-called shader controlling directly each step of the graphics pipeline that generates pixels from the three-dimensional geometry models. As a GPU generates a massive number of pixels on screen simultaneously, shader programming can be understood as conducting the identical computation with many geometric primitives (such as vertices) as the input to determine the output, or the intensities of the pixels.

Thus for a general-purpose computing problem, if the problem has *data-level parallelism*, that is, on many portions of data the computation can be executed at the same time, then by writing an appropriate shader program we can map the data to the geometric primitives and can parallelize the general-purpose computing. Due to this single-instruction multiple-data (SIMD) parallel architecture together with its economic attraction, the use of the GPU has been expanded to general-purpose scientific computing beyond graphics computation. As the demand in general-purpose computing increases, both the architecture and the library for the GPU have been evolved to support such computation, for example, NVIDIA's Compute Unified Device Architecture (CUDA) (Kirk and Hwu 2010; Farber 2011). As mentioned above, applications in which the GPU is effective involve separation of data and parameters, such as nonnegative matrix factorization, positron emission tomography, and multidimensional scaling (Zhou et al. 2010).

For the standard FLR, the PCG step in (33) that is frequently encountered in the inner iteration of Algorithm 2 can be greatly accelerated by the GPU parallelization. The motivation of the parallelization comes from the parallel algorithms such as cyclic reduction (CR), parallel cyclic reduction (PCR), and recursive doubling (RD; Hockney 1965; Stone 1973; Hockney and Jesshope 1981) that can solve a tridiagonal system to which (33) reduces. Although these parallel algorithms require more operations than a basic Gaussian elimination for the tridiagonal system (Zhang et al. 2010), they can be executed more efficiently than the basic Gaussian elimination if many cores are used simultaneously. Moreover, a hybrid algorithm, which combines the CR and the PCR, has been developed for GPUs recently (Zhang et al. 2010) and is implemented in CUDA sparse matrix library (cuSPARSE) as the function `cusparsedgtsv`. We use this function to solve the linear system $\mathbf{M}^{(r)}\mathbf{z}_j = \mathbf{r}_j$ in the GPU implementation of Algorithm 2.

In addition to the tridiagonal system, Algorithm 2 also contains a set of linear operations suitable for GPU parallelization. For instance, lines 7, 11–13, and 15–17 rely on basic linear operations such as matrix–matrix multiplication, matrix–vector multiplication, vector–vector addition, innerproduct, and ℓ_2 -norm computation. These operations involve identical arithmetic for each coordinate, hence the SIMD architecture of GPU is appropriate for parallelizing the algorithm. The CUDA basic linear algebra subroutines library (cuBLAS) provides an efficient implementation of such operations, hence we use cuBLAS for parallelizing lines 7, 11–13, and 15–17 with GPU.

Unfortunately, our parallelization of the PCG step is limited to the standard FLR and at this moment is not extendable to the generalized fusion penalty. In this case, we recommend to solve the linear system (33) with LAPACK functions `dpotrf` and `dpotri` using CPUs as in Algorithm 1. Currently, this is more efficient than using GPUs.

5. NUMERICAL STUDIES

In this section, we compare the performance of the MM algorithm with the other existing algorithms using several datasets that generally fit into the “large p , small n ” setting. The general QP solvers are excluded in the comparison, since they are slower than other specialized algorithms for the FLR. These numerical studies include five scenarios for the FLR problem. In the first three scenarios, we consider the standard FLR with various sparsities of the true coefficients. Since the path algorithms have restrictions on solving the FLR with general design matrices, we only compare the proposed MM algorithms (with and without GPU parallelization, denoted by MMGPU and MM, respectively) with EFLA, SPG, and SB algorithms. In next two scenarios, we consider the two-dimensional FLR with a general design matrix and its application to the image denoising problem. In the two-dimensional FLR with a general design matrix, only MM, SPG, and SB are available. In the image denoising problem, which is equivalent to the two-dimensional FLSA problem, we additionally consider the pathFLSA algorithm. The MM and MMGPU algorithms are implemented in C and CUDA C, and are available in supplementary materials.

The convergence of the algorithm is measured using a relative error of the objective function at each iteration. The relative error at the r th iteration is defined as $\text{RE}(\hat{\beta}^{(r)}) = |f(\hat{\beta}^{(r)}) - f(\hat{\beta}^{(r-1)})|/f(\hat{\beta}^{(r-1)})$, where $\hat{\beta}^{(r)}$ is the estimate in r th iteration (see Algorithms 1 and 2). All the algorithms terminate when the relative error becomes smaller than 10^{-5} or the number of iterations exceeds 10,000. The SB and MM algorithms have additional parameters μ and ϵ , respectively. The details of choosing μ of the SB are given in Appendix B. In the MM and MMGPU, we set the perturbation $\epsilon = 10^{-8}$ to avoid machine precision error. Although setting ϵ a constant does not exactly satisfy the condition for Theorem 1, this choice of constant is sufficiently small within the machine precision and to prevent divide by zero.

The algorithms are compared in the computation of the time and the number of iterations for aforementioned five scenarios. We also investigate the sensitivity of the SB algorithm and the MM algorithm to the choice of their additional parameters μ and ϵ . All algorithms are implemented in MATLAB except for pathFLSA. The computation times are measured in CPU time by using a desktop PC (Intel Core2 extreme X9650 CPU (3.00 GHz) and 8 GB RAM) with NVIDIA GeForce GTX 465 GPU.

5.1 STANDARD FLR

We consider three scenarios for the standard FLR and check the efficiency and the stability of the five algorithms, that is, MMGPU, MM, EFLA, SPG, and SB, with four sets of regularization parameters $(\lambda_1, \lambda_2) \in \{(0.1, 0.1), (0.1, 1), (1, 0.1), (1, 1)\}$. These choices of regularization parameters are similar in Liu et al. (2010), Lin et al. (2011), and Ye and Xie (2011). We generate n samples X_1, X_2, \dots, X_n from a p -dimensional multivariate normal

distribution $N(\mathbf{0}, \mathbf{I}_p)$. The response variable \mathbf{y} is generated from the model

$$\mathbf{y} = \mathbf{X}\beta + \epsilon, \quad (34)$$

where $\epsilon \sim N(\mathbf{0}, \mathbf{I}_n)$ and $\mathbf{X} = (X_1, X_2, \dots, X_n)^T$. To examine the performance with the dimension p , we try $p = 200, 1000, 10,000$, and $20,000$ for sample size $n = 1000$. We generate 10 datasets according to the following scenarios (C1)–(C3).

(C1) Sparse case.

Following Ye and Xie (2011), we set 41 coefficients of β nonzero:

$$\beta_j = \begin{cases} 2 & \text{for } j = 1, 2, \dots, 20, 121, \dots, 125 \\ 3 & \text{for } j = 41 \\ 1 & \text{for } j = 71, \dots, 85 \\ 0 & \text{otherwise} \end{cases}$$

(C2) Moderately sparse case.

Following Lin et al. (2011), we set 30% of the coefficients nonzero:

$$\beta_j = \begin{cases} 1 & \text{for } j = p/10 + 1, p/10 + 2, \dots, 2p/10 \\ 2 & \text{for } j = 2p/10 + 1, 2p/10 + 2, \dots, 4p/10 \\ 0 & \text{otherwise} \end{cases}$$

(C3) Dense case.

We set all the coefficients nonzero:

$$\beta_j = \begin{cases} 1 & \text{for } j = 1, 2, \dots, 5p/10 \\ -1 & \text{for } j = 5p/10 + 1, 5p/10 + 2, \dots, p \end{cases}$$

These true coefficient structures are illustrated in Figure 1.

We report the average computation time and the average number of iterations in Appendix C, Tables C.1–C.3, and summarize the average computation times in Figure 2. Excluding the MMGPU (discussed below), the EFLA is generally the fastest in most of the scenarios considered. For relatively dense case ((C2) and (C3)) with a small penalty ($\lambda_1 = \lambda_2 = 0.1$)







Figure 1. Plots of the true coefficients in (C1)–(C3).




(a) (C1) with $(\lambda_1, \lambda_2) = (0.1, 0.1)$




(b) (C1) with $(\lambda_1, \lambda_2) = (1, 1)$




(c) (C2) with $(\lambda_1, \lambda_2) = (0.1, 0.1)$



(d) (C2) with $(\lambda_1, \lambda_2) = (1, 1)$



(e) (C3) with $(\lambda_1, \lambda_2) = (0.1, 0.1)$



(f) (C3) with $(\lambda_1, \lambda_2) = (1, 1)$

Figure 2. Summary of the results of (C1)–(C3) for $n = 1000$. The vertical lines indicate the range of computation times for 10 datasets. Let p denote the dimension of problem. In (C2) with $(0.1, 0.1)$ at $p = 10,000, 20,000$ and (C3) with $(1, 1)$ at $p = 10,000, 20,000$, EFLA fails to converge the optimal solution at the lower bounds of vertical lines of EFLA (see Appendix C).

for large dimensions ($p = 10,000$ and $20,000$), MM is the fastest. We see that the MM is $1.14 \sim 4.85$ times slower than the EFLA, and is comparable to the SPG and similar to the SB for high dimensions ($p = 10,000$ or $20,000$), each of which is the contender for the second place with the MM in the respective dimensions. Taking into account that the EFLA is only applicable to the standard FLR, we believe that the performance of the MM is acceptable.

The benefit of parallelization is visible in large dimensions. The MMGPU has within about 10% of the computation time of the MM and the fastest among all the algorithms considered when $p = 10,000$ and $20,000$. MMGPU is 2–63 times faster than EFLA, which in turn is the fastest among MM, SPG, and SB. When the dimension is small ($p = 200$), however, MMGPU is the slowest among the five algorithms. This is due to the memory transfer overhead between CPU and GPU. When p is small, memory transfer occupies most of computation time. As the dimension increases, the relative portion of the memory transfer in computation time decreases, and MMGPU becomes efficient.

Focusing on the number of iterations, we see that the MM converges within a few tens of iterations whereas other algorithms require up to a few hundreds of iterations, especially for large p . The number of iterations of MM is also insensitive to the sparsity structure and the choice of λ_1 and λ_2 . (EFLA and SPG need more iterations as the coefficients become dense, together with SB, they also require more iterations for small λ_1 or λ_2 .) This stability in the number of iterations contribute the performance gain in MMGPU. MMGPU speeds up each iteration, hence if the number of iteration does not rely much on the input, we can consistently expect a gain by parallelization.

5.2 TWO-DIMENSIONAL FLR

We consider two scenarios for the two-dimensional FLR, where the coefficients are indexed on $q \times q$ lattice (hence $p = q^2$) and the penalty structure is imposed on the horizontally and vertically adjacent coefficients. The first is the two-dimensional FLR with general design matrix \mathbf{X} . The second is image denoising, which is equivalent to the two-dimensional FLSA ($\mathbf{X} = \mathbf{I}$). Note that EFLA and MMGPU are not applicable for these scenarios. Instead, pathFLSA is considered for the second scenario, since it is known as one of the most efficient algorithms for the two-dimensional FLSA.

For the example of the two-dimensional FLR problem with general \mathbf{X} , we generate 10 datasets according to the following scenario.

(C4) Two-dimensional FLR with general design matrix, moderately sparse case.

We set the $q \times q$ dimensional matrix $B = (b_{ij})$ as follows.

$$b_{ij} = \begin{cases} 2 & \text{for } \frac{q}{4}k + 1 \leq i, j \leq \frac{q}{4}(k+1), k = 0, 1, 2, 3 \\ -2 & \text{for } \frac{q}{4}k + 1 \leq i \leq \frac{q}{4}(k+1), \frac{q}{4}(3-k) + 1 \leq j \leq \frac{q}{4}(4-k) \\ & k = 0, 1, 2, 3 \\ 0 & \text{otherwise} \end{cases}$$

The true coefficient vector β is the vectorization of the matrix B denoted by $\beta = \text{vec}(B)$. The structure of the two-dimensional coefficient matrix B is shown in [Figure 3\(a\)](#).

As the previous scenarios, we generate \mathbf{X} and \mathbf{y} from Equation (34) with true coefficients defined in (C4). The candidate dimensions are $q = 16, 32, 64,$ and 128 for sample size $n = 1000$. We also consider four sets of regularization parameters $(\lambda_1, \lambda_2) \in \{(0.1, 0.1), (0.1, 1), (1, 0.1), (1, 1)\}$ as in Section 5.1.

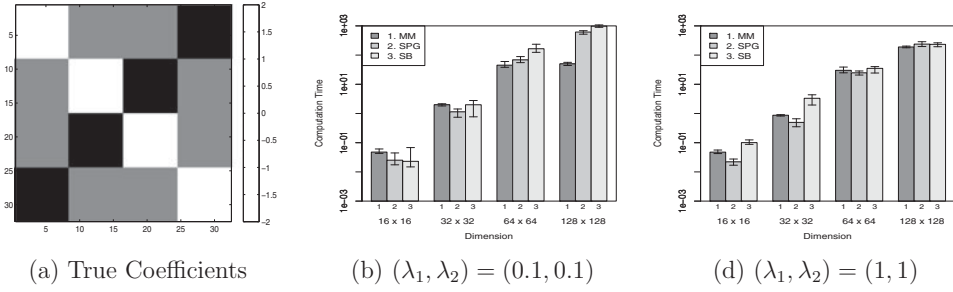


Figure 3. Summary of Case 4 (C4) for sample size $n = 1000$.

The computation times of (C4) are shown in Figures 3(b) and (c). MM gains its competitiveness as the dimension $p (= q^2)$ increases. MM is $1.6 \sim 19.5$ times faster than SB when $p > n$, has similar performance to SPG when $p = 10,000$, and MM is $1.3\text{--}12.3$ times faster than SPG when $p = 20,000$. A detailed report on the computation time can be found in Appendix C, Table C.3. It can be seen that the variation in the number of iterations for both dimensions and values of (λ_1, λ_2) is the smallest for MM. Also, it can be seen that MM tends to converge within a few tens of iterations. This stability is consistent with the observations from (C1)–(C3).

(C5) Image denoising ($\mathbf{X} = \mathbf{I}$)

We use a 256×256 gray scale image (i.e., $q = 256$, $p = q^2 = 65,536$) of R.A. Fisher, as used in Friedman et al. (2007). After standardizing pixel intensities, we add Gaussian noise with a standard variation 0.3. We set $\lambda_1 = 0$ because the application is image denoising.

We consider two regularization parameters $\lambda_2 = 0.1, 1$. We report the average computation times and the average numbers of iterations in Table 2. We present the true image, the noisy image, and the denoising results at $\lambda_2 = 0.1$ in Figure 4. Since the path for all the solutions for two-dimensional FLSA is computationally infeasible, we terminate

Table 2. Summary of the computation time and the numbers of iterations of the MM, SPG, SB, and pathFLSA algorithms for case (C5). N/A denotes the method fails to obtain the optimal solution

(λ_1, λ_2)	Method	Computation time (sec)	No. of iterations*	$f(\hat{\beta})$
(0, 0.1)	MM	21.2082	26	2810.16
	SPG	1.9964	140	2817.31
	SB	217.2167	698	2822.40
	pathFLSA	3.2151	–	2809.96
(0, 1)	MM	48.7812	56	6920.43
	SPG $^\Delta$	N/A	N/A	N/A
	SB	635.8651	4431	7308.07
	pathFLSA	359.2660	–	6919.06

*Since PathFLSA is a path algorithm, the number of iterations of PathFLSA is omitted.

Δ SPG stops after 11 iterations with $f(\hat{\beta}) = 35595.82$.

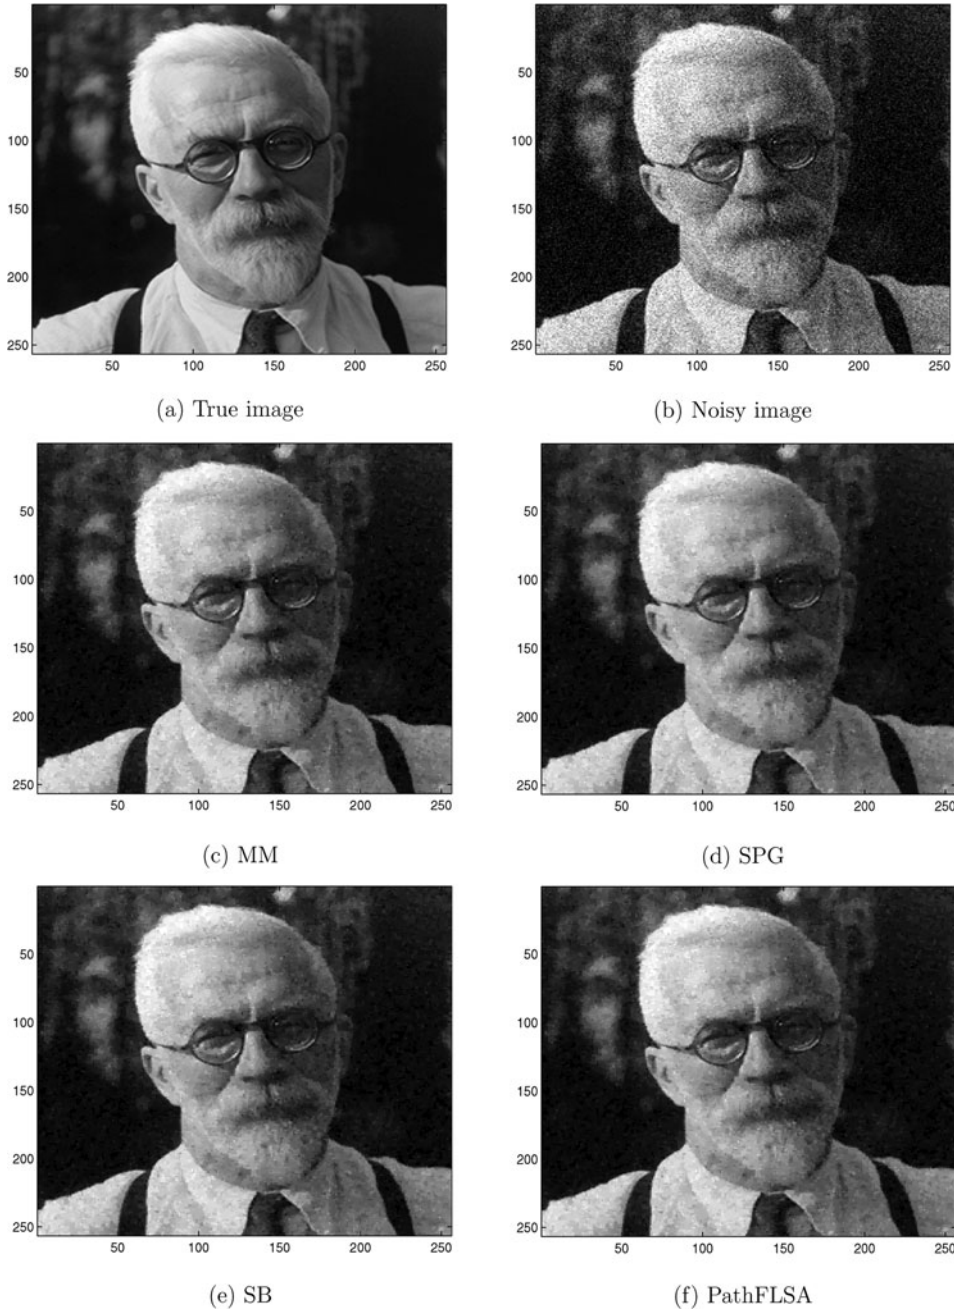


Figure 4. Image denoising with various algorithms for the FLR ($\lambda_2 = 0.1$).

PathFLSA at the desired value of λ_2 . For small λ_2 ($\lambda_2 = 0.1$), MM is slower than SPG and PathFLSA, but much faster than SB. As λ_2 increases, MM becomes faster than PathFLSA since PathFLSA always has to start from $\lambda_2 = 0$. In addition, SPG fails to obtain the solution at $\lambda_2 = 1$ (see Figure 5). MM also exhibits a smaller variation in the number

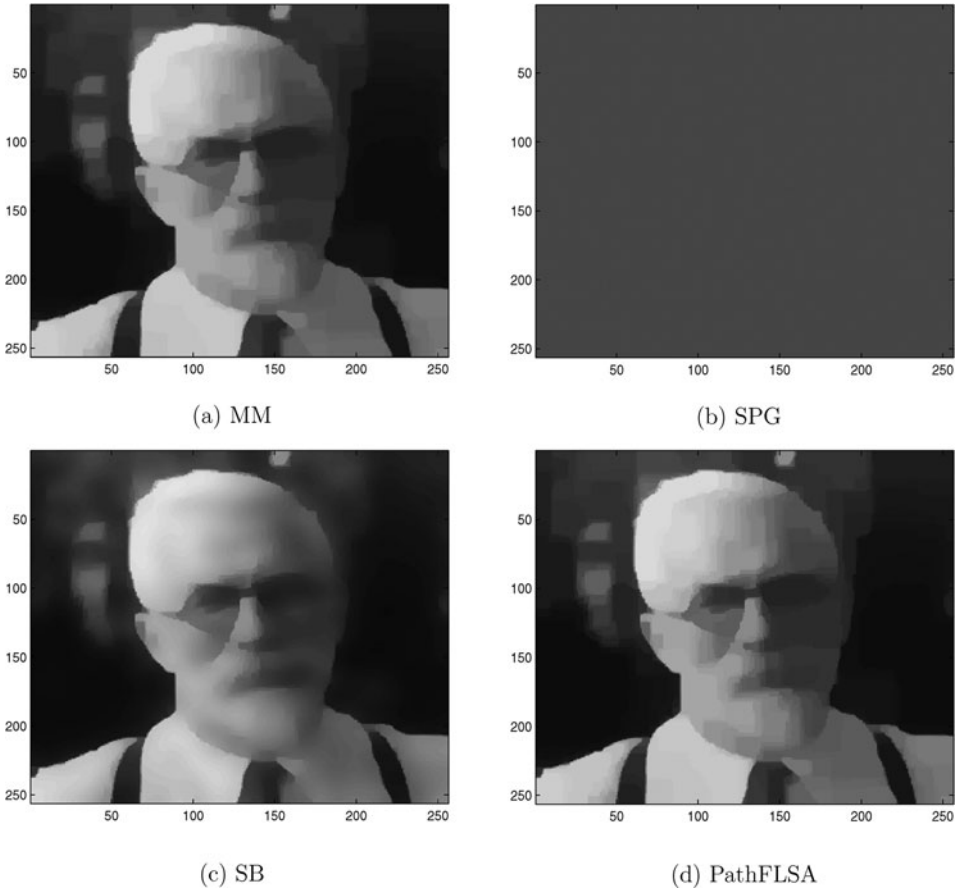


Figure 5. Image denoising with various algorithms for the FLR ($\lambda_2 = 1$). SPG fails to obtain the optimal solution.

of iterations and the objective function value at the obtained solution than SPG and SB. Again, this stability is a consistent behavior of MM throughout (C1)–(C5).

6. CONCLUSION

In this article, we have proposed an MM algorithm for the fused lasso regression (FLR) problem with the generalized fusion penalty. The proposed algorithm is flexible in the sense that it can be applied to the FLR with a wide class of design matrices and penalty structures. It is stable in the sense that the convergence of the algorithm is not sensitive to the dimension of the problem, the choice of the regularization parameters, and the sparsity of the true model. Even when a special structure on the design matrix or the penalty is imposed, the MM algorithm shows comparable performance to the algorithms tailored to the special structure. These features make the proposed algorithm an attractive choice as an off-the-shelf FLR algorithm. Moreover, the performance of the MM algorithm can be improved by parallelizing it with GPU, when the standard fused lasso penalty is imposed

and the dimension increases. Extension of the GPU algorithm to the two-dimensional FLR problem is our future direction of research.

APPENDIX

A. PROOFS

A.1 PROOF OF PROPOSITION 1

For part (i), write

$$f_\epsilon(\beta) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_1 \sum_{j=1}^p a_j(\beta) + \lambda_2 \sum_{(j,k) \in E} b_{jk}(\beta),$$

where

$$a_j(\beta) = |\beta_j| - \epsilon \log \left(1 + \frac{|\beta_j|}{\epsilon} \right)$$

$$b_{jk}(\beta) = |\beta_j - \beta_k| - \epsilon \log \left(1 + \frac{|\beta_j - \beta_k|}{\epsilon} \right).$$

for $j, k = 1, 2, \dots, p$. The functions $a_j(\beta)$ and $b_{jk}(\beta)$ are continuous and convex in $\beta \in \mathbb{R}^p$ and so is $f_\epsilon(\beta)$. In addition, for each $j = 1, 2, \dots, p$, we see that $\lim_{|\beta_j| \rightarrow \infty} f_\epsilon(\beta) = \infty$ held fixed. This, along with the convexity of $f_\epsilon(\beta)$, shows that $\lim_{\|\beta\|_2 \rightarrow \infty} f_\epsilon(\beta) = \infty$.

For part (ii), recall that

$$0 \leq f(\beta) - f_\epsilon(\beta) = \lambda_1 \sum_{j=1}^p \epsilon \log \left(1 + \frac{|\beta_j|}{\epsilon} \right) + \lambda_2 \sum_{(j,k) \in E} \epsilon \log \left(1 + \frac{|\beta_j - \beta_k|}{\epsilon} \right)$$

$$\leq \lambda_1 \sum_{j=1}^p \epsilon \log \left(1 + \frac{\sup_{\beta \in \mathbf{C}} |\beta_j|}{\epsilon} \right) + \lambda_2 \sum_{(j,k) \in E} \epsilon \log \left(1 + \frac{\sup_{\beta \in \mathbf{C}} |\beta_j - \beta_k|}{\epsilon} \right).$$

The suprema in the rightmost side are achieved and finite because \mathbf{C} is compact. Then, the rightmost side monotonically decreases to 0 as ϵ goes to 0.

To see part (iii), for a given $\epsilon > 0$ and $x_0 \in \mathbb{R}$, consider the functions

$$p_\epsilon(x) = |x| - \epsilon \log \left(1 + \frac{|x|}{\epsilon} \right)$$

$$q_\epsilon(x|x^0) = |x^0| - \epsilon \log \left(1 + \frac{|x^0|}{\epsilon} \right) + \frac{x^2 - (x^0)^2}{2(|x^0| + \epsilon)},$$

and their difference $h(y|y^0) = q_\epsilon(x|x^0) - p_\epsilon(x)$, where $y = |x|$ and $y_0 = |x^0|$. The function $h(y|y^0)$ is twice differentiable, and simple algebra shows that it is increasing for $y \geq y_0$ and decreasing for $y < y_0$. We have

$$h(y|y^0) = q_\epsilon(x|x^0) - p_\epsilon(x) \geq h(y^0|y^0) = 0,$$

where the equality holds if and only if $y = y^0$ (equivalently, $|x| = |x^0|$). This shows that $q_\epsilon(x|x^0)$ majorizes $p_\epsilon(x)$ at $x = x^0$, thus $g_\epsilon(\beta|\beta')$ majorizes $f_\epsilon(\beta)$ at $\beta = \beta'$. Finally, the majorizing function

$g_\epsilon(\beta|\beta')$ is strictly convex for fixed β' since

$$\frac{\partial^2 g_\epsilon(\beta|\beta')}{\partial \beta^2} = \mathbf{X}^T \mathbf{X} + \lambda_1 \mathbf{A}' + \lambda_2 \mathbf{B}'$$

is positive definite, where \mathbf{A}' and \mathbf{B}' are defined in (33). Hence, it has a unique global minimum.

A.2 PROOF OF LEMMA 1

Define the level set

$$\Omega(\epsilon, c) = \{\beta \in \mathbb{R}^p | f_\epsilon(\beta) \leq c\},$$

for every $\epsilon > 0$ and $c > 0$. Similarly, define

$$\Omega(0, c) = \{\beta \in \mathbb{R}^p | f(\beta) \leq c\}.$$

Then, from part (i) of Proposition 1, $\Omega(\epsilon, c)$ is compact for every $\epsilon > 0$ and $c > 0$. Compactness of $\Omega(0, c)$ follows from the uniform convergence of $f_\epsilon(\beta)$ to $f(\beta)$.

Note that $\Omega(\epsilon, f(\widehat{\beta}))$ is a nonempty compact set since by the assumption that the solution set $\{\beta | f(\beta) = f(\widehat{\beta})\}$ is nonempty. Then, we have $f_\epsilon(\beta) \leq f(\widehat{\beta})$ for every $\beta \in \Omega(\epsilon, f(\widehat{\beta}))$. By construction, $f_\epsilon(\widehat{\beta}_\epsilon) \leq f_\epsilon(\beta)$. Thus,

$$\min_{\gamma} f_\epsilon(\gamma) = f_\epsilon(\widehat{\beta}_\epsilon) \leq f_\epsilon(\beta) \leq f(\widehat{\beta}) = \min_{\gamma} f(\gamma).$$

In other words, $\widehat{\beta}_\epsilon \in \Omega(\epsilon, f(\widehat{\beta}))$. Similarly, we see $\widehat{\beta} \in \Omega(0, f(\widehat{\beta}))$ that is nonempty.

Since $f(\widehat{\beta}_\epsilon) \geq f(\widehat{\beta})$ by the definition of $\widehat{\beta}$,

$$\begin{aligned} 0 \leq f(\widehat{\beta}_\epsilon) - f(\widehat{\beta}) &\leq f(\widehat{\beta}_\epsilon) - f_\epsilon(\widehat{\beta}_\epsilon) + f_\epsilon(\widehat{\beta}) - f(\widehat{\beta}) \\ &\leq |f(\widehat{\beta}_\epsilon) - f_\epsilon(\widehat{\beta}_\epsilon)| + |f_\epsilon(\widehat{\beta}) - f(\widehat{\beta})|. \end{aligned}$$

The rightmost side of the above inequality goes to zero because $f_\epsilon(\beta)$ converges to $f(\beta)$ uniformly on both $\Omega(\epsilon, f(\widehat{\beta}))$ and $\Omega(0, f(\widehat{\beta}))$. Then, for a limit point β^* of the sequence $\{\widehat{\beta}_{\epsilon_n}\}_{n \geq 1}$ with $\epsilon_n \downarrow 0$, we see

$$\lim_{n \rightarrow \infty} f(\widehat{\beta}_{\epsilon_n}) = f(\beta^*) = f(\widehat{\beta}) = \min_{\gamma} f(\gamma)$$

by the continuity of $f(\beta)$, that is, β^* minimizes $f(\beta)$.

A.3 PROOF OF LEMMA 2

We first claim that the recursive updating rule implicitly defined by the MM algorithm is continuous.

Claim 1. The function $M(\alpha) = \arg \min_{\beta} g_\epsilon(\beta|\alpha)$ is continuous in α , where $g_\epsilon(\beta|\alpha)$ is the proposed majorizing function of $f_\epsilon(\beta)$ at α .

Proof. Consider a sequence $\{\alpha_k\}_{k \geq 0}$ converging to $\tilde{\alpha}$ with finite $f_\epsilon(\tilde{\alpha})$ and satisfying $f_\epsilon(\alpha_k) \leq f_\epsilon(\alpha_0) < \infty$ for $k \geq 1$. We want to show that $M(\alpha_k)$ converges to $M(\tilde{\alpha})$.

The descent property of the MM algorithm states that

$$f_\epsilon(M(\alpha_k)) \leq g_\epsilon(M(\alpha_k)|\alpha_k) \leq g_\epsilon(\alpha_k|\alpha_k) = f_\epsilon(\alpha_k)$$

for every k . Thus, the sequence $\{M(\alpha_k), k = 0, 1, 2, \dots\}$ is a subset of the level set $\Omega(\epsilon, f_\epsilon(\alpha_0)) = \{\beta \in \mathbb{R}^p | f_\epsilon(\beta) \leq f_\epsilon(\alpha_0)\}$.

Since the level set $\Omega(\epsilon, f_\epsilon(\alpha_0))$ is compact (see Section A.2), the sequence $\{M(\alpha_k)\}_{k=0}$ is bounded. Hence, we can find a convergent subsequence $\{M(\alpha_{k_n})\}_{n \geq 1}$; denote its limit point by \tilde{M} . Then, by the continuity of $g_\epsilon(\beta|\alpha)$ in both β and α , for all β ,

$$g_\epsilon(\tilde{M}|\tilde{\alpha}) = \overline{\lim}_{n \rightarrow \infty} g_\epsilon(M(\alpha_{k_n})|\alpha_{k_n}) \leq \overline{\lim}_{n \rightarrow \infty} g_\epsilon(\beta|\alpha_{k_n}) = g_\epsilon(\beta|\tilde{\alpha}).$$

Thus \tilde{M} minimizes $g_\epsilon(\beta|\tilde{\alpha})$. Since the minimizer of $g_\epsilon(\beta|\alpha)$ is unique by Proposition 1, we have $\tilde{M} = M(\tilde{\alpha})$, that is, $M(\alpha_{k_n})$ converges to $M(\tilde{\alpha})$. \square

We then state a result on the global convergence of the updating rule.

Claim 2. (Convergence Theorem A, Zangwill 1969, p. 91) Let the point-to-set map $M : X \rightarrow X$ determine an algorithm that given a point $x_1 \in X$ generates the sequence $\{x_k\}_{k=1}^\infty$. Also let a solution set $\Gamma \subset X$ be given. Suppose that: (i) all points x_k are in a compact set $C \subset X$; (ii) there is a continuous function $u : X \rightarrow \mathbb{R}$ such that (a) if $x \notin \Gamma$, $u(y) > u(x)$ for all $y \in M(x)$, and (b) if $x \in \Gamma$, then either the algorithm terminates or $u(y) \geq u(x)$ for all $y \in M(x)$; (iii) the map M is closed at x if $x \notin \Gamma$. Then either the algorithm stops at a solution, or the limit of any convergent subsequence is a solution.

Closedness of point-to-set maps is an extension of continuity of point-to-point maps: a point-to-set map $A : V \rightarrow V$ on a set V is said to be closed if (a) $z^k \rightarrow z^\infty$, (b) $y^k \in A(z^k)$, and (c) $y^k \rightarrow y^\infty$ imply $y^\infty \in A(z^\infty)$. The map A is closed on a set $X \subset V$ if it is closed at each $z \in X$ (Zangwill 1969, p. 88). The proof of this claim is similar to Vaida (2005).

Finally we are ready to show the main result of this section.

Proof of Lemma 2. We apply Claim 2 with $M(\alpha) = \arg \min_\beta g_\epsilon(\beta|\alpha)$, $\Gamma = \{\beta \mid 0 \in \partial_\beta f_\epsilon(\beta)\}$, and $u(\beta) = -f_\epsilon(\beta)$, where $\partial_\beta f_\epsilon(\beta)$ is the subdifferential of $f_\epsilon(\beta)$ at β . We check the three conditions for Claim 2 as follows. Define the level set of $f_\epsilon(\beta)$ as $\Omega(\epsilon, f_\epsilon(\hat{\beta}^{(0)})) = \{\beta \in \mathbb{R}^p \mid f_\epsilon(\beta) \leq f_\epsilon(\hat{\beta}^{(0)})\}$, which is compact (see Section A.2). Since $f_\epsilon(\hat{\beta}^{(r+1)}) \leq f_\epsilon(\hat{\beta}^{(r)})$ for every $r = 0, 1, 2, \dots$, $\{\hat{\beta}^{(r)}\}_{r \geq 0}$ is a subset of $\Omega(\epsilon, f_\epsilon(\hat{\beta}^{(0)}))$ (see Section A.2) and condition (i) of Claim 2 is satisfied. Lemma 1 states that $M(\alpha)$ is a continuous point-to-point map in α . That is, $M(\alpha)$ is a closed map on any compact set of \mathbb{R}^p , and so on Γ^c . Thus, condition (iii) of Claim 2 is met. Finally, condition (ii) is guaranteed by the strict convexity of $g_\epsilon(\beta|\alpha)$ (see Proposition 1) and the definition of $M(\alpha)$. Therefore, all the limit points of $\{\hat{\beta}^{(r)}\}_{r \geq 0}$ are stationary points of $f_\epsilon(\beta)$. \square

B. CHOICE OF μ IN THE SB ALGORITHM

In the SB algorithm, the current solution $\hat{\beta}^{(r+1)}$ is obtained from the solution of the following linear system:

$$(\mathbf{X}^T \mathbf{X} + \mu \mathbf{I} + \mu \mathbf{D}^T \mathbf{D})\boldsymbol{\beta} = \mathbf{c}^{(r)}(\mu).$$

The additional parameter μ affects the convergence rate but there is no optimal rule for its choice in this problem (Ghadimi et al. 2012). Thus, Ye and Xie (2011) suggested a pretrial procedure to choose μ as one of $\{0.2, 0.4, 0.6, 0.8, 1\} \times \|\mathbf{y}\|_2$ by testing computation times for given (λ_1, λ_2) . In




Figure A.1. The number of iterations in SB algorithm as μ changes for $(\lambda_1, \lambda_2) = (0.1, 0.1)$

In addition to this recipe, we consider $\frac{1}{n} \|\mathbf{y}\|_2$ and $\frac{1}{n^2} \|\mathbf{y}\|_2$ as candidate values of μ . In our design of examples, we observe that the values other than $\frac{1}{n} \|\mathbf{y}\|_2$ makes the SB algorithm prematurely stop before reaching the optimal solution. The value $\frac{1}{n} \|\mathbf{y}\|_2$ also leads the best convergence rate from our pretrials in Figure A.1. Thus, we set $\mu = \frac{1}{n} \|\mathbf{y}\|_2$ in the numerical studies.

C. TABLES OF RESULTS FOR (C1)–(C4)

Appendix C.1. Summary of (C1) with computation times and the numbers of iterations for $n = 1000$

λ_1	λ_2	p	Computation time (sec)					Number of iterations				
			MMGPU	MM	EFLA	SPG	SB	MMGPU	MM	EFLA	SPG	SB
0.1	0.1	200	0.2858	0.0090	0.0093	0.0152	0.0089	3.0	3.0	27.8	42.5	3.0
		1000	2.6552	1.5070	0.3085	0.5799	2.1276	22.6	23.0	258.9	405.3	15.1
		10000	7.2544	66.9573	32.8320	102.5558	99.1599	84.1	83.9	1067.8	1756.1	158.2
		20000	13.3395	133.2365	77.4856	281.4834	209.6247	90.3	90.3	1287.3	2437.4	206.6
0.1	1.0	200	0.2207	0.0092	0.0067	0.0273	0.0656	4.0	4.0	24.3	73.8	28.7
		1000	2.2735	1.1956	0.1987	0.4140	10.0084	26.3	26.4	166.9	285.8	75.3
		10000	5.3980	51.1480	19.3404	58.0435	117.2644	86.3	86.1	618.7	988.9	186.9
		20000	9.8197	101.1593	49.4796	163.0641	238.4468	94.2	93.7	808.7	1411.3	234.9
1.0	0.1	200	0.1469	0.0074	0.0069	0.0168	0.0150	3.2	3.2	26.1	43.0	5.8
		1000	1.6121	0.8684	0.2128	0.4359	9.2411	23.2	23.0	166.5	266.7	69.7
		10000	7.4584	75.5762	15.5698	48.1091	54.5518	82.9	82.9	506.3	822.1	87.7
		20000	14.9428	163.5872	44.3698	133.7532	152.5818	92.8	92.8	737.6	1155.5	152.4
1.0	1.0	200	0.1477	0.0072	0.0068	0.0157	0.0873	4.0	4.0	25.0	42.2	37.8
		1000	1.2295	0.6092	0.1527	0.2986	10.1666	24.4	24.4	128.1	209.6	79.2
		10000	4.5901	44.8417	14.6557	32.6244	67.7480	82.5	82.4	465.8	554.9	110.9
		20000	8.4576	90.2439	35.3691	96.1971	199.6160	88.6	88.5	575.4	830.1	201.5

Appendix C.2. Summary of (C2) with computation times and the numbers of iterations for $n = 1000$. The gray cells denote that the EFLA fails to reach the optimal solution for one or two samples. We report the average computation times of the EFLA removed the failed cases

λ_1	λ_2	p	Computation time (sec.)					Number of iterations				
			MMGPU	MM	EFLA	SPG	SB	MMGPU	MM	EFLA	SPG	SB
0.1	0.1	200	0.2247	0.0108	0.0077	0.0170	0.0088	3.0	3.0	25.0	46.8	3.0
		1000	3.0984	1.4673	0.3732	0.7777	1.6792	22.2	21.8	311.8	508.8	19.9
		10000	4.1122	45.8289	91.7927	263.7253	266.5620	36.3	35.5	2976.3	4527.5	466.7
		20000	4.4272	52.7964	242.0749	821.5516	704.5284	27.2	27.6	4043.3	7108.9	776.4
0.1	1.0	200	0.2276	0.0115	0.0072	0.0168	0.0595	4.2	4.2	25.4	41.9	26.8
		1000	3.6282	1.5755	0.2211	0.5595	4.7977	28.1	28.2	180.5	359.1	57.7
		10000	10.3537	109.7741	47.8311	132.3191	107.4803	53.9	53.9	1526.8	2272.4	188.4
		20000	23.6943	271.3948	236.7397	503.1258	232.9924	63.9	63.7	3915.2	4347.2	256.3
1.0	0.1	200	0.1279	0.0086	0.0061	0.0171	0.0081	3.1	3.1	21.8	44.6	2.7
		1000	1.7817	0.9148	0.1901	0.5022	2.9606	20.3	20.3	160.7	325.7	34.9
		10000	10.8871	127.3424	51.5586	169.2509	142.5682	91.6	91.5	1665.9	2910.1	249.8
		20000	20.1332	244.0547	141.8118	526.4848	313.0687	88.7	89.7	2383.4	4556.9	344.7
1.0	1.0	200	0.1467	0.0084	0.0065	0.0176	0.0633	4.2	4.2	23.4	44.6	28.8
		1000	1.7161	0.8342	0.1476	0.4186	3.0164	22.5	22.5	122.6	269.4	36.0
		10000	7.7781	91.8040	38.7895	120.4112	64.9480	92.1	92.2	1226.0	2077.8	114.0
		20000	17.6899	212.5753	106.2446	339.2643	174.6900	102.3	102.3	1756.6	2929.0	192.6

Appendix C.3. Summary of (C3) with computation times and the numbers of iterations for $n = 1000$. The gray cells denote that the EFLA fails to reach the optimal solution for one sample. We report the average computation times of the EFLA removed failed case

λ_1	λ_2	p	Computation time (sec)					Number of iterations				
			MMGPU	MM	EFLA	SPG	SB	MMGPU	MM	EFLA	SPG	SB
0.1	0.1	200	0.2511	0.0171	0.0081	0.0170	0.0095	3.0	3.0	26.7	45.2	3.0
		1000	2.9360	1.5933	0.3793	0.8082	1.2007	21.3	21.9	327.4	511.3	13.2
		10000	3.7909	42.9165	90.8098	282.7641	286.2912	34.7	34.0	3025.2	4789.1	498.3
		20000	4.0806	48.1163	248.1919	871.9119	756.4445	26.0	25.8	4108.3	7553.7	824.2
0.1	1.0	200	0.3401	0.0133	0.0069	0.0174	0.0662	4.2	4.2	25.3	44.9	27.4
		1000	3.9890	1.8233	0.2031	0.5462	3.8157	27.4	27.3	170.8	352.3	43.5
		10000	15.0492	180.0133	54.7479	144.0007	108.6189	91.8	91.9	1795.4	2446.6	189.3
		20000	28.2846	343.3138	169.7143	505.0450	238.1494	96.8	97.0	2762.1	4376.6	259.7
1.0	0.1	200	0.0890	0.0058	0.0072	0.0178	0.0130	3.4	3.4	26.5	46.3	5.0
		1000	1.6862	0.8142	0.3392	0.8342	2.3927	23.8	23.8	287.1	555.2	24.8
		10000	10.7328	127.3884	50.5020	176.1017	127.5324	88.6	88.6	1681.8	2985.3	222.3
		20000	19.5732	239.8993	150.7861	520.9941	307.9786	85.4	85.8	2497.4	4511.6	335.6
1.0	1.0	200	0.1840	0.0085	0.0067	0.0169	0.0094	4.4	4.4	24.9	44.3	3.4
		1000	1.4819	0.7155	0.1548	0.4869	1.9491	19.7	19.6	124.3	308.7	23.0
		10000	8.1713	97.7493	38.0236	129.9449	89.1512	95.4	95.5	1240.0	2206.9	155.5
		20000	17.7725	219.5605	110.7899	365.3513	173.9217	101.9	101.9	1799.7	3156.2	189.4

Appendix C.4. Summary of (C4) with computation times and the numbers of iterations for $n = 1000$

λ_1	λ_2	$q \times q$	Computation time (sec)			Number of iterations		
			MM	SPG	SB	MM	SPG	SB
0.1	0.1	16×16	0.0477	0.0252	0.0230	5.0	54.9	2.8
		32×32	1.9948	1.1403	1.9887	21.3	603.6	16.5
		64×64	45.1773	68.6246	164.9881	63.7	2745.9	456.0
		128×128	50.2078	615.4028	979.8359	28.8	6414.6	1066.6
0.1	1.0	16×16	0.0549	0.0234	0.0971	8.3	53.3	26.7
		32×32	1.1902	0.5846	4.1854	23.0	294.8	37.9
		64×64	15.1176	23.0897	33.2731	39.6	923.6	90.8
		128×128	72.8389	208.2940	119.7563	49.8	2169.3	123.5
1.0	0.1	16×16	0.0612	0.0232	0.0180	7.3	56.2	3.0
		32×32	1.4630	0.6096	2.3580	20.6	326.3	21.5
		64×64	48.5678	46.7797	92.2762	80.1	1876.4	254.2
		128×128	241.6816	343.2788	657.3494	85.3	3579.7	714.9
1.0	1.0	16×16	0.0476	0.0221	0.1006	9.5	51.5	27.3
		32×32	0.8718	0.4960	3.3169	20.5	273.1	30.3
		64×64	30.0718	24.6371	35.1351	63.5	978.6	95.9
		128×128	189.6904	237.8033	233.3495	109.0	2476.8	248.1

SUPPLEMENTARY MATERIALS

Source codes and datasets: Supplementary materials contain C, CUDA C, and MATLAB source codes and datasets used in the numerical studies. All source codes are run on the 64-bit Windows operating system. To run CUDA C codes, a graphics device supporting CUDA (e.g., NVIDIA's GeForce GTX 465) is required. Detailed descriptions of files are given in `Readme.txt`, that is also enclosed in the supplementary materials. (`FusedGPU_supp.tar.zip`, GNU zipped tar file)

ACKNOWLEDGMENTS

We would like to thank Dr. Gui-Bo Ye for providing the MATLAB code for the Split Bregman algorithm. Joong-Ho Won's research was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIP, No. 2013-R1A1A1057949), and in part and by Research Resettlement Fund for the new faculty of Seoul National University. Johan Lim's research was supported by the NRF grant funded by the Korean government (MSIP, No. 2011-0030810). Sungroh Yoon's research was supported in part by the NRF grant funded by the Korean government (MSIP, No. 2012-R1A2A4A01008475) and in part by Samsung Electronics Co., Ltd.

[Received June 2013. Revised December 2013.]

REFERENCES

- Anderson, E., Bai, Z., Bischof, C., Demmel, J., Dongarra, J., Croz, J. D., Greenbaum, A., Hammarling, S., McKenney, A., Ostrouchov, S., and Sorensen, D. (1995), *LAPACK Users' Guide* (2nd ed.), Philadelphia, PA: SIAM. [134]
- Beck, A., and Teboulle, M. (2009), "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM Journal on Imaging Sciences*, 2, 183–202. [129]

- Chen, X., Lin, Q., Kim, S., Carbonell, J. G., and Xing, E. P. (2012), “Smoothing Proximal Gradient Method for General Structured Sparse Regression,” *Annals of Applied Statistics*, 6, 719–752. [122,128,129]
- Demmel, J. W. (1997), *Applied Numerical Linear Algebra*, Philadelphia, PA: SIAM. [130,135]
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least Angle Regression,” *The Annals of Statistics*, 32, 407–499. [122]
- Farber, R. (2011), *CUDA Application Design and Development*, San Francisco, CA: Morgan Kaufmann Publishers, Inc. [138]
- Ford, L. R., Jr, and Fulkerson, D. R. (1956), “Maximal Flow Through a Network,” *Canadian Journal of Mathematics*, 8, 399–404. [125]
- Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), “Pathwise Coordinate Optimization,” *The Annals of Applied Statistics*, 1, 302–332. [122,123,124,143]
- Ghadimi, E., Teixeira, A., Shames, I., and Johansson, M. (2012), “On the Optimal Step-Size Selection for the Alternating Direction Method of Multipliers,” in *Proceedings IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys)*. [131,148]
- Gill, P. E., Murray, W., and Saunders, M. A. (1997), “User’s Guide for Snopt 5.3: A Fortran Package for Large-Scale Nonlinear Programming,” Technical Report NA 97-4. University of California, San Diego, CA. [122]
- Grant, M., Boyd, S., and Ye, Y. (2008), “CVX: Matlab Software for Disciplined Convex Programming,” Web page and software. Available at <http://stanford.edu/boyd/cvx> [122]
- Hestenes, M. R. (1969), “Multiplier and Gradient Methods,” *Journal Optimization Theory & Applications*, 4, 303–320. [130]
- Hockney, R. (1965), “A Fast Direct Solution of Poisson’s Equation Using Fourier Analysis,” *Journal of the ACM*, 12, 95–113. [138]
- Hockney, R., and Jesshope, C. R. (1981), *Parallel Computers*, Bristol: Adam Hilger. [138]
- Hoeffling, H. (2010), “A Path Algorithm for the Fused Lasso Signal Approximator,” *Journal of Computational and Graphical Statistics*, 19, 984–1006. [122,124,128]
- Hunter, D. R., and Li, R. (2005), “Variable Selection Using mm Algorithm,” *The Annals of Statistics*, 33, 1617–1642. [122,133]
- Kirk, D. B., and Hwu, W. W. (2010), *Programming Massively Parallel Processors: A Hands-on Approach*, San Francisco, CA: Morgan Kaufmann Publishers, Inc. [138]
- Lange, K., Hunter, D. R., and Yang, I. (2000), “Optimization Transfer Using Surrogate Objective Functions,” *Journal of Computational and Graphical Statistics*, 9, 1–20. [122]
- Lin, X., Pham, M., and Ruzsyczynski, A. (2011), “Alternating Linearization for Structured Regularization Problems,” Preprint. Available at <http://arxiv.org/abs/1201.0306> [122,131,139,140]
- Liu, J., Yuan, L., and Ye, J. (2010), “An Efficient Algorithm for a Class of Fused Lasso Problems,” in *The ACM SIG Knowledge Discovery and Data Mining*, Washington, DC: ACM. [122,127,128,139]
- Nesterov, Y. (2004), *Introductory Lectures on Convex Optimization: A Basic Course*, Boston, MA: Kluwer Academic Publishers. [128]
- (2005), “Smooth Minimization of Non-Smooth Functions,” *Mathematical Programming*, 103, 127–152. [129]
- (2007), “Gradient Methods for Minimizing Composite Objective Function,” CORE Discussion Paper. [128]
- Owens, J. D., Houston, M., Luebke, D., Green, S., Stone, J. E., and Phillips, J. C. (2008), “GPU Computing,” *Proceedings of the IEEE*, 96, 879–899. [138]
- Petry, S., Flexeder, C., and Tutz, G. (2011), “Pairwise Fused Lasso,” Technical Reports, No. 102, University of Munich. Available at http://epub.ub.uni-muenchen.de/12164/1/petry_etal_TR102_2011.pdf [122]
- Rockafellar, R. T. (1973), “A Dual Approach to Solving Nonlinear Programming Problems by Unconstrained Optimization,” *Mathematical Programming*, 5, 354–373. [130]

- She, Y. (2010), "Sparse Regression With Exact Clustering," *Electronic Journal of Statistics*, 4, 1055–1096. [122]
- Stone, H. S. (1973), "An Efficient Parallel Algorithm for the Solution of a Tridiagonal Linear System of Equations," *Journal of the ACM*, 20, 27–38. [138]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection Via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [122]
- Tibshirani, R. J., and Taylor, J. (2011), "The Solution Path of the Generalized Lasso," *The Annals of Statistics*, 39, 1335–1371. [122,126,127,128]
- Tibshirani, R. J., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005), "Sparsity and Smoothness Via the Fused Lasso," *Journal of the Royal Statistical Society, Series B*, 67, 91–108. [122]
- Tseng, P. (2001), "Convergence of a Block Coordinate Descent Method for Nondifferentiable Minimization," *Journal of Optimization Theory and Applications*, 109, 475–494. [123]
- Vaida, F. (2005), "Parameter Convergence for EM and MM Algorithms," *Statistica Sinica*, 15, 831–840. [133,148]
- Ye, G. B., and Xie, X. (2011), "Split Bregman Method for Large Scale Fused Lasso," *Computational Statistics & Data Analysis*, 55, 1552–1569. [122,129,130,139,140,148]
- Zangwill, W. I. (1969), *Nonlinear Programming: A Unified Approach*, Englewood Cliffs, NJ: Prentice Hall. [148]
- Zhang, Y., Cohen, J., and Owens, J. D. (2010), "Fast Tridiagonal Solvers on the GPU," *SIGPLAN Not.*, 45, 127–136. [138]
- Zhou, H., Lange, K., and Suchard, M. A. (2010), "Graphics Processing Units and High-Dimensional Optimization," *Statistical Science*, 25, 311–324. [138]