# PATHWISE COORDINATE OPTIMIZATION 

By Jerome Friedman, ${ }^{1}$ Trevor Hastie, ${ }^{2}$ Holger Höfling ${ }^{3}$ and Robert Tibshirani ${ }^{4}$<br>Stanford University


#### Abstract

We consider "one-at-a-time" coordinate-wise descent algorithms for a class of convex optimization problems. An algorithm of this kind has been proposed for the $L_{1}$-penalized regression (lasso) in the literature, but it seems to have been largely ignored. Indeed, it seems that coordinate-wise algorithms are not often used in convex optimization. We show that this algorithm is very competitive with the well-known LARS (or homotopy) procedure in large lasso problems, and that it can be applied to related methods such as the garotte and elastic net. It turns out that coordinate-wise descent does not work in the "fused lasso," however, so we derive a generalized algorithm that yields the solution in much less time that a standard convex optimizer. Finally, we generalize the procedure to the two-dimensional fused lasso, and demonstrate its performance on some image smoothing problems.


1. Introduction. In this paper we consider statistical models that lead to convex optimization problems with inequality constraints. Typically, the optimization for these problems is carried out using a standard quadratic programming algorithm. The purpose of this paper is to explore "one-at-a-time" coordinate-wise descent algorithms for these problems. The equivalent of a coordinate descent algorithm has been proposed for the $L_{1}$-penalized regression (lasso) in the literature, but it is not commonly used. Moreover, coordinate-wise algorithms seem too simple, and they are not often used in convex optimization, perhaps because they only work in specialized problems. We ourselves never appreciated the value of coordinate descent methods for convex statistical problems before working on this paper.

In this paper we show that coordinate descent is very competitive with the wellknown LARS (or homotopy) procedure in large lasso problems, can deliver a path of solutions efficiently, and can be applied to many other convex statistical problems such as the garotte and elastic net. We then go on to explore a nonseparable problem in which coordinate-wise descent does not work-the "fused lasso." We derive a generalized algorithm that yields the solution in much less time that a standard convex optimizer. Finally, we generalize the procedure to

[^0]the two-dimensional fused lasso, and demonstrate its performance on some image smoothing problems.

A key point here: coordinate descent works so well in the class of problems that we consider because each coordinate minimization can be done quickly, and the relevant equations can be updated as we cycle through the variables. Furthermore, often the minimizers for many of the parameters don't change as we cycle through the variables, and hence, the iterations are very fast.

Consider, for example, the lasso for regularized regression [Tibshirani (1996)]. We have predictors $x_{i j}, j=1,2, \ldots, p$, and outcome values $y_{i}$ for the $i$ th observation, for $i=1,2, \ldots, n$. Assume that the $x_{i j}$ are standardized so that $\sum_{i} x_{i j} / n=0, \sum_{i} x_{i j}^{2}=1$. The lasso solves

$$
\begin{align*}
\min _{\beta} & \frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right)^{2} \\
\text { subject to } & \sum_{j=1}^{p}\left|\beta_{j}\right| \leq s . \tag{1}
\end{align*}
$$

The bound $s$ is a user-specified parameter, often chosen by a model selection procedure such as cross-validation. Equivalently, the solution to (1) also minimizes the "Lagrange" version of the problem

$$
\begin{equation*}
f(\beta)=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right)^{2}+\gamma \sum_{j=1}^{p}\left|\beta_{j}\right| \tag{2}
\end{equation*}
$$

where $\gamma \geq 0$. There is a one-to-one correspondence between $\gamma$ and the bound $s-$ if $\hat{\beta}(\gamma)$ minimizes (2), then it also solves (1) with $s=\sum_{j=1}^{p}\left|\hat{\beta}_{j}(\gamma)\right|$. In the signal processing literature, the lasso and $L_{1}$ penalization is known as "basis pursuit" [Chen et al. (1998)].

There are efficient algorithms for solving this problem for all values of $s$ or $\gamma$; see Efron et al. (2004), and the homotopy algorithm of [Osborne et al. (2000)]. There is another, simpler algorithm for solving this problem for a fixed value $\gamma$. It relies on solving a sequence of single-parameter problems, which are assumed to be simple to solve.

With a single predictor, the lasso solution is very simple, and is a softthresholded version [Donoho and Johnstone (1995)] of the least squares estimate $\hat{\beta}$ :

$$
\begin{align*}
\hat{\beta}^{\text {lasso }}(\gamma) & =S(\hat{\beta}, \gamma) \equiv \operatorname{sign}(\hat{\beta})(|\hat{\beta}|-\gamma)_{+}  \tag{3}\\
& = \begin{cases}\hat{\beta}-\gamma, & \text { if } \hat{\beta}>0 \text { and } \gamma<|\hat{\beta}|, \\
\hat{\beta}+\gamma, & \text { if } \hat{\beta}<0 \text { and } \gamma<|\hat{\beta}|, \\
0, & \text { if } \gamma \geq|\hat{\beta}| .\end{cases} \tag{4}
\end{align*}
$$



Fig. 1. The lasso problem with a single standardized predictor leads to soft thresholding. In each case the solid vertical line indicates the lasso estimate, and the broken line the least-squares estimate.

This simple expression arises because the convex-optimization problem (2) reduces to a few special cases when there is a single predictor. Minimizing the criterion (2) with a single standardized $x$ and $\beta$ simplifies to the equivalent problem

$$
\begin{equation*}
\min _{\beta} \frac{1}{2}(\beta-\hat{\beta})^{2}+\gamma|\beta| \tag{5}
\end{equation*}
$$

where $\hat{\beta}=\sum_{i} x_{i} y_{i}$ is the simple least-squares coefficient. If $\beta>0$, we can differentiate (5) to get

$$
\begin{equation*}
\frac{d f}{d \beta}=\beta-\hat{\beta}+\gamma=0 \tag{6}
\end{equation*}
$$

This leads to the solution $\beta=\hat{\beta}-\gamma$ (left panel of Figure 1) as long as $\hat{\beta}>0$ and $\gamma<\hat{\beta}$, otherwise 0 is the minimizing solution (right panel). Similarly, if $\hat{\beta}<0$, if $\gamma<-\hat{\beta}$, then the solution is $\beta=\hat{\beta}+\gamma$, else 0 .

With multiple predictors that are uncorrelated, it is easily seen that once again the lasso solutions are soft-thresholded versions of the individual least squares estimates. This is not the case for general (correlated) predictors. Consider instead a simple iterative algorithm that applies soft-thresholding with a "partial residual" as a response variable. We write (2) as

$$
\begin{equation*}
f(\tilde{\beta})=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{k \neq j} x_{i k} \tilde{\beta}_{k}-x_{i j} \beta_{j}\right)^{2}+\gamma \sum_{k \neq j}\left|\tilde{\beta}_{j}\right|+\gamma\left|\beta_{j}\right| \tag{7}
\end{equation*}
$$

where all the values of $\beta_{k}$ for $k \neq j$ are held fixed at values $\tilde{\beta}_{k}(\gamma)$. Minimizing w.r.t. $\beta_{j}$, we get

$$
\begin{equation*}
\tilde{\beta}_{j}(\gamma) \leftarrow S\left(\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}^{(j)}\right), \gamma\right) \tag{8}
\end{equation*}
$$



FIG. 2. Diabetes data: iterates for each coefficient from algorithm (9). The algorithm converges to the lasso estimates shown on the right side of the plot.
where $\tilde{y}_{i}^{(j)}=\sum_{k \neq j} x_{i k} \tilde{\beta}_{k}(\gamma)$. This is simply the univariate regression coefficient of the partial residual $y_{i}-\tilde{y}_{i}^{(j)}$ on the (unit $L_{2}$-norm) $j$ th variable; hence, this has the same form as the univariate version (3) above. The update (7) is repeated for $j=1,2, \ldots, p, 1,2, \ldots$ until convergence.

An equivalent form of this update is

$$
\begin{equation*}
\tilde{\beta}_{j}(\gamma) \leftarrow S\left(\tilde{\beta}_{j}(\gamma)+\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}\right), \gamma\right), \quad j=1,2, \ldots p, 1,2, \ldots \tag{9}
\end{equation*}
$$

Starting with any values for $\beta_{j}$, for example, the univariate regression coefficients, it can be shown that the $\tilde{\beta}_{j}(\gamma)$ values converge to $\hat{\beta}^{\text {lasso }}$.

Figure 2 shows an example, using the diabetes data from Efron et al. (2004). This data has 442 observations and 10 predictors. We applied algorithm (9) with $\gamma=88$. It produces the iterates shown in the figure and converged after 14 steps to the lasso solution $\hat{\beta}^{\text {lasso }}(88)$.

This approach provides a simple but fast algorithm for solving the lasso, especially useful for large $p$. It was proposed in the "shooting" procedure of Fu (1998) and re-discovered by [Daubechies, Defrise and De Mol (2004)]. Application of the same idea to the elastic net procedure [Zhou and Hastie (2005)] was proposed by [Van der Kooij (2007)].
2. Pathwise coordinatewise optimization algorithms. The procedure described in Section 1 is a successive coordinate-wise descent algorithm for minimizing the function $f(\beta)=\frac{1}{2} \sum_{i}\left(y_{i}-\sum_{j} x_{i j} \beta_{j}\right)^{2}+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right|$. The idea is to apply a coordinate-wise descent procedure for each value of the regularization parameter, varying the regularization parameter along a path. Each solution is used as a warm start for the next problem.

This approach is attractive whenever the single-parameter problem is easy to solve. Some of these algorithms have already been proposed in the literature, and we give appropriate references. Here are some other examples:

- The nonnegative garotte. This method, a precursor to the lasso, was proposed by Breiman (1995) and solves

$$
\begin{equation*}
\min _{c} \frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{p} x_{i j} c_{j} \hat{\beta}_{j}\right)^{2}+\lambda \sum_{j=1}^{p} c_{j} \quad \text { subject to } \quad c_{j} \geq 0 . \tag{10}
\end{equation*}
$$

Here $\hat{\beta}_{j}$ are the usual least squares estimates (we assume that $p \leq n$ ). Using partial residuals as in (7), one can show that the the coordinate-wise update has the form

$$
\begin{equation*}
c_{j} \leftarrow\left(\frac{\tilde{\beta}_{j} \hat{\beta}_{j}-\lambda}{\hat{\beta}_{j}^{2}}\right)_{+}, \tag{11}
\end{equation*}
$$

where $\tilde{\beta}_{j}=\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}^{(j)}\right)$, and $\tilde{y}_{i}^{(j)}=\sum_{k \neq j} x_{i k} c_{k} \hat{\beta}_{k}$.

- Least absolute deviation regression and LAD-lasso. Here the problem is

$$
\begin{equation*}
\min _{\beta} \sum_{i=1}^{n}\left|y_{i}-\beta_{0}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right| \tag{12}
\end{equation*}
$$

We can write this as

$$
\begin{equation*}
\sum_{i=1}^{n}\left|x_{i j}\right|\left|\frac{\left(y_{i}-\tilde{y}_{i}^{(j)}\right)}{x_{i j}}-\beta_{j}\right| \tag{13}
\end{equation*}
$$

holding all but $\beta_{j}$ fixed. This quantity is minimized over $\beta_{j}$ by a weighted median of the values $\left(y_{i}-\tilde{y}_{i}^{(j)}\right) / x_{i j}$. Hence, coordinate-wise minimization is just a repeated computation of medians. This approach is studied and refined by Li and Arce (2004).

The same approach may be used in the LAD-lasso [Wang et al. (2006)]. Here we add an $L_{1}$ penalty to the least absolute deviation loss:

$$
\begin{equation*}
\min _{\beta} \sum_{i=1}^{n}\left|y_{i}-\beta_{0}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right|+\lambda \sum_{j=1}^{p}\left|\beta_{j}\right| \tag{14}
\end{equation*}
$$

This can be converted into an LAD problem (12) by augmenting the dataset with $p$ artificial observations. In detail, let $\mathbf{X}$ denote $n \times(p+1)$ model matrix (with the column of 1 s for the intercept). The extra $p$ observations have response $y_{i}$ equal to zero, and predictor matrix equal to $\lambda \cdot\left(\mathbf{0}: \mathbf{I}_{p}\right)$. Then the above coordinate-wise algorithm for the LAD problem can be applied.

- The elastic net. This method [due to Zhou and Hastie (2005)] adds a second constraint $\sum_{j=1}^{p} \beta_{j}^{2} \leq s_{2}$ to the lasso (1). In Lagrange form, we solve

$$
\begin{equation*}
\min _{\beta} \frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{n} x_{i j} \beta_{j}\right)^{2}+\lambda_{1} \sum_{j=1}^{p}\left|\beta_{j}\right|+\lambda_{2} \sum_{j=1}^{p} \beta_{j}^{2} / 2 \tag{15}
\end{equation*}
$$

The coordinate-wise update has the form

$$
\begin{equation*}
\tilde{\beta}_{j} \leftarrow \frac{S\left(\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}^{(j)}\right), \lambda_{1}\right)_{+}}{1+\lambda_{2}} \tag{16}
\end{equation*}
$$

Thus, we compute the simple least squares coefficient on the partial residual, apply soft-thresholding to take care of the lasso penalty, and then apply a proportional shrinkage for the ridge penalty. This algorithm was suggested by Van der Kooij (2007).

- Grouped lasso [Yuan and Lin (2006)]. This is like the lasso, except variables occur in groups (such as dummy variables for multi-level factors). Suppose $X_{j}$ is an $N \times p_{j}$ orthonormal matrix that represents the $j$ th group of $p_{j}$ variables, $j=1, \ldots, m$, and $\beta_{j}$ the corresponding coefficient vector. The grouped lasso solves

$$
\begin{equation*}
\min _{\beta}\left\|y-\sum_{j=1}^{m} X_{j} \beta_{j}\right\|_{2}^{2}+\sum_{j=1}^{m} \lambda_{j}\left\|\beta_{j}\right\|_{2} \tag{17}
\end{equation*}
$$

where $\lambda_{j}=\lambda \sqrt{p_{j}}$. Other choices of $\lambda_{j} \geq 0$ are possible; this one penalizes large groups more heavily. Notice that the penalty is a weighted sum of $L_{2}$ norms (not squared); this has the effect of selecting the variables in groups. Yuan and Lin (2006) argue for a coordinate-descent algorithm for solving this problem, and show through the Karush-Kuhn-Tucker conditions that the coordinate updates are given by

$$
\begin{equation*}
\tilde{\beta}_{j} \leftarrow\left(\left\|S_{j}\right\|_{2}-\lambda_{j}\right)_{+} \frac{S_{j}}{\left\|S_{j}\right\|_{2}} \tag{18}
\end{equation*}
$$

where $S_{j}=X_{j}^{T}\left(y-\tilde{y}^{(j)}\right)$, and here $\tilde{y}^{(j)}=\sum_{k \neq j} X_{k} \tilde{\beta}_{k}$.

- The "Berhu" penalty. This method is due to Owen (2006), and is another compromise between an $L_{1}$ and $L_{2}$ penalty. In Lagrange form, we solve

$$
\begin{align*}
& \min _{\beta} \frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{n} x_{i j} \beta_{j}\right)^{2} \\
& \quad+\lambda \sum_{j=1}^{p}\left[\left|\beta_{j}\right| \cdot I\left(\left|\beta_{j}\right|<\delta\right)+\frac{\left(\beta_{j}^{2}+\delta^{2}\right)}{2 \delta} \cdot I\left(\left|\beta_{j}\right| \geq \delta\right)\right] \tag{19}
\end{align*}
$$

This penalty is the reverse of a "Huber" function-initially absolute value, but then blending into quadratic beyond $\delta$ from zero. The coordinate-wise update has the form

$$
\tilde{\beta}_{j} \leftarrow \begin{cases}S\left(\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}^{(j)}\right), \lambda\right), & \text { if }\left|\beta_{j}\right|<\delta  \tag{20}\\ \sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}^{(j)}\right) /(1+\lambda / \delta), & \text { if }\left|\beta_{j}\right| \geq \delta\end{cases}
$$

This is a lasso-style soft-thresholding for values less than $\delta$, and ridge-style beyond $\delta$.

Tseng (1988) [see also Tseng (2001)] has established that coordinate descent works in problems like the above. He considers minimizing functions of the form

$$
\begin{equation*}
f\left(\beta_{1}, \ldots, \beta_{p}\right)=g\left(\beta_{1}, \ldots, \beta_{p}\right)+\sum_{i=j}^{p} h_{j}\left(\beta_{j}\right) \tag{21}
\end{equation*}
$$

where $g(\cdot)$ is differentiable and convex, and the $h_{j}(\cdot)$ are convex. Here each $\beta_{j}$ can be a vector, but the different vectors cannot have any overlapping members. He shows that coordinate descent converges to the minimizer of $f$. The key to this result is the separability of the penalty function $\sum_{j=1}^{p} h_{j}\left(\beta_{j}\right)$, a sum of functions of each individual parameter. This result implies that the coordinate-wise algorithms for the lasso, the grouped lasso and elastic net, etc. converge to their optimal solutions.

Next we examine coordinate-wise descent in a more complicated problem, the "fused lasso" [Tibshirani, Saunders, Rosset, Zhu and Knight (2005)], which we represent in Lagrange form:

$$
\begin{equation*}
f(\beta)=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right)^{2}+\lambda_{1} \sum_{j=1}^{p}\left|\beta_{j}\right|+\lambda_{2} \sum_{j=2}^{p}\left|\beta_{j}-\beta_{j-1}\right| . \tag{22}
\end{equation*}
$$

The first penalty encourages sparsity in the coefficients; the second penalty encourages sparsity in their differences; that is, flatness of the coefficient profiles $\beta_{j}$ as a function of the index set $j$. Note that $f(\beta)$ is strictly convex, and hence has a unique minimum.

The left panel of Figure 3 shows an example of an application of the fused lasso, in a special case where the feature matrix $\left\{x_{i j}\right\}$ is the identity matrix-this is called fused lasso signal approximation, discussed in the next section. The data represents Comparative Genomic Hybridization (CGH) measurements from two Glioblastoma Multiforme (GBM) tumors. The measurements are "copy numbers"-logratios of the number of copies of each gene in the tumor versus normal tissue. The data are displayed in the left panel, and the red line in the right panel represents the


Fig. 3. Fused lasso applied to some Glioblastoma Multiforme data. The data are shown in the left panel, and the jagged line in the right panel represents the inferred copy number $\hat{\beta}$ from the fused lasso. The horizontal line is for $y=0$.
smoothed signal $\hat{\beta}$ from the fused lasso. The regions of the nonzero estimated signal can be used to call "gains" and "losses" of genes. Tibshirani and Wang (2007) report excellent results in the application of the fused lasso, finding that the method outperforms other popular techniques for this problem.

Somewhat surprisingly, coordinate-wise descent does not work for the fused lasso. Proposition 2.7.1 of Bertsekas (1999), for example, shows that every limit point of successive coordinate-wise minimization of a continuously differentiable function is a stationary point for the overall minimization, provided that the minimum is uniquely obtained along each coordinate. However, $f(\beta)$ is not continuously differentiable, which means that coordinate-wise descent can get stuck. Looking at Tseng's result, the penalty function for the fused lasso is not separable, and hence, Tseng's theorem cannot be applied in that case.

Figure 4 illustrates the difficulty. We created a fused lasso problem with 100 parameters, with the solutions for two of the parameters, $\beta_{63}$ and $\beta_{64}$, being equal to about -1 . The top panels shows slices of the function $f(\beta)$ varying $\beta_{63}$ and $\beta_{64}$, with the other parameters set to the global minimizers. We see that the coordinatewise descent algorithm has got stuck in a corner of the response surface, and is stationary under single-coordinate moves. In order to advance to the minimum, we have to move both $\beta_{63}$ and $\beta_{64}$ together.

Despite this, it turns out that the coordinate-wise descent procedure can be modified to work for the fused lasso, yielding an algorithm that is much faster than a general quadratic-program solver for this problem.
3. The fused lasso signal approximator. Here we consider a variant of the fused lasso (22) for approximating one- and higher-dimensional signals, which we call the fused-lasso signal approximator (FLSA). For one-dimensional signals we


Fig. 4. Failure of coordinate-wise descent in a fused lasso problem with 100 parameters. The optimal values for two of the parameters, $\beta_{63}$ and $\beta_{64}$, are both -1.05 , as shown by the dot in the right panel. The left and middle panels shows slices of the objective function $f(\beta)$ as a function of $\beta_{63}$ and $\beta_{64}$, with the other parameters set to the global minimizers. The coordinate-wise minimizer over both $\beta_{63}$ and $\beta_{64}$ (separately) is -0.69 , rather than -1.05 . The right panel shows contours of the two-dimensional surface. The coordinate-descent algorithm is stuck at $(-0.69,-0.69)$. Despite being strictly convex, the surface has corners, in which the coordinate-wise procedure can get stuck. In order to travel to the minimum we have to move both $\beta_{63}$ and $\beta_{64}$ together.
solve

$$
\begin{equation*}
\min _{\beta} f(\beta)=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\beta_{i}\right)^{2}+\lambda_{1} \sum_{i=1}^{n}\left|\beta_{i}\right|+\lambda_{2} \sum_{i=2}^{n}\left|\beta_{i}-\beta_{i-1}\right| . \tag{23}
\end{equation*}
$$

The measurements $y_{i}$ are made along a one-dimensional index $i$, and there is one parameter per observation. Later we consider images as well. In the special case of $\lambda_{1}=0$, the fused-lasso signal approximator is equivalent to a discrete version of the "total variation denoising" procedure [Rudin et al. (1992)] used in signal processing. We make this connection clear in Section 6. Thus, the algorithm that we present here also provides a fast implementation for total variation denoising.

Figure 5 illustrates an example of FLSA with 1000 simulated data points, and the fit is shown for $s_{1}=269.2, s_{2}=10.9$.

We now describe a modified coordinate-wise algorithm for the diagonal fused lasso (FLSA) using the Lagrange form (23). The algorithm can also be extended to the general fused lasso problem; details are given in the Appendix. However, it is not guaranteed to give the exact solution for that problem, as we later make clear.

Our algorithm, for a fixed value $\lambda_{1}$, delivers a sequence of solutions corresponding to an increasing sequence of values for $\lambda_{2}$. First we make an observation that makes it easy to obtain the solution over a grid of $\lambda_{1}$ and $\lambda_{2}$ values:

Proposition 1. The solutions to (23) for all $\left(\lambda_{1}^{\prime}>\lambda_{1}, \lambda_{2}\right)$ can be obtained by soft-thresholding the solution for $\left(\lambda_{1}, \lambda_{2}\right)$.

This is proven in the Appendix for the FLSA, two-dimensional penalty fused lasso and even more general penalty functionals. Thus, our overall strategy for


Fig. 5. Fused lasso solution in a constructed example.
obtaining the solution over a grid is to solve the problem for $\lambda_{1}=0$ over a grid of values of $\lambda_{2}$, and then use this result to obtain the solution for all values of $\lambda_{1}$. However, for a single (especially large) value of $\lambda_{1}$, we demonstrate that it is faster to obtain the solution directly for that value of $\lambda_{1}$ (Table 2). Hence, we present our algorithm for fixed but arbitrary values of $\lambda_{1}$.

Two keys for the algorithm are assumptions (A1) and (A2) below, stating that for fixed $\lambda_{1}$, small increments in the value of $\lambda_{2}$ can only cause pairs of parameters to fuse and they do not become unfused for the larger of $\lambda_{2}$. This allows us to efficiently solve the problem for a path of $\lambda_{2}$ values, keeping $\lambda_{1}$ fixed.

The algorithm has three nested cycles:
Descent cycle: Here we run coordinate-wise descent for each parameter $\beta_{j}$, holding all the others fixed.
Fusion cycle: Here we consider the fusion of a neighboring pair of parameters, followed by coordinate-wise descent.
Smoothing cycle: Here we increase the penalty $\lambda_{2}$ a small amount, and rerun the two previous cycles.

We now describe each in more detail.
Descent cycle. Consider the derivative of (23), holding all $\beta_{k}=\tilde{\beta}_{k}, k \neq i$ fixed at their current values:

$$
\begin{align*}
\frac{\partial f(\beta)}{\partial \beta_{i}}= & -\left(y_{i}-\beta_{i}\right)+\lambda_{1} \cdot \operatorname{sign}\left(\beta_{i}\right)  \tag{24}\\
& -\lambda_{2} \cdot \operatorname{sign}\left(\tilde{\beta}_{i+1}-\beta_{i}\right)+\lambda_{2} \cdot \operatorname{sign}\left(\beta_{i}-\tilde{\beta}_{i-1}\right),
\end{align*}
$$

assuming that $\beta_{i} \notin\left\{0, \tilde{\beta}_{i-1}, \tilde{\beta}_{i+1}\right\}$.
The algorithm for coordinate-wise minimization of $f(\beta)$ works as follows. The expression in (24) is piecewise linear in $\beta_{i}$ with breaks at $0, \tilde{\beta}_{i-1}$ and $\tilde{\beta}_{i+1}$. Suppose, for example, at a given iteration we have $0 \leq \tilde{\beta}_{i-1} \leq \tilde{\beta}_{i+1}$. We check for


FIG. 6. Example of the one-dimensional search in the coordinate descent cycle for a FLSA problem with 5 parameters. Shown is the gradient for $\beta_{3}$, with the other 4 parameters set at the global minimizing values. There are discontinuities at $\tilde{\beta}_{2}, \tilde{\beta}_{4}$ and zero. We look for a zero-crossing in each of the intervals $\left(-\infty, \tilde{\beta}_{4}\right),\left(\tilde{\beta}_{4}, 0\right),\left(0, \tilde{\beta}_{2}\right),\left(\tilde{\beta}_{2}, \infty\right)$, and if none is found, take the minimum of $f(\beta)$ over the set of discontinuities. In this case, the minimum is at a discontinuity, with $\tilde{\beta}_{3}=\tilde{\beta}_{4}$.
a zero of (24) in each of the four intervals $(-\infty, 0],\left(0, \tilde{\beta}_{i-1}\right],\left(\tilde{\beta}_{i-1}, \tilde{\beta}_{i+1}\right]$ and $\left(\tilde{\beta}_{i+1}, \infty\right)$. Since the function is piecewise linear, there is an explicit solution, when one exists. If no solution is found, we examine the three active-constraint values for $\beta_{i}: 0, \tilde{\beta}_{i-1}$ and $\tilde{\beta}_{i+1}$, and find the one giving the smallest value of the objective function $f(\beta)$. Figure 6 illustrates the procedure on a simulated example.

Other orderings among $0, \tilde{\beta}_{i-1}$ and $\tilde{\beta}_{i+1}$ are handled similarly, and at the endpoints $i=1$ and $i=p$, there are only three intervals to check, rather than four.

Fusion cycle. The descent cycle moves parameters one at a time. Inspection of Figure 4 shows that this approach can get stuck. One way to get around this is to consider a potential fusion of parameters, when a move of a single $\beta_{i}$ fails to improve the loss criterion. This amounts to enforcing $\left|\beta_{i}-\beta_{i-1}\right|=0$ by setting $\beta_{i}=\beta_{i-1}=\gamma$. With this constraint, we try a descent move in $\gamma$. Equation (24) now becomes

$$
\begin{align*}
\frac{\partial f(\beta)}{\partial \gamma}= & -\left(y_{i-1}-\gamma\right)-\left(y_{i}-\gamma\right) \\
& +2 \lambda_{1} \cdot \operatorname{sign}(\gamma)-\lambda_{2} \cdot \operatorname{sign}\left(\tilde{\beta}_{i+1}-\gamma\right)  \tag{25}\\
& +\lambda_{2} \cdot \operatorname{sign}\left(\gamma-\tilde{\beta}_{i-2}\right)
\end{align*}
$$

If the optimal value for $\gamma$ decreases the criterion, we accept the move setting $\beta_{i}=$ $\beta_{i-1}=\gamma$ 。

Notice that the fusion step is equivalent to temporarily collapsing the problem to one with $p-1$ parameters:

- we replace the pair $y_{i-1}$ and $y_{i}$ with the average response $\bar{y}=\left(y_{i-1}+y_{i}\right) / 2$ and an observation weight of 2 ;
- the pair of parameters $\beta_{i-1}$ and $\beta_{i}$ are replaced by a single $\gamma$, with a penalty weight of 2 for the first penalty.

At the end of the entire process of descent and fusion cycles for a given $\lambda_{2}$, we identify adjacent nonzero solution values that are equal and collapse the data accordingly, maintaining a weight vector to assign weights to the observations averages and the contributions to the first penalty.

Smoothing cycle. Although the fusion step often leads to a decrease in $f(\tilde{\beta})$, it is possible to construct examples where, for a particular value of $\lambda_{2}$, no fusion of two neighbors causes a decrease, but a fusion of three or more can. Our final strategy is to solve a series of fused lasso problems sequentially, fixing $\lambda_{1}$, but varying $\lambda_{2}$ through a range of values increasing in small increments $\delta$ from 0 .

The smoothing cycle is then as follows:

1. Start with $\lambda_{2}=0$, hence, with the lasso solution with penalty parameter $\lambda_{1}$.
2. Increment $\lambda_{2} \leftarrow \lambda_{2}+\delta$, and run the descent and fusion cycles repeatedly until no further changes occur. After convergence of the process for a given value $\lambda_{2}$, identify neighboring solution values that are equal and nonzero and collapse the problem as described above, updating the weights.
3. Repeat step 2 until a target value of $\lambda_{2}$ is reached (or a target bound $s_{2}$ ).

Our strategy relies on the following assumptions:
(A1) If the increments $\delta$ are sufficiently small, fusions will occur between no more than two neighboring points at a time.
(A2) Two parameters that are fused in the solution for $\left(\lambda_{1}, \lambda_{2}\right)$ will be fused for all $\left(\lambda_{1}, \lambda_{2}^{\prime}>\lambda_{2}\right)$.
By collapsing the data after each solution, we can achieve long fusions by a sequence of pairwise fusions. Note that each of the fused parameters can represent more than one parameter in the original problem. For example, if $\beta_{j}$ has a weight of 3 , and $\beta_{i+1}$ a weight of 2 , then the merged parameter has a weight of 5 , and represents 5 neighboring parameters in the original problem.

After $m$ fusions, the problem has the form

$$
\begin{equation*}
C_{m}+\min _{\beta} \frac{1}{2} \sum_{i=1}^{n-m} w_{i}\left(y_{i}-\beta_{i}\right)^{2}+\lambda_{1} \sum_{i=1}^{n-m} w_{i}\left|\beta_{i}\right|+\sum_{i=2}^{n-m}\left|\beta_{i}-\beta_{i-1}\right| \tag{26}
\end{equation*}
$$

Initially, $m=0, w_{i}=1$, and $C_{0}=0$. If the $(m+1)$ st fusion is between $\beta_{i-1}$ and $\beta_{i}$, then the following updates occur:

- $\bar{y} \leftarrow\left(w_{i-1} y_{i-1}+w_{i} y_{i}\right) /\left(w_{i-1}+w_{i}\right)$.
- $w^{+} \leftarrow w_{i-1}+w_{i}$.
- $C_{m+1}=C_{m}+\frac{1}{2}\left[w_{i-1} \cdot\left(y_{i-1}-\bar{y}\right)^{2}+w_{i} \cdot\left(y_{i}-\bar{y}\right)^{2}\right]$.
- $y_{i-1} \leftarrow \bar{y}, w_{i-1} \leftarrow w^{+}$.


Fig. 7. Small example of the fused lasso. $\lambda_{1}$ is fixed at 0.01 ; as $\lambda_{2}$ increases, the number of fused parameters increases.

- Discard observation $i$, and reduce all indices greater than $i$ by 1 .

Note that we don't actually need to carry out the update for $C_{m}$, because no parameters are involved.

Figure 7 shows an example with just 9 data points. We have fixed $\lambda_{1}=0.01$ and show the solutions for four values of $\lambda_{2}$. As $\lambda_{2}$ increases, the number of fused parameters increases.

Assumption (A1) requires that the data have some randomness (i.e., no preexisting flat plateaus exist). Assumption (A2) holds in general. We prove that both assumptions hold for the FLSA procedure in the next section.

Numerical experiments show that (A2) does not always hold for the general fused lasso. Hence, the extension of this algorithm to the general fused lasso (detailed in the Appendix) is not guaranteed to yield the exact solution. Note that each descent and fusion cycle can only decrease the convex objective and, hence, must converge. We terminate this pair of cycles when the change in parameter estimates is less than some threshold. The smoothing cycle is done over a discrete grid of $\lambda_{2}$ values.
4. Optimality conditions. In this section we derive the optimality conditions for the FLSA problem, and use them to show that our algorithms' assumptions (A1) and (A2) are satisfied.

We consider the Lagrangian form (23) for the fused lasso. The standard Karush-Kuhn-Tucker conditions for this problem are fairly complicated, since we need to express each parameter in terms of its positive and negative parts in order to make the penalty differentiable. A more convenient formulation is through the subgradient approach [see, e.g., Bertsekas (1999), Proposition B.24]. The equations for the subgradient have the form

$$
\begin{align*}
-\left(y_{1}-\beta_{1}\right)+\lambda_{1} s_{1}-\lambda_{2} t_{2} & =0,  \tag{27}\\
-\left(y_{j}-\beta_{j}\right)+\lambda_{1} s_{j}+\lambda_{2}\left(t_{j}-t_{j+1}\right) & =0, \quad j=2, \ldots, n,
\end{align*}
$$

with $s_{j}=\operatorname{sign}\left(\beta_{j}\right)$ if $\beta_{j} \neq 0$ and $s_{j} \in[-1,1]$ if $\beta_{j}=0$. Similarly, $t_{j}=\operatorname{sign}\left(\beta_{j}-\right.$ $\beta_{j-1}$ ) if $\beta_{j} \neq \beta_{j-1}$ and $t_{j} \in[-1,1]$ if $\beta_{j}=\beta_{j-1}$. These $n$ equations are necessary and sufficient for the solution. We restate assumptions (A1) and (A2) more precisely, and then prove they hold.

Proposition 2. For the fused-lasso signal approximation algorithm detailed in Section 3:
( $\mathrm{Al}^{\prime}$ ) If the sequence $y_{i}$ are in general position-specifically, no two consecutive $y_{j}$ values are equal-and the increments $\delta$ are sufficiently small, fusions will occur between no more than two neighboring points at a time.
(A2') Two parameters that are fused in the solution for $\left(\lambda_{1}, \lambda_{2}\right)$ will be fused for all $\left(\lambda_{1}, \lambda_{2}^{\prime}>\lambda_{2}\right)$.

Proof. We first prove (A2'). Suppose we have a stretch of nonzero solutions $\hat{\beta}_{j-k}, \hat{\beta}_{j-k+1}, \ldots, \hat{\beta}_{j}$ that are equal for some value $\lambda_{2}$, and $\hat{\beta}_{j-k-1}$ and $\hat{\beta}_{j+1}$ are not equal to this common value. Then $t_{j-k}$ and $t_{j+1}$ each take a value in $\{-1,1\}$; we denote these boundary values by $T_{j-k}$ and $T_{j+1}$. Although the parameters $t_{j-k+1}, \ldots, t_{j}$ can vary in $[-1,1]$ as $\lambda_{2}$ changes (while the fused group remains intact), the values depend on only the $(j-k+1)$ st through $j$ th equations in the system (27), because the boundary values are fixed. Taking pairwise differences, and using the fact that $\hat{\beta}_{j-k}=\hat{\beta}_{j-k+1}=\cdots=\hat{\beta}_{j}$, this subgroup of equations simplifies to

$$
\begin{gather*}
\left(\begin{array}{cccccc}
2 & -1 & 0 & 0 & \cdots & 0 \\
-1 & 2 & -1 & 0 & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & -1 & 2 & -1 \\
0 & 0 & \cdots & 0 & -1 & 2
\end{array}\right)\left(\begin{array}{c}
t_{j-k+1} \\
t_{j-k+2} \\
\vdots \\
t_{j-1} \\
t_{j}
\end{array}\right)  \tag{28}\\
=\frac{1}{\lambda_{2}}\left(\begin{array}{c}
y_{j-k+1}-y_{j+k} \\
y_{j-k+2}-y_{j-k+1} \\
\vdots \\
y_{j-1}-y_{j-2} \\
y_{j}-y_{j-1}
\end{array}\right)+\left(\begin{array}{c}
T_{j-k} \\
0 \\
\vdots \\
0 \\
T_{j+1}
\end{array}\right)
\end{gather*}
$$

Write this system symbolically as $M t=\frac{1}{\lambda_{2}} \Delta y+T$, and let $C=M^{-1}$. The explicit form for $C$ given in Schlegel (1970) gives $C_{\ell 1}=(n-\ell+1) /(n+1), C_{\ell n}=$ $\ell /(n+1)$. It is easy to check for all three possibilities for $T$ that $C T \in[-1,1]$ elementwise. We know that $t=\left(\frac{1}{\lambda_{2}} C \Delta y+C T\right) \in[-1,1]$ elementwise as well, since $t$ is a solution to (23) at $\lambda_{2}$. For $\lambda_{2}^{\prime}>\lambda_{2}$, the elements of the first terms shrink, and hence the values $t\left(\lambda_{2}^{\prime}\right)$ remain in $[-1,1]$. This implies that the fused set remains fused as we increase $\lambda_{2}$. These equations describe the path $t\left(\lambda_{2}\right)$ for $\lambda_{2}$ increasing, and only change when one of the boundary points (fused sets) is fused with the current set, and the argument is repeated. This proves ( $\mathrm{A} 2^{\prime}$ ).

We now address ( $\mathrm{A}^{\prime}$ ). Suppose the data are in general position (e.g., have a random noise component), and we have the lasso-solution $\hat{\beta}_{j}$ for $\lambda_{1}$. Because of the randomness, no neighboring nonzero parameters $\hat{\beta}_{j}$ will be exactly the same. This means for each nonzero value $\hat{\beta}_{j}$, we can write an equation of the form (27) where we know exactly the values for $s_{j}, t_{j}$ and $t_{j+1}$ (each will be one of the values $\{-1,+1\}$ ). This means that we can calculate exactly the path of each such $\beta_{j}$ as we increase $\lambda_{2}$ from zero, until an event occurs that changes the $s_{j}, t_{j}$. By looking at all pairs, we can identify the time of the first fusion of such pairs. The data are then fused together and reduced, and the problem is repeated. Fusions occur one-at-a-time in this fashion, at a distinct sequence of values for $\lambda_{2}$. Hence, for $\delta$ small enough in our smoothing step, we can ensure that we encounter these fusions one at a time.
5. Comparison of run times. In this section we compare the run times of the coordinate-wise algorithm to standard algorithms, for both the lasso and diagonal fused lasso (FLSA) problems. All timings were carried out on a Intel Xeon 2.80GH processor.
5.1. Lasso speed trials. We generated Gaussian data with $n$ observations and $p$ predictors, with each pair of predictors $X_{j}, X_{j^{\prime}}$ having the same population correlation $\rho$. We tried a number of combinations of $n$ and $p$, with $\rho$ varying from zero to 0.95 . The outcome values were generated by

$$
\begin{equation*}
Y=\sum_{j=1}^{p} \beta_{j} X_{j}+k \cdot Z \tag{29}
\end{equation*}
$$

where $\beta_{j}=(-1)^{j} \exp (-2(j-1) / 20), Z \sim N(0,1)$ and $k$ is chosen so that the signal-to-noise ratio is 3.0 . The coefficients are constructed to have alternating signs and to be exponentially decreasing.

Table 1 shows the average CPU timings for the coordinatewise algorithm, two versions of the LARS procedure and lasso2, an implementation of the homotopy algorithm of Osborne et al. (2000). All algorithms are implemented as R language functions. The coordinate-wise algorithm does all of its numerical work in Fortran,

TABLE 1
Run times (CPU seconds) for lasso problems of various sizes $n, p$ and different correlation between the features. Methods are the coordinate-wise optimization (Fortran), LARS (R and Fortran versions) and lasso2 (C language)—the homotopy procedure of Osborne et al. (2000)

| Method | Population correlation between features |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | ---: |
|  |  | $n=100, p=1000$ |  |  |  |  |
|  | 0 | 0.1 | 0.2 | 0.5 | 0.9 | 0.95 |
| coord-Fort | 0.31 | 0.33 | 0.40 | 0.57 | 1.20 | 1.45 |
| LARS-R | 2.18 | 2.46 | 2.14 | 2.45 | 2.37 | 2.10 |
| LARS-Fort | 2.01 | 2.09 | 2.12 | 1.947 | 2.50 | 2.22 |
| lasso2-C | 2.42 | 2.16 | 2.39 | 2.18 | 2.01 | 2.71 |
|  |  |  | $n=100, p=5000$ |  |  |  |
|  | 0 | 0.1 | 0.2 | 0.5 | 0.9 | 0.95 |
| coord-Fort | 4.66 | 4.51 | 3.14 | 5.77 | 4.44 | 5.43 |
| LARS-R | 28.40 | 27.34 | 24.40 | 22.32 | 22.16 | 22.75 |
| LARS-Fort would not run |  |  |  |  |  |  |
| lasso2 would not run |  |  |  |  |  |  |


|  | 0 | 0.1 | 0.2 | 0.5 | 0.9 | 0.95 |
| :--- | :---: | ---: | ---: | ---: | ---: | ---: |
| coord-Fort | 7.03 | 9.34 | 8.83 | 10.62 | 27.46 | 40.37 |
| LARS-R | 116.26 | 122.39 | 121.48 | 104.17 | 100.30 | 107.29 |
| LARS-Fort would not run |  |  |  |  |  |  |
| lasso2 would not run |  |  |  |  |  |  |


|  | $n=1000, p=100$ |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| coord-Fort | 0 | 0.1 | 0.2 | 0.5 | 0.9 | 0.95 |
| LARS-R | 0.03 | 0.04 | 0.04 | 0.04 | 0.06 | 0.08 |
| LARS-Fort | 0.42 | 0.41 | 0.40 | 0.40 | 0.40 | 0.40 |
| lasso2-C | 0.30 | 0.24 | 0.22 | 0.23 | 0.23 | 0.28 |
|  | 0.73 | 0.66 | 0.69 | 0.68 | 0.69 | 0.70 |
|  |  |  | $n=5000, p=100$ |  |  |  |
| coord-Fort | 0 | 0.1 | 0.2 | 0.5 | 0.9 | 0.95 |
| LARS-R | 0.16 | 0.15 | 0.14 | 0.16 | 0.15 | 0.16 |
| LARS-Fort | 1.02 | 1.03 | 1.02 | 1.04 | 1.02 | 1.03 |
| lasso2-C | 1.07 | 1.09 | 1.10 | 1.10 | 1.10 | 1.08 |

while lasso2 does its numerical work in C. LARS-R is the "production version" of LARS (written by Efron and Hastie), doing much of its work in R, calling Fortran routines for some matrix operations. LARS-Fort (due to Ji Zhu) is a version of LARS that does all of its numerical work in Fortran. Comparisons between differ-


FIG. 8. CPU times for coordinate descent, for the same problem as in Table 1, for different values of $n$ and $p$. In each case the times are averaged over five runs and averaged over the set of values of the other parameter ( $n$ or $p$ ).
ent programs are always tricky: in particular, the LARS procedure computes the entire path of solutions, while the coordinate-wise procedure and lasso 2 solve the problems for a set of pre-defined points along the solution path. In the orthogonal case, LARS takes $\min (n, p)$ steps: hence, to make things roughly comparable, we called the latter two algorithms to solve a total of $\min (n, p)$ problems along the path.

Not surprisingly, the coordinate-wise algorithm is fastest when the correlations are small, and gets slower when they are large. It seems to be very competitive with the other two algorithms in general, and offers some potential speedup, especially when $n>p$.

Figure 8 shows the CPU times for coordinate descent, for the same problem as in Table 1. We varied $n$ and $p$, and averaged the times over five runs. We see that the times are roughly linear in $n$ and in $p$.

A key to the success of the coordinate-wise algorithm for lasso is the fact that, for squared error loss, the ingredients needed for each coordinate step can be easily updated as the algorithm proceeds. We can write the second term in (9) as

$$
\begin{equation*}
\sum_{i=1}^{n} x_{i j}\left(y_{i}-\tilde{y}_{i}\right)=\left\langle x_{j}, y\right\rangle-\sum_{k: \tilde{\beta}_{k}>0}\left\langle x_{j}, x_{k}\right\rangle \tilde{\beta}_{k} \tag{30}
\end{equation*}
$$

where $\left\langle x_{j}, y\right\rangle=\sum_{i=1}^{n} x_{i j} y_{i}$, and so on. Hence, we need to compute inner products of each feature with $y$ initially, and then each time a feature $x_{k}$ enters the model, we need to compute its inner product with all the rest. But importantly, $O(n)$ calculations do not have to be made at every step. This is the case for all penalized procedures with squared error loss.

Friedlander and Saunders (2007) do a thorough comparison of the LARS (homotopy) procedure to a number of interior point QP procedures for the lasso problem, and find that LARS is generally much faster. Our finding that coordinate de-
scent is very competitive with LARS therefore suggests that also will outperform interior point methods.

Finally, note that there is another approach to solving the FLSA problem for $\lambda_{1}=0$. We can transform to parameters $\theta_{j}=\beta_{j}-\beta_{j-1}$, and we get a new lasso problem in these new parameters. One can use coordinate descent to solve this lasso problem, and then Proposition 1 gives the FLSA solution for other values of $\lambda_{1}$. However, this new lasso problem has a dense data matrix, and hence, the coordinate descent procedure is many times slower than the procedure described in this section. The procedure developed here exploits the near-diagonal structure of the problem in the original parametrization.
5.2. Fused lasso speed trials. For the example of Figure 5, we compared the pathwise coordinate algorithm to the two-phase active set algorithm sqopt of Gill, Murray and Saunders (1999). Both algorithms are implemented as R functions, but do all but the setup and wrapup computations in Fortran. Table 2 shows

TABLE 2
Run times (CPU seconds) for fused lasso (FLSA) problems of various sizes $n$ for different values of the regularization parameters $\lambda_{1}, \lambda_{2}$. The methods compared are the pathwise coordinate optimization, and "standard"-two-phase active set algorithm sqopt of Gill, Murray and Saunders (1999). The number of active constraints in the solution is shown in each case

| $\boldsymbol{\lambda}_{\mathbf{1}}$ | $\boldsymbol{\lambda} \mathbf{2}$ | \# Active | Coord | Standard |
| :---: | :---: | :---: | :---: | :---: |
|  |  | $n=1000$ |  |  |
| 0.00 | 0.01 | 456 | 0.040 | 2.100 |
| 0.00 | 1.00 | 934 | 0.024 | 0.931 |
| 0.00 | 2.00 | 958 | 0.019 | 0.987 |
| 1.00 | 0.01 | 824 | 0.022 | 1.519 |
| 1.00 | 1.00 | 975 | 0.024 | 1.561 |
| 1.00 | 2.00 | 981 | 0.023 | 1.404 |
| 2.00 | 0.01 | 861 | 0.023 | 1.499 |
| 2.00 | 1.00 | 983 | 0.023 | 1.418 |
| 2.00 | 2.00 | 991 | 0.018 | 1.407 |
|  |  | $n=5000$ |  |  |
| 0.000 | 0.002 | 4063 | 0.217 | 20.689 |
| 0.000 | 0.200 | 3787 | 0.170 | 26.195 |
| 0.000 | 0.400 | 4121 | 0.135 | 29.192 |
| 0.200 | 0.002 | 4305 | 0.150 | 41.105 |
| 0.200 | 0.200 | 4449 | 0.141 | 48.998 |
| 0.200 | 0.400 | 4701 | 0.129 | 45.136 |
| 0.400 | 0.002 | 4301 | 0.108 | 41.062 |
| 0.400 | 0.200 | 4540 | 0.123 | 41.755 |
| 0.400 | 0.400 | 4722 | 0.119 | 38.896 |

TABLE 3
Run times (CPU seconds) for pathwise coordinate optimization applied to fused lasso (FLSA) problems with a large number of parameters $n$ averaged over different values of the regularization parameters $\lambda_{1}, \lambda_{2}$

| $\boldsymbol{n}$ | Average CPU sec |
| :--- | :---: |
| 100,000 | 3.54 |
| 500,000 | 14.93 |
| $1,000,000$ | 29.81 |

the timings for the two algorithms for a range of values of $\lambda_{1}$ and $\lambda_{2}$. The resulting number of active constraints (i.e., $\beta_{j}=0$ or $\beta_{j}-\beta_{j-1}=0$ ) is also shown. In the second part of the table, we increased the sample size to 5000 . We see that the coordinate algorithm offers substantial speedups, by factors of 50 up to 300 or more.

In these tables, each entry for the pathwise coordinate procedure is the computation time for the entire path of solutions leading to the given values $\lambda_{1}, \lambda_{2}$. In practice, one could obtain all of the solutions for a given $\lambda_{1}$ from a single run of the algorithm, and hence, the numbers in the table are very conservative. But we reported the results in this way to make a fair comparison with the standard procedure since it can also exploit warm starts.

In Table 3 we show the run times for pathwise coordinate optimization for larger values of $n$. As in the previous table, these are the averages of run times for the entire path of solutions for a given $\lambda_{1}$, and hence, are conservative. We were unable to run the standard algorithm for these cases.
6. The two-dimensional fused lasso. Suppose we have a total of $n^{2}$ cells, laid out in a $n \times n$ grid (the square grid is not essential). We can generalize the diagonal fused lasso (FLSA) to two-dimensions as follows:

$$
\begin{align*}
& \min _{\beta} \frac{1}{2} \sum_{i=1}^{n} \sum_{i^{\prime}=1}^{n}\left(y_{i i^{\prime}}-\beta_{i i^{\prime}}\right)^{2} \\
& \text { subject to } \sum_{i=1}^{n} \sum_{i^{\prime}=1}^{n}\left|\beta_{i i^{\prime}}\right| \leq s_{1}, \\
& \sum_{i=1}^{n} \sum_{i^{\prime}=2}^{n}\left|\beta_{i, i^{\prime}}-\beta_{i, i^{\prime}-1}\right| \leq s_{2},  \tag{31}\\
& \sum_{i=2}^{n} \sum_{i^{\prime}=1}^{n}\left|\beta_{i, i^{\prime}}-\beta_{i-1, i^{\prime}}\right| \leq s_{3} .
\end{align*}
$$

The penalties encourage the parameter map $\beta_{i i^{\prime}}$ to be both sparse and spatially smooth.

The fused lasso is related to signal processing methods such as "total variation denoising" [Rudin et al. (1992)], which uses a continuous smoothness penalty analogous to the second penalty in the fused lasso. The TV criterion is written in the form

$$
\begin{equation*}
\min _{u} \int_{\Omega}|\nabla u| d u \quad \text { subject to } \quad\|u-y\|^{2}=\sigma^{2} \tag{32}
\end{equation*}
$$

where $y$ is the data, $u$ is the approximation with allowable error $\sigma^{2}, \Omega$ is a bounded convex region in $R^{d},|\cdot|$ denotes Euclidean norm in $R^{d}$ and $\|\cdot\|$ denotes the norm on $L^{2}(\Omega)$. Thus, in $d=1$ dimension, this is a continuous analogue of the fused lasso, but without the (first) $L_{1}$ penalty on the coefficients. In $d=2$ dimensions, the TV approach is different: the discretized version uses the Euclidean norm of the first differences in $u$, rather than the sum of the absolute values of the first differences.

This problem (32) can be solved by a general purpose quadratic-programming algorithm; we give details in the Appendix. However, for a $p \times p$ grid, there are $7 p^{2}$ variables and $3 p^{2}+3$ constraints, in addition to nonnegativity constraints on the variables. For $p=256$, for example, this is 458,752 variables and 196, 611 constraints, so that finding the exact solution is impractical.

Hence, we focus on the pathwise coordinate algorithm. The algorithm has the same form as in the one-dimensional case, except that rather than checking the three active constraint values $0, \beta_{j-1}$ and $\beta_{j+1}$, we check 0 and the four values to the right, left, above and below (its four-neighborhood). The number of constraint values reduces to 4 at the edges and 3 at the corners. The algorithm starts with individual pixels as the groups, and the four-neighborhood pixels are its "distance-1 neighbors." In each fusion cycle we try to fuse a group with its distance-1 neighbors. If the fusion is accepted, then the distance-1 neighbors of the fused group are the union of the distance-1 neighbors of the two groups being joined (with the groups themselves removed). Now one pixel might be the distance-1 neighbor to each of the two groups being fused, and some careful bookkeeping is required to keep track of this through appropriate weights. Full details are given in the Appendix.

We do not provide a proof of the correctness of this procedure. However, in our (limited) experiments we have found that it gives the exact solution to the two-dimensional fused lasso. We guess that a proof along the lines of that in the one-dimensional case can be constructed, although some additional assumptions may be required.

As in the one-dimensional FLSA (Proposition 1), if we write the problem in terms of Lagrange multipliers ( $\lambda_{1}, \lambda_{2}, \lambda_{3}$ ), the solution for $\left(\lambda_{1}^{\prime}>\lambda_{1}, \lambda_{2}, \lambda_{3}\right)$ can be obtained by soft-thresholding the solutions for $\left(\lambda_{1}, \lambda_{2}, \lambda_{3}\right)$.
6.1. Example 1. Figure 9 shows a toy example. The data are in the top left, representing a " + "-shaped image with $N(0,1)$ noise added. The reconstruction by


FIG. 9. A toy example: the data are in the top left, representing a "+"-shaped image with added noise. The reconstructions by the lasso and fused lasso are shown in the other panels. In each case we did a grid search over the tuning parameters using a kind of two-fold validation.
the lasso and fused lasso are shown in the other panels. In each case we did a grid search over the tuning parameters using a kind of two-fold validation. We created a training set of the odd pixels $(1,3,5 \ldots$ in each direction) and tested it on the even pixels. For illustration only, we chose the values that minimized the squared reconstruction error over the test set. We see that the fused lasso has successfully exploited the spatial smoothness and provided a much better reconstruction than the lasso.

Table 4 shows the number of CPU seconds required for the standard and pathwise coordinate descent algorithms, as $n$ increases. We were unable to apply the standard algorithm for $n=256$ (due to memory requirements), and have instead estimated the time by crude quadratic extrapolation.

TABLE 4
2D fused lasso applied to the toy problem. The table shows the number of CPU seconds required for the standard and pathwise coordinate descent algorithms, as $n$ increases. The regularization parameters were set at the values that yielded the solution in the bottom left panel of Figure 9

| $\boldsymbol{n}$ | Standard | Coord |
| ---: | :---: | :---: |
| 8 | 2.0 s | 0.07 s |
| 16 | 3.4 s | 0.13 s |
| 32 | 20.8 s | 0.38 s |
| 256 | 38 min | 7.1 s |

6.2. Example 2. Figure 10 shows another example. The noiseless image (top panel) was randomly generated with $512 \times 512$ pixels. The background pixels are zero, while the signal blocks have constant values randomly chosen between 1 and 4. The top right panel shows the reconstruction by the fused lasso: as expected, it is perfect. In the bottom left we have added Gaussian noise with standard deviation 1.5. The reconstruction by the fused lasso is shown in the bottom right panel, using two-fold validation to estimate $\lambda_{1}, \lambda_{2}$. The reconstruction is still quite good, capturing most of the important features. In this example we did a search over $10 \lambda_{1}$ values. The entire computation for the bottom right panel of Figure 10, including the two-fold validation to estimate the optimal values for $\lambda_{1}$ and $\lambda_{2}$, took 11.3 CPU minutes.
6.3. Example 3. The top left panel of Figure 11 shows a $256 \times 256$ gray scale image of statistician Sir Ronald Fisher. In the top right we have added Gaussian noise with a standard deviation 2.5 . We explore the use of the two-dimensional fused lasso for denoising this image. However, the first (lasso) penalty doesn't make sense here, as zero does not represent a natural baseline. So instead, we tried a pure fusion model, with $\lambda_{1}=0$. We found the best value of $\lambda_{2}$, in terms of reconstruction error from the original noiseless image. The solution shown in the bottom right panel gives a reasonable approximation to the original image and reduces the reconstruction error from 6.18 to 1.15 . In the bottom left panel we have set $\lambda_{2}=0$, and optimized over $\lambda_{1}$. As expected, this pure lasso solution does poorly, and the optimal value of $\lambda_{1}$ turned out to be 0 .
6.4. Applications to higher dimensions and other problems. The general strategy for the two-dimensional fused lasso can be directly applied in higher dimensional problems, the difference being that each cell would have more potential distance-1 neighbors. In fact, the same strategy might be applicable to nonEuclidean problems. All one needs is a notion of distance-1 neighbors and the

Original image


Noisy image


Fused lasso reconstruction


Fused lasso reconstruction


Fig. 10. A second toy example. The $100 \times 100$ noiseless and noisy images are shown on the left, while the corresponding fused lasso reconstructions are shown on the right. In each case we did a grid search over the tuning parameters $\lambda_{1}, \lambda_{2}$ using a kind of two-fold validation.
property that the distance-1 neighbors of a fusion of two groups are the union of the distance-1 neighbors of the two groups, less the fused group members themselves.
7. Discussion. Coordinate-wise descent algorithms deserve more attention in convex optimization. They are simple and well-suited to large problems. We have found that for the lasso, coordinate-wise descent is very competitive with the LARS algorithm, probably the fastest procedure for that problem to-date. Coordinate-wise descent algorithms can be applied to problems in which the constraints decouple, and a generalized version of coordinate-wise descent like the one presented here can handle problems in which each parameter is involved in only a limited number of constraints. This procedure is ideally suited for a special case of


Fig. 11. Top-left panel: $256 \times 256$ gray scale image of statistician Sir Ronald Fisher. Top-right panel: Gaussian noise with standard deviation 2.5 has been added. Bottom-left panel: best solution with $\lambda_{2}$ set to zero (pure lasso penalty); this gives no improvement in reconstruction error. Bottom-right panel: best solution with $\lambda_{1}$ set to zero (pure fusion penalty). This reduces the reconstruction error from 6.18 to 1.15 .
the fused lasso-the fused lasso signal approximator, and runs many times faster than a standard convex optimizer. On the other hand, it is not guaranteed to work for the general fused lasso problem, as it can get stuck away from the solution.
7.1. Software. Both Fortran and R language routines for the lasso, and the one- and two-dimensional fused lasso will be made freely available.

## APPENDIX

A.1. Proof of Proposition 1. Suppose that we are optimizing a function of the form

$$
f(\beta)=\frac{1}{2} \sum_{i=1}^{n}\left(y_{i}-\beta_{i}\right)^{2}+\lambda_{1} \sum_{i=1}^{n}\left|\beta_{i}\right|+\sum_{(i, j) \in \mathcal{C}} \lambda_{i, j}\left|\beta_{i}-\beta_{j}\right|,
$$

where $(i, j) \in \mathcal{C}$ if the difference $\left|\beta_{i}-\beta_{j}\right|$ is $L_{1}$ penalized with penalty parameter $\lambda_{i, j}$. This general form for the penalty includes the following models discussed earlier:

Fused lasso signal approximator: Here, $(i, j) \in \mathcal{C}$ if $i=j-1$. Furthermore, $\lambda_{i, j}=$ $\lambda_{2}$.
Two-dimensional fused lasso: Here $i$ itself is a two-dimensional coordinate $i=$ $\left(i_{1}, i_{2}\right)$. Let $(i, j) \in \mathcal{C}$ if $\left|i_{1}-j_{1}\right|+\left|i_{2}-j_{2}\right|=1$. If $\left|i_{1}-j_{1}\right|=1$, then $\lambda_{i, j}=\lambda_{2}$, otherwise $\lambda_{i, j}=\lambda_{3}$.

Now we prove a soft thresholding result.
Lemma A.1. Assume that the solution for $\lambda_{1}=0$ and $\lambda_{2} \geq 0$ is known and denote it by $\hat{\beta}\left(0, \lambda_{2}\right)$. Then, the solution for $\lambda_{1}>0$ is

$$
\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)=\operatorname{sign}\left(\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right)\left(\left|\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right|-\lambda_{1}\right)^{+} .
$$

Proof. First find the subgradient equations for $\beta_{1}, \ldots, \beta_{n}$, which are

$$
g_{i}=-\left(y_{i}-\beta_{i}\right)+\lambda_{1} s_{i}+\sum_{j:(i, j) \in \mathcal{C}} \lambda_{2} t_{i, j}-\sum_{j:(j, i) \in \mathcal{C}} \lambda_{2} t_{j, i}=0,
$$

where $s_{i}=\operatorname{sign}\left(\beta_{i}\right)$ if $\beta_{i} \neq 0$ and $s_{i} \in[-1,1]$ if $\beta_{i}=0$. Also, $t_{i, j}=\operatorname{sign}\left(\beta_{i}-\right.$ $\beta_{j}$ ) for $\beta_{i} \neq \beta_{j}$ and $t_{i, j} \in[-1,1]$ if $\beta_{i}=\beta_{j}$. These equations are necessary and sufficient for the solution.

As it is assumed that a solution for $\lambda_{1}=0$ is known, let $s(0)$ and $t(0)$ denote the values of the parameters for this solution. Specifically, $s_{i}(0)=\operatorname{sign}\left(\hat{\beta}_{i}(0)\right)$ for $\hat{\beta}_{i}(0) \neq 0$ and for $\hat{\beta}_{i}(0)=0$, it can be chosen arbitrarily, so set $s_{i}(0)=0$. Note that as $\lambda_{2}$ is constant throughout the whole proof, the dependence of $\beta, s$ and $t$ on $\lambda_{2}$ is suppressed for notational convenience.

In order to find $t\left(\lambda_{1}\right)$, observe that soft thresholding of $\beta(0)$ does not change the ordering of pairs $\hat{\beta}_{i}\left(\lambda_{1}\right)$ and $\hat{\beta}_{j}\left(\lambda_{1}\right)$ for those pairs for which at least one of the two is not 0 and, therefore, it is possible to define $t_{i, j}\left(\lambda_{1}\right)=t_{i, j}(0)$. If $\hat{\beta}_{i}\left(\lambda_{1}\right)=\hat{\beta}_{j}\left(\lambda_{1}\right)=0$, then $t_{i, j}$ can be chosen arbitrarily in $[-1,1]$ and, therefore, let $t_{i, j}\left(\lambda_{1}\right)=t_{i, j}(0)$. Thus, without violating restrictions on $t_{i, j}, t\left(\lambda_{1}\right)=t(0)$ for all $\lambda_{1}>0 . s\left(\lambda_{1}\right)$ will be chosen appropriately below so that the subgradient equations hold.

Now insert $\hat{\beta}_{i}\left(\lambda_{1}\right)=\operatorname{sign}\left(\hat{\beta}_{i}(0)\right)\left(\left|\hat{\beta}_{i}(0)\right|-\lambda_{1}\right)^{+}$into the subgradient equations. For $\lambda_{1}>0$, look at 2 cases:

CASE 1. $\left|\hat{\beta}_{i}(0)\right| \geq \lambda_{1}$

$$
\begin{aligned}
g_{i}\left(\lambda_{1}\right)= & -y_{i}+\hat{\beta}_{i}(0)-\lambda_{1} s_{i}(0)+\lambda_{1} s_{i}\left(\lambda_{1}\right) \\
& +\sum_{j:(i, j) \in \mathcal{C}} \lambda_{2} t_{i, j}\left(\lambda_{1}\right)-\sum_{j:(j, i) \in \mathcal{C}} \lambda_{j, i} t_{j, i}\left(\lambda_{1}\right) \\
= & -y_{i}+\hat{\beta}_{i}(0)+\sum_{j:(i, j) \in \mathcal{C}} \lambda_{2} t_{i, j}(0)-\sum_{j:(j, i) \in \mathcal{C}} \lambda_{j, i} t_{j, i}(0)=0
\end{aligned}
$$

by setting $s_{i}\left(\lambda_{1}\right)=s_{i}(0)$, using the definition of $t\left(\lambda_{1}\right)$ and noting that $\hat{\beta}(0)$ was assumed to be a solution.

CASE 2. $\left|\hat{\beta}_{i}(0)\right|<\lambda_{1}$. Here, $\hat{\beta}\left(\lambda_{1}\right)=0$ and, therefore,

$$
\begin{aligned}
g_{i}\left(\lambda_{1}\right) & =-y_{i}+\lambda_{1} s_{i}\left(\lambda_{1}\right)+\sum_{j:(i, j) \in \mathcal{C}} \lambda_{2} t_{i, j}\left(\lambda_{1}\right)-\sum_{j:(j, i) \in \mathcal{C}} \lambda_{j, i} t_{j, i}\left(\lambda_{1}\right) \\
& =-y_{i}+\hat{\beta}_{i}(0)+\sum_{j:(i, j) \in \mathcal{C}} \lambda_{2} t_{i, j}(0)-\sum_{j:(j, i) \in \mathcal{C}} \lambda_{j, i} t_{j, i}(0)=0
\end{aligned}
$$

by choosing $s_{i}\left(\lambda_{1}\right)=\hat{\beta}_{i}(0) / \lambda_{1} \in[-1,1]$ and again using that $\hat{\beta}(0)$ is optimal.
As the subgradient equations hold for every $\lambda_{1}>0$, soft thresholding gives the solution. Note that we have assumed that $\lambda_{i, j}=\lambda_{2}$, but this proof works for any fixed values $\lambda_{i, j}$.

Using this theorem, it is possible to make a more general statement.
Proposition A.1. Let $\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)$ be a minimum of $f(\beta)$ for $\left(\lambda_{1}, \lambda_{2}\right)$. Then the solution for the parameters $\left(\lambda_{1}^{\prime}, \lambda_{2}\right)$ with $\lambda_{1}^{\prime}>\lambda_{1}$ is a soft thresholding of $\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)$, that is,

$$
\hat{\beta}\left(\lambda_{1}^{\prime}, \lambda_{2}\right)=\operatorname{sign}\left(\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)\right)\left(\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)-\left(\lambda_{1}^{\prime}-\lambda_{2}\right)\right)^{+} .
$$

Proof. As a solution for minimizing $f(\beta)$ exists and is unique for all $\lambda_{1}, \lambda_{2} \geq 0$, the solution $\hat{\beta}\left(0, \lambda_{2}\right)$ exists and $\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)$ as well as $\hat{\beta}\left(\lambda_{1}^{\prime}, \lambda_{2}\right)$ are softthresholded versions of it using the previous theorem. Therefore,

$$
\begin{aligned}
& \hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)=\operatorname{sign}\left(\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right)\left(\left|\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right|-\lambda_{1}\right)^{+}, \\
& \hat{\beta}_{i}\left(\lambda_{1}^{\prime}, \lambda_{2}\right)=\operatorname{sign}\left(\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right)\left(\left|\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right|-\lambda_{1}^{\prime}\right)^{+},
\end{aligned}
$$

for $i=1, \ldots, n$. If $\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)=0$, then also $\hat{\beta}_{i}\left(\lambda_{1}^{\prime}, \lambda_{2}\right)=0$. For $\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)>0$, the soft-thresholding implies that the sign did not change and, thus, $\operatorname{sign}\left(\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)\right)=$ $\operatorname{sign}\left(\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right)$. It then follows

$$
\begin{aligned}
\hat{\beta}_{i}\left(\lambda_{1}^{\prime}, \lambda_{2}\right) & =\operatorname{sign}\left(\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right)\left(\left|\hat{\beta}_{i}\left(0, \lambda_{2}\right)\right|-\lambda_{1}^{\prime}\right)^{+} \\
& =\operatorname{sign}\left(\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)\right)\left(\left|\hat{\beta}_{i}\left(\lambda_{1}, \lambda_{2}\right)\right|-\left(\lambda_{1}^{\prime}-\lambda_{1}\right)\right)^{+} .
\end{aligned}
$$

Therefore, $\hat{\beta}\left(\lambda_{1}^{\prime}, \lambda_{2}\right)$ is a soft-thresholded version of $\hat{\beta}\left(\lambda_{1}, \lambda_{2}\right)$.
Quadratic programming solution for the two-dimensional fused lasso. Here we outline the solution to the two-dimensional fused lasso using the general purpose sqopt package of Gill, Murray and Saunders (1999).

Let $\beta_{i i^{\prime}}=\beta_{i i^{\prime}}^{+}-\beta_{i i^{\prime}}^{-}$with $\beta_{i i^{\prime}}^{+}, \beta_{i i^{\prime}}^{-} \geq 0$. Define $\theta_{i i^{\prime}}^{h}=\beta_{i, i^{\prime}}-\beta_{i-1, i^{\prime}}$ for $i>1$, $\theta_{i i^{\prime}}^{v}=\beta_{i, i^{\prime}}-\beta_{i, i^{\prime}-1}$ for $i^{\prime}>1$, and $\theta_{11}=\beta_{11}$. Let $\theta_{i i^{\prime}}^{h}=\theta_{i i^{\prime}}^{h+}-\theta_{i i^{\prime}}^{h-}$ with $\theta_{i^{\prime}}^{h+}, \theta_{i^{\prime}}^{h-} \geq$ 0 , and similarly for $\theta_{i i^{\prime}}^{v}$. We string out each set of parameters into one long vector, starting with the 11 entry in the top left, and going across each row.

Let $L_{1}$ and $L_{2}$ be the $n \times n$ matrices that compute horizontal and vertical differences. Let $e$ be a column $n$-vector of ones, and $I$ be the $n \times n$ identity matrix. Then the constraints can be written as

$$
\left(\begin{array}{c}
-a_{0}  \tag{33}\\
0 \\
0 \\
0
\end{array}\right) \leq\left(\begin{array}{ccccccc}
L_{1} & 0 & 0 & -I & I & 0 & 0 \\
L_{2} & 0 & 0 & 0 & 0 & -I & I \\
I & -I & I & 0 & 0 & 0 & 0 \\
0 & e^{T} & e^{T} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & e^{T} & e^{T} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & e_{0}^{T} & e_{0}^{T}
\end{array}\right)\left(\begin{array}{c}
\beta \\
\beta^{+} \\
\beta^{-} \\
\theta^{h+} \\
\theta^{h-} \\
\theta^{v+} \\
\theta^{v-}
\end{array}\right) \leq\left(\begin{array}{c}
a_{0} \\
0 \\
0 \\
s_{1} \\
s_{2} \\
s_{3}
\end{array}\right)
$$

Here $a_{0}=(\infty, 0,0 \ldots 0)$. Since $\beta_{1 i^{\prime}}=\theta_{1 i^{\prime}}$, setting its bounds at $\pm \infty$ avoids a "double" penalty for $\left|\beta_{1 i^{\prime}}\right|$ and similarly for $\beta_{1 i^{\prime}}$. Similarly, $e_{0}$ equals $e$, with the first component set to zero.

Pathwise coordinate optimization for the general one-dimensional fused lasso. This algorithm has exactly the same form as that for the fused lasso signal approximator given earlier. The form of the basic equations is all that changes.

Equation (24) becomes

$$
\begin{align*}
\frac{\partial f(\beta)}{\partial \beta_{j}}= & -\sum_{i=1}^{n}\left[x_{i j}\left(y_{i}-\sum_{k \neq j}^{p} x_{i k} \tilde{\beta}_{k}-x_{i j} \beta_{j}\right)\right] \\
& +\lambda_{1} \cdot \operatorname{sign}\left(\beta_{j}\right)-\lambda_{2} \cdot \operatorname{sign}\left(\tilde{\beta}_{j+1}-\beta_{j}\right)  \tag{34}\\
& +\lambda_{2} \cdot \operatorname{sign}\left(\beta_{j}-\tilde{\beta}_{j-1}\right)
\end{align*}
$$

assuming that $\beta_{j} \notin\left\{0, \tilde{\beta}_{j-1}, \tilde{\beta}_{j+1}\right\}$.
Similarly, expression (25) becomes

$$
\begin{align*}
\frac{\partial f(\beta)}{\partial \gamma}= & -\sum_{i=1}^{n} z_{i}\left[y_{i}-\sum_{k \notin\{j, j+1\}}^{p} x_{i k} \tilde{\beta}_{k}-z_{i} \gamma\right] \\
& +2 \lambda_{1} \cdot \operatorname{sign}(\gamma)-\lambda_{2} \cdot \operatorname{sign}\left(\tilde{\beta}_{j+2}-\gamma\right)  \tag{35}\\
& +\lambda_{2} \cdot \operatorname{sign}\left(\gamma-\tilde{\beta}_{j-1}\right)
\end{align*}
$$

where $z_{i}=x_{i j}+x_{i, j+1}$. If the optimal value for $\gamma$ decreases the criterion, we accept the move setting $\beta_{j}=\beta_{j-1}=\gamma$.

Pathwise coordinate optimization for two-dimensional fused lasso signal approximator. Consider a set (grid) $G$ of pixels $p=(i, j)$, with $1 \leq i \leq n_{1}$, $1 \leq j \leq n_{2}$. Associated with each pixel is a data value $y_{p}=y_{i j}$. The goal is to obtain smoothed values $\hat{\beta}_{p}=\hat{\beta}_{i j}$ for the pixels that jointly solve

$$
\begin{align*}
& \min _{\left\{\beta_{i j}\right\}} \frac{1}{2} \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}}\left(y_{i j}-\beta_{i j}\right)^{2}+\lambda_{1} \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}}\left|\beta_{i j}\right| \\
& \quad+\lambda_{2} \sum_{i=2}^{n_{1}} \sum_{j=2}^{n_{2}}\left|\beta_{i j}-\beta_{i-1, j}\right|+\sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}}\left|\beta_{i j}-\beta_{i, j-1}\right| . \tag{36}
\end{align*}
$$

Defining the distance between two pixels $p=(i, j)$ and $p^{\prime}=\left(i^{\prime}, j^{\prime}\right)$ as $d\left(p, p^{\prime}\right)=$ $\left|i-i^{\prime}\right|+\left|j-j^{\prime}\right|$, (36) can be expressed as

$$
\begin{align*}
\min _{\left\{\beta_{p}\right\}} & \frac{1}{2} \sum_{p \in G}\left(y_{p}-\beta_{p}\right)^{2} \\
& +\lambda_{1} \sum_{p \in G}\left|\beta_{p}\right|+\left(\lambda_{2} / 2\right) \sum_{d\left(p, p^{\prime}\right)=1}\left|\beta_{p}-\beta_{p^{\prime}}\right| \tag{37}
\end{align*}
$$

Consider a partition of $G$ into $K$ contiguous groups $\left\{G_{k}\right\}, G=\bigcup G_{k}$ and $G_{k} \cap$ $G_{k^{\prime}}=0, k \neq k^{\prime}$. A group $G_{k}$ is contiguous if any $p \in G_{k}$ can be reached from any other $p^{\prime} \in G_{k}$ by a sequence of distance-one steps within $G_{k}$. Define the distance between two groups $G_{k}$ and $G_{k^{\prime}}$ as

$$
D\left(k, k^{\prime}\right)=\min _{\substack{p \in G_{k} \\ p^{\prime} \in G_{k^{\prime}}}} d\left(p, p^{\prime}\right)
$$

Suppose one seeks the solution to (37) under the constraints that all pixels in the same group have the same parameter value. That is, for each $G_{k},\left\{\beta_{p}=\gamma_{k}\right\}_{p \in G_{k}}$. The corresponding optimal group parameter values $\hat{\gamma}_{k}$ are the solution to the unconstrained problem

$$
\begin{align*}
& \min _{\left\{\gamma_{k}\right\}} \frac{1}{2} \sum_{k=1}^{K} N_{k}\left(\bar{y}_{k}-\gamma_{k}\right)^{2}  \tag{38}\\
& \quad+\lambda_{1} \sum_{k=1}^{K} N_{k}\left|\gamma_{k}\right|+\left(\lambda_{2} / 2\right) \sum_{D\left(k, k^{\prime}\right)=1} w_{k k^{\prime}}\left|\gamma_{k}-\gamma_{k^{\prime}}\right|
\end{align*}
$$

where $N_{k}$ is the number of pixels in $G_{k}, \bar{y}_{k}$ is the mean of the pixel data values in $G_{k}$, and

$$
w_{k k^{\prime}}=\sum_{p \in G_{k}} \sum_{p^{\prime} \in G_{k^{\prime}}} I\left[d\left(p, p^{\prime}\right)=1\right]
$$

Note that (38) is equivalent to (37) when all groups contain only one pixel.
Further, suppose that for a given partition one wishes to obtain the solution to (38) with the additional constrain that two adjacent groups $G_{l}$ and $G_{l^{\prime}}, D\left(l, l^{\prime}\right)=1$, have the same parameter value $\gamma_{m}$; that is, $\gamma_{l}=\gamma_{l^{\prime}}=\gamma_{m}$, or equivalently, $\left\{\beta_{p}=\right.$ $\left.\gamma_{m}\right\}_{p \in G_{l} \cup G_{l^{\prime}}}$. This can be accomplished by deleting groups $l$ and $l^{\prime}$ from the sum in (38) and adding the corresponding "fused" group $G_{m}=G_{l} \cup G_{l^{\prime}}$, with $N_{m}=$ $N_{l}+N_{l^{\prime}}, \bar{y}_{m}=\left(N_{l} y_{l}+N_{l^{\prime}} y_{l^{\prime}}\right) / N_{m}$,

$$
\begin{align*}
& \left\{G_{k^{\prime}}\right\}_{D\left(m, k^{\prime}\right)=1}  \tag{39}\\
& \quad=\left\{G_{k^{\prime}}\right\}_{D\left(l, k^{\prime}\right)=1} \cup\left\{G_{k^{\prime}}\right\}_{D\left(l^{\prime}, k^{\prime}\right)=1}-G_{l}-G_{l^{\prime}},
\end{align*}
$$

and $w_{m k^{\prime}}=w_{l k^{\prime}}+w_{l^{\prime} k^{\prime}}$.
The strategy for solving (36) is based on (38). As with FLSA (Section 3), there are three basic operations: descent, fusion and smoothing. For a fixed value of $\lambda_{1}$, we start at $\lambda_{2}=0$ and $n_{1} \cdot n_{2}$ groups each consisting of a single pixel. The initial $\lambda_{2}=0$ solution of each $\gamma_{k}$ is obtained by soft-thresholding $\gamma_{k}=S\left(\bar{y}_{k}, \lambda_{1}\right)$ as in (3). Starting at this solution, the value of $\lambda_{2}$ is incremented by a small amount $\lambda_{2} \leftarrow$ $\lambda_{2}+\delta$. Beginning with $\gamma_{1}$, the descent operation solves (38) for each $\gamma_{k}$ holding all other $\left\{\gamma_{k^{\prime}}\right\}_{k^{\prime} \neq k}$ at their current values. The derivative of the criterion in (38) with respect to $\gamma_{k}$ is piecewise linear with breaks at $0,\left\{\gamma_{k^{\prime}}\right\}_{D\left(k, k^{\prime}\right)=1}$. The solution for $\gamma_{k}$ is thus obtained in the same manner as that for the one-dimensional problem described in Section 3. If this solution fails to change the current value of $\gamma_{k}$, successive provisional fusions of $G_{k}$ with each $G_{k^{\prime}}$ for which $D\left(k, k^{\prime}\right)=1$ are considered, and the solution for the corresponding fused parameter $\gamma_{m}$ is obtained. The derivative of the criterion in (38) with respect to $\gamma_{m}$ is piecewise linear with breaks at $0,\left\{\gamma_{k^{\prime}}\right\}_{D\left(m, k^{\prime}\right)=1}$ (39). If any of these fused solutions for $\gamma_{m}$ improves the criterion, one provisionally sets $\gamma_{k}=\gamma_{k^{\prime}}=\gamma_{m}$. If not, the current value of $\gamma_{k}$ remains unchanged.

One then applies the descent/fusion strategy to the next parameter, $k \leftarrow k+1$, and so on until a complete pass over all parameters $\left\{\gamma_{k}\right\}$ has been made. These passes (cycles) are then repeated until one complete pass fails to change any parameter value, at which point the solution for the current value of $\lambda_{2}$ has been reached.

At this point each current group $G_{k}$ is permanently fused (merged) with all groups $G_{k^{\prime}}$ for which $\gamma_{k}=\gamma_{k^{\prime}}, \gamma_{k^{\prime}} \neq 0$, and $D\left(k, k^{\prime}\right)=1$, producing a new criterion (38) with potentially fewer groups. The value of $\lambda_{2}$ is then further incremented $\lambda_{2} \leftarrow \lambda_{2}+\delta$ and the above process is repeated starting at the solution for the previous $\lambda_{2}$ value. This continues until a specified maximum value of $\lambda_{2}$ has been reached or until only one group remains.

Acknowledgments. We thank Anita van der Kooij for informing us about her work on the elastic net, Michael Saunders and Guenther Walther for helpful discussions, and Balasubramanian Narasimhan for valuable help with our software.

A special thanks to Stephen Boyd for telling us about the subgradient form (28). While writing this paper, we learned of concurrent, independent work on coordinate optimization for the lasso and other convex problems by Ken Lange and Tongtong Wu. We thank the Editors and two referees for comments led to substantial improvements in the manuscript.

## REFERENCES

Bertsekas, D. (1999). Nonlinear Programming. Athena Scientific.
Breiman, L. (1995). Better subset selection using the nonnegative garrote. Technometrics 37 738754. MR1365720

Chen, S. S., Donoho, D. L. and Saunders, M. A. (1998). Atomic decomposition by basis pursuit. SIAM J. Sci. Comput. 33-61. MR1639094
Daubechies, I., Defrise, M. and De Mol, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. Comm. Pure Appl. Math. 57 1413-1457. MR2077704
DONOHO, D. and JOHNSTONE, I. (1995). Adapting to unknown smoothness via wavelet shrinkage. J. Amer. Statist. Assoc. 90 1200-1224. MR1379464

Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004). Least angle regression (with discussion). Ann. Statist. 32 407-499. MR2060166
Friedlander, M. and Saunders, M. (2007). Discussion of "Dantzig selector" by E. Candes and T. Tao. Ann. Statist. 35 2385-2391.

Fu, W. J. (1998). Penalized regressions: The bridge versus the lasso. J. Comput. Graph. Statist. 7 397-416. MR1646710
Gill, P., Murray, W. and Saunders, M. (1999). Users guide for sqopt 5.3: A fortran package for large-scale linear and quadratic programming. Technical report, Stanford Univ.
Li, Y. and Arce, G. (2004). A maximum likelihood approach to least absolute deviation regression. URASIP J. Appl. Signal Processing 2004 1762-1769. MR2131987
Osborne, M., Presnell, B. and Turlach, B. (2000). A new approach to variable selection in least squares problems. IMA J. Numer. Anal. 20 389-404. MR1773265
OwEn, A. (2006). A robust hybrid of lasso and ridge regression. Technical report, Stanford Univ.
Rudin, L. I., Osher, S. and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. Phys. D 60 259-268.
Schlegel, P. (1970). The explicit inverse of a tridiagonal matrix. Math. Comput. 24 665-665. MR0273798
Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. J. Roy. Statist. Soc. Ser. B58 267-288. MR1379242
Tibshirani, R., Saunders, M., Rosset, S., Zhu, J. and Knight, K. (2005). Sparsity and smoothness via the fused lasso. J. Roy. Statist. Soc. Ser. B 67 91-108. MR2136641
Tibshirani, R. and Wang, P. (2007). Spatial smoothing and hot spot detection for CGH data using the fused lasso. Biostatistics. Advance Access published May 18, 2007.
TSENG, P. (1988). Coordinate ascent for maximizing nondifferentiable concave functions. Technical Report LIDS-P, 1840, Massachusetts Institute of Technology, Laboratory for Information and Decision Systems.
TsENG, P. (2001). Convergence of block coordinate descent method for nondifferentiable maximation. J. Opt. Theory Appl. 109 474-494. MR1835069
VAN DER KOOIJ, A. (2007). Prediction accuracy and stability of regresssion with optimal scaling transformations. Technical report, Dept. Data Theory, Leiden Univ.
WANG, H., Li, G. and JiAng, G. (2006). Robust regression shrinkage and consistent variable selection via the lad-lasso. J. Business Econom. Statist. 11 1-6.

YUAN, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables.
J. Roy. Statist. Soc. Ser. B 68 49-67. MR2212574

Zhou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. J. Roy. Statist. Soc. Ser. B 67 301-320. MR2137327

J. Friedman<br>Department of Statistics<br>Stanford University<br>Stanford, CALIFORNIA 94305<br>USA<br>E-MAIL: jhf@stanford.edu<br>H. HöfLING<br>Department of Statistics<br>Stanford University<br>Stanford, CALIFORNIA 94305<br>USA<br>E-MAIL: hhoeflin@gmail.com


[^0]:    Received May 2007; revised August 2007.
    ${ }^{1}$ Supported in part by NSF Grant DMS-97-64431.
    ${ }^{2}$ Supported in part by NSF Grant DMS-05-50676 and NIH Grant 2R01 CA 72028-07.
    ${ }^{3}$ Supported by an Albion Walter Hewlett Stanford Graduate Fellowship.
    ${ }^{4}$ Supported in part by NSF Grant DMS-99-71405 and NIH Contract N01-HV-28183.
    Key words and phrases. Coordinate descent, lasso, convex optimization.

