

A homotopy algorithm for the quantile regression lasso and related piecewise linear problems

M.R. Osborne* B.A. Turlach†

June 18, 2010

Abstract

We show that the homotopy algorithm of Osborne, Presnell and Turlach [10], which has proved such an effective optimal path following method for implementing Tibshirani's "lasso" for variable selection in least squares estimation problems, can be extended to polyhedral objectives in examples such as the quantile regression lasso. The new algorithm introduces the novel feature that it requires two homotopy sequences involving continuation steps with respect to both the constraint bound and the Lagrange multiplier. Performance is illustrated by application to several standard data sets, and these results are compared to calculations made with the original lasso homotopy program. This permits an assessment of the computational complexity to be made both for the new method and for the closely related linear programming post optimality procedures which generate essentially identical solution trajectories. This comparison strongly favours the least squares selection method. However, the new method still provides an effective computational procedure, plus it has distinct implementation advantages over the linear programming approaches to the polyhedral objective problem. The computational difficulty is explained and the problem that needs to be resolved in order to improve performance identified.

*Mathematical Sciences Institute, Australian National University

†School of Mathematics and Statistics, University of Western Australia

1 Introduction

The problem motivating this investigation is the classic variable selection problem of determining a subset of the parameters $\boldsymbol{\beta} \in R^p$ in the linear model

$$\mathbf{r} = \mathbf{y} - X\boldsymbol{\beta}, \quad (1)$$

where $\mathbf{r} \in R^n$ is the residual vector, $\mathbf{y} \in R^n$ is the vector of observations, and $X : R^p \rightarrow R^n$ is the design matrix which is assumed to have full (typically column) rank, in order to provide an adequate but parsimonious representation of the problem data. If an l_1 inequality constraint

$$\sum_{i=1}^p |\beta_i| \leq \kappa \quad (2)$$

is imposed on $\boldsymbol{\beta}$ in the least squares regression formulation of the estimation problem then the number of nonzero components of $\boldsymbol{\beta}$ in the optimum solution increases as the constraint bound is increased away from zero. The problem eventually becomes unconstrained for large enough values of the bound, say $\kappa \geq \kappa_{max}$, so that all parameters can be expected to be involved in the solution when $p \leq n$. Tibshirani suggested in [12] that study of the sets of non-zero parameter estimates as a function of the constraint bound provided a basis for variable selection and gave the acronym “lasso” to the resulting process. Improved algorithms were developed in [10]. In particular, this paper introduced a “homotopy” algorithm which follows the full piecewise linear optimal solution path as a function of the constraint bound. Subsequently significant evidence has accumulated that this algorithm is particularly efficient in many circumstances, frequently taking at most few more steps than the number of variables selected. When this is true then it rivals the cost of an unconstrained least squares algorithm for the full set of variables while providing considerable additional information. This has led to further applications. For example, it has proved distinctly effective in applications in compressed sensing [3]. Generalisations have been made to variable selection in problems with multiple objectives where the form of constraint must be chosen appropriately [13]. The lasso has been used with more general likelihoods by approximating them by piecewise quadratic objectives with continuous first derivative [15], and for robust variable selection using the Huber-M estimator which involves a mixed piecewise linear, quadratic objective [11]. Both these problems have sufficient smoothness to ensure that the necessary conditions for optimality vary continuously at points where the objective function changes its structure. However, this property is not shared

by the homotopy equations in either case so that an extra layer of algorithmic complexity is added. Different selection properties can be addressed by varying the form of constraint. For example, [1] recommend using what is in effect the signed rank objective [5] as a constraint. They claim this has an advantage in suggesting the development of cluster variables as an aid to interpretation when the predictor variables are both relatively highly correlated and make similar contributions to the response. A different suggestion for attacking similar problems is the so-called “elastic net” [17].

The other class of applications that has attracted recent attention is that corresponding to the relaxation of the smoothness of the objective. Examples include the training of support vector machines [16], [14], the piecewise linear objectives in quantile regression [7], [8], and [18] where the form of constraint used in [13] is used to develop a simultaneous variable selection procedure for simultaneous quantile regressions. A suggested compressed sensing application involves the minimization of the l_1 norm of the parameters subject to a maximum norm constraint on the components of the residual vector [2]. This fits the pattern developed here because it follows from the necessary conditions that the roles of the constraint and objective can be reversed.

Here a general homotopy type lasso algorithm for variable selection is developed in a form which is applicable to piecewise linear objective functions of quite general type. There is one striking difference in the properties of the optimal homotopy trajectory between the case when the objective is at least once continuously differentiable and the non-smooth case when the objective is piecewise linear [11]. In the former case the Lagrange multiplier for the l_1 constraint is a piecewise linear, continuous function of the constraint bound κ with the characteristic property that it decreases steadily from its initial positive value at $\kappa = 0$ to 0 for κ large enough. In contrast, the corresponding Lagrange multiplier associated with a piecewise linear objective is a decreasing step function of κ with jumps at the points where the set of non-zero components of β changes, and it is necessary to include an explicit multiplier update phase as these jumps have to be determined as part of the computation. The homotopy algorithm can be developed quite generally for polyhedral convex objective functions given generic forms for the subdifferentials of the objective and the polyhedral constraint [9]. However, here we restrict ourselves to the case of an l_1 objective subject to the standard l_1 form of constraint used in the lasso. This objective corresponds to the special case of parameter $\tau = .5$ in the quantile regression problem

$$\min_{\beta} \sum_{i=1}^n (1 - \tau)(-r_i)_+ + \tau(r_i)_+, \quad 0 < \tau < 1, \quad (3)$$

where the r_i are residuals in a linear model fit. As indicated above, the

application of the lasso to this objective corresponds to one of the more popular applications of the lasso variable selection technology.

The unconstrained quantile regression application already hints that computational complexity is potentially an important consideration in designing algorithms for piecewise linear objectives in the constrained problems. This arises because complexity is already a problem in simplicial algorithms for the l_1 estimation problem where it proves important to incorporate an effective line-search in order to make large correction steps especially in the initial stages of the computation [9]. If this is not done then a sequence of small steps is made as the residual vector adjusts to adapt itself to the form required by the necessary conditions. But every residual changing sign triggers a point of nondifferentiability of the objective. Typically $O(n)$ changes at least are expected in this adaptive process given general initial conditions. A form of this problem occurs also in the homotopy algorithm. The reason for this is that frequently a succession of multiplier update steps occur with the successive parameter estimates contained exactly in the same orthant of a fixed subspace. The corresponding increments in these estimates correspond to descent steps for the objective function in this subspace. However, the objective minimum in this subspace may not be on the homotopy path (may not satisfy the homotopy necessary conditions) so the basic l_1 strategy requires modification.

Remark 1 *The novelty in our approach lies in the manner in which the non-smooth necessary conditions are used explicitly in the basic algorithm structure. This involves switching between the constrained and Lagrangian forms of the problem in order to continue with respect to the constraint bound (increase) and then to update (decrease) the Lagrangian parameter. In particular, decisions about adding and subtracting variables are based on sub-gradient vector components attaining or departing from bounds. This basic pattern generalises to necessary conditions for more complicated piecewise linear structures. The algorithm described here is equivalent to the use of simplex algorithm post-optimality techniques applied to a linear programming formulation of the problem [6], [8]. It has significant advantages both in displaying problem structure and in avoiding the introduction of the slack variables and auxiliary positivity constraints required in the linear programming formulations. This enables the direct use of the design matrix as the basic data structure in the homotopy approach.*

The lasso algorithm is developed in the next section. This is followed by a discussion of numerical experience together with an indication of the source of the additional work required in the non-smooth objective case. The

detailed algebraic calculations to justify the homotopy algorithm have been postponed to the appendix. We describe the current state of our investigation as a “proof of concept” exercise. This means several issues have not been explored in detail. These include:

1. The numerical work is thorough but has not been exhaustively optimised. For example, economised updating and downdating of matrix factorizations has not been implemented. There could be scope for this and it could involve, for example, keeping two lines of factorization for the two types of continuation.
2. Our algorithm requires these factored matrices to be nonsingular. This is a requirement that the current number of variables selected k satisfies $k \leq \text{rank } X \leq \min p, n$. Here this condition is simplified to $\text{rank } X = p$, a condition frequently encountered in practice. However, the more general condition can be important. For example, our original application of the least squares lasso was to the problem of selecting an effective set of knots for a regression spline {reference?}. Here there is no a priori bound on the total number of knots p suggested initially. In such cases use of an extra problem specific condition to ensure the non-singularity requirement is satisfied is recommended .
3. Whenever a change of basis is made in linear programming post-optimality studies (also referred to as parametric linear programming) then a non-unique optimum has occurred in the linear programming formulation. In the typical case, the intersection between the linear program objective and the feasible region is one dimensional in this degeneracy. This is the post-optimality interpretation of Proposition 3 in [11]. This particular degeneracy in the linear program provides the homotopy step with respect to the constraint bound in our formulation. However, there are other ways in which degeneracy can and does occur. Our treatments are expedient and are discussed in subsection 3.2 as one aspect of the computations.
4. There is no doubt of the utility of extracting an intercept term and standardising the design accordingly in ordinary least squares estimation. This is not quite so clear in the least squares lasso as the l_1 norm does not respect orthogonality. It becomes less clear when the objective is polyhedral, and further complicated if a variable rather than its absolute value is entered into the constraint term. The calculations presented here took the simplest option that permitted direct comparison with the original lasso homotopy program [10]. This permitted

the design to be augmented optionally by a column of 1's. Columns of the design and of the response were then scaled to have length 1. The augmented design has been used in the quoted numerical results.

2 Lasso for the l_1 objective

2.1 Necessary conditions

The basic idea of the homotopy algorithm is to follow the evolution of the parameters β as piecewise linear functions of the constraint bound κ . Two forms of the problem prove important in developing the homotopy trajectory. These are

Constrained form This is used to follow β as a piecewise linear function of the constraint bound κ . This problem is

$$\min_{\beta} \sum_{i=1}^n |r_i|, \quad \sum_{i=1}^p |\beta_i| \leq \kappa. \quad (4)$$

If $|y_i| > 0$, $i = 1, 2, \dots, n$, then the problem has the solution $\beta = 0$ when $\kappa = 0$. Increasing κ then permits components of β to move from 0 (become active), typically one at a time. If κ is large enough then β is essentially unconstrained so the problem reduces to an unconstrained l_1 minimization problem. One consequence is that the associated Lagrange multiplier is $\lambda = 0$.

Lagrangian form This is used to update the subgradient vector components in the necessary conditions as a parametric function of the Lagrange multiplier λ with κ fixed. This update is performed whenever there is a change in the set of active β components. The problem Lagrangian is

$$\mathcal{L}(\beta, \lambda) = \sum_{i=1}^n |r_i| + \lambda \left\{ \sum_{i=1}^p |\beta_i| - \kappa \right\}.$$

It is convex if $\lambda \geq 0$. The form used here follows from the necessary condition computed with respect to β expressed in subgradient form

$$0 \in \partial_{\beta} \mathcal{L}(\beta, \lambda) = \partial_{\beta} \sum_{i=1}^n |r_i| + \lambda \partial_{\beta} \sum_{i=1}^p |\beta_i|. \quad (5)$$

The key to its utility is a consequence of κ not appearing explicitly in this expression. Setting the initial value of λ corresponding to $\kappa = 0$ is described in subsection 2.4.

In non-degenerate cases it will be possible to separate zero and non-zero values of both \mathbf{r} and $\boldsymbol{\beta}$ into complementary classes and to take uniquely specified actions in the critical situations that occur when these classes have to be re-defined. Pointer (index) sets are introduced to specify these separations. Let

$$\sigma = \{i : r_i = 0\}, \quad \psi = \{i : \beta_i \neq 0\}, \quad (6)$$

and denote the set complements in $\{1, 2, \dots, n\}$ and $\{1, 2, \dots, p\}$ respectively by

$$\sigma^c = \{i : r_i \neq 0\}, \quad \psi^c = \{i : \beta_i = 0\}. \quad (7)$$

Orderings of the elements in σ and ψ are assumed. Specifically the pointer in the i 'th position in σ is referenced as $\sigma(i)$. In similar fashion the pointer in position j in ψ^c is referenced as $\psi^c(j)$. The computation must provide an implementation of these orderings. In addition define permutation matrices $P_\sigma : R^n \rightarrow R^n$ and $Q_\psi : R^p \rightarrow R^p$ by

$$P_\sigma \mathbf{r} = \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix}, \quad \begin{cases} (\mathbf{r}_1)_i = r_{\sigma(i)} \neq 0, & i = 1, 2, \dots, n - |\sigma|, \\ (\mathbf{r}_2)_i = r_{\sigma(i)} = 0, & i = 1, 2, \dots, |\sigma| \end{cases},$$

$$Q_\psi \boldsymbol{\beta} = \begin{bmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{bmatrix}, \quad \begin{cases} (\boldsymbol{\beta}_1)_i = \beta_{\psi(i)} \neq 0, & i = 1, 2, \dots, |\psi|, \\ (\boldsymbol{\beta}_2)_i = \beta_{\psi^c(i)} = 0, & i = 1, 2, \dots, p - |\psi| \end{cases}.$$

To specify necessary conditions set

$$P_\sigma X Q_\psi^T = \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix}, \quad (8)$$

$$P_\sigma \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix}. \quad (9)$$

Let $[\boldsymbol{\theta}_\sigma^T \quad \mathbf{v}_\sigma^T] \in \partial_r \|P_\sigma \mathbf{r}\|_1$, and $[\boldsymbol{\theta}_\psi^T \quad \mathbf{u}_\psi^T] \in \partial_\beta \|Q_\psi \boldsymbol{\beta}\|_1$, where $\boldsymbol{\theta}_\sigma$ and $\boldsymbol{\theta}_\psi$ are the vectors of signs of the components of \mathbf{r}_1 and $\boldsymbol{\beta}_1$ respectively, be subgradient vectors associated with the corresponding l_1 norm functions. Then the necessary conditions (5) are

$$[\boldsymbol{\theta}_\sigma^T \quad \mathbf{v}_\sigma^T] \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix} = \lambda [\boldsymbol{\theta}_\psi^T \quad \mathbf{u}_\psi^T], \quad \lambda \geq 0, \quad (10)$$

$$-1 \leq v_i \leq 1, \quad i = 1, 2, \dots, |\sigma|, \quad (11)$$

$$-1 \leq u_i \leq 1, \quad i = 1, 2, \dots, |\psi^c|, \quad (12)$$

$$\|\mathbf{r}\|_1 = [\boldsymbol{\theta}_\sigma^T \quad \mathbf{v}_\sigma^T] P_\sigma \mathbf{r} = \boldsymbol{\theta}_\sigma^T \mathbf{r}_1, \quad (13)$$

$$\|\boldsymbol{\beta}\|_1 = [\boldsymbol{\theta}_\psi^T \quad \mathbf{u}_\psi^T] Q_\psi \boldsymbol{\beta} = \boldsymbol{\theta}_\psi^T \boldsymbol{\beta}_1 \leq \kappa. \quad (14)$$

2.2 Varying κ

The evolution of the set of active parameter estimates as functions of κ follows standard homotopy procedures. Consider $\kappa = \kappa^* > 0$, κ^* in an open interval such that the optimal β is determined by a fixed set of zero residuals $r_i = 0$, $i \in \sigma$. These can be identified as l_1 necessary conditions except that here the active bound constraint picks up a degree of freedom so that $|\sigma| = |\psi| - 1$. The resulting equations on this interval have the form

$$\theta_\psi^T \beta_1 = \kappa, \quad l_1 \text{ norm condition}, \quad (15)$$

$$X_{21} \beta_1 = \mathbf{y}_2, \quad \text{zero residual conditions.} \quad (16)$$

Differentiating with respect to κ gives

$$\theta_\psi^T \frac{d\beta_1}{d\kappa} = 1, \quad (17)$$

$$X_{21} \frac{d\beta_1}{d\kappa} = 0, \quad (18)$$

$$\Rightarrow \frac{d\beta_1}{d\kappa} = \begin{bmatrix} \theta_\psi^T \\ X_{21} \end{bmatrix}^{-1} \mathbf{e}_1. \quad (19)$$

This shows that $\frac{d\beta_1}{d\kappa}$ is constant in a neighbourhood of κ^* so that locally β_1 is a linear function of κ . Now it follows that $\lambda = \lambda_0$ is constant on this interval of linear dependence. To show this consider the partitioning of the necessary conditions (10) which gives that

$$\theta_\sigma^T X_{11} + \mathbf{v}_\sigma^T X_{21} = \lambda \theta_\psi^T.$$

Differentiating with respect to κ gives

$$\frac{d\mathbf{v}_\sigma^T}{d\kappa} X_{21} = \frac{d\lambda}{d\kappa} \theta_\psi^T. \quad (20)$$

To show $\frac{d\lambda}{d\kappa} = 0$ multiply by $\frac{d\beta_1}{d\kappa}$ and use (18) and (17) to obtain

$$0 = \frac{d\mathbf{v}_\sigma^T}{d\kappa} X_{21} \frac{d\beta_1}{d\kappa} = \frac{d\lambda}{d\kappa}.$$

Thus λ is constant in intervals of κ with fixed zero residual sets and can have jumps only when it is necessary to vary the partitioning of the optimality conditions. It also follows that $\frac{d\mathbf{v}_\sigma}{d\kappa}$ and $\frac{d\mathbf{u}_\psi}{d\kappa}$ vanish on these intervals.

One connection between the above development and a direct optimisation approach is the following.

Lemma 2 $\frac{d\boldsymbol{\beta}}{d\kappa}$ is a direction of descent for minimising $\|\mathbf{r}\|_1$.

Proof. The calculation of the directional derivative can proceed as follows.

$$\begin{aligned} \|\mathbf{r}\|_1' \left(\boldsymbol{\beta} : \frac{d\boldsymbol{\beta}}{d\kappa} \right) &= - \sup_{\mathbf{z} \in \partial \|\mathbf{r}\|_1} \mathbf{z}^T X \frac{d\boldsymbol{\beta}}{d\kappa}, \\ &= -\boldsymbol{\theta}_\sigma^T X_{11} \frac{d\boldsymbol{\beta}_1}{d\kappa} = -\lambda \boldsymbol{\theta}_\psi^T \frac{d\boldsymbol{\beta}_1}{d\kappa}, \\ &= -\lambda \boldsymbol{\theta}_\psi^T \begin{bmatrix} X_{21} \\ \boldsymbol{\theta}_\psi \end{bmatrix}^{-1} \mathbf{e}_{|\psi|} = -\lambda < 0. \end{aligned}$$

where (18) has been used. Here $\mathbf{e}_{|\psi|}$ is the unit vector with 1 in the $|\psi|$ place.

■

A similar result is given in [8].

The homotopy algorithm proceeds by stepping in the direction defined by $\frac{d\boldsymbol{\beta}_1}{d\kappa}$. There are two possibilities for terminating this step.

- 1 A new zero residual is generated. It is assumed that the change in the characteristic set corresponds to row $\sigma^c(k)$ of X_{11} . Set $\mathbf{x}_k^T = \mathbf{e}_k^T X_{11}$, $y_k = \mathbf{e}_k^T \mathbf{y}_1$, and define r_k similarly. The update of the index set σ to take account of the new zero residual is $\sigma \leftarrow \sigma \cup \{\sigma^c(k)\}$. If σ is reordered such that $\sigma^c(k) \rightarrow \sigma(1)$ then the corresponding subgradient component is given by $v_1(\lambda_0) = \text{sign}(r_k(\kappa^*)) = \theta_k$.
- 2 A component of $\boldsymbol{\beta}_1$ vanishes before a new zero residual is reached. Let this component correspond to index $\psi(j)$. To preserve the necessary conditions it is necessary to downate the index set pointing to the $\boldsymbol{\beta}$ components, $(\psi \leftarrow \psi \setminus \{\psi(j)\})$. It is assumed that $\psi(j) \rightarrow \psi^c(1)$. The corresponding subgradient component is $u_1 = \text{sign}(\beta_j(\kappa^*)) = \theta_j$.

At this point the computation switches to consider continuation with respect to the Lagrange multiplier as parameter. Optimality as λ is reduced requires that the new subgradient components move from their bound values into the interior of their region of feasibility. This is guaranteed by the following results which assume the non-degeneracy condition that exactly one component of \mathbf{r} or $\boldsymbol{\beta}$ is involved.

Lemma 3 *Optimality is preserved in a small enough reduction $\Delta\lambda$ in λ provided $\theta_k v_1(\lambda_0 - \Delta\lambda) < 1$. This is equivalent to the following.*

$$\theta_k \frac{dv_1(\lambda_0)}{d\lambda} \geq 0. \quad (21)$$

Lemma 4 *Optimality is preserved in a small enough reduction $\Delta\lambda$ in λ provided $\theta_j u_1(\lambda_0 - \Delta\lambda) \leq 1$. This result is equivalent to*

$$\theta_j \frac{du_1(\lambda_0)}{d\lambda} \geq 0. \quad (22)$$

These results are proved in appendix A.

2.3 Varying λ

Let the current value of the Lagrange multiplier be $\lambda = \lambda_0$. Also, let the current interval of linear dependence of the parameters $\boldsymbol{\beta}$ on κ be $[\kappa_0 \leq \kappa \leq \kappa_1]$. There are two cases which can interrupt this dependence:

1. Assume the κ homotopy step is terminated by the occurrence of the new zero residual r_k , $r_k(\kappa_1) = 0$, corresponding to the first termination possibility in the κ phase. In this case the updated quantities are given by

$$\bar{X}_{21} \leftarrow \begin{bmatrix} \mathbf{x}_k^T \\ X_{21} \end{bmatrix}, \quad \bar{\mathbf{y}}_2 \leftarrow \begin{bmatrix} y_k \\ \mathbf{y}_2 \end{bmatrix}, \quad (23)$$

\bar{X}_{21} is now square and invertible as the vanishing of r_k as κ is increased requires that the corresponding row \mathbf{x}_k from X_{11} be independent of the rows of X_{21} . Now both $\boldsymbol{\beta}$ and κ are fixed by the condition

$$\bar{X}_{21} \boldsymbol{\beta}_1 = \bar{\mathbf{y}}_2.$$

2. The break in the κ homotopy is triggered by the vanishing of a $\boldsymbol{\beta}$ component. In this case it is necessary to remove the column of X_{21} corresponding to the new zero component of $\boldsymbol{\beta}$ to produce \bar{X}_{21} . Again it is square and nonsingular but now has dimension $|\sigma|$. The vector θ_ψ has its j 'th component added to \mathbf{u}_ψ .

Note that λ does not appear explicitly in either possibility. Now differentiate the repartitioned necessary conditions with respect to λ to obtain

$$\begin{bmatrix} 0 & \dots & 0 & \frac{d\mathbf{v}_\sigma^T}{d\lambda} \end{bmatrix} \begin{bmatrix} \bar{X}_{11} & \bar{X}_{12} \\ \bar{X}_{21} & \bar{X}_{22} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\theta}_\psi^T & \frac{d(\lambda \mathbf{u}_\psi^T)}{d\lambda} \end{bmatrix}.$$

Equating components gives

$$\frac{d\mathbf{v}_\sigma^T}{d\lambda} \bar{X}_{21} = \boldsymbol{\theta}_\psi^T, \quad (24)$$

$$\frac{d\mathbf{v}_\sigma^T}{d\lambda} \bar{X}_{22} = \frac{d(\lambda \mathbf{u}_\psi^T)}{d\lambda}. \quad (25)$$

Remark 5 It follows that $\frac{d\mathbf{v}_\sigma^T}{d\lambda}$ and $\frac{d(\lambda\mathbf{u}_\psi^T)}{d\lambda}$ are constant so the updates of \mathbf{v}_σ and \mathbf{u}_ψ are computed readily. However, \mathbf{u}_ψ is not itself piecewise linear in λ . A similar situation occurs in the lasso homotopy algorithm for the case of a quadratic objective function subject to an l_1 constraint [10]. Here this complicates the checking that the components of \mathbf{u}_ψ are within bound. The condition that $\lambda\mathbf{u}_\psi$ is piecewise linear can be expressed component-wise as

$$\frac{d}{d\lambda} (\lambda\mathbf{u}(\lambda)_i) = \gamma_i, \quad i = 1, 2, \dots, p,$$

where γ_i is constant independent of λ on each linear piece. It follows that

$$(\mathbf{u}_\psi)_i = \frac{\alpha_i}{\lambda} + \gamma_i,$$

where

$$\alpha_i = (\mathbf{u}_\psi(\lambda_0))_i - \gamma_i \lambda_0.$$

Thus

$$\mathbf{u}_\psi(\lambda_0 - \Delta)_i = \frac{\alpha_i}{\lambda_0 - \Delta} + \gamma_i$$

is an increasing function of $\Delta < \lambda_0$ provided $\alpha_i > 0$. This means that $\mathbf{u}_\psi(\lambda_0 - \Delta)_i$ can reach its upper bound as Δ is increased provided

$$1 = \frac{\alpha_i}{\lambda_0 - \Delta} + \gamma_i \Rightarrow \Delta = \lambda_0 - \frac{\alpha_i}{1 - \gamma_i}.$$

This requires $\alpha_i > 0 \Rightarrow \mathbf{u}_\psi(\lambda_0)_i > \gamma_i$. The lower bound is reached if

$$-1 = \frac{\alpha_i}{\lambda_0 - \Delta} + \gamma_i \Rightarrow \Delta = \lambda_0 + \frac{\alpha_i}{1 + \gamma_i}.$$

This requires $\alpha_i < 0 \Rightarrow \mathbf{u}_\psi(\lambda_0)_i < \gamma_i$.

A consequence is that the necessary conditions continue to hold as λ is reduced until either:

1. Let $u_q = \mathbf{e}_q^T \mathbf{u}_\psi$, and assume it is the component of \mathbf{u}_ψ to first reach a bound. Then $\psi \leftarrow \psi \cup \{\psi^c(q)\}$ and the increase κ phase is recommenced. The necessary conditions are maintained if the corresponding component of $\boldsymbol{\beta}$ moves away from zero with the sign of the bound. This is required in order to impose the norm condition correctly on the augmented $\boldsymbol{\beta}$.

2. Let $v_q = \mathbf{e}_q^T \mathbf{v}_\sigma$ and assume it is the component of \mathbf{v}_σ to first reach a bound. The appropriate action is to remove the corresponding index from σ , $\sigma \leftarrow \sigma \setminus \{q\}$, and restart the increase κ phase using the downdated X_{21} . Now $r_q = \mathbf{e}_q^T \mathbf{r}_2$ is no longer bound to zero. The necessary conditions will continue to hold provided its sign agrees with the sign of v_q as κ is increased.

The additional results required are as follows (see appendix A). Again a non-degeneracy condition is assumed. Here it that exactly one component of either \mathbf{v} or \mathbf{u} is involved.

Lemma 6 *Let $v_q(\lambda)$ reach first a bound as λ is decreased. If $\sigma \leftarrow \sigma \setminus \{q\}$ and the κ phase restarted then*

$$\text{sign} \left(\frac{dr_q}{d\kappa} \right) = - \text{sign} \left(\frac{dv_q}{d\lambda} \right). \quad (26)$$

Lemma 7 *Let $u_q(\lambda)$ reach first a bound as λ is decreased. If $\psi \leftarrow \psi \cup \{q\}$ and the κ phase restarted then*

$$\text{sign} \left(\frac{dx_q}{d\kappa} \right) = - \text{sign} \left(\frac{du_q}{d\lambda} \right). \quad (27)$$

2.4 Starting out

Assume $y_i \neq 0$, $i = 1, 2, \dots, n$. Then

$$\mathbf{r} = \mathbf{y}, \quad \sigma = \emptyset, \quad (\boldsymbol{\theta}_\sigma)_i = \text{sign}(y_i), \quad i = 1, 2, \dots, n,$$

when $\kappa = 0$. For simplicity, assume also there is a unique answer to

$$k = \arg \left\{ \max_i |\boldsymbol{\theta}_\sigma^T X_{*i}| \right\},$$

where X_{*i} denotes the i 'th column of X . Set

$$\lambda = |\boldsymbol{\theta}_\sigma^T X_{*k}|, \quad \psi = \{k\}, \quad \theta_\psi = \text{sign}(\boldsymbol{\theta}_\sigma^T X_{*k}), \quad (28)$$

$$\lambda \begin{bmatrix} \theta_\psi & \mathbf{u}_\psi^T \end{bmatrix} = \boldsymbol{\theta}_\sigma^T X Q_\psi^T. \quad (29)$$

Note that initially ψ is a singleton so that θ_ψ is a scalar. Equations (28) and (29) provide a set of quantities satisfying the necessary conditions for small enough κ . This permits the κ phase to be initiated from $\kappa = 0$.

3 Numerical results

3.1 Computational complexity

Results of the homotopy algorithm for several well known data sets are displayed in Table 1. The data sets are

Hald , <http://www.stat.uconn.edu/~nalini/fcilmweb/example13.html>, [4];

Iowa , <http://www.math.unm.edu/splus/node130.html>, [4];

diabetes <http://www4.stat.ncsu.edu/~boos/var.select/diabetes.tab.txt>;

housing <http://lib.stat.cmu.edu/datasets/boston> .

At the end of each step which varies κ either σ is increased or ψ is decreased, and at the end of each step that decreases λ either σ is decreased or ψ is increased. An iteration of our homotopy algorithm involves a step varying κ followed by a step varying λ . Table 1 records the numbers for each of the possible combinations of behaviour in these step pairings that occurred in reaching the unconstrained solution of the problem. Specifically let SA entries correspond to increments and SD to decrements in σ , XA to increments and XD to decrements in ψ . These results show that much of the time the basic step pairing gives entries in the SASD column. This counts descent steps that move one residual r_q away from zero in the κ step (SA) and adds a new residual r_k to the zero residual set in the λ step (SD) while leaving $\beta(\kappa)$ in the same $|\psi|$ dimensional subspace. This is illustrated for the diabetes data set in Figure 1 which shows several examples of this behaviour. Because $\frac{d\beta}{d\kappa}$ is a descent direction (Lemma 2) it follows that the algorithm is making steps that approach the l_1 minimum in the current $|\psi|$ dimensional subspace (actually a fixed orthant as a sign change in β interrupts the SASD sequence).

The smoothness of the least squares objective means that steps that update and down-date σ have no analog in the original lasso algorithm in which the usual step is one that adds a new variable (XA) - here corresponding most closely to the SAXA column in Table 1. The other possibilities include XDSD column entries which correspond to back-track steps and XDXA column entries which often indicate variables changing sign. Both these types of behaviour, which we summarise here under the heading of XD steps, are observed to occur infrequently in the original lasso formulation which has a complexity essentially dominated by variable addition steps. Results for the

	p	n	SASD	SAXA	XDXA	XDSD
Hald	4	13	8	4	1	0
Iowa	8	33	16	14	1	6
diabetes	10	442	514	23	1	13
housing	13	506	783	24	0	11

Table 1: Step counts for homotopy algorithm (R implementation)– l_1 objective




Figure 1: Homotopy algorithm illustrated on the diabetes data. The left panel shows the complete homotopy. The numbers on the right of this panel label the solution components. The right panel is a magnification of the initial part of the homotopy illustrating the large number of SASD steps that are taken.

	XA	XD
Hald	5	0
Iowa	9	0
diabetes	14	3
housing	16	2

Table 2: Step counts for homotopy algorithm – least squares objective

κ	variables at 0	subspace SASD steps
12.30	2,8,9	
13.37	8,9	40
31.79	8,9	
32.04	8,9,10	9
35.47	8,9,10	
36.35	8,9	26
45.13	8,9	
45.82	9	

Table 3: Example of backtrack step, diabetes data

original lasso corresponding to those shown above are given in Table 2. These show an order of magnitude difference in the computational complexity of the l_1 homotopy algorithm which occurs because the large number of the SASD steps turns an $O(n)$ algorithm - the complexity of an update step when n is large and p typically fixed - into at least an $O(n^2)$ algorithm. This problem arises because the process is seeking improvements by taking many “small” SASD steps towards the l_1 minimum in the current subspace. However, this process terminates whenever the homotopy “escapes” corresponding to the occurrence of a step of type SAXA, XDXA, or XDSD, and this can happen before the minimum in the current subspace is reached. This behaviour is a consequence of the role of the currently zero-valued parameter estimates pointed to by ψ^c in the homotopy necessary conditions (10) which have no counterpart in a straight l_1 minimization based on the non-zero parameter estimates alone. This is illustrated in Table 3 which shows a backtrack step occurring in the application of the piecewise linear homotopy algorithm to the diabetes data without an intercept term. The table gives the κ values at the entry and exit, where appropriate, for the subspaces encountered, together with the zero variables, and the number of SASD steps in the current subspace. This example is of interest also because after the backtrack step, which gives a sequence of descent steps in the smaller subspace, there is a return to the larger subspace left previously. It follows that large step methods based on using line search techniques developed for the l_1 minimization problem cannot be applied without modification.

The problem that the sequence of small steps solves can be formulated explicitly. The idea is to maximise κ or minimise λ in the orthant of the subspace determined by θ_ψ subject to constraints which ensure that the homotopy necessary conditions remain satisfied. For example, the problem of determining the range of κ appropriate to the current ψ in a “large step”

homotopy algorithm corresponds to the problem of designing a fast algorithm for solving the problem

$$\max_{\mathbf{v}, \mathbf{w}, \lambda, \beta} \sum_{i \in \psi} (\boldsymbol{\theta}_\psi)_i \beta_i \quad (30)$$

subject to constraints which express a small reformulation of the necessary conditions (1) and (10) – (14)

$$\mathbf{v}^T \begin{bmatrix} X_{11} \\ X_{21} \end{bmatrix} - \lambda \boldsymbol{\theta}_\psi^T = 0, \quad (31)$$

$$-\mathbf{e}_\sigma \leq \mathbf{v} \leq \mathbf{e}_\sigma, \quad (32)$$

$$\mathbf{v}^T \mathbf{r} = \|\mathbf{r}\|_1, \quad (33)$$

$$\mathbf{v}^T \begin{bmatrix} X_{12} \\ X_{22} \end{bmatrix} - \mathbf{w}_\psi^T = 0, \quad (34)$$

$$-\lambda \mathbf{e}_\psi \leq \mathbf{w}_\psi \leq \lambda \mathbf{e}_\psi, \quad (35)$$

$$\lambda \geq 0, \quad (36)$$

$$(\boldsymbol{\theta}_\psi)_i \beta_i \geq 0, \quad i \in \psi, \quad (37)$$

where $\mathbf{w} = \lambda \mathbf{u}$, and \mathbf{e}_σ and \mathbf{e}_ψ are vectors with all components 1 and having the indicated dimension. This is almost a linear program of the right structure. The catch is the single nonlinear constraint (33) which is required to restrict the components of \mathbf{v} to lie in the subdifferential of $\|\mathbf{r}\|_1$. This constraint is satisfied implicitly in the small step homotopy method.

3.2 Dealing with degeneracies

The homotopy algorithm encounters degeneracy when the decision process at the end of either the κ or λ step requires a choice to be made between several possibilities. For example, several residuals could vanish identically to terminate a κ step. One cause encountered in the diabetes and Boston housing data corresponded to identical observations of the response and partial design in sub models when several potential model variables are omitted. This degeneracy could be avoided by repeated inspection at some computational cost. Two alternative approaches to resolving this problem have been employed successfully in our computations:

1. Small random perturbations can be added to the input data. This has been used in a Fortran 95 implementation of the algorithm.
2. An R implementation permits several indices to enter σ during the vary κ step. If this happens then the matrices in (19) and (25) are only

square and invertible once duplicate rows are removed. These equations can be solved using generalised inverse technology. Typically, if several indices enter σ during a κ step then they are later simultaneously removed during a λ step. However, if the λ step ends with an increment in ψ then it may happen that, after the removal of duplicate rows, X_{21} does not have one more column than it has rows which is a requirement for (19) in the next κ step. In this case it is necessary to perform further λ steps, which should all involve decrementing σ , until the required condition on (19) is satisfied.

It is not easy to compare the computational performance of the two procedures. However, we have the impression that the procedure used in the R implementation appears to usually produce smaller iteration counts.

A Proofs

Proofs of the Lemmas that govern the switching between the κ and λ phases are presented. Note that the conditions presented are necessary for the piecewise linear dependence of $\boldsymbol{\beta}$ on κ , and \mathbf{v}_σ and $\lambda\mathbf{u}_\psi$ on λ to be continuous. However, the algebraic proofs presented here have some independent interest.

Remark 8 *It will be convenient to distinguish between the set of zero residuals in the κ phase and the augmented set corresponding to a new zero residual in the subsequent λ phase which starts with λ equal to its value λ_0 in the κ phase. The key components of the augmented system are written*

$$\bar{\mathbf{v}}_\sigma(\lambda_0) = \begin{bmatrix} \theta_k \\ \mathbf{v}_\sigma \end{bmatrix}, \quad \bar{X}_{21} = \begin{bmatrix} \mathbf{x}_k^T \\ X_{21} \end{bmatrix}.$$

As a result of augmenting X_{21} while holding ψ fixed it follows that \bar{X}_{21} is nonsingular. The key component of the necessary conditions becomes

$$\bar{\boldsymbol{\theta}}_\sigma^T \bar{X}_{11} + \bar{\mathbf{v}}_\sigma^T \bar{X}_{21} = \lambda \boldsymbol{\theta}_\psi^T,$$

and

$$\frac{d\bar{v}_1}{d\lambda} = \boldsymbol{\theta}_\psi^T \bar{X}_{21}^{-1} \mathbf{e}_1.$$

A similar notation is used in the case that ψ is downdated to take account of a component of $\boldsymbol{\beta}$ becoming zero. Let this component be $\beta_{\psi(j)}$ then it is

assumed that it is swapped to position $|\psi|$ and then moved to ψ^c in position 1. The repositioning is summarised by

$$\begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} = \begin{bmatrix} \bar{\boldsymbol{\theta}}_\psi^T & \theta_j \\ \bar{X}_{21} & \mathbf{x}^j \end{bmatrix} Q_j, \quad (38)$$

where the permutation matrix that expresses the operation that moves the j 'th component of $\boldsymbol{\theta}_\psi$ to the last position is written Q_j .

Proof of Lemma 3. It is necessary to show that the subgradient component associated with $r_k(\kappa_1) = 0$ satisfies equation (21).

Proof. It is convenient to characterise $\theta_k = \text{sign}(r_k)$ using the condition $r_k(\kappa_1) = 0$. This is

$$\mathbf{x}_k^T \left\{ \boldsymbol{\beta}_1(\kappa_0) + (\kappa_1 - \kappa_0) \frac{d\boldsymbol{\beta}_1}{d\kappa} \right\} - y_k = 0,$$

so that

$$\begin{aligned} \mathbf{x}_k^T \frac{d\boldsymbol{\beta}_1}{d\kappa} &= \mathbf{x}_k^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \mathbf{e}_1, \\ &= \frac{1}{\kappa_1 - \kappa_0} r_k(\kappa_0). \end{aligned} \quad (39)$$

It follows that

$$\theta_k = \text{sign} \left(\mathbf{x}_k^T \frac{d\boldsymbol{\beta}_1}{d\kappa} \right). \quad (40)$$

Let \mathbf{w} be defined by

$$\mathbf{x}_k^T = \mathbf{w}^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}, \quad (41)$$

where $w_1 \neq 0$, and $\mathbf{w}^T = [w_1 \quad \mathbf{w}_2^T]$. Then

$$\begin{aligned}
\bar{X}_{21} &= \begin{bmatrix} \mathbf{w}^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} \\ X_{21} \end{bmatrix}, \\
&= \begin{bmatrix} w_1 \boldsymbol{\theta}_\psi^T + \mathbf{w}_2^T X_{21} \\ (w_1 + (1 - w_1)) X_{21} \end{bmatrix}, \\
&= \begin{bmatrix} w_1 \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} + \begin{bmatrix} \mathbf{w}_2^T \\ (1 - w_1) I_{|\sigma|} \end{bmatrix} X_{21} \end{bmatrix}, \\
&= \begin{bmatrix} w_1 I_{|\psi|} + \begin{bmatrix} \mathbf{w}_2^T \\ (1 - w_1) I_{|\sigma|} \end{bmatrix} X_{21} \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \\ \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} \end{bmatrix}, \\
&= \begin{bmatrix} w_1 I_{|\psi|} + \begin{bmatrix} \mathbf{w}_2^T \\ (1 - w_1) I_{|\sigma|} \end{bmatrix} \begin{bmatrix} 0 & I_{|\sigma|} \end{bmatrix} \\ \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} \end{bmatrix}, \\
&= \begin{bmatrix} w_1 & \mathbf{w}_2^T \\ 0 & I_{|\sigma|} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}.
\end{aligned}$$

Thus

$$\bar{X}_{21}^{-1} = \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \begin{bmatrix} \frac{1}{w_1} & -\frac{1}{w_1} \mathbf{w}_2^T \\ 0 & I_{|\sigma|} \end{bmatrix}. \quad (42)$$

It follows that

$$\begin{aligned}
\frac{d\bar{v}_1}{d\lambda}(\lambda_0) &= \boldsymbol{\theta}_\psi^T \bar{X}_{21}^{-1} \mathbf{e}_1, \\
&= \mathbf{e}_1^T \begin{bmatrix} \frac{1}{w_1} & -\frac{1}{w_1} \mathbf{w}_2^T \\ 0 & I_{|\sigma|} \end{bmatrix} \mathbf{e}_1, \\
&= \frac{1}{w_1}.
\end{aligned} \quad (43)$$

But by (39)

$$\mathbf{x}_k^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \mathbf{e}_1 = \mathbf{w}^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \mathbf{e}_1 = w_1 = \frac{1}{\kappa_1 - \kappa_0} r_k(\kappa_0). \quad (44)$$

Thus $\text{sign}(w_1) = \theta_k$, so that

$$\text{sign} \left(\frac{d\bar{v}_1}{d\lambda}(\lambda_0) \right) \theta_k = 1.$$

■

Proof of Lemma 4

A preliminary result is required.

Lemma 9 *The solution to the linear system*

$$\begin{bmatrix} \mathbf{a}^T & b \\ I & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \gamma \end{bmatrix} = \begin{bmatrix} \mu \\ \mathbf{z} \end{bmatrix}$$

is

$$\begin{bmatrix} \mathbf{x} \\ \gamma \end{bmatrix} = \begin{bmatrix} \mathbf{z} - \frac{\mu - \mathbf{a}^T \mathbf{z}}{b - \mathbf{a}^T \mathbf{c}} \mathbf{c} \\ \frac{\mu - \mathbf{a}^T \mathbf{z}}{b - \mathbf{a}^T \mathbf{c}} \end{bmatrix}. \quad (45)$$

The particular case of interest is

$$\begin{bmatrix} \mu \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

which gives

$$\begin{bmatrix} \mathbf{x} \\ \gamma \end{bmatrix} = \frac{1}{b - \mathbf{a}^T \mathbf{c}} \begin{bmatrix} -\mathbf{c} \\ 1 \end{bmatrix}. \quad (46)$$

Proof. After the reordering the new component of \mathbf{u} will satisfy

$$u_1(\lambda_0) = (\boldsymbol{\theta}_\psi)_j = \theta_j$$

and is at its bound. It is required to move into its feasible region as λ is reduced. This gives the condition

$$\begin{aligned} \theta_j u_1(\lambda_0 - \Delta\lambda) &< 1, \\ \Rightarrow -\Delta\lambda \theta_j \frac{du_1(\lambda_0)}{d\lambda} &< 0, \\ \Rightarrow \theta_j \frac{du_1(\lambda_0)}{d\lambda} &\geq 0 \end{aligned}$$

The starting point is

$$\frac{d\beta_j}{d\kappa} = \mathbf{e}_j^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \mathbf{e}_1.$$

The next step is to highlight the updated quantities. This uses (38):

$$\begin{aligned} \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix} &= \begin{bmatrix} \bar{\boldsymbol{\theta}}_\psi^T & \theta_j \\ \bar{X}_{21} & \mathbf{x}^j \end{bmatrix} Q_j, \\ &= \begin{bmatrix} \bar{\boldsymbol{\theta}}_\psi^T \bar{X}_{21}^{-1} & \theta_j \\ I & \mathbf{x}^j \end{bmatrix} \begin{bmatrix} \bar{X}_{21} & \\ & 1 \end{bmatrix} Q_j, \\ &= \begin{bmatrix} \frac{d\bar{\mathbf{v}}^T}{d\lambda} & \theta_j \\ I & \mathbf{x}^j \end{bmatrix} \begin{bmatrix} \bar{X}_{21} & \\ & 1 \end{bmatrix} Q_j. \end{aligned}$$

It follows that

$$\begin{aligned}
\frac{d\beta_j}{d\kappa} &= \mathbf{e}_{|\psi|}^T \begin{bmatrix} \bar{X}_{21} & \\ & 1 \end{bmatrix}^{-1} \begin{bmatrix} \frac{d\bar{\mathbf{v}}^T}{d\lambda} & \theta_j \\ I & \mathbf{x}^j \end{bmatrix}^{-1} \mathbf{e}_1, \\
&= \mathbf{e}_{|\psi|}^T \begin{bmatrix} \bar{X}_{21} & \\ & 1 \end{bmatrix}^{-1} \frac{1}{\theta_j - \mathbf{x}^{jT} \frac{d\bar{\mathbf{v}}}{d\lambda}} \begin{bmatrix} -\mathbf{x}^j \\ 1 \end{bmatrix}, \\
&= \frac{\mathbf{e}_{|\psi|}^T}{\theta_j - \mathbf{x}^{jT} \frac{d\bar{\mathbf{v}}}{d\lambda}} \begin{bmatrix} \bar{X}_{21}^{-1} \mathbf{x}^j \\ 1 \end{bmatrix}, \\
&= \frac{1}{\theta_j - \mathbf{x}^{jT} \frac{d\bar{\mathbf{v}}}{d\lambda}}, \\
&= \frac{1}{\theta_j - \frac{d(\lambda u_1)}{d\lambda}}.
\end{aligned}$$

The final result is

$$\frac{d\beta_j}{d\kappa} \frac{du_1}{d\lambda} = -\frac{1}{\lambda} \quad (47)$$

showing that u_1 moves into its feasible region. ■

Proof of Lemma 6

Proof. Let the component v_q of \mathbf{v}_σ first reach a bound at $\lambda = \lambda_1 < \lambda_0$. Then

$$\begin{aligned}
\theta_q &= v_q(\lambda_1) = v_q(\lambda_0) + (\lambda_1 - \lambda_0) \frac{dv_q}{d\lambda}, \\
\Rightarrow \text{sign} \left(\frac{dv_q}{d\lambda} \right) &= -\theta_q, \text{ as } |v_q(\lambda_0)| < 1.
\end{aligned}$$

It is convenient to develop a representation of \mathbf{x}_q using (41).

$$\begin{aligned}
\mathbf{x}_k^T &= \mathbf{w}^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}, \\
&= w_1 \boldsymbol{\theta}_\psi^T + w_q \mathbf{x}_q^T + \sum_{j \in \{1, \dots, |\sigma|\} \setminus \{q\}} w_j \mathbf{e}_j^T X_{21}.
\end{aligned}$$

Thus

$$\begin{aligned}
\mathbf{x}_q^T &= \frac{1}{w_q} \mathbf{x}_k^T - \frac{w_1}{w_q} \boldsymbol{\theta}_\psi^T - \sum_{j \in \{1, \dots, |\sigma|\} \setminus \{q\}} \frac{w_j}{w_q} \mathbf{e}_j^T X_{21}, \\
&= \begin{bmatrix} -\frac{w_1}{w_q} & \frac{1}{w_q} & \dots & -\frac{w_j}{w_q} & \dots \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ \mathbf{x}_k^T \\ X_{21}^{-q} \end{bmatrix},
\end{aligned}$$

where X_{21}^{-q} denotes X_{21} with row \mathbf{x}_q removed. This gives, using (19),

$$\frac{dr_q}{d\kappa} = -\mathbf{x}_q \frac{d\boldsymbol{\beta}_1}{d\kappa} = -\mathbf{x}_q^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ \mathbf{x}_k^T \\ X_{21}^{-q} \end{bmatrix}^{-1} \mathbf{e}_1 = \frac{w_1}{w_q}.$$

On the other hand

$$\frac{dv_\sigma^T}{d\lambda} = \boldsymbol{\theta}_\psi^T \bar{X}_{21}^{-1},$$

so that

$$\frac{dv_q}{d\lambda} = \boldsymbol{\theta}_\psi^T \bar{X}_{21}^{-1} \mathbf{e}_q = \boldsymbol{\theta}_\psi^T \begin{bmatrix} \boldsymbol{\theta}_\psi^T \\ X_{21} \end{bmatrix}^{-1} \begin{bmatrix} \frac{1}{w_1} & -\frac{\mathbf{w}_2^T}{w_1} \\ 0 & I_{|\sigma|} \end{bmatrix} \mathbf{e}_q = -\frac{w_q}{w_1}.$$

Thus

$$\text{sign} \left(\frac{dr_q}{d\kappa} \right) = -\text{sign} \left(\frac{dv_q}{d\lambda} \right).$$

■

Proof of Lemma 7

Proof. Assume that u_q , $q \in \psi^c$ is the variable that reaches its bound as λ is decreased. Let $u_q(\lambda_1) = \theta_q$ then the equation determining $\frac{d\boldsymbol{\beta}}{d\kappa}$ is

$$\frac{d\boldsymbol{\beta}}{d\kappa} = \begin{bmatrix} \boldsymbol{\theta}_\psi^T & \theta_q \\ \bar{X}_{21} & \bar{X}_{22} \mathbf{e}_q \end{bmatrix}^{-1} \mathbf{e}_1.$$

From equations (24) and (25) it follows that

$$\begin{bmatrix} 1 & -\frac{dv_\sigma^T}{d\lambda} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_\psi^T & \theta_q \\ \bar{X}_{21} & \bar{X}_{22} \mathbf{e}_q \end{bmatrix} = \begin{bmatrix} 0 & \dots & 0 & \theta_q - \frac{d(\lambda u_q)}{d\lambda} \end{bmatrix}.$$

Thus

$$\begin{bmatrix} 1 & -\frac{dv_\sigma^T}{d\lambda} \end{bmatrix} \mathbf{e}_1 = \begin{bmatrix} 0 & \dots & 0 & \theta_q - \frac{d(\lambda u_q)}{d\lambda} \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_\psi^T & \theta_q \\ \bar{X}_{21} & \bar{X}_{22} \mathbf{e}_q \end{bmatrix}^{-1} \mathbf{e}_1,$$

giving

$$\begin{aligned} 1 &= \left(\theta_q - \frac{d(\lambda u_q)}{d\lambda} \right) \frac{d\beta_q}{d\lambda}, \\ &= -\lambda \frac{du_q}{d\lambda} \frac{d\beta_q}{d\kappa}. \end{aligned}$$

Equation (27) is a direct consequence as $\lambda > 0$. ■

References

- [1] H. D. Bondell and B. J. Reich. Simultaneous regression shrinkage, variable selection, and supervising clustering of predictors with OSCAR. *Biometrics*, 64:115–123, 2007.
- [2] E. Candes and T. Tao. The Dantzig selector: statistical estimation when p is much larger than n . *Annals of Statistics*, 35(6):2313–2351, 2007.
- [3] D. L. Donoho and Y. Tsaig. Fast solution of l_1 -norm minimization problems when the solution may be sparse. *IEEE Transactions on Information Theory*, 54:4789–4812, 2008.
- [4] N. H. Draper and H. Smith. *Applied Regression Analysis*. John Wiley & Sons, 3 edition, 1998.
- [5] T. P. Hettmansperger and J. W. McKean. *Robust Nonparametric Statistical Methods*. John Wiley & Sons, Chichester, 1998.
- [6] Kengo Kato. Solving l_1 regularization problems with piecewise linear losses. *J. Comp. Graph. Statist.*, 2010. to appear.
- [7] R. Koenker. *Quantile Regression*. Cambridge University Press, 2005.
- [8] Y. Li and J. Zhu. L_1 -norm quantile regression. *Journal of Computational and Graphical Statistics*, 17:163–185, 2008.
- [9] M. R. Osborne. *Simplicial Algorithms for Minimizing Polyhedral Functions*. Cambridge University Press, 2001.
- [10] M. R. Osborne, B. Presnell, and B. A. Turlach. A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis*, 20:389–403, 2000.
- [11] S. Rosset and Ji Zhu. Piecewise linear regularised solution paths. *Annals of Statistics*, 35(3):1012–1030, 2007.
- [12] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [13] B. A. Turlach, W. N. Venables, and S. J. Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.
- [14] Y. Yao and Y. Lee. Another look at linear programming for feature selection via methods of regularization. Technical Report 800, Department of Statistics, Ohio State University, 2007.

- [15] Ming Yuan and Hui Zou. Efficient Global Approximation of Generalised Nonlinear ℓ_1 Regularised Solution Paths and its Applications. *J. American Statistical Association*, 104:1562–1573, 2009. DOI:10.1198/jasa.2009.tm08287.
- [16] J. Zhu, T. Hastie, S. Rosset, and R. Tibshirani. l_1 -norm Support Vector Machines. *Advances in Neural Information Processing Systems*, 16:49–56, 2004.
- [17] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67:301–320, 2005.
- [18] H. Zou and M. Yuan. Regularised simultaneous model selection in multiple quantiles regression. *Computational Statistics & Data Analysis*, 52:5296–5304, 2008.