

4-2012

Degrees of freedom in lasso problems

Ryan J. Tibshirani

Carnegie Mellon University, ryantibs@cmu.edu

Jonathan Taylor

Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/statistics>

 Part of the [Statistics and Probability Commons](#)

Published In

The Annals of Statistics , 40, 2, 1198-1232.

This Article is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Statistics by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

DEGREES OF FREEDOM IN LASSO PROBLEMS

BY RYAN J. TIBSHIRANI AND JONATHAN TAYLOR¹

Carnegie Mellon University and Stanford University

We derive the degrees of freedom of the lasso fit, placing no assumptions on the predictor matrix X . Like the well-known result of Zou, Hastie and Tibshirani [*Ann. Statist.* **35** (2007) 2173–2192], which gives the degrees of freedom of the lasso fit when X has full column rank, we express our result in terms of the active set of a lasso solution. We extend this result to cover the degrees of freedom of the generalized lasso fit for an arbitrary predictor matrix X (and an arbitrary penalty matrix D). Though our focus is degrees of freedom, we establish some intermediate results on the lasso and generalized lasso that may be interesting on their own.

1. Introduction. We study degrees of freedom, or the “effective number of parameters,” in ℓ_1 -penalized linear regression problems. In particular, for a response vector $y \in \mathbb{R}^n$, predictor matrix $X \in \mathbb{R}^{n \times p}$ and tuning parameter $\lambda \geq 0$, we consider the lasso problem [Chen, Donoho and Saunders (1998), Tibshirani (1996)]

$$(1) \quad \hat{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1.$$

The above notation emphasizes the fact that the solution $\hat{\beta}$ may not be unique [such nonuniqueness can occur if $\operatorname{rank}(X) < p$]. Throughout the paper, when a function $f: D \rightarrow \mathbb{R}^n$ may have a nonunique minimizer over its domain D , we write $\operatorname{argmin}_{x \in D} f(x)$ to denote the set of minimizing x values, that is, $\operatorname{argmin}_{x \in D} f(x) = \{\hat{x} \in D: f(\hat{x}) = \min_{x \in D} f(x)\}$.

A fundamental result on the degrees of freedom of the lasso fit was shown by Zou, Hastie and Tibshirani (2007). The authors show that if y follows a normal distribution with spherical covariance, $y \sim N(\mu, \sigma^2 I)$, and X, λ are

Received November 2011; revised February 2012.

¹Supported by NSF Grant DMS-09-06801 and AFOSR Grant 113039.

AMS 2000 subject classifications. 62J07, 90C46.

Key words and phrases. Lasso, generalized lasso, degrees of freedom, high-dimensional.

<p>This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in <i>The Annals of Statistics</i>, 2012, Vol. 40, No. 2, 1198–1232. This reprint differs from the original in pagination and typographic detail.</p>

considered fixed with $\text{rank}(X) = p$, then

$$(2) \quad \text{df}(X\hat{\beta}) = \mathbb{E}|\mathcal{A}|,$$

where $\mathcal{A} = \mathcal{A}(y)$ denotes the active set of the unique lasso solution at y , and $|\mathcal{A}|$ is its cardinality. This is quite a well-known result, and is sometimes used to informally justify an application of the lasso procedure, as it says that number of parameters used by the lasso fit is simply equal to the (average) number of selected variables. However, we note that the assumption $\text{rank}(X) = p$ implies that $p \leq n$; in other words, the degrees of freedom result (2) does not cover the important “high-dimensional” case $p > n$. In this case, the lasso solution is not necessarily unique, which raises the questions:

- Can we still express degrees of freedom in terms of the active set of a lasso solution?
- If so, which active set (solution) would we refer to?

In Section 3, we provide answers to these questions, by proving a stronger result when X is a general predictor matrix. We show that the subspace spanned by the columns of X in \mathcal{A} is almost surely unique, where “almost surely” means for almost every $y \in \mathbb{R}^n$. Furthermore, the degrees of freedom of the lasso fit is simply the expected dimension of this column space.

We also consider the generalized lasso problem,

$$(3) \quad \hat{\beta} \in \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|D\beta\|_1,$$

where $D \in \mathbb{R}^{m \times p}$ is a penalty matrix, and again the notation emphasizes the fact that $\hat{\beta}$ need not be unique [when $\text{rank}(X) < p$]. This of course reduces to the usual lasso problem (1) when $D = I$, and Tibshirani and Taylor (2011) demonstrate that the formulation (3) encapsulates several other important problems—including the fused lasso on any graph and trend filtering of any order—by varying the penalty matrix D . The same paper shows that if y is normally distributed as above, and X, D, λ are fixed with $\text{rank}(X) = p$, then the generalized lasso fit has degrees of freedom

$$(4) \quad \text{df}(X\hat{\beta}) = \mathbb{E}[\text{nullity}(D_{-\mathcal{B}})].$$

Here $\mathcal{B} = \mathcal{B}(y)$ denotes the boundary set of an optimal subgradient to the generalized lasso problem at y (equivalently, the boundary set of a dual solution at y), $D_{-\mathcal{B}}$ denotes the matrix D after having removed the rows that are indexed by \mathcal{B} , and $\text{nullity}(D_{-\mathcal{B}}) = \dim(\text{null}(D_{-\mathcal{B}}))$, the dimension of the null space of $D_{-\mathcal{B}}$.

It turns out that examining (4) for specific choices of D produces a number of interpretable corollaries, as discussed in Tibshirani and Taylor (2011). For example, this result implies that the degrees of freedom of the fused lasso fit is equal to the expected number of fused groups, and that the degrees of freedom of the trend filtering fit is equal to the expected number of knots

+ $(k - 1)$, where k is the order of the polynomial. The result (4) assumes that $\text{rank}(X) = p$ and does not cover the case $p > n$; in Section 4, we derive the degrees of freedom of the generalized lasso fit for a general X (and still a general D). As in the lasso case, we prove that there exists a linear subspace $X(\text{null}(D_{-\mathcal{B}}))$ that is almost surely unique, meaning that it will be the same under different boundary sets \mathcal{B} corresponding to different solutions of (3). The generalized lasso degrees of freedom is then the expected dimension of this subspace.

Our assumptions throughout the paper are minimal. As was already mentioned, we place no assumptions whatsoever on the predictor matrix $X \in \mathbb{R}^{n \times p}$ or on the penalty matrix $D \in \mathbb{R}^{m \times p}$, considering them fixed and nonrandom. We also consider $\lambda \geq 0$ fixed. For Theorems 1, 2 and 3 we assume that y is normally distributed,

$$(5) \quad y \sim N(\mu, \sigma^2 I)$$

for some (unknown) mean vector $\mu \in \mathbb{R}^n$ and marginal variance $\sigma^2 \geq 0$. This assumption is only needed in order to apply Stein's formula for degrees of freedom, and none of the other lasso and generalized lasso results in the paper, namely Lemmas 3 through 10, make any assumption about the distribution of y .

This paper is organized as follows. The rest of the Introduction contains an overview of related work, and an explanation of our notation. Section 2 covers some relevant background material on degrees of freedom and convex polyhedra. Though the connection may not be immediately obvious, the geometry of polyhedra plays a large role in understanding problems (1) and (3), and Section 2.2 gives a high-level view of this geometry before the technical arguments that follow in Sections 3 and 4. In Section 3, we derive two representations for the degrees of freedom of the lasso fit, given in Theorems 1 and 2. In Section 4, we derive the analogous results for the generalized lasso problem, and these are given in Theorem 3. As the lasso problem is a special case of the generalized lasso problem (corresponding to $D = I$), Theorems 1 and 2 can actually be viewed as corollaries of Theorem 3. The reader may then ask: why is there a separate section dedicated to the lasso problem? We give two reasons: first, the lasso arguments are simpler and easier to follow than their generalized lasso counterparts; second, we cover some intermediate results for the lasso problem that are interesting in their own right and that do not carry over to the generalized lasso perspective. Section 5 contains some final discussion.

1.1. *Related work.* All of the degrees of freedom results discussed here assume that the response vector has distribution $y \sim N(\mu, \sigma^2 I)$, and that the predictor matrix X is fixed. To the best of our knowledge, Efron et al. (2004) were the first to prove a result on the degrees of freedom of the lasso

fit, using the lasso solution path with λ moving from ∞ to 0. The authors showed that when the active set reaches size k along this path, the lasso fit has degrees of freedom exactly k . This result assumes that X has full column rank and further satisfies a restrictive condition called the “positive cone condition,” which ensures that as λ decreases, variables can only enter, and not leave, the active set. Subsequent results on the lasso degrees of freedom (including those presented in this paper) differ from this original result in that they derive degrees of freedom for a fixed value of the tuning parameter λ , and not a fixed number of steps k taken along the solution path.

As mentioned previously, Zou, Hastie and Tibshirani (2007) established the basic lasso degrees of freedom result (for fixed λ) stated in (2). This is analogous to the path result of Efron et al. (2004); here degrees of freedom is equal to the expected size of the active set (rather than simply the size) because for a fixed λ the active set is a random quantity, and can hence achieve a random size. The proof of (2) appearing in Zou, Hastie and Tibshirani (2007) relies heavily on properties of the lasso solution path. As also mentioned previously, Tibshirani and Taylor (2011) derived an extension of (2) to the generalized lasso problem, which is stated in (4) for an arbitrary penalty matrix D . Their arguments are not based on properties of the solution path, but instead come from a geometric perspective much like the one developed in this paper.

Both of the results (2) and (4) assume that $\text{rank}(X) = p$; the current work extends these to the case of an arbitrary matrix X , in Theorems 1, 2 (the lasso) and 3 (the generalized lasso). In terms of our intermediate results, a version of Lemmas 5, 6 corresponding to $\text{rank}(X) = p$ appears in Zou, Hastie and Tibshirani (2007), and a version of Lemma 9 corresponding to $\text{rank}(X) = p$ appears in Tibshirani and Taylor (2011) [furthermore, Tibshirani and Taylor (2011) only consider the boundary set representation and not the active set representation]. Lemmas 1, 2 and the conclusions thereafter, on the degrees of freedom of the projection map onto a convex polyhedron, are essentially given in Meyer and Woodroffe (2000), though these authors state and prove the results in a different manner.

In preparing a draft of this manuscript, it was brought to our attention that other authors have independently and concurrently worked to extend results (2) and (4) to the general X case. Namely, Dossal et al. (2011) prove a result on the lasso degrees of freedom, and Vaiter et al. (2011) prove a result on the generalized lasso degrees of freedom, both for an arbitrary X . These authors’ results express degrees of freedom in terms of the active sets of special (lasso or generalized lasso) solutions. Theorems 2 and 3 express degrees of freedom in terms of the active sets of any solutions, and hence the appropriate application of these theorems provides an alternative verification of these formulas. We discuss this in detail in the form of remarks following the theorems.

1.2. *Notation.* In this paper, we use $\text{col}(A)$, $\text{row}(A)$ and $\text{null}(A)$ to denote the column space, row space and null space of a matrix A , respectively; we use $\text{rank}(A)$ and $\text{nullity}(A)$ to denote the dimensions of $\text{col}(A)$ [equivalently, $\text{row}(A)$] and $\text{null}(A)$, respectively. We write A^+ for the Moore–Penrose pseudoinverse of A ; for a rectangular matrix A , recall that $A^+ = (A^T A)^+ A^T$. We write P_L to denote the projection matrix onto a linear subspace L , and more generally, $P_C(x)$ to denote the projection of a point x onto a closed convex set C . For readability, we sometimes write $\langle a, b \rangle$ (instead of $a^T b$) to denote the inner product between vectors a and b .

For a set of indices $R = \{i_1, \dots, i_k\} \subseteq \{1, \dots, m\}$ satisfying $i_1 < \dots < i_k$, and a vector $x \in \mathbb{R}^m$, we use x_R to denote the subvector $x_R = (x_{i_1}, \dots, x_{i_k})^T \in \mathbb{R}^k$. We denote the complementary subvector by $x_{-R} = x_{\{1, \dots, m\} \setminus R} \in \mathbb{R}^{m-k}$. The notation is similar for matrices. Given another subset of indices $S = \{j_1, \dots, j_\ell\} \subseteq \{1, \dots, p\}$ with $j_1 < \dots < j_\ell$, and a matrix $A \in \mathbb{R}^{m \times p}$, we use $A_{(R,S)}$ to denote the submatrix

$$A_{(R,S)} = \begin{bmatrix} A_{i_1, j_1} & \cdots & A_{i_1, j_\ell} \\ \vdots & & \\ A_{i_k, j_1} & \cdots & A_{i_k, j_\ell} \end{bmatrix} \in \mathbb{R}^{k \times \ell}.$$

In words, rows are indexed by R , and columns are indexed by S . When combining this notation with the transpose operation, we assume that the indexing happens first, so that $A_{(R,S)}^T = (A_{(R,S)})^T$. As above, negative signs are used to denote the complementary set of rows or columns; for example, $A_{(-R,S)} = A_{(\{1, \dots, m\} \setminus R, S)}$. To extract only rows or only columns, we abbreviate the other dimension by a dot, so that $A_{(R,\cdot)} = A_{(R, \{1, \dots, p\})}$ and $A_{(\cdot, S)} = A_{(\{1, \dots, m\}, S)}$; to extract a single row or column, we use $A_{(i,\cdot)} = A_{(\{i\}, \cdot)}$ or $A_{(\cdot, j)} = A_{(\cdot, \{j\})}$. Finally, and most importantly, we introduce the following shorthand notation:

- For the predictor matrix $X \in \mathbb{R}^{n \times p}$, we let $X_S = X_{(\cdot, S)}$.
- For the penalty matrix $D \in \mathbb{R}^{m \times p}$, we let $D_R = D_{(R, \cdot)}$.

In other words, the default for X is to index its columns, and the default for D is to index its rows. This convention greatly simplifies the notation in expressions that involve multiple instances of X_S or D_R ; however, its use could also cause a great deal of confusion, if not properly interpreted by the reader!

2. Preliminary material. The following two sections describe some background material needed to follow the results in Sections 3 and 4.

2.1. *Degrees of freedom.* If the data vector $y \in \mathbb{R}^n$ is distributed according to the homoskedastic model $y \sim (\mu, \sigma^2 I)$, meaning that the components of y are uncorrelated, with y_i having mean μ_i and variance σ^2 for $i = 1, \dots, n$, then the degrees of freedom of a function $g: \mathbb{R}^n \rightarrow \mathbb{R}^n$ with

$g(y) = (g_1(y), \dots, g_n(y))^T$, is defined as

$$(6) \quad \text{df}(g) = \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}(g_i(y), y_i).$$

This definition is often attributed to Efron (1986) or Hastie and Tibshirani (1990), and is interpreted as the “effective number of parameters” used by the fitting procedure g . Note that for the linear regression fit of $y \in \mathbb{R}^n$ onto a fixed and full column rank predictor matrix $X \in \mathbb{R}^{n \times p}$, we have $g(y) = \hat{y} = XX^+y$, and $\text{df}(\hat{y}) = \text{tr}(XX^+) = p$, which is the number of fitted coefficients (one for each predictor variable). Furthermore, we can decompose the risk of \hat{y} , denoted by $\text{Risk}(\hat{y}) = \mathbb{E}\|\hat{y} - \mu\|_2^2$, as

$$\text{Risk}(\hat{y}) = \mathbb{E}\|\hat{y} - y\|_2^2 - n\sigma^2 + 2p\sigma^2,$$

a well-known identity that leads to the derivation of the C_p statistic [Mallows (1973)]. For a general fitting procedure g , the motivation for the definition (6) comes from the analogous decomposition of the quantity $\text{Risk}(g) = \mathbb{E}\|g(y) - \mu\|_2^2$,

$$(7) \quad \text{Risk}(g) = \mathbb{E}\|g(y) - y\|_2^2 - n\sigma^2 + 2 \sum_{i=1}^n \text{Cov}(g_i(y), y_i).$$

Therefore a large difference between risk and expected training error implies a large degrees of freedom.

Why is the concept of degrees of freedom important? One simple answer is that it provides a way to put different fitting procedures on equal footing. For example, it would not seem fair to compare a procedure that uses an effective number of parameters equal to 100 with another that uses only 10. However, assuming that these procedures can be tuned to varying levels of adaptivity (as is the case with the lasso and generalized lasso, where the adaptivity is controlled by λ), one could first tune the procedures to have the same degrees of freedom, and then compare their performances. Doing this over several common values for degrees of freedom may reveal, in an informal sense, that one procedure is particularly efficient when it comes to its parameter usage versus another.

A more detailed answer to the above question is based the risk decomposition (7). The decomposition suggests that an estimate $\widehat{\text{df}}(g)$ of degrees of freedom can be used to form an estimate of the risk,

$$(8) \quad \widehat{\text{Risk}}(g) = \|g(y) - y\|_2^2 - n\sigma^2 + 2\sigma^2\widehat{\text{df}}(g).$$

Furthermore, it is straightforward to check that an unbiased estimate of degrees of freedom leads to an unbiased estimate of risk; that is, $\text{df}(g) = \mathbb{E}[\widehat{\text{df}}(g)]$ implies $\text{Risk}(g) = \mathbb{E}[\widehat{\text{Risk}}(g)]$. Hence, the risk estimate (8) can be used to choose between fitting procedures, assuming that unbiased estimates of degrees of freedom are available. [It is worth mentioning that bootstrap or Monte Carlo methods can be helpful in estimating degrees of freedom (6)

when an analytic form is difficult to obtain.] The natural extension of this idea is to use the risk estimate (8) for tuning parameter selection. If we suppose that g depends on a tuning parameter $\lambda \in \Lambda$, denoted $g = g_\lambda(y)$, then in principle one could minimize the estimated risk over λ to select an appropriate value for the tuning parameter,

$$(9) \quad \hat{\lambda} = \operatorname{argmin}_{\lambda \in \Lambda} \|g_\lambda(y) - y\|_2^2 - n\sigma^2 + 2\sigma^2 \hat{\operatorname{df}}(g_\lambda).$$

This is a computationally efficient alternative to selecting the tuning parameter by cross-validation, and it is commonly used (along with similar methods that replace the factor of 2 above with a function of n or p) in penalized regression problems. Even though such an estimate (9) is commonly used in the high-dimensional setting ($p > n$), its asymptotic properties are largely unknown in this case, such as risk consistency, or relative efficiency compared to the cross-validation estimate.

Stein (1981) proposed the risk estimate (8) using a particular unbiased estimate of degrees of freedom, now commonly referred to as *Stein's unbiased risk estimate* (SURE). Stein's framework requires that we strengthen our distributional assumption on y and assume normality, as stated in (5). We also assume that the function g is continuous and almost differentiable. (The precise definition of almost differentiability is not important here, but the interested reader may take it to mean that each coordinate function g_i is absolutely continuous on almost every line segment parallel to one of the coordinate axes.) Given these assumptions, Stein's main result is an alternate expression for degrees of freedom,

$$(10) \quad \operatorname{df}(g) = \mathbb{E}[(\nabla \cdot g)(y)],$$

where the function $\nabla \cdot g = \sum_{i=1}^n \partial g_i / \partial y_i$ is called the divergence of g . Immediately following is the unbiased estimate of degrees of freedom,

$$(11) \quad \hat{\operatorname{df}}(g) = (\nabla \cdot g)(y).$$

We pause for a moment to reflect on the importance of this result. From its definition (6), we can see that the two most obvious candidates for unbiased estimates of degrees of freedom are

$$\frac{1}{\sigma^2} \sum_{i=1}^n g_i(y)(y_i - \mu_i) \quad \text{and} \quad \frac{1}{\sigma^2} \sum_{i=1}^n (g_i(y) - \mathbb{E}[g_i(y)])y_i.$$

To use the first estimate above, we need to know μ (remember, this is ultimately what we are trying to estimate!). Using the second requires knowing $\mathbb{E}[g(y)]$, which is equally impractical because this invariably depends on μ . On the other hand, Stein's unbiased estimate (11) does not have an explicit dependence on μ ; moreover, it can be analytically computed for many fitting procedures g . For example, Theorem 2 in Section 3 shows that, except for y in a set of measure zero, the divergence of the lasso fit is equal to $\operatorname{rank}(X_{\mathcal{A}})$ with $\mathcal{A} = \mathcal{A}(y)$ being the active set of a lasso solution at y . Hence, Stein's




FIG. 1. An example of a polyhedron in \mathbb{R}^2 .

formula allows for the unbiased estimation of degrees of freedom (and subsequently, risk) for a broad class of fitting procedures g —something that may have not seemed possible when working from the definition directly.

2.2. Projections onto polyhedra. A set $C \subseteq \mathbb{R}^n$ is called a *convex polyhedron*, or simply a *polyhedron*, if C is the intersection of finitely many half-spaces,

$$(12) \quad C = \bigcap_{i=1}^k \{x \in \mathbb{R}^n : a_i^T x \leq b_i\},$$

where $a_1, \dots, a_k \in \mathbb{R}^n$ and $b_1, \dots, b_k \in \mathbb{R}$. (Note that we do not require boundedness here; a bounded polyhedron is sometimes called a polytope.) See Figure 1 for an example. There is a rich theory on polyhedra; the definitive reference is Grünbaum (2003), and another good reference is Schneider (1993). As this is a paper on statistics and not geometry, we do not attempt to give an extensive treatment of the properties of polyhedra. We do, however, give two properties (in the form of two lemmas) that are especially important with respect to our statistical problem; our discussion will also make it clear why polyhedra are relevant in the first place.

From its definition (12), it follows that a polyhedron is a closed convex set. The first property that we discuss does not actually rely on the special structure of polyhedra, but only on convexity. For any closed convex set

$C \subseteq \mathbb{R}^n$ and any point $x \in \mathbb{R}^n$, there is a unique point $u \in C$ minimizing $\|x - u\|_2$. To see this, note that if $v \in C$ is another minimizer, $v \neq u$, then by convexity $w = (u + v)/2 \in C$, and $\|x - w\|_2 < \|x - u\|_2/2 + \|x - v\|_2/2 = \|x - u\|_2$, a contradiction. Therefore, the projection map onto C is indeed well defined, and we write this as $P_C: \mathbb{R}^n \rightarrow C$,

$$P_C(x) = \operatorname{argmin}_{u \in C} \|x - u\|_2.$$

For the usual linear regression problem, where $y \in \mathbb{R}^n$ is regressed onto $X \in \mathbb{R}^{n \times p}$, the fit $X\hat{\beta}$ can be written in terms of the projection map onto the polyhedron $C = \operatorname{col}(X)$, as in $X\hat{\beta}(y) = XX^+y = P_{\operatorname{col}(X)}(y)$. Furthermore, for both the lasso and generalized lasso problems, (1) and (3), it turns out that we can express the fit as the residual from projecting onto a suitable polyhedron $C \subseteq \mathbb{R}^n$, that is,

$$X\hat{\beta}(y) = (I - P_C)(y) = y - P_C(y).$$

This is proved in Lemma 3 for the lasso and in Lemma 8 for the generalized lasso (the polyhedron C depends on X, λ for the lasso case, and on X, D, λ for the generalized lasso case). Our first lemma establishes that both the projection map onto a closed convex set and the residual map are nonexpansive, hence continuous and almost differentiable everywhere. These are the conditions needed to apply Stein's formula.

LEMMA 1. *For any closed convex set $C \subseteq \mathbb{R}^n$, both the projection map $P_C: \mathbb{R}^n \rightarrow C$ and the residual projection map $I - P_C: \mathbb{R}^n \rightarrow \mathbb{R}^n$ are nonexpansive. That is, they satisfy*

$$\begin{aligned} \|P_C(x) - P_C(y)\|_2 &\leq \|x - y\|_2 && \text{for any } x, y \in \mathbb{R}^n, \quad \text{and} \\ \|(I - P_C)(x) - (I - P_C)(y)\|_2 &\leq \|x - y\|_2 && \text{for any } x, y \in \mathbb{R}^n. \end{aligned}$$

Therefore, P_C and $I - P_C$ are both continuous and almost differentiable.

The proof can be found in Appendix A.1. Lemma 1 will be quite useful later in the paper, as it will allow us to use Stein's formula to compute the degrees of freedom of the lasso and generalized lasso fits, after showing that these fits are indeed the residuals from projecting onto closed convex sets.

The second property that we discuss uses the structure of polyhedra. Unlike Lemma 1, this property will not be used directly in the following sections of the paper; instead, we present it here to give some intuition with respect to the degrees of freedom calculations to come. The property can be best explained by looking back at Figure 1. Loosely speaking, the picture suggests that we can move the point x around a bit and it will still project to the same face of C . Another way of saying this is that there is a neighborhood

of x on which P_C is simply the projection onto an affine subspace. This would not be true if x is in some exceptional set, which is made up of rays that emanate from the corners of C , like the two drawn in the bottom right corner of figure. However, the union of such rays has measure zero, so the map P_C is locally an affine projection, almost everywhere. This idea can be stated formally as follows.

LEMMA 2. *Let $C \subseteq \mathbb{R}^n$ be a polyhedron. For almost every $x \in \mathbb{R}^n$, there is an associated neighborhood U of x , linear subspace $L \subseteq \mathbb{R}^n$ and point $a \in \mathbb{R}^n$, such that the projection map restricted to U , $P_C: U \rightarrow C$, is*

$$P_C(y) = P_L(y - a) + a \quad \text{for } y \in U,$$

which is simply the projection onto the affine subspace $L + a$.

The proof is given in Appendix A.2. These last two properties can be used to derive a general expression for the degrees of freedom of the fitting procedure $g(y) = (I - P_C)(y)$, when $C \subseteq \mathbb{R}^n$ is a polyhedron. [A similar formula holds for $g(y) = P_C(y)$.] Lemma 1 tells us that $I - P_C$ is continuous and almost differentiable, so we can use Stein's formula (10) to compute its degrees of freedom. Lemma 2 tells us that for almost every $y \in \mathbb{R}^n$, there is a neighborhood U of y , linear subspace $L \subseteq \mathbb{R}^n$, and point $a \in \mathbb{R}^n$, such that

$$(I - P_C)(y') = y' - P_L(y' - a) - a = (I - P_L)(y' - a) \quad \text{for } y' \in U.$$

Therefore,

$$(\nabla \cdot (I - P_C))(y) = \text{tr}(I - P_L) = n - \dim(L),$$

and an expectation over y gives

$$\text{df}(I - P_C) = n - \text{E}[\dim(L)].$$

It should be made clear that the random quantity in the above expectation is the linear subspace $L = L(y)$, which depends on y .

In a sense, the remainder of this paper is focused on describing $\dim(L)$ —the dimension of the face of C onto which the point y projects—in a meaningful way for the lasso and generalized lasso problems. Section 3 considers the lasso problem, and we show that L can be written in terms of the equicorrelation set of the fit at y . We also show that L can be described in terms of the active set of a solution at y . In Section 4 we show the analogous results for the generalized lasso problem, namely, that L can be written in terms of either the boundary set of an optimal subgradient at y (the analogy of the equicorrelation set for the lasso) or the active set of a solution at y .

3. The lasso. In this section we derive the degrees of freedom of the lasso fit, for a general predictor matrix X . All of our arguments stem from the Karush–Kuhn–Tucker (KKT) optimality conditions, and we present these first. We note that many of the results in this section can be alternatively

derived using the lasso dual problem. Appendix A.5 explains this connection more precisely. For the current work, we avoid the dual perspective simply to keep the presentation more self-contained. Finally, we remind the reader that X_S is used to extract columns of X corresponding to an index set S .

3.1. *The KKT conditions and the underlying polyhedron.* The KKT conditions for the lasso problem (1) can be expressed as

$$(13) \quad X^T(y - X\hat{\beta}) = \lambda\gamma,$$

$$(14) \quad \gamma_i \in \begin{cases} \{\text{sign}(\hat{\beta}_i)\} & \text{if } \hat{\beta}_i \neq 0, \\ [-1, 1] & \text{if } \hat{\beta}_i = 0. \end{cases}$$

Here $\gamma \in \mathbb{R}^p$ is a subgradient of the function $f(x) = \|x\|_1$ evaluated at $x = \hat{\beta}$. Hence $\hat{\beta}$ is a minimizer in (1) if and only if $\hat{\beta}$ satisfies (13) and (14) for some γ . Directly from the KKT conditions, we can show that $X\hat{\beta}$ is the residual from projecting y onto a polyhedron.

LEMMA 3. *For any X and $\lambda \geq 0$, the lasso fit $X\hat{\beta}$ can be written as $X\hat{\beta}(y) = (I - P_C)(y)$, where $C \subseteq \mathbb{R}^n$ is the polyhedron*

$$C = \{u \in \mathbb{R}^n : \|X^T u\|_\infty \leq \lambda\}.$$

PROOF. Given a point $y \in \mathbb{R}^n$, its projection $\theta = P_C(y)$ onto a closed convex set $C \subseteq \mathbb{R}^n$ can be characterized as the unique point satisfying

$$(15) \quad \langle y - \theta, \theta - u \rangle \geq 0 \quad \text{for all } u \in C.$$

Hence defining $\theta = y - X\hat{\beta}(y)$, and C as in the lemma, we want to show that (15) holds for all $u \in C$. Well,

$$(16) \quad \begin{aligned} \langle y - \theta, \theta - u \rangle &= \langle X\hat{\beta}, y - X\hat{\beta} - u \rangle \\ &= \langle X\hat{\beta}, y - X\hat{\beta} \rangle - \langle X^T u, \hat{\beta} \rangle. \end{aligned}$$

Consider the first term above. Taking an inner product with $\hat{\beta}$ on both sides of (13) gives $\langle X\hat{\beta}, y - X\hat{\beta} \rangle = \lambda\|\hat{\beta}\|_1$. Furthermore, the ℓ_1 norm can be characterized in terms of its dual norm, the ℓ_∞ norm, as in

$$\lambda\|\hat{\beta}\|_1 = \max_{\|w\|_\infty \leq \lambda} \langle w, \hat{\beta} \rangle.$$

Therefore, continuing from (16), we have

$$\langle y - \theta, \theta - u \rangle = \max_{\|w\|_\infty \leq \lambda} \langle w, \hat{\beta} \rangle - \langle X^T u, \hat{\beta} \rangle,$$

which is ≥ 0 for all $u \in C$, and we have hence proved that $\theta = y - X\hat{\beta}(y) = P_C(y)$. To show that C is indeed a polyhedron, note that it can be written as

$$C = \bigcap_{i=1}^p (\{u \in \mathbb{R}^n : X_i^T u \leq \lambda\} \cap \{u \in \mathbb{R}^n : X_i^T u \geq -\lambda\}),$$

which is a finite intersection of half-spaces. \square

Showing that the lasso fit is the residual from projecting y onto a polyhedron is important, because it means that $X\hat{\beta}(y)$ is nonexpansive as a function of y , and hence continuous and almost differentiable, by Lemma 1. This establishes the conditions that are needed to apply Stein's formula for degrees of freedom.

In the next section, we define the equicorrelation set \mathcal{E} , and show that the lasso fit and solutions both have an explicit form in terms of \mathcal{E} . Following this, we derive an expression for the lasso degrees of freedom as a function of the equicorrelation set.

3.2. *The equicorrelation set.* According to Lemma 3, the lasso fit $X\hat{\beta}$ is always unique (because projection onto a closed convex set is unique). Therefore, even though the solution $\hat{\beta}$ is not necessarily unique, the optimal subgradient γ is unique, because it can be written entirely in terms of $X\hat{\beta}$, as shown by (13). We define the unique *equicorrelation set* \mathcal{E} as

$$(17) \quad \mathcal{E} = \{i \in \{1, \dots, p\} : |\gamma_i| = 1\}.$$

An alternative definition for the equicorrelation set is

$$(18) \quad \mathcal{E} = \{i \in \{1, \dots, p\} : |X_i^T(y - X\hat{\beta})| = \lambda\},$$

which explains its name, as \mathcal{E} can be thought of as the set of variables that have equal and maximal absolute inner product (or correlation for standardized variables) with the residual.

The set \mathcal{E} is a natural quantity to work with because we can express the lasso fit and the set of lasso solutions in terms of \mathcal{E} , by working directly from equation (13). First we let

$$(19) \quad s = \text{sign}(\gamma_{\mathcal{E}}) = \text{sign}(X_{\mathcal{E}}^T(y - X\hat{\beta})),$$

the signs of the inner products of the equicorrelation variables with the residual. Since $\hat{\beta}_{-\mathcal{E}} = 0$ by definition of the subgradient, the \mathcal{E} block of the KKT conditions can be rewritten as

$$(20) \quad X_{\mathcal{E}}^T(y - X_{\mathcal{E}}\hat{\beta}_{\mathcal{E}}) = \lambda s.$$

Because $\lambda s \in \text{row}(X_{\mathcal{E}})$, we can write $\lambda s = X_{\mathcal{E}}^T(X_{\mathcal{E}}^T)^+ \lambda s$, so rearranging (20) we get

$$X_{\mathcal{E}}^T X_{\mathcal{E}} \hat{\beta}_{\mathcal{E}} = X_{\mathcal{E}}^T(y - (X_{\mathcal{E}}^T)^+ \lambda s).$$

Therefore, the lasso fit $X\hat{\beta} = X_{\mathcal{E}}\hat{\beta}_{\mathcal{E}}$ is

$$(21) \quad X\hat{\beta} = X_{\mathcal{E}}(X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s),$$

and any lasso solution must be of the form

$$(22) \quad \hat{\beta}_{-\mathcal{E}} = 0 \quad \text{and} \quad \hat{\beta}_{\mathcal{E}} = (X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s) + b,$$

where $b \in \text{null}(X_{\mathcal{E}})$. In the case that $\text{null}(X_{\mathcal{E}}) = \{0\}$ —for example, this holds if $\text{rank}(X) = p$ —the lasso solution is unique and is given by (22) with $b = 0$. But in general, when $\text{null}(X_{\mathcal{E}}) \neq \{0\}$, it is important to note that not every $b \in \text{null}(X_{\mathcal{E}})$ necessarily leads to a lasso solution in (22); the vector b must also preserve the signs of the nonzero coefficients; that is, it must also satisfy

$$(23) \quad \begin{aligned} \text{sign}([(X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s)]_i + b_i) &= s_i \\ \text{for each } i \text{ such that } [(X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s)]_i + b_i &\neq 0. \end{aligned}$$

Otherwise, γ would not be a proper subgradient of $\|\hat{\beta}\|_1$.

3.3. Degrees of freedom in terms of the equicorrelation set. Using relatively simple arguments, we can derive a result on the lasso degrees of freedom in terms of the equicorrelation set. Our arguments build on the following key lemma.

LEMMA 4. *For any y, X and $\lambda \geq 0$, a lasso solution is given by*

$$(24) \quad \hat{\beta}_{-\mathcal{E}} = 0 \quad \text{and} \quad \hat{\beta}_{\mathcal{E}} = (X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s),$$

where \mathcal{E} and s are the equicorrelation set and signs, as defined in (17) and (19).

In other words, Lemma 4 says that the sign condition (23) is always satisfied by taking $b = 0$, regardless of the rank of X . This result is inspired by the LARS work of Efron et al. (2004), though it is not proved in the LARS paper; see Appendix B of Tibshirani (2011) for a proof.

Next we show that, almost everywhere in y , the equicorrelation set and signs are locally constant functions of y . To emphasize their functional dependence on y , we write them as $\mathcal{E}(y)$ and $s(y)$.

LEMMA 5. *For almost every $y \in \mathbb{R}^n$, there exists a neighborhood U of y such that $\mathcal{E}(y') = \mathcal{E}(y)$ and $s(y') = s(y)$ for all $y' \in U$.*

PROOF. Define

$$\mathcal{N} = \bigcup_{\mathcal{E}, s} \bigcup_{i \in \mathcal{E}} \{z \in \mathbb{R}^n : [(X_{\mathcal{E}})^+]_{(i, \cdot)}(z - (X_{\mathcal{E}}^T)^+\lambda s) = 0\},$$

where the first union above is taken over all subsets $\mathcal{E} \subseteq \{1, \dots, p\}$ and sign vectors $s \in \{-1, 1\}^{|\mathcal{E}|}$, but we exclude sets \mathcal{E} for which a row of $(X_{\mathcal{E}})^+$ is entirely zero. The set \mathcal{N} is a finite union of affine subspaces of dimension $n - 1$, and therefore has measure zero.

Let $y \notin \mathcal{N}$, and abbreviate the equicorrelation set and signs as $\mathcal{E} = \mathcal{E}(y)$ and $s = s(y)$. We may assume no row of $(X_{\mathcal{E}})^+$ is entirely zero. (Otherwise, this implies that $X_{\mathcal{E}}$ has a zero column, which implies that $\lambda = 0$, a trivial case for this lemma.) Therefore, as $y \notin \mathcal{N}$, this means that the lasso solution given in (24) satisfies $\hat{\beta}_i(y) \neq 0$ for every $i \in \mathcal{E}$.

Now, for a new point y' , consider defining

$$\hat{\beta}_{-\mathcal{E}}(y') = 0 \quad \text{and} \quad \hat{\beta}_{\mathcal{E}}(y') = (X_{\mathcal{E}})^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s).$$

We need to verify that this is indeed a solution at y' , and that the corresponding fit has equicorrelation set \mathcal{E} and signs s . First notice that, after a straightforward calculation,

$$X_{\mathcal{E}}^T(y' - X\hat{\beta}(y')) = X_{\mathcal{E}}^T(y' - X_{\mathcal{E}}(X_{\mathcal{E}})^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s)) = \lambda s.$$

Also, by the continuity of the function $f: \mathbb{R}^n \rightarrow \mathbb{R}^{p-|\mathcal{E}|}$,

$$f(x) = X_{-\mathcal{E}}^T(x - X_{\mathcal{E}}(X_{\mathcal{E}})^+(x - (X_{\mathcal{E}}^T)^+ \lambda s)),$$

there exists a neighborhood U_1 of y such that

$$\|X_{-\mathcal{E}}^T(y' - X\hat{\beta}(y'))\|_{\infty} = \|X_{-\mathcal{E}}^T(y' - X_{\mathcal{E}}(X_{\mathcal{E}})^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s))\|_{\infty} < \lambda$$

for all $y' \in U_1$. Hence $X\hat{\beta}(y')$ has equicorrelation set $\mathcal{E}(y') = \mathcal{E}$ and signs $s(y') = s$.

To check that $\hat{\beta}(y')$ is a lasso solution at y' , we consider the function $g: \mathbb{R}^n \rightarrow \mathbb{R}^{|\mathcal{E}|}$,

$$g(x) = (X_{\mathcal{E}})^+(x - (X_{\mathcal{E}}^T)^+ \lambda s).$$

The continuity of g implies that there exists a neighborhood U_2 of y such that

$$\begin{aligned} \hat{\beta}_i(y') &= [(X_{\mathcal{E}}^T)^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s)]_i \neq 0 \quad \text{for } i \in \mathcal{E}, \quad \text{and} \\ \text{sign}(\hat{\beta}_{\mathcal{E}}(y')) &= \text{sign}((X_{\mathcal{E}})^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s)) \end{aligned}$$

for each $y' \in U_2$. Defining $U = U_1 \cap U_2$ completes the proof. \square

This immediately implies the following theorem.

THEOREM 1 (Lasso degrees of freedom, equicorrelation set representation). *Assume that y follows a normal distribution (5). For any X and $\lambda \geq 0$, the lasso fit $X\hat{\beta}$ has degrees of freedom*

$$\text{df}(X\hat{\beta}) = \text{E}[\text{rank}(X_{\mathcal{E}})],$$

where $\mathcal{E} = \mathcal{E}(y)$ is the equicorrelation set of the lasso fit at y .

PROOF. By Lemmas 1 and 3 we know that $X\hat{\beta}(y)$ is continuous and almost differentiable, so we can use Stein's formula (10) for degrees of freedom. By Lemma 5, we know that $\mathcal{E} = \mathcal{E}(y)$ and $s = s(y)$ are locally constant for all $y \notin \mathcal{N}$. Therefore, taking the divergence of the fit in (21), we get

$$(\nabla \cdot X\hat{\beta})(y) = \text{tr}(X_{\mathcal{E}}(X_{\mathcal{E}})^+) = \text{rank}(X_{\mathcal{E}}).$$

Taking an expectation over y (and recalling that \mathcal{N} has measure zero) gives the result. \square

Next, we shift our focus to a different subset of variables: the active set \mathcal{A} . Unlike the equicorrelation set, the active set is not unique, as it depends on a particular choice of lasso solution. Though it may seem that such nonuniqueness could present complications, it turns out that all of the active sets share a special property; namely, the linear subspace $\text{col}(X_{\mathcal{A}})$ is the same for any choice of active set \mathcal{A} , almost everywhere in y . This invariance allows us to express the degrees of freedom of lasso fit in terms of the active set (or, more precisely, any active set).

3.4. *The active set.* Given a particular solution $\hat{\beta}$, we define the *active set* \mathcal{A} as

$$(25) \quad \mathcal{A} = \{i \in \{1, \dots, p\} : \hat{\beta}_i \neq 0\}.$$

This is also called the support of $\hat{\beta}$ and written $\mathcal{A} = \text{supp}(\hat{\beta})$. From (22), we can see that we always have $\mathcal{A} \subseteq \mathcal{E}$, and different active sets \mathcal{A} can be formed by choosing $b \in \text{null}(X_{\mathcal{E}})$ to satisfy the sign condition (23) and also

$$[(X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+\lambda s)]_i + b_i = 0 \quad \text{for } i \notin \mathcal{A}.$$

If $\text{rank}(X) = p$, then $b = 0$, so there is a unique active set, and furthermore $\mathcal{A} = \mathcal{E}$ for almost every $y \in \mathbb{R}^n$ (in particular, this last statement holds for $y \notin \mathcal{N}$, where \mathcal{N} is the set of measure zero set defined in the proof of Lemma 5). For the signs of the coefficients of active variables, we write

$$(26) \quad r = \text{sign}(\hat{\beta}_{\mathcal{A}}),$$

and we note that $r = s_{\mathcal{A}}$.

By similar arguments as those used to derive expression (21) for the fit in Section 3.2, the lasso fit can also be written as

$$(27) \quad X\hat{\beta} = (X_{\mathcal{A}})(X_{\mathcal{A}})^+(y - (X_{\mathcal{A}}^T)^+\lambda r)$$

for the active set \mathcal{A} and signs r of any lasso solution $\hat{\beta}$. If we could take the divergence of the fit in the expression above, and simply ignore the dependence of \mathcal{A} and r on y (treat them as constants), then this would give $(\nabla \cdot X\hat{\beta})(y) = \text{rank}(X_{\mathcal{A}})$. In the next section, we show that treating \mathcal{A} and r

as constants in (27) is indeed correct, for almost every y . This property then implies that the linear subspace $\text{col}(X_{\mathcal{A}})$ is invariant under any choice of active set \mathcal{A} , almost everywhere in y ; moreover, it implies that we can write the lasso degrees of freedom in terms of any active set.

3.5. *Degrees of freedom in terms of the active set.* We first establish a result on the local stability of $\mathcal{A}(y)$ and $r(y)$ [written in this way to emphasize their dependence on y , through a solution $\hat{\beta}(y)$].

LEMMA 6. *There is a set $\mathcal{M} \subseteq \mathbb{R}^n$, of measure zero, with the following property: for $y \notin \mathcal{M}$, and for any lasso solution $\hat{\beta}(y)$ with active set $\mathcal{A}(y)$ and signs $r(y)$, there is a neighborhood U of y such that every point $y' \in U$ yields a lasso solution $\hat{\beta}(y')$ with the same active set $\mathcal{A}(y') = \mathcal{A}(y)$ and the same active signs $r(y') = r(y)$.*

The proof is similar to that of Lemma 5, except that it is longer and somewhat more complicated, so it is delayed until Appendix A.3. Combined with expression (27) for the lasso fit, Lemma 6 now implies an invariance of the subspace spanned by the active variables.

LEMMA 7. *For the same set $\mathcal{M} \subseteq \mathbb{R}^n$ as in Lemma 6, and for any $y \notin \mathcal{M}$, the linear subspace $\text{col}(X_{\mathcal{A}})$ is invariant under all sets $\mathcal{A} = \mathcal{A}(y)$ defined in terms of a lasso solution $\hat{\beta}(y)$ at y .*

PROOF. Let $y \notin \mathcal{M}$, and let $\hat{\beta}(y)$ be a solution with active set $\mathcal{A} = \mathcal{A}(y)$ and signs $r = r(y)$. Let U be the neighborhood of y as constructed in the proof of Lemma 6; on this neighborhood, solutions exist with active set \mathcal{A} and signs r . Hence, recalling (27), we know that for every $y' \in U$,

$$X\hat{\beta}(y') = (X_{\mathcal{A}})(X_{\mathcal{A}})^+(y' - (X_{\mathcal{A}}^T)^+\lambda r).$$

Now suppose that \mathcal{A}^* and r^* are the active set and signs of another lasso solution at y . Then, by the same arguments, there is a neighborhood U^* of y such that

$$X\hat{\beta}(y') = (X_{\mathcal{A}^*})(X_{\mathcal{A}^*})^+(y' - (X_{\mathcal{A}^*}^T)^+\lambda r^*)$$

for all $y' \in U^*$. By the uniqueness of the fit, we have that for each $y' \in U \cap U^*$,

$$(X_{\mathcal{A}})(X_{\mathcal{A}})^+(y' - (X_{\mathcal{A}}^T)^+\lambda r) = (X_{\mathcal{A}^*})(X_{\mathcal{A}^*})^+(y' - (X_{\mathcal{A}^*}^T)^+\lambda r^*).$$

Since $U \cap U^*$ is open, for any $z \in \text{col}(X_{\mathcal{A}})$, there is an $\varepsilon > 0$ such that $y + \varepsilon z \in U \cap U^*$. Plugging $y' = y + \varepsilon z$ into the above equation implies that $z \in \text{col}(X_{\mathcal{A}^*})$, so $\text{col}(X_{\mathcal{A}}) \subseteq \text{col}(X_{\mathcal{A}^*})$. A similar argument gives $\text{col}(X_{\mathcal{A}^*}) \subseteq \text{col}(X_{\mathcal{A}})$, completing the proof. \square

Again, this immediately leads to the following theorem.

THEOREM 2 (Lasso degrees of freedom, active set representation). *Assume that y follows a normal distribution (5). For any X and $\lambda \geq 0$, the lasso fit $X\hat{\beta}$ has degrees of freedom*

$$\text{df}(X\hat{\beta}) = \text{E}[\text{rank}(X_{\mathcal{A}})],$$

where $\mathcal{A} = \mathcal{A}(y)$ is the active set corresponding to any lasso solution $\hat{\beta}(y)$ at y .

Note: By Lemma 7, $\text{rank}(X_{\mathcal{A}})$ is an invariant quantity, not depending on the choice of active set (coming from a lasso solution), for almost every y . This makes the above result well defined.

PROOF OF THEOREM 2. We can apply Stein's formula (10) for degrees of freedom, because $X\hat{\beta}(y)$ is continuous and almost differentiable by Lemmas 1 and 3. Let $\mathcal{A} = \mathcal{A}(y)$ and $r = r(y)$ be the active set and active signs of a lasso solution at $y \notin \mathcal{M}$, with \mathcal{M} as in Lemma 7. By this same lemma, there exists a lasso solution with active set \mathcal{A} and signs r at every point y' in some neighborhood U of y , and therefore, taking the divergence of the fit (27), we get

$$(\nabla \cdot X\hat{\beta})(y) = \text{tr}(X_{\mathcal{A}}(X_{\mathcal{A}})^+) = \text{rank}(X_{\mathcal{A}}).$$

Taking an expectation over y completes the proof. \square

REMARK (Equicorrelation set representation). The proof of Lemma 6 showed that, for almost every y , the equicorrelation set \mathcal{E} is actually the active set \mathcal{A} of the particular lasso solution defined in (24). Hence Theorem 1 can be viewed as a corollary of Theorem 2.

REMARK (Full column rank X). When $\text{rank}(X) = p$, the lasso solution is unique, and there is only one active set \mathcal{A} . And as the columns of X are linearly independent, we have $\text{rank}(X) = |\mathcal{A}|$, so the result of Theorem 2 reduces to

$$\text{df}(X\hat{\beta}) = \text{E}|\mathcal{A}|,$$

as shown in Zou, Hastie and Tibshirani (2007).

REMARK (The smallest active set). An interesting result on the lasso degrees of freedom was recently and independently obtained by Dossal et al. (2011). Their result states that, for a general X ,

$$\text{df}(X\hat{\beta}) = \text{E}|\mathcal{A}^*|,$$

where $|\mathcal{A}^*|$ is the smallest cardinality among all active sets of lasso solutions. This actually follows from Theorem 2, by noting that for any y there exists a lasso solution whose active set \mathcal{A}^* corresponds to linear independent

predictors $X_{\mathcal{A}^*}$, so $\text{rank}(X_{\mathcal{A}^*}) = |\mathcal{A}^*|$ [e.g., see Theorem 3 in Appendix B of Rosset, Zhu and Hastie (2004)], and furthermore, for almost every y no active set can have a cardinality smaller than $|\mathcal{A}^*|$, as this would contradict Lemma 7.

REMARK (The elastic net). Consider the elastic net problem [Zou and Hastie (2005)],

$$(28) \quad \hat{\beta} = \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \frac{\lambda_2}{2} \|\beta\|_2^2,$$

where we now have two tuning parameters $\lambda_1, \lambda_2 \geq 0$. Note that our notation above emphasizes the fact that there is always a unique solution to the elastic net criterion, regardless of the rank of X . This property (among others, such as stability and predictive ability) is considered an advantage of the elastic net over the lasso. We can rewrite the elastic net problem (28) as a (full column rank) lasso problem,

$$\hat{\beta} = \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2} \left\| \begin{pmatrix} y \\ 0 \end{pmatrix} - \begin{bmatrix} X \\ \sqrt{\lambda_2} I \end{bmatrix} \beta \right\|_2^2 + \lambda_1 \|\beta\|_1,$$

and hence it can be shown (although we omit the details) that the degrees of freedom of the elastic net fit is

$$\text{df}(X\hat{\beta}) = \text{E}[\text{tr}(X_{\mathcal{A}}(X_{\mathcal{A}}^T X_{\mathcal{A}} + \lambda_2 I)^{-1} X_{\mathcal{A}}^T)],$$

where $\mathcal{A} = \mathcal{A}(y)$ is the active set of the elastic net solution at y .

REMARK (The lasso with intercept). It is often more appropriate to include an (unpenalized) intercept coefficient in the lasso model, yielding the problem

$$(29) \quad (\hat{\beta}_0, \hat{\beta}) \in \underset{(\beta_0, \beta) \in \mathbb{R}^{p+1}}{\text{argmin}} \frac{1}{2} \|y - \beta_0 \mathbf{1} - X\beta\|_2^2 + \lambda \|\beta\|_1,$$

where $\mathbf{1} = (1, 1, \dots, 1) \in \mathbb{R}^n$ is the vector of all 1s. Defining $M = I - \mathbf{1}\mathbf{1}^T/n \in \mathbb{R}^{n \times n}$, we note that the fit of problem (29) can be written as $\hat{\beta}_0 \mathbf{1} + X\hat{\beta} = (I - M)y + MX\hat{\beta}$, and that $\hat{\beta}$ solves the usual lasso problem

$$\hat{\beta} \in \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2} \|My - MX\beta\|_2^2 + \lambda \|\beta\|_1.$$

Now it follows (again we omit the details) that the fit of the lasso problem with intercept (29) has degrees of freedom

$$\text{df}(\hat{\beta}_0 \mathbf{1} + X\hat{\beta}) = 1 + \text{E}[\text{rank}(MX_{\mathcal{A}})],$$

where $\mathcal{A} = \mathcal{A}(y)$ is the active set of a solution $\hat{\beta}(y)$ at y (these are the nonintercept coefficients). In other words, the degrees of freedom is one plus the expected dimension of the subspace spanned by the active variables, once

we have centered these variables. A similar result holds for an arbitrary set of unpenalized coefficients, by replacing M above with the projection onto the orthogonal complement of the column space of the unpenalized variables, and 1 above with the dimension of the column space of the unpenalized variables.

As mentioned in the [Introduction](#), a nice feature of the full column rank result (2) is its interpretability and its explicit nature. The general result is also explicit in the sense that an unbiased estimate of degrees of freedom can be achieved by computing the rank of a given matrix. In terms of interpretability, when $\text{rank}(X) = p$, the degrees of freedom of the lasso fit is $E|\mathcal{A}|$ —this says that, on average, the lasso “spends” the same number of parameters as does linear regression on $|\mathcal{A}|$ linearly independent predictor variables. Fortunately, a similar interpretation is possible in the general case: we showed in [Theorem 2](#) that for a general predictor matrix X , the degrees of freedom of the lasso fit is $E[\text{rank}(X_{\mathcal{A}})]$, the expected dimension of the linear subspace spanned by the active variables. Meanwhile, for the linear regression problem

$$(30) \quad \hat{\beta}_{\mathcal{A}} = \underset{\beta_{\mathcal{A}} \in \mathbb{R}^{|\mathcal{A}|}}{\text{argmin}} \|y - X_{\mathcal{A}}\beta_{\mathcal{A}}\|_2^2,$$

where we consider \mathcal{A} fixed, the degrees of freedom of the fit is $\text{tr}(X_{\mathcal{A}}(X_{\mathcal{A}})^+) = \text{rank}(X_{\mathcal{A}})$. In other words, the lasso adaptively selects a subset \mathcal{A} of the variables to use for a linear model of y , but on average it only “spends” the same number of parameters as would linear regression on the variables in \mathcal{A} , if \mathcal{A} was pre-specified.

How is this possible? Broadly speaking, the answer lies in the shrinkage due to the ℓ_1 penalty. Although the active set is chosen adaptively, the lasso does not estimate the active coefficients as aggressively as does the corresponding linear regression problem (30); instead, they are shrunk toward zero, and this adjusts for the adaptive selection. Differing views have been presented in the literature with respect to this feature of lasso shrinkage. On the one hand, for example, [Fan and Li \(2001\)](#) point out that lasso estimates suffer from bias due to the shrinkage of large coefficients, and motivate the nonconvex *SCAD* penalty as an attempt to overcome this bias. On the other hand, for example, [Loubes and Massart \(2004\)](#) discuss the merits of such shrunk estimates in model selection criteria, such as (9). In the current context, the shrinkage due to the ℓ_1 penalty is helpful in that it provides control over degrees of freedom. A more precise study of this idea is the topic of future work.

4. The generalized lasso. In this section we extend our degrees of freedom results to the generalized lasso problem, with an arbitrary predictor matrix X and penalty matrix D . As before, the KKT conditions play a central role, and we present these first. Also, many results that follow have

equivalent derivations from the perspective of the generalized lasso dual problem; see Appendix A.5. We remind the reader that D_R is used to extract to extract rows of D corresponding to an index set R .

4.1. *The KKT conditions and the underlying polyhedron.* The KKT conditions for the generalized lasso problem (3) are

$$(31) \quad X^T(y - X\hat{\beta}) = D^T\lambda\gamma,$$

$$(32) \quad \gamma_i \in \begin{cases} \{\text{sign}((D\hat{\beta})_i)\} & \text{if } (D\hat{\beta})_i \neq 0, \\ [-1, 1] & \text{if } (D\hat{\beta})_i = 0. \end{cases}$$

Now $\gamma \in \mathbb{R}^m$ is a subgradient of the function $f(x) = \|x\|_1$ evaluated at $x = D\hat{\beta}$. Similar to what we showed for the lasso, it follows from the KKT conditions that the generalized lasso fit is the residual from projecting y onto a polyhedron.

LEMMA 8. *For any X and $\lambda \geq 0$, the generalized lasso fit can be written as $X\hat{\beta}(y) = (I - P_C)(y)$, where $C \subseteq \mathbb{R}^n$ is the polyhedron*

$$C = \{u \in \mathbb{R}^n : X^T u = D^T w \text{ for } w \in \mathbb{R}^m, \|w\|_\infty \leq \lambda\}.$$

PROOF. The proof is quite similar to that of Lemma 3. As in (16), we want to show that

$$(33) \quad \langle X\hat{\beta}, y - X\hat{\beta} \rangle - \langle X^T u, \hat{\beta} \rangle \geq 0$$

for all $u \in C$, where C is as in the lemma. For the first term above, we can take an inner product with $\hat{\beta}$ on both sides of (31) to get $\langle X\hat{\beta}, y - X\hat{\beta} \rangle = \lambda \|D\hat{\beta}\|_1$, and furthermore,

$$\lambda \|D\hat{\beta}\|_1 = \max_{\|w\|_\infty \leq \lambda} \langle w, D\hat{\beta} \rangle = \max_{\|w\|_\infty \leq \lambda} \langle D^T w, \hat{\beta} \rangle.$$

Therefore (33) holds if $X^T u = D^T w$ for some $\|w\|_\infty \leq \lambda$, in other words, if $u \in C$. To show that C is a polyhedron, note that we can write it as $C = (X^T)^{-1}(D^T(B))$ where $(X^T)^{-1}$ is taken to mean the inverse image under the linear map X^T , and $B = \{w \in \mathbb{R}^m : \|w\|_\infty \leq \lambda\}$, a hypercube in \mathbb{R}^m . Clearly B is a polyhedron, and the image or inverse image of a polyhedron under a linear map is still a polyhedron. \square

As with the lasso, this lemma implies that the generalized lasso fit $X\hat{\beta}(y)$ is nonexpansive, and therefore continuous and almost differentiable as a function of y , by Lemma 1. This is important because it allows us to use Stein's formula when computing degrees of freedom.

In the next section we define the boundary set \mathcal{B} , and derive expressions for the generalized lasso fit and solutions in terms of \mathcal{B} . The following section defines the active set \mathcal{A} in the generalized lasso context, and again gives expressions for the fit and solutions in terms of \mathcal{A} . Though neither \mathcal{B} nor \mathcal{A}

are necessarily unique for the generalized lasso problem, any choice of \mathcal{B} or \mathcal{A} generates a special invariant subspace (similar to the case for the active sets in the lasso problem). We are subsequently able to express the degrees of freedom of the generalized lasso fit in terms of any boundary set \mathcal{B} , or any active set \mathcal{A} .

4.2. *The boundary set.* Like the lasso, the generalized lasso fit $X\hat{\beta}$ is always unique (following from Lemma 8, and the fact that projection onto a closed convex set is unique). However, unlike the lasso, the optimal subgradient γ in the generalized lasso problem is not necessarily unique. In particular, if $\text{rank}(D) < m$, then the optimal subgradient γ is not uniquely determined by conditions (31) and (32). Given a subgradient γ satisfying (31) and (32) for some $\hat{\beta}$, we define the *boundary set* \mathcal{B} as

$$\mathcal{B} = \{i \in \{1, \dots, m\} : |\gamma_i| = 1\}.$$

This generalizes the notion of the equicorrelation set \mathcal{E} in the lasso problem [though, as just noted, the set \mathcal{B} is not necessarily unique unless $\text{rank}(D) = m$]. We also define

$$s = \gamma_{\mathcal{B}}.$$

Now we focus on writing the generalized lasso fit and solutions in terms of \mathcal{B} and s . Abbreviating $P = P_{\text{null}(D_{-\mathcal{B}})}$, note that we can expand $PD^T\lambda\gamma = PD_{\mathcal{B}}^T\lambda s + PD_{-\mathcal{B}}^T\lambda\gamma_{-\mathcal{B}} = PD_{\mathcal{B}}^T\lambda s$. Therefore, multiplying both sides of (31) by P yields

$$(34) \quad PX^T(y - X\hat{\beta}) = PD_{\mathcal{B}}^T\lambda s.$$

Since $PD_{\mathcal{B}}^T\lambda s \in \text{col}(PX^T)$, we can write $PD_{\mathcal{B}}^T\lambda s = (PX^T)(PX^T)^+PD_{\mathcal{B}}^T\lambda s = (PX^T)(PX^T)^+D_{\mathcal{B}}^T\lambda s$. Also, we have $D_{-\mathcal{B}}\hat{\beta} = 0$ by definition of \mathcal{B} , so $P\hat{\beta} = \hat{\beta}$. These two facts allow us to rewrite (34) as

$$PX^T X P \hat{\beta} = PX^T(y - (PX^T)^+D_{\mathcal{B}}^T\lambda s),$$

and hence the fit $X\hat{\beta} = X P \hat{\beta}$ is

$$(35) \quad X\hat{\beta} = (X P_{\text{null}(D_{-\mathcal{B}})})(X P_{\text{null}(D_{-\mathcal{B}})})^+(y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s),$$

where we have un-abbreviated $P = P_{\text{null}(D_{-\mathcal{B}})}$. Further, any generalized lasso solution is of the form

$$(36) \quad \hat{\beta} = (X P_{\text{null}(D_{-\mathcal{B}})})^+(y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) + b,$$

where $b \in \text{null}(X P_{\text{null}(D_{-\mathcal{B}})})$. Multiplying the above equation by $D_{-\mathcal{B}}$, and recalling that $D_{-\mathcal{B}}\hat{\beta} = 0$, reveals that $b \in \text{null}(D_{-\mathcal{B}})$; hence $b \in \text{null}(X P_{\text{null}(D_{-\mathcal{B}})}) \cap \text{null}(D_{-\mathcal{B}}) = \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$. In the case that $\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}) = \{0\}$, the generalized lasso solution is unique and is given by (36) with $b = 0$. This occurs when $\text{rank}(X) = p$, for example. Otherwise, any $b \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ gives a generalized lasso solution in (36) as long as

it also satisfies the sign condition

$$(37) \quad \begin{aligned} & \text{sign}(D_i(XP_{\text{null}(D_{-\mathcal{B}})})^+(y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) + D_ib) = s_i \\ & \text{for each } i \in \mathcal{B} \text{ such that } D_i(XP_{\text{null}(D_{-\mathcal{B}})})^+ \\ & \quad \times (y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) + D_ib \neq 0, \end{aligned}$$

necessary to ensure that γ is a proper subgradient of $\|D\hat{\beta}\|_1$.

4.3. *The active set.* We define the *active set* of a particular solution $\hat{\beta}$ as

$$\mathcal{A} = \{i \in \{1, \dots, m\} : (D\hat{\beta})_i \neq 0\},$$

which can be alternatively expressed as $\mathcal{A} = \text{supp}(D\hat{\beta})$. If $\hat{\beta}$ corresponds to a subgradient with boundary set \mathcal{B} and signs s , then $\mathcal{A} \subseteq \mathcal{B}$; in particular, given \mathcal{B} and s , different active sets \mathcal{A} can be generated by taking $b \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ such that (37) is satisfied, and also

$$D_i(XP_{\text{null}(D_{-\mathcal{B}})})^+(y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) + D_ib = 0 \quad \text{for } i \in \mathcal{B} \setminus \mathcal{A}.$$

If $\text{rank}(X) = p$, then $b = 0$, and there is only one active set \mathcal{A} ; however, in this case, \mathcal{A} can still be a strict subset of \mathcal{B} . This is quite different from the lasso problem, wherein $\mathcal{A} = \mathcal{E}$ for almost every y whenever $\text{rank}(X) = p$. [Note that in the generalized lasso problem, $\text{rank}(X) = p$ implies that \mathcal{A} is unique but implies nothing about the uniqueness of \mathcal{B} —this is determined by the rank of D . The boundary set \mathcal{B} is not necessarily unique if $\text{rank}(D) < m$, and in this case we may have $D_i(XP_{\text{null}(D_{-\mathcal{B}})})^+ = 0$ for some $i \in \mathcal{B}$, which certainly implies that $i \notin \mathcal{A}$ for any $y \in \mathbb{R}^n$. Hence some boundary sets may not correspond to active sets at any y .] We denote the signs of the active entries in $D\hat{\beta}$ by

$$r = \text{sign}(D_{\mathcal{A}}\hat{\beta}),$$

and we note that $r = s_{\mathcal{A}}$.

Following the same arguments as those leading up to the expression for the fit (35) in Section 4.2, we can alternatively express the generalized lasso fit as

$$(38) \quad X\hat{\beta} = (XP_{\text{null}(D_{-\mathcal{A}})})(XP_{\text{null}(D_{-\mathcal{A}})})^+(y - (P_{\text{null}(D_{-\mathcal{A}})}X^T)^+D_{\mathcal{A}}^T\lambda r),$$

where \mathcal{A} and r are the active set and signs of any solution. Computing the divergence of the fit in (38), and pretending that \mathcal{A} and r are constants (not depending on y), gives $(\nabla \cdot X\hat{\beta})(y) = \dim(\text{col}(XP_{\text{null}(D_{-\mathcal{A}})})) = \dim(X(\text{null}(D_{-\mathcal{A}})))$. The same logic applied to (35) gives $(\nabla \cdot X\hat{\beta})(y) = \dim(X(\text{null}(D_{-\mathcal{B}})))$. The next section shows that, for almost every y , the quantities \mathcal{A}, r or \mathcal{B}, s can indeed be treated as locally constant in expressions (38) or (35), respectively. We then prove that linear subspaces

$X(\text{null}(D_{-\mathcal{B}})), X(\text{null}(D_{-\mathcal{A}}))$ are invariant under all choices of boundary sets \mathcal{B} , respectively active sets \mathcal{A} , and that the two subspaces are in fact equal, for almost every y . Furthermore, we express the generalized lasso degrees of freedom in terms of any boundary set or any active set.

4.4. *Degrees of freedom.* We call $(\gamma(y), \hat{\beta}(y))$ an *optimal pair* provided that $\gamma(y)$ and $\hat{\beta}(y)$ jointly satisfy the KKT conditions, (31) and (32), at y . For such a pair, we consider its boundary set $\mathcal{B}(y)$, boundary signs $s(y)$, active set $\mathcal{A}(y)$, active signs $r(y)$, and show that these sets and sign vectors possess a kind of local stability.

LEMMA 9. *There exists a set $\mathcal{N} \subseteq \mathbb{R}^n$, of measure zero, with the following property: for $y \notin \mathcal{N}$, and for any optimal pair $(\gamma(y), \hat{\beta}(y))$ with boundary set $\mathcal{B}(y)$, boundary signs $s(y)$, active set $\mathcal{A}(y)$, and active signs $r(y)$, there is a neighborhood U of y such that each point $y' \in U$ yields an optimal pair $(\gamma(y'), \hat{\beta}(y'))$ with the same boundary set $\mathcal{B}(y') = \mathcal{B}(y)$, boundary signs $s(y') = s(y)$, active set $\mathcal{A}(y') = \mathcal{A}(y)$ and active signs $r(y') = r(y)$.*

The proof is delayed to Appendix A.4, mainly because of its length. Now Lemma 9, used together with expressions (35) and (38) for the generalized lasso fit, implies an invariance in representing a (particularly important) linear subspace.

LEMMA 10. *For the same set $\mathcal{N} \subseteq \mathbb{R}^n$ as in Lemma 9, and for any $y \notin \mathcal{N}$, the linear subspace $L = X(\text{null}(D_{-\mathcal{B}}))$ is invariant under all boundary sets $\mathcal{B} = \mathcal{B}(y)$ defined in terms of an optimal subgradient at $\gamma(y)$ at y . The linear subspace $L' = X(\text{null}(D_{-\mathcal{A}}))$ is also invariant under all choices of active sets $\mathcal{A} = \mathcal{A}(y)$ defined in terms of a generalized lasso solution $\hat{\beta}(y)$ at y . Finally, the two subspaces are equal, $L = L'$.*

PROOF. Let $y \notin \mathcal{N}$, and let $\gamma(y)$ be an optimal subgradient with boundary set $\mathcal{B} = \mathcal{B}(y)$ and signs $s = s(y)$. Let U be the neighborhood of y over which optimal subgradients exist with boundary set \mathcal{B} and signs s , as given by Lemma 9. Recalling the expression for the fit (35), we have that for every $y' \in U$

$$X\hat{\beta}(y') = (XP_{\text{null}(D_{-\mathcal{B}})})(XP_{\text{null}(D_{-\mathcal{B}})})^+(y' - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s).$$

If $\hat{\beta}(y)$ is a solution with active set $\mathcal{A} = \mathcal{A}(y)$ and signs $r = r(y)$, then again by Lemma 9 there is a neighborhood V of y such that each point $y' \in V$ yields a solution with active set \mathcal{A} and signs r . [Note that V and U are not necessarily equal unless $\gamma(y)$ and $\hat{\beta}(y)$ jointly satisfy the KKT conditions at y .] Therefore, recalling (35), we have

$$X\hat{\beta}(y') = (XP_{\text{null}(D_{-\mathcal{A}})})(XP_{\text{null}(D_{-\mathcal{A}})})^+(y' - (P_{\text{null}(D_{-\mathcal{A}})}X^T)^+D_{\mathcal{A}}^T\lambda r)$$

for each $y' \in V$. The uniqueness of the generalized lasso fit now implies that

$$\begin{aligned} & (XP_{\text{null}(D_{-\mathcal{B}})})(XP_{\text{null}(D_{-\mathcal{B}})})^+(y' - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) \\ &= (XP_{\text{null}(D_{-\mathcal{A}})})(XP_{\text{null}(D_{-\mathcal{A}})})^+(y' - (P_{\text{null}(D_{-\mathcal{A}})}X^T)^+D_{\mathcal{A}}^T\lambda r) \end{aligned}$$

for all $y' \in U \cap V$. As $U \cap V$ is open, for any $z \in \text{col}(XP_{\text{null}(D_{-\mathcal{B}})})$, there exists an $\varepsilon > 0$ such that $y + \varepsilon z \in U \cap V$. Plugging $y' = y + \varepsilon z$ into the equation above reveals that $z \in \text{col}(XP_{\text{null}(D_{-\mathcal{A}})})$, hence $\text{col}(XP_{\text{null}(D_{-\mathcal{B}})}) \subseteq \text{col}(XP_{\text{null}(D_{-\mathcal{A}})})$. The reverse inclusion follows similarly, and therefore $\text{col}(XP_{\text{null}(D_{-\mathcal{B}})}) = \text{col}(XP_{\text{null}(D_{-\mathcal{A}})})$. Finally, the same strategy can be used to show that these linear subspaces are unchanged for any choice of boundary set $\mathcal{B} = \mathcal{B}(y)$, coming from an optimal subgradient at y and for any choice of active set $\mathcal{A} = \mathcal{A}(y)$ coming from a solution at y . Noticing that $\text{col}(MP_{\text{null}(N)}) = M(\text{null}(N))$ for matrices M, N gives the result as stated in the lemma. \square

This local stability result implies the following theorem.

THEOREM 3 (Generalized lasso degrees of freedom). *Assume that y follows a normal distribution (5). For any X, D and $\lambda \geq 0$, the degrees of freedom of the generalized lasso fit can be expressed as*

$$\text{df}(X\hat{\beta}) = \text{E}[\dim(X(\text{null}(D_{-\mathcal{B}})))],$$

where $\mathcal{B} = \mathcal{B}(y)$ is the boundary set corresponding to any optimal subgradient $\gamma(y)$ of the generalized lasso problem at y . We can alternatively express degrees of freedom as

$$\text{df}(X\hat{\beta}) = \text{E}[\dim(X(\text{null}(D_{-\mathcal{A}})))],$$

with $\mathcal{A} = \mathcal{A}(y)$ being the active set corresponding to any generalized lasso solution $\hat{\beta}(y)$ at y .

Note: Lemma 10 implies that for almost every $y \in \mathbb{R}^n$, for any \mathcal{B} defined in terms of an optimal subgradient, and for any \mathcal{A} defined in terms of a generalized lasso solution, $\dim(X(\text{null}(D_{-\mathcal{B}}))) = \dim(X(\text{null}(D_{-\mathcal{A}})))$. This makes the above expressions for degrees of freedom well defined.

PROOF OF THEOREM 3. First, the continuity and almost differentiability of $X\hat{\beta}(y)$ follow from Lemmas 1 and 8, so we can use Stein's formula (10) for degrees of freedom. Let $y \notin \mathcal{N}$, where \mathcal{N} is the set of measure zero as in Lemma 6. If $\mathcal{B} = \mathcal{B}(y)$ and $s = s(y)$ are the boundary set and signs of an optimal subgradient at y , then by Lemma 10 there is a neighborhood U of y such that each point $y' \in U$ yields an optimal subgradient with boundary

set \mathcal{B} and signs s . Therefore, taking the divergence of the fit in (35),

$$(\nabla \cdot X\hat{\beta})(y) = \text{tr}(P_{X(\text{null}(D_{-\mathcal{B}}))}) = \dim(X(\text{null}(D_{-\mathcal{B}}))),$$

and taking an expectation over y gives the first expression in the theorem.

Similarly, if $\mathcal{A} = \mathcal{A}(y)$ and $r = r(y)$ are the active set and signs of a generalized lasso solution at y , then by Lemma 10 there exists a solution with active set \mathcal{A} and signs r at each point y' in some neighborhood V of y . The divergence of the fit in (38) is hence

$$(\nabla \cdot X\hat{\beta})(y) = \text{tr}(P_{X(\text{null}(D_{-\mathcal{A}}))}) = \dim(X(\text{null}(D_{-\mathcal{A}}))),$$

and taking an expectation over y gives the second expression. \square

REMARK (Full column rank X). If $\text{rank}(X) = p$, then $\dim(X(L)) = \dim(L)$ for any linear subspace L , so the results of Theorem 3 reduce to

$$\text{df}(X\hat{\beta}) = \text{E}[\text{nullity}(D_{-\mathcal{B}})] = \text{E}[\text{nullity}(D_{-\mathcal{A}})].$$

The first equality above was shown in Tibshirani and Taylor (2011). Analyzing the null space of $D_{-\mathcal{B}}$ (equivalently, $D_{-\mathcal{A}}$) for specific choices of D then gives interpretable results on the degrees of freedom of the fused lasso and trend filtering fits as mentioned in the introduction. It is important to note that, as $\text{rank}(X) = p$, the active set \mathcal{A} is unique, but not necessarily equal to the boundary set \mathcal{B} [since \mathcal{B} can be nonunique if $\text{rank}(D) < m$].

REMARK (The lasso). If $D = I$, then $X(\text{null}(D_{-S})) = \text{col}(X_S)$ for any subset $S \subseteq \{1, \dots, p\}$. Therefore the results of Theorem 3 become

$$\text{df}(X\hat{\beta}) = \text{E}[\text{rank}(X_{\mathcal{B}})] = \text{E}[\text{rank}(X_{\mathcal{A}})],$$

which match the results of Theorems 1 and 2 (recall that for the lasso the boundary set \mathcal{B} is exactly the same as equicorrelation set \mathcal{E}).

REMARK (The smallest active set). Recent and independent work of Vaiteer et al. (2011) shows that, for arbitrary X, D and for any y , there exists a generalized lasso solution whose active set \mathcal{A}^* satisfies

$$\text{null}(X) \cap \text{null}(D_{-\mathcal{A}^*}) = \{0\}.$$

(Calling \mathcal{A}^* the “smallest” active set is somewhat of an abuse of terminology, but it is the smallest in terms of the above intersection.) The authors then prove that, for any X, D , the generalized lasso fit has degrees of freedom

$$\text{df}(X\hat{\beta}) = \text{E}[\text{nullity}(D_{-\mathcal{A}^*})],$$

with \mathcal{A}^* the special active set as above. This matches the active set result of Theorem 3 applied to \mathcal{A}^* , since $\dim(X(\text{null}(D_{-\mathcal{A}^*})) = \text{nullity}(D_{-\mathcal{A}^*})$ for this special active set.

We conclude this section by comparing the active set result of Theorem 3 to degrees of freedom in a particularly relevant equality constrained linear

regression problem (this comparison is similar to that made in lasso case, given at the end of Section 3). The result states that the generalized lasso fit has degrees of freedom $E[\dim(X(\text{null}(D_{-\mathcal{A}})))]$, where $\mathcal{A} = \mathcal{A}(y)$ is the active set of a generalized lasso solution at y . In other words, the complement of \mathcal{A} gives the rows of D that are orthogonal to some generalized lasso solution. Now, consider the equality constrained linear regression problem

$$(39) \quad \hat{\beta} \in \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \|y - X\beta\|_2^2 \quad \text{subject to } D_{-\mathcal{A}}\beta = 0,$$

in which the set \mathcal{A} is fixed. It is straightforward to verify that the fit of this problem is the projection map onto $\text{col}(XP_{\text{null}(D_{-\mathcal{A}})}) = X(\text{null}(D_{-\mathcal{A}}))$, and hence has degrees of freedom $\dim(X(\text{null}(D_{-\mathcal{A}})))$. This means that the generalized lasso fits a linear model of y , and simultaneously makes the coefficients orthogonal to an adaptive subset \mathcal{A} of the rows of D , but on average it only uses the same number of parameters as does the corresponding equality constrained linear regression problem (39), in which \mathcal{A} is pre-specified.

This seemingly paradoxical statement can be explained by the shrinkage due to the ℓ_1 penalty. Even though the active set \mathcal{A} is chosen adaptively based on y , the generalized lasso does not estimate the coefficients as aggressively as does the equality constrained linear regression problem (39), but rather, it shrinks them toward zero. Roughly speaking, this shrinkage can be viewed as a “deficit” in degrees of freedom, which makes up for the “surplus” attributed to the adaptive selection. We study this idea more precisely in a future paper.

5. Discussion. We showed that the degrees of freedom of the lasso fit, for an arbitrary predictor matrix X , is equal to $E[\text{rank}(X_{\mathcal{A}})]$. Here $\mathcal{A} = \mathcal{A}(y)$ is the active set of any lasso solution at y , that is, $\mathcal{A}(y) = \text{supp}(\hat{\beta}(y))$. This result is well defined, since we proved that any active set \mathcal{A} generates the same linear subspace $\text{col}(X_{\mathcal{A}})$, almost everywhere in y . In fact, we showed that for almost every y , and for any active set \mathcal{A} of a solution at y , the lasso fit can be written as

$$X\hat{\beta}(y') = P_{\text{col}(X_{\mathcal{A}})}(y') + c$$

for all y' in a neighborhood of y , where c is a constant (it does not depend on y'). This draws an interesting connection to linear regression, as it shows that locally the lasso fit is just a translation of the linear regression fit of on $X_{\mathcal{A}}$. The same results (on degrees of freedom and local representations of the fit) hold when the active set \mathcal{A} is replaced by the equicorrelation set \mathcal{E} .

Our results also extend to the generalized lasso problem, with an arbitrary predictor matrix X and arbitrary penalty matrix D . We showed that degrees of freedom of the generalized lasso fit is $E[\dim(X(\text{null}(D_{-\mathcal{A}})))]$, with $\mathcal{A} = \mathcal{A}(y)$ being the active set of any generalized lasso solution at y , that is,

$\mathcal{A}(y) = \text{supp}(D\hat{\beta}(y))$. As before, this result is well defined because any choice of active set \mathcal{A} generates the same linear subspace $X(\text{null}(D_{-\mathcal{A}}))$, almost everywhere in y . Furthermore, for almost every y , and for any active set of a solution at y , the generalized lasso fit satisfies

$$X\hat{\beta}(y') = P_{X(\text{null}(D_{-\mathcal{A}}))}(y') + c$$

for all y' in a neighborhood of y , where c is a constant (not depending on y). This again reveals an interesting connection to linear regression, since it says that locally the generalized lasso fit is a translation of the linear regression fit on X , with the coefficients β subject to $D_{-\mathcal{A}}\beta = 0$. The same statements hold with the active set \mathcal{A} replaced by the boundary set \mathcal{B} of an optimal subgradient.

We note that our results provide practically useful estimates of degrees of freedom. For the lasso problem, we can use $\text{rank}(X_{\mathcal{A}})$ as an unbiased estimate of degrees of freedom, with \mathcal{A} being the active set of a lasso solution. To emphasize what has already been said, here we can actually choose any active set (i.e., any solution), because all active sets give rise to the same $\text{rank}(X_{\mathcal{A}})$, except for y in a set of measure zero. This is important, since different algorithms for the lasso can produce different solutions with different active sets. For the generalized lasso problem, an unbiased estimate for degrees of freedom is given by $\dim(X(\text{null}(D_{-\mathcal{A}}))) = \text{rank}(XP_{\text{null}(D_{-\mathcal{A}})})$, where \mathcal{A} is the active set of a generalized lasso solution. This estimate is the same, regardless of the choice of active set (i.e., choice of solution), for almost every y . Hence any algorithm can be used to compute a solution.

APPENDIX A: PROOFS AND TECHNICAL ARGUMENTS

A.1. Proof of Lemma 1. The proof relies on the fact that the projection $P_C(x)$ of $x \in \mathbb{R}^n$ onto a closed convex set $C \subseteq \mathbb{R}^n$ satisfies

$$(40) \quad \langle x - P_C(x), P_C(x) - u \rangle \geq 0 \quad \text{for any } u \in C.$$

First, we prove the statement for the projection map. Note that

$$\begin{aligned} & \|P_C(x) - P_C(y)\|_2^2 \\ &= \langle P_C(x) - x + y - P_C(y) + x - y, P_C(x) - P_C(y) \rangle \\ &= \langle P_C(x) - x, P_C(x) - P_C(y) \rangle + \langle y - P_C(y), P_C(x) - P_C(y) \rangle \\ &\quad + \langle x - y, P_C(x) - P_C(y) \rangle \\ &\leq \langle x - y, P_C(x) - P_C(y) \rangle \\ &\leq \|x - y\|_2 \|P_C(x) - P_C(y)\|_2, \end{aligned}$$

where the first inequality follows from (40), and the second is by Cauchy-Schwarz. Dividing both sides by $\|P_C(x) - P_C(y)\|_2$ gives the result.

Now, for the residual map, the steps are similar.

$$\begin{aligned}
& \|(I - P_C)(x) - (I - P_C)(y)\|_2^2 \\
&= \langle P_C(y) - P_C(x) + x - y, x - P_C(x) + P_C(y) - y \rangle \\
&= \langle P_C(y) - P_C(x), x - P_C(x) \rangle + \langle P_C(y) - P_C(x), P_C(y) - y \rangle \\
&\quad + \langle x - y, x - P_C(x) + P_C(y) - y \rangle \\
&\leq \langle x - y, x - P_C(x) + P_C(y) - y \rangle \\
&\leq \|x - y\|_2 \|(I - P_C)(x) - (I - P_C)(y)\|_2.
\end{aligned}$$

Again the two inequalities are from (40) and Cauchy–Schwarz, respectively, and dividing both sides by $\|(I - P_C)(x) - (I - P_C)(y)\|_2$ gives the result.

We have shown that P_C and $I - P_C$ are Lipschitz (with constant 1); they are therefore continuous, and almost differentiability follows from the standard proof of the fact that a Lipschitz function is differentiable almost everywhere.

A.2. Proof of Lemma 2. We write \mathcal{F} to denote the set of faces of C . To each face $F \in \mathcal{F}$, there is an associated normal cone $N(F)$, defined as

$$N(F) = \left\{ x \in \mathbb{R}^n : F = \operatorname{argmax}_{y \in C} x^T y \right\}.$$

The normal cone of F satisfies $N(F) = P_C^{-1}(u) - u$ for any $u \in \operatorname{relint}(F)$. [We use $\operatorname{relint}(A)$ to denote the relative interior of a set A , and $\operatorname{relbd}(A)$ to denote its relative boundary.]

Define the set

$$\mathcal{S} = \bigcup_{F \in \mathcal{F}} (\operatorname{relint}(F) + \operatorname{relint}(N(F))).$$

Because C is a polyhedron, we have that $\dim(F) + \dim(N(F)) = n$ for each $F \in \mathcal{F}$, and therefore each $U_F = \operatorname{relint}(F) + \operatorname{relint}(N(F))$ is an open set in \mathbb{R}^n .

Now let $x \in \mathcal{S}$. We have $x \in U_F$ for some $F \in \mathcal{F}$, and by construction $P_C(U_F) = \operatorname{relint}(F)$. Furthermore, we claim that projecting $x \in U_F$ onto C is the same as projecting x onto the affine hull of F , that is, $P_C(U_F) = P_{\operatorname{aff}(F)}(U_F)$. Otherwise there is some $y \in U_F$ with $P_C(y) \neq P_{\operatorname{aff}(F)}(y)$, and as $\operatorname{aff}(F) \supseteq F$, this means that $\|y - P_{\operatorname{aff}(F)}(y)\|_2 < \|y - P_C(y)\|_2$. By definition of $\operatorname{relint}(F)$, there is some $\alpha \in (0, 1)$ such that $u = \alpha P_C(y) + (1 - \alpha) P_{\operatorname{aff}(F)}(y) \in F$. But $\|y - u\|_2 < \alpha \|y - P_C(y)\|_2 + (1 - \alpha) \|y - P_{\operatorname{aff}(F)}(y)\|_2 < \|y - P_C(y)\|_2$, which is a contradiction. This proves the claim, and writing $\operatorname{aff}(F) = L + a$, we have

$$P_C(y) = P_L(y - a) + a \quad \text{for } y \in U_F,$$

as desired.

It remains to show that $\mathcal{S}^c = \mathbb{R}^n \setminus \mathcal{S}$ has measure zero. Note that \mathcal{S}^c contains points of the form $u + x$, where either:

- (1) $u \in \text{relbd}(F), x \in N(F)$ for some F with $\dim(F) \geq 1$; or
- (2) $u \in \text{relint}(F), x \in \text{relbd}(N(F))$ for some $F \neq C$.

In the first type of points above, vertices are excluded because $\text{relbd}(F) = \emptyset$ when F is a vertex. In the second type, C is excluded because $\text{relbd}(N(C)) = \emptyset$. The lattice structure of \mathcal{F} tells us that for any face $F \in \mathcal{F}$, we can write $\text{relbd}(F) = \bigcup_{G \in \mathcal{F}, G \subsetneq F} \text{relint}(G)$. This, and the fact that the normal cones have the opposite partial ordering as the faces, imply that points of the first type above can be written as $u' + x'$ with $u' \in \text{relint}(G)$ and $x' \in N(G)$ for some $G \subsetneq F$. Note that actually we must have $x' \in \text{relbd}(N(G))$ because otherwise we would have $u' + x' \in \mathcal{S}$. Therefore it suffices to consider points of the second type alone, and \mathcal{S}^c can be written as

$$\mathcal{S}^c = \bigcup_{F \in \mathcal{F}, F \neq C} (\text{relint}(F) + \text{relbd}(N(F))).$$

As C is a polyhedron, the set \mathcal{F} of its faces is finite, and $\dim(\text{relbd}(N(F))) \leq n - \dim(F) - 1$ for each $F \in \mathcal{F}, F \neq C$. Therefore \mathcal{S}^c is a finite union of sets of dimension $\leq n - 1$, and hence has measure zero.

A.3. Proof of Lemma 6. First some notation. For $S \subseteq \{1, \dots, k\}$, define the function $\pi_S: \mathbb{R}^k \rightarrow \mathbb{R}^{|S|}$ by $\pi_S(x) = x_S$. So π_S just extracts the coordinates in S .

Now let

$$\mathcal{M} = \bigcup_{\mathcal{E}, s} \bigcup_{\mathcal{A} \in Z(\mathcal{E})} \{z \in \mathbb{R}^n : P_{[\pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}}))]^\perp} [(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)} (z - (X_{\mathcal{E}}^T)^+ \lambda s) = 0\}.$$

The first union is taken over all possible subsets $\mathcal{E} \subseteq \{1, \dots, p\}$ and all sign vectors $s \in \{-1, 1\}^{|\mathcal{E}|}$; as for the second union, we define for a fixed subset \mathcal{E}

$$Z(\mathcal{E}) = \{\mathcal{A} \subseteq \mathcal{E} : P_{[\pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}}))]^\perp} [(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)} \neq 0\}.$$

Notice that \mathcal{M} is a finite union of affine subspace of dimension $\leq n - 1$, and hence has measure zero.

Let $y \notin \mathcal{M}$, and let $\hat{\beta}(y)$ be a lasso solution, abbreviating $\mathcal{A} = \mathcal{A}(y)$ and $r = r(y)$ for the active set and active signs. Also write $\mathcal{E} = \mathcal{E}(y)$ and $s = s(y)$ for the equicorrelation set and equicorrelation signs of the fit. We know from (22) that we can write

$$\hat{\beta}_{-\mathcal{E}}(y) = 0 \quad \text{and} \quad \hat{\beta}_{\mathcal{E}}(y) = (X_{\mathcal{E}})^+(y - (X_{\mathcal{E}}^T)^+ \lambda s) + b,$$

where $b \in \text{null}(X_{\mathcal{E}})$ is such that

$$\hat{\beta}_{\mathcal{E} \setminus \mathcal{A}}(y) = [(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)} (y - (X_{\mathcal{E}}^T)^+ \lambda s) + b_{-\mathcal{A}} = 0.$$

In other words,

$$[(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)}(y - (X_{\mathcal{E}}^T)^+ \lambda s) = -b_{-\mathcal{A}} \in \pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}})),$$

so projecting onto the orthogonal complement of the linear subspace $\pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}}))$ gives zero,

$$P_{[\pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}}))]^\perp} [(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)}(y - (X_{\mathcal{E}}^T)^+ \lambda s) = 0.$$

Since $y \notin \mathcal{M}$, we know that

$$P_{[\pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}}))]^\perp} [(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)} = 0,$$

and finally, this can be rewritten as

$$(41) \quad \text{col}([(X_{\mathcal{E}})^+]_{(-\mathcal{A}, \cdot)}) \subseteq \pi_{-\mathcal{A}}(\text{null}(X_{\mathcal{E}})).$$

Consider defining, for a new point y' ,

$$\hat{\beta}_{-\mathcal{E}}(y') = 0 \quad \text{and} \quad \hat{\beta}_{\mathcal{E}}(y') = (X_{\mathcal{E}})^+(y' - (X_{\mathcal{E}}^T)^+ \lambda s) + b',$$

where $b' \in \text{null}(X_{\mathcal{E}})$, and is yet to be determined. Exactly as in the proof of Lemma 5, we know that $X_{\mathcal{E}}^T(y' - X\hat{\beta}(y')) = \lambda s$, and $\|X_{-\mathcal{E}}^T(y' - X\hat{\beta}(y'))\|_\infty < \lambda$ for all $y' \in U_1$, a neighborhood of y .

Now we want to choose b' so that $\hat{\beta}(y')$ has the correct active set and active signs. For simplicity of notation, first define the function $f: \mathbb{R}^n \rightarrow \mathbb{R}^{|\mathcal{E}|}$,

$$f(x) = (X_{\mathcal{E}})^+(x - (X_{\mathcal{E}}^T)^+ \lambda s).$$

Equation (41) implies that there is a $b' \in \text{null}(X_{\mathcal{E}})$ such that $b'_{-\mathcal{A}} = -f_{-\mathcal{A}}(y')$, hence $\hat{\beta}_{\mathcal{E} \setminus \mathcal{A}}(y') = 0$. However, we must choose b' so that additionally $\hat{\beta}_i(y') \neq 0$ for $i \in \mathcal{A}$ and $\text{sign}(\hat{\beta}_{\mathcal{A}}(y')) = r$. Write

$$\hat{\beta}_{\mathcal{E}}(y') = (f(y') + b) + (b' - b).$$

By the continuity of $f + b$, there exists a neighborhood of U_2 of y such that $f_i(y') + b_i \neq 0$ for $i \in \mathcal{A}$ and $\text{sign}(f_{\mathcal{A}}(y') + b_{\mathcal{A}}) = r$, for all $y' \in U_2$. Therefore we only need to choose a vector $b' \in \text{null}(X_{\mathcal{E}})$, with $b'_{-\mathcal{A}} = -f_{-\mathcal{A}}(y')$, such that $\|b' - b\|_2$ sufficiently small. This can be achieved by applying the bounded inverse theorem, which says that the bijective linear map $\pi_{-\mathcal{A}}$ has a bounded inverse (when considered a function from its row space to its column space). Therefore there exists some $M > 0$ such that for any y' , there is a vector $b' \in \text{null}(X_{\mathcal{E}})$, $b'_{-\mathcal{A}} = -f_{-\mathcal{A}}(y')$, with

$$\|b' - b\|_2 \leq M \|f_{-\mathcal{A}}(y') - f_{-\mathcal{A}}(y)\|_2.$$

Finally, the continuity of $f_{-\mathcal{A}}$ implies that $\|f_{-\mathcal{A}}(y') - f_{-\mathcal{A}}(y)\|_2$ can be made sufficiently small by restricting $y' \in U_3$, another neighborhood of y .

Letting $U = U_1 \cap U_2 \cap U_3$, we have shown that for any $y' \in U$, there exists a lasso solution $\hat{\beta}(y')$ with active set $\mathcal{A}(y') = \mathcal{A}$ and active signs $r(y') = r$.

A.4. Proof of Lemma 9. Define the set

$$\mathcal{N} = \bigcup_{\mathcal{B}, s} \bigcup_{\mathcal{A} \in Z(\mathcal{B})} \{z \in \mathbb{R}^n : P_{[D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}))]^\perp} \cdot D_{\mathcal{B} \setminus \mathcal{A}}(XP_{\text{null}(D_{-\mathcal{B}})})^+ \times (z - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ D_{\mathcal{B}}^T \lambda s) = 0\}.$$

The first union above is taken over all subsets $\mathcal{B} \subseteq \{1, \dots, m\}$ and all sign vectors $s \in \{-1, 1\}^{|\mathcal{B}|}$. The second union is taken over subsets $\mathcal{A} \subseteq Z(\mathcal{B})$, where

$$Z(\mathcal{B}) = \{\mathcal{A} \subseteq \mathcal{B} : P_{[D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}))]^\perp} D_{\mathcal{B} \setminus \mathcal{A}}(XP_{\text{null}(D_{-\mathcal{B}})})^+ \neq 0\}.$$

Since \mathcal{N} is a finite union of affine subspaces of dimension $\leq n - 1$, it has measure zero.

Now fix $y \notin \mathcal{N}$, and let $(\gamma(y), \hat{\beta}(y))$ be an optimal pair, with boundary set $\mathcal{B} = \mathcal{B}(y)$, boundary signs $s = s(y)$, active set $\mathcal{A} = \mathcal{A}(y)$, and active signs $r = r(y)$. Starting from (34), and plugging in for the fit in terms of \mathcal{B}, s , as in (35) we can show that

$$\begin{aligned} \gamma_{-\mathcal{B}}(y) &= \lambda^{-1} (D_{-\mathcal{B}}^T)^+ (X^T P_{\text{null}(P_{\text{null}(D_{-\mathcal{B}})}X^T)} y \\ &\quad + (X^T (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ - I) D_{\mathcal{B}}^T \lambda s) + c, \end{aligned}$$

where $c \in \text{null}(D_{-\mathcal{B}}^T)$. By (36), we know that

$$\hat{\beta}(y) = (XP_{\text{null}(D_{-\mathcal{B}})})^+ (y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ D_{\mathcal{B}}^T \lambda s) + b,$$

where $b \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$. Furthermore,

$$D_{\mathcal{B} \setminus \mathcal{A}} \hat{\beta}(y) = D_{\mathcal{B} \setminus \mathcal{A}} (XP_{\text{null}(D_{-\mathcal{B}})})^+ (y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ D_{\mathcal{B}}^T \lambda s) + D_{\mathcal{B} \setminus \mathcal{A}} b = 0,$$

or equivalently,

$$\begin{aligned} &D_{\mathcal{B} \setminus \mathcal{A}} (XP_{\text{null}(D_{-\mathcal{B}})})^+ (y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ D_{\mathcal{B}}^T \lambda s) \\ &= -D_{\mathcal{B} \setminus \mathcal{A}} b \in D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}})). \end{aligned}$$

Projecting onto the orthogonal complement of the linear subspace $D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}))$ therefore gives zero,

$$P_{[D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}))]^\perp} D_{\mathcal{B} \setminus \mathcal{A}} (XP_{\text{null}(D_{-\mathcal{B}})})^+ (y - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ D_{\mathcal{B}}^T \lambda s) = 0,$$

and because $y \notin \mathcal{N}$, we know that in fact

$$P_{[D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}}))]^\perp} D_{\mathcal{B} \setminus \mathcal{A}} (XP_{\text{null}(D_{-\mathcal{B}})})^+ = 0.$$

This can be rewritten as

$$(42) \quad \text{col}(D_{\mathcal{B} \setminus \mathcal{A}} (XP_{\text{null}(D_{-\mathcal{B}})})^+) \subseteq D_{\mathcal{B} \setminus \mathcal{A}}(\text{null}(X) \cap \text{null}(D_{-\mathcal{B}})).$$

At a new point y' , consider defining $\gamma_{\mathcal{B}}(y') = s$,

$$\begin{aligned} \gamma_{-\mathcal{B}}(y') &= \lambda^{-1}(D_{-\mathcal{B}}^T)^+(X^T P_{\text{null}(P_{\text{null}(D_{-\mathcal{B}})}X^T)}y' \\ &\quad + (X^T(P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ - I)D_{\mathcal{B}}^T\lambda s) + c, \end{aligned}$$

and

$$\hat{\beta}(y') = (XP_{\text{null}(D_{-\mathcal{B}})})^+(y' - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s) + b',$$

where $b' \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ is yet to be determined. By construction, $\gamma(y')$ and $\hat{\beta}(y')$ satisfy the stationarity condition (31) at y' . Hence it remains to show two parts: first, we must show that this pair satisfies the subgradient condition (32) at y' ; second, we must show this pair has boundary set $\mathcal{B}(y') = \mathcal{B}$, boundary signs $s(y') = s$, active set $\mathcal{A}(y') = \mathcal{A}$ and active signs $r(y') = r$. Actually, it suffices to show the second part alone, because the first part is then implied by the fact that $\gamma(y)$ and $\hat{\beta}(y)$ satisfy the subgradient condition at y . Well, by the continuity of the function $f: \mathbb{R}^n \rightarrow \mathbb{R}^{m-|\mathcal{B}|}$,

$$\begin{aligned} f(x) &= \lambda^{-1}(D_{-\mathcal{B}}^T)^+(X^T P_{\text{null}(P_{\text{null}(D_{-\mathcal{B}})}X^T)}x \\ &\quad + (X^T(P_{\text{null}(D_{-\mathcal{B}})}X^T)^+ - I)D_{\mathcal{B}}^T\lambda s) + c, \end{aligned}$$

we have $\|\gamma_{-\mathcal{B}}(y')\|_{\infty} < 1$ provided that $y' \in U_1$, a neighborhood of y . This ensures that $\gamma(y')$ has boundary set $\mathcal{B}(y') = \mathcal{B}$ and signs $s(y') = s$.

As for the active set and signs of $\hat{\beta}(y')$, note first that $D_{-\mathcal{B}}\hat{\beta}(y') = 0$, following directly from the definition. Next, define the function $g: \mathbb{R}^n \rightarrow \mathbb{R}^p$,

$$g(x) = (XP_{\text{null}(D_{-\mathcal{B}})})^+(x - (P_{\text{null}(D_{-\mathcal{B}})}X^T)^+D_{\mathcal{B}}^T\lambda s),$$

so $\hat{\beta}(y') = g(y') + b'$. Equation (42) implies that there is a vector $b' \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ such that $D_{\mathcal{B}\setminus\mathcal{A}}b' = -D_{\mathcal{B}\setminus\mathcal{A}}g(y')$, which makes $D_{\mathcal{B}\setminus\mathcal{A}}\hat{\beta}(y') = 0$. However, we still need to choose b' such that $D_i\hat{\beta}(y') \neq 0$ for all $i \in \mathcal{A}$ and $\text{sign}(D_{\mathcal{A}}\hat{\beta}(y')) = r$. To this end, write

$$\hat{\beta}(y') = (g(y') + b) + (b' - b).$$

The continuity of $D_{\mathcal{A}}g$ implies that there is a neighborhood U_2 of y such that $D_i g(y') + D_i b \neq 0$ for all $i \in \mathcal{A}$ and $\text{sign}(D_{\mathcal{A}}g(y') + D_{\mathcal{A}}b) = r$, for $y' \in U_2$. Since

$$\begin{aligned} |D_i\hat{\beta}(y)| &\geq |D_i g(y') + D_i b| - |D_i(b' - b)| \\ &\geq |D_i g(y') + D_i b| - \|D^T\|_2 \|b' - b\|_2, \end{aligned}$$

where $\|D^T\|_2$ is the operator norm of the D^T , we only need to choose $b' \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ such that $D_{\mathcal{B}\setminus\mathcal{A}}b' = -D_{\mathcal{B}\setminus\mathcal{A}}g(y')$, and such that $\|b' - b\|_2$ is sufficiently small. This is possible by the bounded inverse theorem applied to the linear map $D_{\mathcal{B}\setminus\mathcal{A}}$: when considered a function from its row space to its

column space, $D_{\mathcal{B}\setminus\mathcal{A}}$ is bijective and hence has a bounded inverse. Therefore there is some $M > 0$ such that for any y' , there is a $b' \in \text{null}(X) \cap \text{null}(D_{-\mathcal{B}})$ with $D_{\mathcal{B}\setminus\mathcal{A}}b' = -D_{\mathcal{B}\setminus\mathcal{A}}g(y')$ and

$$\|b' - b\|_2 \leq M \|D_{\mathcal{B}\setminus\mathcal{A}}g(y') - D_{\mathcal{B}\setminus\mathcal{A}}g(y)\|_2.$$

The continuity of $D_{\mathcal{B}\setminus\mathcal{A}}g$ implies that the right-hand side above can be made sufficiently small by restricting $y' \in U_3$, a neighborhood of y .

With $U = U_1 \cap U_2 \cap U_3$, we have shown for that for $y' \in U$, there is an optimal pair $(\gamma(y'), \hat{\beta}(y'))$ with boundary set $\mathcal{B}(y') = \mathcal{B}$, boundary signs $s(y') = s$, active set $\mathcal{A}(y') = \mathcal{A}$ and active signs $r(y') = r$.

A.5. Dual problems. The dual of the lasso problem (1) has appeared in many papers in the literature; as far as we can tell, it was first considered by Osborne, Presnell and Turlach (2000). We start by rewriting problem (1) as

$$\hat{\beta}, \hat{z} \in \underset{\beta \in \mathbb{R}^p, z \in \mathbb{R}^n}{\text{argmin}} \frac{1}{2} \|y - z\|_2^2 + \lambda \|\beta\|_1 \quad \text{subject to } z = X\beta;$$

then we write the Lagrangian

$$\mathcal{L}(\beta, z, v) = \frac{1}{2} \|y - z\|_2^2 + \lambda \|\beta\|_1 + v^T(z - X\beta),$$

and we minimize \mathcal{L} over β, z to obtain the dual problem

$$(43) \quad \hat{v} = \underset{v \in \mathbb{R}^n}{\text{argmin}} \|y - v\|_2^2 \quad \text{subject to } \|X^T v\|_\infty \leq \lambda.$$

Taking the gradient of \mathcal{L} with respect to β, z , and setting this equal to zero gives

$$(44) \quad \hat{v} = y - X\hat{\beta},$$

$$(45) \quad X^T \hat{v} = \lambda \gamma,$$

where $\gamma \in \mathbb{R}^p$ is a subgradient of the function $f(x) = \|x\|_1$ evaluated at $x = \hat{\beta}$. From (43), we can immediately see that the dual solution \hat{v} is the projection of y onto the polyhedron C as in Lemma 3, and then (44) shows that $X\hat{\beta} = y - \hat{v}$ is the residual from projecting y onto C . Further, from (45), we can define the equicorrelation set \mathcal{E} as

$$\mathcal{E} = \{i \in \{1, \dots, p\} : |X_i^T \hat{v}| = \lambda\}.$$

Noting that together (44), (45) are exactly the same as the KKT conditions (13), (14), and all of the arguments in Section 3 involving the equicorrelation set \mathcal{E} can be translated to this dual perspective.

There is a slightly different way to derive the lasso dual, resulting in a different (but of course, equivalent) formulation. We first rewrite problem (1) as

$$\hat{\beta}, \hat{z} \in \underset{\beta \in \mathbb{R}^p, z \in \mathbb{R}^n}{\text{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|z\|_1 \quad \text{subject to } z = \beta,$$

and by following similar steps to those above, we arrive at the dual problem

$$(46) \quad \hat{v} \in \underset{v \in \mathbb{R}^p}{\operatorname{argmin}} \|P_{\operatorname{col}(X)}y - (X^+)^T v\|_2^2 \quad \text{subject to } \|v\|_\infty \leq \lambda, v \in \operatorname{row}(X).$$

Each dual solution \hat{v} (now no longer unique) satisfies

$$(47) \quad (X^+)^T \hat{v} = P_{\operatorname{col}(X)}y - X\hat{\beta},$$

$$(48) \quad \hat{v} = \lambda\gamma.$$

The dual problem (46) and its relationship (47), (48) to the primal problem offer yet another viewpoint to understand some of the results in Section 3.

For the generalized lasso problem, one might imagine that there are three different dual problems, corresponding to the three different ways of introducing an auxiliary variable z into the generalized lasso criterion:

$$\hat{\beta}, \hat{z} \in \underset{\beta \in \mathbb{R}^p, z \in \mathbb{R}^n}{\operatorname{argmin}} \frac{1}{2} \|y - z\|_2^2 + \lambda \|D\beta\|_1 \quad \text{subject to } z = X\beta;$$

$$\hat{\beta}, \hat{z} \in \underset{\beta \in \mathbb{R}^p, z \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|Dz\|_1 \quad \text{subject to } z = \beta;$$

$$\hat{\beta}, \hat{z} \in \underset{\beta \in \mathbb{R}^p, z \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|z\|_1 \quad \text{subject to } z = D\beta.$$

However, the first two approaches above lead to Lagrangian functions that cannot be minimized analytically over β, z . Only the third approach yields a dual problem in closed-form, as given by Tibshirani and Taylor (2011),

$$(49) \quad \hat{v} \in \underset{v \in \mathbb{R}^m}{\operatorname{argmin}} \|P_{\operatorname{col}(X)}y - (X^+)^T D^T v\|_2^2$$

subject to $\|v\|_\infty \leq \lambda, D^T v \in \operatorname{row}(X)$.

The relationship between primal and dual solutions is

$$(50) \quad (X^+)^T D^T \hat{v} = P_{\operatorname{col}(X)}y - X\hat{\beta},$$

$$(51) \quad \hat{v} = \lambda\gamma,$$

where $\gamma \in \mathbb{R}^m$ is a subgradient of $f(x) = \|x\|_1$ evaluated at $x = D\hat{\beta}$. Directly from (49) we can see that $(X^+)^T D^T \hat{v}$ is the projection of the point $y' = P_{\operatorname{col}(X)}y$ onto the polyhedron

$$K = \{(X^+)^T D^T v : \|v\|_\infty \leq \lambda, D^T v \in \operatorname{row}(X)\}.$$

By (50), the primal fit is $X\hat{\beta} = (I - P_K)(y')$, which can be rewritten as $X\hat{\beta} = (I - P_C)(y')$ where C is the polyhedron from Lemma 8, and finally $X\hat{\beta} = (I - P_C)(y)$ because $I - P_C$ is zero on $\operatorname{null}(X^T)$. By (51), we can define

the boundary set \mathcal{B} corresponding to a particular dual solution \hat{v} as

$$\mathcal{B} = \{i \in \{1, \dots, m\} : |\hat{v}_i| = \lambda\}.$$

(This explains its name, as \mathcal{B} gives the coordinates of \hat{v} that are on the boundary of the box $\{v : \|v\|_\infty \leq \lambda\}$.) As (50), (51) are equivalent to the KKT conditions (31), (32) [following from rewriting (50) using $D^T \hat{v} \in \text{row}(X)$], the results in Section 4 on the boundary set \mathcal{B} can all be derived from this dual setting.

REFERENCES

- CHEN, S. S., DONOHO, D. L. and SAUNDERS, M. A. (1998). Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.* **20** 33–61. [MR1639094](#)
- DOSSAL, C., KACHOUR, M., FADILI, J., PEYRE, G. and CHESNEAU, C. (2011). The degrees of freedom of the lasso for general design matrix. Available at arXiv:[1111.1162](#).
- EFRON, B. (1986). How biased is the apparent error rate of a prediction rule? *J. Amer. Statist. Assoc.* **81** 461–470. [MR0845884](#)
- EFRON, B., HASTIE, T., JOHNSTONE, I. and TIBSHIRANI, R. (2004). Least angle regression (with discussion, and a rejoinder by the authors). *Ann. Statist.* **32** 407–499. [MR2060166](#)
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- GRÜNBAUM, B. (2003). *Convex Polytopes*, 2nd ed. *Graduate Texts in Mathematics* **221**. Springer, New York. [MR1976856](#)
- HASTIE, T. J. and TIBSHIRANI, R. J. (1990). *Generalized Additive Models. Monographs on Statistics and Applied Probability* **43**. Chapman & Hall, London. [MR1082147](#)
- LOUBES, J. M. and MASSART, P. (2004). Discussion to “Least angle regression.” *Ann. Statist.* **32** 460–465.
- MALLOWS, C. (1973). Some comments on C_p . *Technometrics* **15** 661–675.
- MEYER, M. and WOODROOFE, M. (2000). On the degrees of freedom in shape-restricted regression. *Ann. Statist.* **28** 1083–1104. [MR1810920](#)
- OSBORNE, M. R., PRESNELL, B. and TURLACH, B. A. (2000). On the LASSO and its dual. *J. Comput. Graph. Statist.* **9** 319–337. [MR1822089](#)
- ROSSET, S., ZHU, J. and HASTIE, T. (2004). Boosting as a regularized path to a maximum margin classifier. *J. Mach. Learn. Res.* **5** 941–973. [MR2248005](#)
- SCHNEIDER, R. (1993). *Convex Bodies: The Brunn–Minkowski Theory. Encyclopedia of Mathematics and Its Applications* **44**. Cambridge Univ. Press, Cambridge. [MR1216521](#)
- STEIN, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *Ann. Statist.* **9** 1135–1151. [MR0630098](#)
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- TIBSHIRANI, R. J. (2011). The solution path of the generalized lasso, Ph.D. thesis, Dept. Statistics, Stanford Univ.
- TIBSHIRANI, R. J. and TAYLOR, J. (2011). The solution path of the generalized lasso. *Ann. Statist.* **39** 1335–1371. [MR2850205](#)
- VAITER, S., PEYRE, G., DOSSAL, C. and FADILI, J. (2011). Robust sparse analysis regularization. Available at arXiv:[1109.6222](#).
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** 301–320. [MR2137327](#)

ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2007). On the “degrees of freedom” of the lasso. *Ann. Statist.* **35** 2173–2192. [MR2363967](#)

DEPARTMENT OF STATISTICS
CARNEGIE MELLON UNIVERSITY
BAKER HALL
PITTSBURGH, PENNSYLVANIA 15213
USA
E-MAIL: ryantibs@cmu.edu

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
SEQUOIA HALL
STANFORD, CALIFORNIA 94305
USA
E-MAIL: jtaylo@stanford.edu