

Degrees of Freedom and Model Search

Ryan J. Tibshirani

Abstract

Degrees of freedom is a fundamental concept in statistical modeling, as it provides a quantitative description of the amount of fitting performed by a given procedure. But, despite this fundamental role in statistics, its behavior is not completely well-understood, even in somewhat basic settings. For example, it may seem intuitively obvious that the best subset selection fit with subset size k has degrees of freedom larger than k , but this has not been formally verified, nor has it been precisely studied. At large, the current paper is motivated by this problem, and we derive an exact expression for the degrees of freedom of best subset selection in a restricted setting (orthogonal predictor variables). Along the way, we develop a concept that we name “search degrees of freedom”; intuitively, for adaptive regression procedures that perform variable selection, this is a part of the (total) degrees of freedom that we attribute entirely to the model selection mechanism. Finally, we establish a modest extension of Stein’s formula to cover discontinuous functions, and discuss its potential role in degrees of freedom and search degrees of freedom calculations.

Keywords: *degrees of freedom, model search, lasso, best subset selection, Stein’s formula*

1 Introduction

Suppose that we are given observations $y \in \mathbb{R}^n$ from the model

$$y = \mu + \epsilon, \quad \text{with } \mathbb{E}(\epsilon) = 0, \text{Cov}(\epsilon) = \sigma^2 I, \quad (1)$$

where $\mu \in \mathbb{R}^n$ is some fixed, true mean parameter of interest, and $\epsilon \in \mathbb{R}^n$ are uncorrelated errors, with zero mean and common marginal variance $\sigma^2 > 0$. For a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, thought of as a procedure for producing fitted values, $\hat{\mu} = f(y)$, recall that the *degrees of freedom* of f is defined as (Efron 1986, Hastie & Tibshirani 1990):

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

Intuitively, the quantity $\text{df}(f)$ reflects the effective number of parameters used by f in producing the fitted output $\hat{\mu}$. Consider linear regression, for example, where $f(y)$ is the least squares fit of y onto predictor variables $x_1, \dots, x_p \in \mathbb{R}^n$: for this procedure f , our intuition gives the right answer, as its degrees of freedom is simply p , the number of estimated regression coefficients.¹ This, e.g., leads to an unbiased estimate of the risk of the linear regression fit, via Mallows’s C_p criterion (Mallows 1973).

In general, characterizations of degrees of freedom are highly relevant for purposes like model comparisons and model selection; see, e.g., Efron (1986), Hastie & Tibshirani (1990), Tibshirani & Taylor (2012), and Section 1.2, for more motivation. Unfortunately, however, counting degrees of freedom can become quite complicated for nonlinear, adaptive procedures. (By nonlinear, we mean f being nonlinear as a function of y .) Even for many basic adaptive procedures, explicit answers are

¹This is assuming linear independence of x_1, \dots, x_p ; in general, it is the dimension of $\text{span}\{x_1, \dots, x_p\}$.

not known. A good example is best subset selection, in which, for a fixed integer k , we regress on the subset of x_1, \dots, x_p of size at most k giving the best linear fit of y (as measured by the residual sum of squares). Is the degrees of freedom here larger than k ? It seems that the answer should be “yes”, because even though there are k coefficients in the final linear model, the variables in this model were chosen adaptively (based on the data). And if the answer is indeed “yes”, then the natural follow-up question is: how much larger is it? That is, how many effective parameters does it “cost” to search through the space of candidate models? The goal of this paper is to investigate these questions, and related ones.

1.1 A motivating example

We begin by raising an interesting point: though it seems certain that a procedure like best subset selection would suffer an inflation of degrees of freedom, not all adaptive regression procedures do. In particular, the *lasso* (Tibshirani 1996, Chen et al. 1998), which also performs variable selection in the linear model setting, presents a very different story in terms of its degrees of freedom. Stacking the predictor variables x_1, \dots, x_p along the columns of a matrix $X \in \mathbb{R}^{n \times p}$, the lasso estimate can be expressed as:

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

where $\lambda \geq 0$ is a tuning parameter, controlling the level of sparsity. Though not strictly necessary for our discussion, we assume for simplicity that X has columns in general position, which ensures uniqueness of the lasso solution $\hat{\beta}^{\text{lasso}}$ (see, e.g., Tibshirani (2013)). We will write $\mathcal{A}^{\text{lasso}} \subseteq \{1, \dots, p\}$ to denote the indices of nonzero coefficients in $\hat{\beta}^{\text{lasso}}$, called the support or active set of $\hat{\beta}^{\text{lasso}}$, also expressed as $\mathcal{A}^{\text{lasso}} = \operatorname{supp}(\hat{\beta}^{\text{lasso}})$.

The lasso admits a simple formula for its degrees of freedom.

Theorem 1 (Zou et al. 2007, Tibshirani & Taylor 2012). *Provided that the variables (columns) in X are in general position, the lasso fit $\hat{\mu}^{\text{lasso}} = X\hat{\beta}^{\text{lasso}}$ has degrees of freedom*

$$\operatorname{df}(\hat{\mu}^{\text{lasso}}) = \mathbb{E}|\mathcal{A}^{\text{lasso}}|,$$

where $|\mathcal{A}^{\text{lasso}}|$ is the size of the lasso active set $\mathcal{A}^{\text{lasso}} = \operatorname{supp}(\hat{\beta}^{\text{lasso}})$. The above expectation assumes that X and λ are fixed, and is taken over the sampling distribution $y \sim N(\mu, \sigma^2 I)$.

In other words, the degrees of freedom of the lasso fit is the number of selected variables, in expectation. This is somewhat remarkable because, as with subset selection, the lasso uses the data to choose which variables to put in the model. So how can its degrees of freedom be equal to the (average) number of selected variables, and not more? The key realization is that the lasso shrinks the coefficients of these variables towards zero, instead of performing a full least squares fit. This shrinkage is due to the ℓ_1 penalty that appears in (3). Amazingly, the “surplus” from adaptively building the model is exactly accounted for by the “deficit” from shrinking the coefficients, so that altogether (in expectation), the degrees of freedom is simply the number of variables in the model.

Remark 1. An analogous result holds for an entirely arbitrary predictor matrix X (not necessarily having columns in general position), see Tibshirani & Taylor (2012); analogous results also exist for the generalized lasso problem (special cases of which are the fused lasso and trend filtering), see Tibshirani & Taylor (2011, 2012).

Figure 1 shows an empirical comparison between the degrees of freedom of the lasso and best subset selection fits, for a simple example with $n = 20$, $p = 10$. The predictor variables were setup to have a block correlation structure, in that variables 1 through 4 had high pairwise correlation (between 0.6 and 0.9), variables 5 through 10 also had high pairwise correlation (between 0.6 and 0.9), and the two blocks were uncorrelated with each other. The outcome y was drawn by adding

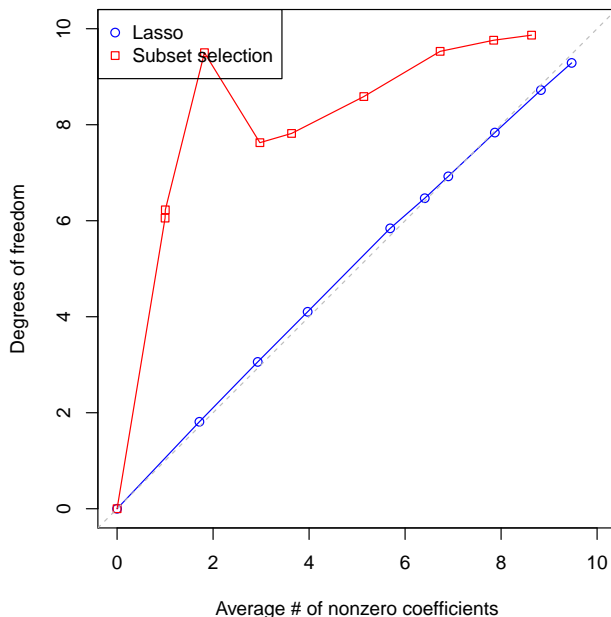


Figure 1: A simulated regression example with $n = 20$, $p = 10$. We drew 100 copies of the outcome y from the same sparse regression setup, and fit the lasso and best subset selection estimates each time, across 10 prespecified tuning parameter values. The plot shows the average number of selected variables by the lasso (in blue) and best subset selection (in red), across the tuning parameter values, versus their (estimated) degrees of freedom. The lasso degrees of freedom lines up with the number of selected variables, but the same is not true for subset selection, with its degrees of freedom being relatively much larger.

independent normal noise to $X\beta^*$, for some true coefficient vector β^* , supported on the first block of variables, and on one variable in the second block. We computed the lasso estimate in (3) over 10 values of the tuning parameter λ , as well as a best subset selection estimate

$$\hat{\beta}^{\text{subset}} \in \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_0, \quad (4)$$

over its own 10 values of λ . Recall that $\|\beta\|_0 = \sum_{j=1}^p 1\{\beta_j \neq 0\}$. We repeated this process 100 times, i.e., drew 100 copies of y from the described regression model, keeping X and β^* fixed, and each time computed fitted values from the lasso and best subset selection across the same 10 values of the tuning parameter. For each method and value of λ , we then:

1. computed the average number of nonzero coefficients over the 100 trials;
2. evaluated the covariance in (2) empirically across the 100 trials, as an (unbiased) estimate of the degrees of freedom.

Figure 1 plots the first quantity versus the second quantity, with the lasso in blue and best subset selection in red. As prescribed by Theorem 1, the (estimated) degrees of freedom of the lasso fit is closely aligned with the average number of nonzero coefficients in its estimate. But subset selection does not follow the same trend; its (estimated) degrees of freedom is much larger than its delivered number of nonzero coefficients. For example, when λ is tuned so that the subset selection estimate has a little less than 3 nonzero coefficients on average, the fit uses about 9 degrees of freedom.

Why does this happen? Again, this can be intuitively explained by shrinkage—this time, a lack thereof. If we denote the support of a best subset selection solution by $\mathcal{A}^{\text{subset}} = \operatorname{supp}(\hat{\beta}^{\text{subset}})$, and

abbreviate $\mathcal{A} = \mathcal{A}^{\text{subset}}$, then it is not hard to see that

$$\hat{\beta}_{\mathcal{A}}^{\text{subset}} = (X_{\mathcal{A}}^T X_{\mathcal{A}})^{-1} X_{\mathcal{A}}^T y,$$

i.e., the active coefficients are given by least squares on the active variables $X_{\mathcal{A}}$ (the submatrix of X formed by taking the columns in \mathcal{A}). Therefore, like the lasso, best subset selection chooses an active set of variables adaptively, but unlike the lasso, it fits their coefficients without shrinkage, using ordinary least squares. It pays for the “surplus” of covariance from the adaptive model search, as well as the usual amount from least squares estimation, resulting in a total degrees of freedom much larger than $|\mathcal{A}|$ (or rather, $\mathbb{E}|\mathcal{A}|$).

A clarifying note: simulations along the lines of that in Figure 1 can be found throughout the literature and we do not mean to claim originality here (e.g., see Figure 4 of Tibshirani & Knight (1999) for an early example, and Figure 2 of Janson et al. (2013) for a recent example). This simulation is instead simply meant to motivate the work that follows, as an aim of this paper is to examine the observed phenomenon in Figure 1 more formally.

1.2 Degrees of freedom and optimism

Degrees of freedom is closely connected to the concept of optimism, and so alternatively, we could have motivated the study of the covariance term on the right-hand side in (2) from the perspective of the optimism, rather than the complexity, of a fitting procedure. Assuming only that y is drawn from the model in (1), and that y' is an independent copy of y (i.e., an independent draw from (1)), it is straightforward to show that for any fitting procedure f ,

$$\mathbb{E}\|y' - f(y)\|_2^2 - \mathbb{E}\|y - f(y)\|_2^2 = 2\sigma^2 \cdot \text{df}(f). \quad (5)$$

The quantity on the left-hand side above is called the *optimism* of f , i.e., the difference in the mean squared test error and mean squared training error. The identity in (5) shows that (for uncorrelated, homoskedastic regression errors as in (1)) the optimism of f is just a positive constant times its degrees of freedom; in other words, fitting procedures with a higher degrees of freedom will have higher a optimism. Hence, from the example in the last section, we know when they are tuned to have the same (expected) number of variables in the fitted model, best subset selection will produce a training error that is generally far more optimistic than that produced by the lasso.

1.3 Lagrange versus constrained problem forms

Recall that we defined the subset selection estimator using the Lagrange form optimization problem (4), instead of the (perhaps more typical) constrained form definition

$$\hat{\beta}^{\text{subset}} \in \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \|y - X\beta\|_2^2 \quad \text{subject to} \quad \|\beta\|_0 \leq k. \quad (6)$$

There are several points now worth making. First, these are nonconvex optimization problems, and so the Lagrange and constrained forms (4) and (6) of subset selection are generally not equivalent. In fact, for all λ , solutions of (4) are solutions of (6) for some choice of k , but the reverse is not true. Second, even in situations in which the Lagrange and constrained forms of a particular optimization problem are equivalent (e.g., this is true under strong duality, and so it is true for most convex problems, under very weak conditions), there is a difference between studying the degrees of freedom of an estimator defined in one problem form versus the other. This is because the map from the Lagrange parameter in one form to the constraint bound in the other generically depends on y , i.e., it is a random mapping (Kaufman & Rosset (2013) discuss this for ridge regression and the lasso).

Lastly, in this paper, we focus on the Lagrange form (4) of subset selection because we find this problem is easier to analyze mathematically. For example, in Lagrange form with $X = I$, the i th

component of the subset selection fit $\hat{\beta}_i^{\text{subset}}$ depends on y_i only (and is given by hard thresholding), for each $i = 1, \dots, n$; in constrained form with $X = I$, each $\hat{\beta}_i^{\text{subset}}$ is a function of the order statistics of $|y_1|, \dots, |y_n|$, and hence depends on the whole sample.

Given the general spirit of our paper, it is important to recall the relevant work of Ye (1998), who studied degrees of freedom for special cases of best subset selection in constrained form. In one such special case (orthogonal predictors with null underlying signal), the author derived a simple expression for degrees of freedom as the sum of the k largest order statistics from a sample of n independent χ_1^2 random variables. This indeed establishes that, in this particular special case, the constrained form of best subset selection with k active variables has degrees of freedom larger than k . It does not, however, imply any results about the Lagrange case for the reasons explained above.

1.4 Assumptions, notation, and outline

Throughout this work, we will assume the model

$$y = \mu + \epsilon, \quad \epsilon \sim N(0, \sigma^2 I). \quad (7)$$

Note that this is stronger than the model in (1), since we are assuming a normal error distribution. While the model in (1) is sufficient to define the notion of degrees of freedom in general, we actually require normality for the calculations to come—specifically, Lemma 1 (on the degrees of freedom of hard thresholding), and all results in Section 5 (on extending Stein’s formula), rely on the normal error model. Beyond this running assumption, we will make any additional assumptions clear when needed.

In terms of notation, we write M^+ to denote the (Moore-Penrose) pseudoinverse of a matrix M , with $M^+ = (M^T M)^+ M^T$ for rectangular matrices M , and we write M_S to denote the submatrix of M whose columns correspond to the set of indices S . We write ϕ for the standard normal density function and Φ for the standard normal cumulative distribution function.

Finally, here is an outline for the rest of this article. In Section 2, we derive an explicit formula for the degrees of freedom of the best subset selection fit, under orthogonal predictors X . We also introduce the notion of search degrees of freedom for subset selection, and study its characteristics in various settings. In Section 3, we define search degrees of freedom for generic adaptive regression procedures, including the lasso and ridge regression as special cases. Section 4 returns to considering best subset selection, this time with general predictor variables X . Because exact formulae for the degrees of freedom and search degrees of freedom of best subset selection are not available in the general X case, we turn to simulation to investigate these quantities. We also examine the search degrees of freedom of the lasso across the same simulated examples (as its analytic calculation is again intractable for general X). Section 5 casts all of this work on degrees of freedom (and search degrees of freedom) in a different light, by deriving an extension of Stein’s formula. Stein’s formula is a powerful tool that can be used to compute the degrees of freedom of continuous and almost differentiable fitting procedures; our extension covers functions that have “well-behaved” points of discontinuity, in some sense. This extended version of Stein’s formula offers an alternative proof of the exact result in Section 2 (the orthogonal X case), and potentially, provides a perspective from which we can formally understand the empirical findings in Section 4 (the general X case). In Section 6, we conclude with some discussion.

2 Best subset selection with an orthogonal X

In the special case that $X \in \mathbb{R}^{n \times p}$ is orthogonal, i.e., X has orthonormal columns, we can compute the degrees of freedom of the best subset selection fit directly.

Theorem 2. Assume that $y \sim N(\mu, \sigma^2 I)$, and that X is orthogonal, meaning that $X^T X = I$. Then the best subset selection fit $\hat{\mu}^{\text{subset}} = X \hat{\beta}^{\text{subset}}$, at any fixed value of λ , has degrees of freedom

$$\text{df}(\hat{\mu}^{\text{subset}}) = \mathbb{E}|\mathcal{A}^{\text{subset}}| + \frac{\sqrt{2\lambda}}{\sigma} \sum_{i=1}^p \left[\phi\left(\frac{\sqrt{2\lambda} - (X^T \mu)_i}{\sigma}\right) + \phi\left(\frac{\sqrt{2\lambda} + (X^T \mu)_i}{\sigma}\right) \right], \quad (8)$$

where ϕ is the standard normal density.

The proof essentially reduces to a calculation on the degrees of freedom of the (componentwise) hard thresholding operator because, in the orthogonal X case, the best subset selection solution is exactly hard thresholding of $X^T y$. Formally, define the hard thresholding operator $H_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$, at a fixed level $t \geq 0$, by its coordinate functions

$$[H_t(y)]_i = y_i \cdot 1\{|y_i| \geq t\}, \quad i = 1, \dots, n.$$

Let \mathcal{A}_t denote the support set of the output, $\mathcal{A}_t = \text{supp}(H_t(y))$. The following result simply comes from the normality of y , and the definition of degrees of freedom in (2).

Lemma 1. Assume that $y \sim N(\mu, \sigma^2 I)$, and $t \geq 0$ is arbitrary but fixed. Then the hard thresholding operator H_t has degrees of freedom

$$\text{df}(H_t) = \mathbb{E}|\mathcal{A}_t| + \frac{t}{\sigma} \sum_{i=1}^n \left[\phi\left(\frac{t - \mu_i}{\sigma}\right) + \phi\left(\frac{t + \mu_i}{\sigma}\right) \right]. \quad (9)$$

Remark 2. This result, on the degrees of freedom of the hard thresholding operator, can be found in both Mazumder et al. (2011) and Deledalle et al. (2013). The former work uses degrees of freedom as a calibration tool in nonconvex sparse regression; the latter derives an estimate of the right hand side in (9) that, although biased, is consistent under some conditions.

The proofs of Lemma 1, and subsequently Theorem 2, involve straightforward calculations, but are deferred until the appendix for the sake of readability.

We have established that, for an orthogonal X , the degrees of freedom of the subset selection fit is equal to $\mathbb{E}|\mathcal{A}^{\text{subset}}|$, plus an “extra” term. We make some observations about this term in the next section.

2.1 Search degrees of freedom

The quantity

$$\text{sdf}(\hat{\mu}^{\text{subset}}) = \frac{\sqrt{2\lambda}}{\sigma} \sum_{i=1}^p \left[\phi\left(\frac{\sqrt{2\lambda} - (X^T \mu)_i}{\sigma}\right) + \phi\left(\frac{\sqrt{2\lambda} + (X^T \mu)_i}{\sigma}\right) \right] \quad (10)$$

appearing in (8) is the amount by which the degrees of freedom exceeds the expected number of selected variables. We will refer to this the *search degrees of freedom* of best subset selection, because roughly speaking, we can think of it as the extra amount of covariance that comes from searching through the space of models. Note that $\text{sdf}(\hat{\mu}^{\text{subset}}) > 0$ for any $\lambda > 0$, because the normal density is supported everywhere, and therefore we can indeed conclude that $\text{df}(\hat{\mu}^{\text{subset}}) > \mathbb{E}|\mathcal{A}^{\text{subset}}|$, as we suspected, in the case of an orthogonal predictor matrix.

How big is $\text{sdf}(\hat{\mu}^{\text{subset}})$? At the extremes: $\text{sdf}(\hat{\mu}^{\text{subset}}) = 0$ when $\lambda = 0$, and $\text{sdf}(\hat{\mu}^{\text{subset}}) \rightarrow 0$ as $\lambda \rightarrow \infty$. In words, searching has no cost when all of the variables, or none of the variables, are in the model. But the behavior is more interesting for intermediate values of λ . The precise shape of the search degrees of freedom curve (10), over λ , depends on the underlying signal μ ; the next three sections study three canonical cases for the underlying signal.




Figure 2: An example with $n = p = 100$, $X = I$, and $\mu = 0$. The left panel plots the curves $\text{df}(\hat{\mu}^{\text{subset}})$, $\text{sdf}(\hat{\mu}^{\text{subset}})$, and $\mathbb{E}|\mathcal{A}^{\text{subset}}|$ as functions of λ , drawn as blue, red, and black lines, respectively. The right panel plots the same quantities with respect to $\mathbb{E}|\mathcal{A}^{\text{subset}}|$.

2.2 Example: null signal

We consider first the case of a null underlying signal, i.e., $\mu = 0$. The best subset selection search degrees of freedom (10), as a function of λ , becomes

$$\text{sdf}(\hat{\mu}^{\text{subset}}) = \frac{2p\sqrt{2\lambda}}{\sigma} \phi\left(\frac{\sqrt{2\lambda}}{\sigma}\right). \quad (11)$$

In Figure 2, we plot the quantities $\text{df}(\hat{\mu}^{\text{subset}})$, $\text{sdf}(\hat{\mu}^{\text{subset}})$, and $\mathbb{E}|\mathcal{A}^{\text{subset}}|$ as functions of λ , for a simple example with $n = p = 100$, underlying signal $\mu = 0$, noise variance $\sigma^2 = 1$, and predictor matrix $X = I$, the 100×100 identity matrix. We emphasize that this figure was produced without any random draws or simulations, and the plotted curves are exactly as prescribed by Theorem 2 (recall that $\mathbb{E}|\mathcal{A}^{\text{subset}}|$ also has an explicit form in terms of λ , given in the proof of Lemma 1). In the left panel, we can see that the search degrees of freedom curve is maximized at approximately $\lambda = 0.5$, and achieves a maximum value of nearly 50. That is, when $\lambda = 0.5$, best subset selection spends nearly 50 (extra) parameters searching through the space of models!

It is perhaps more natural to parametrize the curves in terms of the expected number of active variables $\mathbb{E}|\mathcal{A}^{\text{subset}}|$ (instead of λ), as displayed in the right panel of Figure 2. This parametrization reveals something interesting: the search degrees of freedom curve is maximized at roughly $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 31.7$. In other words, searching is most costly when there are approximately 31.7 variables in the model. This is a bit counterintuitive, because there are more subsets of size 50 than any other size, that is, the function

$$F(k) = \binom{p}{k}, \quad k = 1, 2, \dots, p,$$

is maximized at $k = p/2 = 50$. Hence we might believe that searching through subsets of variables is most costly when $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 50$, because in this case the search space is largest. Instead, the maximum actually occurs at about $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 31.7$. Given the simple form (11) of the search degrees of freedom curve in the null signal case, we can verify this observation analytically: direct calculation shows that the right hand side in (11) is maximized at $\lambda = \sigma^2/2$, which, when plugged

into the formula for the expected number of selected variables in the null case,

$$\mathbb{E}|\mathcal{A}^{\text{subset}}| = 2p\Phi\left(\frac{-\sqrt{2\lambda}}{\sigma}\right),$$

yields $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 2\Phi(-1)p \approx 0.317p$.

Although this calculation may have been reassuring, the intuitive question remains: why is the 31.7 variable model associated with the highest cost of model searching (over, say, the 50 variable model)? At this point, we cannot offer a truly satisfying intuitive answer, but we will attempt an explanation nonetheless. Recall that search degrees of freedom measures the additional amount of covariance in (2) that we attribute to searching through the space of models—additional from the baseline amount $\mathbb{E}|\mathcal{A}^{\text{subset}}|$, which comes from estimating the coefficients in the selected model. The shape of the search degrees of freedom curve, when $\mu = 0$, tells us that there is more covariance to be gained when the selected model has 31.7 variables than when it has 50 variables. As the size of the selected subset k increases from 0 to 50, note that:

1. the number of subsets of size k increases, which means that there are more opportunities to decrease the training error, and so the total degrees of freedom (optimism) increases;
2. trivially, the baseline amount of fitting also increases, as this baseline is just k , the degrees of freedom (optimism) of a fixed model on k variables.

Search degrees of freedom is the difference between these two quantities (i.e., total minus baseline degrees of freedom), and as it turns out, the two are optimally balanced at approximately $k = 31.7$ (at exactly $k = 2\Phi(-1)p$) in the null signal case.

2.3 Example: sparse signal

Now we consider the case in which $\mu = X\beta^*$, for some sparse coefficient vector $\beta^* \in \mathbb{R}^p$. We let $\mathcal{A}^* = \text{supp}(\beta^*)$ denote the true support set, and $k^* = |\mathcal{A}^*|$ the true number of nonzero coefficients, assumed to be small. The search degrees of freedom curve in (10) is

$$\text{sdf}(\hat{\mu}^{\text{subset}}) = \frac{\sqrt{2\lambda}}{\sigma} \sum_{i \in \mathcal{A}^*} \left[\phi\left(\frac{\sqrt{2\lambda} - \beta_i^*}{\sigma}\right) + \phi\left(\frac{\sqrt{2\lambda} + \beta_i^*}{\sigma}\right) \right] + \frac{2(p - k^*)\sqrt{2\lambda}}{\sigma} \phi\left(\frac{\sqrt{2\lambda}}{\sigma}\right). \quad (12)$$

When the nonzero coefficients β_i^* are moderate (not very large), the curve in (12) acts much like the search degrees of freedom curve (11) in the null case. Otherwise, it can behave very differently. We therefore examine two different sparse setups by example, having low and high signal-to-noise ratios. See Figure 3. In both setups, we take $n = p = 100$, $\sigma^2 = 1$, $X = I$, and $\mu = X\beta^*$, with

$$\beta_i^* = \begin{cases} \rho & i = 1, \dots, 10 \\ 0 & i = 11, \dots, 100. \end{cases} \quad (13)$$

The left panel uses $\rho = 1$, and the right uses $\rho = 8$. We plot the total degrees of freedom and search degrees of freedom of subset selection as a function of the expected number of selected variables (and note, as before, that these plots are produced by mathematical formulae, not by simulation). The curves in the left panel, i.e., in the low signal-to-noise ratio case, appear extremely similar to those in the null signal case (right panel of Figure 2). The search degrees of freedom curve peaks when the expected number of selected variables is about $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 31.9$, and its peak height is again just short of 50.

Meanwhile, in the high signal-to-noise ratio case, i.e., the right panel of Figure 3, the behavior is very different. The search degrees of freedom curve is bimodal, and is basically zero when the expected number of selected variables is 10. The intuition: with such a high signal-to-noise ratio in




Figure 3: An example with $n = p = 100$, $X = I$, and $\mu = X\beta^*$ with β^* as in (13). The left panel corresponds to the choice $\rho = 1$ (low signal-to-noise regime) and the right to $\rho = 8$ (high signal-to-noise regime).

the true model (13), best subset selection is able to select the same (true) subset of 10 variables in every random data instance, and therefore the size 10 model produced by subset selection is akin to a fixed model, with no real searching performed whatsoever. Another interesting point is that the cost of model searching is very high when the selected model has average size equal to 5; here the search component contributes over 30 degrees of freedom to the total. Intuitively, with 10 strong variables in the true model (13), there are many competitive subsets of size 5, and hence a lot of searching is needed in order to report the best subset of size 5 (in terms of training error).

It is worth mentioning the interesting, recent works of Kaufman & Rosset (2013) and Janson et al. (2013), which investigate unexpected nonmonotonicities in the (total) degrees of freedom of an estimator, as a function of some underlying parametrization for the amount of imposed regularization. We note that the right panel of Figure 3 portrays a definitive example of this, in that the best subset selection degrees of freedom undergoes a major nonmonotonicity at 10 (expected) active variables, as discussed above.

2.4 Example: dense signal

The last case we consider is that of a dense underlying signal, $\mu = X\beta^*$ for some dense coefficient vector $\beta^* \in \mathbb{R}^p$. For the sake of completeness, in the present case, the expression (10) for the search degrees of freedom of best subset selection is

$$\text{sdf}(\hat{\mu}^{\text{subset}}) = \frac{\sqrt{2\lambda}}{\sigma} \sum_{i=1}^p \left[\phi\left(\frac{\sqrt{2\lambda} - \beta_i^*}{\sigma}\right) + \phi\left(\frac{\sqrt{2\lambda} + \beta_i^*}{\sigma}\right) \right]. \quad (14)$$

The search curve degrees of freedom curve (14) exhibits a very similar behavior to the curve (11) in the null signal case when the coefficients β_i^* are small or moderate, but a very different behavior when some coefficients β_i^* are large. In Figure 4, we take $n = p = 100$, $X = I$, and $\mu = X\beta^*$ with $\beta_i^* = \rho$, $i = 1, \dots, p$. The left panel of the figure corresponds to $\rho = 1$, and the right corresponds to $\rho = 8$. Both panels plot degrees of freedom against the expected number of selected variables (and, as in the last two subsections, these degrees of freedom curves are plotted according to their

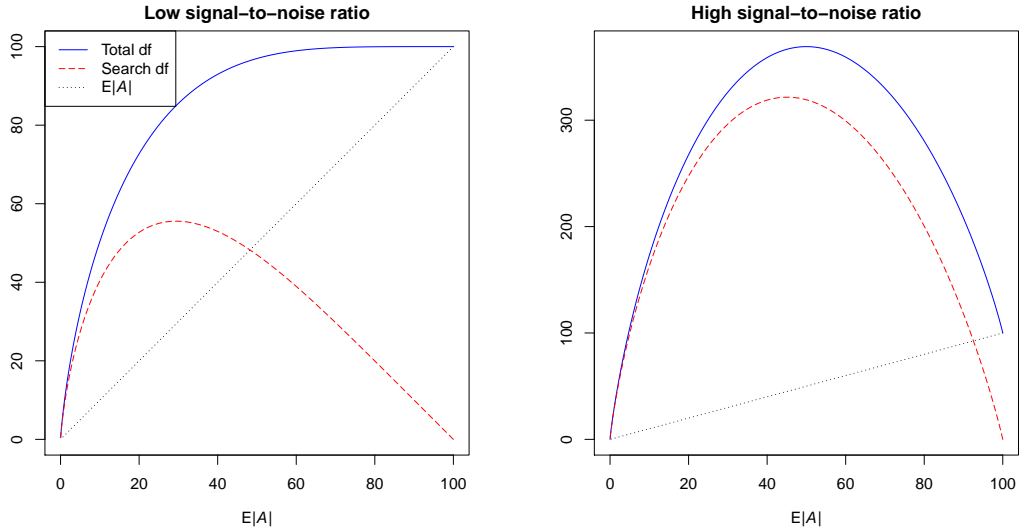


Figure 4: An example with $n = p = 100$, $X = I$, and $\mu = X\beta^*$ with $\beta_i^* = \rho$, $i = 1, \dots, p$. The left panel corresponds to $\rho = 1$ (low signal-to-noise regime) and the right to $\rho = 8$ (high signal-to-noise regime).

closed-form expressions, they are not derived from simulation). We can see that the low signal-to-noise ratio case, in the left panel, yields a set of curves quite similar to those from the null signal case, in the right panel of Figure 2. One difference is that the search degrees of freedom curve has a higher maximum (its value about 56, versus 48 in the null signal case), and the location of this maximum is further to the left (occurring at about $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 29.4$, versus $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 31.7$ in the former case).

On the other hand, the right panel of the figure shows the high signal-to-noise ratio case, where the total degrees of freedom curve is now nonmonotone, and reaches its maximum at an expected number of selected variables (very nearly) $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 50$. The search degrees of freedom curve itself peaks much later than it does in the other cases, at approximately $\mathbb{E}|\mathcal{A}^{\text{subset}}| = 45.2$. Another striking difference is the sheer magnitude of the degrees of freedom curves: at 50 selected variables on average, the total degrees of freedom of the best subset selection fit is well over 300. Mathematically, this makes sense, as the search degrees of freedom curve in (14) is increasing in $|\beta_i^*|$. Furthermore, we can liken the degrees of freedom curves in the right panel of Figure 4 to those in a small portion of the plot in the right panel of Figure 3, namely, the portion corresponding to $\mathbb{E}|\mathcal{A}^{\text{subset}}| \leq 10$. The two sets of curves here appear similar in shape. This is intuitively explained by the fact that, in the high signal-to-noise ratio regime, subset selection over a dense true model is similar to subset selection over a sparse true model, provided that we constrain our attention in the latter case to subsets of size less than or equal to the true model size (since under this constraint, the truly irrelevant variables in the sparse model do not play much of a role).

3 Search degrees of freedom for general procedures

Here we extend the notion of search degrees of freedom to general adaptive regression procedures. Given an outcome $y \in \mathbb{R}^n$ and predictors $X \in \mathbb{R}^{n \times p}$, we consider a fitting procedure $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ of the form

$$f(y) = X\hat{\beta}^{(f)},$$

for some estimated coefficients $\hat{\beta}^{(f)} \in \mathbb{R}^p$. Clearly, the lasso and best subset selection are two examples of such a fitting procedure, with the coefficients as in (3) and (4), respectively. We denote $\mathcal{A}^{(f)} = \text{supp}(\hat{\beta}^{(f)})$, the support set of the estimated coefficients under f . The overall complexity of f is measured by its degrees of freedom, as defined in (2) (just as it is for all fitting procedures), but we may be also interested in a degree of complexity associated solely with its model selection component—i.e., we might ask: how many effective parameters does f spend in simply selecting the active set $\mathcal{A}^{(f)}$?

We propose to address this question by developing a notion of search degrees of freedom for f , in a way that generalizes the notion considered in the last section specifically for subset selection. Abbreviating $\mathcal{A} = \mathcal{A}^{(f)}$, we first define a modified procedure \tilde{f} that returns the least squares fit on the active set \mathcal{A} ,

$$\tilde{f}(y) = P_{\mathcal{A}}y.$$

where $P_{\mathcal{A}} = X_{\mathcal{A}}(X_{\mathcal{A}}^T X_{\mathcal{A}})^+ X_{\mathcal{A}}^T$ is the projection onto the span of active predictors $X_{\mathcal{A}}$ (note the use of the pseudoinverse, as $X_{\mathcal{A}}$ need not have full column rank, depending on the nature of the procedure f). We now define the search degrees of freedom of f as

$$\begin{aligned} \text{sdf}(f) &= \text{df}(\tilde{f}) - \mathbb{E}[\text{rank}(X_{\mathcal{A}})] \\ &= \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}((P_{\mathcal{A}}y)_i, y_i) - \mathbb{E}[\text{rank}(X_{\mathcal{A}})]. \end{aligned} \quad (15)$$

The intuition behind this definition: by construction, \tilde{f} and f are identical in their selection of the active set \mathcal{A} , and only differ in how they estimate the nonzero coefficients once \mathcal{A} has been chosen, with \tilde{f} using least squares, and f using a possibly different mechanism. If \mathcal{A} were fixed, then a least squares fit on $X_{\mathcal{A}}$ would use $\mathbb{E}[\text{rank}(X_{\mathcal{A}})]$ degrees of freedom, and so it seems reasonable to assign the leftover part, $\text{df}(\tilde{f}) - \mathbb{E}[\text{rank}(X_{\mathcal{A}})]$, as the degrees of freedom spent by \tilde{f} in selecting \mathcal{A} in the first place, i.e., the amount spent by f in selecting \mathcal{A} in the first place.

It may help to discuss some specific cases.

3.1 Best subset selection

When f is the best subset selection fit, we have $\tilde{f} = f$, i.e., subset selection already performs least squares on the set of selected variables \mathcal{A} . Therefore,

$$\text{sdf}(f) = \text{df}(f) - \mathbb{E}|\mathcal{A}|, \quad (16)$$

where we have also used the fact that $X_{\mathcal{A}}$ must have linearly independent columns with best subset selection (otherwise, we could strictly decrease the ℓ_0 penalty in (4) while keeping the squared error loss unchanged). This matches our definition (10) of search degrees of freedom for subset selection in the orthogonal X case—it is the total degrees of freedom minus the expected number of selected variables, with the total being explicitly computable for orthogonal predictors, as we showed in the last section.

The same expression (16) holds for any fitting procedure f that uses least squares to estimate the coefficients in its selected model, because then $\tilde{f} = f$. (Note that, in full generality, $\mathbb{E}|\mathcal{A}|$ should be replaced again by $\mathbb{E}[\text{rank}(X_{\mathcal{A}})]$ in case $X_{\mathcal{A}}$ need not have full column rank.) An example of another such procedure is forward stepwise regression.

3.2 Ridge regression

For ridge regression, the active model is $\mathcal{A} = \{1, \dots, p\}$ for any draw of the outcome y , which means that the modified procedure \tilde{f} is just the full regression fit on X , and

$$\text{sdf}(f) = \mathbb{E}[\text{rank}(X)] - \mathbb{E}[\text{rank}(X)] = 0.$$

This is intuitively the correct notion of search degrees of freedom for ridge regression, since this procedure does not perform any kind of variable selection whatsoever. The same logic carries over to any procedure f whose active set \mathcal{A} is almost surely constant.

3.3 The lasso

The lasso case is an interesting one. We know from the literature (Theorem 1) that the degrees of freedom of the lasso fit is $\mathbb{E}|\mathcal{A}|$ (when the predictors are in general position), but how much of this total can we attribute to model searching? The modified procedure f that performs least squares on the lasso active set \mathcal{A} has been called the *relaxed lasso* (see Meinshausen (2007), who uses the same term to refer to a broader family of debiased lasso estimates). We will denote the relaxed lasso fitted values by $\hat{\mu}^{\text{relax}} = P_{\mathcal{A}^{\text{lasso}}}y$. When X has orthonormal columns, it is not hard to see that the relaxed lasso fit is given by hard thresholding, just like best subset selection, but this time with threshold level $t = \lambda$. The following result hence holds by the same arguments as those in Section 2 for subset selection.

Theorem 3. *If $y \sim N(\mu, \sigma^2)$, and $X^T X = I$, then the relaxed lasso fit $\hat{\mu}^{\text{relax}} = P_{\mathcal{A}^{\text{lasso}}}y$, at a fixed value $\lambda \geq 0$, has degrees of freedom*

$$\text{df}(\hat{\mu}^{\text{relax}}) = \mathbb{E}|\mathcal{A}^{\text{lasso}}| + \frac{\lambda}{\sigma} \sum_{i=1}^p \left[\phi\left(\frac{\lambda - (X^T \mu)_i}{\sigma}\right) + \phi\left(\frac{\lambda + (X^T \mu)_i}{\sigma}\right) \right].$$

Therefore the lasso has search degrees of freedom

$$\text{sdf}(\hat{\mu}^{\text{lasso}}) = \frac{\lambda}{\sigma} \sum_{i=1}^p \left[\phi\left(\frac{\lambda - (X^T \mu)_i}{\sigma}\right) + \phi\left(\frac{\lambda + (X^T \mu)_i}{\sigma}\right) \right]. \quad (17)$$

The search degrees of freedom formulae (17) and (10) are different as functions of λ , the tuning parameter, but this is not a meaningful difference; when each is parametrized by their respective expected number of selected variables $\mathbb{E}|\mathcal{A}|$, the two curves are exactly the same, and therefore all examples and figures in Section 2 demonstrating the behavior of the search degrees of freedom of best subset selection also apply to the lasso. In a sense, this is not a surprise, because for orthogonal predictors both the lasso and subset selection fits reduce to a sequence of marginal considerations (thresholds, in fact), and so their search mechanisms can be equated.

But for correlated predictors, we might believe that the search components associated with the lasso and best subset selection procedures are actually quite different. Even though our definition of search degrees of freedom in (15) is not connected to computation in any way, the fact that best subset selection (4) is NP-hard for a general X may seem to suggest (at a very loose level) that it somehow “searches more” than the convex lasso problem (3). For many problem setups, this guess (whether or not properly grounded in intuition) appears to be true in simulations, as we show next.

4 Best subset selection with a general X

We look back at the motivating example given in Section 1.1, where we estimated the degrees of freedom of best subset selection and the lasso by simulation, in a problem with $n = 20$ and $p = 10$. See Section 1.1 for more details about the setup (i.e., correlation structure of the predictor variables X , true coefficients β^* , etc.). Here we also consider the degrees of freedom of the relaxed lasso, estimated from the same set of simulations. Figure 5 plots these degrees of freedom estimates, in green, on top of the existing best subset selection and lasso curves from Figure 1. Interestingly, the relaxed lasso is seen to have a smaller degrees of freedom than best subset selection (when each is parametrized by their own average number of selected variables). Note that this means the search

degrees of freedom of the lasso (i.e., the difference between the green curve and the diagonal) is smaller than the search degrees of freedom of subset selection (the difference between the red curve and the diagonal).




Figure 5: The same simulation setup as in Figure 1, but now including the relaxed lasso degrees of freedom on the left panel, in green. (The relaxed lasso is the fitting procedure that performs least squares on the lasso active set.) We can see that the relaxed lasso has a smaller degrees of freedom than subset selection, as a function of their (respective) average number of selected variables. Hence, the lasso exhibits a smaller search degrees of freedom than subset selection, in this example.

This discrepancy between the search degrees of freedom of the lasso and subset selection, for correlated variables X , stands in contrast to the orthogonal case, where the two quantities were proven to be equal (subject to the appropriate parametrization). Further simulations with correlated predictors show that, for the most part, this discrepancy persists across a variety of cases; consult Figure 6 and the accompanying caption text for details. However, it is important to note that this phenomenon is not universal, and in some instances (particularly, when the computed active set is small, and the true signal is dense) the search degrees of freedom of the lasso can grow quite large and compete with that of subset selection. Hence, we can see that the two quantities do not always obey a simple ordering, and the simulations presented here call for a more formal understanding of their relationship.

Unfortunately, this is not an easy task, since direct calculation of the relevant quantities—the degrees of freedom of best subset selection and the relaxed lasso—is not tractable for a general X . In cases such as these, one usually turns to Stein’s formula as an alternative for calculating degrees of freedom; e.g., the result in Theorem 1 is derived using Stein’s formula. But Stein’s formula only applies to continuous (and almost differentiable) fitting procedures $f = f(y)$, and neither the best subset selection nor the relaxed lasso fit is continuous in y . The next section, therefore, is focused on extending Stein’s result to discontinuous functions.

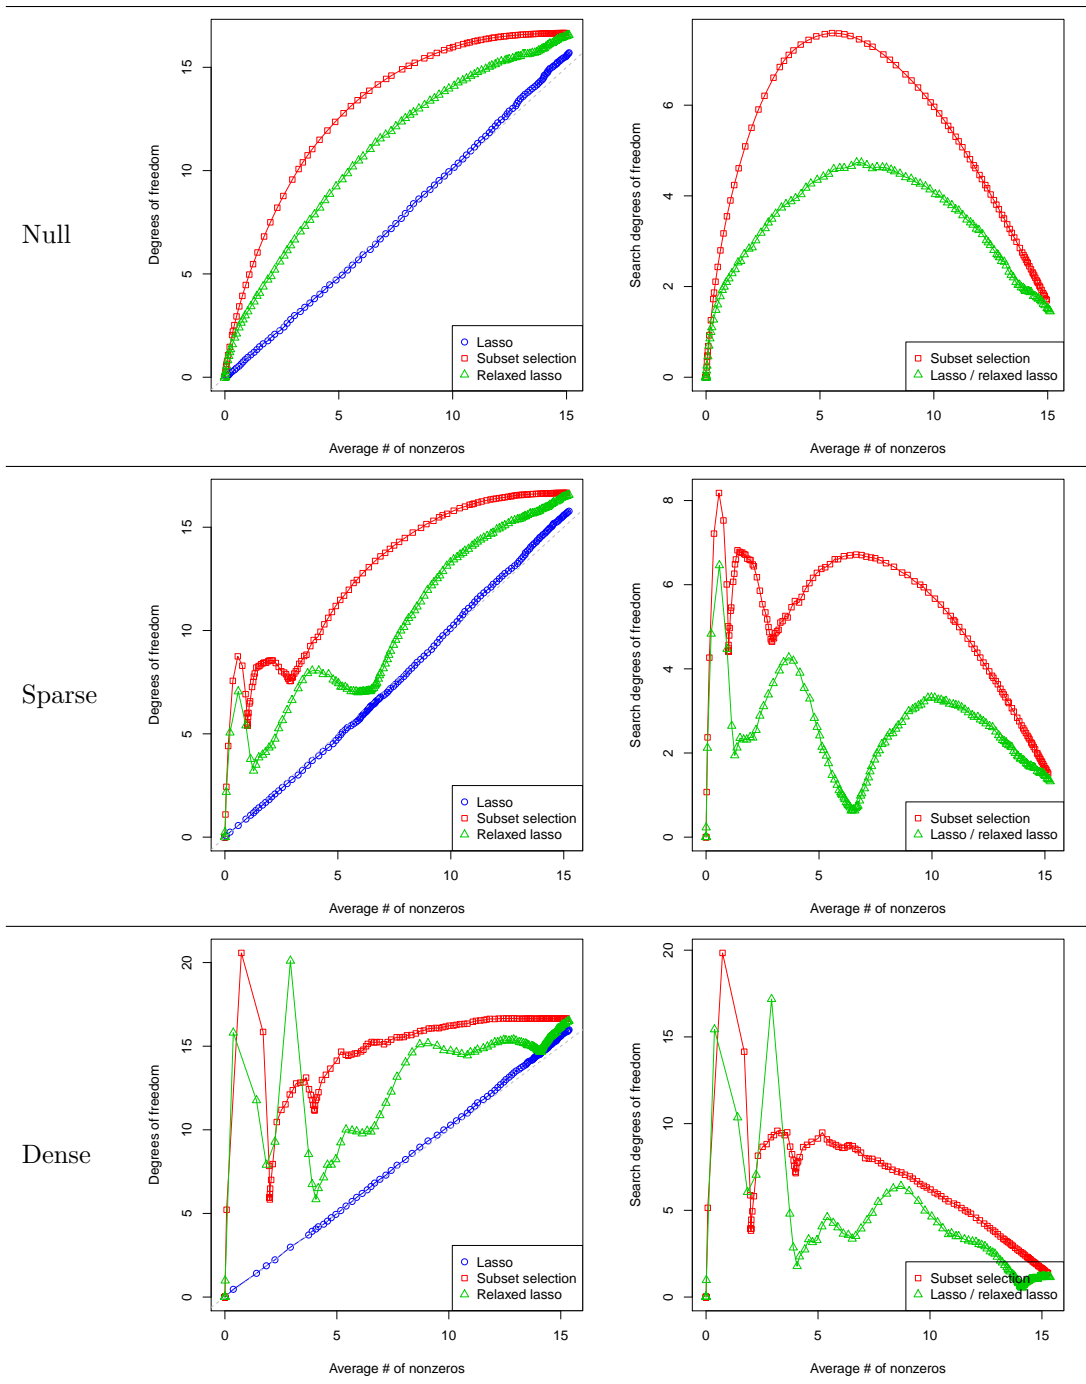


Figure 6: A set of simulation results with $n = 30$, $p = 16$ (we are confined to such a small setup because of the exponential computational complexity of subset selection). The rows of X were drawn i.i.d. from $N(0, \Sigma)$, where Σ is block diagonal with two equal sized (8×8) blocks Σ_1, Σ_2 . All diagonal entries of Σ_1, Σ_2 were set to 1, and the off-diagonal entries were drawn uniformly between 0.4 and 0.9. We considered three cases for the true mean $\mu = X\beta^*$: null ($\beta^* = 0$), sparse (β^* is supported on 3 variables in the first block and 1 in the second, with all nonzero components equal to 1), and dense (β^* has all components equal to 1). In all cases, we drew y around μ with independent standard normal noise, for a total of 100 repetitions. Overall, the search degrees of freedom of subset selection appears to be larger than that of the lasso, but at times the latter can rival the former in magnitude, especially for small active sets, and in the dense signal case.

5 An extension of Stein's formula

This section considers Stein's formula (Stein 1981), and presents an extension that yields an alternative derivation of the degrees of freedom results in Section 2, as well as (potential) insights into the empirical results in Section 4. In his remarkable paper, Stein studies the problem of estimating the mean of a multivariate normal distribution, with a spherical covariance matrix, under the usual squared error loss. The main result is an unbiased estimate of the associated risk for a large class of estimates of the mean. At the root of Stein's arguments lies the following lemma.

Lemma 2 (Stein 1981). *Let $Z \sim N(0, 1)$. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be absolutely continuous, with derivative f' . Assume that $\mathbb{E}|f'(Z)| < \infty$. Then*

$$\mathbb{E}[Zf(Z)] = \mathbb{E}[f'(Z)].$$

In its own right, this lemma (along with a converse statement, which is also true) has a number of important applications that span various areas of probability and statistics. For our purposes, the most relevant application is an alternative and highly useful formula for computing degrees of freedom. This is given by extending the above lemma to a setting in which the underlying normal distribution has an arbitrary mean vector and variance, and is also multivariate.

Lemma 3 (Stein 1981). *Let $X \sim N(\mu, \sigma^2 I)$, for some fixed $\mu \in \mathbb{R}^n$ and $\sigma^2 > 0$. Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be continuous and almost differentiable, and write $\nabla g = (\partial g_1/\partial x_1, \dots, \partial g_n/\partial x_n)$ for the vector of partial derivatives. Assume that $\mathbb{E}\|\nabla g(X)\|_2 < \infty$. Then*

$$\frac{1}{\sigma^2} \mathbb{E}[(X - \mu)g(X)] = \mathbb{E}[\nabla g(X)]. \quad (18)$$

We will delay the definition of almost differentiability until a little while later, but the eager reader can look ahead to Definition 2. Putting aside any concerns about regularity conditions, the result in (18) looks like a statement about degrees of freedom. To complete the connection, consider a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ giving the fit $\hat{\mu} = f(y)$, and assume the usual normal model $y \sim N(\mu, \sigma^2 I)$. Let $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ be the i th coordinate function of f . If f_i satisfies the appropriate conditions (continuity and almost differentiability), then we can apply Lemma 3 with $X = y$ and $g = f_i$, take the i th equality in (18), and sum over i to give

$$df(f) = \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}(f_i(y), y_i) = \mathbb{E} \left[\sum_{i=1}^n \frac{\partial f_i}{\partial y_i}(y) \right], \quad (19)$$

where $\partial f_i/\partial y_i$ denotes the partial derivative of f_i with respect to its i th variable. This is known as *Stein's formula for degrees of freedom*. It can be very useful, because in some cases the divergence $\sum_{i=1}^n \partial f_i/\partial y_i$ on the right-hand side of (19) can be computed explicitly, which yields an unbiased estimate of degrees of freedom. This is true, e.g., of the lasso fit, and as a concrete illustration, we prove the result in Theorem 1, using Stein's formula, in the appendix.

Useful as it can be, Stein's formula (19) is not universally applicable. There are several ways to break its assumptions; our particular interest is in fitting procedures that are discontinuous in y . For example, we showed in the proof of Theorem 2 that, when X is orthogonal, the subset selection solution is given by hard thresholding $X^T y$ at the level $t = \sqrt{2\lambda}$. The hard thresholding function H_t is clearly discontinuous: each one of its coordinate functions is discontinuous at t and $-t$. We therefore derive a modest extension of Stein's formula that allows us to deal with a certain class of (well-behaved) discontinuous functions. We begin with the univariate case, and then move on to the multivariate case.

5.1 An extension of Stein's univariate lemma

We consider functions $f : \mathbb{R} \rightarrow \mathbb{R}$ that are absolutely continuous on a partition of \mathbb{R} . Formally:

Definition 1. We say that a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is *piecewise absolutely continuous*, or *p-absolutely continuous*, if there exist points $\delta_1 < \delta_2 < \dots < \delta_m$ such that f is absolutely continuous on each one of the open intervals $(-\infty, \delta_1), (\delta_1, \delta_2), \dots, (\delta_m, \infty)$.

For a p-absolutely continuous function f , we write $\mathcal{D}(f) = \{\delta_1, \dots, \delta_m\}$ for its discontinuity set. Furthermore, note that such a function f has a derivative f' almost everywhere (because it has a derivative almost everywhere on each of the intervals $(-\infty, \delta_1), (\delta_1, \delta_2), \dots, (\delta_m, \infty)$). We will simply refer to f' as its derivative. Finally, we use the following helpful notation for one-sided limits,

$$f(x)_+ = \lim_{t \downarrow x} f(t) \quad \text{and} \quad f(x)_- = \lim_{t \uparrow x} f(t).$$

We now have the following extension of Stein's univariate lemma, Lemma 2.

Lemma 4. Let $Z \sim N(0, 1)$. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be p-absolutely continuous, and have a discontinuity set $\mathcal{D}(f) = \{\delta_1, \dots, \delta_m\}$. Let f' be its derivative, and assume that $\mathbb{E}|f'(Z)| < \infty$. Then

$$\mathbb{E}[Zf(Z)] = \mathbb{E}[f'(Z)] + \sum_{k=1}^m \phi(\delta_k) [f(\delta_k)_+ - f(\delta_k)_-].$$

The proof is similar to Stein's proof of Lemma 2, and is left to the appendix, for readability. It is straightforward to extend this result to a nonstandard normal distribution.

Corollary 1. Let $X \sim N(\mu, \sigma^2)$. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be p-absolutely continuous, with discontinuity set $\mathcal{D}(h) = \{\delta_1, \dots, \delta_m\}$, and derivative h' satisfying $\mathbb{E}|h'(X)| < \infty$. Then

$$\frac{1}{\sigma^2} \mathbb{E}[(X - \mu)h(X)] = \mathbb{E}[h'(X)] + \frac{1}{\sigma} \sum_{k=1}^m \phi\left(\frac{\delta_k - \mu}{\sigma}\right) [h(\delta_k)_+ - h(\delta_k)_-].$$

With this extension, we can immediately say something about degrees of freedom, though only in a somewhat restricted setting. Suppose that $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ provides the fit $\hat{\mu} = f(y)$, and that f is actually univariate in each coordinate,

$$f(y) = (f_1(y_1), \dots, f_n(y_n)).$$

Suppose also that each coordinate function $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is p-absolutely continuous. We can apply Corollary 1 with $X = y_i$ and $h = f_i$, and sum over i to give

$$\begin{aligned} \text{df}(f) &= \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}(f_i(y_i), y_i) \\ &= \sum_{i=1}^n \mathbb{E}[f'_i(y_i)] + \frac{1}{\sigma} \sum_{i=1}^n \sum_{\delta \in \mathcal{D}(f_i)} \phi\left(\frac{\delta - \mu_i}{\sigma}\right) [f_i(\delta)_+ - f_i(\delta)_-]. \end{aligned} \quad (20)$$

The above expression provides an alternative way of proving the result on the degrees of freedom of hard thresholding, which was given in Lemma 1, the critical lemma for deriving the degrees of freedom of both best subset selection and the relaxed lasso for orthogonal predictors, Theorems 2 and 3. We step through this proof next.

Alternate proof of Lemma 1. For $f(y) = H_t(y)$, the i th coordinate function is

$$f_i(y_i) = [H_t(y_i)]_i = y_i \cdot 1\{|y_i| \geq t\},$$

which has a discontinuity set $\mathcal{D}(f_i) = \{-t, t\}$. The second term in (20) is hence

$$\frac{1}{\sigma} \sum_{i=1}^n \left[\phi\left(\frac{t - \mu_i}{\sigma}\right) \cdot (t - 0) + \phi\left(\frac{-t - \mu_i}{\sigma}\right) \cdot (0 - -t) \right] = \frac{t}{\sigma} \sum_{i=1}^n \left[\phi\left(\frac{t - \mu_i}{\sigma}\right) + \phi\left(\frac{t + \mu_i}{\sigma}\right) \right],$$

while the first term is simply

$$\sum_{i=1}^n \mathbb{E}[1\{|y_i| \geq t\}] = \mathbb{E}|\mathcal{A}_t|.$$

Adding these together gives

$$\text{df}(H_t) = \mathbb{E}|\mathcal{A}_t| + \frac{t}{\sigma} \sum_{i=1}^n \left[\phi\left(\frac{t - \mu_i}{\sigma}\right) + \phi\left(\frac{t + \mu_i}{\sigma}\right) \right],$$

precisely the conclusion of Lemma 1. □

5.2 An extension of Stein's multivariate lemma

The degrees of freedom result (20) applies to functions f for which the i th component function f_i depends only on the i th component of the input, $f_i(y) = f_i(y_i)$, for $i = 1, \dots, n$. Using this result, we could compute the degrees of freedom of the best subset selection and relaxed lasso fits in the orthogonal predictor matrix case. Generally speaking, however, we cannot use this result outside of the orthogonal setting, due to the requirement on f that $f_i(y) = f_i(y_i)$, $i = 1, \dots, n$. Therefore, in the hope of understanding degrees of freedom for procedures like best subset selection and the relaxed lasso in a broader context, we derive an extension of Stein's multivariate lemma.

Stein's multivariate lemma, Lemma 3, is concerned with functions $g : \mathbb{R}^n \rightarrow \mathbb{R}$ that are continuous and almost differentiable. Loosely speaking, the concept of almost differentiability is really a statement about absolute continuity. In words, a function is said to be almost differentiable if it is absolutely continuous on almost every line parallel to the coordinate axis (this notion is different, but equivalent, to that given by Stein). Before translating this mathematically, we introduce some notation. Let us write $x = (x_i, x_{-i})$ to emphasize that $x \in \mathbb{R}^n$ is determined by its i th component $x_i \in \mathbb{R}$ and the other $n - 1$ components $x_{-i} \in \mathbb{R}^{n-1}$. For $g : \mathbb{R}^n \rightarrow \mathbb{R}$, we let $g(\cdot, x_{-i})$ denote g as a function of the i th component alone, with all other components fixed at the value x_{-i} . We can now formally define almost differentiability:

Definition 2. We say that a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is *almost differentiable* if for every $i = 1, \dots, n$ and Lebesgue almost every $x_{-i} \in \mathbb{R}^{n-1}$, the function $g(\cdot, x_{-i}) : \mathbb{R} \rightarrow \mathbb{R}$ is absolutely continuous.

Similar to the univariate case, we propose a relaxed continuity condition. Namely:

Definition 3. We say that a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is *p-almost differentiable* if for every $i = 1, \dots, n$ and Lebesgue almost every $x_{-i} \in \mathbb{R}^{n-1}$, the function $g(\cdot, x_{-i}) : \mathbb{R} \rightarrow \mathbb{R}$ is p-absolutely continuous.

Note that a function g that is p-almost differentiable has partial derivatives almost everywhere, and we write the collection as $\nabla g = (\partial g / \partial x_1, \dots, \partial g / \partial x_n)$.² Also, when dealing with $g(\cdot, x_{-i})$, the function g restricted to its i th variable with all others fixed at x_{-i} , we write its one-sided limits as

$$g(x_i, x_{-i})_+ = \lim_{t \downarrow x_i} g(t, x_{-i}) \quad \text{and} \quad g(x_i, x_{-i})_- = \lim_{t \uparrow x_i} g(t, x_{-i}).$$

We are now ready to present our extension of Stein's multivariate lemma.

²Of course, this does not necessarily mean that g has a well-defined gradient, and so, cumbersome as it may read, we are careful about referring to ∇g as the vector of partial derivatives, instead of the gradient.

Lemma 5. Let $X \sim N(\mu, \sigma^2 I)$, for some fixed $\mu \in \mathbb{R}^n$ and $\sigma^2 > 0$. Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be p -almost differentiable, with vector of partial derivatives $\nabla g = (\partial g / \partial x_1, \dots, \partial g / \partial x_n)$. Then, for $i = 1, \dots, n$,

$$\frac{1}{\sigma^2} \mathbb{E}[(X_i - \mu_i)g(X)] = \mathbb{E} \left[\frac{\partial g}{\partial x_i}(X) \right] + \frac{1}{\sigma} \mathbb{E} \left[\sum_{\delta \in \mathcal{D}(g(\cdot, X_{-i}))} \phi \left(\frac{\delta - \mu_i}{\sigma} \right) [g(\delta, X_{-i})_+ - g(\delta, X_{-i})_-] \right],$$

provided that $\mathbb{E}|\partial g / \partial x_i(X)| < \infty$ and

$$\mathbb{E} \left| \sum_{\delta \in \mathcal{D}(g(\cdot, X_{-i}))} \phi \left(\frac{\delta - \mu_i}{\sigma} \right) [g(\delta, X_{-i})_+ - g(\delta, X_{-i})_-] \right| < \infty.$$

Refer to the appendix for the proof, which utilizes a conditioning argument to effectively reduce the multivariate setup to a univariate one, and then invokes Lemma 4.

The above lemma, Lemma 5, leads to our most general extension of Stein's formula for degrees of freedom. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a fitting procedure, as in $\hat{\mu} = f(y)$, and let $f(y) = (f_1(y), \dots, f_n(y))$. Consider Lemma 5 applied to the i th coordinate function, so that $X = y$ and $g = f_i$. Provided that each f_i is p -almost differentiable and satisfies the regularity conditions

$$\mathbb{E} \left| \frac{\partial f_i}{\partial y_i}(y) \right| < \infty \quad \text{and} \quad \mathbb{E} \left| \sum_{\delta \in \mathcal{D}(f_i(\cdot, y_{-i}))} \phi \left(\frac{\delta - \mu_i}{\sigma} \right) [f_i(\delta, y_{-i})_+ - f_i(\delta, y_{-i})_-] \right| < \infty, \quad (21)$$

we can take the i th equality in the lemma, and sum over i to give

$$\begin{aligned} \text{df}(f) &= \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}(f_i(y), y_i) \\ &= \sum_{i=1}^n \mathbb{E} \left[\frac{\partial f_i}{\partial y_i}(y) \right] + \frac{1}{\sigma} \sum_{i=1}^n \mathbb{E} \left[\sum_{\delta \in \mathcal{D}(f_i(\cdot, y_{-i}))} \phi \left(\frac{\delta - \mu_i}{\sigma} \right) [f_i(\delta, y_{-i})_+ - f_i(\delta, y_{-i})_-] \right]. \end{aligned} \quad (22)$$

Even if we assume that (22) is applicable to subset selection and the relaxed lasso with arbitrary predictors X , the discontinuity sets—and hence the second term in (22)—seem to be quite difficult to calculate in these cases. In other words, unfortunately, the formula (22) does not seem to provide an avenue for exact computation of the degrees of freedom of subset selection or the relaxed lasso in general. However, it may still help us understand these quantities, as we discuss next.

5.3 Potential insights from the multivariate Stein extension

For both of the best subset selection and relaxed lasso fitting procedures, one can show that the requisite regularity conditions (21) indeed hold, which makes the extended Stein formula (22) valid. Here we briefly outline a geometric interpretation for these fits, and describe how it can be used to understand their discontinuity sets, and the formula in (22). For an argument of a similar kind (and one given in more rigorous detail), see Tibshirani & Taylor (2012).

In both cases, we can decompose \mathbb{R}^n into a finite union of disjoint sets, $\mathbb{R}^n = \cup_{i=1}^m U_i$, with each U_i being polyhedral for the relaxed lasso, and each U_i an intersection of quadratic sublevel sets for subset selection. The relaxed lasso and best subset selection now share the property that, on the relative interior of each set U_i in their respective decompositions, the fit is just a linear projection map, and assuming that X has columns in general position, this is just the projection map onto the column space of $X_{\mathcal{A}}$ for some fixed set \mathcal{A} . Hence the discontinuity set of the fitting procedure in each case is contained in $\cup_{i=1}^m \text{relbd}(U_i)$, which has measure zero. In other words, the active set is locally constant for almost every $y \in \mathbb{R}^n$, and only for y crossing the relative boundary of some U_i does it change. From this, we can verify the appropriate regularity conditions in (21).

As f for the relaxed lasso and subset selection is the locally linear projection map $f(y) = P_{\mathcal{A}}y$, almost everywhere in y , the first term $\sum_{i=1}^n \mathbb{E}[\partial f_i(y)/\partial y_i]$ in (22) is simply $\mathbb{E}|\mathcal{A}|$. The second term, then, exactly coincides with the search degrees of freedom of these procedures. (Recall that the same breakdown occurred when using the univariate Stein extension to derive the degrees of freedom of hard thresholding, in Section 5.1.) This suggests a couple potential insights into degrees of freedom and search degrees of freedom that may be gleaned from the extended Stein formula (22), which we discuss below.

- *Positivity of search degrees of freedom.* If one could show that

$$f_i(\delta, y_{-i})_+ - f_i(\delta, y_{-i})_- > 0 \quad (23)$$

for each discontinuity point $\delta \in \mathcal{D}(f_i(\cdot, y_{-i}))$, almost every $y_{-i} \in \mathbb{R}^n$, and each $i = 1, \dots, n$, then this would imply that the second term in (22) is positive. For the relaxed lasso and subset selection fits, this would mean that the search degrees of freedom term is always positive, i.e., the total degrees of freedom of these procedures is always larger than the (expected) number of selected variables. In words, the condition in (23) says that the i th fitted value, at a point of discontinuity, can only increase as the i th component of y increases. Note that this is a sufficient but not necessary condition for positivity of search degrees of freedom.

- *Search degrees of freedom and discontinuities.* The fact that the second term in (22) gives the search degrees of freedom of the best subset selection and the relaxed lasso fits tells us that the search degrees of freedom of a procedure is intimately related to its discontinuities over y . At a high level: the greater the number of discontinuities, the greater the magnitude of these discontinuities, and the closer they occur to the true mean μ , the greater the search degrees of freedom.

This may provide some help in understanding the apparent (empirical) differences in search degrees of freedom between the relaxed lasso and best subset selection fits under correlated setups, as seen in Section 4. The particular discontinuities of concern in (22) arise from fixing all but i th component of the outcome at y_{-i} , and examining the i th fitted value $f_i(\cdot, y_{-i})$ as a function of its i th argument. One might expect that this function $f_i(\cdot, y_{-i})$ generally exhibits more points of discontinuity for best subset selection compared to the relaxed lasso, due to the more complicated boundaries of the elements U_i in the active-set-determining decomposition described above (these boundaries are piecewise quadratic for best subset selection, and piecewise linear for the relaxed lasso). This is in line with the general trend of subset selection displaying a larger search degrees of freedom than the relaxed lasso.

But, as demonstrated in Figure 6, something changes for large values of λ (small active sets, on average), and for $\mu = X\beta^*$ with a sparse or (especially) dense true coefficient vector β^* ; we saw that the search degrees of freedom of both the relaxed lasso and best subset selection fits can grow very large in these cases. Matching search degrees of freedom to the second term in (22), therefore, we infer that both fits must experience major discontinuities here (and these are somehow comparable overall, when measured in number, magnitude, and proximity to μ). This makes sense, especially when we think of taking λ large enough so that these procedures are forced to select an active set that is strictly contained in the true support $\mathcal{A}^* = \text{supp}(\beta^*)$; different values of y , quite close to $\mu = X\beta^*$, will make different subsets of \mathcal{A}^* look more or less appealing according to the criteria in (3), (4).

5.4 Connection to Theorem 2 of Hansen & Sokol (2014)

After completing this work, we discovered the independent and concurrent work of Hansen & Sokol (2014). These authors propose an interesting and completely different geometric approach to study-

ing the degrees of freedom of a metric projection estimator

$$f(y) \in \operatorname{argmin}_{u \in K} \|y - u\|_2^2,$$

where the set $K \subseteq \mathbb{R}^n$ can be nonconvex. Their Theorem 2 gives a decomposition for degrees of freedom that possesses an intriguing tie to ours in (22). Namely, these authors show that the degrees of freedom of any metric projection estimator f can be expressed as its expected divergence plus an “extra” term, this term being the integral of the normal density with respect to a singular measure (dependent on f). Equating this with our expression in (22), we see that the two forms of “extra” terms must match—i.e., our second term in (22), defined by a sum over the discontinuities of the projection f , must be equal to their integral.

This has an immediate implication for the projection operator onto the ℓ_0 ball of radius k , i.e., the best subset selection estimator in constrained form: the search degrees of freedom here must be nonnegative (as the integral of a density with respect to a measure is always nonnegative). The decomposition of Hansen & Sokol (2014) hence elegantly proves that the best subset selection fit, constrained to have k active variables, attains a degrees of freedom larger than or equal k . However, as far as we can tell, their Theorem 2 does not apply to best subset selection in Lagrange form, the estimator considered in our paper, since it is limited to metric projection estimators. To be clear, our extension of Stein’s formula in (22) is not restricted to any particular form of fitting procedure f (though we do require the regularity conditions in (21)).

We find the connections between our work and theirs fascinating, and hope to understand them more deeply in the future.

6 Discussion

In this work, we explored the degrees of freedom of best subset selection and the relaxed lasso (the procedure that performs least squares on the active set returned by the lasso). We derived exact expressions for the degrees of freedom of these fitting procedures with orthogonal predictors X , and investigated by simulation their degrees of freedom for correlated predictors. We introduced a new concept, search degrees of freedom, which intuitively measures the amount of degrees of freedom expended by an adaptive regression procedure in merely constructing an active set of variables (i.e., not counting the degrees of freedom attributed to estimating the active coefficients). Search degrees of freedom has a precise definition for any regression procedure. For subset selection and the relaxed lasso, this reduces to the (total) degrees of freedom minus the expected number of active variables; for the lasso, we simply equate its search degrees of freedom with that of the relaxed lasso, since these two procedures have the exact same search step.

The last section of this paper derived an extension of Stein’s formula for discontinuous functions. This was motivated by the hope that such a formula could provide an alternative lens from which we could view degrees of freedom for discontinuous fitting procedures like subset selection and the relaxed lasso. The application of this formula to these fitting procedures is not easy, and our grasp of the implications of this formula for degrees of freedom is only preliminary. There is much work to be done, but we are hopeful that our extension of Stein’s result will prove useful for understanding degrees of freedom and search degrees of freedom, and potentially, for other purposes as well.

Acknowledgements

The idea for this paper was inspired by a conversation with Jacob Bien. We thank Rob Tibshirani for helpful feedback and encouragement, and the editors and referees who read this paper and gave many useful comments and references.

A Proofs

A.1 Proof of Lemma 1

By definition,

$$\begin{aligned}
df(H_t) &= \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}([H_t(y)]_i, y_i) \\
&= \frac{1}{\sigma^2} \sum_{i=1}^n \mathbb{E} \left[y_i (y_i - \mu_i) \left(1\{y_i \geq t\} + 1\{y_i \leq -t\} \right) \right] \\
&= \frac{1}{\sigma^2} \sum_{i=1}^n \mathbb{E} \left[(z_i + \mu_i) z_i \left(1\{z_i \geq t - \mu_i\} + 1\{z_i \leq -t - \mu_i\} \right) \right], \tag{24}
\end{aligned}$$

where $z = y - \mu \sim N(0, \sigma^2 I)$. To compute the above, we note the identities (the last two can be checked using integration by parts):

$$\mathbb{E}[z_i \cdot 1\{z_i \leq a\}] = -\sigma \phi(a/\sigma), \tag{25}$$

$$\mathbb{E}[z_i \cdot 1\{z_i \geq b\}] = \sigma \phi(b/\sigma), \tag{26}$$

$$\mathbb{E}[z_i^2 \cdot 1\{z_i \leq a\}] = -\sigma a \phi(a/\sigma) + \sigma^2 \Phi(a/\sigma), \tag{27}$$

$$\mathbb{E}[z_i^2 \cdot 1\{z_i \geq b\}] = \sigma b \phi(b/\sigma) + \sigma^2 [1 - \Phi(b/\sigma)], \tag{28}$$

where Φ denotes the standard normal cdf. Plugging these in, the expression in (24) becomes

$$\sum_{i=1}^n \left[1 - \Phi\left(\frac{t - \mu_i}{\sigma}\right) + \Phi\left(\frac{-t - \mu_i}{\sigma}\right) \right] + \frac{t}{\sigma} \sum_{i=1}^n \left[\phi\left(\frac{t - \mu_i}{\sigma}\right) + \phi\left(\frac{-t - \mu_i}{\sigma}\right) \right],$$

and the first sum above is exactly

$$\mathbb{E} \left[\sum_{i=1}^n \left(1\{z_i \geq t - \mu_i\} + 1\{z_i \leq -t - \mu_i\} \right) \right] = \mathbb{E}|\mathcal{A}_t|,$$

as desired. \square

A.2 Proof of Theorem 2

As X is orthogonal, the criterion in (4) can be written as

$$\|y - X\beta\|_2^2 = \|X^T y - \beta\|_2^2 + c,$$

where c is a constant, meaning that it does not depend on β . Hence we can rewrite the optimization problem in (4) as

$$\hat{\beta}^{\text{subset}} \in \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2} \|X^T y - \beta\|_2^2 + \lambda \|\beta\|_0,$$

and from this it is not hard to see that the solution is

$$\hat{\beta}^{\text{subset}} = H_{\sqrt{2\lambda}}(X^T y),$$

hard thresholding the quantity $X^T y$ by the amount $t = \sqrt{2\lambda}$. Finally, we note that

$$df(X\hat{\beta}^{\text{subset}}) = \text{tr}\left(\text{Cov}(X\hat{\beta}^{\text{subset}}, y)\right) = \text{tr}\left(\text{Cov}(\hat{\beta}^{\text{subset}}, X^T y)\right),$$

because the trace operator is invariant under commutation of matrices, and $X^T y \sim N(X^T \mu, \sigma^2 I)$. Applying Lemma 1 completes the proof. \square

A.3 Proof of Theorem 1

We use several facts about the lasso without proof. These are derived in, e.g., Tibshirani & Taylor (2012) and Tibshirani (2013). For fixed X, λ , the lasso fit $f(y) = \hat{\mu}^{\text{lasso}}(y)$ is continuous and almost differentiable in each coordinate, so we can apply Stein's formula (19). As X has columns in general position, there is a unique lasso solution $\hat{\beta}^{\text{lasso}}$, and letting \mathcal{A} denote its active set, and s denote the signs of active lasso coefficients,

$$\mathcal{A} = \text{supp}(\hat{\beta}^{\text{lasso}}) \quad \text{and} \quad s = \text{sign}(\hat{\beta}_{\mathcal{A}}^{\text{lasso}}),$$

the fit can be expressed as

$$\hat{\mu}^{\text{lasso}} = X_{\mathcal{A}}(X_{\mathcal{A}}^T X_{\mathcal{A}})^{-1} X_{\mathcal{A}}^T y - X_{\mathcal{A}}(X_{\mathcal{A}}^T X_{\mathcal{A}})^{-1} \lambda s.$$

For almost every $y \in \mathbb{R}^n$, the set \mathcal{A} and vector s are locally constant (with respect to y), and so they have zero derivative (with respect to y). Hence, for almost every y ,

$$\sum_{i=1}^n \frac{\partial \hat{\mu}_i^{\text{lasso}}}{\partial y_i}(y) = \text{tr}\left(X_{\mathcal{A}}(X_{\mathcal{A}}^T X_{\mathcal{A}})^{-1} X_{\mathcal{A}}^T\right) = |\mathcal{A}|,$$

and taking an expectation gives the result. \square

A.4 Proof of Lemma 4

The result can be shown using integration by parts. We prove it in a different way, mimicking Stein's proof of Lemma 2, which makes the proof for the multivariate case (Lemma 5) easier. We have

$$\begin{aligned} \mathbb{E}[f'(Z)] &= \int_{-\infty}^{\infty} f'(z)\phi(z) dz \\ &= \int_0^{\infty} f'(z) \left\{ \int_z^{\infty} t\phi(t) dt \right\} dz - \int_{-\infty}^0 f'(z) \left\{ \int_{-\infty}^z t\phi(t) dt \right\} dz \\ &= \int_0^{\infty} t\phi(t) \left\{ \int_0^t f'(z) dz \right\} dt - \int_{-\infty}^0 t\phi(t) \left\{ \int_t^0 f'(z) dz \right\} dt. \end{aligned} \quad (29)$$

The second equality follows from $\phi'(t) = -t\phi(t)$, and the third is by Fubini's theorem. Consider the first term in (29); as f is absolutely continuous on each of the intervals $(-\infty, \delta_1), (\delta_1, \delta_2), \dots, (\delta_m, \infty)$, the fundamental theorem of (Lebesgue) integral calculus gives

$$\int_0^t f'(z) dz = f(t) - f(0) - \sum_{k=1}^m [f(\delta_k)_+ - f(\delta_k)_-] \cdot 1(0 \leq \delta_k \leq t).$$

Therefore

$$\int_0^{\infty} t\phi(t) \left\{ \int_0^t f'(z) dz \right\} dt = \int_0^{\infty} t\phi(t)[f(t) - f(0)] dt - \sum_{\delta_k \geq 0} [f(\delta_k)_+ - f(\delta_k)_-] \int_{\delta_k}^{\infty} t\phi(t) dt.$$

The second term in (29) is similar, and putting these together we get

$$\begin{aligned} \mathbb{E}[f'(Z)] &= \mathbb{E}[Zf(Z)] - \mathbb{E}[Z]f(0) - \sum_{\delta_k \geq 0} [f(\delta_k)_+ - f(\delta_k)_-] \cdot \mathbb{E}[Z \cdot 1\{Z \geq \delta_k\}] + \\ &\quad \sum_{\delta_k < 0} [f(\delta_k)_+ - f(\delta_k)_-] \cdot \mathbb{E}[Z \cdot 1\{Z \leq \delta_k\}]. \end{aligned}$$

The result follows by noting that $\mathbb{E}[Z] = 0$ and recalling the identities (25) and (26). \square

A.5 Proof of Corollary 1

Define $Z = (X - \mu)/\sigma$ and $f(z) = h(\sigma z + \mu)$, and apply Lemma 4. \square

A.6 Proof of Lemma 5

We assume that $X \sim N(0, I)$, and then a similar standardization argument to that given in the proof of Corollary 1 can be applied here to prove the result for $X \sim N(\mu, \sigma^2 I)$.

For fixed X_{-i} , the function $g(\cdot, X_{-i})$ is univariate. Hence, following the proof of Lemma 4, and using the independence of X_i and X_{-i} ,

$$\begin{aligned} \mathbb{E}\left[\frac{\partial g}{\partial x_i}(X) \mid X_{-i}\right] &= \int_{-\infty}^{\infty} \frac{\partial g}{\partial x_i}(z, X_{-i}) \phi(z) dz \\ &= \int_0^{\infty} \frac{\partial g}{\partial x_i}(z, X_{-i}) \left\{ \int_z^{\infty} t \phi(t) dt \right\} dz - \int_{-\infty}^0 \frac{\partial g}{\partial x_i}(z, X_{-i}) \left\{ \int_{-\infty}^z t \phi(t) dt \right\} dz \\ &= \int_0^{\infty} t \phi(t) \left\{ \int_0^t \frac{\partial g}{\partial x_i}(z, X_{-i}) dz \right\} dt - \int_{-\infty}^0 t \phi(t) \left\{ \int_t^0 \frac{\partial g}{\partial x_i}(z, X_{-i}) dz \right\} dt. \end{aligned}$$

Consider the first term above. For almost every X_{-i} , the function $g(\cdot, X_{-i})$ is p-absolutely continuous, so the inner integral is

$$\int_0^t \frac{\partial g}{\partial x_i}(z, X_{-i}) dz = g(t, X_{-i}) - g(0, X_{-i}) - \sum_{\delta \in \mathcal{D}_i} [g(z, X_{-i})_+ - g(z, X_{-i})_-] \cdot \mathbf{1}(0 \leq \delta \leq t),$$

where we have abbreviated $\mathcal{D}_i = \mathcal{D}(g(\cdot, X_{-i}))$. The next steps follow the corresponding arguments in the proof of Lemma 4, yielding

$$\mathbb{E}\left[\frac{\partial g}{\partial x_i}(X) \mid X_{-i}\right] = \mathbb{E}[X_i g(X) \mid X_{-i}] - \sum_{\delta \in \mathcal{D}_i} [g(\delta, X_{-i})_+ - g(\delta, X_{-i})_-]$$

for almost every X_{-i} . Taking an expectation over X_{-i} gives the result. \square

References

- Chen, S., Donoho, D. L. & Saunders, M. (1998), ‘Atomic decomposition for basis pursuit’, *SIAM Journal on Scientific Computing* **20**(1), 33–61.
- Deledalle, C.-A., Peyre, G. & Fadili, J. (2013), Stein consistent risk estimator (SCORE) for hard thresholding. arXiv: 1301.5874.
- Efron, B. (1986), ‘How biased is the apparent error rate of a prediction rule?’, *Journal of the American Statistical Association: Theory and Methods* **81**(394), 461–470.
- Hansen, N. & Sokol, A. (2014), Degrees of freedom for nonlinear least squares estimation. arXiv: 1402.2997.
- Hastie, T. & Tibshirani, R. (1990), *Generalized Additive Models*, Chapman and Hall, London.
- Janson, L., Fithian, W. & Hastie, T. (2013), Effective degrees of freedom: A flawed metaphor. arXiv: 1312.7851.
- Kaufman, S. & Rosset, S. (2013), When does more regularization imply fewer degrees of freedom? Sufficient conditions and counter examples from lasso and ridge regression. arXiv: 1311.2791.

- Mallows, C. (1973), ‘Some comments on C_p ’, *Technometrics* **15**(4), 661–675.
- Mazumder, R., Friedman, J. & Hastie, T. (2011), ‘SparseNet: Coordinate descent with nonconvex penalties’, *Journal of the American Statistical Association* **106**(495), 1125–1138.
- Meinshausen, N. (2007), ‘Relaxed lasso’, *Computational Statistics & Data Analysis* **52**, 374–393.
- Stein, C. (1981), ‘Estimation of the mean of a multivariate normal distribution’, *Annals of Statistics* **9**(6), 1135–1151.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *Journal of the Royal Statistical Society: Series B* **58**(1), 267–288.
- Tibshirani, R. J. (2013), ‘The lasso problem and uniqueness’, *Electronic Journal of Statistics* **7**, 1456–1490.
- Tibshirani, R. J. & Taylor, J. (2011), ‘The solution path of the generalized lasso’, *Annals of Statistics* **39**(3), 1335–1371.
- Tibshirani, R. J. & Taylor, J. (2012), ‘Degrees of freedom in lasso problems’, *Annals of Statistics* **40**(2), 1198–1232.
- Tibshirani, R. & Knight, K. (1999), ‘The covariance inflation criterion for adaptive model selection’, *Journal of the Royal Statistical Society: Series B* **61**(3), 529–546.
- Ye, J. (1998), ‘On measuring and correcting the effects of data mining and model selection’, *Journal of the American Statistical Society* **93**(441), 120–131.
- Zou, H., Hastie, T. & Tibshirani, R. (2007), ‘On the “degrees of freedom” of the lasso’, *Annals of Statistics* **35**(5), 2173–2192.