

Inference in Adaptive Regression via the Kac-Rice Formula

Jonathan E. Taylor, Joshua R. Loftus, Ryan J. Tibshirani

*Department of Statistics
Stanford University
Sequoia Hall
Stanford, CA 94305, USA
e-mail: jtaylo@stanford.edu
joftius@stanford.edu*

*Department of Statistics
Carnegie Mellon University
Baker Hall
Pittsburgh, PA 15213, USA
e-mail: ryantibs@cmu.edu*

Abstract: We derive an exact p-value for testing a global null hypothesis in a general adaptive regression setting. Our approach uses the Kac-Rice formula (as described in Adler & Taylor 2007) applied to the problem of maximizing a Gaussian process. The resulting test statistic has a known distribution in finite samples, assuming Gaussian errors. We examine this test statistic in the case of the lasso, group lasso, principal components, and matrix completion problems. For the lasso problem, our test relates closely to the recently proposed covariance test of Lockhart et al. (2014).

In a few specific settings, our proposed tests will be less powerful than other previously-known (and well-established) tests. However, it should be noted that the real strength of our proposal here is its generality. We provide a framework for constructing valid tests across a wide class of regularized regression problems, and as far as we can tell, such a unified view was not possible before this work.

AMS 2000 subject classifications: Primary 62M40; secondary 62J05.
Keywords and phrases: Gaussian processes, convex analysis, regularized regression.

1. Introduction

In this work, we consider the problem of finding the distribution of

$$\max_{\eta \in \mathcal{K}} \eta^T \epsilon, \quad \epsilon \sim N(0, \Theta), \quad (1)$$

for a convex set $\mathcal{K} \subseteq \mathbb{R}^p$. In other words, we study a Gaussian process with a finite Karhunen-Loève expansion (Adler & Taylor 2007), restricted to a convex set in \mathbb{R}^p .

While this is a well-studied topic in the literature of Gaussian processes, our aim here is to describe an implicit formula for both the distribution of (1), as

well as the almost surely unique maximizer

$$\eta^* = \operatorname{argmax}_{\eta \in \mathcal{K}} \eta^T \epsilon. \quad (2)$$

A main point of motivation underlying our work is the application of such a formula for inference in modern statistical estimation problems. We note that a similar (albeit simpler) formula has proven useful in problems related to sparse regression (Lockhart et al. 2014, Taylor 2013). Though the general setting considered in this paper is ultimately much more broad, we begin by discussing the sparse regression case.

1.1. Example: the lasso

As a preview, consider the ℓ_1 penalized regression problem, i.e., the lasso problem (Tibshirani 1996), of the form

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

where $y \in \mathbb{R}^n$, $X \in \mathbb{R}^{n \times p}$, and $\lambda \geq 0$. Very mild conditions on the predictor matrix X ensure uniqueness of the lasso solution $\hat{\beta}$, see, e.g., Tibshirani (2013). Treating X as fixed, we assume that the outcome y satisfies

$$y \sim N(X\beta_0, \Sigma), \quad (4)$$

where $\beta_0 \in \mathbb{R}^p$ is some fixed true (unknown) coefficient vector, and $\Sigma \in \mathbb{R}^{p \times p}$ is a known covariance matrix. In the following sections, we derive a formula that enables a test of a global null hypothesis H_0 in a general regularized regression setting. Our main result, Theorem 1, can be applied to the lasso problem in order to test the null hypothesis $H_0 : \beta_0 = 0$. This test involves the quantity

$$\lambda_1 = \|X^T y\|_\infty,$$

which can be seen as the first knot (i.e., critical value) in the lasso solution path over the regularization parameter λ (Efron et al. 2004). Recalling the duality of the ℓ_1 and ℓ_∞ norms, we can rewrite this quantity as

$$\lambda_1 = \max_{\eta \in \mathcal{K}} \eta^T (X^T y), \quad (5)$$

where $\mathcal{K} = \{\eta : \|\eta\|_1 \leq 1\}$, showing that λ_1 is of the form (1), with $\epsilon = X^T y$ (which has mean zero under the null hypothesis). Assuming uniqueness of the entries of $X^T y$, the maximizer η^* in (5) is

$$\eta_j^* = \begin{cases} \operatorname{sign}(X_j^T y) & |X_j^T y| = \|X^T y\|_\infty \\ 0 & \text{otherwise} \end{cases}, \quad j = 1, \dots, p.$$

Let j^* denote the maximizing index, so that $|X_{j^*}^T y| = \|X^T y\|_\infty$, and also $s^* = \text{sign}(X_{j^*}^T y)$, $\Theta_{jk} = X_j^T \Sigma X_k$. To express our test statistic, we define

$$\mathcal{V}_{\eta^*}^- = \max_{\substack{s \in \{-1, 1\}, k \neq j \\ 1 - s\Theta_{j^*k}/\Theta_{j^*j^*} > 0}} \frac{s(X_k - \Theta_{j^*k}/\Theta_{j^*j^*} X_{j^*})^T y}{1 - s\Theta_{j^*k}/\Theta_{j^*j^*}},$$

$$\mathcal{V}_{\eta^*}^+ = \min_{\substack{s \in \{-1, 1\}, k \neq j \\ 1 - s\Theta_{j^*k}/\Theta_{j^*j^*} < 0}} \frac{s(X_k - \Theta_{j^*k}/\Theta_{j^*j^*} X_{j^*})^T y}{1 - s\Theta_{j^*k}/\Theta_{j^*j^*}}.$$

Then under $H_0 : \beta_0 = 0$, we prove that

$$\frac{\Phi(\mathcal{V}_{\eta^*}^+/\Theta_{j^*j^*}^{1/2}) - \Phi(\lambda_1/\Theta_{j^*j^*}^{1/2})}{\Phi(\mathcal{V}_{\eta^*}^+/\Theta_{j^*j^*}^{1/2}) - \Phi(\mathcal{V}_{\eta^*}^-/\Theta_{j^*j^*}^{1/2})} \sim \text{Unif}(0, 1), \quad (6)$$

where Φ is the standard normal cumulative distribution function. This formula is somewhat remarkable, in that it is exact—not asymptotic in n, p —and relies only on the assumption of normality for y in (4) (with essentially no real restrictions on the predictor matrix X). As mentioned above, it is a special case of Theorem 1, the main result of this paper.

We refer to the above test statistic (6) as the *Kac-Rice test* for the lasso. It may seem complicated, but when the predictors are standardized, $\|X_j\|_2 = 1$ for $j = 1, \dots, p$, and the observations are independent with (say) unit marginal variance, $\Sigma = I$, then $\mathcal{V}_{\eta^*}^-$ is equal to the second knot λ_2 in the lasso path and $\mathcal{V}_{\eta^*}^+$ is equal to ∞ . Therefore (6) simplifies to

$$\frac{1 - \Phi(\lambda_1)}{1 - \Phi(\lambda_2)} \sim \text{Unif}(0, 1). \quad (7)$$

This statistic measures the relative sizes of λ_1 and λ_2 , with values of $\lambda_1 \gg \lambda_2$ being evidence against the null hypothesis.

Figure 1(a) shows the empirical distribution function of a sample of 20,000 p-values from (6), over lasso problems with a variety of different predictor matrices, all under the global null model $\beta_0 = 0$. In particular, for each sample, we drew the matrix X uniformly at random from the following cases:¹

- small case: X is 3×2 , with values (in row-major order) equal to $1, 2, \dots, 6$;
- fat case: X is $100 \times 10,000$, with columns drawn from the compound symmetric Gaussian distribution having correlation 0.5;
- tall case: X is $10,000 \times 100$, with columns drawn from the compound symmetric Gaussian distribution having correlation 0.5;
- lower triangular case: X is 500×500 , a lower triangular matrix of 1s [the lasso problem here is effectively a reparametrization of the 1-dimensional fused lasso problem (Tibshirani et al. 2005)];

¹Here we only drew the matrix X at random across the different cases in order to illustrate that this test holds for many different designs. For each fixed choice of design, the plot would look exactly the same.

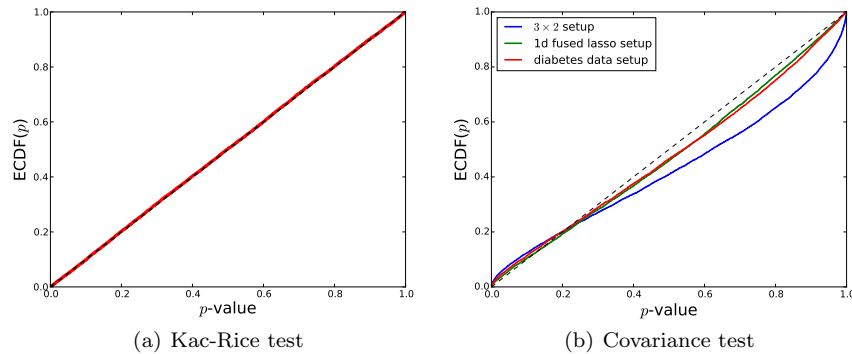


FIG 1. The left panel shows the empirical distribution function of a sample of 20,000 p-values from (6), coming from a variety of different lasso setups. The agreement with uniform here is excellent. The right panel shows the empirical distribution function of a sample of 10,000 covariance test p-values, computed using an $\text{Exp}(1)$ approximation, using three different lasso setups. The $\text{Exp}(1)$ approximation is generally conservative, whereas the Kac-Rice test is exact.

- diabetes data case: X is 442×10 , the diabetes data set studied in Efron et al. (2004).

As is seen in the plot, the agreement with uniform is very strong.

In their proposed *covariance test*, Lockhart et al. (2014) show that under the global null hypothesis $H_0 : \beta_0 = 0$,

$$\lambda_1(\lambda_1 - \lambda_2) \xrightarrow{d} \text{Exp}(1) \quad \text{as } n, p \rightarrow \infty, \quad (8)$$

assuming standardized predictors, $\|X_j\|_2 = 1$ for $j = 1, \dots, p$, independent errors in (4) with unit marginal variance, $\Sigma = I$, and a condition to ensure that λ_2 diverges to ∞ at a sufficient rate.

In finite samples, using $\text{Exp}(1)$ as an approximation to the distribution of the covariance test statistic seems generally conservative, especially for smaller values of n and p . Figure 1(b) shows the empirical distribution function from 10,000 covariance test p-values, in three of the above scenarios. [The predictors were standardized before applying the covariance test, in all three cases; this is not necessary, as the covariance test can be adapted to the more general case of unstandardized predictors, but was done for simplicity, to match the form of the test as written in (8).] Even though the idea behind the covariance test can be conceivably extended to other regularized regression problems (outside of the lasso setting), the $\text{Exp}(1)$ approximation to its distribution is generally inappropriate, as we will see in later examples. Our test, however, naturally extends to general regularization settings, allowing us to attack problems with more complex penalties such as the group lasso and nuclear norm penalties.

It is important now to emphasize the main goal of this paper. We provide a

general recipe for constructing global tests in regularized regression problems. We do not claim these tests to be most powerful across all settings; indeed, in certain settings, our test will have less power than other well-studied alternative approaches. The strength of our proposal is that it applies so broadly, and furthermore, that it leads to selective tests, beyond the global null, for many of these problems as well (pursued in subsequent papers after this one). Power and selective inference are the topics of Sections 6.2 and 6.3 of the discussion.

The rest of this paper is organized as follows. In Section 2, we describe the general framework for regularized regression problems that we consider, and a corresponding global null hypothesis of interest; we also state our main result, Theorem 1, which gives an exact p-value for this null hypothesis. The next two sections are then dedicated to proving Theorem 1. Section 3 characterizes the global maximizer (2) in terms of the related Gaussian process and its gradient. Section 4 applies the Kac-Rice formula to derive the joint distribution of the maximum value of the process (1) and its maximizer (2), which is ultimately used to derive the (uniform) distribution of our proposed test. In Section 5 we broadly consider practicalities associated with our test statistic, revisit the lasso problem, and examine the group lasso, principal components, and matrix completion problems as well. We end with a discussion in Section 6.

2. General regularized regression problems

We examine a class of regularized least squares problems of the form

$$\hat{\beta} \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \cdot \mathcal{P}(\beta), \quad (9)$$

with outcome $y \in \mathbb{R}^n$, predictor matrix $X \in \mathbb{R}^{n \times p}$, and regularization parameter $\lambda \geq 0$. We assume that the penalty function \mathcal{P} satisfies

$$\mathcal{P}(\beta) = \max_{u \in C} u^T \beta, \quad (10)$$

where $C \subseteq \mathbb{R}^p$ is a convex set, i.e., \mathcal{P} is the support function of C . This serves as a very general penalty, as we can represent any seminorm (and hence any norm) in this form with the proper choice of set C . We note that the solution $\hat{\beta} = \hat{\beta}_\lambda$ above is not necessarily unique (depending on the matrix X and set C) and the element notation used in (9) reflects this. A standard calculation, which we give in Appendix A.1, shows that the fitted value $X\hat{\beta}$ is always unique, and hence so is $\mathcal{P}(\hat{\beta})$.

Now define

$$\lambda_1 = \min\{\lambda \geq 0 : \mathcal{P}(\hat{\beta}_\lambda) = 0\}.$$

This is the smallest value of λ for which the penalty term in (9) is zero; any smaller value of the tuning parameter returns a non-trivial solution, according to the penalty. A straightforward calculation involving subgradients, which we give in Appendix A.2, shows that

$$\lambda_1 = \mathcal{Q}(X^T(I - P_{XC^\perp})y), \quad (11)$$

where \mathcal{Q} is the dual seminorm of \mathcal{P} , i.e., $\mathcal{Q}(\beta) = \max_{v \in C^\circ} v^T \beta$, the support function of the polar set C° of C . This set can be defined as $C^\circ = \{v : v^T u \leq 1 \text{ for all } u \in C\}$, or equivalently,

$$C^\circ = \{v : \mathcal{P}(v) \leq 1\},$$

the unit ball in \mathcal{P} . Furthermore, in (11), we use P_{XC^\perp} to denote the projection operator onto the linear subspace

$$XC^\perp = X\{v : v \perp C\} \subseteq \mathbb{R}^n.$$

We recall that for the lasso problem (3), the penalty function is $\mathcal{P}(\beta) = \|\beta\|_1$, so $\mathcal{Q}(\beta) = \|\beta\|_\infty$; also $C = \{u : \|u\|_\infty \leq 1\}$, which means that $C^\perp = 0$, and $P_{XC^\perp} = 0$. Hence $\lambda_1 = \|X^T y\|_\infty$, as claimed in Section 1.1.

Having just given an example of a seminorm in which C is of full dimension p , so that $C^\perp = \{0\}$, we now consider one in which C has dimension less than p , so that C^\perp is nontrivial. In a generalized lasso problem (Tibshirani & Taylor 2011), the penalty is $\mathcal{P}(\beta) = \|D\beta\|_1$ for some chosen penalty matrix $D \in \mathbb{R}^{m \times n}$. In this case, it can be shown that the dual seminorm is $\mathcal{Q}(\beta) = \min_{D^T z = \beta} \|z\|_\infty$. Hence $C = \{u : \min_{D^T z = u} \|z\|_\infty \leq 1\}$, and $C^\perp = \text{null}(D)$, the null space of D . In many interesting cases, this null space is nontrivial; e.g., if D is the fused lasso penalty matrix, then its null space is spanned by the vector of all 1s.

2.1. A null hypothesis

As in the lasso case, we assume that y is generated from the normal model

$$y \sim N(X\beta_0, \Sigma), \quad (12)$$

with X considered fixed. We are interested in the distribution of (1) in order to test the following hypothesis:

$$H_0 : \mathcal{P}(\beta_0) = 0. \quad (13)$$

This can be seen a global null hypothesis, a test of whether the true underlying coefficient vector β_0 has a trivial structure, according to the designated penalty function \mathcal{P} .

Assuming that the set C contains 0 in its relative interior, we have $\mathcal{P}(\beta) = 0 \iff P_C \beta = 0$, where P_C denotes the projection matrix onto $\text{span}(C)$. Therefore we can rewrite the null hypothesis (13) in a more transparent form, as

$$H_0 : P_C \beta_0 = 0. \quad (14)$$

Again, using the lasso problem (3) as a reference point, we have $\text{span}(C) = \mathbb{R}^p$ for this problem, so the above null hypothesis reduces to $H_0 : \beta_0 = 0$, as in Section 1.1. In general, the null hypothesis (14) tests $\beta_0 \in C^\perp$, the orthocomplement of C .

Recalling that

$$\lambda_1 = \max_{v \in C^\circ} v^T X^T (I - P_{XC^\perp}) y,$$

one can check that, under H_0 , the quantity λ_1 is precisely of the form (1), with $\mathcal{K} = C^\circ$, $\epsilon = X^T (I - P_{XC^\perp}) y$, and $\Theta = \text{Cov}(\epsilon) = X^T (I - P_{XC^\perp}) \Sigma (I - P_{XC^\perp}) X$ [as $\mathbb{E}(\epsilon) = X^T (I - P_{XC^\perp}) X \beta_0 = 0$ when $\beta_0 \in C^\perp$].

2.2. Statement of main result and outline of our approach

We now state our main result.

Theorem 1 (Kac-Rice test). *Consider the general regularized regression problem in (9), with $\mathcal{P}(\beta) = \max_{u \in C} u^T \beta$ for a closed, convex set $C \subseteq \mathbb{R}^p$ containing 0 in its relative interior. Denote $\mathcal{K} = C^\circ = \{v : \mathcal{P}(v) \leq 1\}$, the polar set of C , and assume that \mathcal{K} can be stratified into pieces of different dimensions, i.e.,*

$$\mathcal{K} = \bigcup_{j=0}^p \partial_j \mathcal{K}, \quad (15)$$

where $\partial_0 \mathcal{K}, \dots, \partial_p \mathcal{K}$ are smooth disjoint manifolds of dimensions $0, \dots, p$, respectively. Assume also assume that the process

$$f_\eta = \eta^T X^T (I - P_{XC^\perp}) y, \quad \eta \in \mathcal{K}, \quad (16)$$

is Morse for almost every $y \in \mathbb{R}^n$. Finally, assume that $y \in \mathbb{R}^n$ is drawn from the normal distribution in (12).

Now, consider testing the null hypothesis $H_0 : \mathcal{P}(\beta_0) = 0$ [equivalently, $H_0 : P_C \beta_0 = 0$, since we have assumed that $0 \in \text{relint}(C)$]. Define $\Lambda_\eta = G_\eta^{-1} H_\eta$ for G_η, H_η as in (29), (30), $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+$ as in (24), (23), and σ_η^2 as in (32). Finally, let η^* denote the almost sure unique maximizer of the process f_η over \mathcal{K} ,

$$\eta^* = \operatorname{argmax}_{\eta \in \mathcal{K}} f_\eta,$$

and let $\lambda_1 = f_{\eta^*}$ denote the first knot in the solution path of problem (9). Then under H_0 ,

$$\mathbb{S} = \frac{\int_{\lambda_1}^{\mathcal{V}_{\eta^*}^+} \det(\Lambda_{\eta^*} + zI) \phi_{\sigma_{\eta^*}^2}(z) dz}{\int_{\mathcal{V}_{\eta^*}^-}^{\mathcal{V}_{\eta^*}^+} \det(\Lambda_{\eta^*} + zI) \phi_{\sigma_{\eta^*}^2}(z) dz} \sim \text{Unif}(0, 1), \quad (17)$$

where ϕ_{σ^2} denotes the density function of a normal random variable with mean 0 and variance σ^2 .

The quantity (17) is the *Kac-Rice pivot* evaluated at $\mu = 0$. Lemma 4 shows it is a pivotal quantity for the mean μ near η^* , derived via the Kac-Rice formula. Here we give a rough explanation of the result in (17), and the approach

we take to prove it in Sections 3 and 4. The next section, Section 2.3, discusses the assumptions behind Theorem 1; in summary, the assumption that \mathcal{K} separates as in (15) allows us to apply the Kac-Rice formula to each of its strata, and the Morse assumption on the process f_η in (16) ensures the uniqueness of its maximizer η^* . These are very weak assumptions, especially considering the strength of the exact, non-asymptotic conclusion in (17).

Our general approach is based on finding an implicit formula for $\mathbb{P}(\lambda_1 > t)$ under the null hypothesis H_0 , where λ_1 is the first knot in the solution path of problem (9) and can be written as

$$\lambda_1 = \max_{\eta \in \mathcal{K}} f_\eta,$$

where $f_\eta = \eta^T X^T (I - P_{XC^\perp})y$, the process in (16). Our representation for the tail probability of λ_1 has the form

$$\mathbb{P}(\lambda_1 > t) = \mathbb{E}(\mathbb{Q}(1_{(t, \infty)})). \quad (18)$$

Here \mathbb{Q} is a random distribution function and $1_{(t, \infty)}$ is the indicator function for the interval (t, ∞) . The distribution \mathbb{Q} depends on η^* , a maximizer of the process f_η , almost surely unique by the Morse assumption. This maximizer satisfies

$$\eta^* \in \partial \mathcal{Q}(X^T (I - P_{XC^\perp})y) \subseteq \mathcal{K},$$

with $\partial \mathcal{Q}$ the subdifferential of the seminorm \mathcal{Q} . Under the assumption that $\mathcal{K} = \cup_{j=0}^p \partial_j \mathcal{K}$, the main tool we invoke is the Kac-Rice formula (Adler & Taylor 2007), which essentially enables us to compute the expected number² of global maximizers occurring in each stratum $\partial_j \mathcal{K}$. This leads to the distribution of λ_1 , in Theorem 2, as well as the representation in (18), with \mathbb{Q} given in an explicit form. Unfortunately, computing tail probabilities $\mathbb{P}(\lambda_1 > t)$ of this distribution involve evaluating some complicated integrals over \mathcal{K} that depend on X, Σ , and hence the quantity λ_1 as a test statistic does not easily lend itself to the computation of p-values. We therefore turn to the survival function \mathbb{S} associated with the measure \mathbb{Q} , and our main result is that, when carefully evaluated, this (random) survival function can be used to derive a test of H_0 , as expressed in (17) in Theorem 1 above.

2.3. Discussion of assumptions

In terms of the assumptions of Theorem 1, we require that C contains 0 in its relative interior so that we can write the null hypothesis in the equivalent form

²Note that for almost every realization, under the Morse assumption, there is generically only one maximizer overall and hence the number of them is either 0 or 1. We use the term “number of global maximizers” when applying the Kac-Rice formula as it applies to counting different types of points. In our applications of it, however, there is only ever 0 or 1 such points. Similar arguments were used to establish the accuracy of the expected Euler characteristic approximation for the distribution of the global maximum of a smooth Gaussian process in Taylor et al. (2005).

$H_0 : P_C \beta_0 = 0$, which makes the process f_η in (16) have mean zero under H_0 . We additionally assume that C is closed in order to guarantee that f_η has a well-defined (finite) maximum over $\eta \in \mathcal{K} = C^\circ$. See Appendix A.3.

Apart from these rather minor assumptions on C , the main requirements of the theorem are: the polar set $C^\circ = \mathcal{K}$ can be stratified as in (15), the process f_η in (16) is Morse, and y follows the normal distribution in (12). Overall, these are quite weak assumptions. The first assumption, on \mathcal{K} separating as in (15), eventually permits us to apply to the Kac-Rice formula to each stratum $\partial_j \mathcal{K}$. We remark that many convex sets possess such a decomposition, see Adler & Taylor (2007). In particular, note that such an assumption does not limit our consideration to polyhedral \mathcal{K} : a set can be stratifiable but still have a boundary with curvature (e.g., as in \mathcal{K} for the group lasso and nuclear norm penalties).

Further, the property of being a Morse function is truly generic; again, see Adler & Taylor (2007) for a discussion of Morse functions on stratified spaces. If f_η is Morse for almost every y , then its maximizers are almost surely isolated, and the convexity of \mathcal{K} then implies that f_η has an almost surely unique maximizer η^* . From the form of our particular process f_η in (16), the assumption that f_η is Morse can be seen as a restriction on the predictor matrix X (or more generally, how X interacts with the set C). For most problems, this only rules out trivial choices of X . In the lasso case, for example, recall that $f_\eta = (X\eta)^T y$ and \mathcal{K} is equal to the unit ℓ_1 ball, so $f_\eta^* = \|X^T y\|_\infty$, and the Morse property requires $|X_j^T y|$, $j = 1, \dots, p$ to be unique for almost every $y \in \mathbb{R}^p$. This can be ensured by taking X with columns in general position [a weak condition that also ensures uniqueness of the lasso solution; see Tibshirani (2013)].

The assumption of normally distributed errors in the regression model (12) is important for the work that follows in Sections 3 and 4, which is based on Gaussian process theory. Note that we assume a known covariance matrix Σ , but we allow for a dependence between the errors (i.e., Σ need not be diagonal). Empirically, the (uniform) distribution of our test statistic under the null hypothesis appears to quite robust against non-normal errors in many cases of interest; we discuss this in Section 6.1.

Lastly, we make a note about the role of convexity. Aside from computing $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+$ (see Appendix A.6), the convexity of the parameter space \mathcal{K} is not necessary; we only need local convexity as described in Adler & Taylor (2007), i.e., we only need to assume that the support cone of \mathcal{K} is locally convex everywhere. This is essentially the same as positive reach (Federer 1959). To be clear, while convexity is used in connecting the Kac-Rice test to the regularized regression problem in (11) (i.e., it establishes an equality between left- and right-hand sides), the right-hand side is well-defined even if C° is not convex. That is, the set \mathcal{K} in (1) need not be convex. In fact, the issue of convexity is only important for computational purposes, not theoretical purposes. In this sense, the current paper provides an exact conditional test based on the global maximizer of a smooth Gaussian field on a fairly arbitrary set. This is an advance in the theory of smooth Gaussian fields as developed in Adler & Taylor (2007) and will be investigated in future work.

2.4. Notation

Rewrite the process f_η in (16) as

$$f_\eta = \eta^T X^T (I - P_{XC^\perp})(I - P_{XC^\perp})y = \eta^T \tilde{X}^T \tilde{y}, \quad \eta \in \mathcal{K},$$

where $\tilde{X} = (I - P_{XC^\perp})X$ and $\tilde{y} = (I - P_{XC^\perp})y$. The distribution of \tilde{y} is hence $\tilde{y} \sim N(\tilde{X}\beta_0, \tilde{\Sigma})$, where $\tilde{\Sigma} = (I - P_{XC^\perp})\Sigma(I - P_{XC^\perp})$. Furthermore, under the null hypothesis $H_0 : P_C\beta_0 = 0$, we have $\tilde{y} \sim N(0, \tilde{\Sigma})$. For convenience, in Sections 3 and 4, we will drop the tilde notation, and write $\tilde{y}, \tilde{X}, \tilde{\Sigma}$ as simply y, X, Σ , respectively. To be perfectly explicit, this means that we will write the process f_η in (16) as

$$f_\eta = \eta^T X^T y, \quad \eta \in \mathcal{K},$$

where $y \sim N(X\beta_0, \Sigma)$, and the null hypothesis is $H_0 : y \sim N(0, \Sigma)$. Notice that when $\text{span}(C) = \mathbb{R}^p$, we have exactly $\tilde{y} = y, \tilde{X} = X, \tilde{\Sigma} = \Sigma$, since $P_{XC^\perp} = 0$. However, we reiterate that replacing $\tilde{y}, \tilde{X}, \tilde{\Sigma}$ by y, X, Σ in Sections 3 and 4 is done purely for notational convenience, and the reader should bear in mind that the arguments themselves do not portray any loss of generality.

We will write \mathbb{E}_0 to emphasize that an expectation is taken under the null distribution $H_0 : y \sim N(0, \Sigma)$.

3. Characterization of the global maximizer

Near any point $\eta \in \mathcal{K}$, the set \mathcal{K} is well-approximated by the support cone $S_\eta\mathcal{K}$, which is defined as the polar cone of the normal cone $N_\eta\mathcal{K}$. The support cone $S_\eta\mathcal{K}$ contains a largest linear subspace—we will refer to this $T_\eta\mathcal{K}$, the *tangent space* to \mathcal{K} at η . The tangent space plays an important role in what follows.

We study the process f_η in (16), which we now write as $f_\eta = \eta^T X^T y$ over $\eta \in \mathcal{K}$, where $y \sim N(X\beta_0, \Sigma)$, and with the null hypothesis $H_0 : y \sim N(0, \Sigma)$ (see our notational reduction in Section 2.4). We proceed as in Chapter 14 of Adler & Taylor (2007), with an important difference being that here the process f_η does not have constant variance. Aside from the statistical implications of this work that have to do with hypothesis testing, another goal of this paper is to derive analogues of the results in Taylor et al. (2005), Adler & Taylor (2007) for Gaussian processes with nonconstant variance. For each $\eta \in \mathcal{K}$, we define a modified process

$$\tilde{f}_z^\eta = f_z - z^T \alpha_{\eta, X, \Sigma}(\nabla f|_{T_\eta\mathcal{K}}), \quad z \in \mathcal{K},$$

where $\alpha_{\eta, X, \Sigma}(\nabla f|_{T_\eta\mathcal{K}})$ is the vector that, under $H_0 : y \sim N(0, \Sigma)$, computes the expectation of f_z given $\nabla f|_{T_\eta\mathcal{K}}$, the gradient restricted to $T_\eta\mathcal{K}$, i.e.,

$$z^T \alpha_{\eta, X, \Sigma}(\nabla f|_{T_\eta\mathcal{K}}) = \mathbb{E}_0(f_z | \nabla f|_{T_\eta\mathcal{K}}).$$

To check that such a representation is possible, suppose that the tangent space $T_\eta\mathcal{K}$ is j -dimensional, and let $V_\eta \in \mathbb{R}^{p \times j}$ be a matrix whose columns form an

orthonormal basis for $T_\eta \mathcal{K}$. Then $\nabla f|_{T_\eta \mathcal{K}} = V_\eta V_\eta^T X^T y$, and a simple calculation using the properties of conditional expectations for jointly Gaussian random variables shows that

$$\mathbb{E}_0(f_z | \nabla f|_{T_\eta \mathcal{K}}) = z^T X^T P_{\eta, X, \Sigma} y,$$

where

$$P_{\eta, X, \Sigma} = \Sigma X V_\eta (V_\eta^T X^T \Sigma X V_\eta)^\dagger V_\eta^T X^T, \quad (19)$$

the projection onto $X V_\eta$ with respect to Σ (and A^\dagger denoting the Moore-Penrose pseudoinverse of a matrix A). Hence, we gather that

$$\alpha_{\eta, X, \Sigma}(\nabla f|_{T_\eta \mathcal{K}}) = X^T P_{\eta, X, \Sigma} y,$$

and our modified process has the form

$$\tilde{f}_z^\eta = f_z - z^T X^T P_{\eta, X, \Sigma} y = (Xz)^T (I - P_{\eta, X, \Sigma}) y. \quad (20)$$

The key observation, as in Taylor et al. (2005) and Adler & Taylor (2007), is that if η is a critical point, i.e., $\nabla f|_{T_\eta \mathcal{K}} = 0$, then

$$\tilde{f}_z^\eta = f_z \quad \text{for all } z \in \mathcal{K}. \quad (21)$$

Similar to our construction of $\alpha_{\eta, X, \Sigma}(\nabla f|_{T_\eta \mathcal{K}})$, we define $C_{X, \Sigma}(\eta)$ such that

$$\begin{aligned} \mathbb{E}_0(\tilde{f}_z^\eta | \tilde{f}_\eta^\eta) &= \frac{z^T X^T (I - P_{\eta, X, \Sigma}) \Sigma (I - P_{\eta, X, \Sigma}^T) X \eta}{\eta^T X^T (I - P_{\eta, X, \Sigma}) \Sigma (I - P_{\eta, X, \Sigma}^T) X \eta} \cdot \tilde{f}_\eta^\eta \\ &= z^T C_{X, \Sigma}(\eta) \cdot \tilde{f}_\eta^\eta. \end{aligned} \quad (22)$$

and after making three subsequent definitions,

$$\mathcal{V}_\eta^- = \max_{z \in \mathcal{K}: z^T C_{X, \Sigma}(\eta) < 1} \frac{\tilde{f}_z^\eta - z^T C_{X, \Sigma}(\eta) \cdot \tilde{f}_\eta^\eta}{1 - z^T C_{X, \Sigma}(\eta)}, \quad (23)$$

$$\mathcal{V}_\eta^+ = \min_{z \in \mathcal{K}: z^T C_{X, \Sigma}(\eta) > 1} \frac{\tilde{f}_z^\eta - z^T C_{X, \Sigma}(\eta) \cdot \tilde{f}_\eta^\eta}{1 - z^T C_{X, \Sigma}(\eta)}, \quad (24)$$

$$\mathcal{V}_\eta^0 = \max_{z \in \mathcal{K}: z^T C_{X, \Sigma}(\eta) = 1} \tilde{f}_z^\eta - z^T C_{X, \Sigma}(\eta) \cdot \tilde{f}_\eta^\eta, \quad (25)$$

we are ready to state our characterization of the global maximizer η .

Lemma 1. *A point $\eta \in \mathcal{K}$ maximizes f_η over a convex set \mathcal{K} if and only if the following conditions hold:*

$$\nabla f|_{T_\eta \mathcal{K}} = 0, \quad \tilde{f}_\eta^\eta \geq \mathcal{V}_\eta^-, \quad \tilde{f}_\eta^\eta \leq \mathcal{V}_\eta^+, \quad \text{and} \quad \mathcal{V}_\eta^0 \leq 0. \quad (26)$$

The same equivalence holds true even when \mathcal{K} is only locally convex.

Proof. In the forward direction (\Rightarrow), note that $\nabla f|_{T_\eta \mathcal{K}} = 0$ implies that we can replace \tilde{f}_z^η by f_z (and \tilde{f}_η^η by f_η) in the definitions (24), (23), (25), by the key observation (21). As each $z \in \mathcal{K}$ is covered by one of the cases $C_{X,\Sigma}(\eta) < 1$, $C_{X,\Sigma}(\eta) > 1$, $C_{X,\Sigma}(\eta) = 1$, we conclude that

$$f_\eta \geq f_z \quad \text{for all } z \in \mathcal{K},$$

i.e., the point η is a global maximizer.

As for the reverse direction (\Leftarrow), when η is the global maximizer of f_η over \mathcal{K} , the first condition $\nabla f|_{T_\eta \mathcal{K}} = 0$ is clearly true (provided that \mathcal{K} is convex or locally convex), and the other three conditions follow from simple manipulations of the inequalities

$$f_\eta \geq f_z \quad \text{for all } z \in \mathcal{K}.$$

□

Remark 1. *The above lemma does not assume that \mathcal{K} decomposes into strata, or that f_η is Morse for almost all w , or that $y \sim N(X\beta_0, \Sigma)$. It only assumes that \mathcal{K} is convex or locally convex, and its conclusion is completely deterministic, depending only on the process f_η via its covariance function under the null, i.e., via the terms $P_{\eta,X,\Sigma}$ and $C_{X,\Sigma}(\eta)$.*

We note that, under the assumption that f_η is Morse over \mathcal{K} for almost every $y \in \mathbb{R}^p$, and \mathcal{K} is convex, Lemma 1 gives necessary and sufficient conditions for a point $\eta \in \mathcal{K}$ to be the almost sure unique global maximizer. Hence, for convex \mathcal{K} , the conditions in (26) are equivalent to the usual subgradient conditions for optimality, which may be written as

$$\nabla f_\eta \in N_\eta \mathcal{K} \quad \Longleftrightarrow \quad \nabla f|_{T_\eta \mathcal{K}} = 0, \quad \nabla f|_{(T_\eta \mathcal{K})^\perp} \in N_\eta \mathcal{K},$$

where $N_\eta \mathcal{K}$ is the normal cone to \mathcal{K} at η and $\nabla f|_{(T_\eta \mathcal{K})^\perp}$ is the gradient restricted to the orthogonal complement of the tangent space $T_\eta \mathcal{K}$.

Recalling that f_η is a Gaussian process, a helpful independence relationship unfolds.

Lemma 2. *With $y \sim N(X\beta_0, \Sigma)$, for each fixed $\eta \in \mathcal{K}$, the triplet $(\mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0)$ is independent of f_η^η .*

Proof. This is a basic property of conditional expectation for jointly Gaussian variables, i.e., it is easily verified that $\text{Cov}(\tilde{f}_z^\eta - z^T C_{X,\Sigma}(\eta), \tilde{f}_\eta^\eta) = 0$ for all z . □

4. Kac-Rice formulae for the global maximizer and its value

The characterization of the global maximizer from the last section, along with the Kac-Rice formula (Adler & Taylor 2007), allow us to express the joint distribution of

$$\eta^* = \underset{\eta \in \mathcal{K}}{\text{argmax}} f_\eta \quad \text{and} \quad f_{\eta^*} = \max_{\eta \in \mathcal{K}} f_\eta.$$

Theorem 2 (Joint distribution of (η^*, f_{η^*})). *Writing $\mathcal{K} = \cup_{j=0}^p \partial_j \mathcal{K}$ for a stratification of \mathcal{K} , for open sets $A \subseteq \mathbb{R}^p, O \subseteq \mathbb{R}$,*

$$\mathbb{P}(\eta^* \in A, f_{\eta^*} \in O) = \sum_{j=0}^p \int_{\partial_j \mathcal{K} \cap A} \mathbb{E} \left(\det(-\nabla^2 f|_{T_\eta \mathcal{K}}) \cdot \mathbf{1}_{\{\nu_\eta^- \leq \tilde{f}_\eta^\eta \leq \nu_\eta^+, \nu_\eta^0 \leq 0, \tilde{f}_\eta^\eta \in O\}} \left| \nabla f|_{T_\eta \mathcal{K}} = 0 \right. \right) \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \mathcal{H}_j(d\eta), \quad (27)$$

where:

- $\psi_{\nabla f|_{T_\eta \mathcal{K}}}$ is the density of the gradient in some basis for the tangent space $T_\eta \mathcal{K}$, orthonormal with respect to the standard inner product on $T_\eta \mathcal{K}$, i.e., the standard Euclidean Riemannian metric on \mathbb{R}^p ;
- the measure \mathcal{H}_j is the Hausdorff measure induced by the above Riemannian metric on each $\partial_j \mathcal{K}$;
- the Hessian $\nabla^2 f|_{T_\eta \mathcal{K}}$ is evaluated in this orthonormal basis and, for $j = 0$, we take as convention the determinant of a 0×0 matrix to be 1 [in Adler & Taylor (2007), this was denoted by $\nabla^2 f|_{\partial_j \mathcal{K}, \eta}$, to emphasize that it is the Hessian of the restriction of f to $\partial_j \mathcal{K}$].

Proof. This is the Kac-Rice formula, or the “meta-theorem” of Chapter 10 of Adler & Taylor (2007) [see also Azais & Wschebor (2008), Brillinger (1972)], applied to the problem of counting the number of global maximizers in some set $A \subseteq \mathbb{R}^p$ having value in $O \subseteq \mathbb{R}$. That is,

$$\begin{aligned} \mathbb{P}(\eta^* \in A, f_{\eta^*} \in O) &= \mathbb{E} \left(\# \left\{ \eta \in \mathcal{K} \cap A : \nabla f|_{T_\eta \mathcal{K}} = 0, \right. \right. \\ &\quad \left. \left. \nu_\eta^- \leq \tilde{f}_\eta^\eta \leq \nu_\eta^+, \nu_\eta^0 \leq 0, f_\eta \in O \right\} \right) \\ &= \mathbb{E} \left(\# \left\{ \eta \in \mathcal{K} \cap A : \nabla f|_{T_\eta \mathcal{K}} = 0, \right. \right. \\ &\quad \left. \left. \nu_\eta^- \leq \tilde{f}_\eta^\eta \leq \nu_\eta^+, \nu_\eta^0 \leq 0, \tilde{f}_\eta^\eta \in O \right\} \right), \end{aligned}$$

where the second equality follows from (21). Breaking down \mathcal{K} into its separate strata, and then using the Kac-Rice formula, we obtain the result in (27). \square

Remark 2. *As before, the conclusion of Theorem 2 does not actually depend on the convexity of \mathcal{K} . When \mathcal{K} is only locally convex, the Kac-Rice formula [i.e., the right-hand side in (27)] counts the expected total number of global maximizers of f_η lying in some set $A \subseteq \mathbb{R}^p$, with the achieved maximum value in $O \subseteq \mathbb{R}$. For convex \mathcal{K} , our Morse condition on f_η implies an almost surely unique maximizer, and hence the notation $\mathbb{P}(\eta^* \in A, f_{\eta^*} \in O)$ on the left-hand side of (27) makes sense as written. For locally convex \mathcal{K} , one simply needs to interpret the left-hand side as*

$$\mathbb{P} \left(\eta^* \in A \text{ for some } f_{\eta^*} = \max_{\eta \in \mathcal{K}} f_\eta, \max_{\eta \in \mathcal{K}} f_\eta \in O \right).$$

Remark 3. When f_η has constant variance, the distribution of the maximum value f_{η^*} can be approximated extremely well by the expected Euler characteristic (Adler & Taylor 2007) of the excursion set $f_\eta^{-1}(t, \infty) \cap \mathcal{K}$,

$$\sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E} \left(\det(-\nabla^2 f|_{T_\eta \mathcal{K}}) 1_{\{f_\eta > t, \nabla f|_{(T_\eta \mathcal{K})^\perp} \in N_\eta \mathcal{K}\}} \mid \nabla f|_{T_\eta \mathcal{K}} = 0 \right) \cdot \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \mathcal{H}_j(d\eta).$$

This approximation is exact when \mathcal{K} is convex (Takemura & Kuriki 2002), since the Euler characteristic of the excursion set is equal to the indicator that it is not empty.

4.1. Decomposition of the Hessian

In looking at the formula (27), we note that the quantities $\tilde{f}_\eta^\eta, \mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0$ inside the indicator are all independent, by construction, of $\nabla f|_{T_\eta \mathcal{K}}$. It will be useful to decompose the Hessian term similarly. We write

$$-\nabla^2 f|_{T_\eta \mathcal{K}} = -H_\eta + G_\eta \cdot \tilde{f}_\eta^\eta + R_\eta,$$

where

$$R_\eta = -\mathbb{E}_0(\nabla^2 f|_{T_\eta \mathcal{K}} \mid \nabla f|_{T_\eta \mathcal{K}}), \quad (28)$$

$$G_\eta \cdot \tilde{f}_\eta^\eta = -\mathbb{E}_0(\nabla^2 f|_{T_\eta \mathcal{K}} \mid \tilde{f}_\eta^\eta), \quad (29)$$

$$H_\eta = -(\nabla^2 f|_{T_\eta \mathcal{K}} - R_\eta) - G_\eta \cdot \tilde{f}_\eta^\eta. \quad (30)$$

At a critical point of $f|_{\partial_j \mathcal{K}}$, notice that $R_\eta = 0$ (being a linear function of the gradient $\nabla f|_{T_\eta \mathcal{K}}$, which is zero at such a critical point). Furthermore, the pair of matrices $(G_\eta \tilde{f}_\eta^\eta, H_\eta)$ is independent of $\nabla f|_{T_\eta \mathcal{K}}$. Hence, we can rewrite our key formulae for the distribution of the maximizer and its value.

Lemma 3. For each fixed $\eta \in \mathcal{K}$, we have

$$\tilde{f}_\eta^\eta \sim N(\mu_\eta, \sigma_\eta^2),$$

independently of $(\mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0, H_\eta)$, with

$$\mu_\eta = \eta^T X(I - P_{\eta, X, \Sigma})X\beta_0, \quad (31)$$

$$\sigma_\eta^2 = \eta^T X^T(I - P_{\eta, X, \Sigma})\Sigma(I - P_{\eta, X, \Sigma}^T)X\eta, \quad (32)$$

and (recall) $P_{\eta, X, \Sigma} = \Sigma X V_\eta (V_\eta^T X^T \Sigma X V_\eta)^\dagger V_\eta^T X^T$, for an orthonormal basis V_η of $T_\eta \mathcal{K}$.

Moreover, the formula (27) can be equivalently expressed as:

$$\begin{aligned} \mathbb{P}(\eta^* \in A, f_{\eta^*} \in O) &= \sum_{j=0}^p \int_{\partial_j \mathcal{K} \cap A} \mathbb{E} \left(\det(-H_\eta + G_\eta \tilde{f}_\eta^\eta) \cdot \right. \\ &\quad \left. 1_{\{\mathcal{V}_\eta^- \leq \tilde{f}_\eta^\eta \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0, \tilde{f}_\eta^\eta \in O\}} \right) \psi_{\nabla f_{|\mathcal{T}_\eta \mathcal{K}}}(0) \mathcal{H}_j(d\eta) \quad (33) \\ &= \sum_{j=0}^p \int_{\partial_j \mathcal{K} \cap A} \mathbb{E} \left(\mathbb{M}_{\Lambda_\eta, \mathcal{V}_\eta, \mu_\eta, \sigma_\eta^2}(1_O) 1_{\{\mathcal{V}_\eta^- \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0\}} \right) \cdot \\ &\quad \psi_{\nabla f_{|\mathcal{T}_\eta \mathcal{K}}}(0) \det(G_\eta) \mathcal{H}_j(d\eta), \quad (34) \end{aligned}$$

where

$$\mathbb{M}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(h) = \int_{\mathcal{V}^-}^{\mathcal{V}^+} h(z) \det(\Lambda + zI) \frac{e^{-(z-\mu)^2/2\sigma^2}}{\sqrt{2\pi\sigma^2}} dz, \quad (35)$$

and $\Lambda_\eta = G_\eta^{-1} H_\eta$.

Remark 4. Until (34), we had not used the independence of \tilde{f}_η^η and $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0$ (Lemma 2). In (34) we do so, by first integrating over \tilde{f}_η^η (in the definition of \mathbb{M}), and then over $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0$.

4.2. The conditional distribution

The Kac-Rice formula can be generalized further. For a possibly random function h , with $h_{|\partial_j \mathcal{K}}$ continuous for each $j = 0, \dots, p$, we see as a natural extension from (33),

$$\begin{aligned} \mathbb{E}(h(\eta^*)) &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E} \left(h(\eta) \det(-H_\eta + G_\eta \tilde{f}_\eta^\eta) \cdot \right. \\ &\quad \left. 1_{\{\mathcal{V}_\eta^- \leq \tilde{f}_\eta^\eta \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0\}} \right) \psi_{\nabla f_{|\mathcal{T}_\eta \mathcal{K}}}(0) \mathcal{H}_j(d\eta). \quad (36) \end{aligned}$$

This allows us to form a conditional distribution function of sorts. As defined in (35), $\mathbb{M}_{\Lambda, \mathcal{V}, \mu, \sigma^2}$ is not a probability measure, but it can be normalized to yield one:

$$\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g) = \frac{\mathbb{M}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g)}{\mathbb{M}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(1)}. \quad (37)$$

Working form (36), with $h(\eta) = g(f_\eta)$,

$$\begin{aligned} \mathbb{E}(g(f_{\eta^*})) &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E} \left(\mathbb{M}_{\Lambda_\eta, \mathcal{V}_\eta, \mu_\eta, \sigma_\eta^2}(g) 1_{\{\mathcal{V}_\eta^- \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0\}} \right) \\ &\quad \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) \mathcal{H}_j(d\eta) \\ &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E} \left(\mathbb{Q}_{\Lambda_\eta, \mathcal{V}_\eta, \mu_\eta, \sigma_\eta^2}(g) \mathbb{M}_{\Lambda_\eta, \mathcal{V}_\eta, \mu_\eta, \sigma_\eta^2}(1) 1_{\{\mathcal{V}_\eta^- \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0\}} \right) \\ &\quad \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) \mathcal{H}_j(d\eta). \end{aligned} \quad (38)$$

In fact, the quantity $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}$ is a proper conditional distribution. That is, for suitable g

$$\mathbb{E}(g(f_{\eta^*}) \mid \eta^* = \eta, \Lambda_{\eta^*} = \Lambda, \mathcal{V}_{\eta^*} = \mathcal{V}) = \mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g). \quad (39)$$

See Appendix A.4 for details.

4.3. The Kac-Rice pivotal quantity

Suppose that we are interested in testing the null hypothesis $H_0 : y \sim N(0, \Sigma)$. We might look at the observed value of the first knot $\lambda_1 = f_{\eta^*}$, and see if it was larger than we would expect under H_0 . From the results of the last section,

$$\mathbb{P}(f_{\eta^*} > t) = \mathbb{E} \left(\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(1_{(t, \infty)}) \right),$$

and so the most natural strategy seems to be to plug our observed value of the first knot into the above formula. This, however, requires computing the above expectation, i.e., the integral in (38).

In this section, we present an alternative approach that is effectively a conditional test, conditioning on the observed value of η^* , as well as Λ_{η^*} and \mathcal{V}_{η^*} . To motivate our test, it helps to take a step back and think about the measure $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}$ defined in (37). For fixed values of $\Lambda, \mathcal{V}, \mu, \sigma^2$, we can reexpress this (nonrandom) measure as

$$\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g) = \int_{-\infty}^{\infty} g(t) \cdot q_{\Lambda, \mathcal{V}, \mu, \sigma}(t) dt,$$

where $q_{\Lambda, \mathcal{V}, \mu, \sigma}$ is a density function (supported on $[\mathcal{V}^-, \mathcal{V}^+]$). In other words, $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g)$ computes the expectation of g with respect to a density $q_{\Lambda, \mathcal{V}, \mu, \sigma^2}$, so we can write $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(g) = \mathbb{E}(g(W))$ where W is a random variable whose density is $q_{\Lambda, \mathcal{V}, \mu, \sigma^2}$. Now consider the survival function

$$\mathbb{S}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(t) = \mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(1_{(t, \infty)}) = \mathbb{P}(W > t).$$

A classic argument shows that $\mathbb{S}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(W) \sim \text{Unif}(0, 1)$. Why is this useful? Well, according to lemma 7 (or, Remark 7 following the lemma), the first knot

$\lambda_1 = f_{\eta^*}$ almost takes the role of W above, except that there is a further level of randomness in η^* , and $\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}$. That is, instead of the expectation of $g(f_{\eta^*})$ being given by $\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(g)$, it is given by $\mathbb{E}(\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(g))$. The key intuition is that the random variable

$$\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(f_{\eta^*}) = \mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(1_{(f_{\eta^*}, \infty)}) \quad (40)$$

should still be uniformly distributed, since this is true conditional on η^* , $\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}$, and unconditionally, the extra level of randomness in η^* , $\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}$ just gets “averaged out” and does not change the distribution. Our next lemma formalizes this intuition, and therefore provides a test for H_0 based on the (random) survival function in (40).

Lemma 4. [Kac-Rice pivot] *The survival function of $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}$, with $\Lambda = \Lambda_{\eta^*}$, $\mathcal{V} = \mathcal{V}_{\eta^*}$, $\mu = \mu_{\eta^*}$, $\sigma^2 = \sigma_{\eta^*}^2$, and evaluated at $t = f_{\eta^*}$, satisfies*

$$\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(f_{\eta^*}) \sim \text{Unif}(0, 1). \quad (41)$$

Proof. Fix some $h : \mathbb{R} \rightarrow \mathbb{R}$. A standard argument shows that (fixing $\Lambda, \mathcal{V}, \mu, \sigma^2$),

$$\mathbb{Q}_{\Lambda, \mathcal{V}, \mu, \sigma^2}(h \circ \mathbb{S}_{\Lambda, \mathcal{V}, \mu, \sigma^2}) = \int_0^1 h(t) dt.$$

Now we compute, applying (38) with g being a composition of functions,

$$\begin{aligned} & \mathbb{E}\left(h\left(\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(f_{\eta^*})\right)\right) \\ &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E}\left(\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}\left(h \circ \mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}\right) \mathbb{M}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(1) \cdot \right. \\ & \quad \left. 1_{\{\mathcal{V}_{\eta^*}^- \leq \mathcal{V}_{\eta^*}^+, \mathcal{V}_{\eta^*}^0 \leq 0\}}\right) \psi_{\nabla f|_{T_{\eta^*} \mathcal{K}}}(0) \det(G_{\eta^*}) \mathcal{H}_j(d\eta) \\ &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E}\left(\left[\int_0^1 h(t) dt\right] \mathbb{M}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(1) 1_{\{\mathcal{V}_{\eta^*}^- \leq \mathcal{V}_{\eta^*}^+, \mathcal{V}_{\eta^*}^0 \leq 0\}}\right) \cdot \\ & \quad \psi_{\nabla f|_{T_{\eta^*} \mathcal{K}}}(0) \det(G_{\eta^*}) \mathcal{H}_j(d\eta) \\ &= \left[\int_0^1 h(t) dt\right] \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \mathbb{E}\left(\mathbb{M}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(1) 1_{\{\mathcal{V}_{\eta^*}^- \leq \mathcal{V}_{\eta^*}^+, \mathcal{V}_{\eta^*}^0 \leq 0\}}\right) \cdot \\ & \quad \psi_{\nabla f|_{T_{\eta^*} \mathcal{K}}}(0) \det(G_{\eta^*}) \mathcal{H}_j(d\eta) \\ &= \int_0^1 h(t) dt. \end{aligned}$$

□

Remark 5. *In particular, under H_0 , we have $\mu_{\eta^*} = 0$, and Lemma 4 shows that*

$$\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, 0, \sigma_{\eta^*}^2}(f_{\eta^*}) \sim \text{Unif}(0, 1).$$

This proves our main result, Theorem 1, noting that the statistic in (17) is just $\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, 0, \sigma_{\eta^*}^2}(f_{\eta^*})$ written out a little more explicitly.

Remark 6. We have used the survival function of f_{η^*} , conditional on η^* being the global maximizer, as well the additional local information $(\Lambda_{\eta^*}, \mathcal{V}_{\eta^*})$. Conditioning on this extra information makes the test very simple to compute, at least in the lasso, group lasso and nuclear norm cases. If we were to marginalize over these quantities, we would have a more powerful test. In general, it seems difficult to analytically marginalize over these quantities, but perhaps Monte Carlo schemes would be feasible. Further, Lemma 4 holds for any μ_{η^*} , implying that this marginalization over $(\Lambda_{\eta^*}, \mathcal{V}_{\eta^*})$ would need access to the unknown μ_{η^*} . Under the global null, $H_0 : \mu_{\eta^*} = 0$, this would not be an issue, though it would cause a problem for the selective inference work described in Section 6.3.

5. Practicalities and examples

We now summarize the steps required to apply the Kac-Rice test in an instance of the regularized regression problem (9). The test procedure is to compute the statistic

$$\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, 0, \sigma_{\eta^*}^2}(\lambda_1) = \frac{\int_{\lambda_1}^{\mathcal{V}_{\eta^*}^+} \det(\Lambda_{\eta^*} + zI) \phi_{\sigma_{\eta^*}^2}(z) dz}{\int_{\mathcal{V}_{\eta^*}^-}^{\mathcal{V}_{\eta^*}^+} \det(\Lambda_{\eta^*} + zI) \phi_{\sigma_{\eta^*}^2}(z) dz}, \quad (42)$$

and compare this against $\text{Unif}(0, 1)$. Recalling that $\Lambda_{\eta} = H_{\eta}^{-1}G_{\eta}$, this leaves us with essentially 6 quantities to be computed— $\lambda_1, \mathcal{V}_{\eta^*}^+, \mathcal{V}_{\eta^*}^-, G_{\eta^*}, H_{\eta^*}, \sigma_{\eta^*}^2$ from (11), (24), (23), (29), (30), (32), respectively—and the above integral to be calculated.

If we know the dual seminorm \mathcal{Q} of the penalty \mathcal{P} in closed form, then the first knot λ_1 can be found explicitly from (11); otherwise, it can be found numerically by solving the (convex) optimization problem

$$\lambda_1 = \max_{\eta \in \mathbb{R}^p} \eta^T X^T (I - P_{XC^\perp}) y \quad \text{subject to } \mathcal{P}(\eta) \leq 1.$$

The remaining quantities, $\mathcal{V}_{\eta^*}^+, \mathcal{V}_{\eta^*}^-, G_{\eta^*}, H_{\eta^*}, \sigma_{\eta^*}^2$, all depend on η^* and on the tangent space $T_{\eta^*}\mathcal{K}$. Again, depending on \mathcal{Q} , the maximizer η^* can either be found in closed form, or numerically by solving the above optimization problem. Once we know the projection operator onto the tangent space $T_{\eta^*}\mathcal{K}$, (32) is an explicit expression for $\sigma_{\eta^*}^2$; furthermore, $\mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+$ are given by two more tractable (convex) optimization problems (which in some cases admit closed form solutions), see Appendix A.6.

The quantities G_{η^*}, H_{η^*} are different, however; even once we know η^* and the tangent space $T_{\eta^*}\mathcal{K}$, finding G_{η^*}, H_{η^*} involves computing the Hessian $\nabla^2 f|_{T_{\eta^*}\mathcal{K}}$, which requires a geometric understanding of the curvature of f around $T_{\eta^*}\mathcal{K}$. That is, G_{η^*}, H_{η^*} cannot be calculated numerically (say, via an optimization

procedure, as with $\lambda_1, \mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+$), and demand a more problem-specific, mathematical focus. For this reason, computation of G_{η^*}, H_{η^*} can end up being an involved process (depending on the problem). In the examples that follow, we do not give derivation details for the Hessian $\nabla^2 f|_{T_{\eta^*}\mathcal{K}}$, but refer the reader to Adler & Taylor (2007) for the appropriate background material.

To see a high-level summary of the quantities involved, see Algorithm 1.

Algorithm 1 Computing the Kac-Rice pivot

- 1: Solve for λ_1 and η^* {note: we use concise notation throughout, see Section 2.4}
 - 2: Form an orthonormal basis V_{η^*} of the tangent space $T_{\eta^*}\mathcal{K}$
 - 3: Compute the projection $P_{\eta^*, X, \Sigma}$ in (19)
 - 4: Evaluate the conditional variance $\sigma_{\eta^*}^2$ and $C_{X, \Sigma}(\eta^*)$ from (32) and (22)
 - 5: **if** $\nabla^2 f|_{T_{\eta^*}\mathcal{K}} = 0$ **then**
 - 6: Let $\Lambda_{\eta^*} = 0$
 - 7: **else if** $\nabla^2 f|_{T_{\eta^*}\mathcal{K}} \neq 0$ **then**
 - 8: Let $\Lambda_{\eta^*} = G_{\eta^*}^{-1} H_{\eta^*}$ from (29) and (30)
 - 9: **end if**
 - 10: Solve the optimization problems (23) and (24), yielding $\mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+$
 - 11: Evaluate the integrals in (42) to obtain $\mathbb{S} = \mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}^-, 0, \sigma_{\eta^*}^2}(\lambda_1)$
-

We now revisit the lasso example, and then consider the group lasso and nuclear norm penalties, the latter yielding applications to principal components and matrix completion. We remark that in the lasso and group lasso cases, the matrix $\Lambda_{\eta^*} = G_{\eta^*}^{-1} H_{\eta^*}$ is zero, simplifying the computations. In contrast, it is nonzero for the nuclear norm case.

Also, it is important to point out that in all three problem cases, we have $\text{span}(C) = \mathbb{R}^p$, so the notational shortcut that we applied in Sections 3 and 4 has no effect (see Section 2.4), and we can use the formulae from these sections as written.

5.1. Example: the lasso (revisited)

For the lasso problem (3), we have $\mathcal{P}(\beta) = \|\beta\|_1$ and $C = \{u : \|u\|_\infty \leq 1\}$, so $\mathcal{Q}(\beta) = \|\beta\|_\infty$ and $\mathcal{K} = C^\circ = \{v : \|v\|_1 \leq 1\}$. Our Morse assumption on the process $f_\eta = \eta^T X^T y$ over \mathcal{K} (which amounts to an assumption on the predictor matrix X) implies that there is a unique index j^* such that

$$\lambda_1 = |X_{j^*}^T y| = \|X^T y\|_\infty = \max_{\|\eta\|_1 \leq 1} \eta^T X^T y.$$

Then in this notation $\eta^* = \text{sign}(X_{j^*}^T y) \cdot e_{j^*}$ (where e_{j^*} is the j^* th standard basis vector), and the normal cone to \mathcal{K} at η^* is

$$N_{\eta^*}\mathcal{K} = \{v \in \mathbb{R}^p : \text{sign}(v_{j^*}) = \text{sign}(X_{j^*}^T y), |v_j| \leq |v_{j^*}| \text{ for all } j \neq j^*\}.$$

Because this is a full-dimensional set, the tangent space to \mathcal{K} at η^* is $T_{\eta^*}\mathcal{K} = (N_{\eta^*}\mathcal{K})^\perp = \{0\}$. This greatly simplifies our survival function test statistic (40)

since all matrices in consideration here are 0×0 and therefore have determinant 1, giving

$$\mathbb{S}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, 0, \sigma_{\eta^*}^2} = \frac{\Phi(\mathcal{V}_{\eta^*}^- / \sigma_{\eta^*}) - \Phi(\lambda_1 / \sigma_{\eta^*})}{\Phi(\mathcal{V}_{\eta^*}^- / \sigma_{\eta^*}) - \Phi(\mathcal{V}_{\eta^*}^+ / \sigma_{\eta^*})}$$

The lower and upper limits $\mathcal{V}_{\eta^*}^+$, $\mathcal{V}_{\eta^*}^-$ are easily computed by solving two linear fractional programs, see Appendix A.6. The variance $\sigma_{\eta^*}^2$ of f_{η^*} is given by (32), and again simplifies because $T_{\eta^*}\mathcal{K}$ is zero dimensional, becoming

$$\sigma_{\eta^*}^2 = (\eta^*)^T X^T \Sigma X \eta^* = X_{j^*}^T \Sigma X_{j^*}.$$

Plugging in this value gives the test statistic as in (6) in Section 1.1. The reader can return to this section for examples and discussion in the lasso case.

5.2. Example: the group lasso

The group lasso (Yuan & Lin 2006) can be viewed as an extension of the lasso for grouped (rather than individual) variable selection. Given a pre-defined collection \mathcal{G} of groups, with $\cup_{g \in \mathcal{G}} g = \{1, \dots, p\}$, the group lasso penalty is defined as

$$\mathcal{P}(\beta) = \sum_{g=1}^G w_g \|\beta_g\|_2,$$

where $\beta_g \in \mathbb{R}^{|g|}$ denotes the subset of components of $\beta \in \mathbb{R}^p$ corresponding to g , and $w_g > 0$ for all $g \in \mathcal{G}$. We note that

$$C = \{u \in \mathbb{R}^p : \|u_g\|_2 \leq w_g, g \in \mathcal{G}\},$$

so the dual of the penalty is

$$\mathcal{Q}(\beta) = \max_{g \in \mathcal{G}} w_g^{-1} \|\beta_g\|_2,$$

and

$$\mathcal{K} = C^\circ = \left\{ v \in \mathbb{R}^p : \sum_{g \in \mathcal{G}} w_g \|v_g\|_2 \leq 1 \right\}.$$

Under the Morse assumption on $f_\eta = \eta^T X^T y$ over \mathcal{K} (again, this corresponds to an assumption about the design matrix X), there is a unique group g^* such that

$$\lambda_1 = w_{g^*}^{-1} \|X_{g^*}^T y\|_2 = \max_{g \in \mathcal{G}} w_g^{-1} \|X_g^T y\|_2 = \max_{\sum_{g \in \mathcal{G}} w_g \|\eta_g\|_2 \leq 1} \eta^T X^T y,$$

where we write $X_g \in \mathbb{R}^{n \times |g|}$ to denote the matrix whose columns are a subset of those of X , corresponding to g . Then the maximizer η^* is given by

$$\eta_g^* = \begin{cases} \frac{X_g^T y}{w_g \|X_g^T y\|_2} & \text{if } g = g^* \\ 0 & \text{otherwise} \end{cases}, \quad \text{for all } g \in \mathcal{G},$$

and the normal cone $N_{\eta^*}\mathcal{K}$ is seen to be

$$N_{\eta^*}\mathcal{K} = \left\{ v \in \mathbb{R}^p : v_{g^*} = c X_{g^*}^T y, \|v_g\|_2/w_g \leq c \|X_{g^*}^T y\|_2/w_{g^*} \right. \\ \left. \text{for all } g \neq g^*, c \geq 0 \right\}.$$

Hence the tangent space $T_{\eta^*}\mathcal{K} = (N_{\eta^*}\mathcal{K})^\perp$ is

$$T_{\eta^*}\mathcal{K} = \left\{ u \in \mathbb{R}^p : u_{g^*}^T X_{g^*}^T y = 0, u_g = 0 \text{ for all } g \neq g^* \right\},$$

which has dimension $r^* - 1$, with $r^* = \text{rank}(X_{g^*})$. An orthonormal basis V_{η^*} for this tangent space is given by padding an orthonormal basis for $(\text{span}(X_{g^*}^T y))^\perp$ with zeros appropriately. From this we can compute the projection operator

$$P_{\eta^*, X, \Sigma} = \Sigma X V_{\eta^*} (V_{\eta^*}^T X^T \Sigma X V_{\eta^*})^\dagger V_{\eta^*}^T X^T,$$

and the variance of f_{η^*} as

$$\sigma_{\eta^*}^2 = \frac{1}{w_{g^*}^2 \|X_{g^*}^T y\|_2^2} y^T X_{g^*} X_{g^*}^T (I - P_{\eta^*, X, \Sigma}) \Sigma X_{g^*} X_{g^*}^T y.$$

The quantities $\mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+$ can be readily computed by solving two convex programs, see Appendix A.6. Finally, we have $H_{\eta^*} = 0$ in the group lasso case, as the special form of curvature matrix of a sphere implies that $G_{\eta^*} \tilde{f}_{\eta^*}^{\eta^*} = -\nabla^2 f|_{T_{\eta^*}\mathcal{K}}$ in (29). This makes $\Lambda_{\eta^*} = G_{\eta^*}^{-1} H_{\eta^*} = 0$, and the test statistic (42) for the group lasso problem becomes

$$\frac{\int_{\lambda_1}^{\mathcal{V}_{\eta^*}^+} z^{r^*-1} \phi_{\sigma_{\eta^*}^2}(z) dz}{\int_{\mathcal{V}_{\eta^*}^-}^{\mathcal{V}_{\eta^*}^+} z^{r^*-1} \phi_{\sigma_{\eta^*}^2}(z) dz} = \frac{\mathbb{P}(\chi_{r^*} \leq \mathcal{V}_{\eta^*}^+/\sigma_{\eta^*}) - \mathbb{P}(\chi_{r^*} \leq \lambda_1/\sigma_{\eta^*})}{\mathbb{P}(\chi_{r^*} \leq \mathcal{V}_{\eta^*}^+/\sigma_{\eta^*}) - \mathbb{P}(\chi_{r^*} \leq \mathcal{V}_{\eta^*}^-/\sigma_{\eta^*})}. \quad (43)$$

In the above, χ_{r^*} denotes a chi distributed random variable with r^* degrees of freedom, and the equality follows from the fact that the missing multiplicative factor in the χ_{r^*} density [namely, $2^{1-r^*}/2/\Gamma(r^*/2)$] is common to the numerator and denominator, and hence cancels.

Figure 2(a) shows the empirical distribution function of a sample of 20,000 p-values from problem instances sampled randomly from a variety of different group lasso setups (all under the global null model $\beta_0 = 0$):

- small case: X is 3×4 , a fixed matrix slightly perturbed by Gaussian noise; there are 2 groups of size 2, one with weight $\sqrt{2}$, the other with weight 0.1;
- fat case: X is $100 \times 10,000$ with features drawn from the compound symmetric Gaussian distribution having correlation 0.5; here are 1000 groups each of size 10, each having weight $\sqrt{10}$;




FIG 2. The left panel shows the empirical distribution function of a sample of 20,000 p -values (43) computed from various group lasso setups, which agrees very closely with the uniform distribution. The right panel shows the empirical distribution functions in three different group lasso setups, each over 10,000 samples, when p -values are instead computed using an $\text{Exp}(1)$ approximation for the covariance test. This approximation ends up being anti-conservative whereas the Kac-Rice test is exact.

- tall case: X is $10,000 \times 100$ with features drawn from the compound symmetric Gaussian distribution having correlation 0.5; there are 1000 groups each of size 10, each having weight $\sqrt{10}$;
- square case: X is 100×100 with features drawn from the compound symmetric Gaussian distribution having correlation 0.5; there are 10 groups each of size 10, each having weight $\sqrt{10}$;
- diabetes case 1: X is 442×10 , the diabetes data set from Efron et al. (2004); there are 4 (arbitrarily created) groups: one of size 4, one of size 2, one of size 3 and one of size 1, with varying weights;
- diabetes case 2: X is 442×10 , the diabetes data set from Efron et al. (2004); there are now 10 groups of size 1 with i.i.d. random weights drawn from $1 + 0.2 \cdot \text{Unif}(0, 1)$ (generated once for the entire simulation);
- nested case 1: X is 100×10 , with two nested groups (the column space for one group of size 2 is contained in that of the other group of size 8) with the weights favoring inclusion of the larger group first;
- nested case 2: X is 100×10 , with two nested groups (the column space for one group of size 2 is contained in that of the other group of size 8) with the weights favoring inclusion of the smaller group first;
- nested case 3: X is 100×12 , with two sets of two nested groups (in each set, the column space for one group of size 2 is contained in that of the other group of size 4) with the weights chosen according to group size;
- nested case 4: X is 100×120 , with twenty sets of two nested groups (in each set, the column space for one group of size 2 is contained in that of the other group of size 4) with the weights chosen according to group size.

As we can see from the plot, the p-values are extremely close to uniform.

In comparison, arguments similar to those given in Lockhart et al. (2014) for the lasso case would suggest that for the group lasso, under the null hypothesis,

$$\frac{\lambda_1(\lambda_1 - \mathcal{V}_{\eta^*}^-)}{\sigma_{\eta^*}^2} \xrightarrow{d} \text{Exp}(1) \quad \text{as } n, p \rightarrow \infty,$$

under some conditions (one of these being that $\mathcal{V}_{\eta^*}^-$ diverges to ∞ fast enough). Figure 2(b) shows the empirical distribution function of 10,000 samples from three of the above scenarios, demonstrating that, while asymptotically reasonable, the $\text{Exp}(1)$ approximation for the covariance test in the group lasso case can be quite anti-conservative in finite samples.

5.3. Example: nuclear norm

In this setting, we treat the coefficients in (9) as a matrix, instead of a vector, denoted by $B \in \mathbb{R}^{n \times p}$. We consider a nuclear norm penalty on B ,

$$\mathcal{P}(B) = \|B\|_* = \text{tr}(D),$$

where D is the diagonal matrix of singular values in the singular value decomposition $B = UDV^T$. Here the dual seminorm is

$$\mathcal{Q}(B) = \|B\|_{\text{op}} = \max(D),$$

the operator norm (or spectral norm) of B , i.e., its maximum singular value. Therefore we have

$$\begin{aligned} C &= \{A : \|A\|_{\text{op}} \leq 1\}, \\ \mathcal{K} = C^\circ &= \{W : \|W\|_* \leq 1\}. \end{aligned}$$

Examples of problems of the form (9) with nuclear norm penalty $\mathcal{P}(B) = \|B\|_*$ can be roughly categorized according to the choice of linear operator $X = X(B)$. For example,

- *principal components analysis*: if $X : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{n \times p}$ is the identity map, then λ_1 is the largest singular value of $y \in \mathbb{R}^{n \times p}$, and moreover, $\mathcal{V}_{\eta^*}^-$ is the second largest singular value of y ;
- *matrix completion*: if $X : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{n \times p}$ zeros out all of the entries of its argument outside some index set $\mathcal{O} \subseteq \{1, \dots, n\} \times \{1, \dots, p\}$, and leaves the entries in \mathcal{O} untouched, then problem (9) is a noisy version of the matrix completion problem (Candes & Recht 2009, Mazumder et al. 2010).
- *reduced rank regression*: if $X : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{m \times p}$ performs matrix multiplication, $X(B) = XB$, and $y \in \mathbb{R}^{m \times p}$, then problem (9) is often referred to as reduced rank regression (Mukherjee et al. 2015).

The first knot in the solution path is given by $\lambda_1 = \|X^T(y)\|_{\text{op}}$, with X^T denoting the adjoint of the linear operator X . Assuming that $X^T(y)$ has singular value decomposition $X^T(y) = UDV^T$ with $D = \text{diag}(d_1, d_2, \dots)$ for $d_1 \geq d_2 \geq \dots$, and that the process $f_\eta = \langle \eta, X^T(y) \rangle$ is Morse over $\eta \in \mathcal{K}$, there is a unique $\eta^* \in \mathcal{K}$ achieving the value λ_1 ,

$$\eta^* = U_1 V_1^T,$$

where U_1, V_1 are the first columns of U, V , respectively. The normal cone $N_{\eta^*}\mathcal{K}$ is

$$N_{\eta^*}\mathcal{K} = \left\{ cU_1V_1^T + c\tilde{U}\tilde{D}\tilde{V}^T : \tilde{U}^T\tilde{U} = I, \tilde{V}^T\tilde{V} = I, \tilde{D} = \text{diag}(\tilde{d}_1, \tilde{d}_2, \dots), \right. \\ \left. U_1^T\tilde{U} = 0, V_1^T\tilde{V} = 0, \max(\tilde{D}) \leq 1, c \geq 0 \right\},$$

and so the tangent space $T_{\eta^*}\mathcal{K} = (N_{\eta^*}\mathcal{K})^\perp$ is

$$T_{\eta^*}\mathcal{K} = \text{span}\left(\left\{U_1V_j^T, j = 2, \dots, p\right\} \cup \left\{U_jV_1^T, j = 2, \dots, n\right\}\right).$$

From this tangent space, the marginal variance $\sigma_{\eta^*}^2$ in (32) can be easily computed. This leaves $\mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+, G_{\eta^*}, H_{\eta^*}$ to be addressed. As always, the quantities $\mathcal{V}_{\eta^*}^-, \mathcal{V}_{\eta^*}^+$ can be determined numerically, as the optimal values of two convex programs, see Appendix A.6, though these problems admit closed form solutions. We direct the reader to Appendix A.5 for details of the computations.




FIG 3. The left panel shows the empirical distribution function of a sample of 20,000 p-values computed over a variety of problem setups that utilize the nuclear norm penalty. The right panel shows the distribution of 20,000 covariance test p-values when an $\text{Exp}(1)$ approximation is used, which shows the exponential approximation to be clearly inappropriate for the nuclear norm setting.

In Figure 3, we plot the empirical distribution function of a sample of 20,000 p-values computed over problem instances that have been randomly sampled from the following scenarios, all employing the nuclear norm penalty (and all under the null model $B_0 = 0$, with B_0 being the underlying coefficient matrix):

- *principal components analysis*: y is $2 \times 2, 3 \times 4, 50 \times 50, 100 \times 20, 30 \times 1000, 30 \times 5, 1000 \times 1000$;
- *matrix completion*: y is 10×5 with 50% of its entries observed at random, 100×30 with 20% of its entries observed at random, 10×5 with a non-random pattern of observed entries, 20×10 with a nonrandom pattern of observed entries, 200×10 with 10% of its entries observed at random;
- *reduced rank regression*: X is 100×10 whose entries are drawn from the compound symmetric Gaussian distribution with correlation 0.5, and y is 100×5 .

As is evident in the plot, the agreement with uniform is excellent (as above, each particular scenario above produces $\text{Unif}(0, 1)$ p -values regardless of how X was chosen, modulo the Morse assumption).

Again, along the lines of the covariance test, we consider approximation of $\lambda_1(\lambda_1 - \mathcal{V}_{\eta^*}^-)/\sigma_{\eta^*}^2$ by an $\text{Exp}(1)$ distribution under the null hypothesis. Figure 3(b) shows that this approximation is quite far off, certainly much more so than in the other examples. Preliminary calculations confirm mathematically that the $\text{Exp}(1)$ distribution is not the right limiting distribution here; we will pursue this in future work.

6. Discussion

We derived an exact (non-asymptotic) p -value for testing a global null hypothesis in a general regularized regression setup. Our test is based on a geometric characterization of regularized regression estimators and the Kac-Rice formula, and has a close relationship to the covariance test for the lasso problem (Lockhart et al. 2014). In fact, the $\text{Exp}(1)$ limiting null distribution of the covariance test can be derived from the formulae given here. These two tests give similar results for lasso problems, but our new test has exact (not asymptotic) error control under the null, with less assumptions on the predictor matrix X .

Another strength of our approach is that it provably extends well beyond the lasso problem; in this paper we examine tests for both the group lasso and nuclear norm regularization problems. Still, the test can be applied outside of these cases too, and is limited only by the difficulty in evaluating the p -value in practice (which relies on geometric quantities to be computed).

We recall that the covariance test for the lasso can be applied at any knot along the solution path, to test if the coefficients of the predictors not yet in the current model are all zero. In other words, it can be used to test more refined null hypotheses, and not only the global null. Our paper only addresses global testing, but since it was written, there has been much follow-up work in developing tests for the null hypotheses visited further down the regularization path, in various problem settings of interest. We discuss these advances later, and first discuss the important issues of non-Gaussian errors, and power.

6.1. Non-Gaussian errors

Throughout, our calculations have rather explicitly used the fact that $X^T y$ is Gaussian distributed. One could potentially appeal to the central limit theorem if the components of $y - X\beta_0$ are i.i.d. from some error distribution (treating X as fixed), though the calculations in this work focus on the extreme values, so the accuracy of the central limit theorem in the tails may be in doubt. We do not address these theoretical issues here, but instead consider simulations under heavy-tailed, skewed noise. In particular, we drew errors according to a t -distribution with 5 degrees of freedom plus an independent centered $\text{Exp}(1)$ distribution. Figure A.8 in the appendix shows that the lasso and group lasso p-values are relatively well-behaved, while the nuclear norm p-values seem to break down.

6.2. Power

An important issue raised by the editors and referees of this paper concerns the power of our proposed tests. As correctly noted by the editors and referees, the Kac-Rice test is not optimal in terms of power, in a number of practical scenarios. For example, consider the Kac-Rice test in the lasso case, i.e., with the ℓ_1 regularizer $\mathcal{P}(\beta) = \|\beta\|_1$, and the associated the global testing problem $H_0 : \beta_0 = 0$. When the design X is orthogonal, it is known that the most powerful test against sparse alternatives is the max test, which rejects based on $\lambda_1 = \|X^T y\|_\infty$; e.g., see Arias-Castro et al. (2011). The Kac-Rice test, on the other hand, rejects based on a monotone function of $\lambda_1 - \lambda_2$, the first two largest absolute entries of $X^T y$; see (7). Both theoretical and empirical calculations show that a test based on the gap $\lambda_1 - \lambda_2$ is not as powerful as a test based solely on λ_1 , though the former still can have substantial power.

More examples along these lines can be drawn up, showing that the Kac-Rice test, while still providing nontrivial power in a wide variety of specific applications, falls short when compared to the most powerful test against a particular class of alternatives. This is not surprising, and we feel, should not be alarming either. If one has a particular global testing setup in mind, with a particular set of alternatives in mind, then of course one should use the most powerful test when available. The strength of the Kac-Rice framework developed in this paper does not lie in its power for global testing problems such as those considered in the lasso, group lasso, or nuclear norm cases discussed in Sections 5.1, 5.2, and 5.3, respectively. Its real strengths are twofold:

- it provides a single unified framework under which we can derive exact tests for global hypotheses, shown to be applicable in a wide range of settings, and in principle at least, applicable in settings that do not currently possess well-established global testing theory;
- it leads to exact tests beyond the global null, of selective hypotheses along the regularization path, which as far as we can tell was not generically possible before this work.

The second point above is further discussed below.

6.3. *Related work on selective inference*

The global Kac-Rice framework proposed here has spawned new work on selective hypothesis testing, in each of the major example areas considered (lasso, group lasso, and nuclear norm problems). Selective inference is a term used to describe hypothesis testing when the hypotheses themselves are generated or *selected* based on observable data. It is a new field with a small but growing literature, and we believe, has an important place in the modern statistical toolkit full of adaptive regression and classification procedures. We refer the reader to Fithian et al. (2014) for a fairly comprehensive theoretical development of this topic.

After the current paper was written, our Kac-Rice framework has led to selective tests—i.e., tests for hypothesis beyond the global null, at later steps in the regularization path—for least angle regression and the lasso (Taylor et al. 2014), grouped stepwise regression (Loftus & Taylor 2014), and principal components analysis (Choi et al. 2014). To the best of our knowledge, these papers are among the first available rigorous inferential tests for the randomly selected hypotheses visited in these complex procedures. In a sense, each of the aforementioned works were consequences of the Kac-Rice framework, obtained by extending the arguments laid out in Sections 5.1, 5.2, and 5.3, respectively. We are not aware of other global testing frameworks that can also accommodate such extensions, in general.

6.4. *Gaussian random field theory*

Finally, the construction of the process \tilde{f}_η^η is a contribution to the theory of smooth Gaussian random fields as described in Adler & Taylor (2007). Our construction allows earlier proofs that apply only to (centered) Gaussian random fields of constant variance (i.e., marginally stationary) to smooth random fields with arbitrary variance. Even in the marginally stationary case, the conditional distribution \mathbb{Q} defined in (38) provides a new tool for exact selective inference at critical points of such random fields. We leave this, and many other topics, for future work.

Acknowledgements: We are indebted to Robert Tibshirani, without whose help and suggestions this work would not have been possible. Two referees and an associate editor provided feedback on preliminary versions of this work that has helped the authors improve this paper. JT was supported by NSF grant DMS-1208857 and AFOSR grant 113039. RT was supported by NSF grant DMS-1309174.

References

- Adler, R. J. & Taylor, J. E. (2007), *Random fields and geometry*, Springer Monographs in Mathematics, Springer, New York.
- Arias-Castro, E., Candes, E. & Plan, Y. (2011), ‘Global testing under sparse alternatives: ANOVA, multiple comparisons and the higher criticism’, *The Annals of Statistics* **39**(5), 2533–2556.
- Azaïs, J.-M. & Wschebor, M. (2008), ‘A general expression for the distribution of the maximum of a Gaussian field and the approximation of the tail’, *Stochastic Processes and their Applications* **118**(7), 1190–1218.
- Boyd, S., Parikh, N., Chu, E., Peleato, B. & Eckstein, J. (2011), ‘Distributed optimization and statistical learning via the alternative direction method of multipliers’, *Foundations and Trends in Machine Learning* **3**(1), 1–122.
- Boyd, S. & Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press, Cambridge.
- Brillinger, D. R. (1972), ‘On the number of solutions of systems of random equations’, *Annals of Mathematical Statistics* **43**(2), 534–540.
- Candes, E. & Recht, B. (2009), ‘Exact matrix completion via convex optimization’, *Foundations of Computational Mathematics* **9**(6), 717–772.
- Candes, E., Sing-Long, C. & Trzasko, J. (2013), ‘Unbiased risk estimates for singular value thresholding and spectral estimators’, *IEEE Transactions on Signal Processing* **61**(3), 4643–4657.
- Choi, Y., Taylor, J. & Tibshirani, R. (2014), Selecting the number of principal components: estimation of the true rank of a noisy matrix. arXiv: 1410.8260.
- Efron, B., Hastie, T., Johnstone, I. & Tibshirani, R. (2004), ‘Least angle regression’, *Annals of Statistics* **32**(2), 407–499.
- Federer, H. (1959), ‘Curvature measures’, *Transactions of the American Mathematical Society* **93**(3), 418–491.
- Fithian, W., Sun, D. & Taylor, J. (2014), Optimal inference after model selection. arXiv: 1410.2597.
- Lockhart, R., Taylor, J. E., Tibshirani, R. J. & Tibshirani, R. (2014), ‘A significance test for the lasso’, *Annals of Statistics* **42**(2), 413–468.
- Loftus, J. & Taylor, J. (2014), A significance test for forward stepwise model selection. arXiv: 1405.3920.
- Mazumder, R., Hastie, T. & Tibshirani, R. (2010), ‘Spectral regularization algorithms for learning large incomplete matrices’, *Journal of machine learning research* **11**, 2287–2322.
- Mukherjee, A., Chen, K., Wang, N. & Zhu, J. (2015), ‘On the degrees of freedom of reduced-rank estimators in multivariate regression’, *Biometrika* **102**(2), 457–477.
- Takemura, A. & Kuriki, S. (1997), ‘Weights of χ^2 distribution for smooth or piecewise smooth cone alternatives’, *Annals of Statistics* **25**(6), 2368–2387.
- Takemura, A. & Kuriki, S. (2002), ‘On the equivalence of the tube and Euler characteristic methods for the distribution of the maximum of Gaussian fields over piecewise smooth domains’, *Annals of Applied Probability* **12**(2), 768–796.

- Taylor, J. E. (2013), ‘The geometry of least squares in the 21st century’, *Bernoulli* **19**(4), 1449–1464.
- Taylor, J. E., Takemura, A. & Adler, R. J. (2005), ‘Validity of the expected Euler characteristic heuristic’, *Annals of Probability* **33**(4), 1362–1396.
- Taylor, J., Lockhart, R., Tibshirani, R. J. & Tibshirani, R. (2014), Exact post-selection inference for sequential regression procedures. arXiv: 1401.3889.
- 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. E. (2011), ‘The solution path of the generalized lasso’, *The Annals of Statistics* **39**(3).
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J. & Knight, K. (2005), ‘Sparsity and smoothness via the fused lasso’, *Journal of the Royal Statistical Society: Series B* **67**(1), 91–108.
- Yuan, M. & Lin, Y. (2006), ‘Model selection and estimation in regression with grouped variables’, *Journal of the Royal Statistical Society: Series B* **68**(1), 49–67.

Appendix A: Proofs and supplementary details

A.1. Uniqueness of the fitted values and penalty

Consider the strictly convex function $f(u) = \|y - u\|_2^2/2$. Suppose $\hat{\beta}_1, \hat{\beta}_2$ are two solutions of the regularized problem (9) with $X\hat{\beta}_1 \neq X\hat{\beta}_2$. Note

$$f(X\hat{\beta}_1) + \lambda\mathcal{P}(\hat{\beta}_1) = f(X\hat{\beta}_2) + \lambda\mathcal{P}(\hat{\beta}_2) = c^*,$$

the minimum possible value of the criterion in (9). Define $z = \alpha\hat{\beta}_1 + (1 - \alpha)\hat{\beta}_2$ for any $0 < \alpha < 1$. Then by strict convexity of f and convexity of \mathcal{P} ,

$$\begin{aligned} f(Xz) + \lambda\mathcal{P}(z) &< \alpha f(X\hat{\beta}_1) + (1 - \alpha)f(X\hat{\beta}_2) + \lambda(\alpha\mathcal{P}(\hat{\beta}_1) + (1 - \alpha)\mathcal{P}(\hat{\beta}_2)) \\ &= (1 - \alpha)c^* + \alpha c^* \\ &= c^*, \end{aligned}$$

contradicting the fact that c^* is the optimal value of the criterion. Therefore we conclude that $X\hat{\beta}_1 = X\hat{\beta}_2$.

Furthermore, the uniqueness of the fitted value $X\hat{\beta}$ in (9) implies uniqueness of the penalty term $\mathcal{P}(\hat{\beta})$, for any $\lambda > 0$.

A.2. Calculation of λ_1

By the subgradient optimality conditions, $\hat{\beta}$ is a solution in (9) if and only if

$$\frac{X^T(y - X\hat{\beta})}{\lambda} \in \partial\mathcal{P}(\hat{\beta}),$$

where $\partial\mathcal{P}(\hat{\beta})$ denotes the subdifferential (set of subgradients) of \mathcal{P} evaluated at $\hat{\beta}$. Recalling that $\mathcal{P}(\beta) = \max_{u \in C} u^T \beta$, we can rewrite this as

$$\frac{X^T(y - X\hat{\beta})}{\lambda} \in C \quad \text{and} \quad \frac{(X\hat{\beta})^T(y - X\hat{\beta})}{\lambda} = \mathcal{P}(\hat{\beta}). \quad (44)$$

The first statement above is equivalent to $\mathcal{Q}(X^T(y - X\hat{\beta})) \leq \lambda$, where $\mathcal{Q}(\beta) = \max_{v \in C^\circ} v^T \beta$ and C° is the polar body of C . Now define

$$\lambda_1 = \mathcal{Q}(X^T(I - P_{XC^\perp})y),$$

where P_{XC^\perp} denotes projection onto the linear subspace

$$X\{v : v \perp C\} = XC^\perp.$$

Then for any $\lambda \geq \lambda_1$, the conditions in (44) are satisfied by taking $X\hat{\beta} = P_{XC^\perp}y$, and accordingly, $\mathcal{P}(\hat{\beta}) = 0$.

On the other hand, if $\lambda < \lambda_1$, then we must have $\mathcal{P}(\hat{\beta}) \neq 0$, because $\mathcal{P}(\hat{\beta}) = 0$ implies that $X\hat{\beta} \in XC^\perp$, in fact $X\hat{\beta} = P_{XC^\perp}y$, so $\lambda < \mathcal{Q}(X^T(I - P_{XC^\perp})y)$ and the first condition in (44) is not satisfied.

A.3. Bounded support function

We first establish a technical lemma.

Lemma 5. *Suppose that $D \subseteq \mathbb{R}^n$ has $0 \in \text{relint}(D)$. Then for all $u \in \text{span}(D)$,*

$$\max_{v \in D^\circ} v^T u \leq M,$$

for some constant $M < \infty$, with $D^\circ = \{v : u^T v \leq 1 \text{ for all } u \in D\}$ the polar body of D . In other words, $D^\circ \cap \text{span}(D)$ is a bounded set.

Proof. By assumption, we know that $P_D B(r) \subseteq D$ for some $r > 0$, where P_D is the projection matrix onto $\text{span}(D)$, and $B(r)$ is the ℓ_2 ball of radius r centered at 0. Fix $\delta > 0$, and define $S_{D,r^{-1}+\delta} = \partial P_D B(r^{-1} + \delta)$, the boundary of the projected ball of radius $r^{-1} + \delta$. Then for any $v \in S_{D,r^{-1}+\delta} + D^\perp$, we can define $u = r P_D v / \|P_D v\|_2$, and we have $u \in P_D B(r) \subseteq D$ with $u^T v = 1 + \delta r > 1$, hence $v \notin D^\circ$. As this holds for all $\delta > 0$, we must have

$$D^\circ \subseteq \overline{\mathbb{R}^p \setminus \left(\bigcup_{\delta > 0} S_{D,r^{-1}+\delta} + D^\perp \right)} = P_D B(r^{-1}) + D^\perp.$$

This implies that $P_D D^\circ \subseteq P_D B(r^{-1})$, which gives the result. \square

We now use Lemma 5 to show that the process f_η in (16) is bounded over $\eta \in \mathcal{K} = C^\circ$, assuming that C is closed and contains 0 in its relative interior.

Lemma 6. *If $C \in \mathbb{R}^p$ is a closed, convex set with $0 \in \text{relint}(C)$, then*

$$\max_{\eta \in \mathcal{K}} f_\eta \leq M,$$

where f_η is the process in (16), $\mathcal{K} = C^\circ$, and $M < \infty$ is a constant.

Proof. We reparametrize $f_\eta = (X\eta)^T (I - P_{XC^\perp})y$ as

$$g_v = v^T (I - P_{XC^\perp})y, \quad v \in X\mathcal{K},$$

and apply Lemma 5 to the process g_v , with $D = (X^T)^{-1}C$, the inverse image of C under the linear map X^T . It is straightforward to check, using the fact that C is closed, convex, and contains 0 (i.e., using $C^{\circ\circ} = C$), that $D^\circ = XC^\circ = X\mathcal{K}$. By Lemma 5, therefore, we know that $g_v \leq M$ for all $v \in \text{span}(D)$. The proof is completed by noting that $\text{span}(D) = \text{span}((X^T)^{-1}C) = (XC^\perp)^\perp$, which means $(I - P_{XC^\perp})y \in \text{span}(D)$ for any $y \in \mathbb{R}^n$. \square

A.4. The conditional distribution of the global maximum

In Section 4.2, we remarked that the random measure $\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}$ can be viewed as a conditional distribution. We make this precise here.

Lemma 7. *Formally, the measure $\mathbb{Q}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}$ is a conditional distribution function of f_{η^*} , in the sense that on each stratum $\partial_j \mathcal{K}$, $j = 0, \dots, p$, we have*

$$\mathbb{E}(g(f_{\eta^*}) \mid \eta^* = \eta, \Lambda_{\eta^*} = \Lambda, \mathcal{V}_{\eta^*} = \mathcal{V}) = \mathbb{Q}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}(g), \quad (45)$$

for all $\eta \in \partial_j \mathcal{K}$.

Proof. By expanding the right-hand side in (38), we see

$$\begin{aligned} \mathbb{E}(g(f_{\eta^*})) &= \sum_{j=0}^p \int_{\partial_j \mathcal{K}} \int_{\mathbb{R}^j \times \mathbb{R}^3} \mathbb{Q}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}(g) \mathbb{M}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}(1) 1_{\{\mathcal{V}^- \leq \mathcal{V}^+, \mathcal{V}^0 \leq 0\}} \cdot \\ &\quad \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) F_{\Lambda_\eta, \mathcal{V}_\eta}(d\Lambda, d\mathcal{V}) \mathcal{H}_j(d\eta) \end{aligned}$$

where $F_{\Lambda_\eta, \mathcal{V}_\eta}$ denotes the joint distribution of $(\Lambda_\eta, \mathcal{V}_\eta) = (\Lambda_\eta, \mathcal{V}_\eta^-, \mathcal{V}_\eta^+, \mathcal{V}_\eta^0)$ at a fixed value of η . Consider the quantity

$$\mathbb{M}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}(1) 1_{\{\mathcal{V}^- \leq \mathcal{V}^+, \mathcal{V}^0 \leq 0\}} \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) F_{\Lambda_\eta, \mathcal{V}_\eta}(d\Lambda, d\mathcal{V}) \mathcal{H}_j(d\eta). \quad (46)$$

We claim that this is the joint density (modulo differential terms) of η^* , Λ_{η^*} , \mathcal{V}_{η^*} , when η^* is restricted to the smooth piece $\partial_j \mathcal{K}$. This would complete the proof. Hence to verify the claim, we apply Kac-Rice formula with open sets $A \subseteq \partial_j \mathcal{K}$, $B \subseteq \mathbb{R}^j \times \mathbb{R}^3$, and $C \subseteq \mathbb{R}^3$, giving

$$\begin{aligned} \mathbb{P}(\eta^* \in A, \Lambda_{\eta^*} \in B, \mathcal{V}_{\eta^*} \in C) &= \int_A \mathbb{E} \left(\mathbb{M}_{\Lambda_\eta, \mathcal{V}_\eta, \mu_\eta, \sigma_\eta^2}(1) 1_{\{\mathcal{V}_\eta^- \leq \mathcal{V}_\eta^+, \mathcal{V}_\eta^0 \leq 0\}} \cdot \right. \\ &\quad \left. 1_{\{\Lambda \in B, \mathcal{V} \in C\}} \right) \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) \mathcal{H}_j(d\eta) \\ &= \int_A \int_{B \times C} \mathbb{M}_{\Lambda, \mathcal{V}, \mu_\eta, \sigma_\eta^2}(1) 1_{\{\mathcal{V}^- \leq \mathcal{V}^+, \mathcal{V}^0 \leq 0\}} \cdot \\ &\quad \psi_{\nabla f|_{T_\eta \mathcal{K}}}(0) \det(G_\eta) F_{\Lambda_\eta, \mathcal{V}_\eta}(d\Lambda, d\mathcal{V}) \mathcal{H}_j(d\eta). \end{aligned}$$

Taking A, B, C to be open balls around some fixed points $\eta, \Lambda, \mathcal{V}$, respectively, and sending their radii to zero, we see that the joint density of η^* , Λ_{η^*} , \mathcal{V}_{η^*} , with η^* restricted to $\partial_j \mathcal{K}$, is exactly as in (46), as desired. \square

Remark 7. *The analogous result also holds unconditionally, i.e., it is clear that*

$$\mathbb{E}(g(f_{\eta^*})) = \mathbb{E} \left(\mathbb{Q}_{\Lambda_{\eta^*}, \mathcal{V}_{\eta^*}, \mu_{\eta^*}, \sigma_{\eta^*}^2}(g) \right),$$

by taking an expectation on both sides of (39).

A.5. Computations for nuclear norm case

We now discuss computation of the $(n+p-2) \times (n+p-2)$ matrices G_{η^*}, H_{η^*} , and refer the reader to Candes et al. (2013), Takemura & Kuriki (2002, 1997)

for some of the calculations that follow. The entries of G_{η^*} are given by

$$\begin{aligned} G_{\eta^*}(U_1 V_i^T, U_j V_1^T) &= G_{\eta^*}(U_j V_1^T, U_1 V_j^T) \\ &= \text{tr}(V_j U_i^T C_{X,\Sigma}(\eta^*)), \\ G_{\eta^*}(U_1 V_i^T, U_1 V_j^T) &= G_{\eta^*}(U_i V_1^T, U_j V_1^T) \\ &= \delta_{ij} \cdot \text{tr}(V_1 U_1^T C_{X,\Sigma}(\eta^*)), \end{aligned}$$

with $C_{X,\Sigma}(\eta^*)$ as in (22), which has a similar computational form to $\sigma_{\eta^*}^2$, and can be computed from knowledge of $T_{\eta^*}\mathcal{K}$ above. The entries of H_{η^*} are given by

$$\begin{aligned} H_{\eta^*}(U_1 V_i^T, U_j V_1^T) &= H_{\eta^*}(U_i V_1^T, U_1 V_j^T) \\ &= \text{tr}(V_i U_j^T X^T(y)) - d_1 \cdot G_{\eta^*}(U_j V_1^T, U_1 V_i^T) \\ &= \delta_{ij} \cdot d_i - d_1 \cdot G_{\eta^*}(U_j V_1^T, U_1 V_i^T), \\ H_{\eta^*}(U_1 V_i^T, U_1 V_j^T) &= H_{\eta^*}(U_i V_1^T, U_j V_i^T) \\ &= \delta_{ij} [d_1 - \text{tr}(V_1 U_1^T C_{X,\Sigma}(\eta^*))]. \end{aligned}$$

When expressed in a suitable ordering of the above basis of the tangent space, the matrix form of the above expressions are, abbreviating $C = C_{X,\Sigma}(\eta^*)$,

$$\begin{aligned} G_{\eta^*} &= \begin{pmatrix} U_1^T C V_1 \cdot I_{(n-1) \times (n-1)} & U_{-1}^T C V_{-1} \\ V_{-1}^T C^T U_{-1} & U_1^T C V_1 \cdot I_{(p-1) \times (p-1)} \end{pmatrix}, \\ H_{\eta^*} &= \begin{pmatrix} (d_1 - U_1^T C V_1) I_{(n-1) \times (n-1)} & D_{-1} - d_1 U_{-1}^T C V_{-1} \\ D_{-1}^T - d_1 V_{-1}^T C^T U_{-1} & (d_1 - U_1^T C V_1) I_{(p-1) \times (p-1)} \end{pmatrix}, \end{aligned}$$

where U_{-1} denotes all but the first column of U , with V_{-1} similarly defined, and D_{-1} denotes the matrix $\text{diag}(d_2, d_3, \dots)$ with zeros added below in such a way that D_{-1} is an $(n-1) \times (p-1)$ matrix. That is, if $r = \text{rank}(X^T(y))$,

$$D_{-1} = \begin{pmatrix} \text{diag}(d_2, \dots, d_r) \\ 0_{(n-r) \times (p-1)} \end{pmatrix}.$$

In the special case of principal components analysis, in which X is the identity map on $\mathbb{R}^{n \times p}$, the above expressions simplify to $G_{\eta^*} = I_{(n+p-2) \times (n+p-2)}$ and

$$H_{\eta^*} = \begin{pmatrix} 0 & D_{-1} \\ D_{-1}^T & 0 \end{pmatrix},$$

and the $(n+p-2)$ eigenvalues of $\Lambda_{\eta^*} = G_{\eta^*}^{-1} H_{\eta^*}$ are seen to be

$$\{\pm d_j, j = 2, \dots, r\} \cup [0] \cdot (n+p-2r).$$

A.6. Finding \mathcal{V}_{η}^- , \mathcal{V}_{η}^+

Earlier, we noted that convexity of \mathcal{K} was not needed for much of the work in this paper. It is only when we wish to actually compute the test statistic that

we restrict to convex \mathcal{K} . This section describes convex optimization problems that can be solved to find the values of $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+$ which are one of the key stumbling pieces to computing the test statistic. Nevertheless, if we had access to a computer that could compute these quantities, along with Λ_η , then we would not need to restrict interest to convex \mathcal{K} . We would still be assuming positive reach as this is an important assumption to ensure the modified processes \tilde{f}_η^η are not singular. See Chapter 14 of Adler & Taylor (2007) for the case when \mathcal{K} is a smooth locally convex subset of the Euclidean sphere.

Inspecting Lemma 4, we see that in order to compute the p -value in (17), we must find $\mathcal{V}_\eta^-, \mathcal{V}_\eta^+$ at $\eta = \eta^*$. Recall the definition of \mathcal{V}_η^- in (24),

$$\mathcal{V}_\eta^- = \max_{z \in \mathcal{K}: z^T C_{X,\Sigma}(\eta) < 1} \frac{\tilde{f}_z^\eta - z^T C_{X,\Sigma}(\eta) \cdot \tilde{f}_\eta^\eta}{1 - z^T C_{X,\Sigma}(\eta)},$$

where $C_{X,\Sigma}(\eta)$ is defined as in (22), and can be expressed more concisely as

$$z^T C_{X,\Sigma}(\eta) = \frac{z^T X^T (I - P_{\eta,X,\Sigma}) \Sigma X \eta}{\eta^T X^T (I - P_{\eta,X,\Sigma}) \Sigma X \eta}.$$

Recall that at $\eta = \eta^*$ we have $\tilde{f}_z^\eta = f_z$ for all z , and $\tilde{f}_\eta^\eta = \lambda_1$, so

$$\tilde{f}_z^{\eta^*} - z^T C_{X,\Sigma}(\eta^*) \cdot \tilde{f}_{\eta^*}^{\eta^*} = z^T (X^T y - C_{X,\Sigma}(\eta^*) \cdot \lambda_1).$$

Therefore, $\mathcal{V}_{\eta^*}^-$ is given by the maximization problem

$$\begin{aligned} \mathcal{V}_{\eta^*}^- &= \max_{z \in \mathbb{R}^p} \frac{z^T (X^T y - C_{X,\Sigma}(\eta^*) \lambda_1)}{1 - z^T C_{X,\Sigma}(\eta^*)} \\ &\text{subject to } z \in \mathcal{K}, 1 - z^T C_{X,\Sigma}(\eta^*) > 0. \end{aligned} \quad (47)$$

This is a generalized *linear-fractional* problem (Boyd & Vandenberghe 2004). (We say “generalized” here because the constraint $z \in \mathcal{K} \iff \mathcal{P}(z) \leq 1$ can be seen as an infinite number of linear constraints, while the standard linear-fractional setup features a finite number of linear constraints.) Though not convex, problem (47) is quasilinear (i.e., both quasiconvex and quasiconcave), and further, it is equivalent to a convex problem. Specifically,

$$\begin{aligned} \mathcal{V}_{\eta^*}^- &= \max_{u \in \mathbb{R}^p, w \in \mathbb{R}} u^T (X^T y - C_{X,\Sigma}(\eta^*) \lambda_1) \\ &\text{subject to } \mathcal{P}(u) \leq w, w - u^T C_{X,\Sigma}(\eta^*) = 1, w \geq 0. \end{aligned} \quad (48)$$

The proof of equivalence between (47) and (48) follows closely Section 4.3.2 of Boyd & Vandenberghe (2004), and so we omit it here. Similarly, $\mathcal{V}_{\eta^*}^+$ is given by the convex minimization problem

$$\begin{aligned} \mathcal{V}_{\eta^*}^+ &= \min_{u \in \mathbb{R}^p, w \in \mathbb{R}} u^T (X^T y - C_{X,\Sigma}(\eta^*) \lambda_1) \\ &\text{subject to } \mathcal{P}(u) \leq w, w - u^T C_{X,\Sigma}(\eta^*) = -1, w \geq 0. \end{aligned} \quad (49)$$

In the worst case, $\mathcal{V}_{\eta^*}^-$ and $\mathcal{V}_{\eta^*}^+$ can be determined numerically by solving the problems (48) and (49). Depending on the penalty \mathcal{P} , one may favor a particular convex optimization routine over another for this task, but a general purpose solver like ADMM (Boyd et al. 2011) should be suitable for a wide variety of penalties. In several special cases involving the lasso, group lasso, and nuclear norm penalties, $\mathcal{V}_{\eta^*}^-$ and $\mathcal{V}_{\eta^*}^+$ have closed-form expressions. We state these next, but in the interest of space, withhold derivation details. We leave more precise calculations to future work.

A.6.1. Lasso and group lasso

The lasso problem is a special case of the group lasso where all groups have size one; therefore the formulae derived here for the group lasso also apply to the lasso. (In the lasso case, actually, the angles ψ^\pm below are always $\pm\pi$.)

For any group $g \in \mathcal{G}$, let $\theta(\eta, g)$ be the angle between $X_g^T y - C_{X, \Sigma}(\eta)_g$ and $C_{X, \Sigma}(\eta)_g$ [where recall that $C_{X, \Sigma}(\eta)$ is as defined in (22)], and define the angles $\psi^\pm(\eta, g)$ by

$$\sin \psi^\pm(\eta, g) = \frac{\|C_{X, \Sigma}(\eta)_g\|_2}{w_g} \cdot \sin \theta(\eta, g).$$

Define the quantities

$$w^\pm(\eta, g) = \frac{\|X_g^T y - C_{X, \Sigma}(\eta)_g\|_2 \cdot \cos \psi^\pm(\eta, g)}{w_g - \|C_{X, \Sigma}(\eta)_g\|_2 \cdot \cos(\theta(\eta, g) - \psi^\pm(\eta, g))},$$

and

$$v^+(\eta, g) = \begin{cases} \min\{w^\pm(\eta, g)\} & \text{if } \|C_{X, \Sigma}(\eta)_g\|_2 \geq w_g \\ \max\{w^\pm(\eta, g)\} & \text{otherwise} \end{cases},$$

$$v^-(\eta, g) = \begin{cases} \max\{w^\pm(\eta, g)\} & \text{if } \|C_{X, \Sigma}(\eta)_g\|_2 \geq w_g \\ \infty & \text{otherwise} \end{cases}.$$

Then $\mathcal{V}_{\eta^*}^-$, $\mathcal{V}_{\eta^*}^+$ have the explicitly computable form

$$\mathcal{V}_{\eta^*}^- = \max_{g \neq g^*} v^+(\eta^*, g),$$

$$\mathcal{V}_{\eta^*}^+ = \min_{g \neq g^*} v^-(\eta^*, g).$$

A.6.2. Nuclear norm

For principal components analysis, when $X = I$, it is not difficult to show that $\mathcal{V}_{\eta^*}^- = d_2 = -\min(\Lambda_{\eta^*})$ and $\mathcal{V}_{\eta^*}^+ = \infty$. Numerically, this seems to be true even for an arbitrary X (reduced rank regression) or arbitrary patterns of missingness (matrix completion). We remark that computing $\mathcal{V}_{\eta^*}^- = d_2$ requires solving an eigenvalue problem of size at least $\max(n, p) \times \max(n, p)$. (The ADMM approach

for solving problem (48) is no better, as each iteration involves projecting onto the nuclear norm epigraph, which requires an eigendecomposition.) This is quite an expensive computation, and a fast approximation is desirable. We leave this for future work.

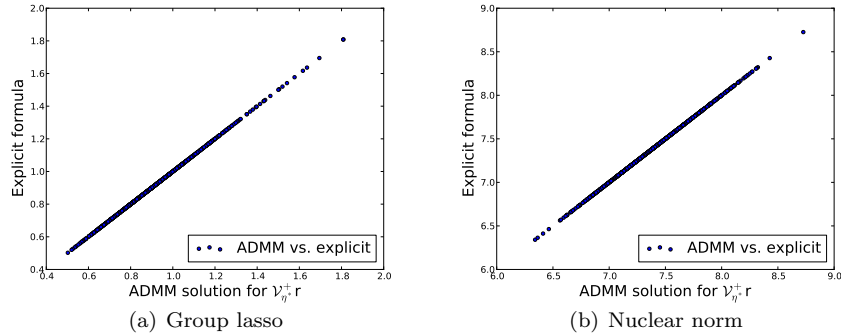


FIG 4. Comparison of the explicit (analytically derived) form of $\mathcal{V}_{\eta^*}^-$ versus the solution of convex program (48) found numerically by ADMM, for the group lasso and nuclear norm penalties.

A.7. Relation to second knot in the solution path

The quantity $\mathcal{V}_{\eta^*}^-$ is always a lower bound on the first knot λ_1 , and possesses an interesting connection to another tuning parameter value of interest along the solution path. In particular, $\mathcal{V}_{\eta^*}^-$ is related to the second knot λ_2 , which can be interpreted more precisely in a problem-specific context, as follows:

- for the lasso, with standardized predictors, $\|X_j\|_2$ for all $j = 1, \dots, p$, we have exactly $\mathcal{V}_{\eta^*}^- = \lambda_2$, the value of the tuning parameter λ at which the solution changes from one nonzero component to two [this follows from a calculation as in Section 4.1 of Lockhart et al. (2014)];
- for the group lasso, as shown in Figure 5(a), $\mathcal{V}_{\eta^*}^-$ is numerically very close to the value λ_2 of the parameter at which the solution changes from one nonzero group of components to two;
- for the nuclear norm penalty, as shown in Figure 5(b), $\mathcal{V}_{\eta^*}^-$ is numerically very close to the value λ_2 of the parameter at which the solution changes from rank 1 to rank 2.

A.8. Non-Gaussian errors

Figure A.8 shows the empirical distribution functions of p-values in the three major example settings of this paper, when the Gaussian noise in observations is replaced by skewed, heavier-tailed noise.




FIG 5. Comparison of $\mathcal{V}_{\eta^*}^-$ with the second knot λ_2 in the solution path, for the group lasso and the nuclear norm problems. For the group lasso, λ_2 is defined to be the largest value of λ such that only one group of variables is active. For the nuclear norm, it is defined as the largest value of λ for which the solution has rank 1.

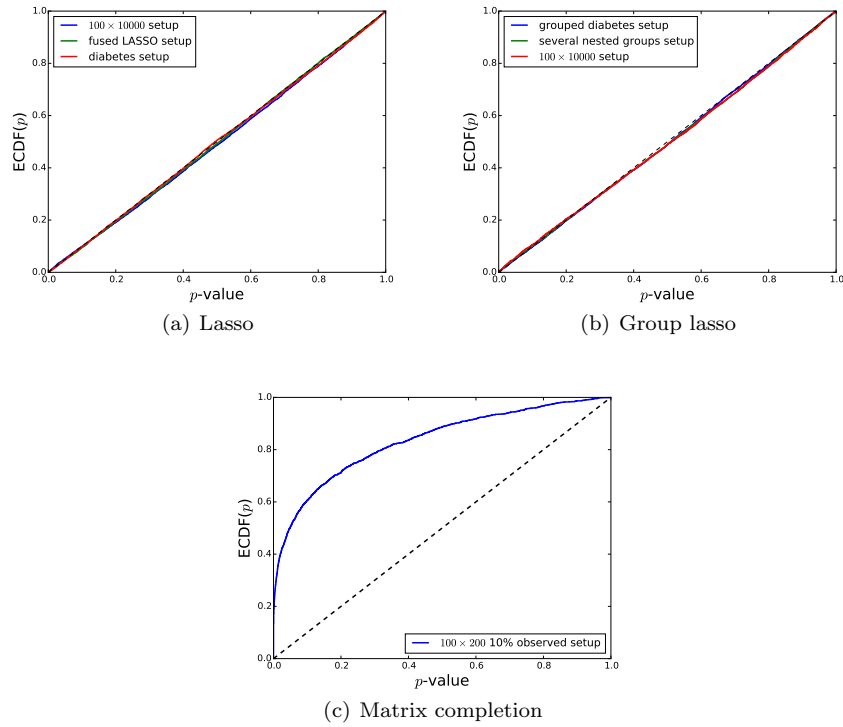


FIG 6. Empirical distribution function of p-values for the lasso, group lasso and matrix completion problems under heavy-tailed and skewed noise: drawn from a t -distribution with 5 degrees of freedom plus an independent centered $\text{Exp}(1)$ distribution.