

High-Dimensional Inference: Confidence Intervals, p -Values and R-Software `hdi`

Ruben Dezeure, Peter Bühlmann, Lukas Meier and Nicolai Meinshausen

Abstract. We present a (selective) review of recent frequentist high-dimensional inference methods for constructing p -values and confidence intervals in linear and generalized linear models. We include a broad, comparative empirical study which complements the viewpoint from statistical methodology and theory. Furthermore, we introduce and illustrate the R-package `hdi` which easily allows the use of different methods and supports reproducibility.

Key words and phrases: Clustering, confidence interval, generalized linear model, high-dimensional statistical inference, linear model, multiple testing, p -value, R-software.

1. INTRODUCTION

Over the last 15 years, a lot of progress has been achieved in high-dimensional statistics where the number of parameters can be much larger than sample size, covering (nearly) optimal point estimation, efficient computation and applications in many different areas; see, for example, the books by Hastie, Tibshirani and Friedman (2009), Bühlmann and van de Geer (2011) or the review article by Fan and Lv (2010). The core task of statistical inference accounting for uncertainty, in terms of frequentist confidence intervals and hypothesis testing, is much less developed. Recently, a few methods for assigning p -values and constructing confidence intervals have been suggested (Wasserman and Roeder, 2009; Meinshausen, Meier and Bühlmann, 2009; Bühlmann, 2013; Zhang and Zhang, 2014; Lockhart et al., 2014; van de Geer et al., 2014; Javanmard and Montanari, 2014; Meinshausen, 2015).

The current paper has three main pillars: (i) a (selective) review of the development in frequentist high-dimensional inference methods for p -values and confidence regions; (ii) presenting the first broad, comparative empirical study among different methods, mainly

for linear models: since the methods are mathematically justified under noncheckable and sometimes non-comparable assumptions, a thorough simulation study should lead to additional insights about reliability and performance of various procedures; (iii) presenting the R-package `hdi` (*high-dimensional inference*) which enables to easily use many of the different methods for inference in high-dimensional generalized linear models. In addition, we include a recent line of methodology allowing to detect significant groups of highly correlated variables which could not be inferred as individually significant single variables (Meinshausen, 2015). The review and exposition in Bühlmann, Kalisch and Meier (2014) is vaguely related to points (i) and (iii) above, but much more focusing on an application oriented viewpoint and covering much less statistical methodology, theory and computational details.

Our comparative study, point (ii), mentioned above, exhibits interesting results indicating that more “stable” procedures based on Ridge-estimation or random sample splitting with subsequent aggregation are somewhat more reliable for type I error control than asymptotically power-optimal methods. Such results cannot be obtained by comparing underlying assumptions of different methods, since these assumptions are often too crude and far from necessary. As expected, we are unable to pinpoint to a method which is (nearly) best in all considered scenarios. In view of this, we also want to offer a collection of useful methods for the

Ruben Dezeure is a Ph.D. student, Peter Bühlmann is Professor, Lukas Meier is Senior Scientist and Nicolai Meinshausen is Professor, Seminar for Statistics, ETH Zürich, CH-8092 Zürich, Switzerland (e-mail: dezeure@stat.math.ethz.ch; buhlmann@stat.math.ethz.ch; meier@stat.math.ethz.ch; meinshausen@stat.math.ethz.ch).

community, in terms of our R-package `hdi` mentioned in point (iii) above.

2. INFERENCE FOR LINEAR MODELS

We consider first a high-dimensional linear model, while extensions are discussed in Section 3:

$$(2.1) \quad Y = \mathbf{X}\beta^0 + \varepsilon,$$

with $n \times p$ fixed or random design matrix \mathbf{X} , $n \times 1$ response and error vectors Y and ε , respectively. The errors are assumed to be independent of \mathbf{X} (for random design) with i.i.d. entries having $\mathbb{E}[\varepsilon_i] = 0$. We allow for high-dimensional settings where $p \gg n$. In further development, the active set or the set of relevant variables

$$S_0 = \{j; \beta_j^0 \neq 0, j = 1, \dots, p\},$$

as well as its cardinality $s_0 = |S_0|$, are important quantities. The main goals of this section are the construction of confidence intervals and p -values for individual regression parameters β_j^0 ($j = 1, \dots, p$) and corresponding multiple testing adjustment. The former is a highly nonstandard problem in high-dimensional settings, while for the latter we can use standard well-known techniques. When considering both goals simultaneously, though, one can develop more powerful multiple testing adjustments. The Lasso (Tibshirani, 1996) is among the most popular procedures for estimating the unknown parameter β^0 in a high-dimensional linear model. It exhibits desirable or sometimes even optimal properties for point estimation such as prediction of $\mathbf{X}\beta^0$ or of a new response Y_{new} , estimation in terms of $\|\hat{\beta} - \beta^0\|_q$ for $q = 1, 2$, and variable selection or screening; see, for example, the book of Bühlmann and van de Geer (2011). For assigning uncertainties in terms of confidence intervals or hypothesis testing, however, the plain Lasso seems inappropriate. It is very difficult to characterize the distribution of the estimator in the high-dimensional setting; Knight and Fu (2000) derive asymptotic results for fixed dimension as sample size $n \rightarrow \infty$ and already for such simple situations, the asymptotic distribution of the Lasso has point mass at zero. This implies, because of noncontinuity of the distribution, that standard bootstrapping and subsampling schemes are delicate to apply and uniform convergence to the limit seems hard to achieve. The latter means that the estimator is exposed to undesirable super-efficiency problems, as illustrated in Section 2.5. All the problems mentioned are expected to apply not only for the Lasso but also for other sparse estimators as well.

In high-dimensional settings and for general fixed design \mathbf{X} , the regression parameter is not identifiable. However, when making some restrictions on the design, one can ensure that the regression vector is identifiable. The so-called compatibility condition on the design \mathbf{X} (van de Geer, 2007) is a rather weak assumption (van de Geer and Bühlmann, 2009) which guarantees identifiability and oracle (near) optimality results for the Lasso. For the sake of completeness, the compatibility condition is described in Appendix A.1.

When assuming the compatibility condition with constant ϕ_0^2 (ϕ_0^2 is close to zero for rather ill-posed designs, and sufficiently larger than zero for well-posed designs), the Lasso has the following property: for Gaussian errors and if $\lambda \asymp \sqrt{\log(p)/n}$, we have with high probability that

$$(2.2) \quad \|\hat{\beta} - \beta^0\|_1 \leq 4s_0\lambda/\phi_0^2.$$

Thus, if $s_0 \ll \sqrt{n/\log(p)}$ and $\phi_0^2 \geq M > 0$, we have $\|\hat{\beta} - \beta^0\|_1 \rightarrow 0$ and, hence, the parameter β^0 is identifiable.

Another often used assumption, although not necessary by any means, is the so-called beta-min assumption:

$$(2.3) \quad \min_{j \in S_0} |\beta_j^0| \geq \beta_{\min},$$

for some choice of constant $\beta_{\min} > 0$. The result in (2.2) immediately implies the screening property: if $\beta_{\min} > 4s_0\lambda/\phi_0^2$, then

$$(2.4) \quad \hat{S} = \{j; \hat{\beta}_j \neq 0\} \supseteq S_0.$$

Thus, the screening property holds when assuming the compatibility and beta-min condition. The power of the screening property is a massive dimensionality reduction (in the original variables) because $|\hat{S}| \leq \min(n, p)$; thus, if $p \gg n$, the selected set \hat{S} is much smaller than the full set of p variables. Unfortunately, the required conditions are overly restrictive and exact variable screening seems rather unrealistic in practical applications (Bühlmann and Mandozzi, 2014).

2.1 Different Methods

We describe here three different methods for construction of statistical hypothesis tests or confidence intervals. Alternative procedures are presented in Sections 2.3 and 2.5.

2.1.1 *Multi sample-splitting.* A generic way for deriving p -values in hypotheses testing is given by splitting the sample with indices $\{1, \dots, n\}$ into two equal halves denoted by I_1 and I_2 , that is, $I_r \subset \{1, \dots, n\}$ ($r = 1, 2$) with $|I_1| = \lfloor n/2 \rfloor$, $|I_2| = n - \lfloor n/2 \rfloor$, $I_1 \cap I_2 = \emptyset$ and $I_1 \cup I_2 = \{1, \dots, n\}$. The idea is to use the first half I_1 for variable selection and the second half I_2 with the reduced set of selected variables (from I_1) for statistical inference in terms of p -values. Such a sample-splitting procedure avoids the over-optimism to use the data twice for selection and inference after selection (without taking the effect of selection into account).

Consider a method for variable selection based on the first half of the sample:

$$\hat{S}(I_1) \subset \{1, \dots, p\}.$$

A prime example is the Lasso which selects all the variables whose corresponding estimated regression coefficients are different from zero. We then use the second half of the sample I_2 for constructing p -values, based on the selected variables $\hat{S}(I_1)$. If the cardinality $|\hat{S}(I_1)| \leq n/2 \leq |I_2|$, we can run ordinary least squares estimation using the subsample I_2 and the selected variables $\hat{S}(I_1)$, that is, we regress Y_{I_2} on $\mathbf{X}_{I_2}^{(\hat{S}(I_1))}$ where the sub-indices denote the sample half and the super-index stands for the selected variables, respectively. Thereby, we implicitly assume that the matrix $\mathbf{X}_{I_2}^{(\hat{S}(I_1))}$ has full rank $|\hat{S}(I_1)|$. Thus, from such a procedure, we obtain p -values $P_{t\text{-test},j}$ for testing $H_{0,j} : \beta_j^0 = 0$, for $j \in \hat{S}(I_1)$, from the classical t -tests, assuming Gaussian errors or relying on asymptotic justification by the central limit theorem. To be more precise, we define (raw) p -values

$$P_{\text{raw},j} = \begin{cases} P_{t\text{-test},j} \text{ based on } Y_{I_2}, \mathbf{X}_{I_2}^{(\hat{S}(I_1))}, & \text{if } j \in \hat{S}(I_1), \\ 1, & \text{if } j \notin \hat{S}(I_1). \end{cases}$$

An interesting feature of such a sample-splitting procedure is the adjustment for multiple testing. For example, if we wish to control the familywise error rate over all considered hypotheses $H_{0,j}$ ($j = 1, \dots, p$), a naive approach would employ a Bonferroni–Holm correction over the p tests. This is not necessary: we only need to control over the considered $|\hat{S}(I_1)|$ tests in I_2 . Therefore, a Bonferroni corrected p -value for $H_{0,j}$ is given by

$$P_{\text{corr},j} = \min(P_{\text{raw},j} \cdot |\hat{S}(I_1)|, 1).$$

In high-dimensional scenarios, $p \gg n > \lfloor n/2 \rfloor \geq |\hat{S}(I_1)|$, where the latter inequality is an implicit assumption which holds for the Lasso (under weak assumptions), and thus, the correction factor employed here is rather small. Such corrected p -values control the familywise error rate in multiple testing when assuming the screening property in (2.4) for the selector $\hat{S} = \hat{S}(I_1)$ based on the first half I_1 only, exactly as stated in Fact 1 below. The reason is that the screening property ensures that the reduced model is a correct model, and hence the result is not surprising. In practice, the screening property typically does not hold exactly, but it is not a necessary condition for constructing valid p -values (Bühlmann and Mandozzi, 2014).

The idea about sample-splitting and subsequent statistical inference is implicitly contained in Wasserman and Roeder (2009). We summarize the whole procedure as follows:

Single sample-splitting for multiple testing of $H_{0,j}$ among $j = 1, \dots, p$:

1. Split (partition) the sample $\{1, \dots, n\} = I_1 \cup I_2$ with $I_1 \cap I_2 = \emptyset$ and $|I_1| = \lfloor n/2 \rfloor$ and $|I_2| = n - \lfloor n/2 \rfloor$.
2. Using I_1 only, select the variables $\hat{S} \subseteq \{1, \dots, p\}$. Assume or enforce that $|\hat{S}| \leq |I_1| = \lfloor n/2 \rfloor \leq |I_2|$.
3. Denote the design matrix with the selected set of variables by $\mathbf{X}^{(\hat{S})}$. Based on I_2 with data $(Y_{I_2}, \mathbf{X}_{I_2}^{(\hat{S})})$, compute p -values $P_{\text{raw},j}$ for $H_{0,j}$, for $j \in \hat{S}$, from classical least squares estimation [i.e., t -test which can be used since $|\hat{S}(I_1)| \leq |I_2|$]. For $j \notin \hat{S}$, assign $P_{\text{raw},j} = 1$.
4. Correct the p -values for multiple testing: consider

$$P_{\text{corr},j} = \min(P_{\text{raw},j} \cdot |\hat{S}|, 1),$$

which is an adjusted p -value for $H_{0,j}$ for controlling the familywise error rate.

A major problem of the single sample-splitting method is its sensitivity with respect to the choice of splitting the entire sample: sample splits lead to wildly different p -values. We call this undesirable phenomenon a p -value lottery, and Figure 1 provides an illustration. To overcome the “ p -value lottery,” we can run the sample-splitting method B times, with B large. Thus, we obtain a collection of p -values for the j th hypothesis $H_{0,j}$:

$$P_{\text{corr},j}^{[1]}, \dots, P_{\text{corr},j}^{[B]} \quad (j = 1, \dots, p).$$

The task is now to do an aggregation to a single p -value. Because of dependence among $\{P_{\text{corr},j}^{[b]}; b =$

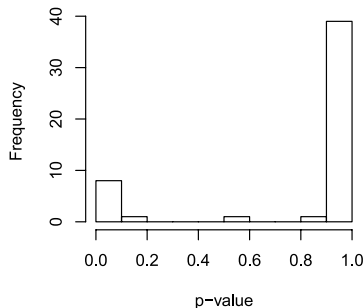


FIG. 1. Histogram of p -values $P_{\text{corr},j}$ for a single covariable, in the riboflavin data set, when doing 50 different (random) sample splits. The figure is taken from Bühlmann, Kalisch and Meier (2014).

$1, \dots, B\}$, because all the different half samples are part of the same full sample, an appropriate aggregation needs to be developed. A simple solution is to use an empirical γ -quantile with $0 < \gamma < 1$:

$$Q_j(\gamma) = \min(\text{emp. } \gamma\text{-quantile}\{P_{\text{corr},j}^{[b]}/\gamma; b = 1, \dots, B\}, 1).$$

For example, with $\gamma = 1/2$, this amounts to taking the sample median $\{P_{\text{corr},j}^{[b]}; b = 1, \dots, B\}$ and multiplying it with the factor 2. A bit more sophisticated approach is to choose the best and properly scaled γ -quantile in the range $(\gamma_{\min}, 1)$ (e.g., $\gamma_{\min} = 0.05$), leading to the aggregated p -value

$$P_j = \min\left((1 - \log(\gamma_{\min})) \inf_{\gamma \in (\gamma_{\min}, 1)} Q_j(\gamma)\right) \quad (j = 1, \dots, p).$$

Thereby, the factor $(1 - \log(\gamma_{\min}))$ is the price to be paid for searching for the best $\gamma \in (\gamma_{\min}, 1)$. This Multi sample-splitting procedure has been proposed and analyzed in Meinshausen, Meier and Bühlmann (2009), and we summarize it below. Before doing so, we remark that the aggregation of dependent p -values as described above is a general principle as described in Appendix A.1.

Multi sample-splitting for multiple testing of $H_{0,j}$ among $j = 1, \dots, p$:

1. Apply the single sample-splitting procedure B times, leading to p -values $\{P_{\text{corr},j}^{[b]}; b = 1, \dots, B\}$. Typical choices are $B = 50$ or $B = 100$.

2. Aggregate these p -values as in (2.5), leading to P_j which are adjusted p -values for $H_{0,j}$ ($j = 1, \dots, p$), controlling the familywise error rate.

The Multi sample-splitting method enjoys the property that the resulting p -values are approximately reproducible and not subject to a “ p -value lottery” anymore, and it controls the familywise error rate under the following assumptions:

(A1) The screening property as in (2.4) for the first half of the sample: $\mathbb{P}[\hat{S}(I_1) \supseteq S_0] \geq 1 - \delta$ for some $0 < \delta < 1$.

(A2) The reduced design matrix for the second half of the sample satisfies $\text{rank}(\mathbf{X}_{I_2}^{\hat{S}(I_1)}) = |\hat{S}(I_1)|$.

FACT 1 [Meinshausen, Meier and Bühlmann (2009)]. Consider a linear model as in (2.1) with fixed design \mathbf{X} and Gaussian errors. Assume (A1)–(A2). Then, for a significance level $0 < \alpha < 1$ and denoting by B the number of sample splits,

$$\mathbb{P}\left[\bigcup_{j \in S_0^c} I(P_j \leq \alpha)\right] \leq \alpha + B\delta,$$

that is, the familywise error rate (FWER) is controlled up to the additional (small) value $B\delta$.

A proof is given in Meinshausen, Meier and Bühlmann (2009). We note that the Multi sample-splitting method can be used in conjunction with any reasonable, sparse variable screening method fulfilling (A1) for very small $\delta > 0$ and (A2); and it does not necessarily rely on the Lasso for variable screening. See also Section 2.1.6. Assumption (A2) typically holds for the Lasso satisfying $|\hat{S}(I_1)| \leq |I_1| = \lfloor n/2 \rfloor \leq |I_2| = n - \lfloor n/2 \rfloor$.

The screening property (A1). The screening property (A1) with very small $\delta > 0$ is not a necessary condition for constructing valid p -values and can be replaced by a zonal assumption requiring the following: there is a gap between large and small regression coefficients and there are not too many small nonzero regression coefficients (Bühlmann and Mandozzi, 2014). Still, such a zonal assumption makes a requirement about the unknown β^0 and the absolute values of its components: but this is the essence of the question in hypothesis testing to infer whether coefficients are sufficiently different from zero, and one would like to do such a test without an assumption on the true values.

The Lasso satisfies (A1) with $\delta \rightarrow 0$ when assuming the compatibility condition (A.1) on the design \mathbf{X} , the sparsity assumption $s_0 = o(\sqrt{n/\log(p)})$ [or $s_0 = o(n/\log(p))$ when requiring a restricted eigenvalue assumption] and a beta-min condition (2.3), as shown in (2.4). Other procedures also exhibit the screening property such as the adaptive Lasso (Zou, 2006), analyzed

in detail in van de Geer, Bühlmann and Zhou (2011), or methods with concave regularization penalty such as SCAD (Fan and Li, 2001) or MC+ (Zhang, 2010). As criticized above, the required beta-min assumption should be avoided when constructing a hypothesis test about the unknown components of β^0 .

Fact 1 has a corresponding asymptotic formulation where the dimension $p = p_n$ and the model depends on sample size n : if (A1) is replaced by $\lim_{n \rightarrow \infty} \mathbb{P}[\hat{S}(I_{1;n}) \supseteq S_{0;n}] \rightarrow 1$ and for a fixed number B , $\limsup_{n \rightarrow \infty} \mathbb{P}[\bigcup_{j \in S_0^c} I(P_j \leq \alpha)] \leq \alpha$. In such an asymptotic setting, the Gaussian assumption in Fact 1 can be relaxed by invoking the central limit theorem (for the low-dimensional part).

The Multi sample-splitting method is very generic: it can be used for many other models, and its basic assumptions are an approximate screening property (2.4) and that the cardinality $|\hat{S}(I_1)| < |I_2|$ so that we only have to deal with a fairly low-dimensional inference problem. See, for example, Section 3 for GLMs. An extension for testing group hypotheses of the form $H_{0,G} : \beta_j = 0$ for all $j \in G$ is indicated in Section 4.1.

Confidence intervals can be constructed based on the duality with the p -values from equation (2.5). A procedure is described in detail in Appendix A.2. The idea to invert the p -value method is to apply a bisection method having a point in and a point outside of the confidence interval. To verify if a point is inside the aggregated confidence interval, one looks at the fraction of confidence intervals from the splits which cover the point.

2.1.2 Regularized projection: De-sparsifying the Lasso. We describe here a method, first introduced by Zhang and Zhang (2014), which does not require an assumption about β^0 except for sparsity.

It is instructive to give a motivation starting with the low-dimensional setting where $p < n$ and $\text{rank}(\mathbf{X}) = p$. The j th component of the ordinary least squares estimator $\hat{\beta}_{\text{OLS};j}$ can be obtained as follows. Do an OLS regression of $\mathbf{X}^{(j)}$ versus all other variables $\mathbf{X}^{(-j)}$ and denote the corresponding residuals by $Z^{(j)}$. Then

$$(2.6) \quad \hat{\beta}_{\text{OLS};j} = Y^T Z^{(j)} / (\mathbf{X}^{(j)})^T Z^{(j)}$$

can be obtained by a linear projection. In a high-dimensional setting, the residuals $Z^{(j)}$ would be equal to zero and the projection is ill-posed.

For the high-dimensional case with $p > n$, the idea is to pursue a regularized projection. Instead of ordinary least squares regression, we use a Lasso regression of $\mathbf{X}^{(j)}$ versus $\mathbf{X}^{(-j)}$ with corresponding residual vector $Z^{(j)}$: such a penalized regression involves

a regularization parameter λ_j for the Lasso, and hence $Z^{(j)} = Z^{(j)}(\lambda_j)$. As in (2.6), we immediately obtain (for any vector $Z^{(j)}$)

$$(2.7) \quad \frac{Y^T Z^{(j)}}{(\mathbf{X}^{(j)})^T Z^{(j)}} = \beta_j^0 + \sum_{k \neq j} P_{jk} \beta_k^0 + \frac{\varepsilon^T Z^{(j)}}{(\mathbf{X}^{(j)})^T Z^{(j)}},$$

$$P_{jk} = (\mathbf{X}^{(k)})^T Z^{(j)} / (\mathbf{X}^{(j)})^T Z^{(j)}.$$

We note that in the low-dimensional case with $Z^{(j)}$ being the residuals from ordinary least squares, due to orthogonality, $P_{jk} = 0$. When using the Lasso-residuals for $Z^{(j)}$, we do not have exact orthogonality and a bias arises. Thus, we make a bias correction by plugging in the Lasso estimator $\hat{\beta}$ (of the regression Y versus \mathbf{X}): the bias-corrected estimator is

$$(2.8) \quad \hat{b}_j = \frac{Y^T Z^{(j)}}{(\mathbf{X}^{(j)})^T Z^{(j)}} - \sum_{k \neq j} P_{jk} \hat{\beta}_k.$$

Using (2.7), we obtain

$$\sqrt{n}(\hat{b}_j - \beta_j^0) = \frac{n^{-1/2} \varepsilon^T Z^{(j)}}{n^{-1} (\mathbf{X}^{(j)})^T Z^{(j)}} + \sum_{k \neq j} \sqrt{n} P_{jk} (\beta_k^0 - \hat{\beta}_k).$$

The first term on the right-hand side has a Gaussian distribution, when assuming Gaussian errors; otherwise, it has an asymptotic Gaussian distribution assuming that $\mathbb{E}|\varepsilon_j|^{2+\kappa} < \infty$ for $\kappa > 0$ (which suffices for the Lyapunov CLT). We will argue in Appendix A.1 that the second term is negligible under the following assumptions:

(B1) The design matrix \mathbf{X} has compatibility constant bounded away from zero, and the sparsity is $s_0 = o(\sqrt{n}/\log(p))$.

(B2) The rows of \mathbf{X} are fixed realizations of i.i.d. random vectors $\sim \mathcal{N}_p(0, \Sigma)$, and the minimal eigenvalue of Σ is bounded away from zero.

(B3) The inverse Σ^{-1} is row-sparse with $s_j = \sum_{k \neq j} I((\Sigma^{-1})_{jk} \neq 0) = o(n/\log(p))$.

FACT 2 (Zhang and Zhang, 2014; van de Geer et al., 2014). Consider a linear model as in (2.1) with fixed design and Gaussian errors. Assume (B1), (B2) and (B3) (or an ℓ_1 -sparsity assumption on the rows of Σ^{-1}). Then

$$\sqrt{n} \sigma_\varepsilon^{-1} (\hat{b} - \beta^0) = W + \Delta, \quad W \sim \mathcal{N}_p(0, \Omega),$$

$$\Omega_{jk} = \frac{n(Z^{(j)})^T Z^{(k)}}{[(\mathbf{X}^{(j)})^T Z^{(j)}][(\mathbf{X}^{(k)})^T Z^{(k)}]},$$

$$\|\Delta\|_\infty = o_P(1).$$

[We note that this statement holds with probability tending to one, with respect to the variables $\mathbf{X} \sim \mathcal{N}_P(0, \Sigma)$ as assumed in (B2)].

The asymptotic implications of Fact 2 are as follows:

$$\sigma_\varepsilon^{-1} \Omega_{jj}^{-1/2} \sqrt{n}(\hat{b}_j - \beta_j^0) \Rightarrow \mathcal{N}(0, 1),$$

from which we can immediately construct a confidence interval or hypothesis test by plugging in an estimate $\hat{\sigma}_\varepsilon$ as briefly discussed in Section 2.1.4. From a theoretical perspective, it is more elegant to use the square root Lasso (Belloni, Chernozhukov and Wang, 2011) for the construction of $Z^{(j)}$; then one can drop (B3) [or the ℓ_1 -sparsity version of (B3)] (van de Geer, 2014). In fact, all that we then need is formula (2.12)

$$\|\hat{\beta} - \beta^0\|_1 = o_P(1/\sqrt{\log(p)}).$$

From a practical perspective, it seems to make essentially no difference whether one takes the square root or plain Lasso for the construction of the $Z^{(j)}$'s.

More general than the statements in Fact 2, the following holds assuming (B1)–(B3) (van de Geer et al., 2014): the asymptotic variance $\sigma_\varepsilon^2 \Omega_{jj}$ reaches the Cramér–Rao lower bound, which equals $\sigma_\varepsilon^2 (\Sigma^{-1})_{jj}$ [which is bounded away from zero, due to (B2)], and the estimator \hat{b}_j is efficient in the sense of semi-parametric inference. Furthermore, the convergence in Fact 2 is uniform over the subset of the parameter space where the number of nonzero coefficients $\|\beta^0\|_0$ is small and, therefore, we obtain *honest* confidence intervals and tests. In particular, both of these results say that all the complications in post-model selection do not arise (Leeb and Pötscher, 2003), and yet \hat{b}_j is optimal for construction of confidence intervals of a single coefficient β_j^0 .

From a practical perspective, we need to choose the regularization parameters λ (for the Lasso regression of Y versus \mathbf{X}) and λ_j [for the nodewise Lasso regressions (Meinshausen and Bühlmann, 2006) of $\mathbf{X}^{(j)}$ versus all other variables $\mathbf{X}^{(-j)}$]. Regarding the former, we advocate a choice using cross-validation; for the latter, we favor a proposal for a smaller λ_j than the one from CV, and the details are described in Appendix A.1.

Furthermore, for a group $G \subseteq \{1, \dots, p\}$, we can test a group hypothesis $H_{0,G} : \beta_j^0 = 0$ for all $j \in G$ by considering the test-statistic

$$\max_{j \in G} \sigma_\varepsilon^{-1} \Omega_{jj}^{-1/2} \sqrt{n} |\hat{b}_j| \Rightarrow \max_{j \in G} \Omega_{jj}^{-1/2} |W_j|,$$

where the limit on the right-hand side occurs if the null-hypothesis $H_{0,G}$ holds true. The distribution of

$\max_{j \in G} |\Omega_{jj}^{-1/2} W_j|$ can be easily simulated from dependent Gaussian random variables. We also remark that sum-type statistics for large groups cannot be easily treated because $\sum_{j \in G} |\Delta_j|$ might get out of control.

2.1.3 Ridge projection and bias correction. Related to the desparsified Lasso estimator \hat{b} in (2.8) is an approach based on Ridge estimation. We sketch here the main properties and refer to Bühlmann (2013) for a detailed treatment.

Consider

$$\hat{\beta}_{\text{Ridge}} = (n^{-1} \mathbf{X}^T \mathbf{X} + \lambda I)^{-1} n^{-1} \mathbf{X}^T Y.$$

A major source of bias occurring in Ridge estimation when $p > n$ comes from the fact that the Ridge estimator is estimating a projected parameter

$$\theta^0 = P_R \beta^0, \quad P_R = \mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-} \mathbf{X},$$

where $(\mathbf{X} \mathbf{X}^T)^{-}$ denotes a generalized inverse of $\mathbf{X} \mathbf{X}^T$. The minor bias for θ^0 then satisfies

$$\max_j |\mathbb{E}[\hat{\beta}_{\text{Ridge};j}] - \theta_j^0| \leq \lambda \|\theta^0\|_2 \lambda_{\min \neq 0}(\hat{\Sigma})^{-1},$$

where $\lambda_{\min \neq 0}(\hat{\Sigma})$ denotes the minimal nonzero eigenvalue of $\hat{\Sigma}$ (Shao and Deng, 2012). The quantity can be made small by choosing λ small. Therefore, for $\lambda \searrow 0^+$ and assuming Gaussian errors, we have that

$$(2.9) \quad \sigma_\varepsilon^{-1} (\hat{\beta}_{\text{Ridge}} - \theta^0) \approx W, \quad W \sim \mathcal{N}_p(0, \Omega_R),$$

where $\Omega_R = (\hat{\Sigma} + \lambda)^{-1} \hat{\Sigma} (\hat{\Sigma} + \lambda)^{-1} / n$. Since

$$\frac{\theta^0}{P_{R;jj}} = \beta_j^0 + \sum_{k \neq j} \frac{P_{R;jk}}{P_{R;jj}} \beta_k^0,$$

the major bias for β_j^0 can be estimated and corrected with

$$\sum_{k \neq j} \frac{P_{R;jk}}{P_{R;jj}} \hat{\beta}_k,$$

where $\hat{\beta}$ is the ordinary Lasso. Thus, we construct a bias-corrected Ridge estimator, which addresses the potentially substantial difference between θ^0 and the target β^0 :

$$(2.10) \quad \hat{b}_{R;j} = \frac{\hat{\beta}_{\text{Ridge};j}}{P_{R;jj}} - \sum_{k \neq j} \frac{P_{R;jk}}{P_{R;jj}} \hat{\beta}_k,$$

$j = 1, \dots, p$.

Based on (2.9), we derive in Appendix A.1 that

$$\begin{aligned}
 & \sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} (\hat{b}_{R;j} - \beta_j^0) \\
 & \approx \Omega_{R;jj}^{-1/2} W_j / P_{R;jj} \\
 (2.11) \quad & + \sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} \Delta_{R;j}, \quad W \sim \mathcal{N}_p(0, \Omega_R), \\
 & |\Delta_{R;j}| \leq \Delta_{R\text{bound};j} \\
 & := \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right| (\log(p)/n)^{1/2-\xi},
 \end{aligned}$$

with the typical choice $\xi = 0.05$. Sufficient conditions for deriving (2.11) are assumption (B1) and that the sparsity satisfies $s_0 = O((n/\log(p))^\xi)$ for ξ as above.

Unlike as in Fact 2, the term $\Delta_{R;j}$ is typically not negligible and we correct the Gaussian part in (2.11) by the upper bound $\Delta_{R\text{bound};j}$. For example, for testing $H_{0,j} : \beta_j^0 = 0$ we use the upper bound for the p -value

$$2(1 - \Phi(\sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} |P_{R;jj}| (|\hat{b}_{R;j}| - \Delta_{R\text{bound};j}))).$$

Similarly, for two-sided confidence intervals with coverage $1 - \alpha$ we use

$$[\hat{b}_{R;j} - c_j, \hat{b}_{R;j} + c_j],$$

$$c_j = \Delta_{R\text{bound};j} + \sigma_\varepsilon \Omega_{R;jj}^{1/2} / |P_{R;jj}| \Phi^{-1}(1 - \alpha/2).$$

For testing a group hypothesis for $G \subseteq \{1, \dots, p\}$, $H_{0,G} : \beta_j^0 = 0$ for all $j \in G$, we can proceed similarly as at the end of Section 2.1.2: under the null-hypotheses $H_{0,G}$, the statistic $\sigma_\varepsilon^{-1} \max_{j \in G} \Omega_{R;jj}^{-1/2} \cdot |\hat{b}_{R;j}|$ has a distribution which is approximately stochastically upper bounded by

$$\max_{j \in G} (\Omega_{R;jj}^{-1/2} |W_j| / |P_{R;jj}| + \sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} |\Delta_{R;j}|);$$

see also (2.11). When invoking an upper bound for $\Delta_{R\text{bound};j} \geq |\Delta_{R;j}|$ as in (2.11), we can easily simulate this distribution from dependent Gaussian random variables, which in turn can be used to construct a p -value; we refer for further details to Bühlmann (2013).

2.1.4 Additional issues: Estimation of the error variance and multiple testing correction. Unlike the Multi sample-splitting procedure in Section 2.1.1, the desparsified Lasso and Ridge projection method outlined in Sections 2.1.2–2.1.3 require to plug-in an estimate of σ_ε and to adjust for multiple testing. The scaled Lasso (Sun and Zhang, 2012) leads to a consistent estimate of the error variance: it is a fully automatic method which does not need any specification of a tuning parameter. In Reid, Tibshirani and Friedman (2013), an empirical comparison of various estimators suggests that

the estimator based on a residual sum of squares of a cross-validated Lasso solution often yields good finite-sample performance.

Regarding the adjustment when doing many tests for individual regression parameters or groups thereof, one can use any valid standard method to correct the p -values from the desparsified Lasso or Ridge projection method. The prime examples are the Bonferroni–Holm procedure for controlling the familywise error rate and the method from Benjamini and Yekutieli (2001) for controlling the false discovery rate. An approach for familywise error control which explicitly takes the dependence among the multiple hypotheses is proposed in Bühlmann (2013), based on simulations for dependent Gaussian random variables.

2.1.5 Conceptual differences between the methods.

We briefly outline here conceptual differences while Section 2.5 presents empirical results.

The Multi sample-splitting method is very generic and in the spirit of Breiman’s appeal for stability (Breiman, 1996a, 1996b), it enjoys some kind of stability due to multiple sample splits and aggregation; see also the discussion in Sections 2.1.6 and 2.4. The disadvantage is that, in the worst case, the method needs a beta-min or a weaker zonal assumption on the underlying regression parameters: this is somewhat unpleasant since a significance test should *find out* whether a regression coefficient is sufficiently large or not.

Both the desparsified Lasso and Ridge projection procedures do not make any assumption on the underlying regression coefficient except sparsity. The former is most powerful and asymptotically optimal if the design were generated from a population distribution whose inverse covariance matrix is sparse. Furthermore, the convergence is uniform over all sparse regression vectors and, hence, the method yields honest confidence regions or tests. The Ridge projection method does not require any assumption on the fixed design but does not reach the asymptotic Cramér–Rao efficiency bound. The construction with the additional correction term in (A.3) leads to reliable type I error control at the cost of power.

In terms of computation, the Multi sample-splitting and Ridge projection method are substantially less demanding than the desparsified Lasso.

2.1.6 Other sparse methods than the Lasso. All the methods described above are used “in default mode” in conjunction with the Lasso (see also Section 2.2). This is not necessary, and other estimators can be used.

For the Multi sample-splitting procedure, assumptions (A1) with $\delta \rightarrow 0$ and (A2) are sufficient for asymptotic correctness; see Fact 1. These assumptions hold for many reasonable sparse estimators when requiring a beta-min assumption and some sort of identifiability condition such as the restricted eigenvalue or the compatibility condition on the design matrix \mathbf{X} ; see also the discussion after Fact 1. It is unclear whether one could gain substantially by using a different screening method than the Lasso. In fact, the Lasso has been empirically found to perform rather well for screening in comparison to the elastic net (Zou and Hastie, 2005), marginal correlation screening (Fan and Lv, 2008) or thresholded Ridge regression; see Bühlmann and Mandozzi (2014).

For the desparsified Lasso, the error of the estimated bias correction can be controlled by using a bound for $\|\hat{\beta} - \beta^0\|_1$. If we require (B2) and (B3) [or an ℓ_1 sparsity assumption instead of (B3)], the estimation error in the bias correction, based on an estimator $\hat{\beta}$ in (2.8), is asymptotically negligible if

$$(2.12) \quad \|\hat{\beta} - \beta^0\|_1 = o_P(1/\sqrt{\log(p)}).$$

This bound is implied by (B1) and (B2) for the Lasso, but other estimators exhibit this bound as well, as mentioned below. When using such another estimator, the wording “desparsified Lasso” does not make sense anymore. Furthermore, when using the square root Lasso for the construction of $Z^{(j)}$, we only need (2.12) to obtain asymptotic normality with the \sqrt{n} convergence rate (van de Geer, 2014).

For the Ridge projection method, a bound for $\|\hat{\beta} - \beta^0\|_1$ is again the only assumption such that the procedure is asymptotically valid. Thus, for the corresponding bias correction, other methods than the Lasso can be used.

We briefly mention a few other methods for which we have reasons that (A1) with very small $\delta > 0$ and (A2), or the bound in (2.12) hold: the adaptive Lasso (Zou, 2006) analyzed in greater detail in van de Geer, Bühlmann and Zhou (2011), the MC+ procedure with its high-dimensional mathematical analysis (Zhang, 2010), or methods with concave regularization penalty such as SCAD (Fan and Li, 2001) analyzed in broader generality and detail in Fan, Xue and Zou (2014). If the assumptions (A1) with small $\delta > 0$ and (A2) fail for the Multi sample-splitting method, the multiple sample splitting still allows to check the stability of the p -values $P_{\text{corr},j}^{[b]}$ across b (i.e., across sample splits). If the variable screening is unstable, many of the $P_{\text{corr},j}^{[b]}$

(across b) will be equal to 1, therefore, the aggregation has a tendency to produce small p -values if most of them, each from a sample split, are stable and small. See also Mandozzi and Bühlmann (2015), Section 5. In connection with the desparsified method, a failure of the single sufficient condition in (2.12), when using, for example, the square root Lasso for construction of the $Z^{(j)}$'s, might result in a too large bias. In absence of resampling or Multi sample splitting, it seems difficult to diagnose such a failure (of the desparsified or Ridge projection method) with real data.

2.2 hdi for Linear Models

In the R-package `hdi`, available on R-Forge (Meier, Meinshausen and Dezeure, 2014), we provide implementations for the Multi sample-splitting, the Ridge projection and the desparsified Lasso method.

Using the R functions is straightforward:

```
> outMssplit
  <- multi.split(x = x, y = y)
> outRidge
  <- ridge.proj(x = x, y = y)
> outLasso
  <- lasso.proj(x = x, y = y)
```

For users that are very familiar with the procedures, we provide flexible options. For example, we can easily use an alternative model selection or another “classical” fitting procedure using the arguments `model.selector` and `classical.fit` in `multi.split`. The default options should be satisfactory for standard usage.

All procedures return p -values and confidence intervals. The Ridge and desparsified Lasso methods return both single testing p -values as well as multiple testing corrected p -values, unlike the Multi sample-splitting procedure which only returns multiple testing corrected p -values. The confidence intervals are for individual parameters only (corresponding to single hypothesis testing).

The single testing p -values and the multiple testing corrected p -values are extracted from the fit as follows:

```
> outRidge$pval
> outRidge$pval.corr
```

By default, we correct for controlling the familywise error rate for the p -values `pval.corr`.

Confidence intervals are acquired through the usual `confint` interface. Below we extract the 95 % confidence intervals for those p -values that are smaller than 0.05:

```
> confint(outMssplit,
  parm = which(outMssplit$pval.corr
```



```
<= 0.05),
level = 0.95)
```

Due to the fact that the desparsified Lasso method is quite computationally intensive, we provide the option to parallelize the method on a user-specified number of cores.

We refer to the manual of the package for more detailed information.

2.3 Other Methods

Recently, other procedures have been suggested for construction of p -values and confidence intervals.

Residual-type bootstrap approaches are proposed and analyzed in Chatterjee and Lahiri (2013) and Liu and Yu (2013). A problem with these approaches is the nonuniform convergence to a limiting distribution and exposure to the super-efficiency phenomenon, that is, if the true parameter equals zero, a confidence region might be the singleton $\{0\}$ (due to a finite amount of bootstrap resampling), while for nonzero true parameter values, the coverage might be very poor or a big length of the confidence interval.

The covariance test (Lockhart et al., 2014) is another proposal which relies on the solution path of the Lasso and provides p -values for conditional tests that all relevant variables enter the Lasso solution path first. It is related to post-selection inference, mentioned in Section 7.1.

In Javanmard and Montanari (2014), a procedure was proposed that is very similar to the one described in Section 2.1.2, with the only difference being that Z is picked as the solution of a convex program rather than using the Lasso. The method is aiming to relax the sparsity assumption (B3) for the design.

A conservative *Group-bound* method which needs no regularity assumption for the design, for example, no compatibility assumption (A.1), has been proposed by Meinshausen (2015). The method has the capacity to automatically determine whether a regression coefficient is identifiable or not, and this makes the procedure very robust against ill-posed designs. The main motivation of the method is in terms of testing groups of correlated variables, and we discuss it in more detail in Section 4.1.

While all the methods mentioned above are considered in a comparative simulation study in Section 2.5, we mention here some others. The idea of estimating a low-dimensional component of a high-dimensional parameter is also worked out in Belloni et al. (2012), Belloni, Chernozhukov and Kato (2015), bearing connections to the approach of desparsifying

the Lasso. Based on stability selection (Meinshausen and Bühlmann, 2010), Shah and Samworth (2013) propose a version which leads to p -values for testing individual regression parameters. Furthermore, there are new and interesting proposals for controlling the false discovery rate, in a “direct way” (Bogdan et al. 2013, 2014; Barber and Candès, 2015).

2.4 Main Assumptions and Violations

We discuss here some of the main assumptions, potential violations and some corresponding implications calling for caution when aiming for confirmatory conclusions.

Linear model assumption. The first one is that the linear (or some other) model is correct. This might be rather unrealistic and, thus, it is important to interpret the output of software or a certain method. Consider a nonlinear regression model

$$\text{random design : } Y_0 = f^0(X_0) + \eta_0,$$

$$\text{fixed design : } Y = f^0(\mathbf{X}) + \eta,$$

where, with some slight abuse of notation, $f^0(\mathbf{X}) = f^0(\mathbf{X}_1), \dots, (f^0(\mathbf{X}_n))^T$. We assume for the random design model, η_0 is independent from X_0 , $\mathbb{E}[\eta_0] = 0$, $\mathbb{E}[f^0(X_0)] = 0$, $\mathbb{E}[X_0] = 0$, and the data are n i.i.d. realizations of (X_0, Y_0) ; for the fixed design model, the $n \times 1$ random vector η has i.i.d. components with $\mathbb{E}[\eta_i] = 0$. For the random design model, we consider

$$\begin{aligned} Y_0 &= (\beta^0)^T X_0 + \varepsilon_0, \\ (2.13) \quad \varepsilon_0 &= f^0(X_0) - (\beta^0)^T X_0 + \eta_0, \\ \beta^0 &= \operatorname{argmin}_{\beta} \mathbb{E}[(f^0(X_0) - \beta^T X_0)^2] \end{aligned}$$

[where the latter is unique if $\operatorname{Cov}(X_0)$ is positive definite]. We note that $\mathbb{E}[\varepsilon_0|X_0] \neq 0$ while $\mathbb{E}[\varepsilon_0] = 0$ and, therefore, the inference should be *unconditional* on \mathbf{X} and is to be interpreted for the projected parameter β^0 in (2.13). Furthermore, for correct asymptotic inference of the projected parameter β^0 , a modified estimator for the asymptotic variance of the estimator is needed; and then both the Multi sample-splitting and the desparsified Lasso are asymptotically correct (assuming similar conditions as if the model were correct). The Multi sample-splitting method is well suited for the random design case because the sample splitting (resampling type) is coping well with i.i.d. data. This is in contrast to fixed design, where the data is not i.i.d. and

the Multi sample-splitting method for a misspecified linear model is typically not working anymore. The details are given in Bühlmann and van de Geer (2015).

For a fixed design model with $\text{rank}(\mathbf{X}) = n$, we can always write

$$Y = \mathbf{X}\beta^0 + \varepsilon, \quad \varepsilon = \eta$$

for many solutions β^0 . For ensuring that the inference is valid, one should consider a sparse β^0 , for example, the basis pursuit solution from compressed sensing (Candes and Tao, 2006) as one among many solutions. Thus, inference should be interpreted for a *sparse* solution β^0 , in the sense that a confidence interval for the j th component would cover this j th component of all sufficiently sparse solutions β^0 . For the high-dimensional fixed design case, there is no misspecification with respect to linearity of the model; misspecification might happen, though, if there is no solution β^0 which fulfills a required sparsity condition. The details are given again in Bühlmann and van de Geer (2015).

The assumption about constant error variance might not hold. We note that in the random design case of a nonlinear model as above, the error in (2.13) has non-constant variance when conditioning on \mathbf{X} , but, unconditionally, the noise is homoscedastic. Thus, as outlined, the inference for a random design linear model is asymptotically valid (unconditional on \mathbf{X}) even though the conditional error distribution given \mathbf{X} has nonconstant variance.

Compatibility or incoherence-type assumption. The methods in Section 2.1 require an identifiability assumption such as the compatibility condition on the design matrix \mathbf{X} described in (A.1). The procedure in Section 4.1 does not require such an assumption: if a component of the regression parameter is not identifiable, the method will not claim significance. Hence, some robustness against nonidentifiability is offered with such a method.

Sparsity. All the described methods require some sparsity assumption of the parameter vector β^0 [if the model is misspecified, this concerns the parameter β^0 as in (2.13) or the basis pursuit solution]; see the discussion of (A1) after Fact 1 or assumption (B1). Such sparsity assumptions can be somewhat relaxed to require weak sparsity in terms of $\|\beta^0\|_r$ for some $0 < r < 1$, allowing that many or all regression parameters are nonzero but sufficiently small (cf. van de Geer, 2015; Bühlmann and van de Geer, 2015).

When the truth (or the linear approximation of the true model) is nonsparse, the methods are expected to break down. With the Multi sample-splitting procedure, however, a violation of sparsity might be detected, since for nonsparse problems, a sparse variable screening method will be typically unstable with the consequence that the resulting aggregated p -values are typically not small; see also Section 2.1.6.

Finally, we note that for the desparsified Lasso, the sparsity assumption (B3) or its weaker version can be dropped when using the square root Lasso; see the discussion after Fact 2.

Hidden variables. The problem of hidden variables is most prominent in the area of causal inference (cf. Pearl, 2000). In the presence of hidden variables, the presented techniques need to be adapted, adopting ideas from, for example, the framework of EM-type estimation (cf. Dempster, Laird and Rubin, 1977), low-rank methods (cf. Chandrasekaran, Parrilo and Willsky, 2012) or the FCI technique from causal inference (cf. Spirtes, Glymour and Scheines, 2000).

2.5 A Broad Comparison

We compare a variety of methods on the basis of multiple testing corrected p -values and single testing confidence intervals. The methods we look at are the multiple sample-splitting method *MS-Split* (Section 2.1.1), the desparsified Lasso method *Lasso-Pro* (Section 2.1.2), the Ridge projection method *Ridge* (Section 2.1.3), the covariance test *Covtest* (Section 2.3), the method by Javanmard and Montanari *Jm2013* (Section 2.3) and the two bootstrap procedures mentioned in Section 2.3 [*Res-Boot* corresponds to Chatterjee and Lahiri (2013) and *liuyu* to Liu and Yu (2013)].

2.5.1 Specific details for the methods. For the estimation of the error variance, for the Ridge projection or the desparsified Lasso method, the scaled Lasso is used as mentioned in Section 2.1.4.

For the choice of tuning parameters for the nodewise Lasso regressions (discussed in Section 2.1.2), we look at the two alternatives of using either cross-validation or our more favored alternative procedure (denoted by Z&Z) discussed in Appendix A.1.

We do not look at the bootstrap procedures in connection with multiple testing adjustment due to the fact that the required number of bootstrap samples grows out of proportion to go far enough in the tails of the distribution; some additional importance sampling might help to address such issues.

Regarding the covariance test, the procedure does not directly provide p -values for the hypotheses we are interested in. For the sake of comparison though, we use the interpretation as in [Bühlmann, Meier and van de Geer \(2014\)](#).

This interpretation does not have a theoretical reasoning behind it and functions more as a heuristic.

Thus, the results of the covariance test procedure should be interpreted with caution.

For the method *Jm2013*, we used our own implementation instead of the code provided by the authors. The reason for this is that we had already implemented our own version when we discovered that code was available and our own version was (orders of magnitude) better in terms of error control. Posed with the dilemma of fair comparison, we stuck to the best performing alternative.

2.5.2 Data used. For the empirical results, simulated design matrices as well as design matrices from real data are used. The simulated design matrices are generated $\sim \mathcal{N}_p(0, \Sigma)$ with covariance matrix Σ of the following three types:

$$\begin{aligned} \text{Toeplitz:} \quad & \Sigma_{j,k} = 0.9^{|j-k|}, \\ \text{Exp.decay:} \quad & (\Sigma^{-1})_{j,k} = 0.4^{|j-k|/5}, \\ \text{Equi.corr:} \quad & \Sigma_{j,k} \equiv 0.8 \quad \text{for all } j \neq k, \\ & \Sigma_{j,j} \equiv 1 \quad \text{for all } j. \end{aligned}$$

The sample size and dimension are fixed at $n = 100$ and $p = 500$, respectively. We note that the Toeplitz type has a banded inverse Σ^{-1} , and, vice-versa, the Exp.decay type exhibits a banded Σ . The design matrix RealX from real gene expression data of *Bacillus Subtilis* ($n = 71$, $p = 4088$) was kindly provided by DSM (Switzerland) and is publicly available ([Bühlmann, Kalisch and Meier, 2014](#)). To make the problem somewhat comparable in difficulty to the simulated designs, the number of variables is reduced to $p = 500$ by taking the variables with highest empirical variance.

The cardinality of the active set is picked to be one of two levels $s_0 \in \{3, 15\}$.

For each of the active set sizes, we look at 6 different ways of picking the sizes of the nonzero coefficients:

Randomly generated : $U(0, 2), U(0, 4), U(-2, 2)$,

A fixed value : 1, 2 or 10.

The positions of the nonzero coefficients as columns of the design \mathbf{X} are picked at random. Results where the nonzero coefficients were positioned to be the first

s_0 columns of \mathbf{X} can be found in the supplemental article ([Dezeure et al., 2015](#)).

Once we have the design matrix \mathbf{X} and coefficient vector β^0 , the responses Y are generated according to the linear model equation with $\varepsilon \sim \mathcal{N}(0, 1)$.

2.5.3 p -values. We investigate multiple testing corrected p -values for two-sided testing of the null hypotheses $H_{0,j} : \beta_j^0 = 0$ for $j = 1, \dots, p$. We report the power and the familywise error rate (FWER) for each method:

$$\text{Power} = \sum_{j \in S_0} \mathbb{P}[H_{0,j} \text{ is rejected}] / s_0,$$

$$\text{FWER} = \mathbb{P}[\exists j \in S_0^c : H_{0,j} \text{ is rejected}].$$

We calculate empirical versions of these quantities based on fitting 100 simulated responses Y coming from newly generated ε .

For every design type, active set size and coefficient type combination we obtain 50 data points of the empirical versions of “Power” and “FWER,” from 50 independent simulations. Thereby, each data point has a newly generated X , β^0 (if not fixed) and active set positions $S_0 \in \{1, \dots, p\}$; thus, the 50 data points indicate the variability with respect to the three quantities in the data generation (for the same covariance model of the design, the same model for the regression parameter and its active set positions). The data points are grouped in plots by design type and active set size.

We also report the average number of false positives $\text{AVG}(\vee)$ over all data points per method next to the FWER plot.

The results, illustrating the performance for various methods, can be found in [Figures 2, 3, 4 and 5](#).

2.5.4 Confidence intervals. We investigate confidence intervals for the one particular setup of the Toeplitz design, active set size $s_0 = 3$ and coefficients $\beta_j^0 \sim U[0, 2]$ ($j \in S_0$). The active set positions are chosen to be the first s_0 columns of \mathbf{X} . The results we show will correspond to a single data point in the p -value results.

In [Figure 6](#), 100 confidence intervals are plotted for each coefficient for each method. These confidence intervals are the results of fitting 100 different responses Y resulting from newly generated ε error terms.

For the Multi sample-splitting method from [Section 2.1.1](#), if a variable did not get selected often enough in the sample splits, there is not enough information to draw a confidence interval for it. This is represented in the plot by only drawing confidence intervals when this was not the case. If the (uncheckable)

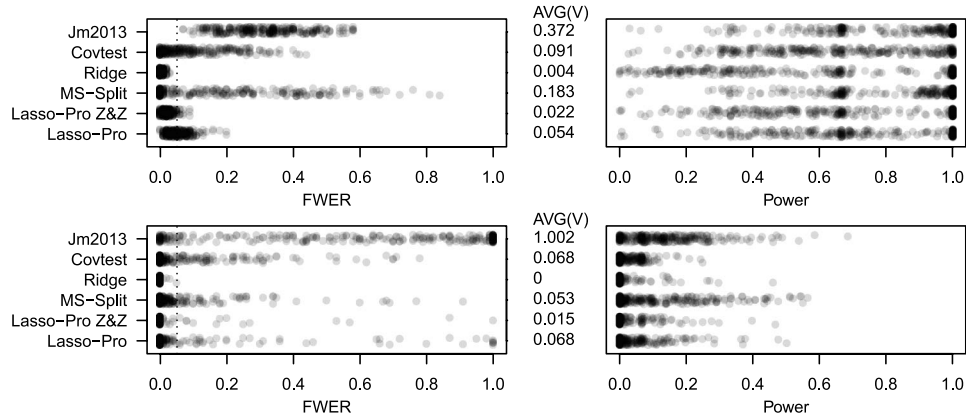


FIG. 2. Familywise error rate (FWER), average number of false positive [AVG(V)] and power for multiple testing based on various methods for a linear model. The desired control level for the FWER is $\alpha = 0.05$. The average number of false positives AVG(V) for each method is shown in the middle. The design matrix is of type Toeplitz, and the active set size being $s_0 = 3$ (top) and $s_0 = 15$ (bottom).

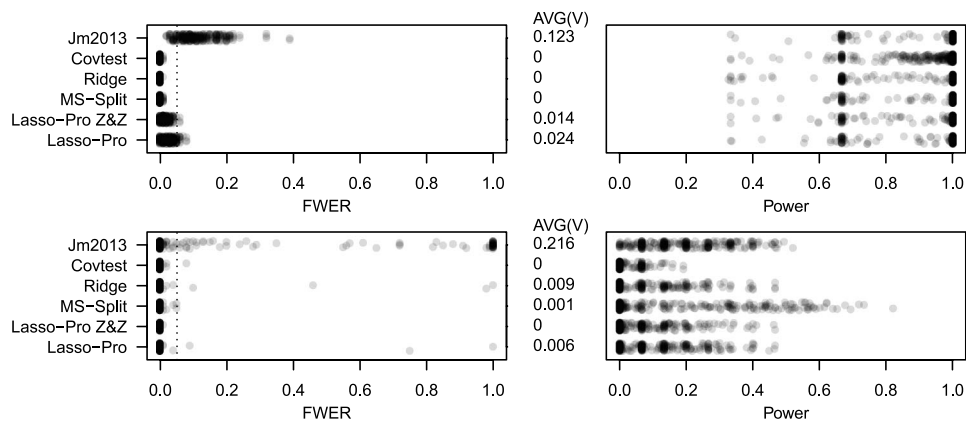


FIG. 3. See caption of Figure 2 with the only difference being the type of design matrix. In this plot, the design matrix type is Exp.decay.

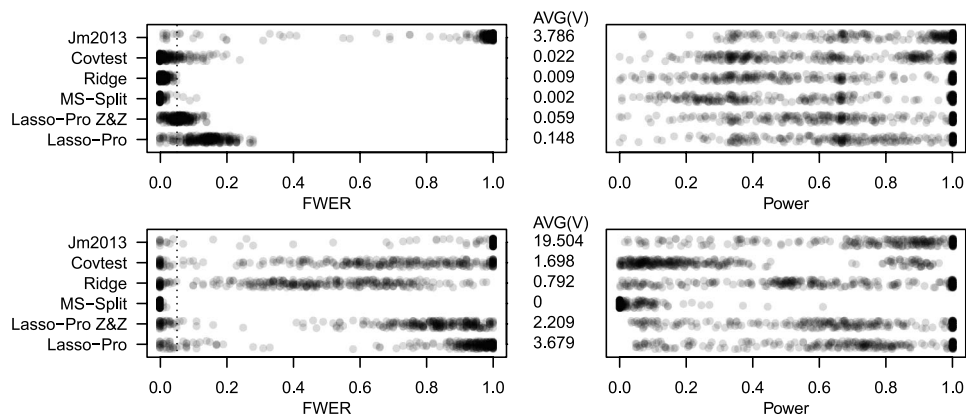


FIG. 4. See caption of Figure 2 with the only difference being the type of design matrix. In this plot, the design matrix type is Equi.corr.

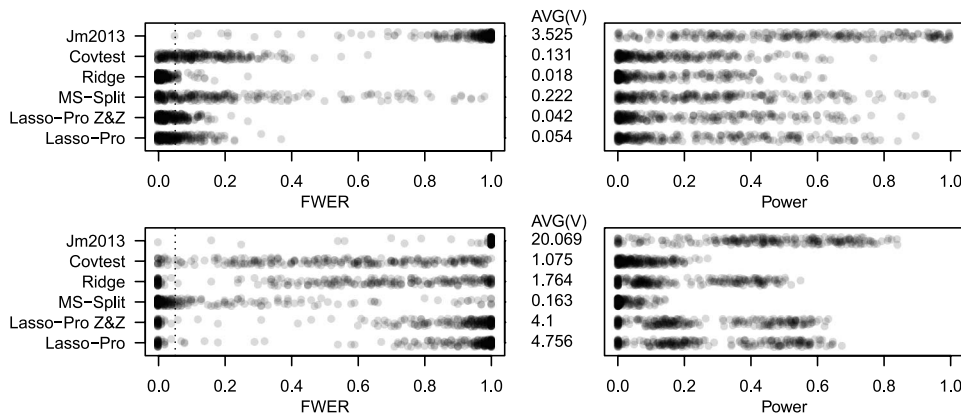


FIG. 5. See caption of Figure 2 with the only difference being the type of design matrix. In this plot, the design matrix type is RealX.

beta-min condition (2.3) would be fulfilled, we would know that those confidence intervals cover zero.

For the bootstrapping methods, an invisible confidence interval is the result of the coefficient being set to zero in all bootstrap iterations.

2.5.5 Summarizing the empirical results. As a first observation, the impact of the sparsity of the problem on performance cannot be denied. The power clearly gets worse for $s_0 = 15$ for the Toeplitz and Exp.decay setups. The FWER becomes too high for quite a few

methods for $s_0 = 15$ in the cases of Equi.corr and Re-alX.

For the sparsity $s_0 = 3$, the Ridge projection method manages to control the FWER as desired for all setups. In the case of $s_0 = 15$, it is the Multi sample-splitting method that comes out best in comparison to the other methods. Generally speaking, good error control tends to be associated with a lower power, which is not too surprising since we are dealing with the trade-off between type I and type II errors. The desparsified Lasso method turns out to be a less conservative alternative

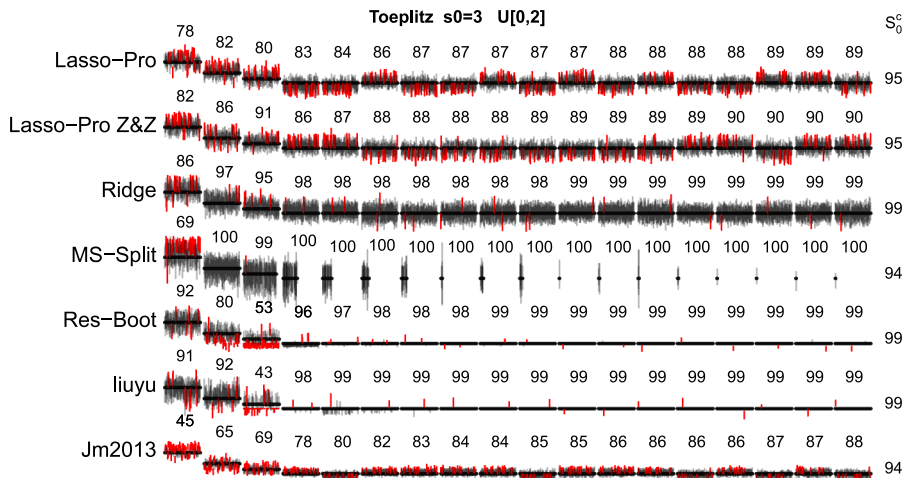


FIG. 6. Confidence intervals and their coverage rates for 100 realizations of a linear model with fixed design of dimensions $n = 100$, $p = 500$. The design matrix was of type Toeplitz and the active set was of size $s_0 = 3$. The nonzero coefficients were chosen by sampling once from the uniform distribution $U[0, 2]$. For each method, 18 coefficients are shown from left to right with the 100 estimated 95%-confidence intervals drawn for each coefficient. The first 3 coefficients are the non-zero coefficients in descending order of value. The other 15 coefficients, to the right of the first 3, were chosen to be those coefficients with the worst coverage. The size of each coefficient is illustrated by the height of a black horizontal bar. To illustrate the coverage of the confidence intervals, each confidence interval is either colored red or black depending on the inclusion of the true coefficient in the interval. Black means the true coefficient was covered by the interval. The numbers written above the coefficients are the number of confidence intervals, out of 100, that covered the truth. All confidence intervals are on the same scale such that one can easily see which methods have wider confidence intervals. To summarize the coverage for all zero coefficients S_0^c (including those not shown on the plot), the rounded average coverage of those coefficients is given to the right of all coefficients.

with not perfect but reasonable FWER control as long as the problem is sparse enough ($s_0 = 3$). The method has a slightly too high FWER for the Equi.corr and RealX setups, but FWER around 0.05 for Toeplitz and Exp.decay designs. Doing the Z&Z tuning procedure helps the error control, as can be seen most clearly in the Equi.corr setup.

The results for the simulations where the positions for the nonzero coefficients were not randomly chosen, presented in the supplemental article (Dezeure et al., 2015), largely give the same picture. In comparison to the results presented before, the Toeplitz setup is easier while the Exp.decay setup is more challenging. The Equi.corr results are very similar to the ones from before, which is to be expected from the covariance structure.

Looking into the confidence interval results, it is clear that the confidence intervals of the Multi sample-splitting method and the Ridge projection method are wider than the rest. For the bootstrapping methods, the super-efficiency phenomenon mentioned in Section 2.3 is visible. Important to note here is that the smallest nonzero coefficient, the third column, has very poor coverage from these methods.

We can conclude that the coverage of the zero coefficients is decent for all methods and that the coverage of the nonzero coefficients is in line with the error rates for the p -values.

Confidence interval results for many other setup combinations are provided in the supplemental article (Dezeure et al., 2015). The observations are to a large extent the same.

3. GENERALIZED LINEAR MODELS

Consider a generalized linear model

$$Y_1, \dots, Y_n \text{ independent,}$$

$$g(\mathbb{E}[Y_i | X_i = x]) = \mu^0 + \sum_{j=1}^p \beta_j^0 x^{(j)},$$

where $g(\cdot)$ is a real-valued, known link function. As before, the goal is to construct confidence intervals and statistical tests for the unknown parameters $\beta_1^0, \dots, \beta_p^0$, and maybe μ^0 as well.

3.1 Methods

The Multi sample-splitting method can be modified for GLMs in an obvious way: the variable screening step using the first half of the data can be based on

the ℓ_1 -norm regularized MLE, and p -values and confidence intervals using the second half of the sample are constructed from the asymptotic distribution of the (low-dimensional) MLE. Multiple testing correction and aggregation of the p -values from multiple sample splits are done exactly as for linear models in Section 2.1.1.

A desparsified Lasso estimator for GLMs can be constructed as follows (van de Geer et al., 2014): The ℓ_1 -norm regularized MLE $\hat{\theta}$ for the parameters $\theta^0 = (\mu^0, \beta^0)$ is desparsified with a method based on the Karush–Kuhn–Tucker (KKT) conditions for $\hat{\theta}$, leading to an estimator with an asymptotic Gaussian distribution. The Gaussian distribution can then be used to construct confidence intervals and hypothesis tests.

3.2 Weighted Squared Error Approach

The problem can be simplified in such a way that we can apply the approaches for the linear model from Section 2. This can be done for all types of generalized linear models (as shown in Appendix A.3), but we restrict ourselves in this section to the specific case of logistic regression. Logistic regression is usually fitted by applying the iteratively reweighted least squares (IRLS) algorithm where at every iteration one solves a weighted least squares problem (Hastie, Tibshirani and Friedman, 2009).

The idea is now to apply a standard ℓ_1 -penalized fitting of the model, build up the weighted least squares problem at the ℓ_1 -solution and then apply our linear model methods on this problem.

We use the notation $\hat{\pi}_i, i = 1, \dots, n$ for the estimated probability of the binary outcome. $\hat{\pi}$ is the vector of these probabilities.

From Hastie, Tibshirani and Friedman (2009), the adjusted response variable becomes

$$Y_{\text{adj}} = \mathbf{X}\hat{\beta} + \mathbf{W}^{-1}(Y - \hat{\pi}),$$

and the weighted least squares problem is

$$\hat{\beta}_{\text{new}} = \operatorname{argmin}_{\beta} (Y_{\text{adj}} - \mathbf{X}\beta)^T \mathbf{W} (Y_{\text{adj}} - \mathbf{X}\beta),$$

with weights

$$\mathbf{W} = \begin{pmatrix} \hat{\pi}_1(1 - \hat{\pi}_1) & 0 & \dots & 0 \\ 0 & \hat{\pi}_2(1 - \hat{\pi}_2) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \hat{\pi}_n(1 - \hat{\pi}_n) \end{pmatrix}.$$

We rewrite $Y_w = \sqrt{\mathbf{W}}Y_{\text{adj}}$ and $X_w = \sqrt{\mathbf{W}}\mathbf{X}$ to get

$$\hat{\beta}_{\text{new}} = \operatorname{argmin}_{\beta} (Y_w - \mathbf{X}_w\beta)^T (Y_w - \mathbf{X}_w\beta).$$

The linear model methods can now be applied to Y_w and \mathbf{X}_w , thereby the estimate $\hat{\sigma}_\varepsilon$ has to be set to the value 1. We note that in the low-dimensional case, the resulting p -values (with unregularized residuals Z_j) are very similar to the p -values provided by the standard R-function `glm`.

3.3 Small Empirical Comparison

We provide a small empirical comparison of the methods mentioned in Sections 3.1 and 3.2. When applying the linear model procedures, we use the naming from Section 2.5. The new GLM-specific methods from Section 3.1 are referred to by their linear model names with a capital G added to them.

For simulating the data, we use a subset of the variations presented in Section 2.5.2. We only look at Toeplitz and Equi.corr and an active set size of $s_0 = 3$. The number of variables is fixed at $p = 500$, but the sample size is varied $n \in \{100, 200, 400\}$. The coefficients were randomly generated:

Randomly generated: $U(0, 1), U(0, 2), U(0, 4)$.

The nonzero coefficient positions are chosen randomly in one case and fixed as the first s_0 columns of \mathbf{X} in the other.

For every combination (of type of design, type of coefficients, sample size and coefficient positions), 100 responses Y are simulated to calculate empirical versions of the “Power” and “FWER” described in Section 2.5.3. In contrast to the p -value results from Section 2.5.3, there is only one resulting data point per setup combination (i.e., no additional replication with new random covariates, random coefficients and random active set). For each method, there are 18 data

points, corresponding to 18 settings, in each plot. The results can be found in Figure 7.

Both the modified GLM methods as well as the weighted squared error approach work adequately. The Equi.corr setup does prove to be challenging for *Lasso-ProG*.

3.4 hdi for Generalized Linear Models

In the `hdi` R-package (Meier, Meinshausen and Dezeure, 2014) we also provide the option to use the Ridge projection method and the desparsified Lasso method with the weighted squared error approach.

We provide the option to specify the family of the response Y as done in the R-package `glmnet`:

```
> outRidge
  <- ridge.proj(x = x, y = y,
               family = 'binomial')
> outLasso
  <- lasso.proj(x = x, y = y,
               family = 'binomial')
```

p -values and confidence intervals are extracted in the exact same way as for the linear model case; see Section 2.2.

4. HIERARCHICAL INFERENCE IN THE PRESENCE OF HIGHLY CORRELATED VARIABLES

The previous sections and methods assume in some form or another that the effects are strong enough to enable accurate estimation of the contribution of *individual variables*.

Variables are often highly correlated for high-dimensional data. Working with a small sample size, it is impossible to attribute any effect to an individual

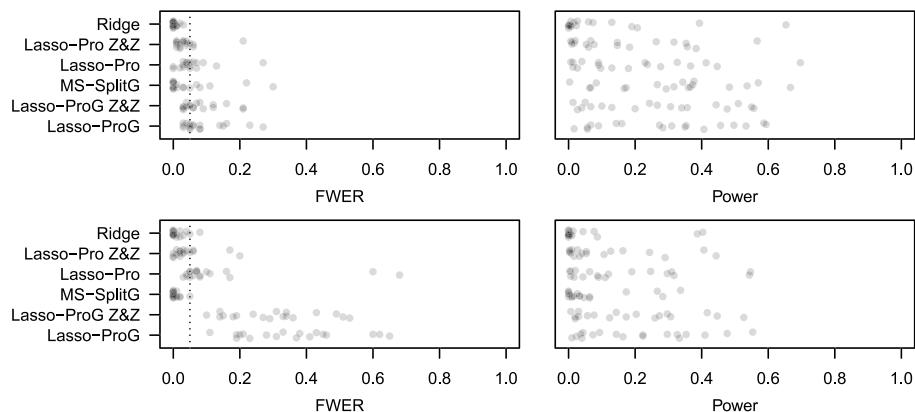


FIG. 7. Familywise error rate (FWER) and power for multiple testing based on various methods for logistic regression. The desired control level for the FWER is $\alpha = 0.05$. The design matrix is of type Toeplitz in the top plot and Equi.corr in the bottom plot. If the method name contains a capital G, it is the modified glm version, otherwise the linear model methods are using the weighted squared error approach.

variable if the correlation between a block of variables is too high. Confidence intervals for individual variables are then very wide and uninformative. Asking for confidence intervals for individual variables thus leads to poor power of all procedures considered so far. Perhaps even worse, under high correlation between variables the coverage of some procedures will also be unreliable as the necessary conditions for correct coverage (such as the compatibility assumption) are violated.

In such a scenario, the individual effects are not granular enough to be resolved. However, it might yet still be possible to attribute an effect to a group of variables. The groups can arise naturally due to a specific structure of the problem, such as in applications of the *group Lasso* (Yuan and Lin, 2006).

Perhaps more often, the groups are derived via hierarchical clustering (Hartigan, 1975), using the correlation structure or some other distance between the variables. The main idea is as follows. A hierarchy \mathcal{T} is a set of clusters or groups $\{C_k; k\}$ with $C_k \subseteq \{1, \dots, p\}$. The root node (cluster) contains all variables $\{1, \dots, p\}$. For any two clusters C_k, C_ℓ , either one cluster is a subset of the other or they have an empty intersection. Usually, a hierarchical clustering has an additional notion of a level such that, on each level, the corresponding clusters build a partition of $\{1, \dots, p\}$. We consider a hierarchy \mathcal{T} and first test the root node cluster $C_0 = \{1, \dots, p\}$ with hypothesis $H_{0,C_0} : \beta_1 = \beta_2 = \dots = \beta_p = 0$. If this hypothesis is rejected, we test the next clusters C_k in the hierarchy (all clusters whose supersets are the root node cluster C_0 only): the corresponding cluster hypotheses are $H_{0,C_k} : \beta_j = 0$ for all $j \in C_k$. For the hypotheses which can be rejected, we consider all smaller clusters whose only supersets are clusters which have been rejected by the method before, and we continue to go down the tree hierarchy until no more cluster hypotheses can be rejected.

With the hierarchical scheme in place, we still need a test for the null hypothesis $H_{0,C}$ of a cluster of variables. The tests have different properties. For example, whether a multiplicity adjustment is necessary will depend on the chosen test. We will describe below some methods that are useful for testing the effect of a group of variables and which can be used in such a hierarchical approach. The nice and interesting feature of the procedures is that they adapt automatically to the level of the hierarchical tree: if a signal of a small cluster of variables is strong, and if that cluster is sufficiently uncorrelated from all other variables or clusters, the cluster will be detected as significant. Vice-versa, if the signal is weak or if the cluster has too high a correlation

with other variables or clusters, the cluster will not become significant. For example, a single variable cannot be detected as significant if it has too much correlation to other variables or clusters.

4.1 Group-Bound Confidence Intervals Without Design Assumptions

The *Group-bound* proposed in Meinshausen (2015) gives confidence intervals for the ℓ_1 -norm $\|\beta_{C_k}^0\|_1$ of a group $C_k \subseteq \{1, \dots, p\}$ of variables. If the lower-bound of the $1 - \alpha$ confidence interval is larger than 0, then the null hypothesis $\beta_{C_k}^0 \equiv 0$ can be rejected for this group. The method combines a few properties:

(i) The confidence intervals are valid without an assumption like the compatibility condition (A.1). In general, they are conservative, but if the compatibility condition holds, they have good “power” properties (in terms of length) as well.

(ii) The test is hierarchical. If a set of variables can be rejected, all supersets will also be rejected. And vice-versa, if a group of variables cannot be rejected, none of its subsets can be rejected.

(iii) The estimation accuracy has an optimal detection rate under the so-called group effect compatibility condition, which is weaker than the compatibility condition necessary to detect the effect of individual variables.

(iv) The power of the test is unaffected by adding highly or even perfectly correlated variables in C_k to the group. The compatibility condition would fail to yield a nontrivial bound, but the group effect compatibility condition is unaffected by the addition of perfectly correlated variables to a group.

The price to pay for the assumption-free nature of the bound is a weaker power than with previously discussed approaches when the goal is to detect the effect of individual variables. However, for groups of highly correlated variables, the approach can be much more powerful than simply testing all variables in the group.

We remark that previously developed tests can be adapted to the context of hierarchical testing of groups with hierarchical adjustment for familywise error control (Meinshausen, 2008); for the Multi sample-splitting method, this is described next.

4.2 Hierarchical Multi Sample-Splitting

The Multi sample-splitting method (Section 2.1.1) can be adapted to the context of hierarchical testing of groups by using hierarchical adjustment of fami-

lywise error control (Meinshausen, 2008). When testing a cluster hypotheses $H_{0,c}$, one can use a modified form of the partial F -test for high-dimensional settings; and the multiple testing adjustment due to the multiple cluster hypotheses considered can be taken care of by a hierarchical adjustment scheme proposed in Meinshausen (2008). A detailed description of the method, denoted here by *Hier. MS-Split*, together with theoretical guarantees is given in Mandozzi and Bühlmann (2015).

4.3 Simultaneous Inference with the Ridge or Desparsified Lasso Method

Simultaneous inference for all possible groups can be achieved by considering p -values P_j of individual

hypotheses $H_{0,j} : \beta_j^0 = 0$ ($j = 1, \dots, p$) and adjusting them for simultaneous coverage, namely, $P_{\text{adjusted},j} = P_j \cdot p$. The individual p -values P_j can be obtained by the Ridge or desparsified Lasso method in Section 2.

We can then test any group hypothesis $H_{0,G} : \beta_j^0 = 0$ for all $j \in G$ by simply looking whether $\min_{j \in G} P_{\text{adjusted},j} \leq \alpha$, and we can consider as many group hypotheses as we want without any further multiple testing adjustment.

4.4 Illustrations

A semi-real data example is shown in Figure 8, where the predictor variables are taken from the Riboflavin data set (Bühlmann, Kalisch and Meier, 2014)

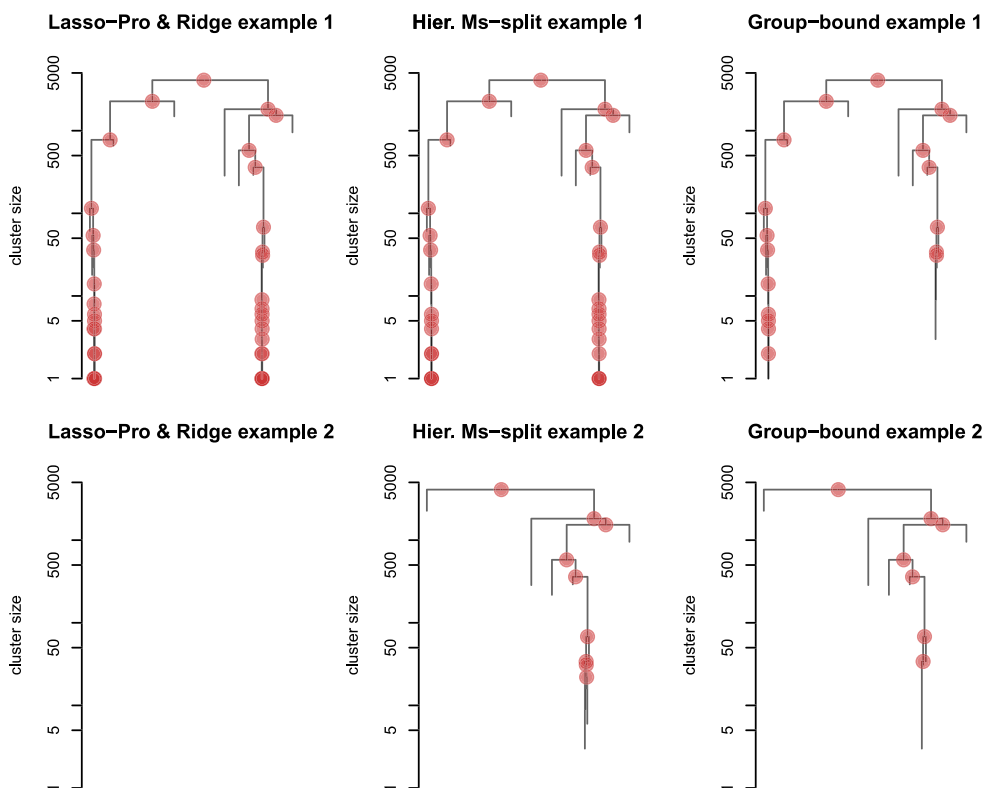


FIG. 8. A visualization of the hierarchical testing scheme as described in the beginning of Section 4, for the examples described in Section 4.4. One moves top-down through the output of a hierarchical clustering scheme, starting at the root node. For each cluster encountered, the null hypothesis that all the coefficients of that particular cluster are 0 is tested. A rejection is visualized by a red semi-transparent circle at a vertical position that corresponds to the size of the cluster. The chosen significance level was $\alpha = 0.05$. The children of significant clusters in the hierarchy are connected by a black line. The process is repeated by testing the null hypotheses for all those children clusters until no more hypotheses could be rejected. The ordering of the hierarchy in the horizontal direction has no meaning and was chosen for a clean separation of children hierarchies. The hierarchical clustering and orderings are the same for all 6 plots since the design matrix was the same. Two different examples were looked at (corresponding to top and bottom row, resp.) and four different methods were applied to these examples. The desparsified Lasso and the Ridge method gave identical results and were grouped in the two plots on the left, while results from the hierarchical Multi sample-splitting method are presented in the middle column and the results for the Group-bound method are shown in the right column. In example 1, the responses were simulated with 2 clusters of highly correlated variables of size 3 having coefficients different from zero. In example 2, the responses were simulated with 2 clusters of highly correlated variables of sizes 11 and 21 having coefficients different from zero. More details about the examples can be found in Section 4.4.

($n = 71$, $p = 4088$) and the coefficient vector is taken to have entries 0, except for 2 clusters of highly correlated variables. In example 1, the clusters both have size 3 with nonzero coefficient sizes equal to 1 for all the variables in the clusters and Gaussian noise level $\sigma = 0.1$. In example 2, the clusters are bigger and have different sizes 11 and 21; the coefficient sizes for all the variables in the clusters is again 1, but the Gaussian noise level here is chosen to be $\sigma = 0.5$.

In the first example, 6 out of the 6 relevant variables are discovered as individually significant by the *Lasso-Pro*, *Ridge* and *MS-Split* methods (as outlined in Sections 2.1.1–2.1.2), after adjusting for multiplicity.

In the second example, the methods cannot reject the single variables individually any longer. The results for the *Group-bound* estimator are shown in the right column. The *Group-bound* can reject a group of 4 and 31 variables in the first example, each containing a true cluster of 3 variables. The method can also detect a group of 2 variables (a subset of the cluster of 4) which contains 2 out of the 3 highly correlated variables. In the second example, a group of 34 variables is rejected with the *Group-bound* estimator, containing 16 of the group of 21 important variables. The smallest group of variables containing the cluster of 21 that the method can detect is of size 360. It can thus be detected that the variables jointly have a substantial effect even though the null hypothesis cannot be rejected for any variable individually. The hierarchical Multi sample-splitting method (outlined in Section 4.2) manages to detect the same clusters as the *Group-bound* method. It even goes one step further by detecting a smaller sub-cluster.

We also consider the following simulation model. The type of design matrix was chosen to be such that the population covariance matrix Σ is a block-diagonal matrix with blocks of dimension 20×20 being of the same type as Σ for *Equi.corr* (see Section 2.5.2) with off-diagonal ρ instead of 0.8. The dimensions of the problem were chosen to be $p = 500$ number of variables, $n = 100$ number of samples and noise level $\sigma = 1$. There were only 3 nonzero coefficients chosen with three different signal levels $U[0, 2]$, $U[0, 4]$ and $U[0, 8]$ being used for the simulations. Aside from varying signal level, we studied the two cases where in one case all the nonzero coefficients were contained in one single highly correlated block and in the other case each of those variables was in a different block. We

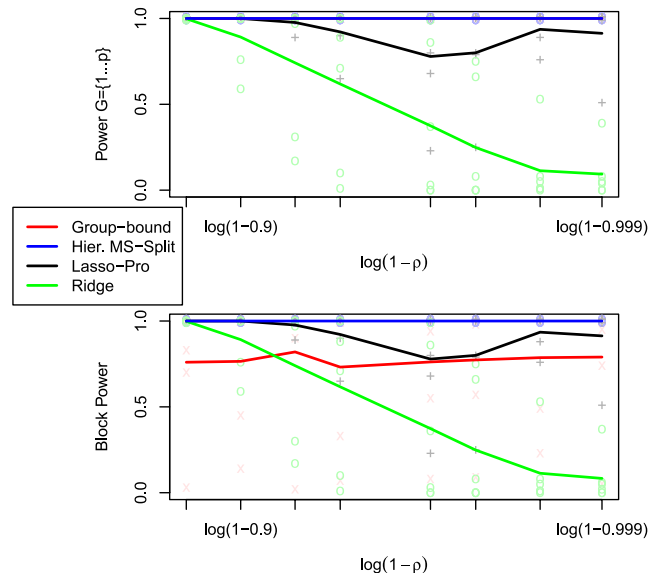


FIG. 9. The power for the rejection of the group-hypothesis of all variables (top) and the power for the rejection of the group-hypothesis of the variables in blocks highly correlated with S_0 variables (bottom). The design matrix used is of type *Block Equi.corr* which is similar to the *Equi.corr* setup in that Σ is block diagonal with blocks (of size 20×20) being the Σ of *Equi.corr*. The power is plotted as a function of the correlations in the blocks, quantified by ρ . The Ridge-based method loses power as the correlation between variables increases, while the group bound, Hier. MS-Split and Lasso-Pro methods can maintain power close to 1 for both measures of power.

look at 3 different measures of power. One can define the power as the fraction of the 100 repeated simulations that the method managed to reject the group of all variables $G = 1, \dots, p$. This is shown at the top in Figure 9. Alternatively, one can look at the rejection rate of the hypothesis for the group G that contains all variables in the highly correlated blocks that contain a variable from S_0 . This is the plot at the bottom in Figure 9. Finally, one can look at the rejection rate of the hypothesis where the group G contains only the variables in S_0 (of size 3 in this case). The type I error we define to be the fraction of the simulations in which the method rejected the group hypothesis H_{0,S_0^c} where all regression coefficients are equal to zero. These last two measures are presented in Figure 10.

The power of the Ridge-based method (Bühlmann, 2013) drops substantially for high correlations. The power of the *Group-bound* stays close to 1 at the level of the highly correlated groups (Block-power) and above (Power $G = 1, \dots, p$) throughout the entire range of correlation values. The *Lasso-Pro* and *MS-Split* perform well here as well. The power of the

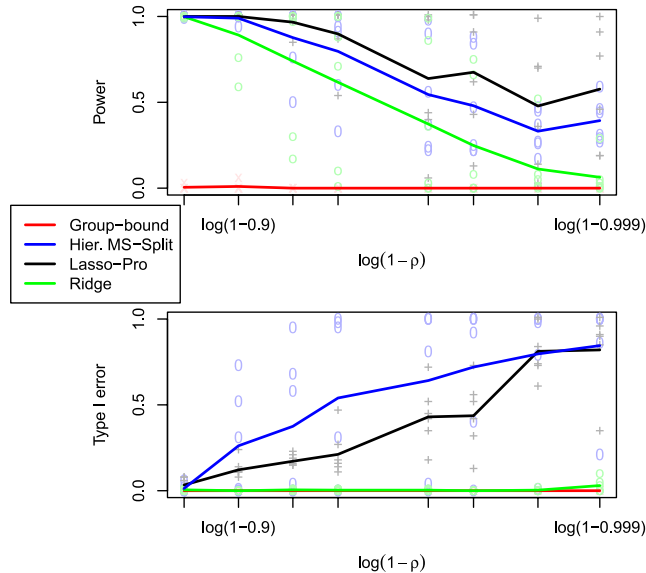


FIG. 10. The power for the rejection of the group-hypothesis of all S_0 variables (top) and type I error rate corresponding to the rejection of the group-hypothesis of all S_0^c variables (bottom) for the design matrix of type Block Equi.corr when changing the correlation ρ between variables. The design matrix type is described in detail in the caption of Figure 9 and in the text. The desparsified Lasso, Hier. MS-Split and the Ridge-based method lose power as the correlation between variables increases, while the Group-bound cannot reject the small group of variables S_0 (3 in this case). The desparsified Lasso and MS-Split methods also exceed the nominal type I error rate for high correlations (as the design assumptions break down), whereas the Ridge-based method and the Group-bound are both within the nominal 5% error rate for every correlation strength.

Group-bound is 0 when attempting to reject the small groups H_{0,S_0} . The type I error rate is supposedly controlled at level $\alpha = 0.05$ with all three methods. However, the *Lasso-Pro* and the hierarchical *MS-Split* methods fail to control the error rates, with the type I error rate even approaching 1 for large values of the correlation. The *Group-bound* and Ridge-based estimator have, in contrast, a type I error rate close to 0 for all values of the correlation.

For highly correlated groups of variables, trying to detect the effect of individual variables has thus two inherent dangers. The power to detect interesting groups of variables might be very low. And the assumptions for the methods might be violated, which invalidates the type I error control. The assumption-free *Group-bound* method provides a powerful test for the group effects even if variables are perfectly correlated, but suffers in power, relatively speaking, when variables are not highly correlated.

4.5 hdi for Hierarchical Inference

An implementation of the *Group-bound* method is provided in the *hdi* R-package (Meier, Meinshausen and Dezeure, 2014).

For specific groups, one can provide a vector or a list of vectors where the elements of the vector specify the desired columns of \mathbf{X} to be tested for. The following code tests the group hypothesis if the group contains all variables:

```
> group
  <- 1:ncol(x)
> outGroupBound
  <- groupBound(x = x, y = y,
               group = group, alpha = 0.05)
> rejection
  <- outGroupBound > 0
```

Note that one needs to specify the significance level α .

One can also let the method itself apply the hierarchical clustering scheme as described at the beginning of Section 4.

This works as follows:

```
> outClusterGroupBound
  <- clusterGroupBound(x = x,
                      y = y, alpha = 0.05)
The output contains all clusters that were tested for
significance in members. The corresponding lower
bounds are contained in lowerBound.
To extract the significant clusters, one can do
> significant.cluster.numbers
  <- which
    (outClusterGroupBound
     $lowerBound > 0)
> significant.clusters
  <- outClusterGroupBound$members
    [[significant.cluster.numbers]]
```

The figures in the style of Figure 8 can be achieved by using the function `plot` on `outClusterGroupBound`.

Note that one can specify the distance matrix used for the hierarchical clustering, as done for `hclust`.

To test group hypotheses $H_{0,G}$ for the Ridge and desparsified Lasso method as described in Section 4.3, one uses the output from the original single parameter fit, as illustrated for the group of all variables:

```
> outRidge
  <- ridge.proj(x = x, y = y)
> outLasso
  <- lasso.proj(x = x, y = y)
> group
```

```
<- 1:ncol(x)
> outRidge$groupTest(group)
> outLasso$groupTest(group)
```

To apply a hierarchical clustering scheme as done in `clusterGroupBound`, one calls `clusterGroupTest`:

```
> outRidge$clusterGroupTest
(alpha = 0.95)
```

To summarize, the R-package provides functions to test individual groups as well as to test according to a hierarchical clustering scheme for the methods *Group-bound*, Ridge and desparsified Lasso. An implementation of the hierarchical Multi sample-splitting method is not provided at this point in time.

5. STABILITY SELECTION AND ILLUSTRATION WITH HDI

Stability selection (Meinshausen and Bühlmann, 2010) is another methodology to guard against false positive selections, by controlling the expected number of false positives $\mathbb{E}[V]$. The focus is on selection of a single or a group of variables in a regression model, or on a selection of more general discrete structures such as graphs or clusters. For example, for a linear model in (2.1) and with a selection of single variables, stability selection provides a subset of variables \hat{S}_{stable} such that for $V = |\hat{S}_{\text{stable}} \cap S_0^c|$ we have that $\mathbb{E}[V] \leq M$, where M is a prespecified number.

For selection of single variables in a regression model, the method does not need a beta-min assumption, but the theoretical analysis of stability selection for controlling $\mathbb{E}[V]$ relies on a restrictive exchangeability condition (which, e.g., is ensured by a restrictive condition on the design matrix). This exchangeability condition seems far from necessary though (Meinshausen and Bühlmann, 2010). A refinement of stability selection is given in Shah and Samworth (2013).

An implementation of the stability selection procedure is available in the `hdi` R-package. It is called in a very similar way as the other methods. If we want to control, for example, $\mathbb{E}[V] \leq 1$, we use

```
> outStability
<- stability
(x = x, y = y, EV = 1)
```

The “stable” predictors are available in the element `select`.

The default model selection algorithm is the Lasso (the first q variables entering the Lasso paths). The

option `model.selector` allows to apply a user defined model selection function.

6. R WORKFLOW EXAMPLE

We go through a possible R workflow based on the Riboflavin data set (Bühlmann, Kalisch and Meier, 2014) and methods provided in the `hdi` R-package:

```
> library(hdi)
> data(riboflavin)
```

We assume a linear model and we would like to investigate which effects are statistically significant on a significance level of $\alpha = 0.05$. Moreover, we want to construct the corresponding confidence intervals.

We start by looking at the individual variables. We want a conservative approach and, based on the results from Section 2.5, we choose the Ridge projection method for its good error control:

```
> outRidge
<- ridge.proj
(x = riboflavin$x,
 y = riboflavin$y)
```

We investigate if any of the multiple testing corrected p -values are smaller than our chosen significance level:

```
> any(outRidge$pval.corr <= 0.05)
[1] FALSE
```

We calculate the 95% confidence intervals for the first 3 predictors:

```
> confint(outRidge, parm=1:3,
level=0.95)
lower upper
AADK_at -0.8848403 1.541988
AAPA_at -1.4107374 1.228205
ABFA_at -1.3942909 1.408472
```

Disappointed with the lack of significance for testing individual variables, we want to investigate if we can find a significant group instead. From the procedure proposed for the Ridge method in Section 4, we know that if the Ridge method can not find any significant individual variables, it would not find a significant group either.

We apply the Group-bound method with its clustering option to try to find a significant group:

```
> outClusterGroupBound
<- clusterGroupBound
(x = riboflavin$x,
 y = riboflavin$y,
 alpha = 0.05)
> significant.cluster.numbers
<- which(outClusterGroupBound
```



```

$lowerBound
> 0)
> significant.clusters
  <- outClusterGroupBound
     $members
     [[significant.cluster.numbers]]
> str(significant.clusters)
num [1:4088] 1 2 3 4 5 6 7 8 9 10...

```

Only a single group, being the root node of the clustering tree, is found significant.

These results are in line with the results achievable in earlier studies of the same data set in [Bühlmann, Kalisch and Meier \(2014\)](#) and [van de Geer et al. \(2014\)](#).

7. CONCLUDING REMARKS

We present a (selective) overview of recent developments in frequentist high-dimensional inference for constructing confidence intervals and assigning p -values for the parameters in linear and generalized linear models. We include some methods which are able to detect significant groups of highly correlated variables which cannot be individually detected as single variables. We complement the methodology and theory viewpoints with a broad empirical study. The latter indicates that more “stable” procedures based on Ridge estimation or sample splitting with subsequent aggregation might be more reliable for type I error control, at the price of losing power; asymptotically, power-optimal methods perform nicely in well-posed scenarios but are more exposed to fail for error control in more difficult settings where the design or degree of sparsity are more ill-posed. We introduce the R-package `hdi` which allows the user to choose from a collection of frequentist inference methods and eases reproducible research.

7.1 Post-Selection and Sample Splitting Inference

Since the main assumptions outlined in Section 2.4 might be unrealistic in practice, one can consider a different route.

The view and “POSI” (Post-Selection Inference) method by [Berk et al. \(2013\)](#) makes inferential statements which are protected against all possible submodels and, therefore, the procedure is not exposed to the issue of having selected an “inappropriate” submodel. The way in which [Berk et al. \(2013\)](#) deal with misspecification of the (e.g., linear) model is closely related to addressing this issue with the Multi sample splitting or desparsified Lasso method; see Section 2.4

and [Bühlmann and van de Geer \(2015\)](#). The method by [Berk et al. \(2013\)](#) is conservative, as it protects against any possible submodel, and it is not feasible yet for high-dimensional problems. [Wasserman \(2014\)](#) briefly describes the “HARNES” (High-dimensional Agnostic Regression Not Employing Structure or Sparsity) procedure: it is based on single data splitting and making inference for the selected submodel from the first half of the data. When giving up on the goal to infer the true or best approximating parameter β^0 in (2.13), one can drop many of the main assumptions which are needed for high-dimensional inference.

The “HARNES” is related to post-selection inference where the inefficiency of sample splitting is avoided. Some recent work includes exact post-selection inference, where the full data is used for selection and inference: it aims to avoid the potential inefficiency of single sample splitting and to be less conservative than “POSI”, thereby restricting the focus to a class of selection procedures which are determined by affine inequalities, including the Lasso and least angle regression ([Lee et al., 2013](#); [Taylor et al., 2014](#); [Fithian, Sun and Taylor, 2014](#)).

Under some conditions, the issue of selective inference can be addressed by using an adjustment factor ([Benjamini and Yekutieli, 2005](#)): this could be done by adjusting the output of our high-dimensional inference procedures, for example, from the `hdi` R-package.

APPENDIX

A.1 Additional Definitions and Descriptions

Compatibility condition ([Bühlmann and van de Geer, 2011](#), page 106). Consider a fixed design matrix \mathbf{X} . We define the following:

The compatibility condition holds if for some $\phi_0 > 0$ and all β satisfying $\|\beta_{S_0^c}\|_1 \leq 3\|\beta_{S_0}\|_1$,

$$(A.1) \quad \|\beta_{S_0}\|_1^2 \leq \beta^T \hat{\Sigma} \beta_{S_0} / \phi_0^2, \quad \hat{\Sigma} = n^{-1} \mathbf{X}^T \mathbf{X}.$$

Here β_A denotes the components $\{\beta_j; j \in A\}$ where $A \subseteq \{1, \dots, p\}$. The number ϕ_0 is called the compatibility constant.

Aggregation of dependent p -values. Aggregation of dependent p -values can be generically done as follows.

LEMMA 1 [Implicitly contained in [Meinshausen, Meier and Bühlmann \(2009\)](#)]. *Assume that we have*

B p -values $P^{(1)}, \dots, P^{(B)}$ for testing a null-hypothesis H_0 , that is, for every $b \in \{1, \dots, B\}$ and any $0 < \alpha < 1$, $\mathbb{P}_{H_0}[P^{(b)} \leq \alpha] \leq \alpha$. Consider for any $0 < \gamma < 1$ the empirical γ -quantile

$$Q(\gamma) = \min(\text{empirical } \gamma\text{-quantile}\{P^{(1)}/\gamma, \dots, P^{(B)}/\gamma\}, 1),$$

and the minimum value of $Q(\gamma)$, suitably corrected with a factor, over the range $(\gamma_{\min}, 1)$ for some positive (small) $0 < \gamma_{\min} < 1$:

$$P = \min\left(\left(1 - \log(\gamma_{\min})\right) \min_{\gamma \in (\gamma_{\min}, 1)} Q(\gamma), 1\right).$$

Then, both $Q(\gamma)$ [for any fixed $\gamma \in (0, 1)$] and P are conservative p -values satisfying for any $0 < \alpha < 1$, $\mathbb{P}_{H_0}[Q(\gamma) \leq \alpha] \leq \alpha$ or $\mathbb{P}_{H_0}[P \leq \alpha] \leq \alpha$, respectively.

Bounding the error of the estimated bias correction in the desparsified Lasso. We will argue now why the error from the bias correction

$$\sum_{k \neq j} \sqrt{n} P_{jk} (\hat{\beta}_k - \beta_k^0)$$

is negligible. From the KKT conditions when using the Lasso of $\mathbf{X}^{(j)}$ versus $\mathbf{X}^{(-j)}$, we have (Bühlmann and van de Geer, 2011, cf. Lemma 2.1)

$$(A.2) \quad \max_{k \neq j} 2|n^{-1}(\mathbf{X}^{(k)})^T \mathbf{Z}^{(j)}| \leq \lambda_j.$$

Therefore,

$$\begin{aligned} & \left| \sqrt{n} \sum_{k \neq j} P_{jk} (\hat{\beta}_k - \beta_k^0) \right| \\ & \leq \sqrt{n} \max_{k \neq j} |P_{jk}| \|\hat{\beta} - \beta^0\|_1 \\ & \leq 2\sqrt{n} \lambda_j \|\hat{\beta} - \beta^0\|_1 (n^{-1}(\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)})^{-1}. \end{aligned}$$

Assuming sparsity and the compatibility condition (A.1), and when choosing $\lambda_j \asymp \sqrt{\log(p)/n}$, one can show that $(n^{-1}(\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)})^{-1} = O_P(1)$ and $\|\hat{\beta} - \beta^0\|_1 = O_P(s_0 \sqrt{\log(p)/n})$ [for the latter, see (2.2)]. Therefore,

$$\begin{aligned} & \left| \sqrt{n} \sum_{k \neq j} P_{jk} (\hat{\beta}_k - \beta_k^0) \right| \\ & \leq O_P(\sqrt{n} s_0 \sqrt{\log(p)/n} \lambda_j) \\ & = O_P(s_0 \log(p) n^{-1/2}), \end{aligned}$$

where the last bound follows by assuming $\lambda_j \asymp \sqrt{\log(p)/n}$. Thus, if $s_0 \ll n^{1/2}/\log(p)$, the error from bias correction is asymptotically negligible.

Choice of λ_j for desparsified Lasso. We see from (A.2) that the numerator of the error in the bias correction term (i.e., the P_{jk} 's) is decreasing as $\lambda_j \searrow 0$; for controlling the denominator, λ_j should not be too small to ensure that the denominator [i.e., $n^{-1}(\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)}$] behaves reasonable (staying away from zero) for a fairly large range of λ_j .

Therefore, the strategy is as follows:

1. Compute a Lasso regression of $\mathbf{X}^{(j)}$ versus all other variables $\mathbf{X}^{(-j)}$ using CV, and the corresponding residual vector is denoted by $\mathbf{Z}^{(j)}$.
2. Compute $\|\mathbf{Z}^{(j)}\|_2^2 / ((\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)})^2$ which is the asymptotic variance of $\hat{b}_j / \sigma_\varepsilon$, assuming that the error in the bias correction is negligible.
3. Increase the variance by 25%, that is, $V_j = 1.25 \|\mathbf{Z}^{(j)}\|_2^2 / ((\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)})^2$.
4. Search for the smallest λ_j such that the corresponding residual vector $\mathbf{Z}^{(j)}(\lambda_j)$ satisfies

$$\|\mathbf{Z}^{(j)}(\lambda_j)\|_2^2 / ((\mathbf{X}^{(j)})^T \mathbf{Z}^{(j)}(\lambda_j))^2 \leq V_j.$$

This procedure is similar to the choice of λ_j advocated in Zhang and Zhang (2014).

Bounding the error of bias correction for the Ridge projection. The goal is to derive the formula (2.11). Based on (2.9), we have

$$\begin{aligned} & \sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} (\hat{b}_{R;j} - \beta_j^0) \\ & \approx \Omega_{R;jj}^{-1/2} W_j / P_{R;jj} \\ & \quad + \sigma_\varepsilon^{-1} \Omega_{R;jj}^{-1/2} \Delta_{R;j}, \quad W \sim \mathcal{N}_p(0, \Omega_R), \\ & |\Delta_{R;j}| \leq \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right| \|\hat{\beta} - \beta^0\|_1. \end{aligned}$$

In relation to the result in Fact 2 for the desparsified Lasso, the problem here is that the behaviors of $\max_{k \neq j} |P_{R;jj}^{-1} P_{R;jk}|$ and of the diagonal elements $\Omega_{R;jj}$ are hard to control, but, fortunately, these quantities are fixed and observed for fixed design \mathbf{X} .

By invoking the compatibility constant for the design \mathbf{X} , we obtain the bound for $\|\hat{\beta} - \beta^0\|_1 \leq s_0 4\lambda / \phi_0$ in (2.2) and, therefore, we can upper-bound

$$|\Delta_{R;j}| \leq 4s_0 \lambda / \phi_0^2 \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right|.$$

Asymptotically, for Gaussian errors, we have with high probability

$$(A.3) \quad \begin{aligned} |\Delta_{R;j}| &= O\left(s_0 \sqrt{\log(p)/n} \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right| \right) \\ &\leq O\left((\log(p)/n)^{1/2-\xi} \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right| \right), \end{aligned}$$

where the last inequality holds due to assuming $s_0 = O((n/\log(p))^\xi)$ for some $0 < \xi < 1/2$. In practice, we use the bound from (A.3) in the form

$$\Delta_{Rbound;j} := \max_{k \neq j} \left| \frac{P_{R;jk}}{P_{R;jj}} \right| (\log(p)/n)^{1/2-\xi},$$

with the typical choice $\xi = 0.05$.

A.2 Confidence Intervals for Multi Sample-Splitting

We construct confidence intervals that satisfy the duality with the p -values from equation (2.5), and, thus, they are corrected already for multiplicity:

$$\begin{aligned} (1-\alpha)\% \text{ CI} &= \text{Those values } c \text{ for which the } p\text{-value} \geq \\ &\quad \alpha \text{ for testing the null hypothesis } H_{0,j} : \beta = c, \\ &= \text{Those } c \text{ for which the } p\text{-value resulting from} \\ &\quad \text{the } p\text{-value aggregation procedure is } \geq \alpha, \\ &= \{c | P_j \geq \alpha\}, \\ &= \left\{ c | (1 - \log \gamma_{\min}) \inf_{\gamma \in (\gamma_{\min}, 1)} Q_j(\gamma) \geq \alpha \right\}, \\ &= \{c | \forall \gamma \in (\gamma_{\min}, 1) : (1 - \log \gamma_{\min}) Q_j(\gamma) \geq \alpha\}, \\ &= \{c | \forall \gamma \in (\gamma_{\min}, 1) : \\ &\quad \min(1, \text{emp. } \gamma \text{ quantile}(P_{\text{corr};j}^{[b]})/\gamma) \geq \\ &\quad \alpha/(1 - \log \gamma_{\min})\}, \\ &= \{c | \forall \gamma \in (\gamma_{\min}, 1) : \\ &\quad \text{emp. } \gamma \text{ quantile}(P_{\text{corr};j}^{[b]})/\gamma \geq \\ &\quad \alpha/(1 - \log \gamma_{\min})\}, \\ &= \left\{ c | \forall \gamma \in (\gamma_{\min}, 1) : \right. \\ &\quad \left. \text{emp. } \gamma \text{ quantile}(P_{\text{corr};j}^{[b]}) \geq \frac{\alpha \gamma}{(1 - \log \gamma_{\min})} \right\}. \end{aligned}$$

We will use the notation $\gamma^{[b]}$ for the position of $P_{\text{corr};j}^{[b]}$ in the ordering by increasing the value of the corrected p -values $P_{\text{corr};j}^{[l]}$, divided by B .

We can now rewrite our former expression in a form explicitly using our information from every sample split

$$\begin{aligned} (1-\alpha)\% \text{ CI} &= \left\{ c | \forall b = 1, \dots, B : (\gamma^{[b]} \leq \gamma_{\min}) \right. \\ &\quad \left. \vee \left(P_{\text{corr};j}^{[b]} \geq \frac{\alpha \gamma^{[b]}}{(1 - \log \gamma_{\min})} \right) \right\} \\ &= \left\{ c | \forall b = 1, \dots, B : (\gamma^{[b]} \leq \gamma_{\min}) \right. \\ &\quad \left. \vee \left(c \in \text{the} \left(1 - \frac{\alpha \gamma^{[b]}}{(1 - \log \gamma_{\min}) |\hat{S}^{[b]}|} \right) \right. \right. \\ &\quad \left. \left. \cdot 100\% \text{ CI for split } b \right) \right\}. \end{aligned}$$

For single testing (not adjusted for multiplicity), the corresponding confidence interval becomes

$$\begin{aligned} (1-\alpha)\% \text{ CI} &= \left\{ c | \forall b = 1, \dots, B : (\gamma^{[b]} \leq \gamma_{\min}) \right. \\ &\quad \left. \vee \left(c \in \text{the} \left(1 - \frac{\alpha \gamma^{[b]}}{(1 - \log \gamma_{\min})} \right) \right. \right. \\ &\quad \left. \left. \cdot 100\% \text{ CI for split } b \right) \right\}. \end{aligned}$$

If one has starting points with one being in the confidence interval and the other one outside of it, one can apply the bisection method to find the bound in between these points.

A.3 Weighted Squared Error Approach for General GLM

We describe the approach presented in Section 3.2 in a more general way. One algorithm for fitting generalized linear models is to calculate the maximum likelihood estimates $\hat{\beta}$ by applying iterative weighted least squares (McCullagh and Nelder, 1983).

As in Section 3.2, the idea is now to apply a standard l1-penalized fitting of the model, then build up the weighted least squares problem at the l1-solution and apply our linear model methods on this problem.

From McCullagh and Nelder (1983), using the notation $\hat{z}_i = g^{-1}((\mathbf{X}\hat{\beta})_i)$, $i = 1, \dots, n$, the adjusted response variable becomes

$$Y_{i,\text{adj}} = (\mathbf{X}\hat{\beta})_i + (Y_i - \hat{z}_i) \left. \frac{\partial g(z)}{\partial z} \right|_{z=\hat{z}_i}, \quad i = 1, \dots, n.$$

- CHATTERJEE, A. and LAHIRI, S. N. (2013). Rates of convergence of the adaptive LASSO estimators to the oracle distribution and higher order refinements by the bootstrap. *Ann. Statist.* **41** 1232–1259. [MR3113809](#)
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **39** 1–38. [MR0501537](#)
- DEZEURE, R., BÜHLMANN, P., MEIER, L. and MEINSHAUSEN, N. (2015). Supplement to “High-Dimensional Inference: Confidence Intervals, p -Values and R-Software hdi.” DOI:10.1214/15-STSS27SUPP.
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **70** 849–911. [MR2530322](#)
- FAN, J. and LV, J. (2010). A selective overview of variable selection in high dimensional feature space. *Statist. Sinica* **20** 101–148. [MR2640659](#)
- FAN, J., XUE, L. and ZOU, H. (2014). Strong oracle optimality of folded concave penalized estimation. *Ann. Statist.* **42** 819–849. [MR3210988](#)
- FITHIAN, W., SUN, D. and TAYLOR, J. (2014). Optimal inference after model selection. Preprint. Available at [arXiv:1410.2597](#).
- HARTIGAN, J. A. (1975). *Clustering Algorithms*. Wiley, New York. [MR0405726](#)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York. [MR2722294](#)
- JAVANMARD, A. and MONTANARI, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.* **15** 2869–2909. [MR3277152](#)
- KNIGHT, K. and FU, W. (2000). Asymptotics for Lasso-type estimators. *Ann. Statist.* **28** 1356–1378. [MR1805787](#)
- LEE, J., SUN, D., SUN, Y. and TAYLOR, J. (2013). Exact post-selection inference, with application to the Lasso. Preprint. Available at [arXiv:1311.6238](#).
- LEEB, H. and PÖTSCHER, B. M. (2003). The finite-sample distribution of post-model-selection estimators and uniform versus nonuniform approximations. *Econometric Theory* **19** 100–142. [MR1965844](#)
- LIU, H. and YU, B. (2013). Asymptotic properties of Lasso + mL and Lasso + Ridge in sparse high-dimensional linear regression. *Electron. J. Stat.* **7** 3124–3169. [MR3151764](#)
- LOCKHART, R., TAYLOR, J., TIBSHIRANI, R. J. and TIBSHIRANI, R. (2014). A significance test for the Lasso. *Ann. Statist.* **42** 413–468. [MR3210970](#)
- MANDOZZI, J. and BÜHLMANN, P. (2015). Hierarchical testing in the high-dimensional setting with correlated variables. *J. Amer. Statist. Assoc.* To appear. DOI:10.1080/01621459.2015.1007209. Available at [arXiv:1312.5556](#).
- MCCULLAGH, P. and NELDER, J. A. (1983). *Generalized Linear Models*, 2nd ed. Chapman & Hall, London. [MR0727836](#)
- MEIER, L., MEINSHAUSEN, N. and DEZEURE, R. (2014). hdi: High-Dimensional Inference. R package version 0.1-2.
- MEINSHAUSEN, N. (2008). Hierarchical testing of variable importance. *Biometrika* **95** 265–278. [MR2521583](#)
- MEINSHAUSEN, N. (2015). Group-bound: Confidence intervals for groups of variables in sparse high-dimensional regression without assumptions on the design. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* To appear. DOI:10.1111/rssb.12094. Available at [arXiv:1309.3489](#).
- MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)
- MEINSHAUSEN, N. and BÜHLMANN, P. (2010). Stability selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **72** 417–473. [MR2758523](#)
- MEINSHAUSEN, N., MEIER, L. and BÜHLMANN, P. (2009). p -values for high-dimensional regression. *J. Amer. Statist. Assoc.* **104** 1671–1681. [MR2750584](#)
- PEARL, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge Univ. Press, Cambridge. [MR1744773](#)
- REID, S., TIBSHIRANI, R. and FRIEDMAN, J. (2013). A study of error variance estimation in Lasso regression. Preprint. Available at [arXiv:1311.5274](#).
- SHAH, R. D. and SAMWORTH, R. J. (2013). Variable selection with error control: Another look at stability selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **75** 55–80. [MR3008271](#)
- SHAO, J. and DENG, X. (2012). Estimation in high-dimensional linear models with deterministic design matrices. *Ann. Statist.* **40** 812–831. [MR2933667](#)
- SPIRITES, P., GLYMOUR, C. and SCHEINES, R. (2000). *Causation, Prediction, and Search*, 2nd ed. MIT Press, Cambridge, MA. [MR1815675](#)
- SUN, T. and ZHANG, C.-H. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898. [MR2999166](#)
- TAYLOR, J., LOCKHART, R., TIBSHIRANI, R. and TIBSHIRANI, R. (2014). Exact post-selection inference for forward stepwise and least angle regression. Preprint. Available at [arXiv:1401.3889](#).
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **58** 267–288. [MR1379242](#)
- VAN DE GEER, S. (2007). The deterministic Lasso. In *JSM Proceedings* 140. American Statistical Association, Alexandria, VA.
- VAN DE GEER, S. (2014). Statistical theory for high-dimensional models. Preprint. Available at [arXiv:1409.8557](#).
- VAN DE GEER, S. (2015). χ^2 -confidence sets in high-dimensional regression. Preprint. Available at [arXiv:1502.07131](#).
- VAN DE GEER, S. A. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.* **3** 1360–1392. [MR2576316](#)
- VAN DE GEER, S., BÜHLMANN, P. and ZHOU, S. (2011). The adaptive and the thresholded Lasso for potentially misspecified models (and a lower bound for the Lasso). *Electron. J. Stat.* **5** 688–749. [MR2820636](#)
- VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.* **42** 1166–1202. [MR3224285](#)
- WASSERMAN, L. (2014). Discussion: “A significance test for the Lasso”. *Ann. Statist.* **42** 501–508. [MR3210975](#)
- WASSERMAN, L. and ROEDER, K. (2009). High-dimensional variable selection. *Ann. Statist.* **37** 2178–2201. [MR2543689](#)

- YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 49–67. [MR2212574](#)
- ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. [MR2604701](#)
- ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 217–242. [MR3153940](#)
- ZOU, H. (2006). The adaptive Lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. [MR2279469](#)
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **67** 301–320. [MR2137327](#)