Published in final edited form as:

Stat Appl Genet Mol Biol. 2016 August 1; 15(4): 305–320. doi:10.1515/sagmb-2015-0043.

The use of vector bootstrapping to improve variable selection precision in Lasso models

Charles Laurin,

Integrative Epidemiology Unit, School of Social and Community Medicine, University of Bristol, Bristol, BS8 2BN, UK. Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, USA

Dorret Boomsma, and

Department of Biological Psychology, VU University Amsterdam, Amsterdam, 1081 HV, Netherlands

Gitta Lubke

Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, USA. Department of Biological Psychology, VU University Amsterdam, Amsterdam, 1081 HV, Netherlands

Abstract

The Lasso is a shrinkage regression method that is widely used for variable selection in statistical genetics. Commonly, K -fold cross-validation is used to fit a Lasso model. This is sometimes followed by using bootstrap confidence intervals to improve precision in the resulting variable selections. Nesting cross-validation within bootstrapping could provide further improvements in precision, but this has not been investigated systematically. We performed simulation studies of Lasso variable selection precision with and without nesting cross-validation within bootstrapping. Data were simulated to represent genomic data under a polygenic model as well as under a model with effect sizes representative of typical GWAS results. We compared these approaches to each other as well as to software defaults for the Lasso. Nested cross-validation had the most precise variable selection at small effect sizes. At larger effect sizes, there was no advantage to nesting. We illustrated the nested approach with empirical data comprising SNPs and SNP-SNP interactions from the most significant SNPs in a GWAS of borderline personality symptoms. In the empirical example, we found that the default Lasso selected low-reliability SNPs and interactions which were excluded by bootstrapping.

Keywords

Variable selection; Lasso; Bootstrap; Additive-by-Additive Epistasis; Association; Polygenic model

1 Introduction

Multiple linear regression is most useful when it is applied to samples in which the number of predictors is small relative to the number of observations. However, if the number of predictors is large relative to the number of observations, there will be considerable sampling variability in the estimated coefficients. Under those conditions, the increased variability of coefficients can be managed by applying shrinkage methods. Shrinkage means estimating coefficients under a constraint that leads them to have reduced absolute values, drawing them toward 0, thus reducing sampling variability. There are many such constraints, hence many shrinkage methods. One of the most important shrinkage methods is the Lasso (Tibshirani, 2011).

When the Lasso is applied to a multiple regression problem, small-valued coefficient estimates are reduced to 0 and the remaining coefficient estimates are shrunk by a fixed amount (Tibshirani, 2013). Because of this property, the Lasso is often used for variable selection (Tibshirani, 1996). In Lasso variable selection, the predictors having nonzero coefficient estimates after shrinkage are selected into the regression model, and those that are shrunk to 0 are excluded from the model.

Variable selection with the Lasso is the task of deciding whether the predictor variables that could be included in a regression model are “unimportant” or as “important.” The unimportant, or “noise,” predictors are not associated with the outcome in the population, but could be strongly associated in a given sample because of sampling fluctuation. Important predictors are consistently associated with the outcome over independent samples. Unimportant predictors should be excluded from the regression model (i.e., their coefficient estimates shrunk to 0), while important predictors should be included in the model.

In practice, using the Lasso for variable selection translates to classifying the excluded predictors as unimportant and the included predictors as important. The accuracy of this claim has been studied in simulation studies, and the Lasso has often been found to have low variable selection precision (VSP), meaning that among the included predictors, a relatively large proportion (> 50%) were false positives (Devlin et al., 2003, Ayers and Cordell, 2010, He and Lin, 2011). In this paper, we propose a method for controlling the Lasso’s VSP.

We build on a variety of research that has investigated how to control the Lasso’s VSP. A main current in this research has been to estimate standard errors (SEs) and confidence intervals (CIs) for Lasso coefficient estimates with the bootstrap (e.g. Minnier et al., 2011, Chatterjee, 2011). Lasso SEs and CIs are used in a secondary variable selection step based on analogy with hypothesis testing (Freedman and Lane, 1983). The SEs are used to generate t -statistics; predictors having t -statistics below a user-set cutoff (e.g. $|t| \geq 2$) are excluded. Similarly, predictors with CIs that contain 0 are excluded.

Successfully applying a second variable selection step requires knowledge of why the Lasso includes too many variables in the model during the first step. Lasso variable selection depends on its degree of shrinkage. The degree of shrinkage in a Lasso model is controlled by a user-set metaparameter called λ .

The shrinkage metaparameter λ is a scalar; its value determines the number of variables that the Lasso selects. The larger the value of λ , the more conservative the model. The choice of λ is thus related to the Lasso's VSP. In general, each value of λ is associated with a single Lasso model. λ is commonly chosen through model-comparison methods: the user proposes a set of candidate λ s and chooses one based on, e.g., BIC or cross-validation indices.

Of these common methods, using cross-validation to choose λ has been associated with overfitting, leading to an excess of false positives and low VSP (Meinshausen and Bühlmann, 2010, James and Radchenko, 2009). Fan and colleagues (2012) attribute this to the large sample correlations that can arise between noise predictors and the outcome. They demonstrate empirically that noise correlations increase in magnitude with increasing numbers of noise predictors. Further, when sample size N is less than number of predictors p , the largest noise correlation can easily exceed true correlations in magnitude. In cross-validation, the sample size used for model-fitting is always smaller than that in the entire sample, exacerbating the problems identified by Fan *et al*, 2012.

There has been little investigation into using bootstrap SEs and CIs to mitigate low VSP associated with cross-validated selection of λ . Asymptotic analyses have found that under bootstrap resampling, Lasso coefficients for noise predictors are expected to fluctuate between positive and negative values (Chatterjee, 2011, Camponovo, 2014). This should lead noise predictors to have larger bootstrap SEs and CIs than important predictors, which suggests that these statistics are useful for identifying false positive associations.

To investigate this expectation, we introduced and evaluated a bootstrap-based method with the goal of improving Lasso VSP by excluding false positive predictor selections under cross-validation. We investigated bootstrap SEs and CIs for the Lasso when cross-validated selection of λ is done before bootstrapping, which is the standard approach (e.g. Sartori 2009, D'Angelo et al. 2009). An alternative is to nest cross-validated selection of λ within each bootstrap replication, which could lead to larger SEs and wider CIs than in the standard approach (Buckland et al., 1997, Bühlmann et al., 2011). Such larger SEs and wider CIs lead to more conservative variable selection, and possibly to improved VSP. We compared the VSP resulting from the standard approach to bootstrapping to the VSP resulting from nested selection of λ .

We did this comparison in simulated and empirical data. The simulated data were generated to resemble data observed in genetic association studies. The simulated data were high-dimensional, with a small number of weak predictors and a large number of noise predictors. The empirical data were drawn from a genome-wide association study (GWAS). We made this choice because of the prominence of genetics, in particular GWAS and the use of polygenic (risk) scores, as a context for the application and development of the Lasso (Waldron et al., 2011, Lange et al., 2014).

2 Approach

Estimation of SEs and CIs for nonzero Lasso coefficient estimates may provide a way to control the Lasso's VSP as well as to assess the relative importance of predictors. In general,

Lasso SEs or CIs cannot be estimated using a closed form (Osborne et al., 2000). A large variety of approaches has been tried to estimate Lasso SEs and CIs: see review paragraphs in e.g. Bühlmann et al. (2014), Chatterjee (2011), Kyung et al. (2010). Bootstrap estimation has received substantial interest, but many issues remain less explored, most importantly, the effects of choosing λ through cross-validation when using the bootstrap.

Next, we briefly review the two most common approaches to bootstrapping Lasso SEs and CIs: vector bootstrapping and residual bootstrapping. We follow this with a selective review of applied and methodological research using these approaches. We review both approaches to show that the behavior of the residual bootstrap has been studied in detail, but that comparatively little methodological research has been done on the vector bootstrap, despite its popularity in applied research. Hence, this paper focuses on the vector bootstrap.

2.1 Vector and Residual Bootstrapping

In nonparametric bootstrapping, samples are repeatedly drawn from observed sample data. The statistic of interest is calculated in each bootstrap sample. In this case, Lasso coefficient estimates are calculated, leading to an approximate sampling distribution. The approximate sampling distribution of Lasso coefficient estimates then permits calculation of SEs and CIs (Efron and Tibshirani, 1994).

We denote an estimate of the sampling distribution for the Lasso coefficient for predictor j as $\hat{R}(\hat{\beta}_j)$. $\hat{R}(\hat{\beta}_j)$ represents the marginal distribution of $\hat{\beta}_j$ values, as estimated using the nonparametric bootstrap. Two ways to use non-parametric bootstrapping to find $\hat{R}(\hat{\beta}_j)$ are vector bootstrapping and residual bootstrapping (Sartori, 2009).

2.1.1 The vector bootstrap—Vector bootstrapping begins with an observed sample of N observations measured on p predictors $\mathbf{x}_1, \dots, \mathbf{x}_p$ and an outcome $\mathbf{y}$. Each observation in the sample is considered as a row vector $\mathbf{z}_i = (x_{i1}, \dots, x_{ip}, y_i)$, which consists of p predictor values x_{ij} and a single outcome value y_i . A Lasso model can be fit to every bootstrap sample of N observations $\mathbf{z}_i$, yielding a set of Lasso coefficient estimates (Camponovo, 2014). $\hat{R}(\hat{\beta}_j)$ is defined as the distribution of $\hat{\beta}_j$ values from each of all possible bootstrap samples of size N . The number of unique bootstrap samples increases faster than exponentially in N ; to save computation time in practice, $\hat{R}(\hat{\beta}_j)$ is estimated using Monte Carlo simulation. The Monte Carlo estimate of $\hat{R}(\hat{\beta}_j)$ is written $\hat{F}^*(\hat{\beta}_j)$.

2.1.2 The residual bootstrap—Residual bootstrapping begins by fitting a linear regression model to a sample of N observations measured on p predictors $\mathbf{x}_1, \dots, \mathbf{x}_p$ with outcome $\mathbf{y}$, and generating the N residuals $\mathbf{e}$. Residual bootstrapping uses the sampling distribution of residuals to simulate the distribution of y values about their conditional means $\mathbf{X}\beta$. Importantly, this requires treating the observed predictor values $\mathbf{x}_1, \dots, \mathbf{x}_p$ as fixed and assuming that only the correct predictors are in the model (Efron and Gong, 1983), likely an inappropriate assumption in the context of variable selection. The residuals are then resampled.

Each bootstrap sample of N residuals, stored in the vector $\mathbf{e}^*$, can be used to generate N outcomes y^* , defined as $y^* = \mathbf{X}\hat{\beta} + \mathbf{e}^*$. The y^* values are regressed on $\mathbf{X}$ using the Lasso,

yielding, as in vector bootstrapping, a set of Lasso coefficient estimates for each resample. This set of coefficient estimates is used to define $\hat{R}(\hat{\beta}_j)$, which is typically approximated through Monte Carlo simulation, as $\hat{F}^*(\hat{\beta}_j)$.

2.2 Previous research in bootstrapping the Lasso

2.2.1 Research with residual bootstrap—Residual bootstrapping of Lasso SEs and CIs has received more methodological research interest than has vector bootstrapping. Detailed investigations of the residual bootstrapped Lasso have been undertaken by Chatterjee (2011), Minnier et al. (2011), Kyung et al. (2010), and Knight and Fu (2000), among others, with the theoretical results in Chatterjee (2011) synthesizing much of the previous work.

By comparison, the behavior of Lasso SEs and CIs under vector bootstrapping has been under-studied, particularly with λ selected through cross-validation.

2.2.2 Research with vector bootstrap—The vector-bootstrapped Lasso has often been applied in statistical genetics, sometimes with λ selected through cross-validation. Further, the vector-bootstrapped Lasso is closely related to several other prominent variable-selection methods proposed in statistical genetics (Valdar et al., 2012, Cho et al., 2010, Motyer et al., 2011).

D'Angelo et al. (2009) proposed using vector resampling to estimate SEs of Lasso coefficients of SNP-SNP and gene-gene interaction terms. Sartori (2009) compared residual and vector bootstrapping of Lasso CIs and SEs in the context of statistical genetics, including the selection of λ through cross-validation before resampling. She found that: 1) residual and vector bootstrap SEs of Lasso coefficient estimates had similar degrees of bias in linear models; and 2) vector bootstrap CIs had superior coverage in linear and in logistic models. Camponovo (2014) used vector bootstrapping to generate simultaneous confidence regions in linear models with random predictors. He found poor coverage rates, and used an asymptotic argument to propose two modified vector bootstrapping procedures. The modified procedures generated confidence regions with adequate coverage.

The present study builds on the above research: in it, we compare the vector bootstrap with and without cross-validation nested within bootstrap replications. An important difference in the present study is that VSP, rather than coverage rate, is the criterion of comparison: if the CI for an important predictor excludes the true coefficient value but also excludes 0, the variable is correctly selected. Both versions of the bootstrapped Lasso are straightforward to implement in existing statistical software; e.g. R package glmnet, PLINK 1.9. (Friedman et al., 2010, Chang et al., 2014). This makes them attractive and approachable to applied researchers, who would benefit from understanding the trade-off in VSP involved in choosing one over the other.

2.2.3 Research with cross-validated selection of λ —Many methods have been proposed to increase the VSP of Lasso regression and related methods (Chatterjee, 2011, Lockhart et al., 2013, Fan et al., 2012). However, little methodological research has addressed the combination of vector bootstrapping and cross-validated selection of λ that is

used in practice (Sartori, 2009, Cho et al., 2010). Despite the suggestion, in a recent textbook, that cross-validated selection of λ should be nested within bootstrap samples when applying the Lasso (Bühlmann et al., 2011), little published work has evaluated this procedure (Okser et al., 2014). The goal of the present paper is to address this deficiency and to investigate the conditions in which nested selection of λ leads to improved VSP. In addition, we address the question whether bootstrapped t -statistics are useful for the identification of false positives.

2.3 Role of λ when fitting Lasso models

The metaparameter λ controls the bias and parsimony of a fitted Lasso model. Equation 1 gives the definition of a Lasso model for predictors X and outcome y when λ is known (Tibshirani, 1996).

$$\hat{\beta} = \arg_{\beta} \min \frac{1}{2} \sum_i \left(y_i - \sum_j x_{ij} \beta_j \right)^2 + \lambda \sum_j |\beta_j| \quad (1)$$

The value of λ determines the degree of shrinkage toward 0 and serves as a threshold for variable selection. A predictor is selected if the absolute value of its covariance with the outcome is larger than λ , and excluded otherwise (Efron et al., 2004). This thresholding property limits the useful range of values that λ can take. The minimum value that λ can take is 0, where the Lasso fit is the same as that of OLS regression. Such solutions are unbiased, but, because of the improbability of any OLS coefficients equaling 0 exactly, they are also unparsimonious.

The maximum value that λ can take depends on the largest sample covariance of any predictor with the outcome. More specifically, when λ is equal to or greater than that covariance, all coefficients are shrunk to 0, and the fitted model is intercept-only, thus parsimonious but biased (Friedman et al., 2007).

2.3.1 Selection of λ through K -fold cross-validation—In general, each value of λ is associated with a single Lasso model (Tibshirani, 2013, Efron et al., 2004).

K -fold cross-validation (Zhang, 1993) is often used to select the best-performing model. Lasso model fitting, λ selection, and K -fold cross-validation has been described in detail for its implementation in the R package glmnet (Friedman et al., 2010).

Following this procedure, the λ value that is associated with the best-performing model is selected. The best-performing model is the one having the minimum cross-validation index, which is computed as the sum of squared residuals averaged over the K cross-validations. The selected λ value is then used to fit a finalized model by solving Equation 1 in the *entire* sample. This produces a set of selected predictors that are then indexed in set s .

Different λ values might be selected in different samples from the same population due to the influence of noise correlations (Fan et al., 2012). In the next section, we interpret selection of λ as a source of variation in Lasso coefficient estimates.

2.3.2 Lasso variance estimates: Contribution of λ selection—The variance of Lasso coefficient estimates depends on the joint distribution of the p predictor variables X and the outcome y , as well as on the value of λ (Pötscher and Leeb, 2009). The conditional distribution of estimates for a single predictor, denoted $g_j(\hat{\beta}_j|\lambda)$, is the distribution of $\hat{\beta}_j$ coefficients at a fixed λ value. The marginal distribution, $h_j(\hat{\beta}_j)$, is the distribution of $\hat{\beta}_j$ averaged over λ values.

The variance of $\hat{\beta}_j$ can be found using either g_j or h_j . Heuristically, h_j treats the selected λ value as a realization of a random variable (Zhang, 1993, Bühlmann et al., 2014). We argue that using h_j might improve VSP because using g_j treats λ as fixed, which can underestimate the variance of coefficients.

To support this claim, consider the inequality:

$$\text{Var}(\hat{\beta}_j) = E_{\lambda} \left\{ \text{Var}_{\beta}(\hat{\beta}_j|\lambda) \right\} + \text{Var}_{\lambda} \left\{ E_{\beta}(\hat{\beta}_j|\lambda) \right\} \geq \text{Var}_{\beta}(\hat{\beta}_j|\lambda) \quad (2)$$

(Chatfield, 1995). If $\hat{\beta}_j$ and λ were independent, then $\text{Var}(\hat{\beta}_j) = \text{Var}_{\beta}(\hat{\beta}_j|\lambda)$ and there would be little difference between the fixed and random λ approaches in practice. However, $\hat{\beta}_j$ and λ are not necessarily independent: the range of possible λ values is bounded by $(0, r_{\max})$. Thus, although using g_j (treating $\hat{\beta}_j$ and λ as independent) has the practical advantage of using fewer computational resources, it will only be acceptable if the resulting underestimate of the standard error of $\hat{\beta}_j$ is small.

2.3.3 λ selection in bootstrapping—In practice, both g_j , the conditional distribution of $\hat{\beta}_j$ given λ , and h_j , the distribution of $\hat{\beta}_j$ averaged over all λ s, are unknown. Both distributions can be estimated using the vector bootstrap.

Finding the bootstrap estimate of the conditional distribution, $\hat{g}_j^*$ is done by fitting Lasso models to resampled X and y values, given the λ value chosen through K -fold cross-validation in the original sample.

Finding the bootstrap estimate of the marginal distribution $\hat{h}_j^*$ requires treating the selected λ value as random. Nesting λ -selection within each bootstrap replication approximates the effect of sampling error on the value of λ selected.

Our simulations compared the fixed- and random- λ approaches with respect to VSP, and suggest effect sizes at which the increased computational burden of the random approach is worthwhile.

3 Methods

The purpose of the current paper is to propose and to evaluate the use of the vector bootstrap, with λ selected through K -fold cross-validation, as a method for estimating Lasso SEs and CIs. In particular, we compared three variants of this approach: a software default approach to variable selection using the Lasso (Method 1, see Figure 1); an approach involving selection of λ before resampling (Method 2, see Figure 2); and a third approach where λ -selection is nested within bootstrap samples (Method 3, see Figure 3). Our evaluation was in terms of VSP and of accuracy of ranking predictors by relative importance. Relative importance was calculated using coefficients of variation, $(|t^{-1}|)$, where t is a bootstrapped t -statistic).

In the first step of each variable selection method, a Lasso model is fit to the entire sample. This requires selection of λ , which is done through K -fold cross-validation. The initial model fit produces a set of selected predictors, which are indexed in the set s . This step is the default application of the Lasso in the R packages `glmnet` and `grpreg`. We denote it Method 1.

Methods 2 and 3 differ from Method 1 by having a second variable selection step. In this step, further reduction of the set of selected predictors is done using bootstrap SEs or CIs. All predictors are used in vector bootstrap resampling, but, to save computational resources, SEs and CIs are not calculated for predictors that were excluded in the initial variable selection step.

Method 2 uses the same value of λ in every bootstrap sample.

Method 3 differs from Method 2 by re-selecting λ in each bootstrap sample. In both methods, after SEs or CIs are calculated, variables that: 1) include 0 in their confidence intervals; or that: 2) have a coefficient of variation greater than a certain cutoff; are excluded.

3.1 Lasso CIs and SEs: Secondary selection or ranking

Bootstrap CI or SE estimates improve Lasso models through a second step of selecting or ranking predictors. CI and SE estimates are both directly related to the sampling variance of a Lasso coefficient estimate, discussed above. A $1 - \alpha$ CI for the coefficient estimate $\hat{\beta}_j$ is generated either using the $\frac{\alpha}{2}, 1 - \frac{\alpha}{2}$ quantiles of the bootstrap distribution $\hat{g}_j^*$, or using an approximate inverted z -test, which gives the interval $\hat{\beta}_j \pm z_{\alpha/2} SE^*(\hat{\beta}_j)$, where $z_{\alpha/2}$ is the $\frac{\alpha}{2}$ quantile of a standard normal distribution and $SE^*(\hat{\beta}_j)$ is the bootstrap estimate of the standard error of $\hat{\beta}_j$.

Using either CI method, predictors that have CIs that contain 0 are excluded since this can be regarded as evidence that predictor x_j is a false positive selection.

SEs can also be used to give information about the relative importance of predictors in addition to improved VSP. We propose the use of the coefficient of variation for each nonzero Lasso regression weight as an index of relative importance and to apply cutoffs to this statistic to exclude false positives. The coefficient of variation of a random variable X , denoted $C_{\text{var}}(X)$, is the ratio of its standard error to the absolute value of its mean.

This index is sensitive to small differences in mean values and thus may be better able to distinguish small true positives from false positives. We used the vector bootstrap to estimate

C_{var} for individual Lasso coefficients, denoted $C_{\text{var}}^* \left(\hat{\beta}_j \right)$.

4 Simulation Studies

We first compared Methods 1–3 using a factorial simulation study. The simulation had two goals: first, evaluating the Methods' ability to distinguish signal from noise; second, evaluating their ability to correctly order signals of differing strengths.

The data generation models were as simple as possible while still representing two empirically interesting scenarios based on statistical genetics: 1) a low probability of selecting important predictors at random; and 2) a spectrum of small true effect sizes. To this end, data were generated under two different linear models: first, a few-important-predictors model with under 5% of predictors having true effects, and with effect sizes (given in r^2) representing 2.5% or less of outcome variance attributable to any important predictor; and second, a polygenic model, in which there were thousands of predictors, each of which had an effect drawn from a normal distribution, all of which together accounted for 60% of the variance in the outcome.

2500 Monte Carlo (MC) replications were used in each cell of the few-important-predictors design. This number of replications was chosen based on pilot studies, in which at least 2500 replications were required in order to generate relatively smooth empirical distributions of coefficient estimates (see also Sartori 2009). However, only 250 MC replications were used when data were simulated under the polygenic model due to the substantial use of computational resources needed to bootstrap such data.

$B = 1000$ bootstrap replications were used within each MC replication. The average performances of the three methods across samples were compared; within each MC replication, each method was employed on an independent sample drawn from the population distribution. This was done in order to avoid creating dependence among results that might arise from fitting the methods to the same data.

4.1 Simulation Design

Three factors were manipulated in the simulations: method, data-generating model, and effect size.

As described above, the methods compared were the fixed- (Method 2) and random- λ (Method 3) variants of the vector bootstrapped Lasso, with the default application, Method 1.

The second factor manipulated in the simulation study was the data-generating model. Three data-generating models were used: a *single* important predictor and 99 noise predictors; a 5 important predictors and 99 noise predictors, with the important predictors having different r^2 values, enabling us to *rank* them; and a *polygenic* model with 3000 predictors with effect

sizes drawn from a $\mathcal{N}(0, 0.60)$ distribution—the three predictors with the largest (absolute) effects were treated as the important predictors.

Each data generating model was a linear regression model having a standard normal outcome and binomial(2, 0.5) distributed predictors; $N = 2500$ was used as the sample size. This was chosen as a rough approximation of the sample size and predictor structure of smaller genome-wide association studies of quantitative phenotypes (Balding, 2006).

The third factor manipulated in the simulation study was effect size. Effect sizes of $r^2 = 0.01, 0.0033, 0.001$ were used in the single-important-predictor analyses. In the five-important-predictors analyses, each important predictor had a different effect size: the set $r^2 = 0.01, 0.0067, 0.0033, 0.0022, 0.001$ was used. In the polygenic model, the three strongest predictors were expected to account for 5.1% of phenotypic variance together. This estimate is based on treating the largest simulated-SNP effects as being drawn from a truncated normal distribution representing the upper 0.15% tail area of the $\mathcal{N}(0, 0.60)$, and then treating the strongest negative effects as being drawn independently from the corresponding part of the lower tail (Barr and Sherrill, 1999).

We manipulated the effect size of important predictors for two reasons: we used effect size as a measure of the “difficulty” of correctly selecting important predictors, giving us a way to use the data to influence the methods’ VSP; and because previous simulation studies (e.g. Tibshirani, 2011, Leng et al., 2006, Meinshausen and Bühlmann, 2010), used effect sizes that are now considered to be unrealistically large in the context of statistical genetics (Stefansson et al., 2009, Park et al., 2011).

In the next section, we describe the evaluation criteria.

4.2 Evaluation Criteria

The main question asked in this paper is: when does nested selection of λ lead to improved VSP over other approaches? A subsidiary question is: can C_{var}^* -statistics give information about the relative importance of predictors that might be used to identify false positives? Addressing these questions requires quantifying the performance of the different methods. We used variable selection precision (VSP) to quantify the methods’ performance and additionally used the False Negative Rate (FNR) to identify the risk that each method might be over-conservative.

$$\text{VSP is } \frac{\# \text{Important Predictors Selected}}{\# \text{Selected Predictors}} = \frac{\# \text{True Positives}}{\# \text{Positives}}.$$

VSP is set to 0 if there are no positives. Thus, in each replication of the few-important-predictors simulation, VSP ranged from $0, \frac{1}{104}, \frac{1}{103}, \dots, \frac{1}{2}, 1$, while in the polygenic model, the denominator was 3000. FNR is the proportion of truly important predictors that have been classified as unimportant by a variable selection method (Fawcett, 2006).

A positive in the vector bootstrap Lasso (Methods 2 and 3) was defined as a predictor for which the $(1 - \alpha) \times 100\%$ bootstrap percentile confidence interval (Efron and Tibshirani, 1994) excluded 0. Confidence level (α) was set to 0.05, and intervals were symmetric. In sensitivity analyses, confidence levels of 0.02, 0.10 and 0.20 were also used.

A positive in the default Lasso (Method 1) was defined as a predictor having a coefficient in the finalized Lasso model (i.e. a predictor indexed in $\mathbf{s}$).

To quantify the performance of C_{var}^* as an importance measure, we used the median rank and median absolute deviation of ranks of each predictor's C_{var}^* over MC replications. We chose this measure to obtain both a typical rank and the variability of rankings for important and for noise predictors. This importance measure was not used with data that had been simulated under a polygenic model because the randomness of the effects meant that there was no consistent mapping between a predictor's index and its effect size.

4.3 Simulation Results

4.3.1 Improved Variable Selection Precision—Use of the vector bootstrap (Methods 2 and 3) was associated with increased VSP at all effect sizes, but also with increased FNR at all effect sizes. This is shown in Table 1, for nominal $\alpha = 0.05$ and percentile bootstrap confidence intervals; normal-theory bootstrap confidence intervals are not shown because their performance was similar to, but slightly worse than that of the percentile intervals. At the smaller effect sizes, the increased precision was only apparent with random λ , and, at the smallest effect size, the magnitude of this advantage was small. The increased FNR suggests that this pattern was due to the bootstrapping procedures being more conservative than Method 1.

In sensitivity analyses, we found that the nominal coverage rate α chosen for confidence intervals interacted with the bootstrapping method used. For example, at $r^2 = 0.0033$, the random approach (Method 3) is more precise than the fixed approach (Method 2) at $\alpha = 0.01, 0.05$ (Table 1), but less precise at $\alpha = 0.10$ (not shown).

The bootstrapped coefficient of variation (C_{var}^*) was in general much larger for false positives than for true positives. Bootstrap means for noise predictors were never 0, preventing the occurrence of division-by-0 errors in computing C_{var}^* values. These observations support the use of a C_{var}^* cutoff as a way of increasing Lasso VSP. However, C_{var}^* cutoffs did not increase VSP to the extent that CIs did. They offered no improvement over the default Method 1 at effect sizes $r^2 = 0.0033, 0.01$ or in multiple-predictor models.

4.3.2 Ranking Predictors—The usefulness of C_{var}^* to rank predictors by relative importance depended on the effect size as well as the number of true predictors. At $r^2 = 0.0033, 0.01$ or with multiple true predictors, using C_{var}^* to rank predictors by relative importance yielded no improvement over ranking predictors by the absolute value of their non-bootstrapped coefficients. This is shown in Table 2. Additionally, true predictor ranks were identical regardless of whether λ was treated as random or as fixed.

For a single true predictor with $r^2 = 0.001$, ranking by C_{var}^* led to better discrimination of the true predictor than did ranking by non-bootstrapped coefficients. The true predictor did not always have the smallest C_{var}^* , but that it was within the top 6 predictors at least half the time. Without bootstrapping, the true predictor was often selected out of the model. Overall, it was in the top 50% of predictors close to half the time but that it was not frequently among the highest ranks.

4.3.3 Variability of Coefficients under Different Methods—Treating the metaparameter λ as random was associated with greater variability of coefficient estimates, as illustrated in Table 3 by larger standard deviations, and wider confidence intervals for Method 3 when compared to Method 2. Table 3 presents summaries for important predictors only; results for noise predictors were very similar (results in the table are averages over MC replications and are not conditional on the important predictors being selected into the model). Confidence interval lengths were similar between quantile-based bootstrap CIs and normal-theory bootstrap CIs.

These results are consistent with inequality (2), which suggests that the increased variability is attributable to the variance in expected coefficient values with respect to the distribution of λ . Table 4 compares five-number summaries of the distribution of λ from Method 1 to those from Method 3. The results for Method 1 are five-number summaries of the λ values that were selected across replications, while those for Method 3 are the averages across replications of the five-number-summary of λ values. Selection of λ *via* Method 3 (cross-validation nested within bootstrap replications) tends to produce lower values of λ (distribution shifted left) which are also less-variable (smaller IQR). The smaller median λ means that Method 3 performed less regularization than did Method 1, hence had coefficients with larger values and may have included more predictors in each bootstrap replication. Thus, bootstrap CIs from Method 3 would have been relatively wide, leading to increased VSP *because of* liberal variable selection within each bootstrap replication.

The results of the simulation studies quantified the methods' relative performance in idealized data. For a more critical evaluation of their practical utility, we applied them in a GWAS data set, using them to select pairwise interactions as well as main effects.

5 Empirical Illustration

We used data gathered for a genome-wide association study (GWAS) of Borderline Personality Disorder features to compare the vector bootstrap with λ -selection nested within bootstrap samples to the standard Lasso. This comparison serves as a representative analysis for possible applications of Method 3. The original GWAS was based on a sample of $N = 7124$ individuals who participated in a twin-family study of mental and somatic health (Boomsma et al., 2006, Willemsen et al., 2013); see Willemsen et al. for detailed methods including IRB approval, genotyping, and quality control procedures. Responses to a psychiatric inventory measuring Borderline Personality features were used as outcomes because they have shown a promising signal that was replicated in an independent sample (Lubke et al., 2014).

Borderline features were measured using total scores on the PAI-BOR inventory, a 24-item test (Morey, 1991). More specifically, the outcome we used in this study was the residual of PAI-BOR score after OLS regression on age, gender, and their interaction, as well as a principal component score representing ancestry (Abdellaoui et al., 2013, Price et al., 2010). Following up on D'Angelo et al. (2009)'s proposal, we fit Lasso multiple regressions of Borderline features on SNP main effects and SNP-SNP interactions. To control the computational resources required, we limited the analysis to pairwise interactions and main effects of the 125 SNPs having the strongest univariate association with the Borderline features phenotype, as listed in Lubke et al. 2014. R's memory limitations limit the application of bootstrapping to interactions between 1500 or fewer variables ($1500^2 \times 1000$ bootstrap samples $\approx 2^{31}$ objects) (R Core Team, 2013). To avoid multicollinearity, these 125 were then pruned to the set of 77 SNPs that had pairwise correlations of less than $r = 0.6$ among each other. From these, 2926 pairwise interaction terms were calculated, resulting in a total $p = 3003$, and hence up to 3003 Lasso partial regression weights needing CIs and C_{var}^* .

The primary purpose of this analysis was to compare the different methods for lasso variable selection in a data set with SNPs having different allele frequencies and unknown effects on the outcome. The secondary purpose was an exploratory analysis of epistatic effects between SNPs as predictors of borderline personality symptoms, generating hypotheses that can be tested in independent samples. We did not emphasize the need for the selected models to be biologically plausible; that is, we did not group SNPs for selection by gene or pathway. We did not rule out models of pure interaction in the absence of main effects (Cordell, 2009). Accordingly, we did not force selection of hierarchical models: that is, it was not the case that any interaction terms considered for selection must have had main effects in the model. Because of this and because of the bias caused by Lasso estimation, the values of the coefficients for (pure) interaction terms cannot be interpreted straightforwardly as magnitudes of effect modifications.

Results from applying the vector bootstrap with random λ (Method 3) were compared to those from the default Lasso (Method 1). Method 3 was used to generate percentile CIs and C_{var}^* . The resulting variable selections and rankings were compared to Method 1's selections and to ranking coefficient estimates from the finalized model by absolute value.

5.1 Empirical illustration: Results

The empirical illustration concerned the application of the bootstrapped Lasso to pairwise interaction effects. The bootstrapped Lasso produced different variable selections and importance rankings than did the default Lasso. The bootstrapped Lasso tended to select better quality predictors than did the default Lasso. The default Lasso selected predictors with low minor allele frequency (MAF), hence had very large standard error estimates. This suggests that the default Lasso can ignore important aspects of the data.

Figure 4 and Table 3 present comparisons of the default Lasso and the approach using vector bootstrapping. Figure 4 plots bootstrap mean estimates on the horizontal axis and bootstrap standard error estimates on the vertical axis. Points falling outside of the dark gray **V** have

C_{var}^* values less than 0.5 (bootstrapped t -statistics greater than 2). Predictors that were selected by the default Lasso are plotted as light gray diamonds. There is no obvious relationship between default Lasso selection of a predictor and its bootstrap moments. dbSNP lookup of the predictors in Table 3 showed that the bootstrapped Lasso was less prone to selecting interactions between SNPs having low MAF than was the default Lasso. The default Lasso, in selecting these interactions, was in effect including interactions between binomial predictors that have low success probabilities. These interactions tended to have large bootstrap standard errors, hence are excluded when a C_{var}^* cutoff is used. Interestingly, the low-MAF SNPs involved these interactions tended to have moderately strong main effects in the conventional GWAS analyses.

Using vector bootstrapping of Lasso coefficients (CI or C_{var}^* cutoff of 0.5) suggested a single interaction for followup.

6 Conclusion

Using vector bootstrap CIs on Lasso regression coefficients offers a valid way to distinguish false positive selections from true positives. Percentile CIs were associated with increased precision of variable selection at all effect sizes. At the smallest effect sizes, including those in a polygenic model, gains were only achieved when using Method 3, which treated the metaparameter λ as random. Additionally, the bootstrapped methods were more conservative in variable selection than was Method 1, the default Lasso. This suggests that, if several small effects are expected and if avoiding false positives is more important than avoiding false negatives, treating λ as random justifies increased computational cost.

The (bootstrapped) coefficient of variation, C_{var}^* does measure the relative importance of Lasso predictors, but offers little to recommend it over using the absolute value of Lasso coefficients.

We observed a “rising tide lifts all boats” effect for all methods, where a predictor having a given small effect size was more likely to be selected when the data were generated to have other important predictors. This was despite the predictors and their effects being independent of one another. The more complex models tended to have lower (more lenient) thresholds λ selected by cross-validation, regardless of the method used. A possible explanation is that a λ causing inclusion of a solitary small effect might not be able to consistently decrease the residual sum of squares in different cross-validation subsamples, but that a λ that admits multiple small effects could.

The low degree of overlap in the distributions of λ estimated by Method 1 and Method 3 suggests a need for explanation. The two distributions should have the same mean value, because the bootstrap distribution of a statistic should approximate the distribution of a statistic across repeated independent samples.

In the empirical analysis, the vector-bootstrapped Lasso excluded unreliable predictors that had been selected by the default Lasso. However, it is possible that residual bootstrapping could have led to better performance with low MAF predictors. Under vector bootstrapping,

“monoallelic” SNPs are possible within the bootstrap samples. The result would be inflation of the intercept term in the regression model, which would affect model-fitting through cross-validation. Further, VSP is dependent on the base rates of positives and negatives, and would have been skewed if important SNPs tended to have low MAFs and noise SNPs high MAFs, or vice-versa, which limits the generalizability of the simulation results to empirical data.

Three follow-up studies are suggested by this result: a simulation study comparing the two Lasso methods after manipulating the reliability of predictors; an attempt to replicate the promising interaction between rs118160379 and rs59194015; and a comparison of VSP from Method 3 to that from other Lasso confidence intervals, e.g. those in Camponovo (2014).

In addition, both our empirical and our simulation results are also relevant to research that uses polygenic scores to predict complex traits and to investigate the polygenic architecture of closely related traits. Lasso regression has recently been implemented in PLINK2, a software package that is very widely used in the analysis of complex traits. In consequence, Lasso variable selections are being used to construct polygenic scores, but have not yielded significant improvements over conventional methods (Warren et al., 2014).

Our simulation results suggest that vector bootstrapping (with nested selection of λ) may be able to yield polygenic scores with greater predictive accuracy, but will require relatively large sample sizes to avoid excluding important variants from the score. Large samples are relatively common in the GWAS context, however using this approach efficiently will require careful application and data management.

There were several limitations to this study. The data generating model had predictors that were independently and identically distributed as well as a normally distributed outcome. These attributes are unlikely to hold in empirical data. We plan to extend the current simulations with correlated and differently scaled predictors as well as skewed outcomes. Second, the simulated effect sizes, while small, were still somewhat larger than those that are typically observed in the statistical genetics of complex traits (Stefansson et al., 2009). The phenotypic variance explained by the polygenic model was consistent with a highly heritable trait, e.g. height, and is larger than would that expected for most complex traits of interest. The number of important predictors used was also much smaller than the number of genetic loci expected to influence complex phenotypes (Sivakumaran et al., 2011). Similarly, our results suggest that bootstrapping is most useful when there are many predictors to consider, to be reduced to a relatively small number of important ones, and that the single-important predictor case (tested here) is perhaps suboptimal for evaluating the precision and conservatism of the bootstrap methods. Finally, the argument used to justify nesting λ -selection within bootstrapping was intuitive. A more rigorous argument might be able to identify specific conditions on $\mathbf{X}$ or $\mathbf{y}$ that would lead to Method 3 consistently outperforming Method 2, or vice-versa. On the other hand, simulations could be used to estimate the components of bootstrapped variance due to the individual terms in Equation (2).

Vector bootstrapping CIs of Lasso coefficients led to increased variable selection precision, especially at small effect sizes. Our illustration with empirical data showed that this is also an effective approach to select important interactions between predictors. In consequence, vector bootstrapping CIs is a very promising approach for identifying sets of SNP-SNP and SNP-environment interactions.

References

- Abdellaoui A, Hottenga JJ, de Knijff P, Nivard MG, Xiao X, Scheet P, Brooks A, Ehli EA, Hu Y, Davies GE, et al. Population structure, migration, and diversifying selection in the Netherlands. *European Journal of Human Genetics*. 2013; 21(11):1277–1285. [PubMed: 23531865]
- Ayers KL, Cordell HJ. SNP Selection in genome-wide and candidate gene studies via penalized logistic regression. *Genetic epidemiology*. 2010; 34(8):879–891. [PubMed: 21104890]
- Balding DJ. A tutorial on statistical methods for population association studies. *Nature Reviews Genetics*. 2006; 7(10):781–791.
- Barr DR, Sherrill ET. Mean and variance of truncated normal distributions. *The American Statistician*. 1999; 53(4):357–361.
- Boomsma DI, de Geus EJC, Vink JM, Stubbe JH, Distel MA, Hottenga JJ, Posthuma D, Van Beijsterveldt TCEM, Hudziak JJ, Bartels M, Willemsen G. Netherlands twin register: From twins to twin families. *Twin Research and Human Genetics*. 2006; 9(6):849–857. [PubMed: 17254420]
- Buckland ST, Burnham KP, Augustin NH. Model selection: an integral part of inference. *Biometrics*. 1997:603–618.
- Bühlmann, PL.; van de Geer, SA.; Van de Geer, S. *Statistics for high-dimensional data*. Springer; 2011.
- Bühlmann P, Meier L, van de Geer S. Discussion: “A significance test for the lasso”. *Ann Statist*. 2014; 42(2):469–477.
- Camponovo, L. *Social Science Research Network Working Paper Series*. 2014. On the Validity of the Pairs Bootstrap for Lasso Estimators.
- Chang CC, Chow CC, Tellier LC, Vattikuti S, Purcell SM, Lee JJ. Second-generation plink: rising to the challenge of larger and richer datasets. 2014 arXiv preprint arXiv:1410.4803.
- Chatfield C. Model uncertainty, data mining and statistical inference. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*. 1995; 158(3):419–466.
- Chatterjee A. Bootstrapping lasso estimators. *Journal of the American Statistical Association*. 2011; 106(494):608–625.
- Cho S, Kim K, Kim YJ, Lee JK, Cho YS, Lee JY, Han BG, Kim H, Ott J, Park T. Joint identification of multiple genetic variants via elastic-net variable selection in a genome-wide association analysis. *Annals of human genetics*. 2010; 74(5):416–428. [PubMed: 20642809]
- Cordell HJ. Detecting gene–gene interactions that underlie human diseases. *Nature Reviews Genetics*. 2009; 10(6):392–404.
- D’Angelo G, Rao DC, Gu CC. Combining least absolute shrinkage and selection operator (LASSO) and principal-components analysis for detection of gene-gene interactions in genome-wide association studies. *BMC Proceedings*. 2009; 3(Suppl 7):S62. [PubMed: 20018056]
- Devlin B, Roeder K, Wasserman L. Analysis of multilocus models of association. *Genetic Epidemiology*. 2003; 25(1):36–47. [PubMed: 12813725]
- Efron B, Gong G. A leisurely look at the bootstrap, the jackknife, and Cross-Validation. *The American Statistician*. 1983; 37(1):36–48.
- Efron B, Hastie T, Johnstone L, Tibshirani R. Least angle regression. *Annals of Statistics*. 2004; 32:407–499.
- Efron, B.; Tibshirani, RJ. *An Introduction to the Bootstrap* (Chapman & Hall/CRC Monographs on Statistics & Applied Probability). Chapman and Hall/CRC; 1994.
- Fan J, Guo S, Hao N. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2012; 74(1):37–65. [PubMed: 22312234]

- Fawcett T. An introduction to roc analysis. *Pattern recognition letters*. 2006; 27(8):861–874.
- Freedman D, Lane D. A nonstochastic interpretation of reported significance levels. *Journal of Business & Economic Statistics*. 1983; 1(4):292–298.
- Friedman J, Hastie T, Höfling H, Tibshirani R. Pathwise coordinate optimization. *Annals of Applied Statistics*. 2007; 1(2):302–332.
- Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*. 2010; 33(1):1. [PubMed: 20808728]
- He Q, Lin DYY. A variable selection method for genome-wide association studies. *Bioinformatics (Oxford, England)*. 2011; 27(1):1–8.
- James GM, Radchenko P. A generalized dantzig selector with shrinkage tuning. *Biometrika*. 2009; 96(2):323–337.
- Knight K, Fu W. Asymptotics for Lasso-Type estimators. *The Annals of Statistics*. 2000; 28(5):1356–1378.
- Kyung M, Gill J, Ghosh M, Casella G. Penalized regression, standard errors, and bayesian lassos. *Bayesian Analysis*. 2010; 5:369–412.
- Lange K, Papp JC, Sinsheimer JS, Sobel EM. Next-generation statistical genetics: Modeling, penalization, and optimization in high-dimensional data. *Annual review of statistics and its applications*. 2014; 1:279–300.
- Leng C, Lin Y, Wahba G. A note on the lasso and related procedures in model selection. *Statistica Sinica*. 2006; 16(4):1273.
- Lockhart R, Taylor J, Tibshirani RJ, Tibshirani R. A significance test for the lasso. 2013 arXiv preprint arXiv:1301.7161.
- Lubke G, Laurin C, Amin N, Hottenga J, Willemsen G, van Grootheest G, Abdel-laoui A, Karssen L, Oostra B, van Duijn C, et al. Genome-wide analyses of borderline personality features. *Molecular psychiatry*. 2014; 19(8):923–929. [PubMed: 23979607]
- Meinshausen N, Bühlmann P. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2010; 72(4):417–473.
- Minnier J, Tian L, Cai T. A perturbation method for inference on regularized regression estimates. *Journal of the American Statistical Association*. 2011; 106(496):1371–1382. [PubMed: 22844171]
- Morey, LC. *Personality Assessment Inventory: Professional manual*. Psychological Assessment Resources; Odessa, FL: 1991.
- Motyer A, McKendry C, Galbraith S, Wilson S. LASSO model selection with post-processing for a genome-wide association study data set. *BMC Proceedings*. 2011; 5(Suppl 9):S24.
- Okser S, Pahikkala T, Airola A, Salakoski T, Ripatti S, Aittokallio T. Regularized machine learning in the genetic prediction of complex traits. *PLoS Genet*. 2014; 10(11):e1004754. [PubMed: 25393026]
- Osborne MR, Presnell B, Turlach BA. On the LASSO and its dual. *Journal of Computational and Graphical Statistics*. 2000; 9(2):319–337.
- Park JHH, Gail MH, Weinberg CR, Carroll RJ, Chung CC, Wang Z, Chanock SJ, Fraumeni JF, Chatterjee N. Distribution of allele frequencies and effect sizes and their interrelationships for common genetic susceptibility variants. *Proceedings of the National Academy of Sciences of the United States of America*. 2011; 108(44):18026–18031. [PubMed: 22003128]
- Pötscher BM, Leeb H. On the distribution of penalized maximum likelihood estimators: The LASSO, SCAD, and thresholding. *Journal of Multivariate Analysis*. 2009; 100(9):2065–2082.
- Price AL, Zaitlen NA, Reich D, Patterson N. New approaches to population stratification in genome-wide association studies. *Nature Reviews Genetics*. 2010; 11(7):459–463.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; Vienna, Austria: 2013.
- Sartori, S. PhD thesis. Università Degli Studi di Milano; 2009. Penalized Regression: bootstrap confidence intervals and variable selection for high dimensional data sets.
- Sivakumaran S, Agakov F, Theodoratou E, Prendergast JG, Zgaga L, Manolio T, Rudan I, McKeigue P, Wilson JF, Campbell H. Abundant pleiotropy in human complex diseases and traits. *The American Journal of Human Genetics*. 2011; 89(5):607–618. [PubMed: 22077970]

- Stefansson H, Ophoff RA, Steinberg S, Andreassen OA, Cichon S, Rujescu D, Werge T, Pietiläinen OP, Mors O, Mortensen PB, et al. Common variants conferring risk of schizophrenia. *Nature*. 2009; 460(7256):744–747. [PubMed: 19571808]
- Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1996; 58(1):267–288.
- Tibshirani R. Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2011; 73(3):273–282.
- Tibshirani RJ. The lasso problem and uniqueness. *Electronic Journal of Statistics*. 2013; 7(0):1456–1490.
- Valdar W, Sabourin J, Nobel A, Holmes CC. Reprioritizing genetic associations in hit regions using LASSO-based resample model averaging. *Genetic Epidemiology*. 2012; 36(5):451–462. [PubMed: 22549815]
- Waldron L, Pintilie M, Tsao MS, Shepherd FA, Huttenhower C, Jurisica I. Optimized application of penalized regression methods to diverse genomic data. *Bioinformatics*. 2011; 27(24):3399–3406. [PubMed: 22156367]
- Warren H, Casas JP, Hingorani A, Dudbridge F, Whittaker J. Genetic prediction of quantitative lipid traits: comparing shrinkage models to gene scores. *Genetic epidemiology*. 2014; 38(1):72–83. [PubMed: 24272946]
- Willemsen G, Vink JM, Abdellaoui A, den Braber A, van Beek JH, Draisma HH, van Dongen J, van't Ent D, Geels LM, van Lien R, et al. The adult netherlands twin register: Twenty-five years of survey and biological data collection. *Twin Research and Human Genetics*. 2013; 16(01):271–281. [PubMed: 23298648]
- Zhang P. Model selection via multifold cross validation. *The Annals of Statistics*. 1993; 21(1):299–313.

Given N individuals measured on p standardized predictors x_j , $j = 1, \dots, p$, with outcome y :

1. Fit Lasso model to select predictors
    - (a) Identify candidate λ values using sample covariances
    - (b) Use K -fold cross-validation ($K = 10$) to select finalized λ
  2. Finalized λ selects predictors, indexes them in $\mathbf{s} \subset \{1, \dots, j\}$
  3. Exclude predictors x_k , where $k \notin \mathbf{s}$
-

Fig. 1.

Method 1—Default approach to Lasso variable selection using K -fold cross-validation

-
1. Fit Lasso models, select set of predictors $\mathbf{s}$ as in Method 1
  2. Estimate marginal sampling distributions $\hat{F}^*(\hat{\beta}_k)$ for $k \in \mathbf{s}$
    - (a) Draw B vector bootstrap samples ($B \geq 1000$)
    - (b) In each bootstrap sample:
      - i. Fit a Lasso model with λ at finalized value from Step 1 (Fixed λ)
  3. Use $\hat{F}^*(\hat{\beta}_k)$ to calculate mean, SE, and CI for each $\hat{\beta}_k$
  4. Exclude x_k s:
    - if CI for $\hat{\beta}_k$ includes 0
    - or
    - if $C_{\text{var}}^* = \left(t_k^*\right)^{-1} = \frac{\text{SE}^*(\hat{\beta}_k)}{|\text{Mean}^*(\hat{\beta}_k)|} > 0.5$ (or other cutoff)
-

Fig. 2.
Method 2–Vector bootstrap for improved Lasso variable selection precision with λ treated as fixed

-
1. Fit Lasso models, select set of predictors $\mathbf{s}$ as in Method 1
  2. Estimate marginal sampling distributions $\hat{F}^*(\hat{\beta}_k)$ for $k \in \mathbf{s}$
    - (a) Draw B vector bootstrap samples ($B \geq 1000$)
    - (b) In each bootstrap sample:
      - i. Fit a Lasso model exactly as in Step 1 (Random λ)
  3. Use $\hat{F}^*(\hat{\beta}_k)$ to calculate mean, SE, and CI for each $\hat{\beta}_k$
  4. Exclude x_k s using rules listed in Method ??
-

Fig. 3.
Method 3–Vector bootstrapping with metaparameter λ treated as random

Figure 4. Bootstrap means and SEs of 77 SNPs, 2926 pairwise interactions; light gray diamonds represent predictors selected without bootstrapping. The lines represent $\text{mean} = \pm 2 \times \widehat{SE}^*$: points and diamonds outside the V-shape (i.e. in lower corners) are promising signals. A cube-root transformation was used on both axes.

Table 1
Variable selection precisions and false negative rates of default lasso vs. bootstrapped percentile CI

The random- λ bootstrap CI (Method 3) is more precise than the default at each effect size, but improvement is marginal for the very smallest effects. In all cases, the default approach showed low false-negative rates.

	$R^2 \backslash$ Method	Default	Fixed λ	Random λ
VSP	Ranks	0.3646	0.9529	0.8956
	0.01	0.3720	0.8304	0.8528
	0.0033	0.2049	0.1536	0.3372
	0.001	0.0503	0.0096	0.0507
	Polygenic	0.0096	0.1154	0.0990
FNR	Ranks	0.2008	0.5784	0.4594
	0.01	0.0084	0.1500	0.0408
	0.0033	0.3776	0.8420	0.6088
	0.001	0.7928	0.9888	0.9404
	Polygenic	0.1343	0.8804	0.8372

Bold text indicates the largest VSP/smallest FNR in each row. Nominal $\alpha = 0.05$.

Table 2
Median and MAD of ranks of 5 important predictors out of 104 total

Bootstrapping resulted in no improvement of predictor ranks. Both bootstrapping methods (2 and 3) showed identical performance.

R^2 \ Method	Default (1)	Bootstrapped (C_{var}^*) (2 and 3)
0.01	1 [0]	1 [0]
0.0067	2 [0]	1 [0]
0.0033	3 [0]	3 [0]
0.0022	4 [1.48]	4 [1.48]
0.001	5 [2.97]	6 [4.48]

Table 3
Standard deviations and confidence interval lengths of Lasso estimates of important predictors

Measures of the variability of coefficient estimates for important predictors, averaged over MC replications, showed that Method 3 had increased variability compared to Method 2.

Model	Method	SD	qCI Length	ntCI Length
Ranks	2 (Fixed λ)	0.0099	0.0351	0.0405
	3 (Nested λ)	0.0177	0.0693	0.0692
Single/ $R^2=0.0033$	2 (Fixed λ)	0.00637	0.02060	0.02500
	3 (Nested λ)	0.0167	0.0650	0.0656
polygenic	2 (Fixed λ)	0.0172	0.0604	0.0674
	3 (Nested λ)	0.0239	0.0860	0.0857

SD = standard deviation; qCI = bootstrap quantile confidence interval; ntCI = bootstrapped normal approximation confidence interval; Nominal α = 0.05

Table 4
Five-number summaries of the distribution of λ values in different simulations

Summaries of the distribution of λ values suggest that Method 3 (selection nested within bootstrap samples) produces a less-variable, left-shifted distribution.

Model	Method	Min	1Q	Med	3Q	Max
Ranks	1 (CV only)	0.0154	0.0282	0.0324	0.0367	0.0582
	3 (Nested λ)	0.00753	0.01410	0.01640	0.01910	0.04030
Single/ $R^2=0.0033$	1 (CV only)	0.0196	0.0365	0.0446	0.0505	0.0685
	3 (Nested λ)	0.00812	0.01540	0.01820	0.02190	0.06020
Polygenic	1 (CV only)	0.00796	0.01320	0.01490	0.01680	0.03180
	3 (Nested λ)	0.00182	0.00463	0.00543	0.00625	0.00968

Table 5

Promising SNP-SNP interactions

Chrs	rsIDs	MAFs	C_{var}^*	CI
16, 1	rs118160379×rs59194015	.05, .27	.49	[-.078, -.002]
<i>Selected by default but rejected by bootstrap:</i>				
Chrs	rsIDs	MAFs	C_{var}^*	CI
9, 4	rs112188788×rs139344595	.02, .01	.73	[0, .501]
6, 1	rs117666484×rs73008417	.01, .01	.94	[0, .562]
9, 1	rs112188788×rs73008417	.02, .01	1.17	[-2.904, 0]
12, 9	rs117256451×rs112188788	.02, .02	1.57	[-.056, .451]