Published in final edited form as:

Stat Sin. 2015 July 1; 25(3): 1185–1206. doi:10.5705/ss.2013.284.

A Small-Sample Choice of the Tuning Parameter in Ridge Regression

Philip S. Boonstra, Bhramar Mukherjee, and Jeremy M. G. Taylor

Department of Biostatistics, University of Michigan, Ann Arbor 48109

Abstract

We propose new approaches for choosing the shrinkage parameter in ridge regression, a penalized likelihood method for regularizing linear regression coefficients, when the number of observations is small relative to the number of parameters. Existing methods may lead to extreme choices of this parameter, which will either not shrink the coefficients enough or shrink them by too much. Within this “small- n , large- p ” context, we suggest a correction to the common generalized cross-validation (GCV) method that preserves the asymptotic optimality of the original GCV. We also introduce the notion of a “hyperpenalty”, which shrinks the shrinkage parameter itself, and make a specific recommendation regarding the choice of hyperpenalty that empirically works well in a broad range of scenarios. A simple algorithm jointly estimates the shrinkage parameter and regression coefficients in the hyperpenalized likelihood. In a comprehensive simulation study of small-sample scenarios, our proposed approaches offer superior prediction over nine other existing methods.

Keywords

Akaike’s information criterion; Cross-validation; Generalized cross-validation; Hyperpenalty; Marginal likelihood; Penalized likelihood

1 Introduction

Suppose we have data, $\{\mathbf{y}, \mathbf{x}\}$, which are n observations of a continuous outcome Y and p covariates $\mathbf{X}$, with the covariate matrix $\mathbf{x}$ regarded as fixed. n is small relative to p . We relate Y and $\mathbf{X}$ by a linear model, $Y = \beta_0 + \mathbf{X}^\top \boldsymbol{\beta} + \sigma \varepsilon$, with $\varepsilon \sim N\{0, 1\}$. Up to an additive constant, the log-likelihood is

$$\ell(\boldsymbol{\beta}, \beta_0, \sigma^2) = -\frac{n}{2} \ln(\sigma^2) - \frac{1}{2\sigma^2} (\mathbf{y} - \beta_0 \mathbf{1}_n - \mathbf{x}\boldsymbol{\beta})^\top (\mathbf{y} - \beta_0 \mathbf{1}_n - \mathbf{x}\boldsymbol{\beta}). \quad (1)$$

We center $\mathbf{y}$ and standardize $\mathbf{x}$ to have unit variance. As a consequence of this, although β_0 is estimated in the fitted model, our notation will implicitly reflect the assumption $\beta_0 = 0$.

We consider penalized estimation of $\boldsymbol{\beta}$, with our primary interest being prediction of future observations, rather than variable selection. Thus, we focus on L_2 -penalization, ie ridge regression (Hoerl and Kennard, 1970), which, from a prediction perspective, has favorable properties even compared to more recently developed penalization methods (eg Frank and

Friedman, 1993; Tibshirani, 1996; Fu, 1998; Zou and Hastie, 2005). Ridge regression may be viewed as a hierarchical linear model, similar to mixed effects modeling. Here the “random effects” are the elements of β . An L_2 -penalty on β implicitly assumes these are jointly and independently Normal with mean zero and variance σ^2/λ , because the penalty term matches the negative Normal log-density, up to a normalizing constant not depending on β .

$$p_\lambda(\beta, \sigma^2) = -\frac{\lambda}{2\sigma^2} \beta^\top \beta - \frac{p}{2} \ln(\lambda) + \frac{p}{2} \ln(\sigma^2). \quad (2)$$

The scalar λ is the ridge parameter, controlling the shrinkage of β toward zero; larger values yield greater shrinkage. Given λ , the maximum penalized likelihood estimate of β is

$$\beta_\lambda = \arg\max_{\beta|\lambda} \left\{ \ell(\beta, \sigma^2) - p_\lambda(\beta, \sigma^2) \right\} = (x^\top x + \lambda I_p)^{-1} x^\top y. \quad (3)$$

When $n - 1 > p$, a key result from Hoerl and Kennard (Theorem 4.3, 1970) is that $\lambda^* = \arg \min_{\lambda} \text{E}[(\beta - \beta_\lambda)^\top (\beta - \beta_\lambda)] > 0$, i.e. there exists $\lambda^* > 0$ for which the mean squared error (MSE) of β_λ decreases relative to $\lambda = 0$. If $x^\top x/n = I_p$, then $\lambda^* = p\sigma^2/\beta^\top \beta$; however, there is no closed-form solution for λ^* in the general $x^\top x$ case. A strictly positive λ introduces bias in β_λ but decreases variance, making a bias-variance tradeoff. A choice of λ which is too small leads to overfitting the data, and one which is too large shrinks β by too much. To contrast these extremes, we will hereafter refer to this latter scenario as “underfitting.” The existence of λ^* is relevant because prediction error, $\text{E}[(\beta - \beta_\lambda)^\top x^\top x (\beta - \beta_\lambda)]$, is closely related to MSE and may correspondingly benefit from such a bias-variance tradeoff.

To approximate λ^* , one cannot simply maximize $\ell(\beta, \sigma^2) - p_\lambda(\beta, \sigma^2)$ jointly with respect to β , σ^2 and λ , because the expression can be made arbitrarily large by plugging in $\beta = 0$ and letting $\lambda \rightarrow \infty$. Typically, λ is selected by optimizing some other objective function. Our motivation for this paper is to investigate selection strategies for λ when n is “small”, by which we informally mean $n < p$ or $n \approx p$, the complement being a more standard $n \gg p$ situation. This small- n situation increasingly occurs in modern genomic studies, whereas common approaches for selecting λ are often justified asymptotically in n .

Our contribution is two-fold. First, we present new ideas for choosing λ , including both a small-sample modification to a common existing approach and novel proposals. Our framework categorizes existing strategies into two classes, based on whether a goodness-of-fit criterion or a likelihood is optimized. Methods in either class may be susceptible to over- or underfitting; a third, new class extends the hierarchical perspective of ridge regression, the first level being $\ell(\beta, \sigma^2)$ and the second $p_\lambda(\beta, \sigma^2)$. Following ideas by Takada (1979), who showed that Stein’s Positive Part Estimator corresponds to a posterior mode given a certain prior, and, more recently, Strawderman and Wells (2012), who place a hyperprior on the Lasso penalty parameter, we add a third level, defining a “hyperpenalty” on λ . This hyper-penalty induces shrinkage on λ itself, thereby protecting against extreme choices of λ .

The second contribution follows naturally, namely, a comprehensive evaluation of all methods, both existing and newly proposed, in this small- n situation via simulation studies.

The remainder of this paper is organized as follows. We review current approaches for choosing λ (the first and second classes discussed above) in Sections 2 and 3 and propose a small-sample modification to one of these methods, generalized cross-validation (GCV, Craven and Wahba, 1979). In Section 4, we define a generic hyperpenalty function and explore a specific choice for the form of hyperpenalty in 4.1. Section 5 conducts a comprehensive simulation study. Our results suggest that the existing approaches for choosing λ can be improved upon in many small- n cases. Section 7 concludes with a discussion, in which we discuss useful extensions of the hyperpenalty framework.

2 Goodness-of-fit-based methods for selection of λ

These methods define an objective function in terms of λ which is to be minimized. Commonly used is K -fold cross-validation, which partitions observations into K groups, $\kappa(1), \dots, \kappa(K)$, and calculates β_λ K times using equation (3), each time leaving out group $\kappa(i)$, to get $\beta_\lambda^{-k(1)}, \beta_\lambda^{-k(2)}$, etc. For $\beta_\lambda^{-k(i)}$, cross-validated residuals are calculated on the observations in $\kappa(i)$, which did not contribute to estimating β . The objective function estimates prediction error and is the sum of the squared cross-validated residuals:

$$\lambda_{K-CV} = \arg\min_{\lambda} \ln \sum_{i=1}^K \left(y_{k(i)} - x_{k(i)} \beta_\lambda^{-k(i)} \right)^\top \left(y_{k(i)} - x_{k(i)} \beta_\lambda^{-k(i)} \right). \quad (4)$$

A suggested choice for K is 5 (Hastie et al., 2009). When $K = n$, some simplification (Golub et al., 1979) gives

$$\lambda_{n-CV} = \arg\min_{\lambda} \ln \sum_{i=1}^n \left(Y_i - \mathbf{X}_i^\top \beta_\lambda \right)^2 / \left(1 - P_{\lambda[ii]} - 1/\eta \right)^2 \quad (5)$$

$$\text{with } \mathbf{P}_\lambda = \mathbf{x} \left(\mathbf{x}^\top \mathbf{x} + \lambda \mathbf{I}_p \right)^{-1} \mathbf{x}. \quad (6)$$

$P_{\lambda[ii]}$ is the i th diagonal element of $\mathbf{P}_\lambda$ and measures the i th observation's influence in estimating β . Further discussion of its interpretation is given in Section 2.1. From (5), observations for which $P_{\lambda[ii]}$ is large, ie influential observations, have greater weight. Re-centering $\mathbf{y}$ at each fold implies β_0 is re-estimated; this is reflected by the “ $-1/n$ ” term in (5). This term does not appear in the derivations by Golub et al. (1979), which assume β_0 is known, but this difference in assumptions is important with regard to GCV, which is discussed next, and our proposed extension of GCV.

GCV multiplies each squared residual in (5) by $(1 - P_{\lambda[ii]} - 1/n)^2 / (1 - \text{Trace}(\mathbf{P}_\lambda)/n - 1/n)^2$, thereby giving equal weight to all observations. Using the equality $\mathbf{y} - \mathbf{x}\beta_\lambda = (\mathbf{I}_n - \mathbf{P}_\lambda)\mathbf{y}$, further simplification yields

$$\lambda_{\text{GCV}} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} - 2 \ln (1 - \text{Trace}(\mathbf{P}_{\lambda}) / n - 1/n) \right\}. \quad (7)$$

Although derived using different principles, other methods reduce to a “model fit + penalty” or “model fit + model complexity” form similar to (7): Akaike’s Information Criterion (aic, Akaike, 1973) and the Bayesian Information Criterion (BIC, Schwarz, 1978). Respectively, each chooses λ as follows:

$$\lambda_{\text{AIC}} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} + 2 (\text{Trace}(\mathbf{P}_{\lambda}) + 2) / n \right\}, \quad (8)$$

$$\lambda_{\text{BIC}} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} + \ln((n - \text{Trace}(\mathbf{P}_{\lambda}) + 2)/n) \right\}. \quad (9)$$

Asymptotically in n , GCV will choose the value of λ which minimizes the prediction criterion $E \left[(\boldsymbol{\beta} - \boldsymbol{\beta}_{\lambda})^{\top} \mathbf{x}^{\top} \mathbf{x} (\boldsymbol{\beta} - \boldsymbol{\beta}_{\lambda}) \right]$ (Golub et al., 1979; Li, 1986). Further, Golub et al. observe that GCV and AIC asymptotically coincide. BIC asymptotically selects the true underlying model from a set of nested candidate models (Sin and White, 1996; Hastie et al., 2009), so its justification for use in selecting λ , which is a shrinkage parameter, is weak. For all of these methods, optimality is based upon the assumption that $n \gg p$. When n is small, extreme overfitting is possible (Wahba and Wang, 1995; Efron, 2001), giving small bias/large variance estimates. A small-sample correction of aic (AIC_C , Hurvich and Tsai, 1989; Hurvich et al., 1998) and a robust version of GCV (RGCV_{γ} , Lukas, 2006) exist:

$$\lambda_{\text{AIC}_C} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} + 2 (\text{Trace}(\mathbf{P}_{\lambda}) + 2) / (n - \text{Trace}(\mathbf{P}_{\lambda}) - 3) \right\}, \quad (10)$$

$$\lambda_{\text{RGCV}_{\gamma}} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} - 2 \ln (1 - \text{Trace}(\mathbf{P}_{\lambda}) / (n - 1/n)) + \ln (\gamma + (1 - \gamma) \text{Trace}(\mathbf{P}_{\lambda}^2) / n) \right\}. \quad (11)$$

For AIC_C , the modified penalty is the product of the original penalty, $2(\text{Trace}(\mathbf{P}_{\lambda}) + 2)/n$, and $n/(n - \text{Trace}(\mathbf{P}_{\lambda}) - 3)$. The authors do not consider the possibility of $n - \text{Trace}(\mathbf{P}_{\lambda}) - 3 < 0$, which would inappropriately change the sign of the penalty, and we have found no discussion of this in the literature. In our implementation of AIC_C , we replace $n - \text{Trace}(\mathbf{P}_{\lambda}) - 3$ with its positive part, $(n - \text{Trace}(\mathbf{P}_{\lambda}) - 3)_+$, effectively making the criterion infinitely large in this case. As a rule of thumb, Burnham and Anderson (2002) suggest to use AIC_C over AIC when $n < 40p$ (their threshold for small n) and thus also when $n \approx p$. RGCV_{γ} subtracts another penalty from GCV based on a tuning parameter $\gamma \in (0, 1]$, as in (11); we use $\gamma = 0.3$ based on Lukas’ recommendation. Small choices of λ are more severely penalized, thereby offering protection against overfitting. To the best of our knowledge, the performance of AIC_C or RGCV_{γ} in the context of selecting λ in ridge regression has not been extensively studied.

2.1 Small-sample GCV

$\text{Trace}(\mathbf{P}_\lambda)$, with $\mathbf{P}_\lambda$ defined in (6), is the effective number of model parameters, excluding β_0 and σ^2 . It decreases monotonically with $\lambda > 0$ and lies in the interval $(0, \min\{n-1, p\})$. The upper bound on $\text{Trace}(\mathbf{P}_\lambda)$ is not $\min\{n, p\}$ because the standardization of $\mathbf{x}$ reduces its rank by one when $n = p$. Although the parameters β_0 and σ^2 are counted (in the literal sense) in the model complexity terms of aic and BIC, they have only an additive effect, being represented by the “+2” expressions in (8) and (9). For this reason, β_0 and σ^2 may be ignored in considering model complexity. However, from (7), GCV counts β_0 , which is given by the “ $-1/n$ ” term, but not σ^2 ; counting both *will* change the penalty, since the model complexity term is on the log-scale. This motivates our proposed small-sample correction to GCV, called GCV_C , which *does* count σ^2 as a parameter:

$$\lambda_{\text{GCV}_C} = \arg\min_{\lambda} \left\{ \ln \mathbf{y}^\top (\mathbf{I}_n - \mathbf{P}_\lambda)^2 \mathbf{y} - 2 \ln((1 - \text{Trace}(\mathbf{P}_\lambda)/n - 2/n)_+) \right\}. \quad (12)$$

As with AIC_C , $1 - \text{Trace}(\mathbf{P}_\lambda)/n - 2/n$ may be negative. In this case, subtracting the log of the positive part of $1 - \text{Trace}(\mathbf{P}_\lambda)/n - 2/n$ makes the objective function infinite. This is only a small-sample correction because the objective functions in (7) and (12) coincide as $n \rightarrow \infty$, and the asymptotic optimality of GCV transfers to GCV_C .

An explanation of why GCV_C corrects the small-sample deficiency of GCV is as follows. If $n-1 = p$, the model-fit term in the objective function of (7), $\ln \mathbf{y}^\top (\mathbf{I}_n - \mathbf{P}_\lambda)^2 \mathbf{y}$, tends to $-\infty$ as λ decreases. When $\lambda = 0$, the fitted values, $\mathbf{P}_\lambda \mathbf{y}$, will perfectly match the observations, $\mathbf{y}$, and the data are overfit. The penalty term, $-2 \ln(1 - \text{Trace}(\mathbf{P}_\lambda)/n - 1/n)$, tends to ∞ as λ decreases, because $\text{Trace}(\mathbf{P}_\lambda)$ approaches $n-1$. The rates of convergence for the model-fit and penalty terms determine whether GCV chooses a too-small λ . If the model-fit term approaches $-\infty$ faster than the penalty approaches ∞ , the objective function is minimized by setting λ as small as possible, which is $\lambda = 0$ when $n-1 = p$. Although this phenomenon is most striking in cases for which $n-1 = p$, as we will see in Section 5, this finding appears to hold when $n-1 < p$. In this case, predictions will *nearly* match observations as λ decreases but remains numerically positive to allow for the matrix inversion in $\mathbf{P}_\lambda$, and the penalty term still approaches ∞ as λ decreases. Like GCV, the penalty function associated with GCV_C also approaches ∞ as λ decreases. In contrast to GCV, however, the GCV_C penalty equals ∞ when $\lambda = \tilde{\lambda} > 0$, where $\tilde{\lambda}$ is the solution to $1 - \text{Trace}(\mathbf{P}_\lambda)/n - 2/n = 0$, or, equivalently, $\text{Trace}(\mathbf{P}_\lambda) = n-2$. In other words, when fitting GCV_C , the effective number of remaining parameters, beyond σ^2 and β_0 , will be less than $n-2$, and perfect fit of the observations to the predictions, i.e. $\lambda = 0$, cannot occur.

Remark 1—A reviewer observed that the GCV_C penalty can be generalized according to $-2 \ln((1 - \text{Trace}(\mathbf{P}_\lambda)/n - c/n)_+)$ for $c \geq 1$; special cases of this include GCV ($c = 1$) and GCV_C ($c = 2$). Extending the interpretation given above, this ensures that the effective number of remaining parameters, beyond σ^2 and β_0 , will be less than $n - c$, rather than $n - 2$. Preliminary results from allowing c to vary did not point to a uniformly better choice of c .

2. Also, using $c = 2$ is consistent with our original motivation for proposing GCV_C , namely properly counting the model parameters.

3 Likelihood-based methods for selection of λ

A second approach treats the ridge penalty in (2) as a negative log-density. One can consider a marginal likelihood, where λ is interpreted as the variance component of a mixed-effects model:

$$m(\lambda, \sigma^2) - \ln \int_{\beta} \exp \left\{ \ell(\beta, \sigma^2) - p_{\lambda}(\beta, \sigma^2) \right\} d\beta \\ = -\frac{1}{2\sigma^2} \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda}) \mathbf{y} - \frac{n}{2} \ln(\sigma^2) + \frac{1}{2} \ln |\mathbf{I}_n - \mathbf{P}_{\lambda}|. \quad (13)$$

From this, $\mathbf{y}/\lambda$, σ^2 is multivariate Normal with mean 0_n ($\mathbf{y}$ is centered) and covariance $\sigma^2(\mathbf{I}_n - \mathbf{P}_{\lambda})^{-1}$. The maximum profile marginal likelihood (MPML) estimate, originally proposed for smoothing splines (Wecker and Ansley, 1983), profiles $m(\lambda, \sigma^2)$ over σ^2 , replacing each instance with $\hat{\sigma}_{\lambda}^2 = \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda}) \mathbf{y} / n$, and maximizes the “concentrated” log-likelihood, $m(\lambda, \hat{\sigma}_{\lambda}^2)$:

$$\lambda_{\text{MPML}} = \arg \min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda}) \mathbf{y} - \frac{1}{n} \ln |\mathbf{I}_n - \mathbf{P}_{\lambda}| \right\}. \quad (14)$$

Closely related is the generalized/restricted MPML (GMPML, Harville, 1977; Wahba, 1985), which adjusts the penalty to account for estimation of regression parameters that are not marginalized. Here, only β_0 is not marginalized, so the adjustment is by one degree of freedom (see Supplement S1):

$$\lambda_{\text{GMPML}} = \arg \min_{\lambda} \left\{ \ln \mathbf{y}^{\top} (\mathbf{I}_n - \mathbf{P}_{\lambda}) \mathbf{y} - \frac{1}{n-1} \ln |\mathbf{I}_n - \mathbf{P}_{\lambda}| \right\}. \quad (15)$$

In a smoothing-spline comparison of GMPML to GCV, Wahba (1985) found mixed results, with neither method offering uniformly better predictions. For scatterplot smoothers, Efron (2001) notes that GMPML may oversmooth, yielding large bias/small variance estimates.

Remark 2

Rather than profiling over σ^2 , one could jointly maximize $m(\lambda, \sigma^2)$ over λ and σ^2 . We have not found this approach previously used as a selection criterion in ridge regression. Our initial investigation of this and its restricted likelihood counterpart gave results similar to MPML and GMPML, and so we do not consider it further.

An alternative to the marginal likelihood methods described above is to treat the objective function in (3) as an h -log-likelihood, or “ h -loglikelihood”, of the type proposed by Lee and Nelder (1996) for hierarchical generalized linear models. The link between penalized likelihoods, like ridge regression, and the h -loglikelihood was noted in the paper’s ensuing discussion. To estimate σ^2 (the dispersion) and λ (the variance component), Lee and Nelder suggested an iterative profiling approach, yielding the maximum adjusted profile h -loglikelihood (maphl) estimate. In Supplement S2, we show one iteration proceeds as follows:

$$\sigma^{2(i)} \leftarrow \frac{\left(\mathbf{y} - \mathbf{x}\boldsymbol{\beta}^{(i-1)}\right)^\top (\mathbf{y} - \mathbf{x}\boldsymbol{\beta}^{(i-1)}) + \lambda^{(i-1)} \boldsymbol{\beta}^{(i-1)\top} \boldsymbol{\beta}^{(i-1)}}{n-1} \quad (16)$$

$$\lambda^{(i)} \leftarrow \operatorname{argmin}_{\lambda} \left\{ \lambda \boldsymbol{\beta}^{(i-1)\top} \boldsymbol{\beta}^{(i-1)} / \sigma^{2(i)} - \ln |I_n - P_{\lambda}| \right\} \quad (17)$$

$$\boldsymbol{\beta}^{(i)} \leftarrow \boldsymbol{\beta}_{\lambda^{(i)}} \quad (18)$$

and $\lambda_{\text{MAPHL}} = \lambda^{(\infty)}$

Finally, Tran (2009) proposed the “Loss-rank” (LR) method for selecting λ . Its derivation, which we do not give, is likelihood-based, but the criterion resembles that of AIC in (8):

$$\lambda_{\text{LR}} = \operatorname{argmin}_{\lambda} \left\{ \ln \mathbf{y}^\top (\mathbf{I}_n - \mathbf{P}_{\lambda})^2 \mathbf{y} - \frac{2}{n} \ln |I_n - \mathbf{P}_{\lambda}| \right\}. \quad (19)$$

Tran also suggested a modified penalty term, which is dependent on $\mathbf{y}$, but this did not give appreciably different results from λ_{LR} in their study.

4 Maximization with Hyperpenalties

As noted previously, some existing methods may choose extreme values of λ , particularly when n is small, suggesting a need for a second level of shrinkage, that is, shrinkage of λ itself. We extend the hierarchical framework of (1) and (2) with a “hyperpenalty” on λ , $h(\lambda)$, which gives non-negligible support for λ over a finite range of values. The “hyperpenalized log-likelihood” is

$$hpl(\boldsymbol{\beta}, \lambda, \sigma^2) = \ell(\boldsymbol{\beta}, \sigma^2) - p_{\lambda}(\boldsymbol{\beta}, \sigma^2) - h(\lambda) - \ln(\sigma^2). \quad (20)$$

From the Bayesian perspective, when $h(\lambda)$ is in the form of a log-density, the hyperpenalty corresponds to a hyperprior on λ , and the hyperpenalized likelihood is the posterior (the expression $-\ln(\sigma^2)$ is the log-density of an improper prior on σ^2). In contrast to fully Bayesian methods, which characterize the entire posterior, we desire a single point estimate of $\boldsymbol{\beta}$, σ^2 and λ and focus on mode finding. Importantly, joint maximization with respect to $\boldsymbol{\beta}$, σ^2 and λ is now possible.

For a general $h(\lambda)$, we find the joint mode of (20):

$\{\hat{\boldsymbol{\beta}}, \hat{\lambda}, \hat{\sigma}^2\} \leftarrow \operatorname{argmax}_{\boldsymbol{\beta}, \lambda, \sigma^2} \{hpl(\boldsymbol{\beta}, \lambda, \sigma^2)\}$. Alternatively, $\{\hat{\boldsymbol{\beta}}, \hat{\lambda}, \hat{\sigma}^2\}$ may be calculated using conditional maximization steps:

$$\begin{aligned} \sigma^{2(i)} &\leftarrow \operatorname{argmax}_{\sigma^2} - \left\{ hpl(\boldsymbol{\beta}^{(i-1)}, \sigma^2, \lambda^{(i-1)}) \right\} \\ &= \frac{(\mathbf{y} - \mathbf{x}\boldsymbol{\beta}^{(i-1)})^\top (\mathbf{y} - \mathbf{x}\boldsymbol{\beta}^{(i-1)}) + \lambda^{(i-1)} \boldsymbol{\beta}^{(i-1)\top} \boldsymbol{\beta}^{(i-1)}}{n+p+2} \end{aligned} \quad (21)$$

$$\lambda^{(i)} \leftarrow \operatorname{argmax}_{\lambda} \left\{ hpl \left(\beta^{(i-1)}, \sigma^{2(i)}, \lambda \right) \right\} \quad (22)$$

$$\beta^{(i)} \leftarrow \operatorname{argmax}_{\beta} \left\{ hpl \left(\beta, \sigma^{2(i)}, \lambda^{(i)} \right) \right\} = \beta_{\lambda^{(i)}} \quad (23)$$

with $\{\hat{\beta}, \hat{\lambda}, \hat{\sigma}^2\} = \{\beta^{(\infty)}, \lambda^{(\infty)}, \sigma^{2(\infty)}\}$. The only step that depends on the choice $h(\lambda)$ is (22); the other steps are available in closed form regardless of $h(\lambda)$.

Based on the expression for $p_{\lambda}(\beta, \sigma^2)$ given in (2), if $\exp\{-h(\lambda)\} = o(\lambda^{-p/2})$, then, upon applying the maximization step in (22), $\lambda^{(i)}$ is guaranteed to be finite, regardless of the values of $\beta^{(i-1)}$ and $\sigma^{2(i)}$. This relates to an earlier comment in the Introduction that one cannot simply maximize $\ell(\beta, \sigma^2) - p_{\lambda}(\beta, \sigma^2)$ alone. For example, using $h(\lambda) = C$, C constant, would yield the same result as maximizing $\ell(\beta, \sigma^2) - p_{\lambda}(\beta, \sigma^2)$, namely an infinite hyperpenalized log-likelihood. We will propose one such choice of $h(\lambda)$ that satisfies $\exp\{-h(\lambda)\} = o(\lambda^{-p/2})$ and is empirically observed to work well for the ridge penalty; different choices of $h(\lambda)$ may be better suited for other penalty functions.

4.1 Choice of hyperpenalty

Crucial to this approach is the determination of an appropriate hyperpenalty and accompanying hyperparameters. Our recommended hyperpenalty is based on the gamma distribution, namely $h(\lambda) = -(a-1) \ln(\lambda) + \lambda/b$. From the Bayesian perspective, this is natural because it is conjugate to the precision of the Normal distribution, which is one possible interpretation of λ (e.g. Tipping, 2001; Armagan and Zaretzki, 2010). From Supplement S3, the update for λ given in (22) becomes

$$\lambda^{(i)} = \frac{p+2a-2}{\beta^{(i-1)\top} \beta^{(i-1)} / \sigma^{2(i)} + 2b}. \quad (24)$$

This additionally requires choosing values for a and b . We will do so by first choosing a desired prior mean for λ , given by a/b , and a value for $\lambda^{(i)}$ in (24), and then solving the two expressions for a and b . Necessary to this strategy is that the chosen value of $\lambda^{(i)}$ must result in a and b that are free of $\sigma^{2(i)}$ and $\beta^{(i-1)}$. To choose a/b , recall the key result from Hoerl and

Kennard (1970): when $\mathbf{x}^{\top} \mathbf{x} / n = \mathbf{I}_p$, $\lambda^* = \arg \min_{\lambda \geq 0} E \left[(\beta - \beta_{\lambda})^{\top} (\beta - \beta_{\lambda}) \right] = p\sigma^2 / \beta^{\top} \beta$. While not of immediate practical use, since σ^2 and β are the parameters to be estimated, we note that $\sigma^2 / \beta^{\top} \beta \approx (1/R^2 - 1)$, where $R^2 = \beta^{\top} \Sigma_X \beta / (\beta^{\top} \Sigma_X \beta + \sigma^2)$ is the coefficient of determination, and the approximation comes from substituting $\mathbf{x}^{\top} \mathbf{x} / n = \mathbf{I}_p$ for Σ_X . In contrast to σ^2 or the individual elements of β , there may be knowledge about R^2 . Alternatively, we will propose two strategies (Section 4.2) that provide an estimate of R^2 . Given an estimate or prior guess of $R^2 \in (0, 1)$, say $\hat{R}^2$, we set $a/b = p \left(1/\hat{R}^2 - 1 \right)$.

In addition to a sensible mean, it is important to have a and b be such that $\lambda^{(i)}$, which is the resulting update for λ , is not extreme. Let the update for λ given in (24) be $\lambda^{(i)} = (p-1)H^{(i)}$,

where $H^{(i)}$ is the harmonic mean of $\sigma^{2(i)}/\beta^{(i-1)T}\beta^{(i-1)}$ and $(1/\hat{R}^2 - 1)$. Being a harmonic mean, the $(1/\hat{R}^2 - 1)$ term, which will typically be less than 10 for most analyses, moderates potentially large values of $\sigma^{2(i)}/\beta^{(i-1)T}\beta^{(i-1)}$, thereby preventing underfitting. Simultaneously, $\lambda^{(i)}$ increases linearly with p , which prevents overfitting in $n < p$ scenarios.

Solving these expressions, $a/b = p(1/\hat{R}^2 - 1)$ and $\lambda^{(i)} = (p - 1)H^{(i)}$, yields $a = p/2$ and $b = (1/\hat{R}^2 - 1)^{-1/2}$. When the covariates are approximately uncorrelated and $\hat{R}^2$ is close to R^2 , a/b will be near λ^* . However, the uncertainty coming from the variance of λ makes this useful in the general $\mathbf{x}^T\mathbf{x}$ case, for which no closed-form solution of λ^* exists. As we will see in the simulation study, this holds true even when $\hat{R}^2 \neq R^2$.

Finally, it is important that the hyperpenalty strategy not be inferior in a standard regression, when $n \gg p$. To establish this, we derive in Supplement S4 a large- n approximation for λ^* in the general $\mathbf{x}^T\mathbf{x}$ case: $\lambda^* \approx \sigma^2 \text{Tr}[(\mathbf{x}^T\mathbf{x})^{-1}]/\beta^T(\mathbf{x}^T\mathbf{x})^{-1}\beta$. As $n \rightarrow \infty$, this converges in probability to $\sigma^2 \text{Tr}[\sum_x^{-1}]/\beta^T \sum_x^{-1} < \infty$, from which we can see that λ^* asymptotes in n . From (3), because the crossproduct term, $\mathbf{x}^T\mathbf{x}$, increases linearly in n , β_{λ^*} , the ridge estimate of β at the optimal value of λ , approaches the standard ordinary least squares (OLS) estimate, $\beta_{\lambda/\lambda=0}$. The $\mathbf{x}^T\mathbf{x}$ expression grows linearly in n , but λ^* asymptotes in n , so the effect of λ^* is reduced. The implication is that any choice of λ induces the same effective shrinkage for large n . The parameters a and b do not depend on n , and so the effect of the hyperpenalty will decrease with n , as desired.

Remark 3—Choosing a “flat” hyperpenalty, namely $h(\lambda) = \ln(\lambda)$, is untenable when $p > 2$, because it does not satisfy $\exp\{-h(\lambda)\} = o(\lambda^{-p/2})$. Specifically, plugging $h(\lambda)$ in (20), the expression $hp\ell(\beta, \lambda, \sigma^2)$ can be made arbitrarily large in λ by setting $\beta = 0_p$.

4.2 Estimating R^2

For analyses in which one is unable or unwilling to make a prior guess of R^2 to use as $\hat{R}^2$, here we describe a strategy to estimate R^2 . We note first that $R^2 = \text{Cor}(Y, \mathbf{X}^T\beta)^2$, where ‘Cor’ denotes the Pearson correlation. It is known that the empirical prediction error from using a vector of fitted values, $\mathbf{x}\hat{\beta}$, corresponding to the observed outcomes $\mathbf{y}$ will be

optimistic when $\hat{\beta}$ depends on $\mathbf{y}$. This means that the empirical R^2 , $\bar{R}^2 = \text{Cor}(\mathbf{y}, \mathbf{x}\hat{\beta})^2$, will also be optimistic, i.e. upwardly biased, for R^2 . In contrast, Efron (1983) showed how an estimate of prediction error using the bootstrap will be pessimistic. Applied to our context, given B bootstrapped datasets, the bootstrap-estimate of R^2 ,

$\hat{R}_{\text{boot}}^2 = (1/B) \sum_{b=1}^B \text{Cor}(\mathbf{y}_*(-b), \mathbf{x}_*(-b)\hat{\beta}_{*}^{(b)})^2$, where the $*(b)$ and $*(-b)$ notation indicate that the training and test datasets do not overlap, will be biased downward from R^2 . Efron (1983) suggested that a particular linear combination of the optimistic and pessimistic prediction error estimates would provide an approximately unbiased estimate of prediction error. Analogizing this to estimating R^2 , the linear combination is given by

$0.632\hat{R}_{\text{boot}}^2 + 0.368\bar{R}^2$. The weight is based on a bootstrapped dataset containing, on average,

about $e^{-1}n \approx 0.632n$ unique observations from the original dataset. When $n > p$, this “632-estimate” of R^2 , using, say, OLS to estimate β in each bootstrap, would provide a reasonable estimate of R^2 for our purposes. However, in the $n < p$ scenarios we are specifically interested in, OLS is not an option, and other methods, e.g. ridge regression, must be used to estimate β in each dataset. In addition, when p is large, the bootstrap may add a non-trivial computational component, if determining $\hat{\beta}_*^{(b)}$ is computationally expensive. So as to minimize any added burden due to estimating R^2 , which is only a preprocessing step before applying the hyperpenalty approach, we propose to modify the 632-estimate by replacing the bootstrap with 5-CV:

$$\hat{R}_{0.632}^2 = 0.632 \times (1/5) \sum_{i=1}^5 \text{C\hat{or}}\left(y_{k(i)}, x_{k(i)}\beta_{\lambda_{5-CV}}\right)^2 + 0.368 \times \text{C\hat{or}}\left(y, x\beta_{\lambda_{5-CV}}\right)^2, \quad (25)$$

where we use the same $\kappa(i)$ notation defined in Section 2. In words: first calculate λ_{5-CV} . Then, calculate a weighted average of the cross-validated empirical correlation and the standard empirical correlation over all the data. Let HYP_{632} denote the hyperpenalty approach using $\hat{R}_{0.632}^2$ as the estimate for R^2 .

5 Simulation Study

Our simulation study is designed to mimic the current reality of the “-omics” era in which many covariates are analyzed but few contribute substantial effect sizes. The relevant quantities are described as follows:

Covariates ($\mathbf{x}$). One simulated dataset consists of training and validation data generated from the same model. The dimension of the training data is $n \times p$, with $n \in \{25, 100, 250, 4000\}$ and $p \in \{100, 4000\}$. The $n \times p$ matrix $\mathbf{x}$ is drawn from $N_p\{\mathbf{0}_p, \Sigma_{\mathbf{x}}\}$. For the validation data, a $2000 \times p$ matrix $\mathbf{x}_{\text{new}}$ is sampled from this same distribution. We construct $\Sigma_{\mathbf{x}}$ according to “approximately uncorrelated” and “correlated” scenarios. The construction of $\Sigma_{\mathbf{x}}$ is described in Supplement S5. Briefly: we begin with a block-wise compound symmetric matrix with 10 blocks. Within blocks, the correlation is $\rho = 0$ (approximately uncorrelated) or $\rho = 0.4$ (correlated), and between blocks, there is zero correlation. We then stochastically perturb the matrix in such a way to generate $\Sigma_{\mathbf{x}}$, using algorithms by Hardin et al. (2013), as to maintain positive-definiteness but mask the underlying structure. This perturbation occurs at each simulation iterate, and, as we outline in the Supplement, it is less extreme for the approximately uncorrelated case, so that the resulting $\Sigma_{\mathbf{x}}$ is close to $\mathbf{I}_p$.

Parameters (β, σ^2). To better account for many plausible configurations of the coefficients, which would be difficult using a single, fixed choice of β , we specify a generating distribution, drawing β once per simulation iterate and making it common to all observations. We draw β from a mixture density. Let $Z_i \in \{1, 2, 3\}$, $i = 1, \dots, p$, be a random variable with $\Pr(Z_i = 1) = \Pr(Z_i = 2) = 0.005$. Then, construct β as follows:

$$\alpha_i | Z_i \stackrel{\text{ind}}{\sim} \begin{cases} t_3 \{ \sigma^2 = 1/3 \}, & Z_i = 1 \\ \text{Exp} \{ 1 \}, & Z_i = 2 \\ N \{ 0, \sigma^2 = 10^{-6} \}, & Z_i = 3 \end{cases} \quad (26)$$

$$\boldsymbol{\beta} = \boldsymbol{\alpha} \times AR_1(\pi), \quad (27)$$

where $AR_1(\pi)$ is a $p \times p$, first-order auto-regressive matrix with correlation coefficient $\pi \in \{0, 0.3\}$. The sampling density for $\boldsymbol{\alpha}$ is a mixture of scaled t_3 , Exponential, and Normal distributions, and 99% of the coefficients have small effect sizes, coming from the Normal component. The vector $\boldsymbol{\alpha}$ is scaled to obtain $\boldsymbol{\beta}$, which encourages neighboring coefficients to have similar effects, depending on π . So that some meaningful signal is present in every dataset, we ensure $\#\{i : Z_i = 3\} \geq 3$, regardless of p . Given $\boldsymbol{\beta}$, $\Sigma_{\mathbf{X}}$, and $R^2 \in \{0.05, 0.2, 0.4, 0.6, 0.8, 0.95\}$, we calculate $\sigma^2 = \boldsymbol{\beta}^\top \Sigma_{\mathbf{X}} \boldsymbol{\beta} / (R^2 - 1)$.

Outcomes ($\mathbf{y}/\mathbf{x}$) The outcomes from the training and validation data are, respectively, $\mathbf{y}/\mathbf{x}$, $\boldsymbol{\beta}$, σ^2 and $\mathbf{y}_{\text{new}}/\mathbf{x}_{\text{new}}$, $\boldsymbol{\beta}$, σ^2 . For each of the 192 combinations of π , p , n , uncorrelated/correlated $\mathbf{x}$, and R^2 , 10,000 (when $p = 100$) or 500 ($p = 4000$) training and validation datasets are sampled. We present results only for $\pi = 0.3$ and $R^2 \in \{0.2, 0.4, 0.8\}$ here, with the remainder being in Supplement S6.

We compare the 11 methods listed in Table 1. n -CV is left out, being approximated by GCV and computationally expensive. aic is replaced with its small-sample correction, AIC_C . The new methods considered are GCV_C and the hyperpenalty-approach, HYP_{632} . We also present HYP_{TRUE} , which is the hyperpenalty method using the true, unknown value of R^2 . The difference in prediction error between HYP_{632} and HYP_{TRUE} quantifies the possible gain, if any, from improving upon the “632” estimation scheme. All methods differ only in λ , which determines the estimate of $\boldsymbol{\beta}$ via (3). The criterion by which we evaluate methods on the validation data is relative MSPE, $\text{rMSPE}(\lambda)$:

$$\text{rMSPE}(\lambda) = 1000 \times (\text{MSPE}(\lambda) / \text{MSPE}(\lambda_{\text{opt}}) - 1), \quad (28)$$

where $\text{MSPE}(\lambda) = (\mathbf{y}_{\text{new}} - \mathbf{x}_{\text{new}}\boldsymbol{\beta}_\lambda)^\top (\mathbf{y}_{\text{new}} - \mathbf{x}_{\text{new}}\boldsymbol{\beta}_\lambda)$ and $\lambda_{\text{opt}} = \arg \min_\lambda \text{MSPE}(\lambda)$. Thus, rMSPE measures the percentage increase above the smallest possible MSPE, and $\text{rMSPE} = 0$ is ideal. Equivalently, rMSPE measures the inefficiency of each method. We used an iterative grid search to calculate λ_{opt} as well as λ for all methods except MAPHL and HYP_{632} , for which explicit maximization steps are available. Table 2 gives the average rMSPE ; values in boldface are the column-wise minima (excluding HYP_{TRUE}), and values with an asterisk are less than twice each column-wise minimum. Figure 1 compares the rMSPE of HYP_{632} to the median rMSPE of remaining methods, excluding GCV_C , as an overall performance comparison. Figure 2 compares the rMSPE of GCV to GCV_C . Finally, Figure 3 gives histograms of $\ln(\lambda/\lambda_{\text{opt}})$ for each of the methods from one scenario.

From Table 2, hyp_{632} and GCV_C achieve the stated goal of being useful in $n \approx p$ or $n < p$ situations, as shown by the frequency of being in boldface or annotated with an asterisks. This is most evident in the smaller R^2 scenarios: either hyp_{632} or GCV_C is frequently the

best performing method when $R^2 = 0.2$ or $R^2 = 0.4$. However, these are not uniformly best across all scenarios considered, i.e. the “4000/0.8” column in the bottom sub-table or “25/0.8” column in the second-from-bottom sub-table. Even HYP_{TRUE} , which uses the true value of R^2 , has large rMSPE in these settings, which suggests that the optimal choice of λ is not well-approximated by the $p(1/R^2 - 1)$ expression discussed in Section 4.1. Also competitive are gmpml and 5-CV; GCV performs well except in the $n = 100$ scenarios, in which $n \approx p$. The remaining methods, BIC, AIC_C , $\text{RGCV}_{0.3}$, MPML, MAPHL, and LR have large rMSPE in some scenarios.

To further explore this, Figure 1 plots the ratio of the rMSPE corresponding to HYP_{632} to the median rMSPE from the existing methods, to represent the performance of a typical method. When this ratio is less than one, HYP_{632} has a smaller rMSPE. When $p = 100$, the $y = 1$ line is sometimes crossed when $\log_{10}(n) = 2.5$, or $n = 250$. When $p = 4000$, the $y = 1$ line is exceeded when $n = 4000$ or, regardless of n , when $R^2 = 0.8$ and the covariates are approximately uncorrelated. Based on the discussion in 4.1, HYP_{632} will become equivalent to the asymptotically optimal methods as $n \rightarrow \infty$.

As evidenced in the table, GCV_C has markedly smaller rMSPE than GCV when n is small. The two values of rMSPE coincide as n increases. We argued in Section 2.1 that the GCV penalty has the potential to elicit undesirable behavior when $p = n - 1$, namely choosing $\lambda = 0$, or as close as possible thereto. Figure 2 gives the ramifications of this in terms of prediction error, plotting locally-smoothed values of rMSPE of GCV and GCV_C over many values of n for the $p = 100$ scenarios; $n - 1$ is less than, equal to, and greater than p . For all six panels, which correspond to different R^2 or correlation combinations, there is a peak in the GCV curve beginning near $n - 1 = p$. GCV_C effectively eliminates this behavior and has almost uniformly smaller rMSPE when $n - 1 < p$ and nearly equal rMSPE when $n - 1 > p$.

Finally, we compare the actual values of λ for each method. Figure 3 plots histograms of $\ln(\lambda/\lambda_{\text{opt}})$ for the eleven methods from one simulation setting in Table 2: $p = 4000$, $n = 100$, $\pi = 0.3$, $R^2 = 0.2$, and correlated covariates. For reference, when $\ln(\lambda/\lambda_{\text{opt}}) = 0$, the method has selected the optimal λ . In this small- n scenario, all of the existing methods, at times, choose a very small or large λ . In contrast, the shrinkage from hyperpenalization is evident: the histogram for HYP_{632} has a considerably smaller range. It has the overall smallest rMSPE in this scenario. Finally, GCV has $\ln(\lambda/\lambda_{\text{opt}})$ as small as -12 , and GCV_C has $\ln(\lambda/\lambda_{\text{opt}})$ slightly less than -2 , providing additional evidence that GCV_C effectively prevents overfitting.

6 Bardet-Biedl Data Analysis

To evaluate these methods in a real dataset, we consider the rat gene-expression data first reported in Scheetz et al. (2006). Tissue from 120 12-week old rats was analyzed using microarrays (Affymetrix GeneChip Rat Genome 230 2.0 Array), normalized, and log-transformed. The goal is to find genes associated with expression of the *BBS11/TRIM32* gene, which is causative for Bardet-Biedl syndrome (Chiang et al., 2006). Copying the strategy of Huang et al. (2008), we considered 18,975 probesets that were sufficiently expressed, and, of these, reduced the number further to the $p = 3,000$ probes displaying the

largest variation. We randomly selected $n = 80$ arrays as our training data, fit all methods to these data, and measured rMSPE based on the remaining 40 arrays, repeating this 1,000 times. Figure 4 gives the boxplots of these 1,000 rMSPEs, ordered from left to right in terms of the average rMSPE. The best-performing method is GCV_C with average rMSPE of 50.2, followed by GMPML (53.0), GCV (65.6), HYP_{632} (66.3), and 5-CV (74.2). These rankings correspond closely to those from the simulation study.

7 Discussion

We have examined strategies for choosing the ridge parameter λ when the sample size n is small relative to p . Our small-sample modification to GCV, called GCV_C , is conceptually trivial but uniformly dominates GCV in our simulation study. This corrected GCV may be applied in other shrinkage or smoothing situations that would otherwise use the standard GCV, such as smoothing splines (Wahba, 1985) or adaptively-weighted linear combinations of linear regression estimates (Boonstra et al., 2013).

We also proposed a novel approach using what we call hyperpenalties, which add another level of shrinkage, that of λ itself, by extending the hierarchical model. A hyperpenalty based on the Gamma density with mean $p(1/\hat{R}^2 - 1)$ was shown to work well in the context of ridge regression. The approach is based on the observation that the optimal tuning parameter λ is approximated by the expression $p(1/\hat{R}^2 - 1)$, where p is known. Furthermore, we proposed a simple strategy for estimating R^2 . Relative to existing methods, our implementation can offer superior prediction and protection against choosing extreme values of λ . One area for improvement of this approach lies in the higher- R^2 scenarios, for which it is clear that $p(1/\hat{R}^2 - 1)$ does not approximate the optimal tuning parameter. However, it is unusual in a high-dimensional regression to expect R^2 larger than 0.6 or 0.7.

Another advantage of the hyperpenalty approach is its applicability in missing data problems: when implementing the Expectation-Maximization (EM) algorithm (Dempster et al., 1977), it is not clear how one might do ridge regression concurrently using goodness-of-fit or marginal likelihood approaches to select λ . On the other hand, by taking advantage of the conditional independence, specified by the hierarchical framework, between λ and any missing data given the remaining parameters, it is conceptually straightforward to embed a maximization step for λ , like expression (23), within a larger EM algorithm. This remains the focus of our current research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by the National Science Foundation [DMS1007494] and the National Institutes of Health [CA129102]. The authors thank Jian Huang for generously sharing the Bardet-Biedl data. Code for the simulation study was written in R (R Core Team, 2012), using the Matrix package (Bates and Maechler, 2013) to construct block-diagonal matrices, and is available at <http://www-personal.umich.edu/~philb>.

References

- Akaike H. Information theory and an extension of the maximum likelihood principle. Second International Symposium on Information Theory. 1973;267–281.
- Armagan A, Zaretzki RL. Model selection via adaptive shrinkage with t priors. Computational Statistics. 2010; 25:441–461.
- Bates D, Maechler M. Matrix: Sparse and Dense Matrix Classes and Methods. 2013 R package version 1.0-12.
- Boonstra PS, Taylor JMG, Mukherjee B. Incorporating auxiliary information for improved prediction in high-dimensional datasets: an ensemble of shrinkage approaches. Biostatistics. 2013; 14:259–272. [PubMed: 23087411]
- Burnham, KP.; Anderson, DR. Model selection and multimodel inference: A practical information-theoretic approach. 2nd. Springer; New York: 2002.
- Chiang AP, Beck JS, Yen HJ, Tayeh MK, Scheetz TE, Swiderski RE, Nishimura DY, Braun TA, Kim KYA, Huang J, et al. Homozygosity mapping with snp arrays identifies TRIM32, an E3 ubiquitin ligase, as a Bardet-Biedl syndrome gene (BBS11). Proceedings of the National Academy of Sciences. 2006; 103:6287–6292.
- Craven P, Wahba G. Smoothing noisy data with spline functions. Numerische Mathematik. 1979; 31:377–403.
- Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society: Series B. 1977; 39:1–38.
- Efron B. Estimating the error rate of a prediction rule: Improvement on cross-validation. Journal of the American Statistical Association. 1983; 78:316–331.
- Efron B. Selection criteria for scatterplot smoothers. Annals of Statistics. 2001; 29:470–504.
- Efron B, Tibshirani R. Improvements on cross-validation: the 632+ bootstrap method. Journal of the American Statistical Association. 1997; 92:548–560.
- Frank IE, Friedman JH. A statistical view of some chemometrics regression tools. Technometrics. 1993; 35:109–135.
- Fu WJ. Penalized regressions: The bridge versus the Lasso. Journal of Computational and Graphical Statistics. 1998; 7:397–416.
- Golub GH, Heath M, Wahba G. Generalized cross-validation as a method for choosing a good ridge parameter. Technometrics. 1979; 21:215–223.
- Hardin J, Garcia SR, Golan D, et al. A method for generating realistic correlation matrices. The Annals of Applied Statistics. 2013; 7:1733–1762.
- Harville DA. Maximum likelihood approaches to variance component estimation and to related problems. Journal of the American Statistical Association. 1977; 72:320–338.
- Hastie, T.; Tibshirani, R.; Friedman, J. The Elements of Statistical Learning. 2nd. Springer; New York: 2009.
- Hoerl AE, Kennard RW. Ridge regression: Biased estimation for nonorthogonal problems. Technometrics. 1970; 12:55–67.
- Huang J, Ma S, Zhang CH. Adaptive lasso for sparse high-dimensional regression models. Statistica Sinica. 2008; 18:1603–1618.
- Hurvich CM, Simonoff JS, Tsai CL. Smoothing parameter selection in nonparametric regression using an improved Akaike information criterion. Journal of the Royal Statistical Society: Series B. 1998; 60:271–293.
- Hurvich CM, Tsai CL. Regression and time series model selection in small samples. Biometrika. 1989; 76:297–307.
- Lee Y, Nelder JA. Hierarchical generalized linear models. Journal of the Royal Statistical Society: Series B. 1996; 58:619–678.
- Li KC. Asymptotic optimality of CL and generalized cross-validation in ridge regression with application to spline smoothing. The Annals of Statistics. 1986; 14:1101–1112.
- Lukas MA. Robust generalized cross-validation for choosing the regularization parameter. Inverse Problems. 2006; 22:1883–1902.

- R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria: 2012.
- Scheetz TE, Kim KYA, Swiderski RE, Philp AR, Braun TA, Knudtson KL, Dorrance AM, DiBona GF, Huang J, Casavant TL, et al. Regulation of gene expression in the mammalian eye and its relevance to eye disease. *Proceedings of the National Academy of Sciences*. 2006; 103:14429–14434.
- Schwarz G. Estimating the dimension of a model. *The Annals of Statistics*. 1978; 6:461–464.
- Sin CY, White H. Information criteria for selecting possibly misspecified parametric models. *Journal of Econometrics*. 1996; 71:207–225.
- Strawderman, RL.; Wells, MT. On hierarchical prior specifications and penal-ized likelihood. In: Fourdrinier, D.; Éric, Marchand; Rukhin, AL., editors. *Contemporary Developments in Bayesian Analysis and Statistical Decision Theory: A Festschrift for William E Strawderman*. Vol. 8. Institute of Mathematical Statistics; 2012. p. 154–180.
- Takada Y. Stein's positive part estimator and Bayes estimator. *Annals of the Institute of Statistical Mathematics*. 1979; 31:177–183.
- Tibshirani R. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B*. 1996; 58:267–288.
- Tipping ME. Sparse Bayesian learning and the relevance vector machine. *J Mach Learn Res*. 2001; 1:211–244.
- Tran MN. Penalized maximum likelihood for choosing ridge parameter. *Communications in Statistics*. 2009; 38:1610–1624.
- Wahba G. A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem. *Annals of Statistics*. 1985; 13:1378–1402.
- Wahba G, Wang Y. Behavior near zero of the distribution of GCV smoothing parameter estimates. *Statistics & Probability Letters*. 1995; 25:105–111.
- Wecker WE, Ansley CF. The signal extraction approach to nonlinear regression and spline smoothing. *Journal of the American Statistical Association*. 1983; 78:81–89.
- Zou H, Hastie T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B*. 2005; 67:301–320.

Figure 1.

Ratio of the rMSPE of hyp_{632} to the median rMSPE of remaining methods, excluding GCV_C , over $\log_{10}(n)$. Values less than one indicate that hyp_{632} has smaller prediction error.

Figure 2.

Locally-smoothed values of $rMSPE$ of GCV and GCV_C over n in scenarios for which $p = 100$. The curve corresponding to GCV changes in behavior near the point $n - 1 = p$, given by the vertical dashed line.

Figure 3.

Histograms of $\ln(\lambda/\lambda_{\text{opt}})$ for $p = 4000$, $n = 100$, $\pi = 0.3$, $R^2 = 0.2$, and correlated covariates. $\ln(\lambda/\lambda_{\text{opt}}) = 0$ means that λ was chosen to yield optimal shrinkage. All methods are described in Table 1.

Figure 4.

Table 1

Annotation of all methods from the simulation study and the corresponding Equations and References.

Abbr.	Name	Eqn.	Reference
5-CV	Five-fold cross-validation	(4)	Hastie et al. (Section 7.10, 2009)
BIC	Bayesian Information Criterion	(9)	Schwarz (1978)
AIC _C	Corrected Akaike's Information Criterion	(10)	Hurvich et al. (1998)
GCV	Generalized cross-validation	(7)	Craven and Wahba (1979)
GCV _C	Corrected generalized cross-validation	(12)	Section 2.1
RGCV _γ	Robust generalized cross-validation	(11)	Lukas (2006)
MPML	Maximum profile marginal likelihood	(14)	Wecker and Ansley (1983)
GMPML	Generalized maximum profile marginal likelihood	(15)	Harville (1977); Wahba (1985)
MAPHL	Maximum adjusted profile <i>h</i> -likelihood	(16)–(18)	Lee and Nelder (1996)
LR	Loss-rank	(19)	Tran (2009)
HYP ₆₃₂	Hyperpenalty, $\hat{R}^2$ based on "632" estimator	(21)–(25)	Section 4, Efron (1983)

Table 2

Average rMSPE as defined in (28) for the 11 methods summarized in Table 1 and also hyp_{true} , corresponding to using the true, unknown value of R^2 , with $\pi = 0.3$ (Equation (27)).

<i>p</i> = 100, Approximately Uncorrelated												
Method/ $\{m/R^2\}$	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
5-CV	66.3	24.1	*8.7	*0.8	70.3	*26.2	*9.0	*0.8	*74.9	*35.2	*10.6	*0.6
BIC	370.9	$> 10^4$	103.5	31.6	246.0	$> 10^4$	250.3	16.4	*63.6	$> 10^4$	258.7	3.4
AIC _C	*22.2	17.9	*8.5	*0.8	70.5	40.9	14.5	*0.8	283.7	212.4	31.6	*0.6
GCV	79.2	271.8	*7.7	*0.8	78.9	257.0	*8.1	*0.8	*63.7	392.8	*9.1	*0.6
GCV _C	*23.9	29.7	*7.6	*0.8	37.5	*31.5	*8.1	*0.8	*86.1	47.6	*9.1	*0.6
RGCV _{0.3}	*19.3	53.9	46.2	13.7	64.2	141.3	114.1	*1.0	278.3	560.2	105.0	*0.6
MPML	311.0	18.4	*6.8	*0.8	228.1	*19.4	*6.8	*0.8	*64.1	*25.3	*7.7	*0.6
GMPL	57.1	17.9	*6.8	*0.8	65.8	*18.8	*6.7	*0.8	*66.8	*23.6	*7.7	*0.6
MAPHL	298.0	25.9	*7.0	*0.8	191.5	*27.6	*7.0	*0.8	*44.4	*42.7	*7.8	*0.6
LR	183.6	16.9	19.1	11.3	160.9	40.8	46.3	15.9	*74.9	174.4	178.8	21.6
HYP ₆₃₂	*23.9	*8.0	*8.0	*0.8	*16.1	*19.2	*12.1	*0.8	*59.8	*45.9	*8.5	*0.6
HYP _{TRUE}	8.4	7.1	5.2	0.7	11.2	7.6	4.2	0.7	12.6	15.1	7.8	0.6

<i>p</i> = 100, Positively Correlated												
Method/ $\{m/R^2\}$	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
5-CV	89.4	*23.8	*6.9	*0.8	105.9	*21.0	*6.9	*0.9	*92.0	*22.1	*7.9	*0.8
BIC	593.7	$> 10^4$	61.6	14.9	446.6	$> 10^4$	73.2	24.3	169.0	$> 10^4$	135.8	17.7
AIC _C	*45.3	*15.0	*5.1	*0.8	99.7	*18.1	*5.6	*0.8	362.1	46.6	*12.9	*0.8
GCV	113.1	113.1	*6.5	*0.8	121.9	107.5	*6.5	*0.8	*87.0	168.0	*7.4	*0.8
GCV _C	*55.4	*22.0	*6.4	*0.8	*66.2	*23.2	*6.4	*0.8	*71.9	*25.5	*7.3	*0.8
RGCV _{0.3}	*49.4	45.2	28.7	5.0	126.4	76.6	29.2	11.4	377.5	95.5	64.6	1.8
MPML	232.0	*15.5	*4.9	*1.0	273.2	*13.6	*6.3	*1.0	165.2	*26.6	*11.8	*0.8
GMPL	*64.7	*15.4	*4.8	*1.0	*81.6	*13.2	*6.2	*1.0	*79.5	*24.5	*11.5	*0.8
MAPHL	244.6	*17.5	*5.7	*1.1	200.0	*18.3	*7.7	*1.0	*77.9	39.6	*13.3	*0.8
LR	*49.4	*24.8	16.4	4.3	86.5	41.0	24.4	7.5	172.1	80.9	57.0	15.5

<i>p</i> = 100, Approximately Uncorrelated												
Method/ $\{n/R^2\}$	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
HYP ₆₃₂	*34.4	*13.9	*6.6	*1.4	*42.0	*16.7	*5.3	1.8	*88.1	*19.6	18.2	*1.2
HYP _{TRUE}	17.9	10.5	4.9	1.4	27.1	10.4	2.9	1.8	32.7	15.8	20.5	1.2
<i>p</i> = 4000, Approximately Uncorrelated												
Method/ $\{n/R^2\}$	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
5-CV	39.7	17.7	*7.7	*1.0	45.8	*19.2	*7.4	*1.1	40.2	*14.3	*6.6	*1.2
BIC	100.4	114.0	91.2	40.3	53.0	68.0	63.8	73.5	*5.6	*12.7	16.1	376.7
AIC _C	34.3	22.6	*11.0	1.9	100.9	49.9	22.9	8.5	329.2	160.4	78.4	89.1
GCV	40.6	16.3	*8.2	*0.9	43.9	*18.5	*8.6	*1.0	33.3	*13.7	*6.4	*1.1
GCV _C	*19.1	13.0	*6.5	*0.9	38.8	*13.7	*5.5	*1.0	110.6	*17.9	*5.4	*1.1
RGCV _{0.3}	29.5	37.8	35.0	5.2	90.0	94.5	73.7	20.7	312.5	273.2	137.1	216.8
MPML	100.3	54.1	*7.1	*1.0	53.0	64.1	11.2	*1.2	*5.6	*12.7	16.1	*1.1
GNPML	38.7	15.3	*6.4	*1.0	43.1	*17.9	*6.1	*1.2	29.6	*12.9	*7.2	*1.1
MAPHL	99.9	113.6	90.7	13.0	52.6	67.6	63.3	12.2	*5.6	*12.5	15.9	11.7
LR	100.3	14.4	12.8	6.0	53.1	*26.9	22.0	14.7	*5.6	55.8	48.3	75.6
HYP ₆₃₂	*10.2	*6.0	*5.6	*1.1	*15.0	*14.3	*7.3	2.8	77.4	35.0	10.9	9.0
HYP _{TRUE}	17.4	13.2	6.9	0.9	29.0	12.8	4.7	1.8	21.9	6.6	2.3	11.9
<i>p</i> = 4000, Positively Correlated												
Method/ $\{n/R^2\}$	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
5-CV	83.5	*15.8	*8.3	*0.7	82.8	*16.5	9.7	*0.8	*72.0	*15.9	*7.2	*1.0
BIC	314.2	95.5	85.3	11.0	261.6	73.9	74.4	14.9	*95.2	33.4	52.1	42.2
AIC _C	*48.9	*11.8	*5.3	*0.6	84.3	*15.5	*6.2	*0.7	239.1	45.8	17.0	5.2
GCV	85.8	*16.7	*8.7	*0.7	84.9	*17.0	9.8	*0.7	*70.8	*16.2	*6.6	*0.9
GCV _C	*49.7	*13.3	*6.4	*0.7	*50.3	*12.9	*6.6	*0.7	*61.5	*12.2	*4.0	*0.9
RGCV _{0.3}	*56.7	30.6	19.7	*0.9	111.7	47.4	25.6	*1.3	254.7	71.7	24.2	12.4
MPML	242.0	*14.1	*5.3	*0.8	249.2	*16.3	*5.4	*1.4	*95.2	32.6	21.5	2.3
GNPML	*60.1	*13.5	*5.3	*0.7	*67.9	*13.5	*4.8	*1.4	*70.3	*13.8	*6.2	2.3
MAPHL	312.4	94.9	83.9	13.3	259.8	73.3	73.0	14.1	*94.2	32.9	50.9	14.5

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

<i>p</i> = 100, Approximately Uncorrelated												
Method/ <i>m/R</i> ²	25/0.2	100/0.2	250/0.2	4000/0.2	25/0.4	100/0.4	250/0.4	4000/0.4	25/0.8	100/0.8	250/0.8	4000/0.8
LR	80.9	*16.9	11.7	2.4	108.3	26.5	17.4	2.9	122.7	59.2	31.5	8.4
HYP ₆₃₂	* 31.0	* 10.9	*5.6	*0.7	* 35.5	* 12.9	* 4.7	2.2	* 55.5	* 12.0	*4.6	32.4
HYP _{TRUE}	18.9	5.0	3.3	0.6	16.2	7.7	3.4	2.2	33.8	9.3	3.7	35.3

Values in **bold** are the column-wise minima, excluding HYP_{TRUE}, and those with an *,** are less than twice the column-wise minimum. HYP_{TRUE}.