

MEASUREMENT ERROR IN LASSO: IMPACT AND LIKELIHOOD BIAS CORRECTION

Øystein Sørensen, Arnaldo Frigessi and Magne Thoresen

University of Oslo

Abstract: Regression with the lasso penalty is a popular tool for performing dimension reduction when the number of covariates is large. In many applications of the lasso, like in genomics, covariates are subject to measurement error. We study the impact of measurement error on linear regression with the lasso penalty, both analytically and in simulation experiments. A simple method of correction for measurement error in the lasso is then considered. In the large sample limit, the corrected lasso yields sign consistent covariate selection under conditions very similar to the lasso with perfect measurements, whereas the uncorrected lasso requires much more stringent conditions on the covariance structure of the data. Finally, we suggest methods to correct for measurement error in generalized linear models with the lasso penalty, which we study empirically in simulation experiments with logistic regression, and also apply to a classification problem with microarray data. We see that the corrected lasso selects less false positives than the standard lasso, at a similar level of true positives. The corrected lasso can therefore be used to obtain more conservative covariate selection in genomic analysis.

Key words and phrases: Conditional score, generalized linear model, lasso, measurement error.

1. Introduction

Due to rapid technological progress, complex, high-dimensional data sets are now commonplace in a range of fields, e.g., genomics and finance. Various penalization schemes have been proposed, which shrink the parameter space, including the Dantzig selector (Candès and Tao (2007)), the lasso (Tibshirani (1996)), ridge regression (Hoerl and Kennard (1970)), and the SCAD penalty (Fan and Li (2001)). The lasso has been extensively used in applied problems, and its statistical scope and limitations are well understood (e.g., Benjamini and Speed (2012) and the references cited therein). A common assumption is *sparsity*, i.e., only a small number of covariates influence the outcome. Several refinements have been proposed, in particular the adaptive lasso (Zou (2006)), which relaxes the rather strict conditions required for consistent covariate selection by the standard lasso.

Measurement error in the covariates is a problem in various high-dimensional data sets. In genomics, examples include gene expression microarray data (Purdom and Holmes (2005), Rocke and Durbin (2001)) and high-throughput sequencing (Benjamini and Speed (2012)). In classical regression models, measurement error is known to cause biased parameter estimates and lack of power (Carroll et al. (2006)). Measurement error in SCAD regression has been studied by, e.g., Liang and Li (2009), Ma and Li (2010), Xu and You (2007). Rosenbaum and Tsybakov (2010) introduced the matrix uncertainty (MU) selector, a modification of the Dantzig selector that handles measurement error and missing data. Through analytical results for the finite sample case, the MU selector is shown to give good parameter estimation and covariate selection. An improved MU selector is presented in Rosenbaum and Tsybakov (2013). Loh and Wainwright (2012) consider generalized M-estimators with lasso regularization, of which special cases include correction for additive measurement error or missing data. The method is shown to yield estimates close to the true parameters, as measured in the ℓ_1 - or ℓ_2 -norm, and is computationally feasible in the high-dimensional case, despite its non-convexity. We also mention Chen and Caramanis (2013), who consider high-dimensional measurement error problems with independent covariates, and develop a modified orthogonal matching pursuit algorithm yielding correct sparsity recovery with high probability, also when estimates of the measurement error do not exist.

Since the standard lasso is widely used despite the presence of measurement error, it is of interest to study the impact measurement error has on the analysis. In the first half of this paper, we thus ask: *Under which conditions can the standard lasso (naive approach) be safely used, and when are correction methods required?* For a linear model with additive measurement error, we demonstrate how measurement error affects estimation and prediction error. In the fixed p , large n setting, we also show that the naive lasso yields asymptotically sign consistent covariate selection only under very stringent conditions on the noise. Next, a correction of the lasso loss function for linear models is considered, to compensate for measurement error in the covariates. The estimation error of this correction has been studied earlier by Loh and Wainwright (2012). Here, we derive finite sample conditions under which this corrected lasso yields sign consistent covariate selection, and show that it performs asymptotically as well as the lasso without measurement error. We then go on to consider the lasso for generalized linear models (GLMs), and suggest ways to correct for additive measurement error in GLMs using the conditional score method of Stefanski and Carroll (1987), together with an efficient projection algorithm of Duchi et al. (2008). The analytical results for linear regression are illustrated through simulations, and the statistical and computational properties of the suggested

correction method for GLMs are studied in simulation experiments. We also illustrate the use of measurement error correction in the lasso for logistic regression in an example with microarray data. Proofs and additional conditions are given in the Supplementary Material (Sørensen, Frigessi, and Thoresen (2014)).

2. Model Setup

In Sections 3 and 4, we consider a linear regression model with additive measurement error,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^0 + \boldsymbol{\epsilon} \quad \text{and} \quad \mathbf{W} = \mathbf{X} + \mathbf{U}, \quad (2.1)$$

with observations of p covariates and a continuous response $\mathbf{y} \in \mathbb{R}^n$ on n individuals. The true covariates \mathbf{X} are not observed, and instead we have noisy measurements \mathbf{W} . The matrix of measurement errors $\mathbf{U} \in \mathbb{R}^{n \times p}$ is assumed to have normally distributed rows, with mean zero and covariance $\boldsymbol{\Sigma}_{uu}$. The model errors $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)'$ are i.i.d. normally with mean zero and variance σ^2 . In Section 5, the linear model on the left-hand side of (2.1) is replaced by a GLM, but additive measurement error is still assumed.

Let $S_0 = \{j : \beta_j^0 \neq 0\}$ be the index set of non-zero components of the true coefficient vector $\boldsymbol{\beta}^0 \in \mathbb{R}^p$, and denote the number of relevant covariates by $s_0 = \text{card}\{S_0\}$. Under the sparsity assumption, most components of $\boldsymbol{\beta}^0$ are zero, such that $s_0 \ll p$. Direct use of error-prone measurements is referred to as the *naïve approach* in the measurement error literature, and the naive lasso for a linear model takes the form

$$\hat{\boldsymbol{\beta}}(\lambda) = \arg \min_{\boldsymbol{\beta}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{W}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1 \right\}, \quad (2.2)$$

where $\lambda > 0$ is a regularization parameter. For any $\lambda \geq 0$, define the *active set* of the lasso, $\hat{S}(\lambda) = \{j : \hat{\beta}_j(\lambda) \neq 0\}$. Given $\boldsymbol{\beta}^0$, we order the covariates such that $S_0 = \{1, \dots, s_0\}$, $S_0^c = \{s_0 + 1, \dots, p\}$, and introduce the partitioning $\mathbf{W} = (\mathbf{W}_{S_0}, \mathbf{W}_{S_0^c})$, where $\mathbf{W}_{S_0} \in \mathbb{R}^{n \times s_0}$ contains the n measurements of the s_0 relevant covariates, and $\mathbf{W}_{S_0^c} \in \mathbb{R}^{n \times (p-s_0)}$ contains the n measurements of the $(p-s_0)$ irrelevant covariates. The same notation is used for \mathbf{X} and \mathbf{U} . Sample covariance matrices are denoted by \mathbf{C} , and subscripts show which covariates are involved. For example, the empirical covariance of the measurements is given by $\mathbf{C}_{ww} = (1/n)\mathbf{W}'\mathbf{W}$. Using \mathbf{C}_{ww} as an example, we partition the covariance matrices in the form

$$\mathbf{C}_{ww} = \begin{pmatrix} \mathbf{C}_{ww}(S_0, S_0) & \mathbf{C}_{ww}(S_0, S_0^c) \\ \mathbf{C}_{ww}(S_0^c, S_0) & \mathbf{C}_{ww}(S_0^c, S_0^c) \end{pmatrix}.$$

Population covariance matrices are denoted by $\boldsymbol{\Sigma}$, and indexed by subscripts and superscripts in the same way as described for the sample covariance matrices.

The true coefficient vector is written as $\beta^0 = ((\beta_{S_0}^0)', (\beta_{S_0^c}^0)')$, where $\beta_{S_0}^0 \in \mathbb{R}^{s_0}$ are the non-zero coefficients and $\beta_{S_0^c}^0 \in \mathbb{R}^{(p-s_0)}$ is a vector of zeros. The lasso estimates $\hat{\beta}$ are divided according to the same pattern, with implicit dependence on λ . The elements of $\hat{\beta}_{S_0}$ are not necessarily non-zero, neither are the elements of $\hat{\beta}_{S_0^c}$ necessarily zero.

Finally, vectors and matrices are written in boldface, and we use the notation $|\mathbf{v}| \leq |\mathbf{w}|$ for vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^p$ if $|v_i| \leq |w_i|$ for $i = 1, \dots, p$, and equivalently for other relational operators.

3. Impact of Ignoring Measurement Error

Using the error-free case as a reference, we show in this section how known results for estimation, screening, and selection are affected by additive measurement error.

3.1. Estimation error

In the absence of measurement error, the lasso is consistent for estimation and prediction under certain conditions. In particular, the design \mathbf{X} must satisfy a compatibility condition, and the noise must satisfy $(2/n)\|\epsilon'\mathbf{X}\|_\infty \leq \lambda_0$ for some constant λ_0 . If the regularization parameter is chosen large enough to rule out the noise, the lasso is consistent. For $\lambda \geq 2\lambda_0$, the bound

$$\frac{1}{n} \left\| \mathbf{X} (\hat{\beta} - \beta^0) \right\|_2^2 + \lambda \left\| \hat{\beta} - \beta^0 \right\|_1 \leq \frac{4\lambda^2 s_0}{\phi_0^2}$$

holds, where ϕ_0 is a compatibility constant (Benjamini and Speed (2012, Chap. 6)). Then, e.g., for Gaussian errors, $(2/n)\|\epsilon'\mathbf{X}\|_\infty \leq \lambda_0$ holds with high probability for $\lambda_0 \asymp \sqrt{\log p/n}$. Hence, as long as $n \rightarrow \infty$ faster than $s_0 \log p$, lasso is consistent for prediction and, if $n \rightarrow \infty$ faster than $s_0^2 \log p$, lasso is consistent for estimation in the ℓ_1 -norm.

When the covariates are subject to measurement error, there are two noise terms which need to be bounded: the model error ϵ and the measurement error \mathbf{U} . In order to bound the estimation error, we need a compatibility condition involving the observed covariates.

Definition 1. The compatibility condition holds for the index set S_0 if, for some $\phi_0 > 0$ and all $\gamma \in \mathbb{R}^p$ such that $\|\gamma_{S_0^c}\|_1 \leq 3\|\gamma_{S_0}\|_1$, it holds that

$$\|\gamma_{S_0}\|_1^2 \leq \frac{s_0 \|\mathbf{W}\gamma\|_2^2}{n\phi_0^2}.$$

Proposition 1. Assume the compatibility condition with constant ϕ_0 , and that there exists a constant λ_0 such that

$$\frac{2}{n} \left\| (\boldsymbol{\epsilon} - \mathbf{U}\boldsymbol{\beta}^0)' \mathbf{W} \right\|_{\infty} \leq \lambda_0. \tag{3.1}$$

Then, with a regularization parameter $\lambda \geq 2\lambda_0$,

$$\frac{1}{n} \left\| \mathbf{W} \left(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0 \right) \right\|_2^2 + \lambda \left\| \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^0 \right\|_1 \leq \frac{4\lambda^2 s_0}{\phi_0^2}. \tag{3.2}$$

The result (3.2) shows us that in the presence of measurement error, the estimation error of the lasso can be bounded. However, the bound (3.1) contains a term which is quadratic in the measurement error. By the triangle inequality, the bound (3.1) is implied by

$$\frac{2}{n} \left\| \boldsymbol{\epsilon}' \mathbf{W} \right\|_{\infty} + \frac{2}{n} \left\| (\boldsymbol{\beta}^0)' \mathbf{U}' \mathbf{X} \right\|_{\infty} + \frac{2}{n} \left\| \boldsymbol{\beta}^0 \right\|_1 \left\| \mathbf{U}' \mathbf{U} \right\|_{\infty} \leq \lambda_0.$$

Hence, if all three terms in the expression above converge to zero, the lasso with measurement error is consistent. However, the term $\mathbf{U}' \mathbf{U}$ converges to $n \boldsymbol{\Sigma}_{uu}$ as $n \rightarrow \infty$. Since $\left\| \boldsymbol{\Sigma}_{uu} \right\|_{\infty} \neq 0$, we do not obtain consistency.

Proposition 2. Assume $\lambda \rightarrow 0$ as $n \rightarrow \infty$. Then, as $n \rightarrow \infty$ with fixed p , $\hat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\Sigma}_{ww}^{-1} \boldsymbol{\Sigma}_{xx} \boldsymbol{\beta}^0$.

In the absence of measurement error, the lasso estimates converge in probability to $\boldsymbol{\beta}^0$ under the same conditions (Knight and Fu (2000)). Hence, with a proper scaling of λ , the bias induced by additive measurement error is the same as for a multivariate linear model (Carroll et al. (2006)).

3.2. Covariate selection

We now consider exact recovery of the sign pattern of $\boldsymbol{\beta}^0$, an important goal in high-throughput genomics. In the absence of measurement error, such *sign consistent covariate selection* requires an irrepresentable condition (IC) (Meinshausen and Bühlmann (2010), Zhao and Yu (2006)). In the presence of measurement error, the IC has a new form.

Definition 2. The IC with Measurement Error (IC-ME) holds if there exists a constant $\theta \in [0, 1)$ such that $\left\| \mathbf{C}_{ww} (S_0^c, S_0) \mathbf{C}_{ww} (S_0, S_0)^{-1} \text{sign} (\boldsymbol{\beta}_{S_0}^0) \right\|_{\infty} \leq \theta$.

We refer to Zhao and Yu (2006) for a thorough interpretation of the IC. In the presence of measurement error, we need an additional condition to obtain sign consistent covariate selection with high probability.

Definition 3. The Measurement Error Condition (MEC) is satisfied if

$$\boldsymbol{\Sigma}_{ww} (S_0^c, S_0) \boldsymbol{\Sigma}_{ww} (S_0, S_0)^{-1} \boldsymbol{\Sigma}_{uu} (S_0, S_0) - \boldsymbol{\Sigma}_{uu} (S_0^c, S_0) = \mathbf{0}.$$

The MEC applies to population covariance matrices, whereas the IC-ME applies to sample covariance matrices. Let

$$\begin{aligned}\mathbf{Z}_1 &= \mathbf{C}_{ww}(S_0, S_0)^{-1} \frac{\mathbf{W}'_{S_0}}{\sqrt{n}}, \\ \mathbf{Z}_2 &= \sqrt{n} \mathbf{C}_{ww}(S_0, S_0)^{-1} \mathbf{C}_{wu}(S_0, S_0), \\ \mathbf{Z}_3 &= \mathbf{C}_{ww}(S_0^c, S_0) \mathbf{C}_{ww}(S_0, S_0)^{-1} \frac{\mathbf{W}'_{S_0}}{\sqrt{n}} - \frac{\mathbf{W}'_{S_0^c}}{\sqrt{n}}, \\ \mathbf{Z}_4 &= \sqrt{n} \left(\mathbf{C}_{ww}(S_0^c, S_0) \mathbf{C}_{ww}(S_0, S_0)^{-1} \mathbf{C}_{wu}(S_0, S_0) - \mathbf{C}_{wu}(S_0^c, S_0) \right).\end{aligned}$$

We have a result for covariate selection with the lasso in the presence of measurement error.

Theorem 1. *If the IC-ME holds with constant θ , then $P(\text{sign}(\hat{\beta}) = \text{sign}(\beta^0)) \geq P(A \cap B)$, for the events*

$$\begin{aligned}A &= \left\{ \left| \mathbf{Z}_1 \boldsymbol{\epsilon} - \mathbf{Z}_2 \beta_{S_0}^0 \right| < \sqrt{n} \left(\left| \beta_{S_0}^0 \right| - \frac{\lambda}{2} \left| \mathbf{C}_{ww}(S_0, S_0)^{-1} \text{sign}(\beta_{S_0}^0) \right| \right) \right\}, \\ B &= \left\{ \left| \mathbf{Z}_3 \boldsymbol{\epsilon} - \mathbf{Z}_4 \beta_{S_0}^0 \right| < \frac{\lambda \sqrt{n}}{2} (1 - \theta) \mathbf{1} \right\}.\end{aligned}$$

If the MEC is satisfied and $|\beta_{S_0}^0| > |\boldsymbol{\Sigma}_{ww}(S_0, S_0)^{-1} \boldsymbol{\Sigma}_{uu}(S_0, S_0) \beta_{S_0}^0|$, then $P(\text{sign}(\hat{\beta}) = \text{sign}(\beta^0)) = 1 - o(\exp(-n^c))$, for some $c \in [0, 1)$, if $\lambda \rightarrow 0$ and $\lambda n^{(1-c)/2} \rightarrow \infty$ as $n \rightarrow \infty$ with fixed p .

Event A has the relevant covariates estimated with correct sign and, given A , event B has the coefficients of the irrelevant covariates correctly set to zero. As with perfectly measured covariates (Zhao and Yu (2006)), the left-hand sides of A and B involve the model error $\boldsymbol{\epsilon}$, which needs to be bounded, and terms involving $\beta_{S_0}^0$. The same dependence on β^0 is seen in the results of Chen and Caramanis (2013) for covariate selection by orthogonal matching pursuit. The events A and B illustrate the trade-off between choosing λ small enough to include the relevant covariates (increasing $P(A)$) and large enough to discard the irrelevant covariates (increasing $P(B)$).

Useful insight into necessary and sufficient conditions for sign consistent covariate selection can be obtained by considering the case of no model error, $\boldsymbol{\epsilon} = \mathbf{0}$. In the absence of measurement error, the IC is known to be a sharp condition when $\boldsymbol{\epsilon} = \mathbf{0}$: for a finite sample, the lasso will estimate the signs correctly if and only if a version of the IC holds (Benjamini and Speed (2012, Chap. 7)). Our next result states necessary and sufficient conditions for sign consistent covariate selection when $\boldsymbol{\epsilon} = \mathbf{0}$ and the covariates are subject to measurement error. We use the shorthand

$$\mathbf{Z}_5 = \mathbf{C}_{ww}(S_0^c, S_0) \mathbf{C}_{ww}(S_0, S_0)^{-1} \mathbf{C}_{wu}(S_0, S_0) - \mathbf{C}_{wu}(S_0^c, S_0).$$

Proposition 3. *For the naive lasso in the case of no model error, let the set of detectable covariates be*

$$S_0^{det} = \left\{ j : |\beta_j^0| > \frac{\lambda}{2} \left(\sup_{\|\tau_{S_0}\|_\infty \leq 1} \left\| \mathbf{C}_{ww}(S_0, S_0)^{-1} \tau_{S_0} \right\|_\infty \right) + |v_j| \right\}, \quad (3.3)$$

where $\mathbf{v} = (v_1, \dots, v_p)' = \mathbf{C}_{ww}(S_0, S_0)^{-1} \mathbf{C}_{wu}(S_0, S_0) \beta_{S_0}^0$. If the IC-ME is satisfied and $\mathbf{Z}_5 \beta_{S_0}^0 = \mathbf{0}$, then $S_0^{det} \subseteq \hat{S}(\lambda) \subseteq S_0$. Conversely, if $\hat{S}(\lambda) = S_0 = S_0^{det}$, then

$$\left\| \mathbf{C}_{ww}(S_0^c, S_0) \mathbf{C}_{ww}(S_0, S_0)^{-1} \text{sign}(\beta_{S_0}^0) + \frac{2}{\lambda} \mathbf{Z}_5 \beta_{S_0}^0 \right\|_\infty \leq 1. \quad (3.4)$$

In (3.3), the first term is the same as in the absence of measurement error, with \mathbf{C}_{ww} in place of \mathbf{C}_{xx} , while the second term involves the measurement errors and $\beta_{S_0}^0$; the lasso cannot detect arbitrarily small coefficients in the presence of measurement error.

4. Correction for Measurement Error in Lasso: Linear Case

We consider a corrected lasso, which yields sign consistent covariate selection under an IC-type condition. The correction we use is motivated by the fact that the loss function of the naive lasso is biased:

$$E(\|\mathbf{y} - \mathbf{W}\beta\|_2^2 | \mathbf{X}, \mathbf{y}) = \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + n\beta' \Sigma_{uu} \beta.$$

This suggests the regularized corrected lasso (RCL),

$$\hat{\beta}_{\text{RCL}} = \arg \min_{\beta: \|\beta\|_1 \leq R} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{W}\beta\|_2^2 - \beta' \Sigma_{uu} \beta + \lambda \|\beta\|_1 \right\}, \quad (4.1)$$

introduced by Loh and Wainwright (2012). The loss function of the RCL is always non-convex when $p > n$, and its parameter space must be restricted to the ℓ_1 -ball $\{\beta : \|\beta\|_1 \leq R\}$ with some finite radius R to avoid trivial solutions. There are thus two regularization parameters, λ and R . A related problem is the constrained corrected lasso (CCL),

$$\hat{\beta}_{\text{CCL}} = \arg \min_{\beta: \|\beta\|_1 \leq \kappa} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{W}\beta\|_2^2 - \beta' \Sigma_{uu} \beta \right\}, \quad (4.2)$$

where κ is to be chosen by some model selection procedure. Unless a distinction is necessary, we refer to both as the corrected lasso. The same correction has been proposed for linear regression with the SCAD penalty (Liang and Li (2009)). Since the lasso does not possess the oracle property of the SCAD (Fan and Li (2001)), the results of those papers do not immediately hold for the lasso. The corrected lasso has already been shown to yield good estimation bounds (Loh and Wainwright (2012)), and we now study its capacity for sign consistent selection.

Definition 4. The Irrepresentable Condition for the Corrected lasso (IC-CL) holds if the matrix $\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0)$ is invertible, and there exists a constant $\theta \in [0, 1)$ such that

$$\left\| (\mathbf{C}_{ww}(S_0^c, S_0) - \boldsymbol{\Sigma}_{uu}(S_0^c, S_0)) (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} \text{sign}(\boldsymbol{\beta}_{S_0}^0) \right\|_{\infty} \leq \theta.$$

When the empirical covariance matrices are replaced by population covariance matrices, the IC-CL reduces to the standard IC without measurement error. We use the shorthands,

$$\begin{aligned} \mathbf{Z}_6 &= (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} \frac{\mathbf{W}'_{S_0}}{\sqrt{n}}, \\ \mathbf{Z}_7 &= \sqrt{n} (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} (\mathbf{C}_{wu}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0)), \\ \mathbf{Z}_8 &= (\mathbf{C}_{ww}(S_0^c, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0)) (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} \frac{\mathbf{W}'_{S_0}}{\sqrt{n}} \\ &\quad - \frac{\mathbf{W}'_{S_0^c}}{\sqrt{n}}, \\ \mathbf{Z}_8 &= \sqrt{n} \left((\mathbf{C}_{ww}(S_0^c, S_0) - \boldsymbol{\Sigma}_{uu}(S_0^c, S_0)) (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} \right. \\ &\quad \left. (\mathbf{C}_{wu}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0)) - (\mathbf{C}_{wu}(S_0^c, S_0) - \boldsymbol{\Sigma}_{uu}(S_0^c, S_0)) \right), \end{aligned}$$

in a result for covariate selection with the corrected lasso in the presence of measurement error.

Theorem 2. Assume the IC-CL holds with constant θ . If a local optimum of the RCL, $\hat{\boldsymbol{\beta}}$, satisfies $\|\hat{\boldsymbol{\beta}}\|_1 < R$, then

$$P\left(\text{sign}(\hat{\boldsymbol{\beta}}) = \text{sign}(\boldsymbol{\beta}^0)\right) \geq P(A \cap B), \quad (4.3)$$

for the events

$$A = \left\{ |\mathbf{Z}_6 \boldsymbol{\epsilon} - \mathbf{Z}_7 \boldsymbol{\beta}_{S_0}^0| < \sqrt{n} \left(|\boldsymbol{\beta}_{S_0}^0| - \frac{\lambda}{2} \left| (\mathbf{C}_{ww}(S_0, S_0) - \boldsymbol{\Sigma}_{uu}(S_0, S_0))^{-1} \text{sign}(\boldsymbol{\beta}_{S_0}^0) \right| \right) \right\}, \quad (4.4)$$

$$B = \left\{ |\mathbf{Z}_8 \boldsymbol{\epsilon} - \mathbf{Z}_9 \boldsymbol{\beta}_{S_0}^0| < \frac{\lambda \sqrt{n}}{2} (1 - \theta) \mathbf{1} \right\}. \quad (4.5)$$

Furthermore, if $\lambda_n \rightarrow 0$ and $\lambda_n n^{(1-c)/2} \rightarrow \infty$ as $n \rightarrow \infty$ with fixed p ,

$$P\left(\text{sign}(\hat{\boldsymbol{\beta}}) = \text{sign}(\boldsymbol{\beta}^0)\right) = 1 - o(\exp(-n^c)), \text{ for some } c \in [0, 1). \quad (4.6)$$

As pointed out by a referee, the condition $\|\hat{\beta}\|_1 < R$ is required since the KKT conditions do not characterize critical points on the boundary of the feasible set. However, a consequence of Theorem 2 in Loh and Wainwright (2012) is that the distance $\|\hat{\beta} - \beta^0\|_1$ for any local optimum $\hat{\beta}$ is $O(s_0\sqrt{\log(p)/n})$. For sufficiently large n , we can choose R such that the feasible set contains all these optima.

5. Conditional Scores Lasso for GLMs

We consider a generalized linear model (GLM), for which Y given \mathbf{X} has density

$$f(y|\mathbf{x}, \Theta) = \exp \left\{ \frac{y\eta - \mathcal{D}(\eta)}{\phi} + c(y, \phi) \right\},$$

where $\eta = \mu + \mathbf{x}'\beta$, $c(\cdot)$ and $\mathcal{D}(\cdot)$ are functions, and $\Theta = (\mu, \beta, \phi)$ is the vector of unknown parameters: μ is the intercept and ϕ is the dispersion parameter. When covariates are subject to additive measurement error, unbiased score functions can be constructed using the conditional scores method of Stefanski and Carroll (1987), yielding consistent estimators of Θ . The method is reviewed by Carroll et al. (2006, Chap. 7), and we follow their notation. The basic idea is to introduce the sufficient statistic for \mathbf{x} , $\delta = \mathbf{w} + y\boldsymbol{\Sigma}_{uu}\beta/\phi$, to obtain the conditional density

$$f(y|\delta, \Theta, \boldsymbol{\Sigma}_{uu}) = \exp \left\{ \frac{y\eta_* - \mathcal{D}_*(\eta_*, \phi, \beta'\boldsymbol{\Sigma}_{uu}\beta)}{\phi} + c_*(y, \phi, \beta'\boldsymbol{\Sigma}_{uu}\beta) \right\}, \quad (5.1)$$

where η_* , $c_*(\cdot)$ and $\mathcal{D}_*(\cdot)$ are modifications of the functions used in the absence of measurement error. Assuming the dispersion parameter ϕ is known, as it is for logistic and Poisson regression, consistent estimators of $(\mu, \beta)'$ are now obtained by solving the estimating equation

$$\sum_{i=1}^n \left(y_i - \frac{\partial}{\partial \eta_*} \mathcal{D}_* \right) \begin{pmatrix} 1 \\ \delta_i \end{pmatrix} = \mathbf{0}. \quad (5.2)$$

This suggests a way of obtaining corrected lasso estimates for GLMs with measurement error, by plugging the estimating equation (5.2) into the projected gradient algorithm used by Loh and Wainwright (2012). In particular, the iteration scheme

$$\mu^{s+1} = \mu^s + \alpha \sum_{i=1}^n \left(y_i - \frac{\partial}{\partial \eta_*} \mathcal{D}_*(\eta_{*i}^s, (\beta^s)'\boldsymbol{\Sigma}_{uu}\beta^s) \right), \quad (5.3)$$

$$\beta^{s+1} = \Pi_{\mathcal{B}(\kappa)} \left\{ \beta^s + \alpha \sum_{i=1}^n \left(y_i - \frac{\partial}{\partial \eta_*} \mathcal{D}_*(\eta_{*i}^s, (\beta^s)'\boldsymbol{\Sigma}_{uu}\beta^s) \right) \delta_i^s \right\}, \quad (5.4)$$

for $s = 1, 2, \dots$, until convergence, where $\Pi_{\mathcal{B}(\kappa)}(\cdot)$ denotes projection onto $\mathcal{B}(\kappa)$, α is the stepsize, and η_{*i}^s is the value of η_* for subject i at iteration s , gives regression coefficients constrained to $\mathcal{B}(\kappa) = \{\boldsymbol{\beta} \in \mathbb{R}^p : \|\boldsymbol{\beta}\|_1 \leq \kappa\}$. The projection $\Pi_{\mathcal{B}(\kappa)}(\cdot)$ can be performed by an efficient algorithm proposed by Duchi et al. (2008). The theoretical results of Loh and Wainwright (2012) do not necessarily apply to GLMs, so we cannot guarantee that the local optimum found by iteration (5.3)-(5.4) is close to the global optimum. However, we provide empirical results that suggest that this algorithm is indeed useful.

Selection of the constraint parameter κ by standard cross-validation requires a loss function. Hanfelt and Liang (1997) construct an approximate likelihood for this model by path-dependent integration, but a simpler alternative may be to use stability selection (Meinshausen and Bühlmann (2010)) or the ‘elbow rule’ (Rosenbaum and Tsybakov (2010, Fig. 1)), for which no loss function is required. Conditional score functions for GLMs with measurement error can be straightforwardly derived (Carroll et al. (2006, Chap. 7)). We will here consider logistic and Poisson regression, for which $\phi = 1$ and corrected lasso estimates are easily obtained.

5.1. Logistic regression

Logistic regression with the lasso penalty has been used, e.g., in detection of differentially expressed genetic markers in case/control studies (Ayers and Cordell (2010), Wu et al. (2009)). Here, we consider binomial logistic regression, with response $y_i \sim B(1, H(\eta))$, $i = 1, \dots, n$, where $H(\eta) = \{1 + \exp(-\eta)\}^{-1}$ is the logit function. When the covariates are subject to additive measurement error, the terms in the conditional density (5.1) are

$$\begin{aligned} \eta_* &= \mu + \boldsymbol{\beta}'(\mathbf{w} + y \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}), \\ c_*(y, \boldsymbol{\beta}' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}) &= \left(\frac{-y^2}{2}\right) \boldsymbol{\beta}' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}, \\ \mathcal{D}_*(\eta_*, \boldsymbol{\beta}' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}) &= \log \left\{ 1 + \exp \left(\eta_* - \frac{1}{2} \boldsymbol{\beta}' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta} \right) \right\}, \\ \frac{\partial \mathcal{D}_*}{\partial \eta_*} &= H \left\{ \eta_* - \frac{1}{2} \boldsymbol{\beta}' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta} \right\}. \end{aligned}$$

Hence, the iteration scheme (5.3)-(5.4) is

$$\begin{aligned} \mu^{s+1} &= \mu^s + \alpha \sum_{i=1}^n \left(y_i - H \left\{ \mu^s + (\boldsymbol{\beta}^s)' \mathbf{w}_i + (y_i - 1/2) (\boldsymbol{\beta}^s)' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}^s \right\} \right), \\ \boldsymbol{\beta}^{s+1} &= \\ \Pi_{\mathcal{B}(\kappa)} &\left\{ \boldsymbol{\beta}^s + \alpha \sum_{i=1}^n \left(y_i - H \left\{ \mu^s + (\boldsymbol{\beta}^s)' \mathbf{w}_i + (y_i - \frac{1}{2}) (\boldsymbol{\beta}^s)' \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}^s \right\} \right) (\mathbf{w}_i + y_i \boldsymbol{\Sigma}_{uu} \boldsymbol{\beta}^s) \right\}, \end{aligned}$$

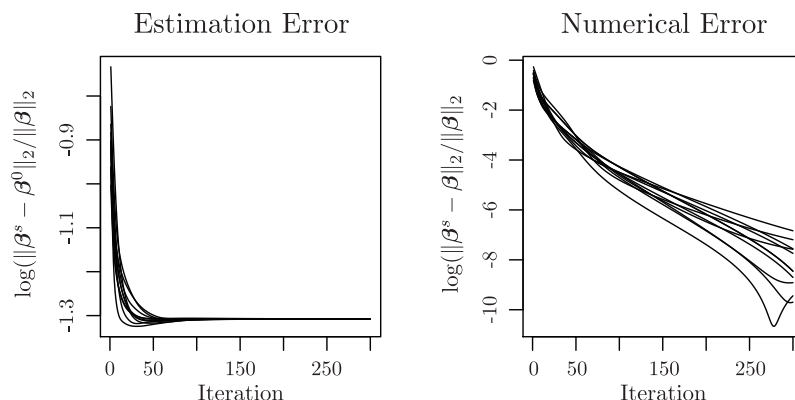


Figure 1. Results of experiments performed to assess the convergence properties of the conditional scores iteration scheme for logistic regression. The left plot shows the logarithm of the ℓ_2 distance to the true regression coefficient β^0 as a function of the iteration count, and the right plot shows the distance to the estimate obtained in the first run, for 10 consecutive runs with random starting points.

for $s = 1, 2, \dots$ until convergence.

We performed a simple experiment, similar to Loh and Wainwright (2012, Fig. 2), to assess the convergence properties of this scheme. Setting $n = 100$, $p = 500$, and $\beta^0 = (1, 1, 1, 1, 1, 0, \dots, 0)$, we generated a matrix $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. A response vector \mathbf{y} with elements $y_i \sim B(1, H(\mathbf{x}'_i \beta^0))$, $i = 1, \dots, n$, was then generated, as well as a measurement matrix $\mathbf{W} = \mathbf{X} + \mathbf{U}$, where $\mathbf{U} \sim \mathcal{N}(\mathbf{0}, (0.2)\mathbf{I}_p)$. Setting $\kappa = \|\beta^0\|_1/2$, we ran 300 iterations with stepsize $\alpha = 0.01$, obtaining an estimate of β . We then repeated this procedure 10 times, each time with a random initial value β^1 . The left plot in Figure 1 shows the logarithm of the relative estimation error in each of the 10 runs. Starting out at different values, we see that they all converge to a value around -1.3 . The right plot shows the logarithm of the ℓ_2 distance between each of the 10 iterates with random starting points, and the estimate obtained in the first run. This numerical error gets small as the number of iterations increases, while the estimation error stabilizes. Similar results were obtained for different values of κ , and for different problem dimensions n and p . This suggests that local optima do not pose a problem here.

5.2. Poisson regression

Poisson regression is used when the outcome can be modeled by a Poisson process, $y_i \sim \text{Pois}(e^{\eta})$, $i = 1, \dots, n$. An example with high-dimensional data is given by Huang et al. (2010), who define the *spatial lasso* that is applied with a

Poisson regression model to study the distribution of tree species in a geographic area. For additive measurement error, the terms in the modified density (5.1) are

$$\begin{aligned}\eta_* &= \mu + \boldsymbol{\beta}'(\mathbf{w} + y\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta}), \\ c_*(y, \boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta}) &= -\log(y!) - \left(\frac{y^2}{2}\right)\boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta}, \\ \mathcal{D}_*(\eta_*, \boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta}) &= \log \left\{ \sum_{z=0}^{\infty} (z!)^{-1} \exp \left\{ z\eta_* - \left(\frac{z^2}{2}\right)\boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta} \right\} \right\}, \\ \frac{\partial \mathcal{D}_*}{\partial \eta_*} &= \frac{\sum_{z=0}^{\infty} z(z!)^{-1} \exp \{ z\eta_* - (z^2/2)\boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta} \}}{\sum_{z=0}^{\infty} (z!)^{-1} \exp \{ z\eta_* - (z^2/2)\boldsymbol{\beta}'\boldsymbol{\Sigma}_{uu}\boldsymbol{\beta} \}}.\end{aligned}$$

Hence, the iteration scheme (5.3)–(5.4) for Poisson regression involves numerical approximation of the infinite sums in $\partial \mathcal{D}_*/\partial \eta_*$, but is otherwise straightforward.

6. Experiments

6.1. Linear regression

We present the results of simulations comparing the naive lasso (2.2) to the corrected lasso for linear models. The constrained version of the corrected lasso (4.2) was used to avoid dealing with more than one regularization parameters. In all simulations, the number of samples was $n = 100$, the number of covariates $p = 500$, and the sparsity index was either $s_0 = 5$ or $s_0 = 10$. The measurement error covariance $\boldsymbol{\Sigma}_{uu}$ is assumed known here. Several different covariance matrices $\boldsymbol{\Sigma}_{xx}$ and $\boldsymbol{\Sigma}_{uu}$ were used, and the overall simulation procedure was as follows.

A random set of indices $S_0 \subset \{1, \dots, p\}$ with cardinality s_0 was generated. The corresponding nonzero entries of $\boldsymbol{\beta}^0$ were then generated by drawing s_0 i.i.d. values from $\mathcal{N}(0, 2^2)$. The matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with rows distributed according to $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{xx})$ was generated. The response $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^0 + \boldsymbol{\epsilon}$ was sampled with $\boldsymbol{\epsilon}$ i.i.d. drawn from $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$ with $\sigma = 0.1$, and \mathbf{y} had its mean subtracted to avoid estimating the intercept. A measurement matrix $\mathbf{W} = \mathbf{X} + \mathbf{U}$ was generated with the rows of \mathbf{U} i.i.d. distributed according to $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{uu})$, and \mathbf{W} had its mean subtracted.

The naive lasso estimate $\hat{\boldsymbol{\beta}}_L$ was computed using the R package GLMNET (Friedman, Hastie, and Tibshirani (2010)), choosing the regularization level λ corresponding to the minimum of the 10-fold cross-validation curve using the `cv.glmnet` function with default parameters. The corrected lasso estimate $\hat{\boldsymbol{\beta}}_{CL}$ was computed by 10-fold cross-validation with 100 candidate constraint parameters κ equally spaced in the range $[10^{-3}R, R]$, where $R = 2\|\hat{\boldsymbol{\beta}}_L\|_1$. The final value of κ was chosen to minimize the cross-validated loss.

The whole procedure was repeated 200 times for each experiment. Tables 1–3 summarize the simulation results. TP (true positives) denotes the number of

Table 1. Comparison of naive and corrected lasso for linear regression, when $\Sigma_{xx} = \mathbf{I}_p$ and $\Sigma_{uu} = \sigma_u^2 \mathbf{I}_p$.

		TP	FP	$\ \hat{\beta} - \beta^0\ _2$	$\ \hat{\beta} - \beta^0\ _1$
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.0$	Naive	4.95 (0.02)	5.18 (0.52)	0.11 (0.00)	0.25 (0.01)
	Corrected	4.13 (0.06)	17.57 (0.47)	0.98 (0.03)	3.23 (0.12)
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	4.18 (0.06)	24.82 (1.35)	1.56 (0.04)	5.03 (0.20)
	Corrected	3.62 (0.06)	11.81 (0.32)	1.45 (0.04)	4.20 (0.14)
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	3.78 (0.07)	23.11 (1.37)	2.23 (0.06)	6.44 (0.21)
	Corrected	3.62 (0.06)	11.81 (0.32)	1.45 (0.04)	4.20 (0.14)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.0$	Naive	9.82 (0.03)	4.82 (0.32)	0.20 (0.00)	0.64 (0.01)
	Corrected	7.16 (0.09)	19.77 (0.39)	2.05 (0.05)	7.91 (0.20)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	7.46 (0.09)	34.27 (1.26)	2.80 (0.05)	11.67 (0.29)
	Corrected	7.16 (0.09)	19.77 (0.39)	2.05 (0.05)	7.91 (0.20)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	6.35 (0.10)	27.98 (1.21)	3.80 (0.07)	14.29 (0.31)
	Corrected	5.76 (0.10)	11.97 (0.28)	3.01 (0.08)	10.28 (0.28)

nonzero covariates which were correctly selected by the procedure and FP (false positives) denotes the number of irrelevant covariates which were selected by the procedure, and the two rightmost columns denote the estimation error as measured in the ℓ_2 - and ℓ_1 -norm, respectively. All results are averages over the 200 Monte Carlo simulations, and the numbers in parentheses are the corresponding standard errors.

In our first simulation experiment, all elements of \mathbf{X} and \mathbf{U} were i.i.d. Gaussian, with $\Sigma_{xx} = \mathbf{I}_p$ and $\Sigma_{uu} = \sigma_u^2 \mathbf{I}_p$, where the measurement error variance σ_u^2 was either 0.2 or 0.4. For comparison, the standard lasso in the absence of measurement error ($\sigma_u^2 = 0.0$) was also computed. Table 1 summarizes the results for both $s_0 = 5$ and $s_0 = 10$. In all cases shown, the naive lasso makes a very large number of false selections compared to the corrected lasso. The estimation errors of the corrected lasso are also consistently smaller than the naive lasso. The naive lasso is slightly better than the corrected lasso in detecting the relevant covariates. The lasso without measurement error is also seen to have a higher true positive rate and lower false positive rate than any of the cases with measurement error, indicating that measurement error makes the covariate selection problem considerably harder.

Next, a block diagonal Σ_{xx} was considered, with 10 blocks $\mathbf{B}_1, \dots, \mathbf{B}_{10} \in \mathbb{R}^{50 \times 50}$ along the diagonal, and all other elements equal to zero. Each block had a Toeplitz structure, with (j, k) th element given by $(\mathbf{B}_l)_{(j,k)} = 0.8^{|j-k|}$, $l = 1, \dots, 10$, $j, k = 1, \dots, 50$. Hence, 10 groups of 50 covariates each were correlated with each other, but not with the covariates outside the group. The measurement errors were assumed i.i.d. Gaussian with $\Sigma_{uu} = \sigma_u^2 \mathbf{I}_p$. This is a plausible model for gene expression data, with the blocks corresponding to genes within a functional group or pathway having strong correlation (Tai and Pan (2007)), and the measurement error corresponding to noisy measurements.

Table 2. Comparison of naive and corrected lasso for linear regression, when Σ_{xx} is block structured and $\Sigma_{uu} = \sigma_u^2 \mathbf{I}_p$.

		TP	FP	$\ \hat{\beta} - \beta^0\ _2$	$\ \hat{\beta} - \beta^0\ _1$
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.0$	Naive	4.95 (0.02)	6.45 (0.48)	0.13 (0.00)	0.31 (0.01)
	Corrected	3.75 (0.06)	15.50 (0.40)	1.32 (0.05)	4.02 (0.17)
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	4.03 (0.06)	31.58 (1.18)	2.22 (0.06)	7.41 (0.23)
	Corrected	3.17 (0.07)	10.43 (0.26)	1.85 (0.07)	5.08 (0.19)
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	3.82 (0.06)	34.56 (1.28)	2.79 (0.07)	9.29 (0.28)
	Corrected	3.17 (0.07)	10.43 (0.26)	1.85 (0.07)	5.08 (0.19)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.0$	Naive	9.73 (0.04)	11.13 (0.51)	0.27 (0.01)	0.92 (0.02)
	Corrected	6.07 (0.11)	17.27 (0.34)	2.62 (0.07)	9.62 (0.26)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	7.01 (0.09)	41.85 (1.18)	3.57 (0.06)	15.63 (0.32)
	Corrected	6.07 (0.11)	17.27 (0.34)	2.62 (0.07)	9.62 (0.26)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	6.40 (0.10)	43.88 (1.31)	4.51 (0.08)	19.38 (0.40)
	Corrected	4.38 (0.10)	11.38 (0.24)	4.00 (0.12)	13.43 (0.40)

Table 2 shows the simulation results. It is clear that the correlations imposed make this a harder problem. In particular the corrected lasso shows a weaker performance in detecting the relevant covariates, but the naive lasso makes an even larger number of false positive selections.

Finally, the entries of \mathbf{X} were i.i.d. Gaussian, $\Sigma_{xx} = \mathbf{I}_p$, while the measurement errors had correlations $(\Sigma_{uu})_{(j,k)} = \rho_u^{1+|j-k|}$. The diagonal $\rho_u = \sigma_u^2$ took either the value 0.2 or 0.4. Table 3 shows the simulation results. The setting without measurement error is here equivalent to the one in Table 1, and therefore omitted. In this setting, it is really clear that the corrected lasso performs better than the naive approach. At the cost of a slight reduction in the number of correct selections, the corrected lasso substantially reduces the number of false positive selections. The estimation errors of the corrected lasso are consistently smaller than those of the naive lasso.

Overall, the reduction in false positive selections when using the corrected lasso compared to the naive lasso, was between 24% and 74%.

6.2. Logistic regression

We performed simulation experiments to investigate the merit of the conditional scores lasso for logistic regression, outlined in Section 5.1. The setup was similar to the one described for linear regression in the last section, except that cross-validation was not performed, due to the lack of a loss function. We compared the naive lasso solution for logistic regression

$$\hat{\beta}_{naive}(\kappa) = \arg \min_{\beta: \|\beta\|_1 \leq \kappa} \left\{ \sum_{i=1}^n y_i \mathbf{w}'_i \beta + \log(1 - H(\mathbf{w}'_i \beta)) \right\} \quad (6.1)$$

to the corrected estimate $\hat{\beta}_{corr}(\kappa)$ obtained using the conditional scores algorithm, over a range of candidate κ values.

Table 3. Comparison of naive and corrected lasso for linear regression, when $\Sigma_{xx} = \mathbf{I}_p$ and Σ_{uu} is a Toeplitz matrix with elements $(\Sigma_{uu})_{j,k} = \rho_u^{1+|j-k|}$, $j, k = 1, \dots, p$, and $\rho_u = \sigma_u^2$.

		TP	FP	$\ \hat{\beta} - \beta^0\ _2$	$\ \hat{\beta} - \beta^0\ _1$
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	4.19 (0.06)	23.86 (1.12)	1.63 (0.04)	5.09 (0.16)
	Corrected	4.13 (0.06)	18.02 (0.49)	1.04 (0.03)	3.53 (0.11)
$s_0 = 5, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	3.79 (0.06)	24.37 (1.46)	2.30 (0.06)	6.88 (0.25)
	Corrected	3.58 (0.07)	11.87 (0.30)	1.50 (0.05)	4.34 (0.16)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.2$	Naive	7.36 (0.10)	34.77 (1.46)	2.89 (0.05)	12.08 (0.29)
	Corrected	7.16 (0.09)	19.98 (0.39)	2.11 (0.04)	8.16 (0.18)
$s_0 = 10, \sigma_{\mathbf{U}}^2 = 0.4$	Naive	6.49 (0.10)	32.38 (1.42)	3.80 (0.06)	14.91 (0.32)
	Corrected	5.79 (0.09)	11.86 (0.25)	2.98 (0.07)	10.14 (0.25)

We used a sample size $n = 100$, with $p = 500$ covariates, of which either $s_0 = 5$ or $s_0 = 10$ were nonzero. We considered i.i.d. Gaussian covariates and measurement errors, with $\Sigma_{xx} = \mathbf{I}_p$ and $\Sigma_{uu} = (0.2)\mathbf{I}_p$. Since both covariance matrices were diagonal, we did not randomize over the indices in S_0 , and simply let $\beta_j \neq 0$ for $j = 1, \dots, s_0$. The matrices \mathbf{X} and \mathbf{U} were drawn from $\mathcal{N}(0, \Sigma_x)$ and $\mathcal{N}(0, \Sigma_{uu})$, respectively. The values of $\beta_{S_0}^0$ were generated by drawing s_0 i.i.d. values from $\mathcal{N}(0, 5^2)$, and the responses y_i , $i = 1, \dots, n$ were sampled from a binomial distribution with mean $H(\mathbf{x}_i' \beta^0)$. The procedure was repeated 200 times.

Figure 2 shows the receiver operating characteristics (ROC) curve for the simulations. With $s_0 = 5$, the conditional scores lasso ('Corrected') is seen to perform better covariate selection than the standard lasso for logistic regression ('Naive'). In the $s_0 = 10$ case, the conditional scores lasso is also better, but only marginally. Figure 3 shows the ℓ_1 estimation error over a range of values of the regularization parameter κ . Here, the corrected lasso clearly has a lower estimation error than the naive approach when $s_0 = 5$ or $s_0 = 10$. The simulations suggest that the conditional scores lasso is a useful method for measurement error correction in logistic regression when $p > n$.

6.3. Microarray data

We present an application of the conditional scores lasso for logistic regression to an Affymetrix microarray data set publicly available from the Array-Express database (www.ebi.ac.uk/arrayexpress) under accession number E-GEOD-10320. The data set contains gene expression measurements of 144 favorable histology Wilms tumors (FHWT), 53 of which did relapse (cases) and 91 of which did not relapse (controls). For the Affymetrix microarrays used, each gene expression was measured by multiple probes. The Bayesian Gene Expression




Figure 2. The plots show ROC curves over a range over regularization parameter for the conditional scores lasso ('Corrected') and the standard lasso for logistic regression ('Naive'), for the $s_0 = 5$ case (left) and the $s_0 = 10$ case (right).




Figure 3. The plots show ℓ_1 estimation error over a range over regularization parameter for the conditional scores lasso ('Corrected') and the standard lasso for logistic regression ('Naive'), for the $s_0 = 5$ case (left) and the $s_0 = 10$ case (right).

(BGX) Bioconductor package (Hein et al. (2005)) utilizes these replicate measurements to form posterior distributions of the mean gene expression, measured on the log scale, of each gene for each sample.

In our additive measurement error model, the mean posterior gene expres-

sions $\hat{\mu}_{ij}$ are used as estimates of the ‘true’ gene expressions x_{ij} , for $i = 1, \dots, n$, $j = 1, \dots, p$. Letting $\hat{\boldsymbol{\mu}}_j = (\hat{\mu}_{1j}, \dots, \hat{\mu}_{nj})'$, $\bar{\mu}_j = (1/n) \sum_{i=1}^n \hat{\mu}_{ij}$ and $\hat{\sigma}_j^2 = (1/n) \sum_{i=1}^n (\hat{\mu}_{ij} - \bar{\mu}_j)^2$, the standardized design matrix \mathbf{W} now has entries

$$w_{ij} = \frac{\hat{\mu}_{ij} - \bar{\mu}_j}{\hat{\sigma}_j}, \quad i = 1, \dots, n, \quad j = 1, \dots, p.$$

We let $\text{var}(\hat{\mu}_{ij})$ denote the posterior variance of gene expression estimate $\hat{\mu}_{ij}$. Assuming equal measurement error variance across samples, but not across covariates, we estimate the measurement error of gene j by $\hat{\sigma}_{u,j}^2 = (1/n) \sum_{i=1}^n \text{var}(\hat{\mu}_{ij})$, $j = 1, \dots, p$. For simplicity, covariance of measurement errors are not considered, so the final estimate $\hat{\boldsymbol{\Sigma}}_{uu}$ on the scale of the standardized \mathbf{W} has entries

$$\left(\hat{\boldsymbol{\Sigma}}_{uu}\right)_{j,k} = \begin{cases} \frac{\hat{\sigma}_{u,j}^2}{\hat{\sigma}_j^2}, & \text{if } j = k, \\ 0, & \text{if } j \neq k, \end{cases}$$

for $j, k = 1, \dots, p$. In many cases, the estimated measurement error variance $\hat{\sigma}_{u,j}^2$ is large compared to the between-sample variance of the means $\hat{\sigma}_j^2$, and for these cases little can be done. We therefore chose to analyze only the $p = 1,857$ genes for which $\hat{\sigma}_{u,j}^2 < (1/2)\hat{\sigma}_j^2$, out of the original 20,931. For the 1,857 selected genes, the naive lasso estimate $\hat{\boldsymbol{\beta}}_L$ was computed by ten-fold cross-validation using GLMNET (Friedman, Hastie, and Tibshirani (2010)), yielding 22 nonzero coefficients.

Since the conditional scores lasso lacks a well-defined loss function, the elbow rule (Rosenbaum and Tsybakov (2010, Fig. 1)) was used to choose the regularization level. The conditional scores solution was computed for a grid of constraint values between $3\|\hat{\boldsymbol{\beta}}_L\|$ and $(0.1)\|\hat{\boldsymbol{\beta}}_L\|$, with spacing $(0.1)\|\hat{\boldsymbol{\beta}}_L\|$. Figure 4 shows the number of nonzero coefficient estimates plotted versus the constraint level; the elbow rule now amounts to selecting κ where the curve begins to be flat. The plot in Figure 4 is surprisingly good: the number of selected covariates is between 11 and 10 for all constraint values between $\kappa = (1.5)\|\hat{\boldsymbol{\beta}}_L\|$ and $\kappa = (0.5)\|\hat{\boldsymbol{\beta}}_L\|$.

Based on Figure 4, $\kappa_1 = (1.5)\|\hat{\boldsymbol{\beta}}_L\|$ was chosen as our optimal constraint level. At κ_1 , the ℓ_1 norm of the estimated coefficient vector is 1.5 times that of the naive estimate, while selecting only half as many covariates. Figure 5 illustrates this by plotting the coefficient estimates of the naive lasso and the conditional scores lasso. The impact of the measurement error correction clearly is to amplify the coefficients of some seemingly important genes, while the naive lasso has many coefficients with magnitudes of the same order. This is an analogue to the measurement error attenuation in standard linear regression.

7. Discussion




Figure 4. Illustration of the elbow rule for the conditional scores lasso. The number of nonzero coefficients is plotted against the constraint level, here as a fraction of the ℓ_1 norm of the naive lasso estimate.

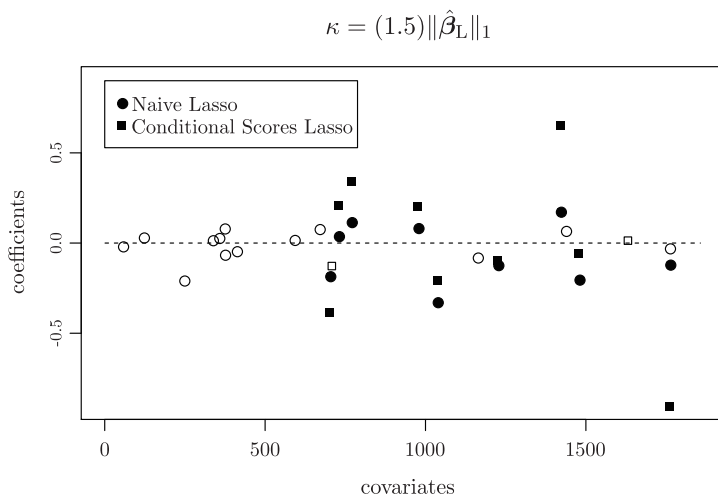


Figure 5. Comparison of the coefficient estimates of the 22 covariates selected by the cross-validated naive lasso and the 11 covariates selected by the conditional scores lasso. The covariates selected by both methods are plotted with filled circles/squares, whereas those selected by only one of the methods are plotted with empty circles/squares.

In this paper, we have shown how linear regression with the lasso is affected by additive measurement error. In particular, standard results for consistency of estimation and covariate selection no longer hold when the covariates are subject to measurement error. A simple correction method was considered, studied earlier by Loh and Wainwright (2012). Our finite sample results show conditions

under which this corrected lasso is a sign consistent covariate selector. Asymptotically, sign consistent covariate selection with the corrected lasso requires conditions very similar to the lasso in the absence of measurement error. In contrast, asymptotically sign consistent covariate selection with the naive lasso essentially requires the relevant and the irrelevant covariates to be uncorrelated. We also suggest a conditional scores approach for correcting for measurement error in ℓ_1 -constrained GLMs; it shows promising empirical results. Using the iteration scheme suggested by Loh and Wainwright (2012) for linear models, corrected lasso estimates for GLMs are computed efficiently even when $p \gg n$.

Simulation results confirm that ignoring measurement error can yield a large number of false positive selections. The same was observed by Rosenbaum and Tsybakov (2010) for censored and missing data. Correction for measurement yields a much sparser fit, while finding almost as many of the relevant covariates. An application to microarray data agrees well with the simulations. Measurement error correction yields a substantially sparser model, and the elbow rule (Figure 4) works well for finding a good constraint level for the conditional scores lasso.

Acknowledgement

The authors thank Po-Ling Loh and Martin J. Wainwright for sharing the Matlab code used to compute the corrected lasso estimates. Thanks to Sylvia Richardson for discussions and for suggesting the use of BGX output data, and to Bin Yu for discussions. This work has been supported by funding from the Norwegian Research Council through the centre Statistics for Innovation in Oslo.

References

- Ayers, K. L. and Cordell, H. J. (2010). SNP selection in genome-wide and candidate gene studies via penalized logistic regression. *Genet. Epidemiol.* **34**, 879-891.
- Benjamini, Y. and Speed, T. P. (2012). Summarizing and correcting the GC content bias in high-throughput sequencing. *Nucl. Acids Res.* **40**, e72.
- Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-dimensional Data*. Springer, Heidelberg.
- Candès, E. and Tao, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.* **35**, 2313-2351.
- Carroll, R. J., Ruppert, D., Stefanski, L. A. and Crainiceanu, C. M. (2006). *Measurement Error in Nonlinear Models*. Chapman and Hall/CRC, Boca Raton.
- Chen, Y. and Caramanis, C. (2013). Noisy and missing data regression: Distribution-oblivious support recovery. *J. Mach. Learn. Res. W&CP* **28**, 383-391.
- Duchi, J., Shalev-Shwartz, S., Singer, Y. and Chandra, T. (2008). Efficient projections onto the ℓ_1 -ball for learning in high dimensions. *Proceedings of the 25th International Conference on Machine Learning*, 272-279.

- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96**, 1348-1360.
- Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**, 1-22.
- Hanfelt, J. J. and Liang, K.-Y. (1997). Approximate likelihoods for generalized linear errors-in-variables models. *J. Roy. Statist. Soc. Ser. B* **59**, 627-637.
- Hein, A.-M. K., Richardson, S., Causton, H. C., Ambler, G. K. and Green, P. J. (2005). BGX: a fully Bayesian integrated approach to the analysis of Affymetrix GeneChip data. *Bio-statistics* **6**, 349-373.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12**, 55-67.
- Huang, H.-C., Hsu, N.-J., Theobald, D. M. and Breidt, F. J. (2010). Spatial lasso with applications to GIS model selection. *J. Comput. Graph. Statist.* **19**, 963-983.
- Knight, K. and Fu, W. (2000). Asymptotics of lasso-type estimators. *Ann. Statist.* **28**, 1356-1378.
- Liang, H. and Li, R. (2009). Variable selection for partially linear models with measurement errors. *J. Amer. Statist. Assoc.* **104**, 234-248.
- Loh, P.-L. and Wainwright, M. J. (2012). High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. *Ann. Statist.* **40**, 1637-1664.
- Ma, Y. and Li, R. (2010). Variable selection in measurement error models. *Bernoulli* **16**, 274-300.
- Meinshausen, N. and Bühlmann, P. (2010). Stability selection. *J. Roy. Statist. Soc. Ser. B* **72**, 417-473.
- Purdom, E. and Holmes, S. P. (2005). Error distribution for gene expression data. *Stat. Appl. Genet. Mol. Biol.* **4**.
- Rocke, D. M. and Durbin, B. (2001). A model for measurement error for gene expression arrays. *J. Comput. Biol.* **8**, 557-569.
- Rosenbaum, M. and Tsybakov, A. B. (2010). Sparse recovery under matrix uncertainty. *Ann. Statist.* **38**, 2620-2651.
- Rosenbaum, M. and Tsybakov, A. B. (2013). Improved matrix uncertainty selector. *From Probability to Statistics and Back: High-dimensional Models and Processes* **9**, 276-290.
- Stefanski, L. A. and Carroll, R. J. (1987). Conditional scores and optimal scores for generalized linear measurement-error models. *Biometrika* **74**, 703-716.
- Sørensen, Ø, Frigessi, A. and Thoresen, M. (2014). Online supplementary material to “Measurement error in lasso: Impact and likelihood bias correction”. Available at <http://www.stat.sinica.edu.tw/statistica>.
- Tai, F. and Pan, W. (2007). Incorporating prior knowledge of predictors into penalized classifiers with multiple penalty terms. *Bioinformatics*, **23**, 1775-1782.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58**, 267-288.
- Wu, T. T., Chen, Y. F., Hastie, T., Sobel, E. and Lange, K. (2009). Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics* **25**, 714-721.
- Xu, Q. and You, J. (2007). Covariate selection for linear errors-in-variables regression models. *Comm. Statist. Theory Methods* **36**, 375-386.

- Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *J. Mach. Learn. Res.* **7**, 2541-2563.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101**, 1418-1429.

Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, P.O box 1110 Blindern 0317 OSLO, Norway.

E-mail: oystein.sorensen@medisin.uio.no

Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, and Research Support Services, Oslo University Hospital, P.O box 1110 Blindern 0317 OSLO, Norway.

E-mail: arnoldo.frigessi@medisin.uio.no

Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, P.O box 1110 Blindern 0317 OSLO, Norway.

E-mail: magne.thoresen@medisin.uio.no

(Received July 2013; accepted May 2014)