

Estimation Stability with Cross Validation (*ESCV*)

Chinghway Lim

Department of Statistics and Applied Probability
National University of Singapore

Bin Yu

Departments of Statistics and EECS
University of California, Berkeley

March 14, 2013

Abstract

Cross-validation (*CV*) is often used to select the regularization parameter in high dimensional problems. However, when applied to the sparse modeling method Lasso, *CV* leads to models that are unstable in high-dimensions, and consequently not suited for reliable interpretation. In this paper, we propose a model-free criterion *ESCV* based on a new *estimation stability* (*ES*) metric and *CV*. Our proposed *ESCV* finds a smaller and locally *ES*-optimal model smaller than the *CV* choice so that the it fits the data and also enjoys estimation stability property. We demonstrate that *ESCV* is an effective alternative to *CV* at a similar easily parallelizable computational cost. In particular, we compare the two approaches with respect to several performance measures when applied to the Lasso on both simulated and real data sets. For dependent predictors common in practice, our main finding is that, *ESCV* cuts down false positive rates often by a large margin, while sacrificing little of true positive rates. *ESCV* usually outperforms *CV* in terms of parameter estimation while giving similar performance as *CV* in terms of prediction. For the two real data sets from neuroscience and cell biology, the models found by *ESCV* are less than half of the model sizes by *CV*. Judged based on subject knowledge, they are more plausible than those by *CV* as well. We also discuss some regularization parameter alignment issues that come up in both approaches.

Keywords: Lasso, model selection, parameter estimation, prediction.

1 Introduction

1.1 Regularization Methods

There is an ever increasing amount of data in all fields of science and engineering. Often, this data comes in high dimensions relative to the sample size, posing a new challenge to scientists, engineers, and decision makers. These problems, plagued by the curse of dimensionality, suffer from overfitting when classical methods are applied. Regularization methods are used to tackle this problem of overfitting head on, usually by imposing a penalty on the complexity of the solution or through early stopping. For example, in fitting the usual linear regression model, the Lasso (Tibshirani 1996) and ridge regression (Tikhonov 1943; Hoerl 1962) adds a L_1 and L_2 penalty on the coefficient estimates respectively to the usual least squares fit objective function. Regularization methods can also take the form of early stopping iterative algorithms like classical forward selection or L_2 -Boosting (Friedman 2001; Bühlmann and Yu 2003; Zhang and Yu 2012; Zhang 2011). Common to these methods is that they provide a family of possible estimators instead of just one estimator, with the unregularized solution at one end of the spectrum. This family is indexed by a regularization parameter and is commonly referred to as the solution path. For the Lasso and ridge regression, this regularization parameter determines the extent of the respective penalties. For the iterative algorithms, this parameter corresponds to the number of steps they take. Despite the difference in nature, numerous works have shown these regularization methods, at least in the context of the linear model, are intrinsically related (Efron et al. 2004; Zhao and Yu 2007; Meinshausen et al. 2007). In that light, we will not focus on the distinction between the different types of regularization parameters but instead simply use λ as a catch-all representation for them. In the same vein, we focus on the Lasso in this chapter even though we believe the method we present will work in the general framework.

1.2 Selecting the Regularization Parameter λ

Much work has been done to show that regularization methods yield desirable solutions in high dimensional problems. For example, the popular Lasso has been shown to be L_2 -consistent (Meinshausen and Yu 2009; Bickel et al. 2009) and model selection consistent (Meinshausen and Bühlmann 2006; Zhao and Yu 2006; Tropp 2006; Wainwright 2009) in the high dimensional setting

when respective conditions are met. These results guarantee the existence of the λ needed, but offer little guidance on how to find the desired λ in practice. Indeed, data-driven regularization parameter selection with guaranteed theoretical performance turns out to be a particularly difficult problem.

One can rely on traditional model selection criteria like Akaike's information criterion (*AIC*) (Akaike 1974) and Bayesian information criterion (*BIC*) (Schwarz 1978). They are easy to compute but their validity rely on model assumptions. Furthermore, they are derived from asymptotic results, so even when model assumptions are satisfied, they may not work well in the finite sample case.

More commonly used today are model-free approaches like cross-validation (*CV*) (Allen 1974; Stone 1974) and bootstrap methods (Efron 1979; Zhang 1993; Shao 1996). They have become computationally feasible for increasingly large data sets with the rapid advancements in computing power, especially the parallel computing paradigm that is currently the platform for dealing with big data. These methods rely on data resampling to assess prediction error of candidate solutions and can be found in various statistics and machine learning literature (Hastie et al. 2002; Breiman 1995, 1996, 2001). In particular, it is the most popular approach for regularization methods to select λ . Doing so often leads to estimators with good predictive performance when the sample size is not small. However, there are other performance metrics that are also of interest in statistics, among them parameter estimation and variable selection metrics, with important practical connections. Unsurprisingly, optimizing predictive performance does not necessarily translate to having success with respect to these other performance metrics.

1.3 Estimation Stability

Statistical estimation is often tied to the optimization of an empirical loss or a random function based on data. Take for example, when fitting a linear model for random variables $X \in \mathbb{R}^p, Y \in \mathbb{R}$, one might want to minimize the predictive L_2 loss,

$$f(\beta) = E_{X,Y}(Y - X'\beta)^2.$$

However, since the underlying joint distribution of (X, Y) is unknown, we instead minimize the empirical loss

$$\hat{f}(\beta) = \frac{1}{n} \sum_{i=1}^n (y_i - x_i' \beta)^2,$$

where (x_i, y_i) for $i = 1, \dots, n$, are the observed samples of (X, Y) . By minimizing \hat{f} instead of f , we incur a random estimation error dependent on the sample we observed. In the classical ideal scenario, e.g. when the sample contain independent and identically distributed observations and the sample size n is large and p is small, this estimation error incurred is small. If we draw multiple samples from (X, Y) , each resulting estimate will be close to that of minimizing f , and consequently close to each other. This closeness across different samples can be seen as a form of stability in the estimation procedure, and we call it *estimation stability*.

When the differences across different samples are measured by the L_2 error, the estimation stability is obviously related to variance. We opt to use the term “stability” rather than the more commonly used term “variability” in statistics. This is to recognize the fact that stability is a concept broader than variance or variability and that it is used in other quantitative fields such as numerical analysis, dynamical systems, and linear analysis (Higham 1996; Salle 1976; Ellis 1998). Stability is also not associated with a particular metric (unlike variance) and thus allows its consideration under different metrics. In a recent paper (Yu 2013 (to appear)), we advocate an enhanced emphasis on stability in statistical inference, especially for large and high dimensional data for which instability of statistical methods is much more common than in the domain of classical statistics.

It is clear that estimation stability is a necessary property for a reasonable estimation procedure: the solution is not meaningful if it varies considerably from sample to sample. The converse certainly cannot be true in general: an arbitrary constant estimate will not vary but is certainly meaningless. When we add a data faithfulness requirement through cross-validation, we are able to devise a model-free criterion based on estimation stability for the selection of the regularization parameter λ . Specifically, our proposed new criterion of estimation stability cross validation (*ESCV*) combines a new metric of estimation stability (*ES*) with *CV*. For a given L_1 norm τ or a regularization parameter λ , our new $ES(\tau)$ metric is the reciprocal of a test statistic for testing the null hypothesis that the regression function is zero. The test statistic is an estimate of the regression function standardized by an approximate delete-d Jackknife standard error estimate

based on the same pseudo data sets as in CV , and both estimates are functions of τ . The proposed $ESCV$ criterion chooses a local minimum of $ES(\tau)$ which is smaller than the selection of τ by CV . It is worth noting that the computational cost of $ESCV$ is similar to that of CV and that they are both well suited to parallel computation, the dominant computing platform for big data.

We demonstrate that our criterion $ESCV$ provides a viable alternative to CV and BIC . We compare the three approaches with respect to several performance metrics when applied to the Lasso on both simulated data sets with different predictor dependence set-ups and two real data sets. These performance metrics are L_2 error for parameter estimation, prediction error, F -measure and model size for model selection performance. We find that our criterion compares favorably with CV and BIC where they are known to excel, and outperforms them in other scenarios over different performance criteria. In particular, $ESCV$ obtains excellent model selection results that are substantially better than those from CV , both in simulations and our real data sets, where the results are validated by subject knowledge. When the predictors are correlated, which is often the case in practice, $ESCV$ also often outperforms CV for parameter estimation while at same time provides prediction errors comparable to those of CV .

We note that previous works based on stability of solutions have shown positive results in terms of model selection (Breiman 1996; Bach 2008; Meinshausen and Bühlmann 2010). The work here differs from them in three substantial ways. Firstly, we develop a different measure of stability ES that is closely related to estimation than model selection, even though our $ESCV$ does have desirable model selection properties such as the best F -measures across all simulation set-ups in Section 3. Secondly, we restrict our attention to selecting the regularization parameter. Even though we evaluate our choice by the performance of the corresponding solution, our focus remains on determining the right amount of regularization. We do not introduce any further tuning parameters as in Meinshausen and Bühlmann (2010). Concurrent with and independent of our work, recent follow-up papers on (Meinshausen and Bühlmann 2010) use model selection stability to select edges in graphical models (Liu et al. 2010; Haury et al. 2012) or modify stability model selection to improve its false discovery rate theoretical properties (Shah and Samworth 2013). The former two papers introduce further tuning parameters and they recommend fixed values for them. Shah and Samworth (2013) employs the complementary half-sample data perturbation scheme.

ESCV can work on such a scheme, but doing so would depart from the usual implementation of *CV* for comparison purposes. Thirdly, as in Meinshausen and Bühlmann (2010), these three papers apply data perturbation schemes such as bootstrap and subsampling with hundreds or thousands runs of model fitting. On the contrary, the *CV* (and *ESCV*) data perturbation scheme often works well based on 5-10 runs of model fitting.

2 Methodology

2.1 Lasso and Pseudo Solutions

Let $X \in \mathbb{R}^{n \times p}$, $Y \in \mathbb{R}^n$ be our data set. The Lasso generates a family of solutions,

$$\hat{\beta}[\lambda] = \arg \min_{\beta} \{ \|Y - X\beta\|_2^2 + \lambda \|\beta\|_1 \}.$$

$\hat{\beta}[\lambda]$, as a function of $\lambda \geq 0$ is also known as the Lasso solution path for β_j ($j = 1, \dots, p$). We want to select a solution from this solution path; that is, choose a λ and take its corresponding solution in the solution path. As alluded to earlier, we would like to make this choice based on estimation stability and fit.

Since the notion of estimation stability is tied to the sampling distribution of the data, it is unavoidable that we need multiple solution paths to make such an assessment. Of course, it is often costly and infeasible to obtain extra data in practice. Thankfully, this problem is not new, and there are well-established ways to get around it. The key is to exploit the existing data by employing data perturbation schemes, parlaying it into multiple data sets. Let $(X^*[k], Y^*[k])$ represents our k th pseudo data set, derived from (X, Y) . In our case, these are the cross-validation folds: we randomly partition the data into V groups and form V pseudo data sets by leaving out one group at a time. (See Section 2.6 for other data perturbation schemes.) We then get pseudo solutions,

$$\hat{\beta}[k; \lambda] = \arg \min_{\beta} \{ \|Y^*[k] - X^*[k]\beta\|_2^2 + \lambda \|\beta\|_1 \}$$

for $k = 1 \dots V$.

2.2 Alignment

For many regularization methods, there are multiple representations for the regularization parameter λ . In the case of the Lasso above, λ refers to the L_1 penalty parameter. Other popular choices to index the solution path are the L_1 -norm of the coefficient estimate, and the L_1 -norm expressed as a fraction of the L_1 -norm of the unregularized solution. Each of these representations for the solution path has its own merits, and is equivalent to the others (when non-trivial) for any single solution path.

However, care must be taken on how to most meaningfully align our solution paths, when we reference the same λ across different (pseudo) solution paths. In particular, when $n < p$, the L_1 -norm of the unregularized solution corresponds to the saturated fit and can vary a lot depending on which data points were sampled. This makes L_1 -fraction a poor choice, as the same index may correspond to very different amounts of regularization. The effect is more pronounced when the features are more correlated. Figure 1 shows a histogram of the maximum L_1 -norms for 10,000 bootstrap Lasso estimates of the base case Gaussian simulation (with $n = 100$, $p = 150$, $\sigma = 1$, $\rho = 0.5$) in Section 3.1.1. There is considerable spread: in this case, the upper decile is more than 20% more than the lower decile.

To highlight the effect of alignment on estimation performance, we compared the performance of cross-validation with the three alignments for the low noise scenarios detailed in Section 3.1.1. As shown in Table 1, aligning the solution paths with L_1 -fraction does comparatively worse than aligning with L_1 -norm or the penalty parameter. Notably, in the popular R package “lars” used in solving the Lasso efficiently, the included cross-validation code aligns with L_1 -fraction.

For *ESCV* to be proposed later, we find that there is little difference in performance when aligning with either the penalty parameter, λ or the L_1 -norm τ . In this work, since the L_1 norm of the solution is more comparable than the regularization parameter across different pseudo data sets, we opt to align with τ . Note that our main comparison target *CV* performs best with the τ alignment, as shown in Table 1. To be clear,

$$\tau = \tau(\lambda) = \|\hat{\beta}(\lambda)\|_1,$$

a one-to-one function for $0 \leq \lambda < \lambda_0$, where $\lambda_0 = \inf \left\{ \lambda \mid \|\hat{\beta}(\lambda)\|_1 = 0 \right\}$.

Indexing by τ also has the benefit of visualizing the solution growing with the index. $\hat{\beta}(\tau)$ starts at the origin when $\tau = 0$ and moves towards the unregularized solution as τ increases.




Figure 1: Empirical bootstrap distribution of maximum L_1 -norms of Lasso estimates on a typical simulated data set: a base case Gaussian simulation with $n = 100$, $p = 150$, $\sigma = 1$, $\rho = 0.5$ in Section 3.1.1.

2.3 Convergence of Pseudo Solutions

Given p -dimensional pseudo solutions $\hat{\beta}[k; \tau]$ for $k = 1, \dots, V$, we want to measure their differences or see how similar or stable they are. Computing their pair-wise L_2 errors was a natural first thing that we tried. However, we found that these errors vary too wildly to be useful even after normalization by means when there is high dependence between the components in the vector and this happens often especially when p is large. Notice that the components of an estimate of β are combined in a linear fashion through $X\beta$ to achieve our primary goal of estimating the linear regression function. Therefore we propose to compute the estimates

$$\hat{Y}[k; \tau] = X\hat{\beta}[k; \tau],$$

and study their stability.

To evaluate such stability, as mentioned earlier we need a measure for how far apart the estimates are at each τ : stable pseudo solutions should give similar estimates. One possibility is

Cross-Validation Estimation Error			
ρ	Regularization parameter	L_1 -norm	L_1 -fraction
0	0.794	0.787	0.817
0.2	0.781	0.773	0.821
0.5	0.969	0.956	1.04
0.9	1.83	1.80	1.95

Table 1: Effect of alignment on cross validation performance on the base case Gaussian simulation with $n = 100$, $p = 150$, $\sigma = 0.5$ in Section 3.1.1. The first column corresponds to the alignment based on λ , the second based on τ and the third based on the L_1 fraction. Cross-Validation performs worst when aligning with L_1 -fraction.

to look at the average pairwise squared Euclidean distance between the V estimates:

$$A(\tau) := \frac{1}{\binom{V}{2}} \sum_{k \neq j} \|\hat{Y}[k; \tau] - \hat{Y}[j; \tau]\|_2^2.$$

It is not hard to see that this is proportional to the more familiar “sample variance” formulation,

$$\widehat{\mathbf{Var}}(\hat{Y}[\tau]) = \frac{1}{V} \sum_{k=1}^V \|\hat{Y}[k; \tau] - \bar{\hat{Y}}[\tau]\|_2^2,$$

where $\bar{\hat{Y}}[\tau] = \frac{1}{V} \sum_{i=1}^V \hat{Y}[i; \tau]$.

Figure 2 shows two examples of this sample variance metric. The left panel is particularly illuminating: the pseudo solutions diverge as they grow at first but converge somewhat before diverging again. Here, *convergence* and *divergence* simply refer to the sample variance metric (which is really just the average pairwise distance) decreasing and increasing respectively. Heuristically, this behavior is exactly what one would expect if there is a “correct” amount of regularization. Different samples would take different paths towards the “correct” solution before moving away from one another due to overfitting. Hence, we might select the τ corresponding to the minimum point *after* the first negative slope. That is, we want to choose τ corresponding to the “dip”.

By doing this, we incorporate fit into our selection even though our criterion is based on stability. The *convergence* of the solution paths is key: not only does it suggest we are close to the truth, we are also gifted with estimation stability. Note that this helps us automatically exclude




Figure 2: Examples of the sample variance metric. The left panel shows an example where the metric exhibits a “dip”, representing the “convergence” of the pseudo solutions. The right panel shows an example with a much muted “dip”. It is difficult to use the sample variance metric to select a solution on the right.

τ 's where the solution paths trivially agree. We see this trivial effect in Figure 2, where the global minimum for the sample variance metric occurs where the solutions are close to zero.

However, this convergence effect is not always clear. The “dip” is not always present as shown in the example on the right panel. There you can still see the drop in gradient, but it is not clear which τ we should pick. Notice, however, that in a solution path, the norm of the solution varies with the amount of regularization (by definition in our case). Since larger solutions naturally varies more, using the sample variance metric skews the choice towards solutions with small norms. We need to bring in the concept of normalization to account for this effect.

2.4 Hypothesis Testing and the Estimation Stability Metric

In hypothesis testing, a test statistic based on data is computed and its corresponding p -value is calculated by matching the test statistic with its model-specific theoretical distribution. This test statistic often takes the form of a mean value over its estimated standard deviation, e.g. the

student's t -test. The desired outcome for the t -test, as is often the case regardless of the assumed model and p -value computation, is to have the test-statistic away from 0. The heuristic there is clear: if the hypothesized effect is real, the size of the mean value should be large compared to its estimated standard deviation.

In the same vein, our sample variance metric should be relative to the squared mean size of the corresponding solution. We define the *estimation stability metric*,

$$ES(\tau) := \frac{\widehat{\mathbf{Var}}(\hat{Y}[\tau])}{\|\hat{Y}[\tau]\|_2^2},$$

the normalized version of the sample variance metric. Figure 3 shows the corresponding ES metrics in dashed lines superimposed on the old sample variance metric. On the left, the “dip” from the sample variance metric is preserved by the ES metric. On the right, there is now a pronounced minimum we can select.




Figure 3: Examples of the sample variance metric (as in Figure 2) and the corresponding ES metric. We see that the ES metric preserves the local minimum from the sample variance metric on the left panel, and introduces one on the right panel where there was no local minimum from the sample variance metric.

The ES metric's reciprocal has exactly the form of a test-statistic. We can view the ES

selection of τ as a set of hypothesis tests. For each τ , we are testing if the fit ($\hat{Y}[\tau]$) is statistically different from fitting the null model ($E(Y) = 0$), albeit without a specified theoretical distribution. Our ES criterion of choosing the τ corresponding to the convergence of pseudo solutions, is exactly choosing $\hat{Y}[\tau]$ with locally minimal normalized variance. This in turn, is exactly choosing the solution whose ES metric has the largest reciprocal, or in our analogy, the most statistically significant solution along the path.

2.5 $ESCV$: Incorporating Cross-Validation

There is no guarantee that our ES metric would have only one local minimum. Note that unless the multiple solution paths match up perfectly, there will be a local minimum or multiple local minima. Hence, even in the case where Y bears no relation to X at all, an inadvertent minimum on the ES metric will falsely suggest the pseudo solutions are converging towards a meaningful solution. To prevent scenarios like this where ES fails, we incorporate cross-validation into our selection. We have already limited our choice of minimum ES to local minima. Here we further limit it to the local minimum of τ that is smaller than the cross-validation choice. We call this improved criterion *estimation stability with cross validation* ($ESCV$). In Section 3 on experimental results, we use a grid-search algorithm to find such a local minimum of ES as commonly done for CV . Thus $ESCV$'s computational cost is similar to that of CV and they are both easily parallelable.

We are exploiting the fact that cross-validation overselects (Leng et al. 2006; Wasserman and Roeder 2009). When ES gives a meaningful local minimum, cross-validation will likely overselect. Hence, $ESCV$ behaves like ES above. However, when Y bears no relation to X , or when the noise overwhelms the signal, cross-validation will likely choose the trivial solution correctly. In this case, $ESCV$ will follow suit and pick up the trivial solution.

Note that this has negligible additional computation cost, as we are essentially getting the cross-validation choice for free. The bulk of the computation lies in computing the multiple solution paths we already have.

2.6 Discussion on *ESCV*

Our *ES* metric is based on assessing the stability of the fitted values $\hat{Y}[\tau] = X\hat{\beta}[\tau]$ instead of the estimates $\hat{\beta}[\tau]$. This seems counter-intuitive since we are interested in a variety of performance measures, most of which are based on the quality of $\hat{\beta}[\tau]$ itself. However, we note that these performance measures only make sense if the underlying β is identifiable. To that end, there is a large volume of work showing the Lasso is model selection consistent under regularity conditions including that the smallest non-zero true parameter value is not too small compared a rate decaying in n (Meinshausen and Bühlmann 2006; Tropp 2006; Zhao and Yu 2006; Wainwright 2009). In particular, it assures us the asymptotic recovery of the underlying true β under appropriate conditions.

However, in the finite sample case, and especially when the features are highly correlated, different linear combinations of features (of a given sparsity) may give approximately equivalent fits. Under data perturbation, it is not surprising that the different solution paths choose different features. This makes any metric based on $\hat{\beta}[\tau]$ statistically unstable since V is small. Note that this does not contradict the assessment of the eventual $\hat{\beta}[\tau]$ picked since *ESCV* and *CV*, picking from the same solution path, would both suffer from any failure of the original Lasso.

In *ESCV*, we have used cross-validation folds to compute our pseudo-solutions. There are of course many other ways to generate pseudo datasets. One related approach would be to apply bootstrap sampling (Bach 2008). Here, simply sample with replacement from the original data set to generate multiple data sets. These two approaches are obvious choices, and can be applied to any estimation procedure (even those without an optimization formulation). A third choice, which applies only to penalized M -estimators such as the Lasso, is based on perturbations of the penalty (Meinshausen and Bühlmann 2010). Note that such perturbations of the penalty amount to perturbing (indirectly) the samples, but in a different way than bootstrapping. Finally, we can simply perturb the data directly by adding noise to X and/or Y . For example, we can add random Gaussian noise to the response (Breiman 1996). We find in the experimental results section that the choice of data perturbation scheme does not affect the results much.

With high dimensional data, computation can be costly. In the case of the Lasso, the computation quickly gets expensive with larger data sets (Efron et al. 2004; Mairal and Yu 2012). Using the estimation stability metric to select the regularization parameter incurs only as much

computation as using cross validation. This is because the bulk of the computation in both cases rests in computing the solution paths of the V perturbed data sets. V in this case can be small as demonstrated in Section 3. This is in contrast to related work (Bach 2008; Meinshausen and Bühlmann 2010) which requires a much larger V .

3 Experimental Results for Lasso

In this section, we evaluate *ESCV*'s performance relative to the cross validation (*CV*) across a variety of data examples. In each problem, we fit a linear model using the Lasso. We focus our attention on the comparison with *CV* as it is the most popular criterion in practice.

In all the data examples, we use the same grid-search algorithm to find a minimum of τ for *ESCV* and that for *CV*. For our algorithm, we determine the domains for τ , $[0, \tau_{\max}]$ in each of our pseudo-solutions $\hat{\beta}[l; \tau]$, where τ_{\max} is the L_1 -norm of the saturated pseudo-solution. We limit our choice of τ to the intersection of the domains. This potentially truncates some large values of τ from the original lasso path, but we view this as inconsequential as we are often after sparse solutions corresponding to relatively small values of τ .

The fact that a τ value is larger than the upper limit of the intersection implies that for at least one pseudo data set, the solution at this τ value is not unique so this τ value corresponds to instability and we would not want to consider it in our search for a stable solution. We evaluate each criterion in a equally spaced grid of 1000 values of τ in the resulting domain.

We start with simple sparse gaussian linear model simulations with our focus on the high dimensional data set up. We will vary the simulation parameters such as correlation strength within features and signal strength, as well as explore popular correlation structures of the design matrix, to cover a wide range of data scenarios in practice. We compare the solutions picked by *ESCV* and *CV* with regard to parameter estimation, prediction, and model selection performance measures such as F -measure and model size. We also include the *BIC* choice, but note that it performs poorly as expected in our high dimensional setting.

We also explore the performance of our method on two real data sets from neuroscience and bioinformatics. We use a combination of objective predictive performance and subject knowledge on plausible models to illustrate the efficacy of *ESCV* over *CV*. In all cases, note that we are comparing different choices of τ on the same solution path (from the original data). Furthermore,

we use the same data splits to make comparable results of *CV* and *ESCV*. We find that the number of cross validation folds V do not affect the relative performance between *CV* and *ESCV*. For the simulations and data below, we use commonly used $V = 8$ to obtain the pseudo solutions.

3.1 Gaussian Simulation

Let $X_i \in \mathbb{R}^p$ for $i = 1, \dots, n$ be independent identically distributed Gaussian variables with mean 0 and covariance Σ . We have the usual linear model $Y_i = X_i' \beta + \epsilon_i$, where $\beta \in \mathbb{R}^p$ is the unknown parameter, and $\epsilon_i \in \mathbb{R}$ is independent Gaussian noise with standard deviation σ . β_j are drawn from $U[\frac{1}{3}, 1]$ for $j = 1, \dots, 10$ and 0 otherwise. The separation from zero is for model selection to make sense. This is a common assumption in theoretical work.

The reported estimation and prediction errors are defined as

$$\|\hat{\beta} - \beta\|_2 \text{ and } \sqrt{E_X(\|X\hat{\beta} - X\beta\|_2^2)} = \sqrt{(\hat{\beta} - \beta)' \Sigma (\hat{\beta} - \beta)}$$

respectively. For model selection, we use the F -measure which balances false positive and false negative rates of identifying non-zero coefficients of β . The higher the F -measure the better. Each simulation is repeated 1000 times and the performance measures are aggregated across them.

3.1.1 A Base Case

Within the Gaussian linear model setup, there are many problem scenarios that favor one method over others. In particular, the following problem settings are known to affect the performance of the Lasso: correlation strength between features, strength of signal (size of coefficients) relative to the noise levels, dimension of the problem (p), and the correlation structure of the features. This is of course not an exhaustive list but is sufficient to cover a wide range of problems. As the strength of the correlation and signal are key to the behavior of the Lasso solution, we will include a full complement of these problem settings to illustrate when and why *ESCV* works well.

We start with a base case scenario. Here, Σ has entries 1 down the diagonal and constant ρ on the off-diagonal. We vary $\rho = 0, 0.2, 0.5, 0.9$ and $\sigma = 0.5, 1, 2$. We set $n = 100$ and $p = 150$ to emulate the high dimensional data setting. Note that this implies that the columns of X are empirically correlated even when the features they represent are independent.

As expected, *CV* does well in terms of prediction error (see Table 2). However, observe that this does not necessarily translate to success in terms of other performance measures. With estimation

error, we find that once we leave the orthogonal case $\rho = 0$ where estimation and prediction error are equivalent, *ESCV* has lower estimation error than *CV* despite having comparable prediction error.

For model selection, we use the *F*-measure, the harmonic mean of the precision and recall rates, which are inversely proportional to false positive rate and false negative rate respectively. A high *F*-measure is achieved when both false positive and false negative rates are low. Recall that we are selecting solutions from the same solution path. The Lasso solution path corresponds roughly to a nested family of models in terms of features picked since features seldom gets dropped as the we relax the penalty term. Hence, having a low false negative rate (high recall) typically comes at the cost of a high false positive rate (low precision). The *F*-measure balances these two objectives.

By this measure, *ESCV* often outscores *CV* by a considerable margin. *CV* picks more true variables, but in the process picks up a disproportionately large number of noise variables. This is in line with theory that *CV* often overselects (Wasserman and Roeder 2009). *ESCV* cuts down the false positive rate, but not too much at the expense of the false negative rate.

The results are summarized in Table 2 and the standard errors (SE) are given in Table 3. Note that the performance measures are highly correlated since for each simulation run, the selections by *ESCV*, *CV* and *BIC* are from the same solution path. Hence, the SEs for paired differences in performance measures are actually lower than the SEs for each of the values as reported in Table 3.

3.1.2 Effect Of Ambient Dimension

We repeat the simulations but this time for $p = 50$ and $p = 500$ to investigate the effect of the ambient dimension. Note that only the number of non-relevant features is changing; the number of non-zero coefficients remain at 10, the sample size n remains at 100. The comparison of *ESCV* and *CV* from the base case extends here: *CV* does well in prediction error, especially in the independent predictors case, but loses out to *ESCV* in the other scenarios with dependence more relevant to practice and in terms of parameter estimation and model selection metrics that are important for scientific applications. The results are summarized in Table 5 and 6.

For the low dimensional case $p = 50$, we see that *BIC* performance is much improved compared

		Estimation error			Prediction error			Model Selection F -measure			Model Size		
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>
0	0.5	0.470	0.399	0.632	0.470	0.399	0.632	0.673	0.402	0.225	19.7	39.7	78.7
0	1	0.885	0.787	1.27	0.885	0.787	1.27	0.594	0.413	0.225	21.7	37.8	77.4
0	2	1.53	1.47	2.29	1.53	1.47	2.29	0.434	0.403	0.207	22.8	29.8	58.3
0.2	0.5	0.386	0.391	0.687	0.373	0.360	0.623	0.535	0.440	0.223	27.4	35.4	79.6
0.2	1	0.757	0.773	1.36	0.719	0.711	1.23	0.522	0.445	0.223	27.8	34.5	78.3
0.2	2	1.37	1.43	2.56	1.31	1.33	2.34	0.475	0.414	0.206	25.6	31.6	73.4
0.5	0.5	0.465	0.479	0.871	0.341	0.348	0.624	0.493	0.456	0.222	30.6	33.8	80.0
0.5	1	0.927	0.956	1.72	0.682	0.695	1.23	0.485	0.445	0.214	29.9	33.6	80.2
0.5	2	1.62	1.67	3.18	1.21	1.23	2.30	0.427	0.396	0.19	26.2	29.5	74.6
0.9	0.5	1.02	1.04	1.92	0.330	0.339	0.615	0.466	0.444	0.211	31.1	33.2	80.4
0.9	1	1.75	1.80	3.55	0.572	0.587	1.14	0.396	0.377	0.183	26.9	28.9	75.1
0.9	2	2.46	2.58	5.50	0.858	0.861	1.81	0.274	0.262	0.158	19.0	21.3	58.3

Table 2: Performance of *ESCV*, *CV* and *BIC* in picking the regularization parameter for the Lasso for our base case design: constant correlation ρ , $n = 100$, $p = 150$. We see that *ESCV* performs best in parameter estimation (when different from prediction) and model selection, while doing comparably to *CV* in prediction.

		Estimation error SE			Prediction error SE			Model Selection <i>F</i> -measure SE		
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>
0	0.5	0.003	0.002	0.004	0.003	0.002	0.004	0.004	0.003	0.002
0	1	0.007	0.005	0.008	0.007	0.005	0.008	0.004	0.003	0.002
0	2	0.009	0.008	0.02	0.009	0.008	0.02	0.004	0.003	0.002
0.2	0.5	0.002	0.003	0.005	0.002	0.002	0.004	0.002	0.003	0.002
0.2	1	0.004	0.005	0.01	0.004	0.004	0.009	0.002	0.003	0.002
0.2	2	0.007	0.008	0.02	0.007	0.008	0.02	0.003	0.003	0.002
0.5	0.5	0.003	0.003	0.006	0.002	0.002	0.004	0.002	0.003	0.002
0.5	1	0.006	0.006	0.01	0.004	0.004	0.009	0.002	0.003	0.002
0.5	2	0.008	0.009	0.03	0.007	0.007	0.02	0.003	0.003	0.002
0.9	0.5	0.006	0.006	0.01	0.002	0.002	0.004	0.002	0.003	0.002
0.9	1	0.008	0.01	0.03	0.003	0.003	0.01	0.003	0.003	0.002
0.9	2	0.01	0.02	0.08	0.01	0.005	0.03	0.003	0.003	0.002

Table 3: Standard errors (SE) for performance numbers in Table 2.

to the $p = 150$ case. This is to be expected since BIC was developed to tackle problems in the classical regime of $n > p$. It is interesting to note that the performance of $ESCV$ and BIC are very comparable in this case. Indeed, in this low dimensional case, they are very close to (empirically) optimal with respect to the F -measure. Of course, once we go to $p = 500$, BIC 's performance falls off the cliff. $ESCV$ continues to beat out CV by a wide margin.

3.1.3 Other Correlation Structures

The constant correlation structure can be seen as a simple one latent variable model. Here we introduce other correlation structures corresponding to more complex models and run the same simulations ($n = 100, p = 150$, and varying σ and ρ). First, block correlation: all p features are randomly grouped into 10 blocks, and within each block, the features have correlation ρ while features from separate blocks are independent. Here, we let $\rho = 0.3, 0.5, 0.9$. Second, Toeplitz design: $\Sigma_{ij} = \rho^{|i-j|}$, with $\rho = 0.5, 0.9, 0.99$. In this case, the ten true variables indices are randomly distributed among the p variables so that they are not all strongly correlated with each other. The results for the two designs are summarized in Tables 7 and 8 respectively.

Despite the different correlation structures, the qualitative results from the prior section holds again in both variations. For prediction error, CV almost always outperforms $ESCV$, but $ESCV$'s predictive performance can be quite close to CV 's when $\rho \neq 0$. For estimation error, $ESCV$ gains on and eventually outperforms CV with increasing correlation levels. And for model selection, $ESCV$ almost always has a higher F -measure than CV . Digging deeper, Table 9 shows the breakdown of the F -measure into the true positive and false positive rates. We can see that $ESCV$ has much lower false positive rates while sacrificing relatively little on the true positive rates. Also unsurprising, BIC continues to do poorly in these $p > n$ regimes.

3.2 fMRI Data

This data is from the Gallant Neuroscience Lab at University of California, Berkeley. In this experiment, a subject is shown a series of randomly selected natural images and the fMRI response from his primary visual cortex is recorded. The fMRI response is recorded at the voxel level, where each voxel corresponds to a tiny volume of the visual cortex. The task is to model each voxel's response to the $n = 1500$ images. The image features are approximately 10000 transformed Gabor

wavelet coefficients. We evaluate the prediction performance by looking at correlation scores against an untouched validation set of 120 images with 10-13 replicates. There are 1250 voxels in all. We ranked them according to their predictive performance under a different procedure from a previous study (Kay et al. 2008). Not all of them are informative, so we only look at the top 500.

We find that while the prediction performance are nearly identical for *ESCV* and *CV*, *ESCV* selects much fewer features. The results are in Table 4. For the sake of brevity, they are averaged across groups of 100 voxels. For example, for the top 100 voxels, on average, the correlation scores are similar, but *ESCV* selects 30 features compared to *CV*'s 70 features - a close to 60% reduction. That is, *ESCV* selects a much simpler and also more reliable model that predicts just as well as *CV*. Figure 4 shows how close the correlation scores are.

Voxels	Correlation Score		Model Size	
	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>
1-100	0.730	0.735	30.1	70.2
101-200	0.653	0.655	27.0	61.8
201-300	0.567	0.566	22.6	49.6
301-400	0.455	0.459	16.7	40.3
401-500	0.347	0.347	16.5	33.6

Table 4: Performance on fMRI data set. The numbers are averaged across the respective hundred voxels. *ESCV* cuts down the model size by more than half compared to *CV*, while largely preserving prediction accuracy.

We note again that *ESCV* picks fewer features than *CV* by design (Section 2.5). That being said, the reduction is huge here: *ESCV* picks less than *half* the number of features as *CV* across the different voxels. Furthermore, this was with little or no loss in predictive performance. To understand the results better, we look at the individual voxels and examine the features selected. In almost all the cases, *ESCV* selects a subset of the features selected by *CV*. This is because they both select from the same Lasso solution path and features are rarely dropped after being added to the solution as we relax the regularization.

Now, each feature corresponds to a Gabor wavelet characterized by its location, frequency, and orientation. We plot the features selected by both *CV* and *ESCV* as well as the extra features




Figure 4: Scatterplot of predictive correlation scores of *ESCV* and *CV* for the top 500 voxels in the fMRI data set. We see that for almost all 500 voxels, the predictive performances are similar for *ESCV* and *CV*.

selected by *CV*. The points in the plot represent the location and size of the Gabor wavelet selected. Figure 5 shows four randomly selected voxels.

We can see quite clearly that the features selected by *ESCV* are clustered in one area whereas the features selected by *CV* but not *ESCV* are scattered across the image. Biologically, we expect each voxel to respond only to a particular area of the visual receptive field. This confirms that the extra features selected by *CV* are most likely not meaningful. Note that the location information of the Gabor wavelets were *not* used in fitting the model.

3.3 Cytokine Data

This data is from experiments performed by the Alliance for Cellular Signaling (AfCS), archived and made available at the Signaling Gateway, a comprehensive and free resource supported by the University of California, San Diego (UCSD). Pradervand et al. (2006) from the Bioinformatics

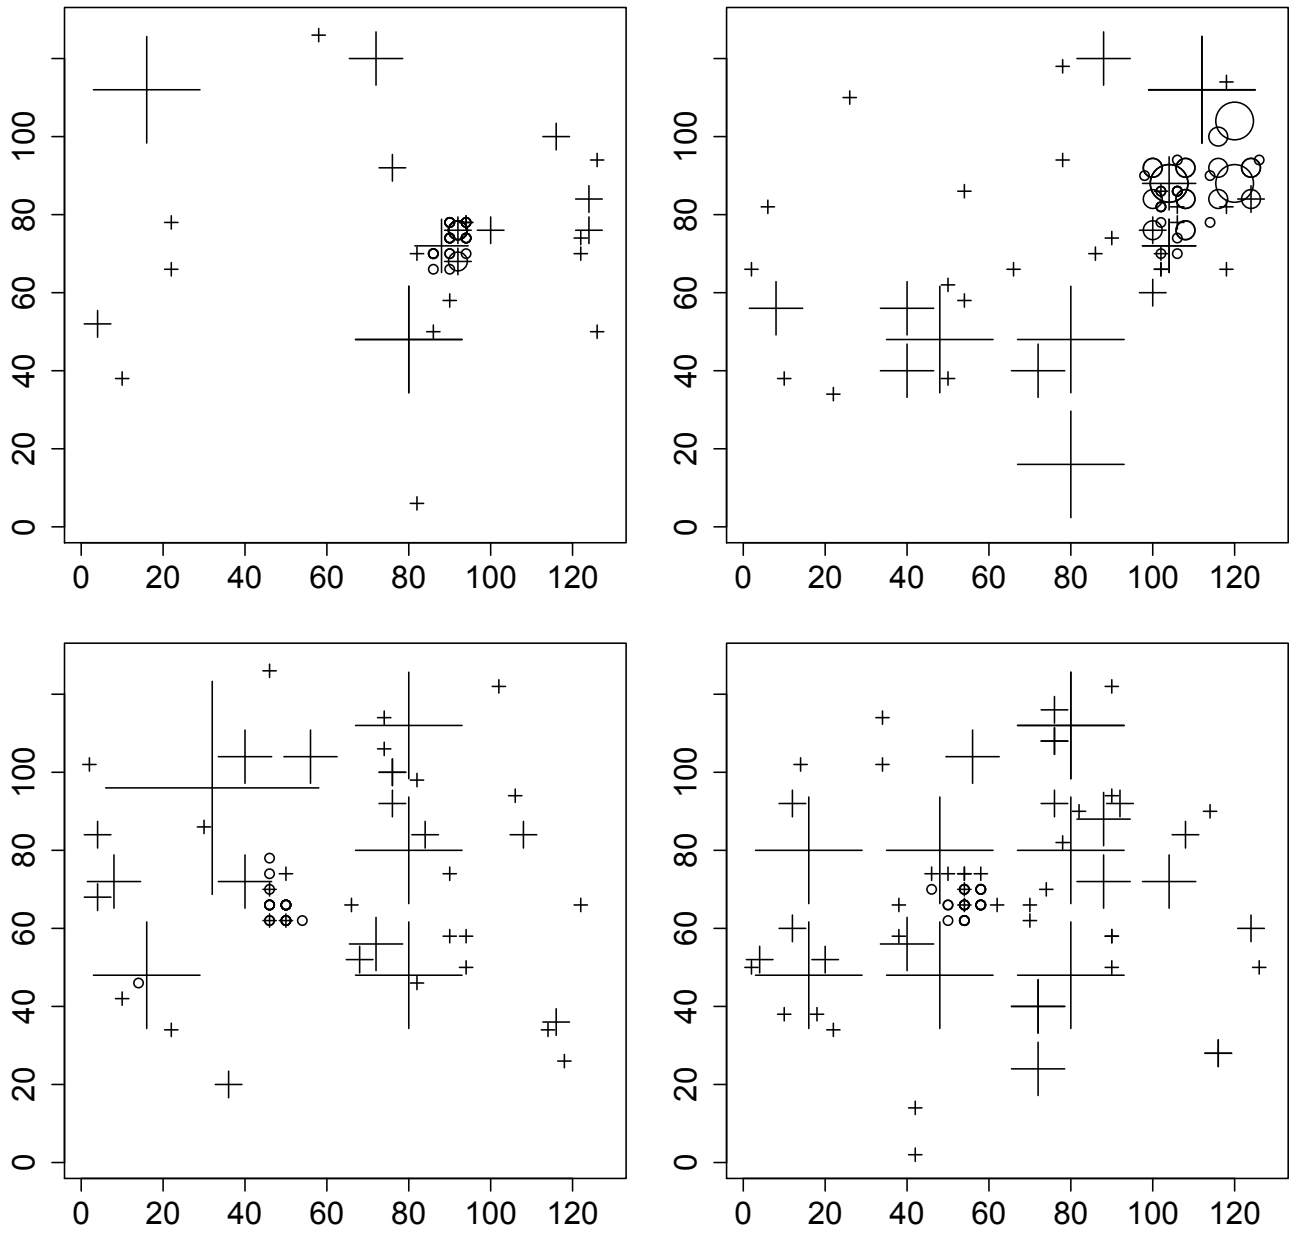


Figure 5: Feature selection by *ESCV* and *CV* on four randomly selected voxels. The “o”s represent features selected by both methods, while the “+”s represent features selected only by *CV*. The axes represent the pixel location of the images. The position and size of the points represents the wavelet location and wavelet scale respectively. Note that most of the extra features *CV* select are scattered and less biologically plausible.

and Data Coordination Laboratory at UCSD processed and analyzed this data in an attempt to identify signal pathways responsible for regulating cytokine release. There are 7 cytokines, 22 signal pathway predictors. The signal pathways cannot be directly manipulated. Instead, ligands are stimulated to elicit responses from the signal pathway predictors and cytokines. For each cytokine, we have about 100 samples, each corresponding to average measured responses of the cytokine and signal pathways when a specific ligand pair is stimulated.

In the original study (Pradervand et al. 2006), principal component regression (*PCR*) is used to fit the data to a linear model and select the significant signal pathways. The selection is done by thresholding the estimated coefficients via a pseudo-bootstrap method. They do this for each of the seven cytokines. That is, they solve seven linear regression problems, each with $n \approx 100$ and $p = 22$, and apply thresholding to select the relevant signal pathways. These *PCR* results are then merged with other data and analysis to derive a final minimal model (*MM*).

We run Lasso with *ESCV* and *CV* on the seven linear regression problems and compare our results with the results from *PCR* and *MM*. Fig 6 shows the feature selection results for the four methods. We regard *MM* as the benchmark for feature selection performance because it encompasses extra data and is not directly restricted by the linear model.

We can see from Fig 6 that Lasso with *CV* does poorly. It selects the most features for every cytokine, often by a large margin. Lasso with *ESCV* on the other hand, selects the same or slightly larger number of features than *MM*. Moreover, with the exception of cytokine *TNFa*, *ESCV* always includes the features *PCR* selected which survived to the minimal model. In the case of *TNFa*, *PCR* barely selects (close to threshold) the one feature that *ESCV* missed. *ESCV* in general selects only about half the number of features *PCR* selects. There are far fewer false positives with respect to *MM*. At the same time, it rarely misses out any of the important features that *PCR* picked up.

4 Conclusion

Regularization methods are employed to deal with problems in the increasingly common high dimensional setting. However, the difficult problem of selecting the associated regularization parameter for interpretation or parameter estimation, is not well studied. Our method *ESCV* is based on estimation stability but also takes into account model fit via *CV*. With a similar paral-

		cAMP	AKT	ERK1	ERK2	Ezr/Rdx	GSK3A	GSK3B	JNK Ig	JNK sh	Msn	p38	p40Phox	NFkB p65	PKCd	PKCmu2	RSK	Rps6	SMAD2	STAT1a	STAT1b	STAT3	STAT5	
G-CSF	MM																							
	PCR																							
	ESCV																							
	CV																							
IL-1a	MM																							
	PCR																							
	ESCV																							
	CV																							
IL-6	MM																							
	PCR																							
	ESCV																							
	CV																							
IL-10	MM																							
	PCR																							
	ESCV																							
	CV																							
MIP-1a	MM																							
	PCR																							
	ESCV																							
	CV																							
RANTES	MM																							
	PCR																							
	ESCV																							
	CV																							
TNFa	MM																							
	PCR																							
	ESCV																							
	CV																							

Figure 6: Feature selection results on cytokine data. The columns represent signal pathways predictors and each block of four rows correspond to a cytokine. The four rows within each block represent the selections of the four methods: the final minimal model (*MM*) and principal component regression (*PCR*) from the original study, and Lasso with *ESCV* and *CV*. The white squares corresponds to selected predictors. With only one exception, *ESCV* always selects the pathways that *MM* (which we regard as ground truth) does, while having much smaller models than *CV*.

lelizable computational cost as CV , we have demonstrated that $ESCV$ is an effective alternative to the popular CV for choosing the regularization parameter for the Lasso. For the practical situation of dependent predictors, $ESCV$ has an overall performance better than CV for parameter estimation and significantly better for model selection. Their prediction performances are comparable, unless the predictors are independent. In particular, we found much sparser models of less than half the size in both the real data sets from neuroscience and cell biology without sacrificing prediction accuracy, and these models are more plausible biologically based on subject knowledge. We believe this result is not restricted to the Lasso but holds for other regularization methods as well.

We also believe that this method can also be readily extended to the classification problem through the generalized linear model, and leave this to future work.

Acknowledgements

Lim would like to thank Jing Lei and Garvesh Raskutti for helpful discussions, and Yu would like to thank Arnaldo Frigessi and Ingrid Glad for helpful comments. This research is supported in part by US NSF grants DMS-1107000, SES-0835531 (CDI), ARO grant W911NF-11-1-0114, and the Center for Science of Information (CSoI), an US NSF Science and Technology Center, under grant agreement CCF-0939370.

References

- Akaike, H. (1974), “A new look at the statistical model identification,” *IEEE transactions on automatic control*.
- Allen, D. M. (1974), “The relationship between variable selection and data argumentation and a method of prediction,” *Technometrics*, 16, 125–127.
- Bach, F. (2008), “Bolasso: Model Consistent Lasso Estimation through the Bootstrap,” *Proceedings of the 25th international conference on Machine learning*.
- Bickel, P., Ritov, Y., and Tsybakov, A. (2009), “Simultaneous Analysis of Lasso and Dantzig Selector,” *Annals of Statistics*, 37, 1705–1732.

		Estimation error			Prediction error			Model Selection F -measure			Model Size		
ρ	σ	$ESCV$	CV	BIC	$ESCV$	CV	BIC	$ESCV$	CV	BIC	$ESCV$	CV	BIC
0	0.5	0.357	0.303	0.327	0.357	0.303	0.327	0.776	0.556	0.740	15.8	26.0	17.0
0	1	0.665	0.597	0.654	0.665	0.597	0.654	0.712	0.559	0.746	17.6	25.7	16.4
0	2	1.23	1.17	1.39	1.23	1.17	1.39	0.586	0.549	0.664	18.8	23.4	9.96
0.2	0.5	0.294	0.296	0.291	0.294	0.276	0.284	0.702	0.603	0.706	18.5	23.2	18.3
0.2	1	0.581	0.589	0.584	0.566	0.548	0.564	0.687	0.606	0.701	19.0	22.9	18.4
0.2	2	1.10	1.13	1.12	1.07	1.06	1.12	0.647	0.593	0.682	18.6	21.8	16.4
0.5	0.5	0.352	0.364	0.355	0.266	0.268	0.268	0.660	0.614	0.681	20.3	22.6	19.4
0.5	1	0.704	0.727	0.712	0.531	0.534	0.538	0.658	0.613	0.679	20.1	22.4	19.1
0.5	2	1.31	1.36	1.33	1.00	1.01	1.04	0.613	0.580	0.628	18.7	20.7	17.2
0.9	0.5	0.784	0.799	0.784	0.275	0.262	0.262	0.632	0.614	0.659	21.0	22.1	19.8
0.9	1	1.41	1.47	1.42	0.473	0.485	0.485	0.597	0.572	0.612	18.9	20.4	17.7
0.9	2	2.21	2.31	2.22	0.793	0.775	0.803	0.461	0.449	0.455	14.4	15.9	13.0

Table 5: Performance of $ESCV$, CV and BIC in picking the regularization parameter for the Lasso for $p = 50$ with the base case Gaussian simulation $n = 100$, constant correlation ρ .

		Estimation error			Prediction error			Model Selection F -measure			Model Size		
ρ	σ	$ESCV$	CV	BIC	$ESCV$	CV	BIC	$ESCV$	CV	BIC	$ESCV$	CV	BIC
0	0.5	0.626	0.531	0.593	0.626	0.531	0.593	0.546	0.309	0.204	26.0	54.7	88.0
0	1	1.17	1.04	1.17	1.17	1.04	1.17	0.458	0.317	0.197	27.3	49.7	87.7
0	2	1.80	1.76	2.08	1.80	1.76	2.08	0.300	0.291	0.159	27.1	33.1	78.9
0.2	0.5	0.511	0.513	0.624	0.479	0.469	0.563	0.408	0.342	0.203	39.0	48.5	88.5
0.2	1	1.01	1.02	1.24	0.945	0.931	1.12	0.392	0.332	0.195	38.8	47.9	88.6
0.2	2	1.67	1.71	2.19	1.58	1.58	1.99	0.326	0.287	0.154	33.6	40.9	87.8
0.5	0.5	0.626	0.636	0.783	0.454	0.458	0.559	0.376	0.347	0.203	43.1	47.5	88.4
0.5	1	1.21	1.23	1.52	0.888	0.888	1.09	0.356	0.325	0.184	41.1	46.3	88.9
0.5	2	1.89	1.93	2.62	1.41	1.41	1.87	0.281	0.256	0.133	33.1	38.5	88.0
0.9	0.5	1.32	1.35	1.67	0.430	0.433	0.533	0.335	0.315	0.178	43.0	46.3	89.0
0.9	1	2.02	2.06	2.87	0.667	0.670	0.917	0.251	0.235	0.122	34.4	38.0	88.4
0.9	2	2.58	2.72	4.85	0.891	0.903	1.56	0.135	0.125	0.0689	22.6	26.9	85.8

Table 6: Performance of $ESCV$, CV and BIC in picking the regularization parameter for the Lasso for $p = 500$ with the base case Gaussian simulation $n = 100$, constant correlation ρ .

- Breiman, L. (1995), “Better Subset Regression Using the Nonnegative Garrote,” *Technometrics*, 37.
- (1996), “Heuristics of instability and stabilization in model selection,” *The Annals of Statistics*, 24, 2350–2383.
- (2001), “Statistical Modeling: The Two Cultures,” *Statistical Science*, 16, 199–231.
- Bühlmann, P. and Yu, B. (2003), “Boosting with the L2 Loss: Regression and Classification,” *Journal of the American Statistical Association*, 98, 324–339.
- Efron, B. (1979), “Bootstrap methods: another look at the jackknife,” *The Annals of Statistics*, 7, 1–26.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least Angle Regression,” *The Annals of Statistics*, 32, 407–499.
- Ellis, S. P. (1998), “Instability of Least Squares, Least Absolute Deviation and Least Median of Squares Linear Regression,” *Statistical Science*, 13, 337–350.
- Friedman, J. H. (2001), “Greedy Function Approximation: A Gradient Boosting Machine,” *The Annals of Statistics*, 29, 1189–1232.
- Hastie, T., Tibshirani, R., and Friedman, J. (2002), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer.
- Haury, A.-C., Mordelet, F., Vera-Licona, P., and Vert, J.-P. (2012), “TIGRESS: Trustful Inference of Gene REgulation using Stability Selection,” *BMC Systems Biology*, 6.
- Higham, N. J. (1996), *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia.
- Hoerl, A. (1962), “Application of ridge analysis to regression problems,” *Chemical Engineering Progress*.
- Kay, K., Naselaris, T., Prenger, R., and Gallant, J. (2008), “Identifying natural images from human brain activity,” *Nature*, 452, 352–355.

- Leng, C., Lin, Y., and Wahba, G. (2006), “A Note on the Lasso and Related Procedures in Model Selection,” *Statistica Sinica*, 16, 1273–1284.
- Liu, H., Roeder, K., and Wasserman, L. (2010), “Stability Approach to Regularization Selection (StARS) for High Dimensional Graphical Models,” *Advances in Neural Information Processing Systems*, 23.
- Mairal, J. and Yu, B. (2012), “Complexity Analysis of the Lasso Regularization Path,” *Proceedings of the 29th International Conference on Machine Learning (ICML), Edinburgh, Scotland, UK*.
- Meinshausen, N. and Bühlmann, P. (2006), “High-Dimensional Graphs and Variable Selection with the Lasso,” *The Annals of Statistics*, 34, 1436–1462.
- (2010), “Stability Selection,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Meinshausen, N., Rocha, G., and Yu, B. (2007), “Discussion: A tale of three cousins: Lasso, L2Boosting and Dantzig,” *The Annals of Statistics*, 35, 2373–2384.
- Meinshausen, N. and Yu, B. (2009), “Lasso-type recovery of sparse representations for high-dimensional data,” *Annals of Statistics*, 37, 246–270.
- Pradervand, S., Maurya, M., and Subramaniam, S. (2006), “Identification of signaling components required for the prediction of cytokine release in RAW 264.7 macrophages,” *Genome Biology*, 7.
- Salle, J. P. L. (1976), *The stability of dynamical systems*, CBMS-NSF Regional Conference Series in Applied Mathematics.
- Schwarz, G. (1978), “Estimating the dimension of a model,” *The Annals of Statistics*.
- Shah, R. D. and Samworth, R. J. (2013), “Variable selection with error control: another look at stability selection,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75, 55–80.
- Shao, J. (1996), “Bootstrap model selection,” *Journal of the American Statistical Association*, 91, 353–360.

- Stone, M. (1974), “Cross-validation choice and assessment of statistical prediction,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 39, 44–47.
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58, 267–288.
- Tikhonov, A. N. (1943), “[On the stability of inverse problems],” *Doklady Akademii Nauk SSSR*, 39(5), 195–198.
- Tropp, J. (2006), “Just relax: Convex programming methods for identifying sparse signals in noise,” *IEEE Transactions on Information Theory*, 52.
- Wainwright, M. (2009), “Sharp Thresholds for High-Dimensional and Noisy Sparsity Recovery Using l_1 -Constrained Quadratic Programming (Lasso),” *IEEE Transactions on Information Theory*, 55, 2183–2202.
- Wasserman, L. and Roeder, K. (2009), “High-Dimensional Variable Selection,” *The Annals of Statistics*, 37, 2178–2201.
- Yu, B. (2013 (to appear)), “Stability,” *Bernoulli*.
- Zhang, P. (1993), “Model selection via multifold cross-validation,” *The Annals of Statistics*, 21, 299–313.
- Zhang, T. (2011), “Adaptive Forward-Backward Greedy Algorithm for Learning Sparse Representations,” *IEEE Trans. Info. Th.*, 57, 4689–4708.
- Zhang, T. and Yu, B. (2012), “Boosting with early stopping: Convergence and Consistency,” *The Annals of Statistics*, 33, 1538–1579.
- Zhao, P. and Yu, B. (2006), “On Model Selection Consistency of Lasso,” *Journal of Machine Learning Research*, 7, 2541–2563.
- (2007), “Stagewise Lasso,” *The Journal of Machine Learning Research*, 8, 2701–2726.

		Estimation error			Prediction error			Model Selection F -measure			Model Size		
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>
0.3	0.5	0.437	0.406	0.697	0.428	0.378	0.626	0.636	0.422	0.229	21.4	37.4	77.4
0.3	1	0.845	0.798	1.36	0.824	0.747	1.23	0.594	0.438	0.235	22.3	35.0	73.3
0.3	2	1.48	1.47	2.50	1.45	1.39	2.30	0.465	0.407	0.211	22.1	30.6	62.1
0.5	0.5	0.457	0.445	0.811	0.395	0.363	0.625	0.586	0.441	0.229	24.1	35.3	77.4
0.5	1	0.908	0.899	1.62	0.783	0.735	1.25	0.544	0.433	0.222	25.2	35.1	77.2
0.5	2	1.56	1.57	2.82	1.39	1.32	2.27	0.454	0.403	0.204	22.2	29.6	62.7
0.9	0.5	0.876	0.904	1.77	0.346	0.345	0.618	0.501	0.452	0.218	28.8	33.1	77.8
0.9	1	1.54	1.59	3.30	0.648	0.629	1.16	0.444	0.403	0.192	25.9	30.5	74.2
0.9	2	2.17	2.33	5.18	1.09	1.03	1.98	0.310	0.291	0.164	19.4	23.9	57.7

Table 7: Performance of *ESCV*, *CV* and *BIC* in picking the regularization parameter for the Lasso for the block correlation design. $n = 100$, $p = 150$.

		Estimation error			Prediction error			Model Selection F -measure			Model Size		
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>	<i>ESCV</i>	<i>CV</i>	<i>BIC</i>
0.5	0.5	0.460	0.411	0.726	0.444	0.383	0.648	0.649	0.424	0.230	20.8	37.2	76.8
0.5	1	0.877	0.813	1.43	0.845	0.761	1.28	0.578	0.424	0.228	22.9	36.5	75.4
0.5	2	1.53	1.48	2.49	1.48	1.40	2.30	0.444	0.409	0.211	22.7	29.5	55.2
0.9	0.5	0.696	0.706	1.45	0.353	0.345	0.625	0.532	0.467	0.238	27.2	32.4	71.8
0.9	1	1.28	1.31	2.79	0.676	0.663	1.22	0.488	0.433	0.209	26.2	31.2	71.4
0.9	2	1.93	1.99	4.43	1.18	1.14	2.11	0.372	0.346	0.184	21.6	25.7	56.3
0.99	0.5	1.72	1.76	4.14	0.290	0.296	0.572	0.400	0.388	0.185	27.0	28.1	68.9
0.99	1	2.31	2.39	6.51	0.476	0.481	0.954	0.290	0.282	0.162	21.6	22.6	55.9
0.99	2	2.81	2.96	8.72	0.860	0.758	1.45	0.188	0.185	0.142	15.6	16.8	38.9

Table 8: Performance of *ESCV*, *CV* and *BIC* in picking the regularization parameter for the Lasso for the Toeplitz correlation design. $n = 100$, $p = 150$.

Constant correlation design													
		$p = 150$				$p = 50$				$p = 500$			
		True Positive		False Positive		True Positive		False Positive		True Positive		False Positive	
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>
0	0.5	9.98	10.0	9.69	29.7	10.0	10.0	5.77	16.0	9.84	9.99	16.2	44.7
0	1	9.41	9.88	12.3	27.9	9.83	9.97	7.80	15.7	8.56	9.48	18.8	40.2
0	2	7.12	8.02	15.7	21.8	8.43	9.15	10.3	14.2	5.57	6.28	21.6	26.9
0.5	0.5	10.0	10.0	20.6	23.8	10.0	10.0	10.3	12.6	9.98	9.98	33.1	37.5
0.5	1	9.68	9.70	20.2	23.9	9.91	9.93	10.2	12.4	9.08	9.16	32.0	37.2
0.5	2	7.72	7.82	18.5	21.7	8.79	8.90	9.87	11.8	6.06	6.20	27.0	32.3
0.9	0.5	9.59	9.59	21.5	23.7	9.81	9.85	11.2	12.2	8.86	8.87	34.1	37.4
0.9	1	7.29	7.33	19.6	21.6	8.63	8.71	10.3	11.7	5.56	5.63	28.8	32.4
0.9	2	3.96	4.10	15.0	17.2	5.62	5.81	8.76	10.1	2.21	2.31	20.4	24.6

		Block design, $p = 150$			
		True Positive		False Positive	
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>
0.3	0.5	10.0	10.0	11.4	27.4
0.3	1	9.60	9.87	12.7	25.2
0.3	2	7.46	8.26	14.6	22.4
0.5	0.5	9.99	10.0	14.1	25.3
0.5	1	9.57	9.78	15.6	25.4
0.5	2	7.31	7.99	14.9	21.6
0.9	0.5	9.73	9.75	19.1	23.4
0.9	1	7.96	8.16	17.9	22.3
0.9	2	4.56	4.94	14.8	19.0

		Toeplitz design, $p = 150$			
		True Positive		False Positive	
ρ	σ	<i>ESCV</i>	<i>CV</i>	<i>ESCV</i>	<i>CV</i>
0.5	0.5	9.99	10.0	10.8	27.2
0.5	1	9.51	9.87	13.4	26.7
0.5	2	7.25	8.06	15.4	21.4
0.9	0.5	9.90	9.92	17.3	22.5
0.9	1	8.84	8.93	17.4	22.3
0.9	2	5.87	6.18	15.7	19.5
0.99	0.5	7.41	7.41	19.6	20.7
0.99	1	4.59	4.60	17.0	18.0
0.99	2	2.41	2.48	13.2	14.3

Table 9: Breakdown of the F -measure: the true positive and false positive rates of *ESCV* and *CV* for all the simulation scenarios. In all the cases above, there are 10 true variables and $p - 10$ noise variables.