
Trace Lasso: a trace norm regularization for correlated designs

Édouard Grave
INRIA, Sierra Project-team
École Normale Supérieure, Paris
edouard.grave@inria.fr

Guillaume Obozinski
INRIA, Sierra Project-team
École Normale Supérieure, Paris
guillaume.obozinski@inria.fr

Francis Bach
INRIA, Sierra Project-team
École Normale Supérieure, Paris
francis.bach@inria.fr

Abstract

Using the ℓ_1 -norm to regularize the estimation of the parameter vector of a linear model leads to an unstable estimator when covariates are highly correlated. In this paper, we introduce a new penalty function which takes into account the correlation of the design matrix to stabilize the estimation. This norm, called the trace Lasso, uses the trace norm of the selected covariates, which is a convex surrogate of their rank, as the criterion of model complexity. We analyze the properties of our norm, describe an optimization algorithm based on reweighted least-squares, and illustrate the behavior of this norm on synthetic data, showing that it is more adapted to strong correlations than competing methods such as the elastic net.

1 Introduction

The concept of parsimony is central in many scientific domains. In the context of statistics, signal processing or machine learning, it takes the form of variable or feature selection problems, and is commonly used in two situations: first, to make the model or the prediction more interpretable or cheaper to use, i.e., even if the underlying problem does not admit sparse solutions, one looks for the best sparse approximation. Second, sparsity can also be used given prior knowledge that the model should be sparse. Many methods have been designed to learn sparse models, namely methods based on greedy algorithms [1, 2], Bayesian inference [3] or convex optimization [4, 5].

In this paper, we focus on the regularization by sparsity-inducing norms. The simplest example of such norms is the ℓ_1 -norm, leading to the Lasso, when used within a least-squares framework. In recent years, a large body of work has shown that the Lasso was performing optimally in high-dimensional low-correlation settings, both in terms of prediction [6], estimation of parameters or estimation of supports [7, 8]. However, most data exhibit strong correlations, with various correlation structures, such as clusters (i.e., close to block-diagonal covariance matrices) or sparse graphs, such as for example problems involving sequences (in which case, the covariance matrix is close to a Toeplitz matrix [9]). In these situations, the Lasso is known to have stability problems: although its predictive performance is not disastrous, the selected predictor may vary a lot (typically, given two correlated variables, the Lasso will only select one of the two, at random).

Several remedies have been proposed to this instability. First, the elastic net [10] adds a strongly convex penalty term (the squared ℓ_2 -norm) that will stabilize selection (typically, given two correlated variables, the elastic net will select the two variables). However, it is blind to the exact

correlation structure, and while strong convexity is required for some variables, it is not for other variables. Another solution is to consider the group Lasso, which will divide the predictors into groups and penalize the sum of the ℓ_2 -norm of these groups [11]. This is known to accommodate strong correlations within groups [12]; however it requires to know the groups in advance, which is not always possible. A third line of research has focused on sampling-based techniques [13, 14, 15].

An ideal regularizer should thus take into account the design (like the group Lasso, with oracle groups), but without requiring human intervention (like the elastic net); it should thus add strong convexity only where needed, and not modifying variables where things behave correctly. In this paper, we propose a new norm towards this end.

More precisely we make the following contributions:

- We propose in Section 2 a new norm based on the trace norm (a.k.a. nuclear norm) that interpolates between the ℓ_1 -norm and the ℓ_2 -norm depending on correlations.
- We show that there is a unique minimum when penalizing with this norm in Section 2.2.
- We provide optimization algorithms based on reweighted least-squares in Section 3.
- We study the second-order expansion around independence and relate it to existing work on including correlations in Section 4.
- We perform synthetic experiments in Section 5, where we show that the trace Lasso outperforms existing norms in strong-correlation regimes.

Notations. Let $\mathbf{M} \in \mathbb{R}^{n \times p}$. We use superscripts for the columns of \mathbf{M} , i.e., $\mathbf{M}^{(i)}$ denotes the i -th column, and subscripts for the rows, i.e., \mathbf{M}_i denotes the i -th row. For $\mathbf{M} \in \mathbb{R}^{p \times p}$, $\text{diag}(\mathbf{M}) \in \mathbb{R}^p$ is the diagonal of the matrix \mathbf{M} , while for $\mathbf{u} \in \mathbb{R}^p$, $\text{Diag}(\mathbf{u}) \in \mathbb{R}^{p \times p}$ is the diagonal matrix whose diagonal elements are the u_i . Let S be a subset of $\{1, \dots, p\}$, then \mathbf{u}_S is the vector \mathbf{u} restricted to the support S , with 0 outside the support S . We denote by \mathbb{S}_p the set of symmetric matrices of size p . We will use various matrix norms, here are the notations we use: $\|\mathbf{M}\|_*$ is the trace norm, i.e., the sum of the singular values of the matrix \mathbf{M} , $\|\mathbf{M}\|_{op}$ is the operator norm, i.e., the maximum singular value of the matrix \mathbf{M} , $\|\mathbf{M}\|_F$ is the Frobenius norm, i.e., the ℓ_2 -norm of the singular values, which is also equal to $\sqrt{\text{tr}(\mathbf{M}^\top \mathbf{M})}$ and $\|\mathbf{M}\|_{2,1}$ is the sum of the ℓ_2 -norm of the columns of \mathbf{M} : $\|\mathbf{M}\|_{2,1} = \sum_{i=1}^p \|\mathbf{M}^{(i)}\|_2$.

2 Definition and properties of the trace Lasso

We consider the problem of predicting $y \in \mathbb{R}$, given a vector $\mathbf{x} \in \mathbb{R}^p$, assuming a linear model

$$y = \mathbf{w}^\top \mathbf{x} + \varepsilon,$$

where ε is an additive (typically Gaussian) noise with mean 0 and variance σ^2 . Given a training set $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times p}$ and $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$, a widely used method to estimate the parameter vector \mathbf{w} is penalized empirical risk minimization

$$\hat{\mathbf{w}} \in \underset{\mathbf{w}}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n \ell(y_i, \mathbf{w}^\top \mathbf{x}_i) + \lambda f(\mathbf{w}), \quad (1)$$

where ℓ is a loss function used to measure the error we make by predicting $\mathbf{w}^\top \mathbf{x}_i$ instead of y_i , while f is a regularization term used to penalize complex models. This second term helps avoiding overfitting, especially in the case where we have many more parameters than observation, i.e., $n \ll p$.

2.1 Related work

We will now present some classical penalty functions for linear models which are widely used in the machine learning and statistics community. The first one, known as Tikhonov regularization [16] or ridge [17], is the squared ℓ_2 -norm. When used with the square loss, estimating the parameter vector \mathbf{w} is done by solving a linear system. One of the main drawbacks of this penalty function is the fact

that it does not perform variable selection and thus does not behave well in sparse high-dimensional settings.

Hence, it is natural to penalize linear models by the number of variables used by the model. Unfortunately, this criterion, sometimes denoted by $\|\cdot\|_0$ (ℓ_0 -penalty), is not convex and solving the problem in Eq. (1) is generally NP-Hard [18]. Thus, a convex relaxation for this problem was introduced, replacing the size of the selected subset by the ℓ_1 -norm of \mathbf{w} . This estimator is known as the Lasso [4] in the statistics community and basis pursuit [5] in signal processing. Under some assumptions, the two problems are in fact equivalent (see for example [19] and references therein).

When two predictors are highly correlated, the Lasso has a very unstable behavior: it often only selects the variable that is the most correlated with the residual. On the other hand, Tikhonov regularization tends to shrink coefficients of correlated variables together, leading to a very stable behavior. In order to get the best of both worlds, stability and variable selection, Zou and Hastie introduced the elastic net [10], which is the sum of the ℓ_1 -norm and squared ℓ_2 -norm. Unfortunately, this estimator needs two regularization parameters and is not adaptive to the precise correlation structure of the data. Some authors also proposed to use pairwise correlations between predictors to interpolate more adaptively between the ℓ_1 -norm and squared ℓ_2 -norm, by introducing a method called pairwise elastic net [20] (see comparisons with our approach in Section 5).

Finally, when one has more knowledge about the data, for example clusters of variables that should be selected together, one can use the group Lasso [11]. Given a partition (S_i) of the set of variables, it is defined as the sum of the ℓ_2 -norms of the restricted vectors \mathbf{w}_{S_i} :

$$\|\mathbf{w}\|_{GL} = \sum_{i=1}^k \|\mathbf{w}_{S_i}\|_2.$$

The effect of this penalty function is to introduce sparsity at the group level: variables in a group are selected all together. One of the main drawback of this method, which is also sometimes one of its quality, is the fact that one needs to know the partition of the variables, and so one needs to have a good knowledge of the data.

2.2 The ridge, the Lasso and the trace Lasso

In this section, we show that Tikhonov regularization and the Lasso penalty can be viewed as norms of the matrix $\mathbf{X} \text{Diag}(\mathbf{w})$. We then introduce a new norm involving this matrix.

The solution of empirical risk minimization penalized by the ℓ_1 -norm or ℓ_2 -norm is not equivariant to rescaling of the predictors $\mathbf{X}^{(i)}$, so it is common to normalize the predictors. When normalizing the predictors $\mathbf{X}^{(i)}$, and penalizing by Tikhonov regularization or by the Lasso, people are implicitly using a regularization term that depends on the data or design matrix \mathbf{X} . In fact, there is an equivalence between normalizing the predictors and not normalizing them, using the two following reweighted ℓ_2 and ℓ_1 -norms instead of Tikhonov regularization and the Lasso:

$$\|\mathbf{w}\|_2^2 = \sum_{i=1}^p \|\mathbf{X}^{(i)}\|_2^2 w_i^2 \quad \text{and} \quad \|\mathbf{w}\|_1 = \sum_{i=1}^p \|\mathbf{X}^{(i)}\|_2 |w_i|. \quad (2)$$

These two norms can be expressed using the matrix $\mathbf{X} \text{Diag}(\mathbf{w})$:

$$\|\mathbf{w}\|_2 = \|\mathbf{X} \text{Diag}(\mathbf{w})\|_F \quad \text{and} \quad \|\mathbf{w}\|_1 = \|\mathbf{X} \text{Diag}(\mathbf{w})\|_{2,1},$$

and a natural question arises: are there other relevant choices of functions or matrix norms? A classical measure of the complexity of a model is the number of predictors used by this model, which is equal to the size of the support of \mathbf{w} . This penalty being non-convex, people use its convex relaxation, which is the ℓ_1 -norm, leading to the Lasso.

Here, we propose a different measure of complexity which can be shown to be more adapted in model selection settings [21]: the dimension of the subspace spanned by the selected predictors. This is equal to the rank of the selected predictors, or also to the rank of the matrix $\mathbf{X} \text{Diag}(\mathbf{w})$. Like the size of the support, this function is non-convex, and we propose to replace it by a convex surrogate, the *trace norm*, leading to the following penalty that we call ‘‘trace Lasso’’:

$$\Omega(\mathbf{w}) = \|\mathbf{X} \text{Diag}(\mathbf{w})\|_*.$$

The trace Lasso has some interesting properties: if all the predictors are orthogonal, then, it is equal to the ℓ_1 -norm. Indeed, we have the decomposition:

$$\mathbf{X} \text{Diag}(\mathbf{w}) = \sum_{i=1}^p \left(\|\mathbf{X}^{(i)}\|_2 w_i \right) \frac{\mathbf{X}^{(i)}}{\|\mathbf{X}^{(i)}\|_2} \mathbf{e}_i^\top,$$

where \mathbf{e}_i are the vectors of the canonical basis. Since the predictors are orthogonal and the \mathbf{e}_i are orthogonal too, this gives the singular value decomposition of $\mathbf{X} \text{Diag}(\mathbf{w})$ and we get

$$\|\mathbf{X} \text{Diag}(\mathbf{w})\|_* = \sum_{i=1}^p \|\mathbf{X}^{(i)}\|_2 |w_i| = \|\mathbf{X} \text{Diag}(\mathbf{w})\|_{2,1}.$$

On the other hand, if all the predictors are equal to $\mathbf{X}^{(1)}$, then

$$\mathbf{X} \text{Diag}(\mathbf{w}) = \mathbf{X}^{(1)} \mathbf{w}^\top,$$

and we get $\|\mathbf{X} \text{Diag}(\mathbf{w})\|_* = \|\mathbf{X}^{(1)}\|_2 \|\mathbf{w}\|_2 = \|\mathbf{X} \text{Diag}(\mathbf{w})\|_F$, which is equivalent to Tikhonov regularization. Thus when two predictors are strongly correlated, our norm will behave like Tikhonov regularization, while for almost uncorrelated predictors, it will behave like the Lasso.

Always having a unique minimum is an important property for a statistical estimator, as it is a first step towards stability. The trace Lasso, by adding strong convexity exactly in the direction of highly correlated covariates, always has a unique minimum, and is thus much more stable than the Lasso.

Proposition 1. *If the loss function ℓ is strongly convex with respect to its second argument, then the solution of the empirical risk minimization penalized by the trace Lasso, i.e., Eq. (1), is unique.*

The technical proof of this proposition can be found in [22], and consists in showing that in the flat directions of the loss function, the trace Lasso is strongly convex.

2.3 A new family of penalty functions

In this section, we introduce a new family of penalties, inspired by the trace Lasso, allowing us to write the ℓ_1 -norm, the ℓ_2 -norm and the newly introduced trace Lasso as special cases. In fact, we note that $\|\text{Diag}(\mathbf{w})\|_* = \|\mathbf{w}\|_1$ and $\|p^{-1/2} \mathbf{1}^\top \text{Diag}(\mathbf{w})\|_* = \|\mathbf{w}^\top\|_* = \|\mathbf{w}\|_2$. In other words, we can express the ℓ_1 and ℓ_2 -norms of \mathbf{w} using the trace norm of a given matrix times the matrix $\text{Diag}(\mathbf{w})$. A natural question to ask is: what happens when using a matrix \mathbf{P} other than the identity or the line vector $p^{-1/2} \mathbf{1}^\top$, and what are good choices of such matrices? Therefore, we introduce the following family of penalty functions:

Definition 1. *Let $\mathbf{P} \in \mathbb{R}^{k \times p}$, all of its columns having unit norm. We introduce the norm $\Omega_{\mathbf{P}}$ as*

$$\Omega_{\mathbf{P}}(\mathbf{w}) = \|\mathbf{P} \text{Diag}(\mathbf{w})\|_*.$$

Proof. The positive homogeneity and triangle inequality are direct consequences of the linearity of $\mathbf{w} \mapsto \mathbf{P} \text{Diag}(\mathbf{w})$ and the fact that $\|\cdot\|_*$ is a norm. Since all the columns of \mathbf{P} are not equal to zero, we have

$$\mathbf{P} \text{Diag}(\mathbf{w}) = 0 \Leftrightarrow \mathbf{w} = 0,$$

and so, $\Omega_{\mathbf{P}}$ separates points and thus is a norm. □

As stated before, the ℓ_1 and ℓ_2 -norms are special cases of the family of norms we just introduced. Another important penalty that can be expressed as a special case is the group Lasso, with non-overlapping groups. Given a partition (S_j) of the set $\{1, \dots, p\}$, the group Lasso is defined by

$$\|\mathbf{w}\|_{GL} = \sum_{S_j} \|\mathbf{w}_{S_j}\|_2.$$

We define the matrix \mathbf{P}^{GL} by

$$\mathbf{P}_{ij}^{GL} = \begin{cases} 1/\sqrt{|S_k|} & \text{if } i \text{ and } j \text{ are in the same group } S_k, \\ 0 & \text{otherwise.} \end{cases}$$

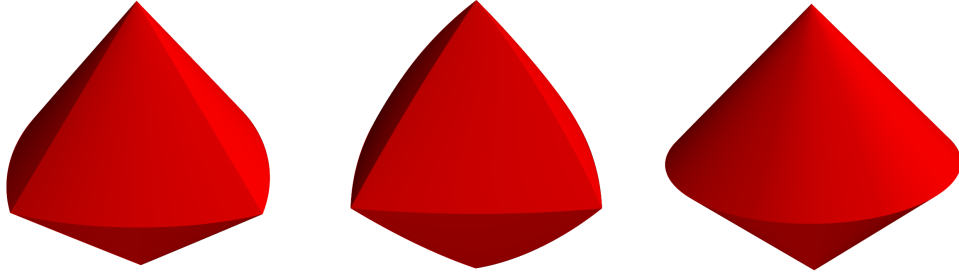


Figure 1: Unit balls for various value of $\mathbf{P}^\top \mathbf{P}$. See the text for the value of $\mathbf{P}^\top \mathbf{P}$. (Best seen in color).

Then,

$$\mathbf{P}^{GL} \text{Diag}(\mathbf{w}) = \sum_{S_j} \frac{\mathbf{1}_{S_j}}{\sqrt{|S_j|}} \mathbf{w}_{S_j}^\top. \quad (3)$$

Using the fact that (S_j) is a partition of $\{1, \dots, p\}$, the vectors $\mathbf{1}_{S_j}$ are orthogonal and so are the vectors \mathbf{w}_{S_j} . Hence, after normalizing the vectors, Eq. (3) gives a singular value decomposition of $\mathbf{P}^{GL} \text{Diag}(\mathbf{w})$ and so the group Lasso penalty can be expressed as a special case of our family of norms:

$$\|\mathbf{P}^{GL} \text{Diag}(\mathbf{w})\|_* = \sum_{S_j} \|\mathbf{w}_{S_j}\|_2 = \|\mathbf{w}\|_{GL}.$$

In the following proposition, we show that our norm only depends on the value of $\mathbf{P}^\top \mathbf{P}$. This is an important property for the trace Lasso, where $\mathbf{P} = \mathbf{X}$, since it underlies the fact that this penalty only depends on the correlation matrix $\mathbf{X}^\top \mathbf{X}$ of the covariates.

Proposition 2. *Let $\mathbf{P} \in \mathbb{R}^{k \times p}$, all of its columns having unit norm. We have*

$$\Omega_{\mathbf{P}}(\mathbf{w}) = \|(\mathbf{P}^\top \mathbf{P})^{1/2} \text{Diag}(\mathbf{w})\|_*.$$

We plot the unit ball of our norm for the following value of $\mathbf{P}^\top \mathbf{P}$ (see figure 1):

$$\begin{pmatrix} 1 & 0.9 & 0.1 \\ 0.9 & 1 & 0.1 \\ 0.1 & 0.1 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 0.7 & 0.49 \\ 0.7 & 1 & 0.7 \\ 0.49 & 0.7 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

We can lower bound and upper bound our norms by the ℓ_2 -norm and ℓ_1 -norm respectively. This shows that, as for the elastic net, our norms interpolate between the ℓ_1 -norm and the ℓ_2 -norm. But the main difference between the elastic net and our norms is the fact that our norms are *adaptive*, and require a single regularization parameter to tune. In particular for the trace Lasso, when two covariates are strongly correlated, it will be close to the ℓ_2 -norm, while when two covariates are almost uncorrelated, it will behave like the ℓ_1 -norm. This is a behavior close to the one of the pairwise elastic net [20].

Proposition 3. *Let $\mathbf{P} \in \mathbb{R}^{k \times p}$, all of its columns having unit norm. We have*

$$\|\mathbf{w}\|_2 \leq \Omega_{\mathbf{P}}(\mathbf{w}) \leq \|\mathbf{w}\|_1.$$

2.4 Dual norm

The dual norm is an important quantity for both optimization and theoretical analysis of the estimator. Unfortunately, we are not able in general to obtain a closed form expression of the dual norm for the family of norms we just introduced. However we can obtain a bound, which is exact for some special cases:

Proposition 4. *The dual norm, defined by $\Omega_{\mathbf{P}}^*(\mathbf{u}) = \max_{\Omega_{\mathbf{P}}(\mathbf{v}) \leq 1} \mathbf{u}^\top \mathbf{v}$, can be bounded by:*

$$\Omega_{\mathbf{P}}^*(\mathbf{u}) \leq \|\mathbf{P} \text{Diag}(\mathbf{u})\|_{op}.$$

Proof. Using the fact that $\text{diag}(\mathbf{P}^\top \mathbf{P}) = \mathbf{1}$, we have

$$\begin{aligned} \mathbf{u}^\top \mathbf{v} &= \text{tr}(\text{Diag}(\mathbf{u}) \mathbf{P}^\top \mathbf{P} \text{Diag}(\mathbf{v})) \\ &\leq \|\mathbf{P} \text{Diag}(\mathbf{u})\|_{op} \|\mathbf{P} \text{Diag}(\mathbf{v})\|_*, \end{aligned}$$

where the inequality comes from the fact that the operator norm $\|\cdot\|_{op}$ is the dual norm of the trace norm. The definition of the dual norm then gives the result. \square

As a corollary, we can bound the dual norm by a constant times the ℓ_∞ -norm:

$$\Omega_{\mathbf{P}}^*(\mathbf{u}) \leq \|\mathbf{P} \text{Diag}(\mathbf{u})\|_{op} \leq \|\mathbf{P}\|_{op} \|\text{Diag}(\mathbf{u})\|_{op} = \|\mathbf{P}\|_{op} \|\mathbf{u}\|_\infty.$$

Using proposition (3), we also have the inequality $\Omega_{\mathbf{P}}^*(\mathbf{u}) \geq \|\mathbf{u}\|_\infty$.

3 Optimization algorithm

In this section, we introduce an algorithm to estimate the parameter vector \mathbf{w} when the loss function is equal to the square loss: $\ell(y, \mathbf{w}^\top \mathbf{x}) = \frac{1}{2}(y - \mathbf{w}^\top \mathbf{x})^2$ and the penalty is the trace Lasso. It is straightforward to extend this algorithm to the family of norms indexed by \mathbf{P} . The problem we consider is thus

$$\min_{\mathbf{w}} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \|\mathbf{X} \text{Diag}(\mathbf{w})\|_*.$$

We could optimize this cost function by subgradient descent, but this is quite inefficient: computing the subgradient of the trace Lasso is expensive and the rate of convergence of subgradient descent is quite slow. Instead, we consider an iteratively reweighted least-squares method. First, we need to introduce a well-known variational formulation for the trace norm [23]:

Proposition 5. *Let $\mathbf{M} \in \mathbb{R}^{n \times p}$. The trace norm of \mathbf{M} is equal to:*

$$\|\mathbf{M}\|_* = \frac{1}{2} \inf_{\mathbf{S} \succeq 0} \text{tr}(\mathbf{M}^\top \mathbf{S}^{-1} \mathbf{M}) + \text{tr}(\mathbf{S}),$$

and the infimum is attained for $\mathbf{S} = (\mathbf{M}\mathbf{M}^\top)^{1/2}$.

Using this proposition, we can reformulate the previous optimization problem as

$$\min_{\mathbf{w}} \inf_{\mathbf{S} \succeq 0} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{\lambda}{2} \mathbf{w}^\top \text{Diag}(\text{diag}(\mathbf{X}^\top \mathbf{S}^{-1} \mathbf{X})) \mathbf{w} + \frac{\lambda}{2} \text{tr}(\mathbf{S}).$$

This problem is jointly convex in (\mathbf{w}, \mathbf{S}) [24]. In order to optimize this objective function by alternating the minimization over \mathbf{w} and \mathbf{S} , we need to add a term $\frac{\lambda \mu_i}{2} \text{tr}(\mathbf{S}^{-1})$. Otherwise, the infimum over \mathbf{S} could be attained at a non invertible \mathbf{S} , leading to a non convergent algorithm. The infimum over \mathbf{S} is then attained for $\mathbf{S} = (\mathbf{X} \text{Diag}(\mathbf{w})^2 \mathbf{X}^\top + \mu_i \mathbf{I})^{1/2}$.

Optimizing over \mathbf{w} is a least-squares problem penalized by a reweighted ℓ_2 -norm equal to $\mathbf{w}^\top \mathbf{D}\mathbf{w}$, where $\mathbf{D} = \text{Diag}(\text{diag}(\mathbf{X}^\top \mathbf{S}^{-1} \mathbf{X}))$. It is equivalent to solving the linear system

$$(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{D}) \mathbf{w} = \mathbf{X}^\top \mathbf{y}.$$

This can be done efficiently by using a conjugate gradient method. Since the cost of multiplying $(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{D})$ by a vector is $O(np)$, solving the system has a complexity of $O(knp)$, where $k \leq n+1$ is the number of iterations needed to converge (see theorem 10.2.5 of [9]). Using warm restarts, k can be even smaller than n , since the linear system we are solving does not change a lot from an iteration to another. Below we summarize the algorithm:

ITERATIVE ALGORITHM FOR ESTIMATING \mathbf{w}

Input: the design matrix \mathbf{X} , the initial guess \mathbf{w}^0 , number of iteration N , sequence μ_i .

For $i = 1 \dots N$:

- Compute the eigenvalue decomposition $\mathbf{U} \text{Diag}(s_k) \mathbf{U}^\top$ of $\mathbf{X} \text{Diag}(\mathbf{w}^{i-1})^2 \mathbf{X}^\top$.
 - Set $\mathbf{D} = \text{Diag}(\text{diag}(\mathbf{X}^\top \mathbf{S}^{-1} \mathbf{X}))$, where $\mathbf{S}^{-1} = \mathbf{U} \text{Diag}(1/\sqrt{s_k + \mu_i}) \mathbf{U}^\top$.
 - Set \mathbf{w}^i by solving the system $(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{D}) \mathbf{w} = \mathbf{X}^\top \mathbf{y}$.
-

For the sequence μ_i , we use a decreasing sequence converging to ten times the machine precision.

3.1 Choice of λ

We now give a method to choose the initial parameter λ of the regularization path. In fact, we know that the vector $\mathbf{0}$ is solution if and only if $\lambda \geq \Omega^*(\mathbf{X}^\top \mathbf{y})$ [25]. Thus, we need to start the path at $\lambda = \Omega^*(\mathbf{X}^\top \mathbf{y})$, corresponding to the empty solution $\mathbf{0}$, and then decrease λ . Using the inequalities on the dual norm we obtained in the previous section, we get

$$\|\mathbf{X}^\top \mathbf{y}\|_\infty \leq \Omega^*(\mathbf{X}^\top \mathbf{y}) \leq \|\mathbf{X}\|_{op} \|\mathbf{X}^\top \mathbf{y}\|_\infty.$$

Therefore, starting the path at $\lambda = \|\mathbf{X}\|_{op} \|\mathbf{X}^\top \mathbf{y}\|_\infty$ is a good choice.

4 Approximation around the Lasso

We recall that when $\mathbf{P} = \mathbf{I} \in \mathbb{R}^{p \times p}$, our norm is equal to the ℓ_1 -norm, and we want to understand its behavior when \mathbf{P} departs from the identity. Thus, we compute a second order approximation of our norm around the Lasso: we add a small perturbation $\Delta \in \mathbb{S}_p$ to the identity matrix, and using Prop. 6 of [22], we obtain the following second order approximation:

$$\begin{aligned} \|(\mathbf{I} + \Delta) \text{Diag}(\mathbf{w})\|_* &= \|\mathbf{w}\|_1 + \text{diag}(\Delta)^\top |\mathbf{w}| + \\ &\sum_{|w_i| > 0} \sum_{|w_j| > 0} \frac{(\Delta_{ji}|w_i| - \Delta_{ij}|w_j|)^2}{4(|w_i| + |w_j|)} + \sum_{|w_i|=0} \sum_{|w_j| > 0} \frac{(\Delta_{ij}|w_j|)^2}{2|w_j|} + o(\|\Delta\|^2). \end{aligned}$$

We can rewrite this approximation as

$$\|(\mathbf{I} + \Delta) \text{Diag}(\mathbf{w})\|_* = \|\mathbf{w}\|_1 + \text{diag}(\Delta)^\top |\mathbf{w}| + \sum_{i,j} \frac{\Delta_{ij}^2 (|w_i| - |w_j|)^2}{4(|w_i| + |w_j|)} + o(\|\Delta\|^2),$$

using a slight abuse of notation, considering that the last term is equal to 0 when $w_i = w_j = 0$. The second order term is quite interesting: it shows that when two covariates are correlated, the effect of the trace Lasso is to shrink the corresponding coefficients toward each other. We also note that this term is very similar to pairwise elastic net penalties, which are of the form $|\mathbf{w}|^\top \mathbf{P} |\mathbf{w}|$, where \mathbf{P}_{ij} is a decreasing function of Δ_{ij} .

5 Experiments

In this section, we perform experiments on synthetic data to illustrate the behavior of the trace Lasso and other classical penalties when there are highly correlated covariates in the design matrix. The support S of \mathbf{w} is equal to $\{1, \dots, k\}$, where k is the size of the support. For i in the support of \mathbf{w} , w_i is independently drawn from a uniform distribution over $[-1, 1]$. The observations \mathbf{x}_i are drawn from a multivariate Gaussian with mean $\mathbf{0}$ and covariance matrix Σ . For the first setting, Σ is set to the identity, for the second setting, Σ is block diagonal with blocks equal to $0.2\mathbf{I} + 0.8\mathbf{1}\mathbf{1}^\top$ corresponding to clusters of four variables, finally for the third setting, we set $\Sigma_{ij} = 0.95^{|i-j|}$, corresponding to a Toeplitz design. For each method, we choose the best λ . We perform a first series of experiments ($p = 1024$, $n = 256$) for which we report the estimation error. For the second series of experiments ($p = 512$, $n = 128$), we report the Hamming distance between the estimated support and the true support.

In all six graphs of Figure 2, we observe behaviors that are typical of Lasso, ridge and elastic net: the Lasso performs very well on very sparse models but its performance degrades for denser models. The elastic net performs better than the Lasso for settings where there are strongly correlated covariates, thanks to its strongly convex ℓ_2 term. In setting 1, since the variables are uncorrelated, there is no reason to couple their selection. This suggests that the Lasso should be the most appropriate convex regularization. The trace Lasso approaches the Lasso when n is much larger than p , but the weak coupling induced by empirical correlations is sufficient to slightly decrease its performance compared to that of the Lasso. By contrast, in settings 2 and 3, the trace Lasso outperforms other methods (including the pairwise elastic net) since variables that should be selected together are indeed correlated. As for the penalized elastic net, since it takes into account the correlations between variables, it is not surprising that in experiments 2 and 3 it performs better than methods that do not. We do not have a compelling explanation for its superior performance in experiment 1.

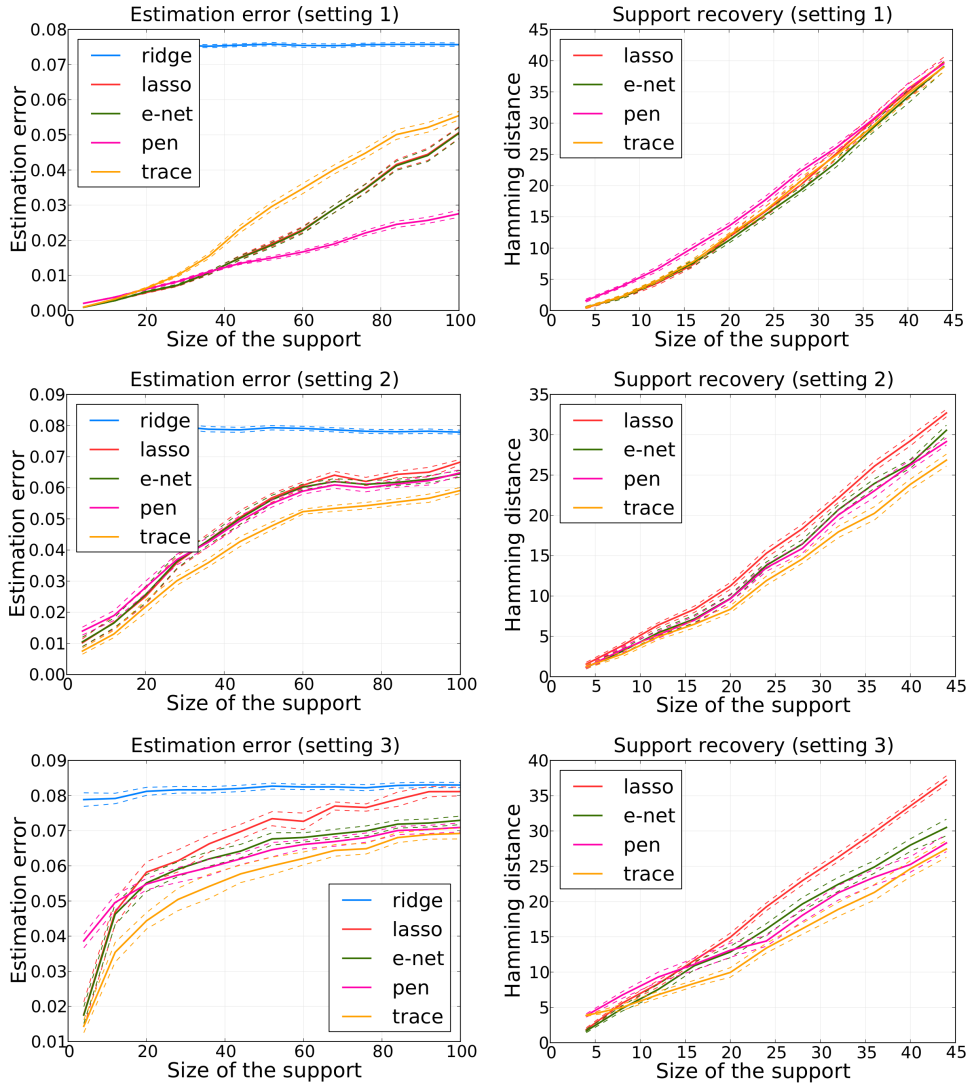


Figure 2: Left: estimation error ($p = 1024, n = 256$), right: support recovery ($p = 512, n = 128$). (Best seen in color. e-net stands for elastic net, pen stands for pairwise elastic net and trace stands for trace Lasso. Error bars are obtained over 20 runs.)

6 Conclusion

We introduce a new penalty function, the trace Lasso, which takes advantage of the correlation between covariates to add strong convexity exactly in the directions where needed, unlike the elastic net for example, which blindly adds a squared ℓ_2 -norm term in every directions. We show on synthetic data that this adaptive behavior leads to better estimation performance. In the future, we want to show that if a dedicated norm using prior knowledge such as the group Lasso can be used, the trace Lasso will behave similarly and its performance will not degrade too much, providing theoretical guarantees to such adaptivity. Finally, we will seek applications of this estimator in inverse problems such as deblurring, where the design matrix exhibits strong correlation structure.

Acknowledgments

Guillaume Obozinski and Francis Bach are supported in part by the European Research Council (SIERRA ERC-239993).

References

- [1] S.G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41(12):3397–3415, 1993.
- [2] T. Zhang. Adaptive forward-backward greedy algorithm for sparse learning with linear models. *Advances in Neural Information Processing Systems*, 22, 2008.
- [3] M.W. Seeger. Bayesian inference and optimal design for the sparse linear model. *The Journal of Machine Learning Research*, 9:759–813, 2008.
- [4] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- [5] S.S. Chen, D.L. Donoho, and M.A. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1999.
- [6] P.J. Bickel, Y. Ritov, and A.B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [7] P. Zhao and B. Yu. On model selection consistency of Lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.
- [8] M.J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *Information Theory, IEEE Transactions on*, 55(5):2183–2202, 2009.
- [9] G.H. Golub and C.F. Van Loan. *Matrix computations*. Johns Hopkins Univ Pr, 1996.
- [10] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [11] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [12] F. Bach. Consistency of the group Lasso and multiple kernel learning. *The Journal of Machine Learning Research*, 9:1179–1225, 2008.
- [13] F. Bach. Bolasso: model consistent Lasso estimation through the bootstrap. In *Proceedings of the 25th international conference on Machine learning*, pages 33–40. ACM, 2008.
- [14] H. Liu, K. Roeder, and L. Wasserman. Stability approach to regularization selection (stars) for high dimensional graphical models. *Advances in Neural Information Processing Systems*, 23, 2010.
- [15] N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [16] A. Tikhonov. Solution of incorrectly formulated problems and the regularization method. In *Soviet Math. Dokl.*, volume 5, page 1035, 1963.
- [17] A.E. Hoerl and R.W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [18] G. Davis, S. Mallat, and M. Avellaneda. Adaptive greedy approximations. *Constructive approximation*, 13(1):57–98, 1997.
- [19] E.J. Candès and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.
- [20] A. Lorbort, D. Eis, V. Kostina, D. M. Blei, and P. J. Ramadge. Exploiting covariate similarity in sparse regression via the pairwise elastic net. *JMLR - Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 9:477–484, 2010.
- [21] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. 2001.
- [22] E. Grave, G. Obozinski, and F. Bach. Trace lasso: a trace norm regularization for correlated designs. *Technical report, arXiv:1109.1990*, 2011.
- [23] A. Argyriou, T. Evgeniou, and M. Pontil. Multi-task feature learning. *Advances in neural information processing systems*, 19:41, 2007.
- [24] S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge Univ Pr, 2004.
- [25] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Convex optimization with sparsity-inducing norms. *S. Sra, S. Nowozin, S. J. Wright., editors, Optimization for Machine Learning*, 2011.