
Independently Interpretable Lasso: A New Regularizer for Sparse Regression with Uncorrelated Variables

Masaaki Takada^{1,2}
mtakada@ism.ac.jp

Taiji Suzuki^{3,4,5}
taiji@mist.i.u-tokyo.ac.jp

Hironori Fujisawa^{1,6}
fujisawa@ism.ac.jp

¹School of Advanced Sciences, The Graduate University for Advanced Studies

²Toshiba Corporation

³Graduate School of Information Science and Technology, The University of Tokyo

⁴PRESTO, Japan Science and Technology Agency, Japan

⁵Center for Advanced Integrated Intelligence Research, RIKEN, Tokyo, Japan

⁶The Institute of Statistical Mathematics

Abstract

Sparse regularization such as ℓ_1 regularization is a quite powerful and widely used strategy for high dimensional learning problems. The effectiveness of sparse regularization has been supported practically and theoretically by several studies. However, one of the biggest issues in sparse regularization is that its performance is quite sensitive to correlations between features. Ordinary ℓ_1 regularization can select variables correlated with each other, which results in deterioration of not only its generalization error but also interpretability. In this paper, we propose a new regularization method, “Independently Interpretable Lasso” (IILasso). Our proposed regularizer suppresses selecting correlated variables, and thus each active variable independently affects the objective variable in the model. Hence, we can interpret regression coefficients intuitively and also improve the performance by avoiding overfitting. We analyze theoretical property of IILasso and show that the proposed method is much advantageous for its sign recovery and achieves almost minimax optimal convergence rate. Synthetic and real data analyses also indicate the effectiveness of IILasso.

1 Introduction

High dimensional data appears in many fields such as biology, economy and industry. A common approach for high dimensional regression is sparse regularization strategy such as Lasso (Least absolute shrinkage and selection operator) [25]. Since the sparse regularization performs both parameter estimation and feature selection simultaneously, (i) it offers *interpretable* results by identifying informative variables, and (ii) it can effectively *avoid overfitting* by discarding redundant variables. Because of these properties, sparse regularization has shown huge success in wide range of data analysis in science and engineering. Moreover, several theoretical studies have been developed to support the effectiveness of sparse regularization, and efficient optimization methods also have been proposed so that sparse learning is easily executed.

However, the performance of sparse regularization is guaranteed only under “small correlation” assumptions, that is, the features are not much correlated with each other. Actually, typical theoretical supports are based on a kind of small correlation assumptions such as *restricted eigenvalue condition* [3, 2]. In the situation where the features are highly correlated, the selected variables are likely to be jointly correlated. As a result, each coefficient cannot be seen as independent variable contribution so that the model is no longer easy to interpret. This kind of interpretability is called “decomposability” in [14], which represents the ability whether we can decompose a model into some parts and interpret each component. In our experience, this kind of interpretability is quite important in many practical modeling for decision-making. In addition to this, selecting several correlated variables

results in worse generalization error because the redundant representation tends to give overfitting. Thus, it is favorable to construct the model with uncorrelated variables both for interpretability and generalization ability.

Several methods have been proposed to resolve the problem induced by correlations among variables. The first line of research is based on a strategy in which correlated variables are either all selected or not selected at all. Examples of this line are Elastic Net [30], Pairwise Elastic Net [15] and Trace Lasso [11]. These methods select not only important variables but also variables strongly correlated. Although these methods often give better generalization error, this strategy makes it hard to interpret the model, because many correlated variables are incorporated into the final model.

The second line of research including our proposed method is based on another strategy in which we select uncorrelated variables, and thus obtain decomposability. Uncorrelated Lasso [5] intends to construct a model with uncorrelated variables. However, it still tends to select “negatively” correlated variables and hence the correlation problem is not resolved. Exclusive Group Lasso [13] is also in this line, but it is necessary to group correlated variables beforehand. They suggest that we group variables whose correlations are greater than a certain threshold. However, determination of the threshold is not a trivial problem and practically it causes unstable results.

In this paper, we propose a new regularization method, named “Independently Interpretable Lasso” (IILasso), which offers efficient variable selection, does not select negatively correlated variables, and is free from a specific pre-processing such as grouping. Our proposed regularization formulation more aggressively induces the sparsity of the active variables and reduces the correlations among them. Hence, we can independently interpret the effects of active variables in the output model, and the generalization performance is also much improved. To support the usefulness of our proposal, we give the following contributions:

- We show the necessary and sufficient condition for the sign consistency of variables selection. As a result, it is shown that our method achieves the sign consistency under a milder condition than Lasso.
- The convergence rate of the estimation error is analyzed. We show that the estimation error achieves the almost minimax optimal rate and gives better performance than Lasso in some situations.

- We propose a coordinate descent algorithm to find a local optimum of the objective function. This is guaranteed to converge to a stationary point.
- It will be shown that every local optimal solution achieves the same statistical error rate as the global optimal solution and thus is almost minimax optimal though the objective function is not necessarily convex for the sake of better statistical properties.

The rest of this paper is as follows: In Section 2, we propose a new regularization formulation and introduce its optimization method. We also state the relationship with existing works. In Section 3, we show theoretical results on the sign recovery and the convergence rate. We can see that IILasso is superior to Lasso for correlated design. In Section 4, both synthetic and real-world data experiments, including 10 microarray datasets, are illustrated. In Section 5, we summarize the properties of IILasso.

Notations. Let $M \in \mathbb{R}^{n \times p}$. We use subscripts for the columns of M , i.e., M_j denotes the j -th column. Let $v \in \mathbb{R}^p$. $\text{Diag}(v) \in \mathbb{R}^{p \times p}$ is the diagonal matrix whose j -th diagonal element is v_j . $|v|$ is the element-wise absolute vector whose j -th element is $|v_j|$. $\text{sgn}(v)$ is the sign vector whose elements are 1 for $v_j > 0$, -1 for $v_j < 0$, and 0 for $v_j = 0$. $\text{supp}(v)$ is the support set of v , i.e., $\{j \in \{1, \dots, p\} | v_j \neq 0\}$. Let S be a subset of $\{1, \dots, p\}$. $|S|$ is the number of the elements in S . S^c is the complement subset of S , i.e., $S^c = \{1, \dots, p\} \setminus S$. v_S is the vector v restricted to the index set S . $M_{S_1 S_2}$ is the matrix whose row indexes are restricted to S_1 and column indexes are restricted to S_2 .

2 Proposed Method

2.1 IILasso: A New Regularization Formulation

Consider the problem of predicting $y \in \mathbb{R}^n$, given a design matrix $X \in \mathbb{R}^{n \times p}$, assuming a linear model

$$y = X\beta + \epsilon,$$

where $\epsilon \in \mathbb{R}^n$ is a noise and $\beta \in \mathbb{R}^p$ is a regression coefficient. We assume without loss of generality that the features are standardized such that $\sum_{i=1}^n X_{ij} = 0$, $\sum_{i=1}^n X_{ij}^2/n = 1$ and $\sum_{i=1}^n y_i = 0$. Then, Lasso solves the following optimization problem:

$$\min_{\beta} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1,$$

where $\lambda > 0$ is a regularization parameter. Since the last penalty term induces the sparsity of estimated values, the output model contains few variables and hence




Figure 1: Contours of penalties such that $\|\beta\|_1 + \frac{1}{2}|\beta|^T R|\beta| = 0.5, 1, 1.5$ with $\beta = [\beta_1, \beta_2]$

is tractable. This is critically important for interpretability. However, as we have described in Section 1, when there are highly correlated variables, Lasso can select correlated variables especially for a small λ resulting in worse interpretability and poor generalization error.

To overcome this problem, we propose a new regularization formulation as follows:

$$\min_{\beta} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \left(\|\beta\|_1 + \frac{\alpha}{2} |\beta|^T R |\beta| \right), \quad (1)$$

where $\alpha > 0$ is a regularization parameter for the new regularization term, and $R \in \mathbb{R}^{p \times p}$ is a symmetric matrix whose component $R_{jk} \geq 0$ represents the similarity between X_j and X_k . The last term of (1) is also written as $\frac{\lambda\alpha}{2} \sum_{j=1}^p \sum_{k=1}^p R_{jk} |\beta_j| |\beta_k|$. We define R_{jk} for $j \neq k$ as a monotonically increasing function of the absolute correlation $r_{jk} = \frac{1}{n} |X_j^T X_k|$, so that correlated variables are hard to be selected simultaneously. In particular, when X_j and X_k are strongly correlated, the squared error does not change under the condition that $\beta_j + \beta_k$ is constant, but the penalty $R_{jk} |\beta_j| |\beta_k|$ strongly induces either $\beta_j = 0$ or $\beta_k = 0$. On the contrary, if X_j and X_k are uncorrelated, i.e., R_{jk} is small, then the penalty of selecting both β_j and β_k is negligible.

We can see the exclusive effect of our regularization term by the constraint regions corresponding to the penalties. Figure 1 illustrates the constraint regions

of $\|\beta\|_1 + \frac{1}{2}|\beta|^T R|\beta|$ for the case $p = 2$. As diagonal elements of R increases (from the top to the bottom panel), the contours become smooth at the axis of coordinates. Because of this, the minimizer tends to select both variables if two variables are strongly correlated. This is the grouping effect of Elastic Net as we describe later. On the other hand, as off-diagonal elements of R increases (from the left to the right panel), the contours become pointed at the axis of coordinates and the minimizer tends to be sparser. This is the exclusive effect for correlated variables. Although the shape of contours resemble that of ℓ_q ($0 < q < 1$) penalty, we would emphasize that we use correlation information R_{jk} and multiplication term $|\beta_j| |\beta_k|$. Since our regularization term is adaptive for correlations, our penalty achieves both sparse and stable solutions.

Some definition variations of the similarity matrix R can be considered. One of the natural choices is $R_{jk} = r_{jk}^2$. R is positive semidefinite in this case because the Hadamard product of positive semidefinite matrices is also positive semidefinite. Hence, the problem (1) turns to be convex and easy to solve the global optimum. However, it may not reduce correlations enough. Another choice is $R_{jk} = |r_{jk}|$, which reduces correlations more strongly. Another effective choice is $R_{jk} = |r_{jk}| / (1 - |r_{jk}|)$ for $j \neq k$ and $R_{jk} = 0$ for $j = k$. In this case, if a correlation between two certain variables becomes higher, i.e., $r_{jk} \rightarrow 1$, then the penalty term diverges infinitely and Lasso cannot si-




Figure 2: Uncorrelated model vs correlated model

multaneously select both of them. We use the last one in our numerical experiments, because it is favorable from theoretical studies, as described in Section 3.

Now, let us emphasize the advantage of uncorrelated models for interpretability with a simple example (Figure 2). Suppose $X = [X_1, X_2, X_3] \in \mathbb{R}^{n \times 3}$ is standardized, X_1 and X_2 are orthogonal, and $X_3 = (X_1 + X_2)/\sqrt{2}$. Consider two models: (A) $y = 2X_1 + X_2$ and (B) $y = X_1 + \sqrt{2}X_3$. Both models output the same prediction. Which model do you think is more interpretable? We believe that the model (A) is more interpretable than (B), because active variables are uncorrelated so that we can see each coefficient as independent variable contribution. Although Lasso selects (B) because ℓ_1 norm of its coefficients is small, IILasso for α large enough selects (A) because our regularization term excludes correlations.

2.2 Optimization

To solve (1), we introduce Coordinate Descent Algorithm (CDA), which was originally proposed for Lasso ($\alpha = 0$ for IILasso) [8, 9]. CDA is a simple and efficient algorithm, particularly in high dimensions. CDA basically follows simply: For each $j \in \{1, \dots, p\}$, we optimize the objective function with respect to β_j with the remaining elements of β fixed at their most recently updated values.

CDA is applicable even when the quadratic penalty is included. Let $L(\beta)$ denote the objective function in (1). To derive the update equation, when $\beta_j \neq 0$, differentiating $L(\beta)$ with respect to β_j yields

$$\begin{aligned} \partial_{\beta_j} L(\beta) = & -\frac{1}{n} X_j^\top (y - X_{-j} \beta_{-j}) + (1 + \lambda \alpha R_{jj}) \beta_j \\ & + \lambda (1 + \alpha R_{j,-j} |\beta_{-j}|) \operatorname{sgn}(\beta_j), \end{aligned}$$

where β_{-j} denotes β without the j -th component, X_{-j}

Algorithm 1 CDA for IILasso

```

for  $\lambda = \lambda_{\max}, \dots, \lambda_{\min}$  do
  initialize  $\beta$ 
  while until convergence do
    for  $j = 1, \dots, p$  do
       $\beta_j \leftarrow \frac{1}{1 + \lambda \alpha R_{jj}} S\left(\frac{1}{n} X_j^\top (y - X_{-j} \beta_{-j}), \lambda (1 + \alpha R_{j,-j} |\beta_{-j}|)\right)$ 
    end for
  end while
end for

```

denotes X without j -th column and $R_{j,-j}$ denotes the j -th row vector without j -th column of R . Solving $\partial_{\beta_j} L(\beta) = 0$, we obtain the update rule as

$$\beta_j \leftarrow \frac{1}{1 + \lambda \alpha R_{jj}} S\left(\frac{1}{n} X_j^\top (y - X_{-j} \beta_{-j}), \lambda (1 + \alpha R_{j,-j} |\beta_{-j}|)\right), \quad (2)$$

where $S(z, \gamma)$ is a soft thresholding function

$$\begin{aligned} S(z, \gamma) &:= \operatorname{sgn}(z)(|z| - \gamma)_+ \\ &= \begin{cases} z - \gamma & \text{if } z > 0 \text{ and } \gamma < |z|, \\ z + \gamma & \text{if } z < 0 \text{ and } \gamma < |z|, \\ 0 & \text{if } |z| \leq \gamma. \end{cases} \end{aligned}$$

The whole algorithm for solving IILasso is described in Algorithm 1. We search several λ from λ_{\max} to λ_{\min} . β is initialized at each λ in some way such as a) zeros for all elements, b) the solution of previous λ , or c) the solution of ordinary Lasso.

In Algorithm 1, we can see that the objective function monotonically decreases at each update and the estimate converges a stationary point.

Proposition 1. *Let $\{\beta^t\}_{t=0,1,\dots}$ be a sequence of β in Algorithm 1. Then, every cluster point of $\{\beta^t\}_{t \equiv (p-1) \bmod p}$ is a stationary point.*

Proof. The proof is based on Theorem 4.1 in [26]. First, we can see that the level set $\{\beta | L(\beta) \leq L(\beta^0)\}$ is compact and $L(\beta)$ is continuous. Moreover, $L(\beta)$ has a unique minimum with (2) in terms of β_j . Therefore, every cluster point of $\{\beta^t\}_{t \equiv (p-1) \bmod p}$ is a coordinate-wise minimum point. In addition, since $L(\beta)$ can be seen as a locally quadratic function in any directions, $L(\beta)$ is *regular* at the cluster point. These and Theorem 4.1 (c) in [26] concludes the assertion. \square

2.3 Related Work

There are some existing works that take correlations among variables into account. We can divide them into

mainly two directions: 1) grouping selection and 2) exclusive selection. The former groups correlated variables and selects correlated variables together. The latter excludes correlated variables and selects uncorrelated variables.

The representative method of grouping selection is Elastic Net [30]. The objective function is constructed by squared ℓ_2 penalty in addition to ℓ_1 penalty:

$$\min_{\beta} \|y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2.$$

Due to ℓ_2 penalty, if the variables X_j and X_k are strongly correlated tending to be 1, then estimated values of β_j and β_k get closer. If X_j and X_k are equal, especially, then β_j and β_k must be equal. This behavior is called grouping effect. Pairwise Elastic Net [15] and Trace Lasso [11] are the same direction of research (and improve the prediction accuracy). These methods tend to include many correlated variables and each coefficient no longer indicates independent variable contribution. As a result, it is hard to interpret which variables are truly active and how variables affect the objective variable.

Another direction of research is exclusive selection. Uncorrelated Lasso (ULasso) [5] aims to reduce correlations among active variables. It optimizes the following objective function:

$$\min_{\beta} \|y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \beta^\top R \beta, \quad (3)$$

where $R \in \mathbb{R}^{p \times p}$ with each element $R_{jk} = (\frac{1}{n} X_j^\top X_k)^2$. ULasso quite resembles our formulation, but there exists a critical difference that they use β instead of $|\beta|$ in the objective function (3). We found that ULasso does not necessarily select uncorrelated variables. For example, consider the case $X = [X_1, X_2]$. The last term of (3) is $\lambda_2(\beta_1^2 + \beta_2^2 + 2R_{12}\beta_1\beta_2)$. If $R_{12} \neq 0$, then the term $R_{12}\beta_1\beta_2$ encourages $|\beta_1\beta_2|$ larger with $\beta_1\beta_2 < 0$. This implies that ULasso tends to select correlated variables and set coefficients to the opposite sign. In particular, X_1 and X_2 are strongly correlated, then it reduces $\lambda_2(\beta_1 + \beta_2)^2$, which induces $\beta_1 = -\beta_2$. It is not a major problem when X_1 and X_2 are positively correlated, but is a significant problem when X_1 and X_2 are negatively correlated. This problem is overcome in our method, as described in Section 2.1. Therefore, the difference between their ULasso and our IILasso is essential and crucial.

Exclusive Group Lasso (EGLasso) [13] is also the same direction of exclusive selection. It optimizes the following objective function:

$$\min_{\beta} \|y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \sum_{k=1}^K \|\beta^{(k)}\|_1^2, \quad (4)$$

where $\beta^{(k)}$ consists of the variables of β within a group of predictors $g_k \subset \{1, \dots, p\}$ and K is the number of groups. The last term is ℓ_1/ℓ_2 penalty, which acts on exclusive feature selection. For example, when $p = 3, g_1 = \{1, 2\}$ and $g_2 = \{3\}$, then the last term becomes $\lambda_2((|\beta_1| + |\beta_2|)^2 + |\beta_3|^2)$. This enforces sparsity over each intra-group. They suggest that we put highly correlated variables into the same group in order to select uncorrelated variables. They use $|r_{ij}| > \theta$ with $\theta \approx 0.90$ as a threshold. EGLasso can be seen as a special case of IILasso. Let R be a group indicator matrix such as $R_{jk} = 1$ if X_j and X_k belong to the same group and $R_{jk} = 0$ otherwise. Then IILasso is reduced to EGLasso. For the above example, if we define similarity matrix $R = [1, 1, 0; 1, 1, 0; 0, 0, 1]$, then the last term of IILasso objective function (1) becomes $\lambda(\beta_1^2 + 2|\beta_1||\beta_2| + \beta_2^2 + \beta_3^2)$, which is the same as the last term of (4). As we see, EGLasso needs to determine the threshold θ and group variables beforehand. This can cause severely unstable estimation.

3 Theoretical Results

In this section, we show theoretical properties of IILasso. We first show the sign recovery condition of IILasso. Then, we derive the convergence rate of IILasso. These results are significantly important for interpretability and generalization ability. In addition, we show the property of local minimum, which implies that every local optimal solution achieves the same statistical error rate as the global optimal solution. In this section, let β^* denote the true parameter, S denote the true active sets, i.e., $S = \{j \in \{1, \dots, p\} | \beta_j^* \neq 0\}$, and $s = |S|$.

3.1 Sign Recovery

We give the necessary and sufficient condition of sign recovery.

Theorem 2. *Define*

$$U := \frac{1}{n} X_S^\top X_S + \lambda \alpha \text{Diag}(\text{sgn}(\beta_S^*)) R_{SS} \text{Diag}(\text{sgn}(\beta_S^*)),$$

$$V := \lambda \text{sgn}(\beta_S^*) + \lambda \alpha \text{Diag}(\text{sgn}(\beta_S^*)) R_{SS} \text{Diag}(\text{sgn}(\beta_S^*)) \beta_S^* - \frac{1}{n} X_S^\top \epsilon.$$

Assume U is invertible. Then, there exists a critical point $\hat{\beta}$ of (1) with correct sign recovery $\text{sgn}(\hat{\beta}) = \text{sgn}(\beta^)$ if and only if the following two conditions hold:*

$$\text{sgn}(\beta_S^* - U^{-1}V) = \text{sgn}(\beta_S^*), \quad (5)$$

$$\left| \frac{1}{n} X_{S^c}^\top X_S U^{-1}V + \frac{1}{n} X_{S^c}^\top \epsilon \right| \leq \lambda (1 + \alpha R_{S^c S} |\beta_S^* - U^{-1}V|), \quad (6)$$

where both of these vector inequalities are taken elementwise.

The proof is given in the supplementary material. The sign recovery condition is derived from the standard conditions for optimality. We note that $\alpha = 0$ reduces the condition into the ordinary Lasso condition in [27]. The invertible assumption of U is not restrictive because it is true for almost all λ if $X_S^T X_S$ is invertible, which is the same assumption as standard analysis of Lasso.

The condition of IILasso is milder than that of Lasso when R_{SS} is small enough, since (5) is the same as Lasso and (6) is easier to be satisfied unless $\alpha R_{S^c S} = 0$. This implies that IILasso is more favorable than Lasso from the viewpoint of sign recovery. In addition, we can see that a large value of $R_{S^c S}$ is favorable for sign recovery. That is, IILasso tends to succeed the sign recovery if the true active variables have small correlations and the inactive variables are strongly correlated with the true active variables.

Theorem 2 is *not* the condition of a global optimal solution. However, if global optimal solutions are finite, they must be a critical point. Hence, there exists a global optimal solution with correct sign recovery only if (5) and (6) hold.

3.2 Convergence Rate

Here we give the convergence rate of the estimation error of our method. Before we give the statement, the assumption and definition are prepared.

Assumption 1. $(\epsilon_i)_{i=1}^n$ is an i.i.d. sub-Gaussian sequence: $E[e^{t\epsilon_i}] \leq e^{\frac{\sigma^2 t^2}{2}}$ ($\forall t \in \mathbb{R}$) for $\sigma > 0$.

Definition 1 (Generalized Restricted Eigenvalue Condition (GRE(S, C, C'))). *Let a set of vectors $\mathcal{B}(S, C, C')$ for $C > 0, C' > 0$ be*

$$\mathcal{B}(S, C, C') := \left\{ \beta \in \mathbb{R}^p \mid \|\beta_{S^c}\|_1 + \frac{C'\alpha}{2} |\beta_{S^c}|^\top R_{S^c S^c} |\beta_{S^c}| + C'\alpha |\beta_{S^c}|^\top R_{S^c S} |\beta_S + \beta_S^*| \leq C \|\beta_S\|_1 \right\}.$$

Then, we assume $\phi_{\text{GRE}} > 0$ where

$$\phi_{\text{GRE}} = \phi_{\text{GRE}}(S, C, C') := \inf_{v \in \mathcal{B}(S, C, C')} \frac{v^\top \frac{1}{n} X^\top X v}{\|v\|_2^2}.$$

Definition 1 is a generalized notion of the *restricted eigenvalue condition* [3, 2] tailored for our regularization. One can see that, if $\alpha = 0$ or $C' = 0$, then ϕ_{GRE} is reduced to the ordinary restricted eigenvalue [3, 19] for the analysis of Lasso. Since there are additional terms related to $|\beta|^\top R |\beta|$, the set $\mathcal{B}(S, C, C')$ is smaller than that for the ordinary restricted eigenvalue if the same

C is used. In particular, the term $|\beta_{S^c}|^\top R_{S^c S} |\beta_S + \beta_S^*|$ on the left-hand side strongly restricts the amplitude of coefficients for unimportant variables (especially the variables with large $R_{S^c S}$), and thus promote independence among the selected variables.

GRE condition holds in quite general class of Gaussian design, because RE(S, C) condition is satisfied in general class of Gaussian design [18], and Assumption GRE(S, C, C') is milder than Assumption RE(S, C). Therefore, GRE condition is not so restrictive.

Here, for any $\delta > 0$, we define γ_n as

$$\gamma_n = \gamma_n(\delta) := \sigma \sqrt{\frac{2 \log(2p/\delta)}{n}}.$$

Then, we obtain the convergence rate of IILasso as follows.

Theorem 3. *Suppose that Assumption 1 is satisfied. Suppose R satisfy for all $j, k \in S$,*

$$0 \leq R_{jk} \leq D,$$

for some positive constant D , and the estimator $\hat{\beta}$ of (1) is approximately minimizing the objective function so that

$$L_{\lambda_n}(\hat{\beta}) \leq L_{\lambda_n}(\beta^*). \quad (7)$$

Fix any $0 < \delta < 1$, let the regularization parameters satisfy

$$3\gamma_n \leq \lambda_n \text{ and } \alpha \leq \frac{1}{4D \|\beta_S^*\|_1}.$$

Suppose that Assumption GRE($S, 3, \frac{3}{2}$) (Definition 1) is satisfied. Then, it holds that

$$\|\hat{\beta} - \beta^*\|_2^2 \leq \frac{16s\lambda_n^2}{\phi_{\text{GRE}}^2},$$

with probability $1 - \delta$.

The proof is given in the supplementary material. The obtained convergence rate is roughly evaluated as

$$\|\hat{\beta} - \beta^*\|_2^2 = O_p \left(\frac{s \log(p)}{n} \right),$$

which is almost the minimax optimal rate [20].

As is obvious from the proof of Theorem 3, we also have a little bit stricter bound

$$\|\hat{\beta} - \beta^*\|_2^2 \leq \frac{\left(\frac{8}{3} + 5\alpha D \|\beta_S^*\|_1 + \frac{3}{4} (\alpha D \|\beta_S^*\|_1)^2 \right)^2 s \lambda_n^2}{\phi_{\text{GRE}}^2},$$

under Assumption GRE($S, C, \frac{3}{2}$) (Definition 1) where $C = 2 + \frac{15}{4} \alpha D \|\beta_S^*\|_1 + \frac{9}{16} (\alpha D \|\beta_S^*\|_1)^2$, with high probability. This proof is also given in the supplementary material. We can easily see that, when $\alpha = 0$, then

the convergence rate analysis is reduced to the standard one for ordinary Lasso [3, 2]. Under well “interpretable” cases where the true non-zero components S are independent, i.e., $R_{SS} = O$, the error bounds for Lasso and IILasso are the same except for the term ϕ_{GRE}^2 . Since $R_{S^c S}$ and $R_{S^c S^c}$ shrink the set of vectors $\mathcal{B}(S, \mathcal{C}, \mathcal{C}')$ in Definition 1, ϕ_{GRE} of IILasso is larger than that of Lasso. Therefore, our method achieves a better error bound than the ordinary ℓ_1 regularization in this setting. In addition, our method has more advantageous when the variables are correlated between informative and non-informative variables.

3.3 Local Optimality of estimator

The objective function of our method is *not* necessarily convex in exchange for better statistical properties as observed above. Our next theoretical interest is about the global optimality of our optimization algorithm (Algorithm 1). Since our optimization method is a greedy one, there is no confirmation that it achieves the global optimum. However, as we see in this section, the local solution achieves almost the same estimation error as the global optimum satisfying (7). For theoretical simplicity, we assume the following a bit stronger condition.

Assumption 2. *There exists $\phi > 0$ and $q_n \leq p$ such that, for all $V \subset \{1, \dots, p\}$ satisfying $|V| \leq q_n$ and $V \cap S = \emptyset$, it holds that*

$$\frac{1}{n} X_{SUV}^\top X_{SUV} \succ \phi I.$$

Moreover, there exists \bar{D} such that the maximum absolute value of the eigenvalue of R is bounded as

$$\sup_{u \in \mathbb{R}^{S \cup V}} u^\top (R_{SUV, SUV}) u \leq \bar{D} \|u\|_2^2.$$

Theorem 4. *Suppose Assumptions 1 and 2 are satisfied. Assume that $\tilde{\beta}$ is a local optimal solution of (1) satisfying $|\text{supp}(\tilde{\beta})| \leq |S| + q_n$. Let the regularization parameters satisfy*

$$\gamma_n < \lambda_n \text{ and } \alpha < \min \left\{ \frac{\sqrt{s}}{2\bar{D}\|\beta^*\|_2}, \frac{\phi}{2\bar{D}\lambda_n} \right\},$$

for any $\delta > 0$. Then, $\tilde{\beta}$ should satisfy

$$\|\tilde{\beta} - \beta^*\|_2^2 \leq \frac{25s\lambda_n^2}{\phi^2},$$

with probability $1 - \delta$.

The proof is given in the supplementary material. Theorem 4 indicates that every local optimum $\tilde{\beta}$ achieves the same convergence rate with the ideal optimal solution $\hat{\beta}$. In other words, there is *no local optimal*

solution with sparsity level $|S| + q_n$ far from the true vector β^* . In the theorem, we assumed the sparsity of the local optimal solution $\tilde{\beta}$. Such a sparse solution can be easily obtained by running CDA multiple times from different initial solutions.

4 Numerical Experiments

4.1 Synthetic Data

We consider the case in which the true active variables are uncorrelated and many inactive variables are strongly correlated with the active variable. If all of active and inactive variables are uncorrelated, it is easy to estimate which is active or inactive. On the other hand, if the inactive variables are strongly correlated with the active variables, it is hard to distinguish which one is active. We simulate such a situation and validate the effectiveness of IILasso.

First, we generated a design matrix $X \in \mathbb{R}^{n \times p}$ from the Gaussian distribution of $\mathcal{N}(0, \Sigma)$ where $\Sigma = \text{Diag}(\Sigma^{(1)}, \dots, \Sigma^{(b)})$ was a block diagonal matrix whose element $\Sigma^{(l)} \in \mathbb{R}^{q \times q}$ was $\Sigma_{jk}^{(l)} = 0.95$ for $j \neq k$ and $\Sigma_{jk}^{(l)} = 1$ for $j = k$. We set $n = 50$, $p = 100$, $b = 10$ and $q = 10$. Thus, there were 10 groups containing 10 strongly correlated variables. Next, we generated an objective variable y by the true active variables $X_1, X_{11}, X_{21}, \dots, X_{91}$, such that $y = 10X_1 - 9X_{11} + 8X_{21} - 7X_{31} + \dots + 2X_{81} - X_{91} + \epsilon$, with a standard Gaussian noise ϵ . Each group included one active variable. We generated three datasets for training, validation and test as above.

Then, we compared the performance of Lasso, SCAD [7], MCP [29], EGLasso and IILasso. Evaluation criteria are prediction error (mean squared error), estimation error (ℓ_2 norm between the true and estimated coefficients) and model size (the number of non-zero coefficients). SCAD and MCP are representative methods of folded concave penalty, so their objective functions are non-convex, which are the same as our method. They have a tuning parameter γ ; we set $\gamma = 2.5, 3.7, 10, 20, 100, 1000$ for SCAD and $\gamma = 1.5, 3, 10, 20, 100, 1000$ for MCP. EGLasso has a parameter λ_2 ; we set $\lambda_2/\lambda_1 = 0.01, 0.1, 1, 10, 100, 1000$. For EGLasso, we used the true group information beforehand. We used R packages `ncvreg` [4] for Lasso, SCAD and MCP, and `sparsenet` [16] for MCP. One can solve MCP using either `ncvreg` or `sparsenet`, but they differ in their optimization algorithms and ways of initialization. For IILasso, we defined $R_{jk} = |r_{jk}|/(1 - |r_{jk}|)$ for $j \neq k$ and $R_{jk} = 0$ for $j = k$. Hence, R_{SS} takes small values if active variables are independent, and $R_{S^c S^c}$ and $R_{S^c S}$ take large values if inactive variables are strongly correlated with other variables, which is

Table 1: Results of synthetic data

	prediction error	estimation error	model size
Lasso(<code>ncvreg</code>)	2.67(0.05)	4.44(0.06)	34.1(0.46)
SCAD(<code>ncvreg</code>)	1.52(0.02)	1.79(0.04)	14.6(0.23)
MCP(<code>ncvreg</code>)	1.53(0.02)	1.79(0.04)	14.6(0.24)
MCP(<code>sparsenet</code>)	2.41(0.11)	3.15(0.13)	13.4(0.28)
EGLasso	2.60(0.04)	4.36(0.05)	33.3(0.32)
IILasso (ours)	1.45(0.02)	1.40(0.04)	13.5(0.23)

favorable from the theoretical results obtained in Section 3. We set $\alpha = 0.01, 0.1, 1, 10, 100, 1000$. We tuned the above parameters using validation data and calculated errors using test data. We iterated this procedure 500 times and evaluated the averages and standard errors.

Table 1 shows the performances with their standard error in parentheses. IILasso achieved the best prediction and estimation among all of them. This was because our penalty term excluded the correlations and avoided overfitting. Moreover, the model size of IILasso was much less than those of Lasso and EGLasso, and comparable to MCP. As a whole, IILasso could distinguish the true active variables.

4.2 Real Data: Gene Expression Data

We applied our method to various gene expression data to validate its effectiveness for real applications. We used the following 10 datasets: ‘alon’ [1] (colon cancer), ‘chiaretti’ [6] (leukemia), ‘gordon’ [10] (lung cancer), ‘gravier’ [12] (breast cancer), ‘pomeroy’ [17] (central nervous system disorders), ‘shipp’ [21] (lymphoma), ‘singh’ [22] (prostate cancer), ‘subramanian’ [23] (miscellaneous), ‘tian’ [24] (myeloma), ‘west’ [28] (breast cancer). All of these data are provided by R package `datamicroarray`. The abstract of these datasets is described in Table 2. All datasets are small-sample high-dimensional DNA microarray data.

Since the objective variable is binary, logistic regression was applied. Logistic regression of Lasso, SCAD and MCP is supported by `ncvreg` (`sparsenet` does not support logistic regression). EGLasso and IILasso can also be formulated as logistic regression. For details, see the supplementary material. We used the same settings on regularization parameters as described in Section 4.1. We evaluated the log-likelihood, misclassification error and model size using ten-fold cross validation.

The results are given in Figure 3. From the viewpoint of log-likelihood, IILasso won in 5 out of 10 cases. MCP showed similar performance to IILasso, but only won 2 cases. Lasso, EGLasso and SCAD showed sim-

Table 2: Abstract of microarray data

data	# samples	# dimensions
alon	62	2000
chiaretti	111	12625
gordon	181	12533
gravier	168	2905
pomeroy	60	7128
shipp	58	6817
singh	102	12600
subramanian	50	10100
tian	173	12625
west	49	7129

ilar performance, but fell behind IILasso and MCP. We can see the similar tendency of misclassification errors. IILasso won in 9 out of 10 cases including 5 ties. IILasso lost only 1 case, in which MCP won. In addition, the model size of IILasso was always smaller than Lasso and the smallest in 6 out of 10 cases including 1 tie. In particular, IILasso showed a much smaller misclassification error with a much smaller model size for the dataset ‘subramanian’, and smaller model sizes for the datasets ‘gordon’, ‘gravier’, ‘pomeroy’ and ‘west’ among the comparable methods with almost the same misclassification errors. As a whole, IILasso could construct accurate models with a small number of variables.

5 Conclusion

In this paper, we proposed a new regularization method “IILasso”. IILasso reduces correlations among the active variables, hence it is easy to decompose and interpret the model. We showed that the sign recovery condition of IILasso is milder than that of Lasso when the true active variables are uncorrelated with each other. The convergence rate of IILasso also has a better performance compared to that of Lasso. In addition, we showed that every local optimum by coordinate descent algorithm has the same order convergence rate as the global optimum. Finally, we verified the effectiveness of IILasso by synthetic and real data analyses using 10 gene expression data, and we saw that IILasso was superior in many cases on high-dimensional data.

Acknowledgement

TS was partially supported by MEXT kakenhi (25730013, 25120012, 26280009, 15H01678 and 15H05707), JST-PRESTO and JST-CREST. HF was partially supported by MEXT kakenhi 17K00065.




Figure 3: Results of 10 microarray datasets

References

- [1] U. Alon, N. Barkai, D. A. Notterman, K. Gish, S. Ybarra, D. Mack, and A. J. Levine. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *Proceedings of the National Academy of Sciences*, 96(12):6745–6750, 1999.
- [2] Peter Bhlmann and Sara van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, 2011.
- [3] Peter J Bickel, Ya’acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, pages 1705–1732, 2009.
- [4] P. Breheny and H. Jian. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *The annals of applied statistics*, 5(1):232–253, 2011.
- [5] S. B. Chen, C. Ding, B. Luo, and Y. Xie. Uncorrelated lasso. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 166–172, 2013.
- [6] S. Chiaretti, X. Li, R. Gentleman, A. Vitale, M. Vignetti, F. Mandelli, J. Ritz, and R. Foa. Gene expression profile of adult T-cell acute lymphocytic leukemia identifies distinct subsets of patients with different response to therapy and survival. *Blood*, 103(7):2771–2778, 2004.
- [7] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- [8] J. Friedman, T. Hastie, H. Hfling, and R. Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.
- [9] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2010.
- [10] Gordon G. J, R. V. Jensen, L. L. Hsiao, S. R. Gullans, J. E. Blumenstock, S. Ramaswamy, W. G. Richards, D. J. Sugarbaker, and R. Bueno. Translation of Microarray Data into Clinically Relevant Cancer Diagnostic Tests Using Gene Expression Ratios in Lung Cancer and Mesothelioma. *Cancer Research*, 62(17):4963–4967, 2002.
- [11] E. Grave, G. R. Obozinski, and F. R. Bach. Trace lasso: a trace norm regularization for correlated designs. In *Advances in Neural Information Processing Systems*, pages 2187–2195, 2011.
- [12] E. Gravier, G. Pierron, A. Vincent-Salomon, N. gruel, V. Raynal, A. Savignoni, Y. De Rycke, J. Y. Pierga, C. Lucchesi, F. Reyat, A. Fourquet, S. Roman-Roman, F. Radvanyi, X. Sastre-Garau, B. Asselain, and O. Delattre. A prognostic DNA signature for T1T2 node-negative breast cancer patients. *Genes, Chromosomes and Cancer*, 49(12):1125–1125, 2010.
- [13] D. Kong, R. Fujimaki, J. Liu, F. Nie, and C. Ding. Exclusive feature learning on arbitrary structures via $\ell_{1,2}$ -norm. In *Advances in Neural Information Processing Systems*, pages 1655–1663, 2014.
- [14] Z. C. Lipton. The mythos of model interpretability. In *ICML Workshop on Human Interpretability in Machine Learning*, pages 96–100, 2016.

- [15] A. Lobert, D. Eis, V. Kostina, D. Blei, and P. Ramage. Exploiting covariate similarity in sparse regression via the pairwise elastic net. In *Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 477–484, 2010.
- [16] R. Mazumder, J. H. Friedman, and T. Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106(495):1125–1138, 2011.
- [17] S. L. Pomeroy, P. Tamayo, M. Gaasenbeek, L. M. Sturla, M. Angelo, M. E. McLaughlin, J. Y. H. Kim, L. C. Goumnerova, P. M. Black, C. Lau, J. C. Allen, D. Zagzag, J. M. Olson, T. Curran, C. Wetmore, J. A. Biegel, T. Poggio, S. Mukherjee, R. Rifkin, A. Califano, G. Stolovitzky, D. N. Louis, J. P. Mesirov, E. S. Lander, and T. R. Golub. Prediction of central nervous system embryonal tumour outcome based on gene expression. *Nature*, 415(6870):436–442, 2002.
- [18] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- [19] Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11(Aug):2241–2259, 2010.
- [20] Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE transactions on information theory*, 57(10):6976–6994, 2011.
- [21] M. A. Shipp, K. N. Ross, P. Tamayo, A. P. Weng, J. L. Kutok, R. C. T. Aguiar, M. Gaasenbeek, M. Angelo, M. Reich, G. S. Pinkus, T. S. Ray, M. A. Koval, K. W. Last, A. Norton, T. A. Lister, J. Mesirov, D. S. Neuberg, E. S. Lander, J. C. Aster, and T. R. Golub. Diffuse large B-cell lymphoma outcome prediction by gene-expression profiling and supervised machine learning. *Nature Medicine*, 8(1):68–74, 2002.
- [22] D. Singh, P. G. Febbo, K. Ross, D. G. Jackson, J. Manola, C. Ladd, P. Tamayo, A. Renshaw, A. V. D’Amico, J. P. Richie, E. S. Lander, M. Loda, P. W. Kantoff, T. R. Golub, and W. R. Sellers. Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell*, 1(2):203–209, 2002.
- [23] A. Subramanian, P. Tamayo, V. K. Mootha, S. Mukherjee, B. L. Ebert, M. A. Gillette, A. Paulovich, S. L. Pomeroy, T. R. Golub, E. S. Lander, and J. P. Mesirov. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 102(43):15545–15550, 2005.
- [24] E. Tian, F. Zhan, R. Walker, E. Rasmussen, Y. Ma, B. Barlogie, and J. D. Shaughnessy Jr. The Role of the Wnt-Signaling Antagonist DKK1 in the Development of Osteolytic Lesions in Multiple Myeloma. *New England Journal of Medicine*, 349(26):2483–2494, 2003.
- [25] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)*, 58:267–288, 1996.
- [26] P. Tseng. Convergence of a Block Coordinate Descent Method for Nondifferentiable Minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, 2001.
- [27] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *The Annals of statistics*, 35(5):2183–2202, 2009.
- [28] M. West, C. Blanchette, H. Dressman, E. Huang, S. Ishida, R. Spang, H. Zuzan, J. A. Olson, J. R. Marks, and J. R. Nevins. Predicting the clinical status of human breast cancer by using gene expression profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 98(20):11462–11467, 2001.
- [29] C. H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics*, 38(2):894–942, 2010.
- [30] 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.