

Sparse Estimation with Strongly Correlated Variables using Ordered Weighted ℓ_1 Regularization

Mário A. T. Figueiredo

Instituto de Telecomunicações and Instituto Superior Técnico, Universidade de Lisboa, Portugal

Robert D. Nowak

Department of Electrical and Computer Engineering, University of Wisconsin-Madison, USA

September 13, 2014

Abstract

This paper studies *ordered weighted ℓ_1* (OWL) norm regularization for sparse estimation problems with strongly correlated variables. We prove sufficient conditions for clustering based on the correlation/colinearity of variables using the OWL norm, of which the so-called OSCAR [4] is a particular case. Our results extend previous ones for OSCAR in several ways: for the squared error loss, our conditions hold for the more general OWL norm and under weaker assumptions; we also establish clustering conditions for the absolute error loss, which is, as far as we know, a novel result. Furthermore, we characterize the statistical performance of OWL norm regularization for generative models in which certain clusters of regression variables are strongly (even perfectly) correlated, but variables in different clusters are uncorrelated. We show that if the true p -dimensional signal generating the data involves only s of the clusters, then $O(s \log p)$ samples suffice to accurately estimate the signal, regardless of the number of coefficients within the clusters. The estimation of s -sparse signals with completely independent variables requires just as many measurements. In other words, using the OWL we pay no price (in terms of the number of measurements) for the presence of strongly correlated variables.

1 Introduction

1.1 Definitions, Problem Formulation, and Preview of Main Results

The OWL (*ordered weighted ℓ_1*) regularizer is defined as

$$\Omega_{\mathbf{w}}(\mathbf{x}) = \sum_{i=1}^p w_i |x|_{[i]}, \quad (1)$$

where $|x|_{[i]}$ is the i -th largest component in magnitude of $\mathbf{x} \in \mathbb{R}^p$, and $\mathbf{w} \in \mathbb{R}_+^p$ is a vector of non-negative weights. If $w_1 \geq w_2 \geq \dots \geq w_p$ and $w_1 > 0$ (which we will assume to be always true), then $\Omega_{\mathbf{w}}$ is a norm (as shown in [3], [22]), which satisfies $w_1 \|\mathbf{x}\|_{\infty} \leq \Omega_{\mathbf{w}}(\mathbf{x}) \leq w_1 \|\mathbf{x}\|_1$. The OWL regularizer generalizes the OSCAR (*octagonal shrinkage and clustering algorithm for regression*) [4], which is obtained by setting $w_i = \lambda_1 + \lambda_2 (p - i)$, where $\lambda_1, \lambda_2 \geq 0$. Notice also that if $w_1 > 0$, and $w_2 = \dots = w_p = 0$, the OWL is simply (w_1 times) the ℓ_{∞} norm, whereas for $w_1 = w_2 = \dots = w_p$, the OWL becomes (w_1 times) the ℓ_1 norm.

In this paper, we will study the use of the OWL norm as a regularizer in linear regression with strongly correlated variables, both under the squared error loss and the absolute error loss, *i.e.*, the two following optimization problems:

$$\min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_2^2 + \Omega_{\mathbf{w}}(\mathbf{x}), \quad (2)$$

where $\mathbf{A} \in \mathbb{R}^{n \times p}$ is the design matrix, and

$$\min_{\mathbf{x} \in \mathbb{R}^p} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 + \Omega_{\mathbf{w}}(\mathbf{x}). \quad (3)$$

We also consider constrained versions of these problems; see (7) and (8) below.

The first of our two main results gives sufficient conditions for OWL norm regularization (with either the squared or the absolute error loss) to automatically cluster strongly correlated variables, in the sense that the coefficients associated with such variables have equal estimated values (in magnitude). The result for the squared error loss extends the main theorem about OSCAR in [4], since not only it applies to the more general case of OWL (of which OSCAR is a particular case), but it also holds under weaker conditions. Furthermore, the result for the absolute error loss is, as far as we know, novel.

Our second main result is a finite sample bound for the OWL regularization procedure, which includes the standard LASSO and OSCAR as special cases. To the best of our knowledge, these are the first finite sample error bounds for sparse regression with strongly correlated columns. To preview this result, consider the following special case (which we generalize further in the paper): assume that we observe

$$\mathbf{y} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\nu}. \quad (4)$$

where $\boldsymbol{\nu} \in \mathbb{R}^n$ is the measurement error satisfying

$$\frac{1}{n} \|\boldsymbol{\nu}\|_1 \leq \varepsilon, \quad (5)$$

and about which we make no other assumptions. The measurement/design matrix \mathbf{A} is Gaussian distributed. For the purposes of this introduction, assume that each column of \mathbf{A} has i.i.d. $\mathcal{N}(0, 1)$ entries, but that the columns may be correlated. Specifically, assume that the columns can be grouped so that columns within each group are identical (apart from a possible sign flip) and columns in different groups are uncorrelated. This models cases in which certain variables are perfectly correlated with each other, but uncorrelated with all others. The vector $\mathbf{x}^* \in \mathbb{R}^p$ is assumed to satisfy $\|\mathbf{x}^*\|_1 \leq \sqrt{s}$. Note, for example, that this condition is met if $\|\mathbf{x}^*\|_2 \leq 1$ and \mathbf{x}^* has at most s non-zero components (is s -sparse), which we assume to be true. Note that since certain columns of \mathbf{A} may be identical, in general there may be many sparse vectors \mathbf{x} such that $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}^*$. Thus, for now, assume that if two columns of \mathbf{A} are identical (up to a sign flip), then so are (in magnitude) the corresponding coefficients in \mathbf{x}^* . The following theorem essentially shows that the number of measurements sufficient to estimate an s -sparse signal (i.e., a signal with s nonzero groups of identical coefficients corresponding to identical columns in \mathbf{A}), with a given precision, grows like

$$n \sim s \log p. \quad (6)$$

This agrees with well-known sample complexity bounds for sparse recovery under stronger assumptions such as the restricted isometry property or i.i.d. measurements [6, 7, 10, 12, 21]. Moreover, this shows that *by using OWL we pay no price (in terms of the sufficient number of measurements) for colinearity of some columns of \mathbf{A} .*

Theorem 1.1. *Let \mathbf{y} , \mathbf{A} , \mathbf{x}^* , and ε be as defined above. Let $\Delta := \min\{w_l - w_{l+1}, l = 1, \dots, p-1\}$ be the minimum gap between two consecutive components of vector \mathbf{w} , and assume $\Delta > 0$. Let $\hat{\mathbf{x}}$ be a solution to either of the two following optimization problems:*

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_{\mathbf{w}}(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2, \quad (7)$$

or

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_{\mathbf{w}}(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon. \quad (8)$$

Then,

(i) for every pair of columns (i, j) for which¹ $\mathbf{a}_i = \mathbf{a}_j$, we have $\hat{x}_i = \hat{x}_j$;

(ii) the solution $\hat{\mathbf{x}}$ satisfies

$$\mathbb{E} \|\hat{\mathbf{x}} - \mathbf{x}^*\|_2 \leq \sqrt{2\pi} \left(4\sqrt{2} \frac{w_1}{\bar{w}} \sqrt{\frac{s \log p}{n}} + \varepsilon \right), \quad (9)$$

where $\bar{w} = p^{-1} \sum_{i=1}^p w_i$.

The expectation above (and elsewhere in the paper) is with respect to the Gaussian distribution of \mathbf{A} . Part (i) of this theorem is proved in Subsection 2.5, and part (ii) in Subsection 3.2. As mentioned above, in general there may be many sparse \mathbf{x} that yield the same value of $\mathbf{A}\mathbf{x}$. This is where the OWL norm becomes especially important. If the columns are colinear, then the OWL solution will select a representation including all the columns associated with the true model, rather than an arbitrary subset of them. We generalize Theorem 1.1 in the paper, and show the OWL norm yields similar clustering and recovery conditions for problems with strongly correlated, but not necessarily colinear, variables. Notice that the constant factor in the bound is typically a small constant; for example, in the OSCAR case we have $w_i = \lambda_1 + \lambda_2(p - i)$, thus $\bar{w} = \lambda_1 + \lambda_2(p - 1)/2$ and therefore $w_1/\bar{w} \leq 2$.

1.2 Related Work

Estimates obtained with the LASSO (*i.e.*, ℓ_1) regularizer can be difficult to interpret when columns of the measurement matrix \mathbf{A} are strongly correlated, because it may select only one of a group of highly correlated columns. For scientific and engineering purposes, one is often interested in identifying all of the columns that are important for modeling the data, rather than just a subset of them. Many researchers have proposed alternatives to the LASSO that aim at dealing with this problem. For example, Jia and Yu [14] study the *elastic net* regularizer (a combination of the ℓ_1 and the squared ℓ_2 norms), showing that it can consistently select the true model for certain correlated design matrices \mathbf{A} , when LASSO cannot. Marginal regression methods have also been shown by Genovese *et al* to perform better than the LASSO, in the presence of strongly correlated columns [11]. Stability selection procedures, can also aid in the selection of correlated columns, as shown by Meinshausen and Bühlmann [15], and Shah and Samworth [19]. Recently, Bühlmann *et al* [5] proposed and analyzed a two-stage approach called *cluster-LASSO*, which first identifies clusters of correlated columns, then groups them, and finally applies LASSO or group-LASSO to the groups; the cluster-LASSO is shown to be statistically consistent in certain cases. Adaptive grouping methods based on nonconvex optimizations have also been proposed by Shen and Huang [20], and shown to be asymptotically consistent under certain conditions.

Most closely related to this paper, is the so-called OSCAR (*octagonal shrinkage and clustering algorithm for regression*), proposed and analyzed by Bondell and Reich [4]. As mentioned above, OSCAR is a special case of the OWL regularizer, obtained with $w_i = \lambda_1 + \lambda_2(p - i)$. The OSCAR method has been shown to perform well in practice, but prior work has not addressed its statistical consistency or convergence properties. Motivated by this formulation of OSCAR, the OWL regularizer was recently proposed by Zeng *et al* [22], as a generalization thereof. The OWL norm was also independently proposed by Bodgan *et al* [2, 3], who set the weights to $w_i = F^{-1}(1 - iq/(2p))$, where F is the cumulative distribution function of the error variables, and $0 < q < 1$ is a parameter. Those authors showed that if \mathbf{A} is orthogonal, the solution to (2) with these weights has a false discovery rate for variable selection bounded by $q(p - k)/p$, where k is the number of non-zero coefficients in the true \mathbf{x} that generates \mathbf{y} .

On the computational side, a key tool for solving problems of the form (2), (3), (7), or (8), is the Moreau proximity operator of $\Omega_{\mathbf{w}}$ [1], defined as

$$\text{prox}_{\Omega_{\mathbf{w}}}(\mathbf{u}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 + \Omega_{\mathbf{w}}(\mathbf{x}).$$

¹It is trivial to extend the proof of this result to show that if $\mathbf{a}_i = -\mathbf{a}_j$, then $\hat{x}_i = -\hat{x}_j$

Efficient $O(p \log p)$ algorithms to compute prox_{Ω_w} have been recently proposed by Bodgan *et al* [2, 3], and by Zeng *et al* [22], who generalize to the OWL case the algorithm proposed by Zhong and Kwok [24]. Even more recently, Zeng *et al* [23] have show how Ω_w can be written explicitly as an *atomic norm* (see [8], for definitions), opening the door to the efficient use of the conditional gradient (also known as Frank-Wolfe) algorithm [13].

Notation

We denote (column) vectors by lower-case bold letters, *e.g.*, \mathbf{x} , \mathbf{y} , their transposes by \mathbf{x}^T , \mathbf{y}^T , the corresponding i -th and j -th components as x_i and y_j , and matrices by upper case bold letters, *e.g.*, \mathbf{A} , \mathbf{B} . A vector with all elements equal to 1 is denoted as $\mathbf{1}$ and $|\mathbf{x}|$ denotes the vector with the absolute values of the components of \mathbf{x} . Given some vector \mathbf{x} , $x_{[i]}$ is its i -th largest component (*i.e.*, for $\mathbf{x} \in \mathbb{R}^p$, $x_{[1]} \geq x_{[2]} \geq \dots \geq x_{[p]}$, with ties broken by some arbitrary rule); consequently, $|x|_{[i]}$ is the i -th largest component of x in magnitude. The vector obtained by sorting (in non-increasing order) the components of \mathbf{x} is denoted as \mathbf{x}_\downarrow , thus $|\mathbf{x}|_\downarrow$ denotes the vector obtained by sorting the components of \mathbf{x} in non-increasing order of magnitude (allowing to write $\Omega_w(\mathbf{x}) = \mathbf{w}^T |\mathbf{x}|_\downarrow$).

2 Sufficient Conditions for OWL Variable Clustering

2.1 Introduction

In this section, we study the solutions of (2) and (3) in the case where the design matrix \mathbf{A} has strongly correlated columns, and give corollaries for the particular case of OSCAR. The results presented below extend the main theorem of [4] in several ways: for the squared error loss, our result applies to the more general case of the OWL (of which OSCAR is a particular case) and it holds under weaker conditions; the result for the absolute error regression case is, as far as we know, novel.

2.2 Squared Error Loss with Correlated Columns

Consider the regression problem (2), and let $\mathbf{a}_i \in \mathbb{R}^n$ denote the i -th column (for $i = 1, \dots, p$) of matrix \mathbf{A} . The following theorem shows that (2) clusters (in the sense that the corresponding components of the solution are exactly equal in magnitude) the columns that are correlated enough.

Theorem 2.1. *Consider the objective function in (2) and assume, as is common practice in linear regression, that the columns of the matrix are normalized to a common norm, that is, $\|\mathbf{a}_k\|_2 = c$, for $k = 1, \dots, p$. Let $\hat{\mathbf{x}}$ be any minimizer of the objective function in (2). Then, for every pair of columns (i, j) for which $\|\mathbf{y}\| \|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2 < \Delta$ (where $\Delta := \min\{w_l - w_{l+1}, l = 1, \dots, p-1\}$ is the minimum gap between two consecutive components of vector \mathbf{w}), we have $|\hat{x}_i| = |\hat{x}_j|$.*

Notice that if two columns (affected by the signs of the corresponding regression coefficients) are identical, *i.e.*, if $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2 = 0$, any strictly positive value of Δ is sufficient to guarantee that these two columns will be clustered, that is, that the corresponding coefficients will be equal in magnitude.

The following corollary addresses the case where the columns of \mathbf{A} have zero mean and unit norm.

Corollary 2.1. *Let the columns of \mathbf{A} be normalized to zero sample mean and unit norm: $\mathbf{1}^T \mathbf{a}_k = 0$ and $\|\mathbf{a}_k\|_2 = 1$, for $k = 1, \dots, p$. Denote their inner products (*i.e.*, the sample correlation of the corresponding explanatory variables) as $\rho_{ij} = \mathbf{a}_i^T \mathbf{a}_j / (\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2) = \mathbf{a}_i^T \mathbf{a}_j$. Then, the condition in Theorem 2.1 becomes $\|\mathbf{y}\|_2 \sqrt{2 - 2 \rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j)} < \Delta$.*

Proof. The corollary results trivially from inserting the normalization assumption and the definition of ρ_{ij} into the equality $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2 = \sqrt{(\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j)^T (\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j)}$. \square

The following corollary results from observing that, in the OSCAR case, $w_i = \lambda_1 + \lambda_2(p - i)$, thus $\Delta = \lambda_2$.

Corollary 2.2. *In the particular case of OSCAR, and for the case of normalized columns (as in Corollary 2.1), the condition is $\|\mathbf{y}\|_2 \sqrt{2 - 2 \rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j)} < \lambda_2$.*

Corollary 2.2 is closely related to Theorem 2.1 of [4], but has weaker conditions: unlike in [4], our result does not require that both x_i and x_j are different from zero and from all other x_k , for $k \neq i, j$. Furthermore, Theorem 2.1 applies to the more general class of OWL norms, not just to OSCAR. Note also that the results in [4] assume that columns are signed so that $\hat{x}_i \geq 0$ for all i . Our result could also be stated with this assumption, in which case $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2$ simplifies to $\|\mathbf{a}_i - \mathbf{a}_j\|_2$. Finally, observe that, in the extreme case of perfectly correlated columns ($\rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j) = 1$), the condition for OSCAR simplifies to $\lambda_2 > 0$.

2.3 Absolute Error Loss with Similar Columns

Consider the regression problem under absolute error loss in (3). The following theorem shows that, also in this case, the OWL regularizer clusters (in the sense that the corresponding components of the solution are exactly equal in magnitude) the columns that are similar enough.

Theorem 2.2. *Let $\hat{\mathbf{x}}$ be any minimizer of the objective function in (3). Then, for every pair of columns (i, j) for which $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_1 < \Delta$ (where Δ is as defined in Theorem 2.1), we have $|\hat{x}_i| = |\hat{x}_j|$.*

Under the normalization assumptions on matrix \mathbf{A} that were used in Corollary 2.1, another (weaker) sufficient condition can be obtained which depends on the sample correlations, as stated in the following corollary.

Corollary 2.3. *Let $\hat{\mathbf{x}}$ be any minimizer of the objective function in (3) and assume that the columns of \mathbf{A} are normalized, that is, $\mathbf{1}^T \mathbf{a}_k = 0$ and $\|\mathbf{a}_k\|_2 = 1$, for $i = k, \dots, p$. Denote their inner products (i.e., the sample correlation of the corresponding explanatory variables) as $\rho_{ij} = \mathbf{a}_i^T \mathbf{a}_j / (\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2) = \mathbf{a}_i^T \mathbf{a}_j$. Then, for every pair of columns (i, j) for which $\sqrt{n(2 - 2 \rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j))} < \Delta$, we have $|\hat{x}_i| = |\hat{x}_j|$.*

Proof. The corollary results simply from noticing that, under the assumed normalization of the columns of \mathbf{A} , $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_1 \leq \sqrt{n} \|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2 = \sqrt{n(2 - 2 \rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j))}$. \square

Finally, a simple corollary results from the fact that, for OSCAR, $\Delta = \lambda_2$.

Corollary 2.4. *In the particular case of OSCAR, the condition is $\|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_1 < \lambda_2$, in general, and $\sqrt{n(2 - 2 \rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j))} < \lambda_2$, in the case of normalized columns.*

Finally, as above, in the extreme case of perfectly correlated columns ($\rho_{ij} \text{sign}(\hat{x}_i \hat{x}_j) = 1$), the condition for OSCAR simplifies to $\lambda_2 > 0$.

2.4 Proofs of Theorems 2.1 and 2.2

The proofs of both Theorems 2.1 and 2.2 are based on a useful lemma about the OWL norm, which we state and prove before proceeding to the proofs of the theorems.

Lemma 2.1. *Consider a vector $\mathbf{x} \in \mathbb{R}_+^p$ and any two of its components x_i and x_j , such that $x_i > x_j$. Let $\mathbf{z} \in \mathbb{R}_+^p$ be obtained by applying a so-called Pigou-Dalton² transfer of size $\varepsilon \in (0, (x_i - x_j)/2)$ to \mathbf{x} , that is: $z_i = x_i - \varepsilon$, $z_j = x_j + \varepsilon$, and $z_k = x_k$, for $k \neq i, j$. Let \mathbf{w} be a vector of non-increasing non-negative real values, $w_1 \geq w_2 \geq \dots \geq w_p \geq 0$, and Δ be the minimum gap between two consecutive components of vector \mathbf{w} , that is, $\Delta = \min\{w_l - w_{l+1}, l = 1, \dots, p - 1\}$. Then,*

$$\Omega_{\mathbf{w}}(\mathbf{x}) - \Omega_{\mathbf{w}}(\mathbf{z}) \geq \Delta \varepsilon. \quad (10)$$

²The Pigou-Dalton transfer, also known as a Robin Hood transfer, is used in the study of measures of economic inequality [9], [16].

Proof. Let l and m be the rank orders of x_i and x_j , respectively, *i.e.*, $x_i = x_{[l]}$ and $x_j = x_{[m]}$; of course, $m > l$, because $x_i > x_j$. Now let $l+a$ and $m-b$ be the rank orders of z_i and z_j , respectively, *i.e.*, $x_i - \varepsilon = z_i = z_{[l+a]}$ and $x_j + \varepsilon = z_j = z_{[m-b]}$. Of course, it may happen that a or b (or both) are zero, if ε is small enough not to change the rank orders of one (or both) of the affected components of \mathbf{x} . Furthermore, the condition $\varepsilon < (x_i - x_j)/2$ implies that $x_i - \varepsilon > x_j + \varepsilon$, thus $l+a < m-b$. A key observation is that x_{\downarrow} and z_{\downarrow} only differ in positions l to $l+a$ and $m-b$ to m , thus we can thus write

$$\Omega_{\mathbf{w}}(\mathbf{x}) - \Omega_{\mathbf{w}}(\mathbf{z}) = \sum_{k=l}^{l+a} w_k (x_{[k]} - z_{[k]}) + \sum_{k=m-b}^m w_k (x_{[k]} - z_{[k]}). \quad (11)$$

In the range from l to $l+a$, the relationship between z_{\downarrow} and x_{\downarrow} is

$$z_{[l]} = x_{[l+1]}, z_{[l+1]} = x_{[l+2]}, \dots, z_{[l+a-1]} = x_{[l+a]}, z_{[l+a]} = x_{[l]} - \varepsilon, \quad (12)$$

whereas in the range from $m-b$ to m , we have

$$z_{[m-b]} = x_{[m]} + \varepsilon, z_{[m-b+1]} = x_{[m-b]}, \dots, z_{[m]} = x_{[m-1]}. \quad (13)$$

Plugging these equalities into (11) yields

$$\begin{aligned} \Omega_{\mathbf{w}}(\mathbf{x}) - \Omega_{\mathbf{w}}(\mathbf{z}) &= \sum_{k=l}^{l+a-1} w_k (x_{[k]} - x_{[k+1]}) + \sum_{k=m-b+1}^m w_k (x_{[k]} - x_{[k-1]}) \\ &\quad + w_{l+a} (x_{[l+a]} - x_{[l]} + \varepsilon) + w_{m-b} (x_{[m-b]} - x_{[m]} - \varepsilon) \\ &\stackrel{(a)}{\geq} w_{l+a} \sum_{k=l}^{l+a-1} (x_{[k]} - x_{[k+1]}) + w_{m-b} \sum_{k=m-b+1}^m (x_{[k]} - x_{[k-1]}) \\ &\quad + w_{l+a} (x_{[l+a]} - x_{[l]} + \varepsilon) + w_{m-b} (x_{[m-b]} - x_{[m]} - \varepsilon) \\ &= w_{l+a} \left(\sum_{k=l}^{l+a-1} (x_{[k]} - x_{[k+1]}) + (x_{[l+a]} - x_{[l]} + \varepsilon) \right) \\ &\quad + w_{m-b} \left(\sum_{k=m-b+1}^m (x_{[k]} - x_{[k-1]}) + (x_{[m-b]} - x_{[m]} - \varepsilon) \right) \\ &= \varepsilon (w_{l+a} - w_{m-b}) \\ &\geq \varepsilon \Delta, \end{aligned}$$

where inequality (a) results from $x_{[k]} - x_{[k+1]} \geq 0$, $x_{[k]} - x_{[k-1]} \leq 0$, and the components of \mathbf{w} forming a non-increasing sequence. \square

Armed with Lemma 2.1, we now proceed to prove Theorems 2.1 and 2.2.

Proof. (Theorem 2.1) Let us denote $L(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ and take some pair of columns (i, j) . Suppose that $\hat{\mathbf{x}}$ is a minimizer of the objective function in (2), satisfying the condition of the theorem ($\Delta > \|\mathbf{y}\|_2 \|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_2$), but in contradiction to the theorem's claim, *i.e.*, for which $|\hat{x}_i| \neq |\hat{x}_j|$. Without loss of generality, assume that $|\hat{x}_i| > |\hat{x}_j|$, and define the residual vector

$$\mathbf{g} = \mathbf{y} - \sum_{k=1, k \neq i, k \neq j}^p \hat{x}_k \mathbf{a}_k. \quad (14)$$

Now consider a Pigou-Dalton transfer of size $\varepsilon < \min\{|\hat{x}_i|, (|\hat{x}_i| - |\hat{x}_j|)/2\}$ applied to the magnitudes of \hat{x}_i and \hat{x}_j , *i.e.*, take an alternative candidate solution $\mathbf{v} \in \mathbb{R}^p$, such that, $v_i = \text{sign}(\hat{x}_i)(|\hat{x}_i| - \varepsilon)$, $v_j = \text{sign}(\hat{x}_i)(|\hat{x}_i| + \varepsilon)$,

and $v_k = \hat{x}_k$, for $k \neq i, j$. Denoting $\tilde{\mathbf{a}}_i = \text{sign}(\hat{x}_i) \mathbf{a}_i$ and $\tilde{\mathbf{a}}_j = \text{sign}(\hat{x}_j) \mathbf{a}_j$, the difference in loss function that results from this transfer is

$$L(\mathbf{v}) - L(\hat{\mathbf{x}}) = \frac{1}{2} \|\mathbf{g} - (|\hat{x}_i| - \varepsilon) \tilde{\mathbf{a}}_i - (|\hat{x}_j| + \varepsilon) \tilde{\mathbf{a}}_j\|_2^2 - \frac{1}{2} \|\mathbf{g} - |\hat{x}_i| \tilde{\mathbf{a}}_i - |\hat{x}_j| \tilde{\mathbf{a}}_j\|_2^2. \quad (15)$$

Expanding the squared ℓ_2 norms, cancelling out the common $\frac{1}{2} \|\mathbf{g}\|_2^2$ term, and using the common norm of the columns ($\|\mathbf{a}_k\|_2 = c$, for $k = 1, \dots, p$) leads to

$$\begin{aligned} L(\mathbf{v}) - L(\hat{\mathbf{x}}) &= \frac{1}{2} (|\hat{x}_i| - \varepsilon)^2 c^2 + \frac{1}{2} (|\hat{x}_j| + \varepsilon)^2 c^2 - (|\hat{x}_i| - \varepsilon) \mathbf{g}^T \tilde{\mathbf{a}}_i - (|\hat{x}_j| + \varepsilon) \mathbf{g}^T \tilde{\mathbf{a}}_j \\ &\quad + (|\hat{x}_i| - \varepsilon) (|\hat{x}_j| + \varepsilon) \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j - \frac{1}{2} |\hat{x}_i|^2 c^2 - \frac{1}{2} |\hat{x}_j|^2 c^2 \\ &\quad + |\hat{x}_i| \mathbf{g}^T \tilde{\mathbf{a}}_i + |\hat{x}_j| \mathbf{g}^T \tilde{\mathbf{a}}_j - |\hat{x}_i| |\hat{x}_j| \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j. \end{aligned} \quad (16)$$

Expanding the terms $(|\hat{x}_i| - \varepsilon)^2$, $(|\hat{x}_j| + \varepsilon)^2$, and $(|\hat{x}_i| - \varepsilon) (|\hat{x}_j| + \varepsilon)$ and making some further cancellations yields

$$\begin{aligned} L(\mathbf{v}) - L(\hat{\mathbf{x}}) &= \varepsilon \mathbf{g}^T (\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j) + \varepsilon^2 (c^2 - \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j) - \varepsilon c^2 (|\hat{x}_i| - |\hat{x}_j|) + \varepsilon (|\hat{x}_i| - |\hat{x}_j|) \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j \\ &= \varepsilon \mathbf{g}^T (\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j) - \varepsilon (c^2 - \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j) (|\hat{x}_i| - |\hat{x}_j| - \varepsilon) \\ &\stackrel{(a)}{\leq} \varepsilon \mathbf{g}^T (\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j) \\ &\stackrel{(b)}{\leq} \varepsilon \|\mathbf{y}\|_2 \|\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j\|_2, \end{aligned}$$

where inequality (a) results from the facts that (by the Cauchy-Schwartz inequality) $c^2 \geq \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j$ and both ε and $(|\hat{x}_i| - |\hat{x}_j| - \varepsilon)$ are (by assumption) positive, whereas (b) is again Cauchy-Schwartz together with the fact that $\|\mathbf{g}\|_2 \leq \|\mathbf{y}\|_2$. Finally, since $|\mathbf{v}| \in \mathbb{R}_+^p$ results from the same Pigou-Dalton transfer of size ε applied to $|\hat{\mathbf{x}}| \in \mathbb{R}_+^p$, and $\Omega_{\mathbf{w}}$ only depends on the absolute values of its arguments, we are in condition to invoke Lemma 2.1, which yields

$$L(\mathbf{v}) + \Omega_{\mathbf{w}}(\mathbf{v}) - (L(\hat{\mathbf{x}}) + \Omega_{\mathbf{w}}(\hat{\mathbf{x}})) \leq \varepsilon (\|\mathbf{y}\|_2 \|\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j\|_2 - \Delta) < 0, \quad (17)$$

contradicting the assumption that $\hat{\mathbf{x}}$ is a minimizer of $L(\mathbf{x}) + \Omega_{\mathbf{w}}(\mathbf{x})$, thus completing the proof. \square

Proof. (Theorem 2.2) Denote $G(\mathbf{x}) = \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_1$ and take some pair of columns (i, j) . Assume that $\hat{\mathbf{x}}$ is a minimizer of the objective function in (3), satisfying the condition of the theorem ($\Delta > \|\text{sign}(\hat{x}_i) \mathbf{a}_i - \text{sign}(\hat{x}_j) \mathbf{a}_j\|_1$, but in contradiction to the theorem's claim, *i.e.*, for which $|\hat{x}_i| \neq |\hat{x}_j|$). Define the residual vector \mathbf{g} as in (14) and consider a Pigou-Dalton transfer of size $\varepsilon < \min\{|\hat{x}_i|, (|\hat{x}_i| - |\hat{x}_j|)/2\}$ applied to the magnitudes of \hat{x}_i and \hat{x}_j , *i.e.*, take an alternative candidate solution $\mathbf{v} \in \mathbb{R}^p$, such that, $v_i = \text{sign}(\hat{x}_i)(|\hat{x}_i| - \varepsilon)$, $v_j = \text{sign}(\hat{x}_i)(|\hat{x}_i| + \varepsilon)$, and $v_k = \hat{x}_k$, for $k \neq i, j$. Denoting $\tilde{\mathbf{a}}_i = \text{sign}(\hat{x}_i) \mathbf{a}_i$ and $\tilde{\mathbf{a}}_j = \text{sign}(\hat{x}_j) \mathbf{a}_j$, the difference in loss function that results from this transfer satisfies

$$\begin{aligned} G(\mathbf{v}) - G(\hat{\mathbf{x}}) &= \|\mathbf{g} - (|\hat{x}_i| - \varepsilon) \tilde{\mathbf{a}}_i - (|\hat{x}_i| + \varepsilon) \tilde{\mathbf{a}}_j\|_1 - \|\mathbf{g} - |\hat{x}_i| \tilde{\mathbf{a}}_i - |\hat{x}_j| \tilde{\mathbf{a}}_j\|_1 \\ &= \|\mathbf{g} - |\hat{x}_i| \tilde{\mathbf{a}}_i - |\hat{x}_i| \tilde{\mathbf{a}}_j + \varepsilon (\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j)\|_1 - \|\mathbf{g} - |\hat{x}_i| \tilde{\mathbf{a}}_i - |\hat{x}_j| \tilde{\mathbf{a}}_j\|_1 \\ &\leq \varepsilon \|\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j\|_1, \end{aligned} \quad (18)$$

as a direct consequence of the triangle inequality. Finally, since $|\mathbf{v}| \in \mathbb{R}_+^p$ results from the same Pigou-Dalton transfer of size ε applied to $|\hat{\mathbf{x}}| \in \mathbb{R}_+^p$, and $\Omega_{\mathbf{w}}$ only depends on the absolute values of its arguments, we are in condition to invoke Lemma 2.1, thus

$$G(\mathbf{v}) + \Omega_{\mathbf{w}}(\mathbf{v}) - (G(\hat{\mathbf{x}}) + \Omega_{\mathbf{w}}(\hat{\mathbf{x}})) \leq \varepsilon (\|\tilde{\mathbf{a}}_i - \tilde{\mathbf{a}}_j\|_1 - \Delta) < 0, \quad (19)$$

which contradicts the assumption that $\hat{\mathbf{x}}$ is a minimizer of $G(\mathbf{x}) + \Omega_{\mathbf{w}}(\mathbf{x})$, thus completing the proof. \square

2.5 Proof of Theorem 1.1 (i)

The proof of item (i) in Theorem 1.1 follows the same general structure as the proofs of Theorems 2.1 and 2.2.

Proof. (Theorem 1.1 (i)) Let us define the functions $L(\mathbf{x}) = \frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_2^2$ and $G(\mathbf{x}) = \frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_1$, and the residual \mathbf{g} as in (14). Since $\mathbf{a}_i = \mathbf{a}_j$, the functions

$$L(\hat{\mathbf{x}}) = \frac{1}{n} \|\mathbf{g} - (\hat{x}_i + \hat{x}_j) \mathbf{a}_i\|_2^2 \quad \text{and} \quad G(\hat{\mathbf{x}}) = \frac{1}{n} \|\mathbf{g} - (\hat{x}_i + \hat{x}_j) \mathbf{a}_i\|_1, \quad (20)$$

are both invariant under a transformation that adds any quantity ε to \hat{x}_i and subtracts the same quantity from \hat{x}_j .

We first prove that if $\mathbf{a}_i = \mathbf{a}_j$, then $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$. Assume, by contradiction, that $\text{sign}(\hat{x}_i) \neq \text{sign}(\hat{x}_j)$, and, without loss of generality, that $\hat{x}_i > 0$. We need to consider two cases:

- a) if $\hat{x}_j < 0$, take an alternative feasible solution \mathbf{v} , with: $v_k = \hat{x}_k$, for $k \neq i, j$, $v_i = \hat{x}_i - \varepsilon$, and $v_j = \hat{x}_j + \varepsilon$, for some $\varepsilon \in (0, \min\{|\hat{x}_i|, |\hat{x}_j|\})$. Since $\hat{x}_i > 0$ and $\hat{x}_j < 0$, it's true that $|v_i| = |\hat{x}_i| - \varepsilon$ and $|v_j| = |\hat{x}_j| - \varepsilon$. Finally, the definition of Δ implies that $w_{p-1} \geq \Delta > 0$, thus $\Omega_{\mathbf{w}}(\hat{\mathbf{x}}) - \Omega_{\mathbf{w}}(\mathbf{v}) > \varepsilon \Delta > 0$, contradicting the optimality of $\hat{\mathbf{x}}$, thus proving the claim that $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$.
- b) if $\hat{x}_j = 0$, consider a feasible \mathbf{v} resulting from a Pigou-Dalton transfer of size $\varepsilon \in (0, |\hat{x}_i|/2]$. From Lemma 2.1, $\Omega_{\mathbf{w}}(\hat{\mathbf{x}}) - \Omega_{\mathbf{w}}(\mathbf{v}) > \varepsilon \Delta > 0$, negating the optimality of $\hat{\mathbf{x}}$, thus proving that $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$.

Once it is established that $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$, we proceed to prove that $|\hat{x}_i| = |\hat{x}_j|$. To this end, notice that

$$L(\hat{\mathbf{x}}) = \frac{1}{n} \|\mathbf{g} - (|\hat{x}_i| + |\hat{x}_j|) \text{sign}(\hat{x}_i) \mathbf{a}_i\|_2^2 \quad \text{and} \quad G(\hat{\mathbf{x}}) = \frac{1}{n} \|\mathbf{g} - (|\hat{x}_i| + |\hat{x}_j|) \text{sign}(\hat{x}_i) \mathbf{a}_i\|_1. \quad (21)$$

Proceeding again by contradiction, suppose (without loss of generality) that $|\hat{x}_i| > |\hat{x}_j|$, and define \mathbf{u} via a Pigou-Dalton transfer on the magnitudes of \mathbf{x}_i and \mathbf{x}_j , that is: $u_k = \hat{x}_k$, for $k \neq i, j$, $u_i = (|\hat{x}_i| - \delta) \text{sign}(\hat{x}_i)$, and $u_j = (|\hat{x}_j| + \delta) \text{sign}(\hat{x}_i)$, for some $\delta \in (0, \min\{|\hat{x}_i|, (|\hat{x}_i| - |\hat{x}_j|)/2\})$. Of course, \mathbf{u} is feasible and Lemma 2.1 shows that $\Omega_{\mathbf{w}}(\hat{\mathbf{x}}) - \Omega_{\mathbf{w}}(\mathbf{u}) > \delta \Delta > 0$, contradicting the optimality of $\hat{\mathbf{x}}$, thus concluding the proof. \square

3 Statistical Analysis of OWL Regularization

3.1 Introduction

In this section, we characterize the statistical performance of OWL regularization with both the squared and absolute error losses, by proving finite sample bounds, which apply to the standard LASSO and OSCAR as special cases. At the basis of our approach is the following model for correlated measurement matrices. Recall that \mathbf{A} has dimensions $n \times p$. Assume that the rows of \mathbf{A} are independently and identically distributed $\mathcal{N}(\mathbf{0}, \mathbf{C}^T \mathbf{C})$, the multivariate Gaussian distribution with covariance $\mathbf{C}^T \mathbf{C}$ (i.e., the columns of \mathbf{A} are not necessarily independent). Assume that the matrix \mathbf{C} is $q \times p$ with $q \geq n$. Note that \mathbf{A} can be factorized as $\mathbf{A} = \mathbf{B} \mathbf{C}$, where \mathbf{B} is an $n \times q$ Gaussian random matrix, whose entries are i.i.d. $\mathcal{N}(0, 1)$ random variables. The role of matrix \mathbf{C} is to mix, or even replicate, columns of \mathbf{B} . The next simple example illustrates this construction.

Example 1. Suppose $q = 3$, $p = 4$, and

$$\mathbf{C} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}; \quad (22)$$

then, if $\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3]$, where $\mathbf{b}_i \in \mathbb{R}^n$ is the i -th column of \mathbf{B} , matrix \mathbf{A} has the form $\mathbf{A} = [\mathbf{b}_1, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3]$.

Assume that we observe

$$\mathbf{y} = \mathbf{A} \mathbf{x}^* + \boldsymbol{\nu}, \quad (23)$$

where $\boldsymbol{\nu} \in \mathbb{R}^n$ is the measurement error satisfying

$$\frac{1}{n} \|\boldsymbol{\nu}\|_1 \leq \varepsilon, \quad (24)$$

and about which we make no other assumptions. The signal $\mathbf{x}^* \in \mathbb{R}^p$ is assumed to satisfy $\|\mathbf{x}^*\|_1 \leq \sqrt{s}$. Note, for example, that this condition is met if $\|\mathbf{x}^*\|_2 \leq 1$ and \mathbf{x}^* has at most s non-zero components.

Example 2. To illustrate the model, consider the matrix \mathbf{C} defined in (22) and suppose that $\mathbf{A} \mathbf{x}^* = \mathbf{b}_1$. There are many $\mathbf{x} \in \mathbb{R}^4$ satisfying $\mathbf{A} \mathbf{x} = \mathbf{b}_1$, including the vectors $[1, 0, 0, 0]^T$, $[0, 1, 0, 0]^T$, and all convex combinations of the two. LASSO regularization (i.e., the ℓ_1 norm) does not differentiate these equivalent representations, since $\|[\alpha, (1 - \alpha), 0, 0]^T\|_1 = 1$, for any $\alpha \in [0, 1]$. However the OWL norm (as claimed in Theorem 1.1) prefers the solution $[\frac{1}{2}, \frac{1}{2}, 0, 0]^T$, selecting both colinear columns of \mathbf{A} for the representation.

3.2 Main Result and Corollaries

The main result of this section is stated in the following theorem, whose proof is given in the following subsections, based on the techniques developed by Vershynin [21]. The results are stated in terms of constrained optimization, which, under certain conditions, is equivalent to the Lagrangian formulation studied in Section 2. We also present a corollary for the particular case where \mathbf{C} simply replicates columns of \mathbf{B} , that is, when \mathbf{A} includes groups of identical columns; this corollary is shown to imply part (ii) of Theorem 1.1. Expectations are with respect to the Gaussian distribution of \mathbf{A} .

Theorem 3.1. Let \mathbf{y} , \mathbf{A} , \mathbf{x}^* , and ε be as defined above, and let $\hat{\mathbf{x}}$ be a solution to one of the two following optimization problems:

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_w(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2, \quad (25)$$

or

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_w(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon. \quad (26)$$

Then

$$\mathbb{E} \sqrt{(\hat{\mathbf{x}} - \mathbf{x}^*)^T \mathbf{C}^T \mathbf{C} (\hat{\mathbf{x}} - \mathbf{x}^*)} \leq \sqrt{2\pi} \left(4\sqrt{2} \|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{\frac{s \log q}{n}} + \varepsilon \right), \quad (27)$$

where $\bar{w} = p^{-1} \sum_{i=1}^p w_i$ and $\|\mathbf{C}\|_1$ is the matrix norm induced by the ℓ_1 norm: $\|\mathbf{C}\|_1 = \max_j \|\mathbf{c}_j\|_1$, with \mathbf{c}_j denoting the j -th column of \mathbf{C} .

Note that the error bound in (27) holds for optimizations based on the squared ℓ_2 and ℓ_1 losses. In fact, since $\frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2$ implies $\frac{1}{n} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon$, the ℓ_1 constraint is less restrictive. In both cases, the theorem shows that the number of samples sufficient to estimate an s -sparse signal with a given precision grows like

$$n \sim s \log q.$$

This agrees with well-known sample complexity bounds for sparse recovery under stronger assumptions such as the restricted isometry property or i.i.d. measurements [6, 7, 10, 12, 21]. However, note that the error measure of the theorem is insensitive to components of $\hat{\mathbf{x}}$ in the nullspace of \mathbf{C} , which is to be expected since in general

there may be many sparse \mathbf{x} that yield the same value of $\mathbf{A}\mathbf{x}$ (see Example 2). This is where the OWL norm becomes especially important. With strictly decreasing weights ($\Delta > 0$), OWL prefers solutions that select all colinear columns associated with the model. In other words, if the columns are colinear (or strongly correlated, per the characterizations of given in Section 2), then the OWL solution will select a representation including all the columns associated with the sparse model, rather than an arbitrary subset of them.

The OSCAR norm is a special case of OWL, with $w_i = \lambda_1 + \lambda_2(p - i)$ and $\lambda_1, \lambda_2 > 0$. In this case, $\bar{w} = \lambda_1 + \lambda_2(p - 1)/2$ and therefore $w_1/\bar{w} \leq 2$. Note that the conventional ℓ_1 norm (used in the LASSO) is the special case of OWL with uniform weights (or OSCAR with $\lambda_2 = 0$); thus, all our results apply to ℓ_1 minimization as well, in which case $w_1/\bar{w} = 1$.

To illustrate Theorem 3.1, let \mathbf{C} be an $q \times p$ matrix that replicates each column of \mathbf{B} one or more times. Note that each column of \mathbf{C} is 1-sparse and has unit ℓ_1 norm, thus $\|\mathbf{C}\|_1 = 1$. Let G_1, \dots, G_q denote the groups of replicated columns in $\mathbf{A} = \mathbf{B}\mathbf{C}$; these groups are a partition of the set $\{1, \dots, p\}$. Example 1 is a special case of this scenario, with $G_1 = \{1, 2\}$, $G_2 = \{3\}$, and $G_3 = \{4\}$. Assume that there are s non-zero components in \mathbf{x}^* , each in one of s distinct groups. Let \mathbf{x}_G denote the vector that is zero except on the the subset of entries in $G \subset \{1, \dots, p\}$, where it takes the same values as \mathbf{x} . Then note that for any $\hat{\mathbf{x}}$ we have

$$(\hat{\mathbf{x}} - \mathbf{x}^*)^T \mathbf{C}^T \mathbf{C} (\hat{\mathbf{x}} - \mathbf{x}^*) = \sum_{i=1}^q |\mathbf{1}^T (\hat{\mathbf{x}}_{G_i} - \mathbf{x}_{G_i}^*)|^2, \quad (28)$$

where $\mathbf{1} = [1 \ 1 \ \dots \ 1]^T$. This produces the following corollary to Theorem 3.1.

Corollary 3.1. *Assume that each column of \mathbf{C} is 1-sparse and unit norm. Let $\hat{\mathbf{x}}$ be a solution to the optimization*

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_w(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2, \quad (29)$$

or

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_w(\mathbf{x}) \quad \text{subject to} \quad \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon. \quad (30)$$

Then

$$\mathbb{E} \sqrt{\sum_{i=1}^q |\mathbf{1}^T (\hat{\mathbf{x}}_{G_i} - \mathbf{x}_{G_i}^*)|^2} \leq \sqrt{2\pi} \left(4\sqrt{2} \frac{w_1}{\bar{w}} \sqrt{\frac{s \log q}{n}} + \varepsilon \right). \quad (31)$$

In this case, since the correlated columns are colinear, the OWL norm will select all or none of the columns in each group and each $\hat{\mathbf{x}}_{G_i}$ will have identical non-zero values (if any). If we let $z_i^* = \mathbf{1}^T \mathbf{x}_{G_i}^*$ and $\hat{z}_i = \mathbf{1}^T \hat{\mathbf{x}}_{G_i}$ for $i = 1, \dots, q$, and let $\mathbf{z}^* = [z_1^*, \dots, z_q^*]^T$ and $\hat{\mathbf{z}} = [\hat{z}_1, \dots, \hat{z}_q]^T$; then (9) can be expressed as

$$\mathbb{E} \|\hat{\mathbf{z}} - \mathbf{z}^*\|_2 \leq \sqrt{2\pi} \left(4\sqrt{2} \frac{w_1}{\bar{w}} \sqrt{\frac{s \log q}{n}} + \varepsilon \right), \quad (32)$$

which is the type of result obtained in the compressed sensing literature for the ideal sparse observation model $\mathbf{y} = \mathbf{B}\mathbf{z}^* + \boldsymbol{\nu}$ (based on an i.i.d. observation model) [21]. This shows that *by using OWL we pay no price for colinearity in \mathbf{A}* . Also, as shown next, Corollary 3.1 implies claim (ii) in Theorem 1.1 in Section 1.

Proof. (Theorem 1.1 (ii)) Taking into account the group structure of $\hat{\mathbf{x}}$ and \mathbf{x}^* , we have that

$$\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2 = \sqrt{\sum_{i=1}^q \|\hat{\mathbf{x}}_{G_i} - \mathbf{x}_{G_i}^*\|_2^2} = \sqrt{\sum_{i=1}^q \frac{1}{|G_i|} (\hat{z}_i - z_i^*)^2} \leq \|\hat{\mathbf{z}} - \mathbf{z}^*\|_2;$$

this inequality, together with (32) and the fact that $q \leq p$ yields (9). \square

Before moving on to the proof of Theorem 3.1, notice that since the ℓ_1 norm is a special case of OWL with uniform weights, the same bounds in Theorem 3.1 and Corollary 3.1 hold for ℓ_1 minimization. The difference is that the LASSO solution generally will not select all correlated or colinear columns selected by OWL, making the estimated model less interpretable.

3.3 Proof of Theorem 3.1

The proof of Theorem 3.1 is based on the approach developed by Vershynin [21]. The key ingredient is the so-called *general M^* bound* (Theorem 5.1 in [21]), which applies to the special case when \mathbf{A} is a i.i.d. Gaussian matrix (i.e., when \mathbf{C} is identity in our set-up). We extend the bound to cover our model $\mathbf{A} = \mathbf{BC}$, for general \mathbf{C} . Recall that \mathbf{A} is $n \times p$, \mathbf{B} is $n \times q$, with $n \leq q$, and \mathbf{C} is $q \times p$.

3.3.1 Extended General M^* Bound

Theorem 3.2. (Extended general M^* bound). *Let T be a bounded subset of \mathbb{R}^p . Let \mathbf{B} be an $n \times q$ Gaussian random matrix (with i.i.d. $\mathcal{N}(0, 1)$ entries) and let $\mathbf{A} = \mathbf{BC}$, where \mathbf{C} is a deterministic $q \times p$ matrix. Fix $\varepsilon \geq 0$ and consider the set*

$$T_\varepsilon := \left\{ \mathbf{u} \in T : \frac{1}{n} \|\mathbf{A}\mathbf{u}\|_1 \leq \varepsilon \right\}. \quad (33)$$

Then

$$\mathbb{E} \sup_{\mathbf{u} \in T_\varepsilon} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \leq \sqrt{\frac{2\pi}{n}} \mathbb{E} \sup_{\mathbf{u} \in T} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| + \sqrt{\frac{\pi}{2}} \varepsilon, \quad (34)$$

where $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_q)$ is a standard Gaussian random vector in \mathbb{R}^q .

The proof follows in a straightforward fashion from the proof of Theorem 5.1 in [21], with modifications made to account for \mathbf{C} . For the sake of completeness we include a proof here.

Proof. The bound (34) follows from the deviation inequality

$$\mathbb{E} \sup_{\mathbf{u} \in T} \left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{a}_i, \mathbf{u} \rangle| - \sqrt{\frac{2}{\pi}} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \right| \leq \frac{2}{\sqrt{n}} \mathbb{E} \sup_{\mathbf{u} \in T} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|, \quad (35)$$

where \mathbf{a}_i denotes the i th row of \mathbf{A} . To see this, note that the inequality holds if we replace T by the smaller set T_ε . For $\mathbf{u} \in T_\varepsilon$, and for such \mathbf{u} we have by assumption that $\frac{1}{n} \sum_{i=1}^n |\langle \mathbf{a}_i, \mathbf{u} \rangle| = \frac{1}{n} \|\mathbf{A}\mathbf{u}\|_1 \leq \varepsilon$, and the bound (34) follows by the triangle inequality.

To prove (35), the first thing to note is that

$$\mathbb{E} |\langle \mathbf{a}_i, \mathbf{u} \rangle| = \mathbb{E} |\langle \mathbf{C}^T \mathbf{b}_i, \mathbf{u} \rangle| = \mathbb{E} |\langle \mathbf{b}_i, \mathbf{C}\mathbf{u} \rangle|, \quad (36)$$

where \mathbf{b}_i is the i th row of \mathbf{B} . Because the Gaussian distribution of \mathbf{b}_i is rotationally invariant, it follows that

$$\mathbb{E} |\langle \mathbf{b}_i, \mathbf{C}\mathbf{u} \rangle| = \sqrt{\frac{2}{\pi}} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2}. \quad (37)$$

Using the symmetrization and contraction inequalities from Proposition 5.2 in [21], we have the bound

$$\mathbb{E} \sup_{\mathbf{u} \in T} \left| \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{a}_i, \mathbf{u} \rangle| - \sqrt{\frac{2}{\pi}} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \right| \leq 2 \mathbb{E} \sup_{\mathbf{u} \in T} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i \langle \mathbf{b}_i, \mathbf{C}\mathbf{u} \rangle \right| \quad (38)$$

$$= 2 \mathbb{E} \sup_{\mathbf{u} \in T} \left| \left\langle \frac{1}{n} \sum_{i=1}^n \varepsilon_i \mathbf{b}_i, \mathbf{C}\mathbf{u} \right\rangle \right|, \quad (39)$$

where each ε_i independently takes values -1 and $+1$ with probabilities $1/2$. Note that vector $\mathbf{g} := \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i \mathbf{b}_i \sim \mathcal{N}(0, \mathbf{I}_n)$, thus,

$$2 \mathbb{E} \sup_{\mathbf{u} \in T} \left| \left\langle \frac{1}{n} \sum_{i=1}^n \varepsilon_i \mathbf{b}_i, \mathbf{C} \mathbf{u} \right\rangle \right| = \frac{2}{\sqrt{n}} \mathbb{E} \sup_{\mathbf{u} \in T} |\langle \mathbf{g}, \mathbf{C} \mathbf{u} \rangle| = \frac{2}{\sqrt{n}} \mathbb{E} \sup_{\mathbf{u} \in T} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|. \quad (40)$$

This completes the proof. □

3.3.2 Estimation from Noisy Linear Observations

Theorem 3.2 can be used to derive error bounds for estimating signals known to belong to a certain subset (sparse sets are a special case we will consider in the next section). Let $\mathcal{K} \subset \mathbb{R}^p$ be given. Suppose that we observe

$$\mathbf{y} = \mathbf{A} \mathbf{x}^* + \boldsymbol{\nu}, \quad \frac{1}{n} \|\boldsymbol{\nu}\|_1 \leq \varepsilon, \quad (41)$$

where $\mathbf{x}^* \in \mathcal{K}$. The following theorems are straightforward extensions of Theorems 6.1 and 6.2 in [21]. We include the proofs for the sake of completeness.

Theorem 3.3. (Estimation from noisy linear observations: feasibility program). *Choose $\hat{\mathbf{x}}$ to be any vector satisfying*

$$\hat{\mathbf{x}} \in \mathcal{K} \text{ and } \frac{1}{n} \|\mathbf{A} \hat{\mathbf{x}} - \mathbf{y}\|_1 \leq \varepsilon. \quad (42)$$

Then

$$\mathbb{E} \sup_{\mathbf{x}^* \in \mathcal{K}} \{(\hat{\mathbf{x}} - \mathbf{x}^*)^T \mathbf{C}^T \mathbf{C} (\hat{\mathbf{x}} - \mathbf{x}^*)\}^{1/2} \leq \sqrt{2\pi} \left(\frac{\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|}{\sqrt{n}} + \varepsilon \right). \quad (43)$$

Proof. We apply Theorem 3.2 to the set $T = \mathcal{K} - \mathcal{K}$ with 2ε instead of ε , which yields

$$\mathbb{E} \sup_{\mathbf{u} \in T_{2\varepsilon}} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \leq \sqrt{\frac{2\pi}{n}} \mathbb{E} \sup_{\mathbf{u} \in T} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| + \sqrt{2\pi} \varepsilon,$$

From here, all we need to show is that for any $\mathbf{x}^* \in \mathcal{K}$

$$\hat{\mathbf{x}} - \mathbf{x}^* \in T_{2\varepsilon}. \quad (44)$$

To see this, note that $\hat{\mathbf{x}}, \mathbf{x}^* \in \mathcal{K}$, so $\hat{\mathbf{x}} - \mathbf{x}^* \in \mathcal{K} - \mathcal{K} = T$. By the triangle inequality,

$$\frac{1}{n} \|\mathbf{A}(\hat{\mathbf{x}} - \mathbf{x}^*)\|_1 = \frac{1}{n} \|\mathbf{A} \hat{\mathbf{x}} - \mathbf{y} + \boldsymbol{\nu}\|_1 \leq \frac{1}{n} \|\mathbf{A} \hat{\mathbf{x}} - \mathbf{y}\|_1 + \frac{1}{n} \|\boldsymbol{\nu}\|_1 \leq 2\varepsilon,$$

showing that $\mathbf{u} = \hat{\mathbf{x}} - \mathbf{x}^*$ indeed satisfies the constraints that define $T_{2\varepsilon}$ in (33). □

Next we derive an optimization program for the solution. The *Minkowski functional* of \mathcal{K} is defined as

$$\|\mathbf{x}\|_{\mathcal{K}} = \inf\{\lambda > 0 : \lambda^{-1} \mathbf{x} \in \mathcal{K}\}.$$

If \mathcal{K} is a compact and origin-symmetric convex set with non-empty interior, then $\|\mathbf{x}\|_{\mathcal{K}}$ is a norm on \mathbb{R}^p [18]. Note that $\mathbf{x} \in \mathcal{K}$ if and only if $\|\mathbf{x}\|_{\mathcal{K}} \leq 1$.

Theorem 3.4. (Estimation from noisy linear observations: optimization program). *Choose $\hat{\mathbf{x}}$ to be a solution to the optimization*

$$\min \|\mathbf{x}\|_{\mathcal{K}} \text{ subject to } \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon. \quad (45)$$

Then

$$\mathbb{E} \sup_{\mathbf{x}^* \in \mathcal{K}} \{(\hat{\mathbf{x}} - \mathbf{x}^*)^T \mathbf{C}^T \mathbf{C} (\hat{\mathbf{x}} - \mathbf{x}^*)\}^{1/2} \leq \sqrt{2\pi} \left(\frac{\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|}{\sqrt{n}} + \varepsilon \right). \quad (46)$$

Proof. If we show that $\hat{\mathbf{x}} \in \mathcal{K}$, then the result follows from Theorem 3.3. Note that the constraint of the program guarantees that $\frac{1}{n} \|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_1 \leq \varepsilon$ and by assumption we have $\frac{1}{n} \|\mathbf{A}\mathbf{x}^* - \mathbf{y}\|_1 = \frac{1}{n} \|\boldsymbol{\nu}\|_1 \leq \varepsilon$. Thus we have

$$\|\hat{\mathbf{x}}\|_{\mathcal{K}} \leq \|\mathbf{x}^*\|_{\mathcal{K}} \leq 1,$$

since $\mathbf{x}^* \in \mathcal{K}$. The inequality $\|\hat{\mathbf{x}}\|_{\mathcal{K}} \leq 1$ implies that $\hat{\mathbf{x}} \in \mathcal{K}$. \square

3.3.3 Sparse Recovery via OWL

Recall the definition of the OWL norm

$$\Omega_{\mathbf{w}}(\mathbf{x}) = \sum_{i=1}^p w_i |x|_{[i]},$$

where $|x|_{[1]}, \dots, |x|_{[p]}$ are the magnitudes of the elements of \mathbf{x} in decreasing order and $w_1 \geq w_2 \geq \dots \geq w_p$ is a non-increasing sequence of weights. The OWL norm satisfies

$$\bar{w} \|\mathbf{x}\|_1 \leq \Omega_{\mathbf{w}}(\mathbf{x}) \leq w_1 \|\mathbf{x}\|_1,$$

where $\bar{w} := \frac{1}{N} \sum_{i=1}^N w_i$. This is easily verified by minimizing or maximizing the OWL norm subject to a fixed ℓ_1 norm. We now prove Theorem 3.1.

Proof. (Theorem 3.1) Since the signal generating the measurements is assumed to satisfy $\|\mathbf{x}^*\|_1 \leq \sqrt{s}$, we first need to construct an OWL ball that contains all $\mathbf{x} \in \mathbb{R}^p$ with $\|\mathbf{x}\|_1 \leq \sqrt{s}$. Let $\mathcal{K} = \{\mathbf{x} \in \mathbb{R}^p : \Omega_{\mathbf{w}}(\mathbf{x}) \leq w_1 \sqrt{s}\}$. Because $\Omega_{\mathbf{w}}(\mathbf{x}) \leq w_1 \|\mathbf{x}\|_1$, all vectors satisfying $\|\mathbf{x}\|_1 \leq \sqrt{s}$ belong to \mathcal{K} . Also note that because $\Omega_{\mathbf{w}}(\mathbf{x})$ is a norm, and \mathcal{K} is a ball of this norm, the Minkowski functional $\|\mathbf{x}\|_{\mathcal{K}}$ is proportional to $\Omega_{\mathbf{w}}(\mathbf{x})$.

The quantity $\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|$ in (46), called the width of \mathcal{K} , satisfies

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| = \mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle|.$$

Note that

$$\|\mathbf{C}\mathbf{u}\|_1 \leq \|\mathbf{C}\|_1 \|\mathbf{u}\|_1 \leq \|\mathbf{C}\|_1 \frac{1}{\bar{w}} \Omega_{\mathbf{w}}(\mathbf{u}).$$

The triangle inequality and the definition of \mathcal{K} imply that for any $\mathbf{u} \in \mathcal{K} - \mathcal{K}$, $\Omega_{\mathbf{w}}(\mathbf{u}) \leq 2w_1 \sqrt{s}$, thus we have

$$\|\mathbf{C}\mathbf{u}\|_1 \leq 2 \|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{s}.$$

The width can be then bounded as

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle| \leq \mathbb{E} \sup_{\{\mathbf{v} : \|\mathbf{v}\|_1 \leq 2 \|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{s}\}} |\langle \mathbf{g}, \mathbf{v} \rangle|.$$

The \mathbf{v} that maximizes the right hand side places mass $2\|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{s}$ on the largest element of \mathbf{g} (in magnitude) and zero on every other element. This yields the bound

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K}-\mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle| \leq 2\|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{s} \mathbb{E} \max_{i=1,\dots,q} |g_i|. \quad (47)$$

Using Jensen’s inequality, the square of the expectation in (47) can be bounded as

$$\left(\mathbb{E} \max_{i=1,\dots,q} |g_i| \right)^2 \leq \mathbb{E} \max_{i=1,\dots,q} |g_i|^2 \leq \left(\sqrt{2 \log q} + 1 \right)^2,$$

where the second inequality comes from a chi-square tail bound (see Lemma 3.2 in [17]). Note that since $q > 1$, $\sqrt{2 \log q} + 1 < 2\sqrt{2 \log q}$. Putting everything together, we obtain the bound

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K}-\mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle| \leq 4\sqrt{2} \|\mathbf{C}\|_1 \frac{w_1}{\bar{w}} \sqrt{s \log q}$$

Theorem 3.1 now follows immediately from Theorem 3.4, above. \square

4 Conclusion

In this paper, we have studied *ordered weighted ℓ_1* (OWL) regularization for sparse estimation problems with strongly correlated variables. We have proved sufficient conditions under which the OWL regularizer clusters the coefficient estimates, based on the correlation/collinearity of the variables in the design matrix. We have also characterized the statistical performance of OWL regularization for generative models in which certain clusters of regression variables are strongly (even perfectly) correlated, but variables in different clusters are uncorrelated. Essentially, we showed that, by using OWL regularization, we pay no price (in terms of the number of measurements) for the presence of strongly correlated variables. Future work will include the experimental evaluation of OWL regularization and its application to other problems, such as logistic regression.

References

- [1] H. Bauschke and P. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*, Springer, 2011.
- [2] M. Bogdan, E. van den Berg, W. Su, and E. Candès, “Statistical estimation and testing via the sorted ℓ_1 norm,” available at <http://arxiv.org/abs/1310.1969>, 2013.
- [3] M. Bogdan, E. van den Berg, C. Sabatti, W. Su, and E. Candès, “SLOPE – adaptive variable selection via convex optimization”, available at <http://arxiv.org/abs/1407.3824>, 2014.
- [4] H. Bondell and B. Reich, “Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR,” *Biometrics*, vol. 64, pp. 115–123, 2007.
- [5] P. Bühlmann, P. Rüttimean, S. van de Geer, and C.-H. Zhang, “Correlated variables in regression: Clustering and sparse estimation,” *Journal of Statistical Planning and Inference*, pp. 1835–1858, 2013.
- [6] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, pp. 489–509, 2006.
- [7] E. Candès and T. Tao, “The Dantzig selector: Statistical estimation when p is much larger than n ,” *Annals of Statistics*, vol. 35, pp. 2313–2351, 2007.

- [8] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational Mathematics*, vol. 12, pp. 805–849, 2012.
- [9] H. Dalton, “The measurement of the inequality of incomes,” *The Economic Journal*, vol. 30, pp. 348–361, 1920.
- [10] D. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, 2006.
- [11] C. Genovese, J. Jin, L. Wasserman, and Z. Yao, “A Comparison of the LASSO and marginal regression”, *Journal of Machine Learning Research*, pp. 2107–2143, 2012.
- [12] J. Haupt and R. Nowak, “Signal reconstruction from noisy random projections,” *IEEE Transactions on Information Theory*, vol. 52, pp. 4036–4048, 2006.
- [13] M. Jaggi, “Revisiting Frank-Wolfe: projection-free sparse convex optimization”, *Proceedings of the 30th International Conference on Machine Learning*, pp. 427–435, 2013.
- [14] J. Jia and B. Yu, “On model selection consistency of the elastic net when $p \gg n$,” *Statistica Sinica*, vol. 20, pp. 595–611, 2010.
- [15] N. Meinshausen and P. Bühlmann, “Stability selection,” *Journal of the Royal Statistical Society (B)*, pp. 417–473, 2010.
- [16] A. Pigou, *Wealth and Welfare*, Macmillan, London, 1912.
- [17] N. Rao, B. Recht, and R. Nowak, “Tight Measurement Bounds for Exact Recovery of Structured Sparse Signals,” *Proceedings of AISTATS*, 2012.
- [18] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.
- [19] R. Shah and R. Samworth, “Variable selection with error control: another look at stability selection.” *Journal of the Royal Statistical Society (B)*, pp. 55–80, 2013.
- [20] X. Shen and H.-C. Huang, “Grouping pursuit through a regularization solution surface,” *Journal of the American Statistical Association*, pp. 727–739, 2010.
- [21] R. Vershynin, “Estimation in High Dimensions: A Geometric Perspective,” available at <http://arxiv.org/abs/1405.5103>, 2014.
- [22] X. Zeng, M. Figueiredo, “Decreasing Weighted Sorted ℓ_1 Regularization”, *IEEE Signal Processing Letters*, vol. 21, pp. 1240–1244, 2014.
- [23] X. Zeng, M. Figueiredo, “The atomic norm formulation of OSCAR regularization with application to the Frank-Wolfe algorithm”, *Proceedings of the European Signal Processing Conference*, Lisbon, Portugal, 2014.
- [24] L. Zhong and J. Kwok, “Efficient sparse modeling with automatic feature grouping”, *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, pp. 1436–1447, 2012.