
Ordered Weighted ℓ_1 Regularized Regression with Strongly Correlated Covariates: Theoretical Aspects

Mário A. T. Figueiredo

Instituto de Telecomunicações
Instituto Superior Técnico
Universidade de Lisboa, Portugal

Robert D. Nowak

Dept. of Electrical and Computer Engineering
University of Wisconsin, Madison, USA

Abstract

This paper studies the *ordered weighted ℓ_1* (OWL) family of regularizers for sparse linear regression with strongly correlated covariates. We prove sufficient conditions for clustering correlated covariates, extending and qualitatively strengthening previous results for a particular member of the OWL family: OSCAR (*octagonal shrinkage and clustering algorithm for regression*). We derive error bounds for OWL with correlated Gaussian covariates: for cases in which clusters of covariates are strongly (even perfectly) correlated, but covariates in different clusters are uncorrelated, we show that if the true p -dimensional signal involves only s clusters, then $O(s \log p)$ samples suffice to accurately estimate it, regardless of the number of coefficients within the clusters. Since the estimation of s -sparse signals with completely independent covariates also requires $O(s \log p)$ measurements, this shows that by using OWL regularization, we pay no price (in the number of measurements) for the presence of strongly correlated covariates.

1 Introduction

In high-dimensional linear regression problems, it is likely that several covariates (also referred to as predictors or variables) are highly correlated; *e.g.*, in gene expression data, it is common to find groups of highly co-regulated genes. Using standard sparsity-inducing regularization (ℓ_1 , also known as *LASSO* (Tibshirani, 1996)) in such scenarios is known to be unsatisfactory, as it leads to the selection of one of a set of highly correlated covariates, or an

arbitrary convex combination thereof. For engineering purposes and/or scientific interpretability, it is often desirable to explicitly identify **all** of the covariates that are relevant for modelling the data, rather than just a subset thereof. Several approaches have been proposed to deal with this type of problems (Bühlmann et al., 2013; Genovese et al., 2012; Jia and Yu, 2010; Meinshausen and Bühlmann, 2010; Shah and Samworth, 2013; Shen and Huang, 2010; Zou and Hastie, 2005), the best known of which is arguably the *elastic net* (EN) introduced by Zou and Hastie (2005).

This paper is motivated by the OSCAR (*octagonal shrinkage and clustering algorithm for regression*), proposed by Bondell and Reich (2007) to address regression problems with correlated covariates, which has been shown to perform well in practice, but lacks a theoretical characterization of its error performance. The regularizer underlying OSCAR was recently shown to belong to the more general family of the *ordered weighted ℓ_1* (OWL) norms (Bogdan et al., 2013; Zeng and Figueiredo, 2014a), which also includes the ℓ_1 and ℓ_∞ norms. The goal of this paper is to provide a theoretical characterization of linear regression under OWL regularization, in the presence of highly correlated covariates. Our main contributions are the following.

- a) We prove sufficient conditions for exact covariate clustering, considerably extending the results by Bondell and Reich (2007). In particular, for the squared error loss, our result holds for the general OWL family and, more importantly, under qualitatively weaker conditions: whereas our result shows that OWL can cluster groups of more than 2 covariates, the proof by Bondell and Reich (2007) explicitly excludes that case. Furthermore, we also give clustering conditions under the absolute error loss, which we believe are novel.
- b) We derive error bounds for OWL regularization with correlated Gaussian covariates. For cases in which clusters of covariates are strongly (even perfectly) correlated, but covariates in different clusters are uncorrelated, we show that if the true p -dimensional signal involves only s clusters, then $O(s \log p)$ samples suffice to accurately estimate it, regardless of the number of coefficients within

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 41. Copyright 2016 by the authors.

clusters. Since estimating s -sparse vectors with independent variables requires just as many measurements, this shows that by using OWL regularization no price is paid (in terms of the number of measurements) for the presence of those strongly correlated covariates.

This paper includes no experimental results, as its goal is to theoretically characterize OWL regularization. The particular case of OSCAR was experimentally studied by Bondell and Reich (2007) and Zhong and Kwok (2012); the main conclusion from their experiments is not that OSCAR clearly outperforms EN in terms of accuracy, but that while typically requiring fewer degrees of freedom due to its exact clustering behavior, it is still competitive with EN. In other words, their claim is not that OSCAR achieves higher accuracy, but that its ability to identify clusters of correlated covariates improves interpretability. In very recent work on machine translation, Clark (2015) found that the ability of OSCAR to cluster coefficients brings a significant performance gain. This paper provides theoretical support to these experimental observations.

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 components equal to 1 is written as $\mathbf{1}$, and $|\mathbf{x}|$ is the vector with the absolute values of the components of \mathbf{x} . For $\mathbf{x} \in \mathbb{R}^p$, $x_{[i]}$ is its i -th largest component (i.e., $x_{[1]} \geq x_{[2]} \geq \dots \geq x_{[p]}$), and \mathbf{x}_\downarrow is the vector obtained by sorting the components of \mathbf{x} in non-increasing order. Finally, given $\mathbf{w} \in \mathbb{R}_+^p$, such that $w_1 \geq w_2 \geq \dots \geq w_p \geq 0$, $\Delta_{\mathbf{w}} = \min\{w_l - w_{l+1}, l = 1, \dots, p-1\}$ is the minimum gap between consecutive components of \mathbf{w} and $\bar{w} = \|\mathbf{w}\|_1/p$ is their average.

1.1 Definitions and Problem Formulation

The OWL norm (Bogdan et al., 2014; Zeng and Figueiredo, 2014a), is defined as

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

where $\mathbf{w} \in \mathbb{R}_+^p$ is a vector of weights, such that $w_1 \geq w_2 \geq \dots \geq w_p \geq 0$ and $w_1 > 0$. Clearly, $\Omega_{\mathbf{w}}$ satisfies $w_1 \|\mathbf{x}\|_\infty \leq \Omega_{\mathbf{w}}(\mathbf{x}) \leq w_1 \|\mathbf{x}\|_1$ (with equalities if $w_2 = \dots = w_p = 0$ or $w_1 = w_2 = \dots = w_p$, respectively). It is also easy to show (using Chebyshev’s sum inequality) that $\Omega_{\mathbf{w}}(\mathbf{x}) \geq \bar{w} \|\mathbf{x}\|_1$. The OSCAR regularizer (Bondell and Reich, 2007) is a special case of $\Omega_{\mathbf{w}}$, obtained by setting $w_i = \lambda_1 + \lambda_2 (p - i)$, where $\lambda_1, \lambda_2 \geq 0$.

This paper studies OWL-regularized linear regression under the squared error and absolute error losses. We consider the classical unconstrained formulations

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

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

where $\mathbf{A} \in \mathbb{R}^{n \times p}$ is the design matrix, as well as the following constrained formulations:

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_{\mathbf{w}}(\mathbf{x}) \text{ s.t. } \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2, \tag{4}$$

$$\min_{\mathbf{x} \in \mathbb{R}^p} \Omega_{\mathbf{w}}(\mathbf{x}) \text{ s.t. } \frac{1}{n} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon. \tag{5}$$

Notice that (since all the involved functions are convex) the constrained and unconstrained formulations are equivalent in the following sense (Lorenz and Worliczek, 2013): given a (non-zero) solution $\hat{\mathbf{x}}$ of (4) (respectively, of (5)), there is a choice of λ that makes $\hat{\mathbf{x}}$ also a solution of (2) (respectively, of (3)). Conversely, if $\hat{\mathbf{x}}$ is the unique solution of (2) (respectively, of (3)), then $\hat{\mathbf{x}}$ also solves (4) (respectively, (5)), with $\varepsilon^2 = \frac{1}{n} \|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_2^2$ (respectively, with $\varepsilon = \frac{1}{n} \|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_1$). Regardless of these equivalences, certain angles of analysis are more convenient in the unconstrained formulations (2)–(3), while others are more convenient in the constrained form (4)–(5). The equivalences mentioned above mean that results concerning the solutions of (2) and (3) are, in principle, translatable to results about the solutions of (4) and (5), and *vice-versa*.

On the algorithmic side, the key tool for solving regularization problems involving the OWL norm (such as (2) – (5)) is its Moreau proximity operator, which can be computed in $O(p \log p)$ operations (Bogdan et al., 2014; Zeng and Figueiredo, 2014b), the same being true about the Euclidean projection onto an OWL norm ball (Davis, 2015).

1.2 Preview of the Main Results and Related Work

The first of our two main results (detailed in Section 2) gives sufficient conditions for OWL regularization to cluster strongly correlated covariates, in the sense that the coefficient estimates associated with such covariates are exactly equal (in magnitude). Our result for the squared error loss significantly extends and strengthens the main theorem for OSCAR presented by Bondell and Reich (2007), since our proof involves qualitatively weaker conditions and applies to the general OWL family. Furthermore, the result for the absolute error loss is, as far as we know, novel.

Our second main result (presented in Section 3) is a finite sample bound for formulations (4)–(5). To the best of our knowledge, these are the first finite sample error bounds for sparse regression with strongly correlated columns in the design matrix. To preview this result, consider the following special case (generalized below): assume we observe

$$\mathbf{y} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\nu}, \tag{6}$$

where $\mathbf{x}^* \in \mathbb{R}^p$ is s -sparse (i.e., at most s nonzero components) and $\boldsymbol{\nu} \in \mathbb{R}^n$ is the measurement error, with

$\|\nu\|_1/n \leq \varepsilon$, and about which we make no other assumptions. The design matrix \mathbf{A} is Gaussian distributed; for the purposes of this introduction, assume the entries in each column of \mathbf{A} are i.i.d. $\mathcal{N}(0, 1)$, but different columns may be strongly correlated. Specifically, assume there are groups of identical columns, but those in different groups are uncorrelated; this models scenarios where groups of covariates are perfectly correlated, but uncorrelated with all others. Since certain columns of \mathbf{A} are identical, in general there may be many sparse vectors \mathbf{x} such that $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}^*$. Among these, let $\bar{\mathbf{x}}^*$ denote the vector with identical coefficients for replicated columns (i.e., if two columns of \mathbf{A} are identical, so are the corresponding coefficients in $\bar{\mathbf{x}}^*$).

Theorem 1.1 below states that the sufficient number of measurements n to estimate an s -sparse signal, at a given precision, grows like $n \sim s \log p$, agreeing with well-known bounds for sparse recovery under stronger conditions on \mathbf{A} , e.g., restricted isometry property, incoherence, or fully i.i.d. measurements (Candès et al., 2006; Donoho, 2006; Haupt and Nowak, 2006; Vershynin, 2014). This also shows that, by using OWL, we pay no price (in terms of number of measurements) for the colinearity of some columns of \mathbf{A} .

Theorem 1.1. *Let \mathbf{y} , \mathbf{A} , \mathbf{x}^* , \mathbf{w} , and $\Delta_{\mathbf{w}}$ be as defined above and assume $\Delta_{\mathbf{w}} > 0$. Let $\hat{\mathbf{x}}$ be a solution to either of the two problems (4) or (5). Then,*

- (i) *for every pair of columns such that $\mathbf{a}_i = \mathbf{a}_j$ (respectively, $\mathbf{a}_i = -\mathbf{a}_j$), we have $\hat{x}_i = \hat{x}_j$ (respectively, $\hat{x}_i = -\hat{x}_j$);*
- (ii) *the solution $\hat{\mathbf{x}}$ satisfies (where the expectation is w.r.t. the random \mathbf{A}):*

$$\mathbb{E} \|\hat{\mathbf{x}} - \bar{\mathbf{x}}^*\|_2 \leq \sqrt{8\pi} \left(\sqrt{32} \|\mathbf{x}^*\|_2 \frac{w_1}{\bar{w}} \sqrt{\frac{s \log p}{n}} + \varepsilon \right). \tag{7}$$

Theorem 1.1 (i) shows that OWL regularization automatically identifies and groups the colinear columns in \mathbf{A} . As mentioned above, in general there may be many sparse \mathbf{x} yielding the same $\mathbf{A}\mathbf{x}$. This is where OWL becomes important: its solution includes all the colinear columns associated with the model, rather than an arbitrary subset thereof. This result is proved in Section 2, and generalized to cases where columns are not necessarily identical, but correlated enough.

Part (ii) of Theorem 1.1 is proved in Section 3, and also generalized to strongly correlated (rather than identical) covariates. (Notice that the factor w_1/\bar{w} in (7) is typically small; e.g., for OSCAR, $\bar{w} = \lambda_1 + \lambda_2(p-1)/2$ and $w_1/\bar{w} \leq 2$, whereas for ℓ_1 , $w_1/\bar{w} = 1$.)

It is worth mentioning that these bounds for OWL are fundamentally different than those obtained for the LASSO with correlated Gaussian designs (Raskutti et al., 2010),

which do not cover the case of exactly replicated columns, and essentially require a full-rank design matrix.

Although OWL does bear similarity to the *elastic net* (EN), in the sense that they both aim at handling highly correlated covariates, OWL yields exact covariate clustering, whereas EN does not. In terms of theoretical analysis, the consistency results for the EN proved by Jia and Yu (2010) are asymptotic, not finite sample bounds, and require the so-called *elastic irrepresentability condition* (EIC), which is stronger than our assumptions. We are not aware of finite sample error bounds for EN that come close to those for OWL that we prove in this paper. Finally, although our bounds are for Gaussian designs, generalization to the sub-Gaussian case can be obtained using the tools proposed by Vershynin (2014). We thus feel that our assumptions are less restrictive and more relevant than the EIC.

It is also interesting to observe that the error bound in (7) is essentially the same holding for group-LASSO (Rao et al., 2012), assuming the groups are known a priori rather than automatically identified.

Finally, a particular member of the OWL family (for a specific choice of the weights \mathbf{w}) was recently studied, namely in terms of *false discovery rate* (FDR) control, adaptivity, and asymptotic minimaxity (Bogdan et al., 2013, 2014; Su and Candès, 2015); however, these results are only for orthogonal or uncorrelated covariates, which is not the scenario to which this paper is devoted.

We conclude this section with a simple toy example (Fig. 1) illustrating the qualitatively different behaviour of OWL and LASSO regularization. In this example, $p = 100$, $n = 10$, and \mathbf{x}^* has 20 non-zero components in 2 groups of size 10, with the corresponding columns of \mathbf{A} being highly correlated. Clearly, $n = 10$ is insufficient to allow LASSO to recover \mathbf{x}^* , which is 20-sparse, while OWL successfully recovers its structure.

2 OWL Clustering

This section studies the clustering behaviour of OWL, extending the results of Bondell and Reich (2007) in several ways: for the squared error loss, our results apply to the more general case of OWL and, more importantly, hold under qualitatively weaker conditions; the result for the absolute error case is novel.

2.1 Squared Error Loss

The following theorem shows that criterion (2) *clusters* (i.e., yields coefficient estimates of equal magnitude) the columns that are correlated enough.

Theorem 2.1. *Let $\hat{\mathbf{x}}$ be a solution of (2), and \mathbf{a}_i and \mathbf{a}_j be*




Figure 1: Toy example illustrating the qualitatively different behaviour of OWL and LASSO.

two columns of \mathbf{A} . Then,

$$(a) \quad \|\mathbf{a}_i - \mathbf{a}_j\|_2 < \Delta_w / \|\mathbf{y}\|_2 \Rightarrow \hat{x}_i = \hat{x}_j \quad (8)$$

$$(b) \quad \|\mathbf{a}_i + \mathbf{a}_j\|_2 < \Delta_w / \|\mathbf{y}\|_2 \Rightarrow \hat{x}_i = -\hat{x}_j. \quad (9)$$

Clearly, part (b) of the theorem is a simple corollary of part (a), which results from flipping the signs of either \mathbf{a}_i or \mathbf{a}_j and the corresponding coefficient. Notice that if two columns are identical, $\mathbf{a}_i = \mathbf{a}_j$, or symmetrical, $\mathbf{a}_i = -\mathbf{a}_j$, any $\Delta_w > 0$ is sufficient to guarantee that these two columns are clustered, *i.e.*, the corresponding coefficient estimates are equal in magnitude.

The following corollary addresses the case where the columns of \mathbf{A} have zero mean and unit norm (as is common practice), and results from denoting $\rho_{ij} = \mathbf{a}_i^T \mathbf{a}_j$ as the sample correlation between the i -th and j -th covariates, in which case $\|\mathbf{a}_i \pm \mathbf{a}_j\|_2 = \sqrt{2 \pm 2\rho_{ij}}$.

Corollary 2.1. *Let the columns of \mathbf{A} satisfy $\mathbf{1}^T \mathbf{a}_k = 0$ and $\|\mathbf{a}_k\|_2 = 1$, for $k = 1, \dots, p$. Denote $\rho_{ij} = \mathbf{a}_i^T \mathbf{a}_j \in [-1, 1]$. Then,*

$$(a) \quad \sqrt{2 - 2\rho_{ij}} < \Delta_w / \|\mathbf{y}\|_2 \Rightarrow \hat{x}_i = \hat{x}_j \quad (10)$$

$$(b) \quad \sqrt{2 + 2\rho_{ij}} < \Delta_w / \|\mathbf{y}\|_2 \Rightarrow \hat{x}_i = -\hat{x}_j. \quad (11)$$

Corollary 2.1 recovers the main theorem of Bondell and Reich (2007), since in the case of OSCAR, $\Delta_w = \lambda_2$. However, our result holds under qualitatively weaker conditions: unlike in their proof, we do not require that both \hat{x}_i and \hat{x}_j are different from zero and from all other \hat{x}_k , for $k \neq i, j$, neither that \mathbf{A} is such that $\hat{x}_k \geq 0$, for all k . Furthermore, our results apply to the general family of OWL regularizers, of which OSCAR is only a particular case.

2.2 Absolute Error Loss

The following theorem parallels Theorem 2.1, now for absolute error loss regression (eq. (3)).

Theorem 2.2. *Let $\hat{\mathbf{x}}$ be a solution of (3). Then*

$$(a) \quad \|\mathbf{a}_i - \mathbf{a}_j\|_1 < \Delta_w \Rightarrow \hat{x}_i = \hat{x}_j \quad (12)$$

$$(b) \quad \|\mathbf{a}_i + \mathbf{a}_j\|_1 < \Delta_w \Rightarrow \hat{x}_i = -\hat{x}_j \quad (13)$$

As above, part (b) of the theorem results directly from part (a). Under the normalization assumptions used in Corollary

2.1, another (weaker) sufficient condition can be obtained, which depends on the sample correlations, as stated in the following corollary (the proof of which simply amounts to using the well-known inequality $\|\mathbf{a}\|_1 \leq \sqrt{n}\|\mathbf{a}\|_2$ together with the assumed column normalization):

Corollary 2.2. *Let $\hat{\mathbf{x}}$ be any minimizer of the objective function in (3) and assume the columns of \mathbf{A} are normalized, that is, $\mathbf{1}^T \mathbf{a}_k = 0$ and $\|\mathbf{a}_k\|_2 = 1$, for $i = k, \dots, p$. As above, let $\rho_{ij} = \mathbf{a}_i^T \mathbf{a}_j$. Then,*

$$(a) \quad \sqrt{n(2 - 2\rho_{ij})} < \Delta_w \Rightarrow \hat{x}_i = \hat{x}_j \quad (14)$$

$$(b) \quad \sqrt{n(2 + 2\rho_{ij})} < \Delta_w \Rightarrow \hat{x}_i = -\hat{x}_j. \quad (15)$$

2.3 Proofs of Theorems 2.1 and 2.2

The proofs of Theorems 2.1 and 2.2 are based on the following two new lemmas about the OWL norm (the proofs of which are provided in the supplementary material).

Lemma 2.1. *Consider a vector $\mathbf{x} \in \mathbb{R}_+^p$ and two of its components x_i and x_j , such that $x_i > x_j$ (if they exist). Let $\mathbf{z} \in \mathbb{R}_+^p$ be obtained by applying to \mathbf{x} a so-called Pigou-Dalton¹ transfer of size $\varepsilon \in (0, (x_i - x_j)/2)$, that is: $z_i = x_i - \varepsilon$, $z_j = x_j + \varepsilon$, and $z_k = x_k$, for $k \neq i, j$. Then,*

$$\Omega_w(\mathbf{x}) - \Omega_w(\mathbf{z}) \geq \Delta_w \varepsilon. \quad (16)$$

Lemma 2.2. *Consider a vector $\mathbf{x} \in \mathbb{R}_+^p$ and two of its non-zero components x_i and x_j (if they exist). Let now $\mathbf{z} \in \mathbb{R}_+^p$ be obtained by subtracting $\varepsilon \in (0, \min\{x_i, x_j\})$ from x_i and x_j , that is: $z_i = x_i - \varepsilon$, $z_j = x_j - \varepsilon$, and $z_k = x_k$, for $k \neq i, j$. Then, (16) also holds.*

It is worth pointing out that what Lemma 2.1 states about the OWL norm Ω_w can be seen as a property of *strong Schur convexity*, which, as far as we know, didn't exist in the literature on majorization theory and Schur convexity (Marshall et al., 2011). For more details about this observation, which is tangential to the topic of this paper, but potentially useful in other contexts, see Appendix A.

The proof of Theorem 2.1 also uses a basic result in convex analysis relating minimizers of a convex function with its

¹The Pigou-Dalton (a.k.a. Robin Hood) transfer is a fundamental quantity used in the mathematical study of economic inequality (Dalton, 1920; Pigou, 1912).

directional derivatives (Rockafellar, 1970). Given a proper function f , its *directional derivative* at $\mathbf{x} \in \text{dom}(f)$ (i.e., $f(\mathbf{x}) \neq \infty$), in the direction \mathbf{u} , is defined as

$$f'(\mathbf{x}; \mathbf{u}) = \lim_{\alpha \rightarrow 0^+} (f(\mathbf{x} + \alpha \mathbf{u}) - f(\mathbf{x})) / \alpha.$$

Lemma 2.3. *Let f be a real-valued, proper, convex function, and $\mathbf{x} \in \text{dom}(f)$. Then, $\mathbf{x} \in \arg \min f$, if and only if $f'(\mathbf{x}; \mathbf{u}) \geq 0$, for any \mathbf{u} .*

Proof. (of Theorem 2.1) Let $L_2(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2$ and $f(\mathbf{x}) = L_2(\mathbf{x}) + \Omega_{\mathbf{w}}(\mathbf{x})$ (i.e., the objective function in (2)). Assume that the condition $\|\mathbf{y}\|_2 \|\mathbf{a}_i - \mathbf{a}_j\|_2 < \Delta_{\mathbf{w}}$ is satisfied for some pair of columns and consider some $\hat{\mathbf{x}}$ such that $\hat{x}_i \neq \hat{x}_j$ (w.l.o.g., let $\hat{x}_i > \hat{x}_j$). The directional derivative of L_2 at $\hat{\mathbf{x}}$, in the direction \mathbf{u} , where $u_i = -1$, $u_j = 1$, and $u_k = 0$, for $k \neq i, j$, is

$$\begin{aligned} L_2'(\hat{\mathbf{x}}; \mathbf{u}) &= \lim_{\alpha \rightarrow 0^+} \frac{\|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}} + \alpha(\mathbf{a}_i - \mathbf{a}_j)\|_2^2 - \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}\|_2^2}{2\alpha} \\ &= \mathbf{g}^T(\mathbf{a}_i - \mathbf{a}_j), \end{aligned} \quad (17)$$

where $\mathbf{g} = \mathbf{y} - \mathbf{A}\hat{\mathbf{x}}$. Consider now the directional derivative of $\Omega_{\mathbf{w}}$ at $\hat{\mathbf{x}}$, in the same direction \mathbf{u} :

$$\Omega_{\mathbf{w}}'(\hat{\mathbf{x}}; \mathbf{u}) = \lim_{\alpha \rightarrow 0^+} \frac{\Omega_{\mathbf{w}}(\hat{\mathbf{x}} + \alpha \mathbf{u}) - \Omega_{\mathbf{w}}(\hat{\mathbf{x}})}{\alpha}.$$

If \hat{x}_i and \hat{x}_j are both non-negative or non-positive, $|\hat{\mathbf{x}} + \alpha \mathbf{u}|$ corresponds to a Pigou-Dalton transfer of size α applied to $|\hat{\mathbf{x}}|$, thus Lemma 2.1 (recalling $\Omega_{\mathbf{w}}(\mathbf{v}) = \Omega_{\mathbf{w}}(|\mathbf{v}|)$) guarantees that

$$\begin{aligned} \Omega_{\mathbf{w}}'(\hat{\mathbf{x}}; \mathbf{u}) &= \lim_{\alpha \rightarrow 0^+} \frac{\Omega_{\mathbf{w}}(|\hat{\mathbf{x}} + \alpha \mathbf{u}|) - \Omega_{\mathbf{w}}(|\hat{\mathbf{x}}|)}{\alpha} \\ &\leq \lim_{\alpha \rightarrow 0^+} \frac{-\Delta_{\mathbf{w}} \alpha}{\alpha} = -\Delta_{\mathbf{w}}. \end{aligned} \quad (19)$$

If $\text{sign}(\hat{x}_i) \text{sign}(\hat{x}_j) = -1$, then $|\hat{\mathbf{x}} + \alpha \mathbf{u}|$ corresponds to subtracting α from both $|\hat{x}_i|$ and $|\hat{x}_j|$, thus Lemma 2.2 also yields (19). Finally, adding (17) and (19), and using the Cauchy-Schwarz inequality,

$$\begin{aligned} f'(\mathbf{x}; \mathbf{u}) &\leq \mathbf{g}^T(\mathbf{a}_i - \mathbf{a}_j) - \Delta_{\mathbf{w}} \\ &\leq \|\mathbf{g}\|_2 \|\mathbf{a}_i - \mathbf{a}_j\|_2 - \Delta_{\mathbf{w}} \\ &\leq \|\mathbf{y}\|_2 \|\mathbf{a}_i - \mathbf{a}_j\|_2 - \Delta_{\mathbf{w}} < 0, \end{aligned}$$

(after noticing $\|\mathbf{g}\|_2 \leq \|\mathbf{y}\|_2$), showing that $\hat{\mathbf{x}}$ is not a minimizer. \square

The following lemma (proved in the supplementary material) will be used in proving Theorem 2.2.

Lemma 2.4. *Let $L_1(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1$, consider any \mathbf{x} and two of its components, x_i and x_j , and define \mathbf{v} according to $v_i = x_i - \varepsilon$, $v_j = x_j + \varepsilon$, for some $\varepsilon \in \mathbb{R}$, and $v_k = x_k$, for $k \neq i, j$. Then, $L_1(\mathbf{v}) - L_1(\mathbf{x}) \leq |\varepsilon| \|\mathbf{a}_i - \mathbf{a}_j\|_1$.*

Proof. (of Theorem 2.2) Assume the condition $\Delta_{\mathbf{w}} > \|\mathbf{a}_i - \mathbf{a}_j\|_1$ is satisfied, and that $\hat{\mathbf{x}}$ is a solution of (3). To prove that $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$, suppose that $\text{sign}(\hat{x}_i) \neq \text{sign}(\hat{x}_j)$, which implies that at least one of \hat{x}_i or \hat{x}_j is non-zero; without loss of generality, let $\hat{x}_i > 0$, thus $\hat{x}_j \leq 0$. We need to consider two cases: $\hat{x}_j < 0$ and $\hat{x}_j = 0$.

- If $\hat{x}_j < 0$, take an alternative solution \mathbf{v} , with $v_i = \hat{x}_i - \varepsilon$, $v_j = \hat{x}_j + \varepsilon$, where $\varepsilon \in (0, \min\{\hat{x}_i, -\hat{x}_j\})$, and $v_k = \hat{x}_k$, for $k \neq i, j$. Since $\hat{x}_i > 0$ and $\hat{x}_j < 0$, we have $|v_i| = |\hat{x}_i| - \varepsilon$ and $|v_j| = |\hat{x}_j| + \varepsilon$, thus Lemma 2.2 yields

$$\Omega_{\mathbf{w}}(\hat{\mathbf{x}}) - \Omega_{\mathbf{w}}(\mathbf{v}) = \Omega_{\mathbf{w}}(|\hat{\mathbf{x}}|) - \Omega_{\mathbf{w}}(|\mathbf{v}|) \geq \Delta_{\mathbf{w}} \varepsilon. \quad (20)$$

Combining this inequality with Lemma 2.4 contradicts the optimality of $\hat{\mathbf{x}}$, since

$$\begin{aligned} L_1(\mathbf{v}) + \Omega_{\mathbf{w}}(\mathbf{v}) - (L_1(\hat{\mathbf{x}}) + \Omega_{\mathbf{w}}(\hat{\mathbf{x}})) \\ \leq (\|\mathbf{a}_i - \mathbf{a}_j\|_1 - \Delta_{\mathbf{w}}) \varepsilon < 0. \end{aligned} \quad (21)$$

- If $\hat{x}_j = 0$, follow the same argument with $\varepsilon \in (0, x_i/2)$, i.e., $v_i = \hat{x}_i - \varepsilon$ and $v_j = \hat{x}_j + \varepsilon = \varepsilon$. Since, in this case, $|v_i| = |\hat{x}_i| - \varepsilon$ and $|v_j| = |\hat{x}_j| + \varepsilon$, Lemma 2.1 also yields inequality (20), which combined with Lemma 2.4 contradicts again the optimality of $\hat{\mathbf{x}}$.

Knowing $\text{sign}(\hat{x}_i) = \text{sign}(\hat{x}_j)$, let $\hat{\mathbf{x}}$ be a solution of (3) such that $\hat{x}_i \neq \hat{x}_j$; without loss of generality, consider that both \hat{x}_i and \hat{x}_j are non-negative, and that $\hat{x}_i > \hat{x}_j$. Consider an alternative solution \mathbf{u} such that $u_i = \hat{x}_i - \varepsilon$, $u_j = \hat{x}_j + \varepsilon$, for some $\varepsilon \in (0, (x_i - x_j)/2)$, and $u_k = \hat{x}_k$, for $k \neq i, j$. Combining Lemmas 2.1 and 2.4, yields precisely the same inequality as in (21), contradicting the optimality of $\hat{\mathbf{x}}$, thus concluding the proof. \square

3 OWL Error Bounds

Consider the observation model in (6) and the other assumptions about \mathbf{x} and \mathbf{v} made in Section 1.1. Moreover, assume that $\mathbf{x}^* \in \mathbb{R}^p$ satisfies $\|\mathbf{x}^*\|_1 \leq \sqrt{s} \|\mathbf{x}^*\|_2$; this is true, e.g., if \mathbf{x}^* is s -sparse. At the heart of our analysis is the following model for correlated measurement matrices. Assume that the rows of $\mathbf{A} \in \mathbb{R}^{n \times p}$ are i.i.d. $\mathcal{N}(\mathbf{0}, \mathbf{C}^T \mathbf{C})$, i.e., its columns are not necessarily independent. Let matrix \mathbf{C} be $r \times p$, with $n \leq r \leq p$, so that $\text{rank}(\mathbf{C}) \leq r$. Note that \mathbf{A} can be written as $\mathbf{A} = \mathbf{B}\mathbf{C}$, where $\mathbf{B} \in \mathbb{R}^{n \times r}$ has i.i.d. $\mathcal{N}(0, 1)$ entries. The role of \mathbf{C} is to mix, or even replicate, columns of \mathbf{B} . Figure 2 illustrates this in a case where every column is one of three identical replicates.

3.1 General OWL Error Bound

The main result of this section is stated in the following theorem, the proof of which is based on the techniques introduced by Vershynin (2014). We also present a corollary for the particular case where \mathbf{C} simply replicates columns of \mathbf{B} , i.e., if \mathbf{A} includes groups of identical columns; this

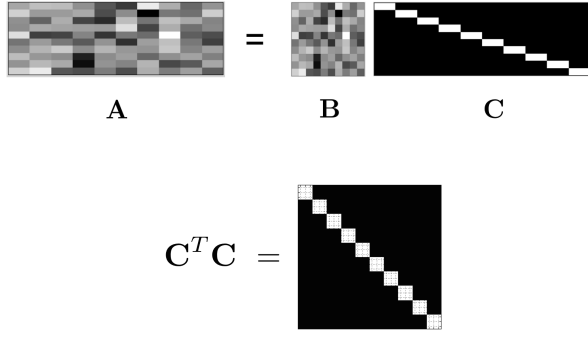


Figure 2: Matrix $A = BC$ with 10 groups of 3 identical replicates and the corresponding covariance matrix $C^T C$.

corollary is shown to imply part (ii) of Theorem 1.1. All expectations are with respect to the Gaussian distribution of the design matrix A .

Theorem 3.1. *Let \mathbf{y} , A , \mathbf{x}^* , ε , and \mathbf{w} be as defined above, and let $\hat{\mathbf{x}}$ be a solution to one of the optimization problems (4) or (5). Then,*

$$\begin{aligned} & \mathbb{E} \sqrt{(\hat{\mathbf{x}} - \mathbf{x}^*)^T C^T C (\hat{\mathbf{x}} - \mathbf{x}^*)} \\ & \leq \sqrt{8\pi} \left(4\sqrt{2} \min_{\ell=1,2} \|C\|_{\ell} \|\mathbf{x}^*\|_2 \frac{w_1}{\bar{w}} \sqrt{\frac{s \log p}{n}} + \varepsilon \right), \end{aligned} \quad (22)$$

where (recall) $\bar{w} = p^{-1} \|\mathbf{w}\|_1$ and $\min_{\ell=1,2} \|C\|_{\ell}$ is the min of the matrix 1-norm and 2-norm of C .

To help understand this bound, consider a special case recovering known results. If $r = p$ and $C = I$, matrix A is i.i.d. $\mathcal{N}(0, 1)$, which corresponds to the well known compressive sensing case. The bound (22) recovers the usual type of result in this situation (Vershynin, 2014), *i.e.*,

$$\mathbb{E} \|\hat{\mathbf{x}} - \mathbf{x}^*\|_2 = O \left(\|\mathbf{x}^*\|_2 \sqrt{\frac{s \log p}{n}} \right).$$

Notably, in this setting the OWL error bound is the same (up to small constant factors) as that of LASSO. Slightly more generally, if $C \neq I$, but has full rank, and λ_{\max} and $\lambda_{\min} > 0$ are its largest and smallest singular values, then using the $\|C\|_2$ factor in (22) yields bounds similar to those proved by Raskutti et al. (2010):

$$\mathbb{E} \|\hat{\mathbf{x}} - \mathbf{x}^*\|_2 = O \left(\frac{\lambda_{\max}}{\lambda_{\min}} \|\mathbf{x}^*\|_2 \sqrt{\frac{s \log p}{n}} \right),$$

Since for $w_1 = w_2 = \dots = w_p$, $\Omega_{\mathbf{w}}(\mathbf{x}) = w_1 \|\mathbf{x}\|_1$, (22) also holds for the LASSO.

The bound in (22) is more novel and interesting if C is rank deficient. Consider $r < p$, with C leading to exactly replicated columns, as in Fig. 2. In this case, the covariance $C^T C$ has a block diagonal structure (Fig. 2), with each

block, corresponding to a group of replicated columns, being equal to a rank-1 matrix with all entries equal. In this case, $(\hat{\mathbf{x}} - \mathbf{x}^*)^T C^T C (\hat{\mathbf{x}} - \mathbf{x}^*)$ is the sum of squared errors between the averages of $\hat{\mathbf{x}}$ and \mathbf{x}^* within each group. This is very reasonable because both $A\hat{\mathbf{x}}$ and $A\mathbf{x}^*$ are functions of only these averages, since the columns corresponding to the groups are identical. Recall that Theorems 2.1 and 2.2 imply that $\hat{\mathbf{x}}$ is constant-valued in each group of replicated columns. Also, in this case $\|C\|_1 = 1$, whereas $\|C\|_2$ is equal to the square-root of the the largest group size. Theorem 1.1 follows directly from these observations.

Proof. (Theorem 1.1 (ii)) Recall that $\bar{\mathbf{x}}^*$ satisfies $A\bar{\mathbf{x}}^* = A\mathbf{x}^*$ and that if two columns of A are identical, then so are the corresponding components of $\bar{\mathbf{x}}^*$. Because of the group structure of $\bar{\mathbf{x}}^*$ and $\hat{\mathbf{x}}$ (assuming strictly decreasing weights), and the special form of C in the case of exactly replicated columns,

$$\|\hat{\mathbf{x}} - \bar{\mathbf{x}}^*\|_2^2 \leq (\hat{\mathbf{x}} - \bar{\mathbf{x}}^*)^T C^T C (\hat{\mathbf{x}} - \bar{\mathbf{x}}^*),$$

from which (7) results. \square

More generally, if C is *approximately* like that of Fig. 2 (each column is *approximately* 1-sparse), then the same reasoning and interpretation apply *approximately*. For example, if each column of C is sufficiently close to one of the canonical unit vectors, then Theorems 2.1–2.2 imply that $\hat{\mathbf{x}}$ is constant on each group of (nearly) replicated columns, effectively averaging the corresponding columns in the prediction $A\hat{\mathbf{x}}$, helping to mitigate the effects of noise in these features and improving prediction.

The bound (22) in Theorem 3.1 holds for both (4) and (5). In fact, since $\frac{1}{n} \|A\mathbf{x} - \mathbf{y}\|_2^2 \leq \varepsilon^2$ implies $\frac{1}{n} \|A\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon$, the ℓ_1 constraint is less restrictive. In both cases, Theorem 3.1 shows that the number of samples sufficient to estimate an s -sparse signal with a given precision grows like $n \sim s \log p$; 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 (Candès et al., 2006; Donoho, 2006; Haupt and Nowak, 2006; Vershynin, 2014). In the case of groups of (nearly) replicated columns, the number of samples grows linearly with the number of nonzero groups, rather than the total number of nonzero components in $\bar{\mathbf{x}}^*$. This is where OWL regularization becomes important, by selecting all colinear columns associated with the model; *i.e.*, if the columns are colinear (or correlated enough), the OWL solution selects a representation including all the columns associated with the sparse model, rather than a subset.

3.2 Proof of Theorem 3.1

The proof of Theorem 3.1 is based on the approach developed by Vershynin (2014). The key ingredient is the

so-called *general M^* bound* (Vershynin, 2014, Theorem 5.1), which applies to the case where \mathbf{A} is i.i.d. Gaussian ($\mathbf{C} = \mathbf{I}$, in our set-up). The following theorem extends that bound to cover our model $\mathbf{A} = \mathbf{BC}$, for general \mathbf{C} .

Theorem 3.2 (extended general M^* bound). *Let \mathcal{T} be a bounded subset of \mathbb{R}^p , $\mathbf{B} \in \mathbb{R}^{n \times r}$ an i.i.d. Gaussian matrix, $\mathbf{C} \in \mathbb{R}^{r \times p}$ a fixed matrix, and $\mathbf{A} = \mathbf{BC} \in \mathbb{R}^{n \times p}$. Fix $\varepsilon \geq 0$ and consider the set*

$$\mathcal{T}_\varepsilon := \{\mathbf{u} \in \mathcal{T} : \|\mathbf{A}\mathbf{u}\|_1/n \leq \varepsilon\}. \quad (23)$$

Then, with $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})$ being a standard Gaussian random vector in \mathbb{R}^q ,

$$\begin{aligned} & \mathbb{E} \sup_{\mathbf{u} \in \mathcal{T}_\varepsilon} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \\ & \leq \sqrt{\frac{8\pi}{n}} \mathbb{E} \sup_{\mathbf{u} \in \mathcal{T}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| + \varepsilon \sqrt{\frac{\pi}{2}}. \end{aligned} \quad (24)$$

The proof (in the supplementary materials) is based on symmetrization and contraction inequalities, modifying the proof by Vershynin (2014) to account for \mathbf{C} . Theorem 3.2 can be used to derive error bounds for estimating signals known to belong to some subset (sparsity is a special case considered below). Let $\mathcal{K} \subset \mathbb{R}^p$ be given and suppose that we observe $\mathbf{y} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\nu}$, with $\frac{1}{n}\|\boldsymbol{\nu}\|_1 \leq \varepsilon$, where $\mathbf{x}^* \in \mathcal{K}$. Recall the *Minkowski gauge* of \mathcal{K} , defined as

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

which is a norm if \mathcal{K} is a compact and origin-symmetric convex set with non-empty interior (Rockafellar, 1970). The following theorem, which extends one by Vershynin (2014, Theorem 6.2), is then used to prove Theorem 3.1.

Theorem 3.3. *Let $\hat{\mathbf{x}} \in \arg \min_{\mathbf{x}} \|\mathbf{x}\|_{\mathcal{K}}$, subject to $\frac{1}{n}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon$, then*

$$\begin{aligned} & \mathbb{E} \sup_{\mathbf{x}^* \in \mathcal{K}} \{(\hat{\mathbf{x}} - \mathbf{x}^*)^T \mathbf{C}^T \mathbf{C}(\hat{\mathbf{x}} - \mathbf{x}^*)\}^{1/2} \\ & \leq \sqrt{8\pi} \left(\frac{\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|}{\sqrt{n}} + \varepsilon \right). \end{aligned} \quad (25)$$

Proof. The constraint guarantees that $\frac{1}{n}\|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_1 \leq \varepsilon$, whereas by assumption, $\frac{1}{n}\|\mathbf{A}\mathbf{x}^* - \mathbf{y}\|_1 = \frac{1}{n}\|\boldsymbol{\nu}\|_1 \leq \varepsilon$. Thus, $\|\hat{\mathbf{x}}\|_{\mathcal{K}} \leq \|\mathbf{x}^*\|_{\mathcal{K}} \leq 1$, since $\mathbf{x}^* \in \mathcal{K}$. The inequality $\|\hat{\mathbf{x}}\|_{\mathcal{K}} \leq 1$ implies that $\hat{\mathbf{x}} \in \mathcal{K}$.

Next, apply Theorem 3.2 to the set $\mathcal{T} = \mathcal{K} - \mathcal{K}$, with 2ε instead of ε , yielding

$$\begin{aligned} & \mathbb{E} \sup_{\mathbf{u} \in \mathcal{T}_{2\varepsilon}} (\mathbf{u}^T \mathbf{C}^T \mathbf{C} \mathbf{u})^{1/2} \\ & \leq \sqrt{2\pi/n} \mathbb{E} \sup_{\mathbf{u} \in \mathcal{T}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| + \sqrt{8\pi}\varepsilon, \end{aligned}$$

From here, all we need to show is that for any $\mathbf{x}^* \in \mathcal{K}$, $\hat{\mathbf{x}} - \mathbf{x}^* \in \mathcal{T}_{2\varepsilon}$. To see this, simply note that $\hat{\mathbf{x}}, \mathbf{x}^* \in \mathcal{K}$, so $\hat{\mathbf{x}} - \mathbf{x}^* \in \mathcal{K} - \mathcal{K} = \mathcal{T}$. By the triangle inequality,

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

showing that $\mathbf{u} = \hat{\mathbf{x}} - \mathbf{x}^* \in \mathcal{T}_{2\varepsilon}$ (see (23)). \square

Proof. (Theorem 3.1) Since \mathbf{x}^* is assumed to satisfy $\|\mathbf{x}^*\|_1 \leq \sqrt{s}\|\mathbf{x}^*\|_2$, we first need to construct an OWL ball that contains all $\mathbf{x} \in \mathbb{R}^p$ with $\|\mathbf{x}\|_1 \leq \sqrt{s}\|\mathbf{x}^*\|_2$. Let

$$\mathcal{K} = \{\mathbf{x} \in \mathbb{R}^p : \Omega_w(\mathbf{x}) \leq w_1 \sqrt{s}\|\mathbf{x}^*\|_2\}.$$

Because $\Omega_w(\mathbf{x}) \leq w_1\|\mathbf{x}\|_1$, all vectors satisfying $\|\mathbf{x}\|_1 \leq \sqrt{s}\|\mathbf{x}^*\|_2$ belong to \mathcal{K} . Also, because $\Omega_w(\mathbf{x})$ is a norm, and \mathcal{K} a ball thereof, $\|\mathbf{x}\|_{\mathcal{K}}$ is proportional to $\Omega_w(\mathbf{x})$.

The quantity $\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|$ in (25) is called the width of \mathcal{K} and satisfies

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

Noting that $\|\mathbf{C}\mathbf{u}\|_1 \leq \|\mathbf{C}\|_1\|\mathbf{u}\|_1 \leq \|\mathbf{C}\|_1\Omega_w(\mathbf{u})/w$ (see the first paragraph of Subsection 1.1), and the fact that the triangle inequality and the definition of \mathcal{K} imply that, for any $\mathbf{u} \in \mathcal{K} - \mathcal{K}$, $\Omega_w(\mathbf{u}) \leq 2w_1\sqrt{s}\|\mathbf{x}^*\|_2$, yields

$$\|\mathbf{C}\mathbf{u}\|_1 \leq 2\|\mathbf{C}\|_1 \frac{w_1}{w} \sqrt{s}\|\mathbf{x}^*\|_2 =: \rho. \quad (26)$$

The width can be then bounded as

$$\begin{aligned} \mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle| &\leq \mathbb{E} \sup_{\{\mathbf{v} : \|\mathbf{v}\|_1 \leq \rho\}} |\langle \mathbf{g}, \mathbf{v} \rangle| \\ &\leq \rho \mathbb{E} \max_{i=1, \dots, r} |g_i|, \end{aligned} \quad (27)$$

where the second inequality results from the fact that the \mathbf{v} achieving the supremum places all its mass ρ on the largest component of \mathbf{g} (in magnitude). The classical Gaussian tail bound² together with the *union bound* yield $\mathbb{P}(\max_{i=1, \dots, r} |g_i| > t) \leq re^{-t^2/2}$; consequently,

$$\begin{aligned} \mathbb{E} \max_{i=1, \dots, r} |g_i| &= \int_0^\infty \mathbb{P}\left(\max_{i=1, \dots, r} |g_i| > t\right) dt \\ &\leq \sqrt{2\log r} + r \int_{\sqrt{2\log r}}^\infty e^{-t^2/2} dt \\ &\leq \sqrt{2\log r} + \sqrt{\pi/2} \\ &< 2\sqrt{2\log r}, \end{aligned} \quad (28)$$

where the second inequality results from applying the Gaussian tail bound², and the third from assuming $r > 2$. Plugging ρ (as defined in (26)) back in, leads to

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K} - \mathcal{K}} |\langle \mathbf{g}, \mathbf{C}\mathbf{u} \rangle| \leq 4\sqrt{2}\|\mathbf{C}\|_1 \frac{w_1}{w} \|\mathbf{x}^*\|_2 \sqrt{s \log r}.$$

²If $g \sim \mathcal{N}(0, 1)$, then $\mathbb{P}(g > t) \leq e^{-t^2/2}/2$.

We can also modify the argument above to obtain a different bound, in terms of $\|\mathbf{C}\|_2$ instead of $\|\mathbf{C}\|_1$, which can be tighter in certain cases. Recall that we must bound the width $\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K}-\mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle|$ and that, as shown above (at the very beginning of the proof),

$$\|\mathbf{u}\|_1 \leq 2 \frac{w_1}{\bar{w}} \sqrt{s} \|\mathbf{x}^*\|_2 =: \rho'. \tag{29}$$

Consequently,

$$\begin{aligned} \mathbb{E} \sup_{\mathbf{u} \in \mathcal{K}-\mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| &\leq \mathbb{E} \sup_{\mathbf{v}: \|\mathbf{v}\|_1 \leq \rho'} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{v} \rangle| \\ &\leq \rho' \mathbb{E} \max_{i=1, \dots, p} |\mathbf{c}_i^T \mathbf{g}|, \end{aligned}$$

where \mathbf{c}_i is the i -th column of \mathbf{C} ; the second inequality stems from the fact that the \mathbf{v} achieving the supremum places its total mass ρ' on the largest component of $\mathbf{C}^T \mathbf{g}$ (in magnitude). Note that $\mathbf{c}_i^T \mathbf{g} \sim \mathcal{N}(0, \|\mathbf{c}_i\|_2^2)$ and that $\|\mathbf{c}_i\|_2 \leq \|\mathbf{C}\|_2$, since $\mathbf{c}_i = \mathbf{C} \mathbf{e}_i$, where \mathbf{e}_i is a canonical unit vector. From this and the bounding argument in (28), it follows that

$$\begin{aligned} \mathbb{E} \max_{i=1, \dots, p} |\mathbf{c}_i^T \mathbf{g}| &\leq \|\mathbf{C}\|_2 \mathbb{E} \max_{i=1, \dots, p} \left| \frac{\mathbf{c}_i^T \mathbf{g}}{\|\mathbf{c}_i\|_2} \right|, \\ &\leq 2 \|\mathbf{C}\|_2 \sqrt{2 \log p}, \end{aligned}$$

where we assume $p > 2$. Plugging ρ' (as defined in (29)) back in yields the bound

$$\mathbb{E} \sup_{\mathbf{u} \in \mathcal{K}-\mathcal{K}} |\langle \mathbf{C}^T \mathbf{g}, \mathbf{u} \rangle| \leq 4\sqrt{2} \|\mathbf{C}\|_2 \frac{w_1}{\bar{w}} \|\mathbf{x}^*\|_2 \sqrt{s \log p}.$$

Theorem 3.1 now follows directly from Theorem 3.3. \square

4 Conclusion

In this paper, we have studied sparse linear regression with strongly correlated covariates under the recently proposed *ordered weighted ℓ_1 (OWL) regularization*, which generalizes the *octagonal shrinkage and clustering algorithm for regression (OSCAR)* (Bondell and Reich, 2007). We have proved sufficient conditions for OWL regularization to cluster the coefficient estimates, extending and qualitatively strengthening a previous result by Bondell and Reich (2007). We have also characterized the statistical performance of OWL regularization for generative models with clusters of strongly correlated covariates. Essentially, we have shown that, by using OWL regularization, no price is paid (in terms of the number of measurements) for the presence of strongly correlated covariates.

Future work will include the experimental evaluation of OWL regularization and its combination with other loss functions, such as logistic and hinge. An important open problem concerns the choice of the weight vector \mathbf{w} , in order to fully exploit the flexibility of the OWL family.

Appendix A: Strong Schur Convexity of $\Omega_{\mathbf{w}}$

This appendix briefly reviews the main concepts of majorization and Schur convexity (Marshall et al., 2011), and introduces a new notion of *strong Schur convexity*, showing that Lemma 2.1 is nothing but a statement about the strong Schur convexity of $\Omega_{\mathbf{w}}$.

Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$. Vector \mathbf{y} is said to *majorize* \mathbf{x} (denoted $\mathbf{y} \succ \mathbf{x}$) if

$$\mathbf{1}^T \mathbf{x} = \mathbf{1}^T \mathbf{y} \text{ and } \sum_{i=1}^k y_{[i]} \geq \sum_{i=1}^k x_{[i]}, \text{ for } k = 1, \dots, p-1.$$

Intuitively, $\mathbf{y} \succ \mathbf{x}$ if the two vectors have the same sum, and the components of \mathbf{x} have a more homogenous distribution than those of \mathbf{y} . If \mathbf{y} is a permutation of \mathbf{x} , then $\mathbf{y} \succ \mathbf{x}$ and $\mathbf{x} \succ \mathbf{y}$. The majorization relation is a *preorder* (i.e., it is reflexive and transitive).

Let $\mathcal{A} \subseteq \mathbb{R}^p$. Function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ is said to be *Schur-convex* on \mathcal{A} , if $\mathbf{y} \succ \mathbf{x} \Rightarrow \phi(\mathbf{y}) \geq \phi(\mathbf{x})$. Furthermore, if

$$(\mathbf{y} \succ \mathbf{x}) \wedge (\mathbf{y} \text{ is not a permutation of } \mathbf{x}) \Rightarrow \phi(\mathbf{y}) > \phi(\mathbf{x}),$$

then ϕ is said to be *strictly Schur-convex*. Intuitively, Schur-convex functions “prefer” (i.e., yield lower values) vector arguments with more uniformly distributed components.

A definition of *strong Schur convexity* requires a measure of “amount of majorization” of \mathbf{y} with respect to \mathbf{x} . A natural choice for this purpose is the so-called *Pigou-Dalton (a.k.a. Robin Hood) transfer* (Dalton, 1920; Marshall et al., 2011; Pigou, 1912). Specifically, given \mathbf{y} , consider two of its components y_i and y_j , such that $y_i > y_j$. We say that \mathbf{y} ε -majorizes \mathbf{x} , denoted $\mathbf{y} \succ_{\varepsilon} \mathbf{x}$, if \mathbf{x} results from a Pigou-Dalton transfer of size $\varepsilon \in (0, (y_i - y_j)/2]$ applied to \mathbf{y} , i.e., $x_i = y_i - \varepsilon$, $x_j = y_j + \varepsilon$, and $x_k = y_k$, for $k \neq i, j$.

Based on the notion of ε -majorization introduced in the previous paragraph, we propose the following definition of *strong Schur convexity*. A function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ is said to be δ -strongly Schur convex on \mathcal{A} if

$$\mathbf{y} \succ_{\varepsilon} \mathbf{x} \Rightarrow \phi(\mathbf{y}) - \phi(\mathbf{x}) \geq \delta \varepsilon.$$

Finally, given this definition of strong Schur convexity, it is clear that what Lemma 2.1 shows is that the OWL norm $\Omega_{\mathbf{w}}$ is $\Delta_{\mathbf{w}}$ -strongly Schur convex on \mathbb{R}_+^p . In contrast, it is easy to show that neither the ℓ_1 norm nor the EN regularizer are strongly Schur convex.

Acknowledgment

MF’s work was partially supported by the *Fundação para a Ciência e Tecnologia* (Portugal), grant UID/EEA/50008/2013. RN’s work was partially supported by the National Science Foundation (USA), grant CCF-1218189.

References

- M. Bogdan, E. van den Berg, W. Su, and E. Candès. Statistical estimation and testing via the ordered ℓ_1 norm. Technical report, <http://arxiv.org/pdf/1310.1969v1.pdf>, 2013.
- M. Bogdan, E. van den Berg, C. Sabatti, W. Su, and E. Candès. SLOPE – adaptive variable selection via convex optimization. Technical report, arxiv.org/abs/1407.3824, 2014.
- H. Bondell and B. Reich. Regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics*, 64:115–123, 2007.
- 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*, 143:1835–1858, 2013.
- E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489–509, 2006.
- J. Clark. *Locally Non-Linear Learning via Feature Induction and Structured Regularization in Statistical Machine Translation*. PhD thesis, Language Technologies Institute, School of Computer Science, Carnegie Mellon University, 2015.
- H. Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30:348–361, 1920.
- D. Davis. An $o(n \log(n))$ algorithm for projecting onto the ordered weighted ℓ_1 norm ball. Technical report, arxiv.org/abs/1505.00870, 2015.
- D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 2006.
- C. Genovese, J. Jin, L. Wasserman, and Z. Yao. A comparison of the LASSO and marginal regression. *Journal of Machine Learning Research*, 13:2107–2143, 2012.
- J. Haupt and R. Nowak. Signal reconstruction from noisy random projections. *IEEE Transactions on Information Theory*, 52:4036–4048, 2006.
- J. Jia and B. Yu. On model selection consistency of the elastic net when $p \gg n$. *Statistica Sinica*, 20:595–611, 2010.
- D. Lorenz and N. Worliczek. Necessary conditions for variational regularization schemes. *Inverse Problems*, 29(7):075016, 2013.
- A. Marshall, I. Olkin, and B. Arnold. *Inequalities: Theory of Majorization and Its Applications*. Springer, New York, 2011.
- N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society (B)*, 72(4):417–473, 2010.
- A. Pigou. *Wealth and Welfare*. Macmillan, London, 1912.
- N. Rao, B. Recht, and R. Nowak. Tight measurement bounds for exact recovery of structured sparse signals. In *Proceedings of AISTATS*, 2012.
- G. Raskutti, M. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- R. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- R. Shah and R. Samworth. Variable selection with error control: another look at stability selection. *Journal of the Royal Statistical Society (B)*, 75(1):55–80, 2013.
- X. Shen and H.-C. Huang. Grouping pursuit through a regularization solution surface. *Journal of the American Statistical Association*, 105:727–739, 2010.
- W. Su and E. Candès. SLOPE is adaptive to unknown sparsity and asymptotically minimax. Technical report, arxiv.org/abs/1503.08393, 2015.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society (B)*, 58(1):267–288, 1996.
- R. Vershynin. Estimation in high dimensions: A geometric perspective. Technical report, <http://arxiv.org/abs/1405.5103>, 2014.
- X. Zeng and M. Figueiredo. Decreasing weighted sorted ℓ_1 regularization. *IEEE Signal Processing Letters*, 21:1240–1244, 2014a.
- X. Zeng and M. Figueiredo. The ordered weighted ℓ_1 norm: atomic formulation, dual norm, and projections. Technical report, arxiv.org/abs/1409.4271, 2014b.
- L. Zhong and J. Kwok. Efficient sparse modeling with automatic feature grouping. *IEEE Transactions on Neural Networks and Learning Systems*, 23:1436–1447, 2012.
- 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.