Signal Recovery in Unions of Subspaces with Applications to Compressive Imaging

Nikhil Rao[†], Benjamin Recht^{*} and Robert D. Nowak[†]

* Computer Sciences Department, University of Wisconsin-Madison

 † Electrical and Computer Engineering, University of Wisconsin-Madison

September 2012

Abstract

In applications ranging from communications to genetics, signals can be modeled as lying in a union of subspaces. Under this model, signal coefficients that lie in certain subspaces are active or inactive together. The potential subspaces are known in advance, but the particular set of subspaces that are active (i.e., in the signal support) must be learned from measurements. We show that exploiting knowledge of subspaces can further reduce the number of measurements required for exact signal recovery, and derive universal bounds for the number of measurements needed. The bound is universal in the sense that it only depends on the number of subspaces under consideration, and their orientation relative to each other. The particulars of the subspaces (e.g., compositions, dimensions, extents, overlaps, etc.) does not affect the results we obtain. In the process, we derive sample complexity bounds for the special case of the group lasso with overlapping groups (the latent group lasso), which is used in a variety of applications. Finally, we also show that wavelet transform coefficients of images can be modeled as lying in groups, and hence can be efficiently recovered using group lasso methods.

Keywords. Union of Subspaces, Group Sparsity, Convex Optimization, Structured Sparsity, Compressed Sensing

1 Introduction

In many fields such as genetics, image processing, and machine learning, one is faced with the task of recovering very high dimensional signals from relatively few measurements. In general this is not possible, but fortunately many real world signals are, or can be transformed to be, sparse, meaning that only a small fraction of signal coefficients is non-zero. Compressed Sensing [5,10] allows us to recover sparse, high dimensional signals with very few measurements as compared to the ambient signal dimension. In fact, results indicate that one only needs $\mathcal{O}(s \cdot \log p)$ random measurements to exactly recover an s sparse signal of length p.

In many applications however, one not only has knowledge about the sparsity of the signal, but some additional information about the structure of the sparsity pattern as well:

- 1. In genetics, the genes are arranged into pathways/clusters, and genes belonging to the same pathway are often active/inactive in a group [31, 37].
- 2. In image processing, the wavelet transform coefficients can be modeled as belonging to a tree, with parent-child coefficients simultaneously being large or small [8,34].

- 3. In wideband spectrum sensing applications, the spectrum typically displays clusters of nonzero frequency coefficients, each corresponding to a narrowband transmission [26]
- 4. In applications in analog compressed sensing [15,16], the signals can be expressed as lying in a union of shift invariant subspaces.
- 5. In reconstruction of signals having a finite rate of innovation [11,40], non zeros are known to be clustered spatially, corresponding to objects in a scene for example.

In cases such as these, the sparsity pattern can be represented lying in a union of certain subspaces (e.g., coefficients in certain pathways, tree branches, frequency bands, or clusters). This knowledge about the signal structure can help further reduce the number of measurements one needs to exactly recover the signal. In this paper, we derive bounds on the number of random i.i.d. Gaussian measurements needed to exactly recover a sparse signal when its pattern of sparsity lies in a union of subspaces, based on solving a *convex* recovery algorithm. This characterization specializes to the latent group lasso, introduced in [20, 30, 32], wherein the sparsity pattern can be expressed as lying in a union of groups.

We analyze the recovery problem using a random Gaussian measurement model. We emphasize that although the derivation assumes the measurement matrix to be Gaussian, it can be extended to *any* subgaussian case, by paying a small constant penalty, as shown in [25]. We restrict ourselves to the Gaussian case here since it highlights the main ideas and keeps the analysis as simple as possible.

1.1 Prior Work

To the best of our knowledge, these results are new and distinct from prior theoretical characterizations of group lasso and the general union of subspace methods. Sampling theorems for unions of subspaces have been considered in [4, 24], where the authors show that the sample complexity depends logarithmically on the number of subspaces under consideration. In [14], the authors also propose a greedy scheme to recover signals that lie in such unions. The authors in [19] derive information theoretic bounds for the number of measurements needed for a variety of signal ensembles, including trees. In [2, 12], the authors show that one needs far fewer measurements when the signal can be expressed as lying in a union of subspaces, and explicit bounds are derived when using a modified version of CoSaMP [28] to recover the signal. Asymptotic consistency results are derived for the group lasso [42] when the groups partition the space of variables in [1]. Similarly, in [18], the authors again consider the groups to partition the space, and derive conditions for recovery using the group lasso. In [21, 22], the authors derive consistency results for the group lasso under arbitrary groupings of variables. Also, in [29], the authors consider overlapping groups and derive sample bounds. The authors in [20] derive consistency results in an asymptotic setting, for the group lasso with overlap, but do not provide exact recovery results. The group lasso with overlap (called the latent group lasso in [30]) is analyzed in detail in [30, 32].

Non-greedy schemes have not been developed to handle the case of the union of subspaces, nor the latent group lasso. Although group lasso with overlapping groups have been considered in the past, it yields vectors whose support can be expressed as a complement of a union of groups, while we consider cases where we require the support to lie in a union of groups, a distinction made in [20]. In the applications considered above, however, it is imperative that we recover a sparsity pattern that lies in a union of groups (or subspaces in more generality).

1.2 Our contributions

To derive our results, we appeal to the notion of restricted minimum singular values of an operator. The restricted minimum singular value (or equivalently the restricted eigenvalues of the gram matrix of the operator) conditions are weaker than the well known Restricted Isometry Conditions of operators, and have been used and studied in [3,38], among others.

We bound number of measurements needed for exact recovery with two terms. One term (kB) grows linearly in the total number of non-zero coefficients (with a constant of proportionality). This is close to the bare minimum of one measurement per non-zero component. Intuitively, this term corresponds to the magnitude estimation once the locations of the non zero components have been determined. The other term logarithmically depends on the number of subspaces under consideration and their relative orientations, and not the particulars of the subspaces (e.g., compositions, dimensions, extents, etc.). In particular, the subspaces need not be disjoint. Intuitively, this term corresponds to the price we pay for the detection of the non zeros in the signal (the active subspaces). The degree to which subspaces overlap, remarkably, has no effect on our bounds. In this regard, our bounds can be termed to be *universal*. This is somewhat surprising since overlapping subspaces are strongly coupled in the observations, tempting one to suppose that overlap may make recovery more challenging.

Our main theoretical result shows that for signals with support on k of M possible subspaces, exact recovery is possible from $C_1 * (\sqrt{2 \log(M-k)} + \sqrt{B})^2 k + C_2 * kB$ measurements using a latent group lasso type algorithm, B being the maximum subspace dimension. The constants C_1 and C_2 depend on the relative angle between subspaces, and will be explicitly derived in the sequel. Note that the bound depends on the sparsity s of the signal via the kB term. The latent group lasso reduces to a special case of our result. We will routinely compare the performance of the group lasso to the standard lasso, to study the effects of overlap between subspaces on the actual number of measurements needed to exactly recover a signal. For the lasso bound, we will use the one derived in [6]: $(2s+1)\log(p-s)$. Assuming that $M = \mathcal{O}(poly(p))$, our bound is roughly $k \log(p) + kB$. For the same problems, the lasso which ignores the group structure of the sparse signal components would require approximately $kB \log(p)$ measurements. Hence, taking advantage of the subspace structure will allow us to take fewer measurements to reconstruct the signal.

Note that in this work, the subspaces can be arbitrary, and we make *no* assumptions about their nature, except that they are known in advance. In short, we derive bounds for any generic union of subspaces, whether they overlap or form a partition of the ambient high dimensional space. Note that, when we say "do not overlap", we mean that the intersection of the subspaces is $\{0\}$, the all zeros vector.

We then propose a novel way to model wavelet coefficients of images in this framework, and show that we perform at least as well as several other state of the art methods in compressive imaging.

To summarize, our contributions in this paper are as follows:

- 1. We derive non asymptotic sample complexity bounds in a compressive-sensing framework when the measurement matrix is i.i.d. Gaussian, and the signal can be expressed as lying in a sparse union of finite dimensional subspaces.
- 2. We show that our bound holds regardless of the nature of overlap between subspaces. In this sense, the bounds we derive are universal.

- 3. We show that the group lasso with overlapping groups is a special case of the general result that we derive
- 4. We propose a new method to model wavelet coefficients of signals, so that one can use convex optimization algorithms to recover the signal exactly, or with high fidelity in the presence of noise.
- 5. An extensive series of experiments verify the theory.

The rest of this paper is organized as follows: In section 2 we set up the problem of recovering a signal lying in a union of subspaces, and present some preliminaries that will be useful in deriving the bounds. Our result for the sample complexity for the union of subspaces is derived in section 3. We then extend the framework for the case of group lasso with overlapping groups in section 4. Section 5 extends our results to approximately sparse signals. In section 6 we propose a novel framework for modeling wavelet transform coefficients of images, that makes use of the concepts we present in the sections that precede it. Experimental validation is provided in section 7. Finally, we conclude our paper in section 8, and present avenues for future work.

2 The Union of Subspaces Model

In the course of the next sections, we derive measurement bounds for the exact recovery of a signal lying in a sparse union of finite dimensional subspaces, and the robust recovery of one that is approximately sparse. In this section, we will argue as to why exact recovery of the signal corresponds to the minimization of the atomic norm of the signal, with the atoms obeying certain properties governed by the signal structure. Before we do so, we dispense with the notations.

2.1 Notations

Consider a signal of length p, that is s sparse. Note here that in case of multidimensional signals like images, we assume they are vectorized to have length p.

Suppose we are given a set of bases for M subspaces

$$\mathcal{K} = \{K_1, K_2, \dots, K_M\}$$

Let the dimensions of each subspace be given by d_1, d_2, \ldots, d_M , with $B = \max_i d_i$. So, $K_i \in \mathbb{R}^{p \times d_i}$. We assume that $span[K_1K_2 \ldots K_M] = \mathbb{R}^p$. Without loss of generality, assume each K_i to be orthonormal. If not, we can perform the Gram Schmidt procedure to orthonormalize them. The subspaces can be overlapping or non overlapping. When we say two subspaces do not overlap, we mean that

$$\operatorname{span}(K_i) \cap \operatorname{span}(K_i) = \{0\}$$

Also, the subspaces may or may not be perpendicular to each other. Two subspaces are perpendicular if

$$|\langle k_a, k_b \rangle| = 0 \quad \forall k_a \in A_i \text{ and } \forall k_b \in A_j \tag{1}$$

where

$$A_i = \operatorname{span}(K_i) \setminus (\operatorname{span}(K_i) \cap \operatorname{span}(K_j))$$

$$A_j = \operatorname{span}(K_j) \setminus (\operatorname{span}(K_i) \cap \operatorname{span}(K_j))$$

The notion of subspaces being perpendicular to each other is clarified in Fig. 1. It is fairly obvious that the more perpendicular the subspaces are to each other, the more separated they are, and hence the easier to is to distinguish between which among the subspaces is active.



Figure 1: Three subspaces, commonly encountered in subspace clustering methods. Subspaces S2 and S3 are perpendicular to each other. However, S1 and S2 are not perpendicular, nor are S1 and S3

We denote vectors by bold lowercase letters (a, v etc.), and matrices by bold uppercase letters (M, etc.). Subscripts following a vector denotes a particular index of the vector. For any vector v, we will routinely use the following decomposition:

$$oldsymbol{v} = \sum_{i=1}^M K_i oldsymbol{ar{v}}^i$$

where $\bar{v}^i \in \mathbb{R}^{d_i}$. The decomposition holds since we assume $\operatorname{span}([K_1K_2, \ldots, K_M]) = \mathbb{R}^p$. Superscripts following a matrix will denote the sub matrix whose columns are the columns indexed by the superscript. $|\cdot|$ denotes the cardinality of a set.

We let \boldsymbol{x}^{\star} be the (subspace sparse) signal to be recovered, whose non zero coefficients lie in k of the M subspaces $\mathcal{K}^{\star} \subset \mathcal{K}$, with $k \ll M$. Formally, noting that $\boldsymbol{x}^{\star} = \sum_{i} K_{i} \bar{\boldsymbol{x}}^{\star i}$,

$$\mathcal{K}^{\star} = \left\{ K_i \in \mathcal{K} : \|\bar{x}^{\star i}\| \neq 0 \right\}$$

We then have $|\mathcal{K}^{\star}| = k$. Let the indices of then active subspaces be given by J. That is,

$$J \subset \{1, 2, \dots, p\} : \{j \in J \iff K_j \in \mathcal{K}^\star\}$$

Later in the paper, we will also consider approximately sparse signals. We let $\Phi_{n \times p}$ be a measurement matrix consisting of i.i.d. Gaussian entries of mean 0 and variance $\frac{1}{n}$ so that every column is a realization of an i.i.d. Gaussian length n vector with covariance matrix $\frac{1}{n}I$. We denote the observed vector by $\boldsymbol{y} \in \mathbb{R}^n$: $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}^*$. The absence of a subscript following a norm $\|\cdot\|$ implies the ℓ_2 norm. The dual norm of $\|\cdot\|_p$ is denoted by $\|\cdot\|_p^*$. The convex hull of a set of points S is denoted by $\operatorname{conv}(S)$. We let $\sigma_i(\boldsymbol{M})$ denote the i^{th} singular value of a matrix \boldsymbol{M} . We define

$$\boldsymbol{K^{\star}} = [K_{j_1}K_{j_2}\ldots] \ \forall j_i \in J$$

and

and

$$\boldsymbol{K} = [K_1 K_2 \dots K_M]$$

Finally, let $\kappa(\mathbf{K})$ be the condition number of \mathbf{K} .

2.2 Atoms, Atomic Set and the Atomic Norm

To begin with, let us formalize the notion of atoms and the atomic norm of a signal (or vector). We will restrict our attention to signals in \mathbb{R}^p that can be expressed as lying in a sparse union of subspaces, though the same concepts can be extended to other spaces as well. We assume that $\boldsymbol{x} \in \mathbb{R}^p$ can be decomposed as :

$$oldsymbol{x} = \sum_{i=1}^k c_i oldsymbol{a}^i, \ \ c_i \ge 0$$

The vectors $\mathbf{a}^i \in \mathbb{R}^p$ are called *atoms*, and form the basic building blocks of any signal, which can be represented as a conic combination of the atoms. Note that the sum notation, rather than the integral notation, implies that only a countable number of coefficients can be non-zero. We denote $\mathcal{A} = \{\mathbf{a}\}$ to be the *atomic set*. Given a vector $\mathbf{x} \in \mathbb{R}^p$ and an atomic set, we define the *atomic norm* as

$$||\boldsymbol{x}||_{\mathcal{A}} = \inf\left\{\sum_{\boldsymbol{a}\in\mathcal{A}} c_{\boldsymbol{a}} : \boldsymbol{x} = \sum_{\boldsymbol{a}\in\mathcal{A}} c_{\boldsymbol{a}}\boldsymbol{a}, \ c_{\boldsymbol{a}} \ge 0 \ \forall \boldsymbol{a}\in\mathcal{A}\right\}$$
(2)

The atomic decomposition of the signal yields a representation of a signal in terms of some predefined atoms. Usually, few atoms used in a representation indicates a "simpler" representation. Hence, to obtain a "simple" representation of a vector, we look to minimize the atomic norm subject to constraints (3):

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{R}^p} ||\boldsymbol{x}||_{\mathcal{A}} \text{ s.t. } \boldsymbol{y} = \boldsymbol{\Phi} \boldsymbol{x}$$
 (3)

When the measurements are corrupted with noise θ , such that $\|\theta\| \leq \epsilon$, the atomic norm minimization problem becomes:

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{R}^p} \| \|\boldsymbol{x}\|_{\mathcal{A}} \text{ s.t. } \| \boldsymbol{y} - \boldsymbol{\Phi} \boldsymbol{x} \|^2 \le \epsilon^2$$
 (4)

Indeed, when the atoms are merely the canonical basis in \mathbb{R}^p , the atomic norm reduces to the standard ℓ_1 norm, and minimization of the atomic norm yields the well known *lasso* procedure [39].

Assuming we are aware of the subspaces \mathcal{K} , we now proceed to define the atomic set and the corresponding atomic norm for our framework. Let

$$A_{i} = \{ \boldsymbol{a} \in \mathbb{R}^{p} : \exists \boldsymbol{\alpha}^{i} : \boldsymbol{a} = K_{i} \boldsymbol{\alpha}^{i}, \quad \|\boldsymbol{a}\| = \|\boldsymbol{\alpha}^{i}\| = 1 \}$$
$$\mathcal{A} = \bigcup_{i=1}^{M} A_{i}$$
(5)

The sub vectors $\mathbf{a} \in A_i$ form the boundary of the unit sphere restricted to the span of K_i . The atomic norm of a vector $x \in \mathbb{R}^p$ is given by

$$\|\boldsymbol{x}\|_{\mathcal{A}} = \min_{\boldsymbol{x} = \sum_{\boldsymbol{a}} c_i \boldsymbol{a}} \sum c_i \quad c_i \ge 0$$
(6)

Lemma 2.1 The atomic norm for a signal $x \in \mathbb{R}^p$ lying in a union of subspaces is given by

$$\|oldsymbol{x}\|_{\mathcal{A}} = \min_{oldsymbol{x} = \sum_i K_i oldsymbol{lpha}^i} \sum_i \|oldsymbol{lpha}^i|$$

Proof The result follows by substituting

$$c_i = \|K_i \boldsymbol{\alpha}^i\| = \|\boldsymbol{\alpha}^i\|$$

and

$$\boldsymbol{a} = K_i \left(\frac{\boldsymbol{\alpha}_i}{\|K_i \boldsymbol{\alpha}^i\|} \right) = K_i \left(\frac{\boldsymbol{\alpha}_i}{\|\boldsymbol{\alpha}^i\|} \right)$$

The notion of the atomic set and the corresponding atomic norm for the union of subspaces model is made clear in Fig 2. The figure shows the atomic norm balls, given by the convex hull of the atomic set \mathcal{A} .



(a) Norm ball for Perpenducu- (b) Norm ball for Nonlar subspaces perpendicular subspaces

Figure 2: Atomic norm balls for a pair of perpendicular and non-perpendicular subspaces. The atomic set corresponds to the union of boundaries of the unit disks. Each disk corresponds to a particular A_i in (5)

Also note that we can directly compute the dual of the atomic norm from the set of atoms

$$\|\boldsymbol{x}\|_{\mathcal{A}}^{*} = \sup_{\boldsymbol{a}} \langle \boldsymbol{x}, \boldsymbol{a} \rangle$$
$$= \max_{i=1,2,\dots,M} \|(K_{i})^{T} \boldsymbol{x}\|$$
(7)

That is, the dual norm is the maximum over the norms of the projections of \boldsymbol{x} onto the different subspaces, noting that $[(K_i)^T K_i]^{-1} = \boldsymbol{I}$. The dual norm will be useful in our derivations below.

2.3 Gaussian Widths and Exact Recovery

Following [6], we define the *tangent cone* and *normal cone* at x^* with respect to $conv(\mathcal{A})$ under $||\boldsymbol{x}||_{\mathcal{A}}$ as [33]:

$$\mathcal{T}_{\mathcal{A}}(\boldsymbol{x}^{\star}) = \operatorname{cone}\{\boldsymbol{z} - \boldsymbol{x}^{\star} : ||\boldsymbol{z}||_{\mathcal{A}} \le |||\boldsymbol{x}^{\star}||_{\mathcal{A}}\}$$
(8)

$$\mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{\star}) = \{ \boldsymbol{u} : \langle \boldsymbol{u}, \boldsymbol{z} \rangle \leq 0, \quad \forall \boldsymbol{z} \in \mathcal{T}_{\mathcal{A}}(\boldsymbol{x}^{\star}) \}$$

$$= \{ \boldsymbol{u} : \langle \boldsymbol{u}, \boldsymbol{x}^{\star} \rangle = \gamma \| \boldsymbol{x} \|_{\mathcal{A}}$$

$$(9)$$

and
$$\|\boldsymbol{u}\|_{\mathcal{A}}^* \leq \gamma$$
 for some $\gamma \geq 0$ }

We note that, from [6] (Prop. 2.1), $\hat{\boldsymbol{x}} = \boldsymbol{x}^{\star}$ (3) is unique *iff*

$$\operatorname{null}(\boldsymbol{\Phi}) \cap \mathcal{T}_{\mathcal{A}}(\boldsymbol{x}^{\star}) = \{0\}$$
(10)

Hence, we require that the tangent cone at x^* intersects the nullspace of Φ only at the origin, to guarantee exact recovery.

Before we state the main recovery result from [6], we define the *Gaussian width* of a set:

Definition 2.2 Let \mathbb{S}^{p-1} denote the unit sphere in \mathbb{R}^p . The Gaussian width $\omega(S)$ of a set $S \in \mathbb{S}^{p-1}$ is

$$\omega(S) = \mathbb{E}_{\boldsymbol{g}} \left[\sup_{\boldsymbol{z} \in S} \boldsymbol{g}^T \boldsymbol{z} \right]$$

where $\boldsymbol{g} \sim \mathcal{N}(0, I)$ is an i.i.d. standard Gaussian vector.

Gordon used the Gaussian width to provide bounds on the probability that a random subspace of a certain dimension misses a subset of the sphere [17]. In [6], these results are specialized to the case of atomic norm recovery. In particular, we will make use of the following:

Proposition 2.3 [[6], Corollary 3.2] Let $\Phi : \mathbb{R}^p \to \mathbb{R}^n$ be a random map with i.i.d. zero-mean Gaussian entries having variance 1/n. Further let $\Omega = T_A(\mathbf{x}^*) \cap \mathbb{S}^{p-1}$ denote the spherical part of the tangent cone $T_A(\mathbf{x}^*)$. Suppose that we have measurements $\mathbf{y} = \Phi \mathbf{x}^*$, and we solve the convex program (3). Then \mathbf{x}^* is the unique optimum of (3) with high probability provided that

$$n \ge \omega(\Omega)^2 + \mathcal{O}(1)$$

To complete our problem setup we will also restate Proposition 3.6 in [6]:

Proposition 2.4 ([6], Proposition 3.6) Let C be any non-empty convex cone in \mathbb{R}^p , and let $g \sim \mathcal{N}(0, I)$ be a Gaussian vector. Then:

$$\omega(C \cap \mathbb{S}^{p-1}) \le \mathbb{E}_{\boldsymbol{g}}[\operatorname{dist}(\boldsymbol{g}, C^*)]$$
(11)

where dist(.,.) denotes the Euclidean distance between a point and a set, and C^* is the dual cone of C

We can then square (11) use Jensen's inequality to obtain

$$\omega(C \cap \mathbb{S}^{p-1})^2 \le \mathbb{E}_g[\operatorname{dist}(\boldsymbol{g}, C^*)^2]$$
(12)

We note here that the dual cone of the tangent cone is the normal cone, and vice-versa.

Thus, to derive measurement bounds, we only need to calculate the square of the Gaussian width of the intersection of the tangent cone at x^* with respect to the atomic norm and the unit sphere. This value can be bounded by the distance of a Gaussian random vector to the normal cone at the same point, as implied by (12). In the next section, we derive bounds on this quantity.

3 Gaussian Width of the Normal Cone for Unions of Subspaces

For generic subspaces \mathcal{K} , we have

$$\boldsymbol{c} \in \mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{\star}) \iff \exists \gamma \geq 0 \quad : \langle \boldsymbol{c}, \boldsymbol{x}^{\star} \rangle = \gamma \| \boldsymbol{x}^{\star} \|_{\mathcal{A}}, \\ \| \bar{\boldsymbol{c}}^{i} \| = \gamma \text{ if } K_{i} \in \mathcal{K}^{\star}, \quad \| \bar{\boldsymbol{c}}^{i} \| \leq \gamma \text{ if } K_{i} \notin \mathcal{K}^{\star}.$$

$$(13)$$

We now prove the main result of this paper, a sufficient number of Gaussian measurements needed to recover a signal lying in a union of subspaces:

Theorem 3.1 To exactly recover a k-subspace sparse signal decomposed into M subspaces in \mathbb{R}^p ,

$$2\frac{\sigma_p^2(\boldsymbol{K}^{\star})}{\sigma_p^2(\boldsymbol{K})}(\sqrt{2\log(M-k)}+\sqrt{B})^2k+2kB\kappa^2(\boldsymbol{K})$$

i.i.d. Gaussian measurements are sufficient.

To prove this result, we need two lemmas:

Lemma 3.2 Let q_1, \ldots, aq_L be L, χ -squared random variables with d-degrees of freedom. Then

$$\mathbb{E}[\max_{1 \le i \le L} q_i] \le (\sqrt{2\log(L)} + \sqrt{d})^2.$$

We defer the proof to appendix 9.1.

Lemma 3.3 Suppose $v \in \mathbb{R}^p$ is supported on some set of groups $\mathcal{K}^* \subset \mathcal{K}$. Then,

$$\|oldsymbol{v}\| \leq \sqrt{|\mathcal{K}^{\star}|} \hspace{0.2cm} \sigma_p(oldsymbol{K}^{\star}) \hspace{0.2cm} \|oldsymbol{v}\|_{\mathcal{A}}^{st}.$$

We defer the proof of this lemma to appendix 9.2

Proof [Proof of Theorem 3.1]

Intuition: Note that, from (12), the Gaussian width of the intersection of the tangent cone at x^* with the unit sphere is bounded above by the expected euclidean distance between a random Gaussian vector and the normal cone at x^* (13). We can further bound this distance by the distance between a random Gaussian vector g and a particular vector $r \in \mathcal{N}_{\mathcal{A}}(x^*)$, as shown in (14). We proceed to construct such a vector r and prove the result

$$\mathbb{E}_{g}[\operatorname{dist}(\boldsymbol{g}, C^{*})^{2}] \leq \mathbb{E}_{g}[\operatorname{dist}(\boldsymbol{g}, \boldsymbol{r})^{2}], \quad \boldsymbol{r} \in \mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{*})$$
(14)

Now, since we assume span(\mathcal{K}) = \mathbb{R}^p , we can write, for any vector $\boldsymbol{\eta}$,

$$oldsymbol{\eta} = oldsymbol{K}oldsymbol{ar{\eta}} \;\; ar{oldsymbol{\eta}} \in \mathbb{R}^{\sum d_i}$$

So, we have for x^* ,

$$oldsymbol{x}^{\star} = oldsymbol{K}oldsymbol{x}^{\star} = oldsymbol{K}^Soldsymbol{x}^{\star S^c} + oldsymbol{K}^{S^c}oldsymbol{x}^{\star S^c}$$

where $S \subset \{1, 2, \dots, \sum_i d_i\}$ is the indices of active coefficients of \bar{x}^{\star} . It is important to note that $\forall i \notin J, x^{\star i}$ is a sub vector of $x^{\star S^c}$ Also, note here that $K^{S^c} \bar{x}^{\star S^c} = 0$.

Since the normal cone is nonempty, there exists a $\boldsymbol{v} \in \mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{\star})$ with $\|\boldsymbol{v}\|_{\mathcal{A}}^{\star} = 1$ and $\boldsymbol{v}^{i} = 0 \ \forall i \notin J$. Since v is in the normal cone, it will also satisfy $\langle v, x^* \rangle = \|x^*\|_{\mathcal{A}}$. We will use this v in our bound below.

Suppose $\boldsymbol{w} \sim \mathcal{N}(0, \boldsymbol{I}_p)$ is a vector with i.i.d. Gaussian entries. We then have

$$egin{array}{lll} egin{array}{lll} egin{array}{llll} egin{array}{lll} egin{array}{lll} egin{arr$$

Let $t(\boldsymbol{w}) = \max_{i \notin J} \|\boldsymbol{w}^i\|$. since $\boldsymbol{w} = \sum_{i=i}^M K_i \bar{\boldsymbol{w}}^i$, we have $\bar{\boldsymbol{w}}^i = \boldsymbol{K}_i^T (\boldsymbol{K}\boldsymbol{K}^T)^{-1} \boldsymbol{w}$, giving us $\bar{\boldsymbol{w}}^i \sim \mathcal{N}(0, \boldsymbol{K}_i^T (\boldsymbol{K}\boldsymbol{K}^T)^{-2} \boldsymbol{K}_i)$. This means that

$$\|\bar{\boldsymbol{w}}^{i}\|^{2} \sim \|\left(\boldsymbol{K}\boldsymbol{K}^{T}\right)^{-2}\|\chi_{d_{i}}^{2}$$

$$\tag{15}$$

So, $\|\bar{w}^i\|^2$ is a scaled χ^2 random variable with d_i degrees of freedom. The scaling factor is merely $\sigma_p^{-2}(\boldsymbol{K}).$

Let us now construct a vector $\boldsymbol{r} \in \mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{\star})$. We can write, as for \boldsymbol{w}

$$r = K^S \bar{r}^S + K^{S^c} \bar{r}^{S^c}$$

Now let $\bar{\boldsymbol{r}}^{S} = t(\boldsymbol{w})\bar{\boldsymbol{v}}^{S}$, and $\bar{\boldsymbol{r}}^{S^{c}} = \bar{\boldsymbol{w}}^{S^{c}}$

From (13), and from our definition of $t(\boldsymbol{w})$, we have $\boldsymbol{r} \in \mathcal{N}_{\mathcal{A}}(x^{\star})$. Referring to (12), we now consider the expected squared distance between $\mathcal{N}_{\mathcal{A}}(x^{\star})$ and w:

$$\begin{split} \mathbb{E}[\operatorname{dist}(\boldsymbol{w}, C^*)^2] \\ &\leq \mathbb{E}[||\boldsymbol{r} - \boldsymbol{w}||^2] \\ &= \mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{r}}^S + \boldsymbol{K}^{S^c} \bar{\boldsymbol{r}}^{S^c} - \boldsymbol{K}^S \bar{\boldsymbol{w}}^S + \boldsymbol{K}^{S^c} \bar{\boldsymbol{w}}^{S^c}\right\|^2\right] \\ \stackrel{(i)}{=} \mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{r}}^S - \boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2 + \left\|\boldsymbol{K}^{S^c} \bar{\boldsymbol{r}}^{S^c} - \boldsymbol{K}^{S^c} \bar{\boldsymbol{w}}^{S^c}\right\|^2\right] \\ &= \mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{r}}^S - \boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ \stackrel{(ii)}{=} 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{r}}^S\right\|^2\right] + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ &= 2\mathbb{E}\left[\left\|\boldsymbol{K}^S t(\boldsymbol{w}) \bar{\boldsymbol{v}}^S\right\|^2\right] + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ \stackrel{(iii)}{=} 2\mathbb{E}[t(\boldsymbol{w})^2] \left\|\boldsymbol{K} \bar{\boldsymbol{v}}\right\|^2 + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ &= 2\mathbb{E}[t(\boldsymbol{w})^2] \left\|\boldsymbol{k} \bar{\boldsymbol{v}}\right\|^2 + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ &= 2\mathbb{E}[t(\boldsymbol{w})^2] \left\|\boldsymbol{v}\right\|^2 + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ \stackrel{(iv)}{=} 2\mathbb{E}[t(\boldsymbol{w})^2] \left\|\boldsymbol{v}\right\|^2 + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ &\stackrel{(iv)}{\leq} 2k\left(\frac{\sigma_p(\boldsymbol{K}^{\star})}{\sigma_p(\boldsymbol{K})}\right)^2 (\sqrt{2\log(M-k)} + \sqrt{B})^2 + 2\mathbb{E}\left[\left\|\boldsymbol{K}^S \bar{\boldsymbol{w}}^S\right\|^2\right] \\ \stackrel{(vi)}{\leq} 2k\left(\frac{\sigma_p(\boldsymbol{K}^{\star})}{\sigma_p(\boldsymbol{K})}\right)^2 (\sqrt{2\log(M-k)} + \sqrt{B})^2 + 2kB\kappa^2(\boldsymbol{K}) \end{split}$$

Where (i) trivially follows because the indices in S and S^c are disjoint, (ii) follows from the result $||a - b||^2 \le 2(||a||^2 + ||b||^2)$ (iii) follows from the fact that v is deterministic. (iv) follows from the fact that \bar{v} is only supported on S (v) follows from Lemma 3.2, Lemma 3.3 and (15). Finally, (vi) follows from bounding the last term as shown in appendix 9.3, and noting that $|S| \le kB$.

3.1 Remarks

- 1. The most important thing to note from our result is that we pay no extra penalty in terms of the number of measurements needed when the subspaces overlap. Hence, we term our result "universal".
- 2. The kB term in the bound is an upper-bound on the signal sparsity. In the case of highly overlapping subspaces, this value may be much larger than the signal sparsity, but such cases seldom arise in real-world applications. If the subspace dimensions are vastly different, then it is pessimistic to bound the quantity with the maximum dimension B, but this yields a simple expression for the measurements needed. It is of course possible to obtain tighter bounds using the techniques in our work for cases where the groups are of varying sizes.
- 3. It can be seen from Theorem 3.1 that the number of measurements is linear in k and B. Hence, the number of measurements that are sufficient for signal recovery grows linearly with the number of active subspaces in the signal, and also the maximum subspace dimension. This can be seen analogous to the linear dependence of the lasso bound on the sparsity s of the signal.

- 4. We note that although we pay no extra price to measure the signal when there is significant overlap between subspaces, there is an additional cost in the recovery process of the signal, in that the subspaces need to first be separated by replication of the coefficients [20], or resort to a primal-dual method to solve the problem [27].
- 5. In the bound we get, $\sigma_p(\mathbf{K})$ captures the price we pay if the subspaces are non perpendicular. Indeed, if subspaces are nearly aligned with each other, then it becomes nearly impossible to distinguish between them. This is reflected by the fact that $\sigma_p(\mathbf{K}) \to 0$ as the subspaces become more aligned. Fig. 3(a) shows this phenomenon for the case of two subspaces of one dimension each that become more and more aligned with each other. Similarly, in the second term, the condition number of the matrix \mathbf{K} , $\kappa(\mathbf{K})$ is determined by the angle between subspaces. The more the subspaces are close to each other, the higher the condition number, and subsequently the more measurements we need. This can be seen from Fig. 3(b).



(a) $\sigma_p(\mathbf{K})$ as the angle between 2 subspaces (b) $\kappa(\mathbf{K})$ as the angle between 2 subspaces is varies varied ied

Figure 3: As the two subspaces become more separated $(\theta \to 90^{\circ})$, both the quantities approach 1. As $(\theta \to 0^{\circ})$, $\sigma_p(\mathbf{K}) \to 0$ and $\kappa(\mathbf{K}) \to \infty$, indicating that it becomes impossible to distinguish between active subspaces

3.2 Perpendicular Subspaces

When the subspaces are perpendicular to each other, as defined in (1), we can make the bound we obtained much tighter. To see this, note that when the subspaces are perpendicular, equation (ii) in the proof of Theorem 3.1 can be replaced by

$$\mathbb{E}[\|oldsymbol{K}^Soldsymbol{ar{r}}^S\|^2] + \mathbb{E}[\|oldsymbol{K}^Soldsymbol{ar{w}}^S\|^2]$$

This follows since, in the case of perpendicular subspaces, \bar{r}^S is independent of \bar{w}^S . Also, since the subspaces are perpendicular, we have

$$\sigma_p(\mathbf{K}^{\star}) = \sigma_p(\mathbf{K}) = \kappa(\mathbf{K}) = 1$$

Substituting the values in the bound we get from Theorem 3.1, we have

$$\mathbb{E}[\operatorname{dist}(\boldsymbol{w}, C^*)] \le k(\sqrt{2\log(M-k)} + \sqrt{B})^2 + kB$$
(16)

for perpendicular subspaces. This is a much smaller quantity than the general result we obtained in Theorem 3.1, underscoring the fact that recovery becomes easier as the subspaces get more and more separated. In the next section, we show that the group lasso with overlapping groups (also called the latent group lasso) is a special case of recovery when subspaces are perpendicular.

4 The Group Lasso with Overlapping Groups

The group lasso with overlapping groups, [20, 30, 32] can be formulated as an atomic norm minimization problem. In the group lasso problem, we are given a set of groups $\mathcal{G} = \{G_1, G_2, \ldots, G_M\}$, $G_i \subset \{1, 2, \ldots, p\}$ and we wish to recover a k-group sparse vector \boldsymbol{x}^* from compressive measurements. In this case, we can define the subspaces $K_i \in \mathcal{K}$ as follows:

$$\forall i \in \{1, 2, \dots, M\}, K_i = \boldsymbol{I}^{G_i} \tag{17}$$

where I^{G_i} is the sub matrix of the identity matrix, consisting of columns indexed by the group G_i .

We now show that minimizing the atomic norm under the atomic set arising out of these subspaces yields the group lasso with overlapping groups. Note that, under the definition of the subspaces as in (17), and referring to (5), we have that A_i is merely the unit sphere restricted to the dimensions indexed by group G_i .

Lemma 4.1 Suppose the atomic set is given as in (5). Let $K_i = \mathbf{I}^{G_i} \quad \forall i = \{1, 2, \dots, M\}$, where $G_i \subset \{1, 2, \dots, p\}$. Then,

$$|m{x}||_{\mathcal{A}} = \Omega^{\mathcal{G}}_{overlap}(m{x})$$

where $\Omega_{overlap}^{\mathcal{G}}(\boldsymbol{x})$ is the overlapping group lasso norm defined in [20].

Proof In (6), we can substitute $\mathbf{v}_G = c_G \mathbf{a}$, giving us $c_G = |c_G| \cdot ||\mathbf{a}|| = ||c_G \mathbf{a}|| = ||\mathbf{v}_G||$. Hence,

$$\begin{aligned} ||\boldsymbol{x}||_{\mathcal{A}} &= \inf \left\{ \sum_{\boldsymbol{a} \in \mathcal{A}} c_{\boldsymbol{a}} : \boldsymbol{x} = \sum_{\boldsymbol{a} \in \mathcal{A}} c_{\boldsymbol{a}} \boldsymbol{a} \ c_{\boldsymbol{a}} \ge 0 \ \forall \boldsymbol{a} \in \mathcal{A} \right\} \\ &= \inf \left\{ \sum_{G \in \mathcal{G}} ||\boldsymbol{v}_{G}|| \ : \boldsymbol{x} = \sum_{G \in \mathcal{G}} \boldsymbol{v}_{G} \right\} \\ &= \Omega_{overlap}^{\mathcal{G}}(\boldsymbol{x}) \end{aligned}$$

Corollary 4.2 Under the atomic set defined in (5), when $K_i = \mathbf{I}^{G_i} \quad \forall i = \{1, 2, \dots, M\}$,

$$||m{x}||_{\mathcal{A}} = \sum_{G\in\mathcal{G}} ||m{x}_G||$$

Proof $\Omega_{overlap}^{\mathcal{G}}(\boldsymbol{x}) = \sum_{G \in \mathcal{G}} ||\boldsymbol{x}_G||$ in the non overlapping case.

Thus, (3) yields:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x} \in \mathbb{R}^p}{\operatorname{argmin}} \quad \Omega^{\mathcal{G}}_{overlap}(\boldsymbol{x}) \quad \text{s.t.} \quad \boldsymbol{y} = \boldsymbol{\Phi} \boldsymbol{x}$$
(18)

which can be solved using [20].

It is not hard to see that, in the case of disjoint groups,

$$\mathcal{N}_{\mathcal{A}}(\boldsymbol{x}^{\star}) = \{ \boldsymbol{z} \in \mathbb{R}^{p} : \boldsymbol{z}_{i} = \gamma \frac{(\boldsymbol{x}^{\star})_{i}}{||\boldsymbol{x}_{G}^{\star}||} \quad \forall G \in \mathcal{G}^{\star},$$

$$||\boldsymbol{z}_{G}|| \leq \gamma \quad \forall G \notin \mathcal{G}^{\star}, \gamma \geq 0 \}$$

$$(19)$$

However, in the case of overlapping groups, as in the case of generic subspaces, no such closed form exists.

Under the group sparsity model defined above, it is not hard to see that $\sigma_p(\mathbf{K}^*) = \kappa(\mathbf{K}) = 1$. Also, note that the subspaces are perpendicular to each other since the basis vectors of each subspace is aligned with one or more of the coordinate axes in \mathbb{R}^p . As a consequence, we obtain the following result:

Theorem 4.3 To exactly recover a k-group sparse signal decomposed into M groups in \mathbb{R}^p , the following is a sufficient number of Gaussian measurements needed:

$$(\sqrt{2\log(M-k)} + \sqrt{B})^2k + kB$$

4.1 Remark: Comparison with the lasso

We compare the group lasso bound we obtain to the standard lasso measurement bound:

$$(2s+1)\log(p-s) \tag{20}$$

The bound we obtain in Theorem 4.3 can be upper bounded by

$$2k\max\{2\log(M), B\} + kB \tag{21}$$

Noting that $s \leq kB$ with equality when the groups do not overlap. In this case, (21) evaluates to

$$\frac{2s}{B}\max\{2\log(M), B\} + s$$
$$= (2s+1)\frac{\max\{2\log(M), B\}}{B}$$

which is smaller than the lasso bound (20) by a factor of roughly $\frac{\log(M)}{B\log(p)}$. So, in most cases, our bound shows that the we can perform better than the conventional lasso by exploiting the additional group structured information that is available.

5 Approximately Sparse Signals

The result we proved in Theorem 3.1 apply to exact reconstruction of k- subspace sparse signals. In many cases however, the signals are not exactly subspace (or group) sparse, but approximately so. Specifically, in a context we are especially interested in, the ordered wavelet coefficients of natural images exponentially decay to zero, and hence can be modeled as *approximately* sparse, the "sparsity" meaning that only few coefficients are of significant magnitude.

In such cases, we can model the approximately sparse signal f^{\star} as

$$oldsymbol{f}^{\star} = oldsymbol{x}^{\star} + oldsymbol{h}^{\star}$$

where x^* is a k- subspace sparse approximation of f^* , retaining the k subspaces having largest norm and h^* corresponds to the remaining coefficients that are small in magnitude. Clearly, we can bound $\|h^*\|$ above by some constant, say c_h .

Now, measuring the approximately sparse signal f^{\star} using a Gaussian measurement matrix amounts to

$$egin{aligned} \Phi oldsymbol{f}^\star &= \Phi oldsymbol{x}^\star + \Phi oldsymbol{h}^\star \ &= oldsymbol{\Phi} oldsymbol{x}^\star + oldsymbol{ heta} \end{aligned}$$

Since the norm of Φ is bounded, we can write

$$\|\boldsymbol{\theta}\| \leq \delta$$

The results we have obtained thus far can be easily extended to the case where we obtain such bounded noisy observations. In the noisy case, we observe

$$\boldsymbol{y} = \boldsymbol{\Phi} \boldsymbol{x}^{\star} + \boldsymbol{\theta}, \ \|\boldsymbol{\theta}\| \leq \delta$$

We then solve the atomic norm minimization problem, with a relaxed constraint to take into account the bounded noise:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x} \in \mathbb{R}^p}{\operatorname{argmin}} ||\boldsymbol{x}||_{\mathcal{A}} \quad \text{s.t.} \quad ||\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{x}|| \le \delta$$
(22)

We restate corollary 3.3 from [6]:

Proposition 5.1 [[6], Corollary 3.3] Let $\mathbf{\Phi} : \mathbb{R}^p \to \mathbb{R}^n$ be a random map with i.i.d. zero-mean Gaussian entries having variance 1/n. Further let $\Omega = T_{\mathcal{A}}(\mathbf{x}^*) \cap \mathbb{S}^{p-1}$ denote the spherical part of the tangent cone $T_{\mathcal{A}}(\mathbf{x}^*)$. Suppose that we have measurements $\mathbf{y} = \mathbf{\Phi}\mathbf{x}^* + \mathbf{\theta}$, and $\|\mathbf{\theta}\| \leq \delta$. Suppose we solve the convex program (22). Let $\hat{\mathbf{x}}$ denote the optimum of (22). Also, suppose $\|\mathbf{\Phi}\mathbf{z}\| \geq \epsilon \|\mathbf{z}\| \quad \forall \mathbf{z} \in T_{\mathcal{A}}(\mathbf{x}^*)$. Then $\|\mathbf{x}^* - \hat{\mathbf{x}}\| \leq \frac{2\delta}{\epsilon}$ with high probability provided that

$$n \ge \frac{\omega(\Omega)^2}{(1-\epsilon)^2} + \mathcal{O}(1).$$

Substituting the result of Theorem 3.1 in Proposition 5.1, we have the following corollary yielding a sufficient condition to accurately recover a signal when the measurements are corrupted with bounded noise:

Corollary 5.2 Suppose we wish to recover a signal that lies in k out of M arbitrarily defined subspaces, such that the maximum subspace dimension is B. Let the set of active subspaces be denoted by \mathcal{K}^* . Let $\hat{\boldsymbol{x}}$ be the optimum of the convex program (22). To have $\|\hat{\boldsymbol{x}} - \boldsymbol{x}^*\| \leq \frac{2\delta}{\epsilon}$ with high probability,

$$\frac{2\frac{\sigma_p^2(\boldsymbol{K}^{\star})}{\sigma_p^2(\boldsymbol{K})}(\sqrt{2\log(M-k)}+\sqrt{B})^2k+2kB\kappa^2(\boldsymbol{K})}{(1-\epsilon)^2}$$

i.i.d. Gaussian measurements are sufficient.

Note that we merely need to set $\sigma_p(\mathbf{K}^*) = \kappa(\mathbf{K}) = 1$, and remove the '2' from both terms to obtain the corresponding result for the latent group lasso:

$$\frac{(\sqrt{2\log(M-k)} + \sqrt{B})^2k + kB}{(1-\epsilon)^2}$$

6 Compressive Imaging with Group Sparsity

We consider the compressive imaging problem, that is to recover an image from a small number of random measurements. Here "small" is used relative to the ambient dimension of the image. The standard lasso [39] formulation is given by

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{x}\|^2 + \lambda \|\boldsymbol{x}\|_1$$
(23)

The ℓ_1 norm acts as a surrogate for the sparsity of the signal. The lasso aims to recover a signal that is sparse, by setting most coefficients of \boldsymbol{x} to be zero. For the exact recovery case, the lasso problem is equivalent to the Basis Pursuit [7]

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{s.t. } \boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}$$
 (24)

The lasso penalty reflects the fact that the wavelet coefficients are approximately sparse, but in reality not all patterns of sparsity are equally plausible/probable. For example, Fig (4(b)) shows the DWT coefficients of the barbara image, and Fig. (4(c)) shows the same coefficients, but randomly scrambled. Clearly, the ℓ_1 norm of both sets of coefficients will be the same. This shows that the lasso penalty in itself is invariant to any structure present in the sparse coefficients.

To model this structure that is inherently present between wavelet transform coefficients of images, [8,13,34] propose making use of graphical models such as Hidden Markov Trees. HMT's, while providing good performance in image denoising applications (where $\Phi = I$) in (23), cannot provide acceptable reconstruction for other, more general inverse problems. This is because the presence of a (non identity) sensing matrix Φ (randomly) mixes up the coefficients for every measurement y_i obtained.

To overcome this mixing between the coefficients, many alternatives have been proposed. [35] propose using a version of loopy belief propagation to solve the recovery problem. The authors in [2, 12] generalize the notion of restricted isometry properties to signals that lie in unions of subspaces, and use a modified version of CoSAMP [28] to solve the inverse problem. Greedy and/or suboptimal iterative reconstruction schemes are used in [13,23]. Finally, the authors in [36] propose modeling the coefficients using an HMT, and using the Approximate Message Passing algorithm [9] to solve the compressed sensing problem.



Figure 4: The ℓ_1 norms of both (a) and (b) are exactly equal, since they do not take structure into account

ized

All the works mentioned above sacrifice the recovery guarantees and the easy analysis that convex optimization algorithms provide, for the sake of modeling the dependencies between DWT coefficients. This motivates out problem: can we on the one hand model the dependencies among wavelet transform coefficients, while at the same time propose to solve a convex optimization problem similar to (23)?

To this end, we model the parent-child coefficients into groups. Parent-child pairs of wavelet transform coefficients across scales and at similar locations tend to be simultaneously high or low. Hence, we can take advantage of this dependency and use group lasso methods to recover the image. Fig 5 shows a representative example.



Figure 5: Quadtree corresponding to the 2-d DWT. At each scale, parent coefficients can be grouped with child coefficients.

We wish to recover the non zero coefficients lying on the wavelet tree shown in Fig 5. When coefficients are modeled into groups, one can use the group lasso [42] to recover the coefficients

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{x}\|^2 + \lambda \sum_{i=1}^{M} \|\boldsymbol{x}_{G_i}\|$$
(25)

where x_{G_i} is the vector x whose coefficients not indexed by group G_i are set to zero. The

group lasso as shown in (25) suffers from a drawback however. It was recently argued in [20] that the sparsity pattern recovered by the group lasso can be expressed as a complement of a union of groups. One look at Fig. 5 tells us that we are interested in the recovery of sparsity patterns that can be expressed as a union of (overlapping) groups. To this end, the authors in [20] propose the latent group lasso [30,32]

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{x}\|^2 + \lambda \Omega_{overlap}^{\mathcal{G}}(\boldsymbol{x})$$
(26)

where $\Omega^{\mathcal{G}}_{overlap}(\boldsymbol{x})$ is the latent group lasso norm.

The latent group lasso lends itself well to the sort of problems we are concerned about in this paper. For a thorough analysis of the various properties of the latent group lasso penalty, we refer the interested reader to [30]. In the next section, we show how the latent group lasso can be effectively used to recover images, when we model the DWT coefficients to lie in parent-child groups.

7 Experiments

In this section we aim to show two things:

- 1. The bound we derived in Theorems 3.1 and 4.3 holds for a wide variety of cases, and is invariant to the grouping observed in the signal
- 2. By modeling the DWT coefficients of images (and 1D signals of course) into parent-child groups, we can recover the signal efficiently and exactly/robustly

Also, henceforth we refer to the latent group lasso method as Glasso.

7.1 Sampling Bounds for Subspace Sparse Signals

We extensively tested our method against the standard lasso procedure. In the case where the groups overlap, we use the replication method outlined in [20], to reduce the optimization problem to that of non overlapping groups. We compare the number of measurements needed for our method with that needed for the lasso. For the lasso, it would be instructive to keep in mind the bound derived in [6], viz. $(2s + 1)\log(p - s)$. In the case of non overlapping groups, the bound evaluates to $(2kB + 1)\log(kM - kB)$. We generate length p = 2000 signals, made up of M = 100 non-overlapping groups of size B = 20. We set k = 5 groups to be "active", and the values within the groups are drawn from a uniform [0, 1] distribution. The active groups are assigned uniformly at random. The sparsity of the signal will thus be s = 100

We use SpaRSA [41] for the lasso and the group lasso with overlap, learning λ over a grid. Fig. 6 displays the mean reconstruction error $||\hat{x} - x^*||_2^2/p$ as a function of the number of random measurements taken. The errors have been averaged over 100 tests, and each time a new random signal was generated with the above mentioned parameters.

From the parameters considered, we conclude that ≈ 380 measurements are sufficient to recover the signal. When we have 380 measurements, the lasso does not recover the signal exactly, as seen in Fig 6, but the latent group lasso does.

To show that the bound we compute holds regardless of the complexity of groupings, we consider the following scenario: Suppose we have M = 100 groups, each of size B = 40. k = 5 of those



Figure 6: The group lasso (red) compared with the lasso (blue). The vertical line indicates our bound. Note that our bound (380) predicts exact recovery of the signal, while at the same value, the lasso does not recover the signal

groups are active, and the values within each group are assigned from a uniform [-1, 1] distribution. We arrange these groups in three configurations:

- 1. The groups do not overlap, yielding a signal of length p = 4000, and signal sparsity s = 200.
- 2. A partial overlapping scenario, where apart from the first and last group, every group has 20 elements in common with a group above it, and 20 common with the group below, giving $p = 2020, s \in [120, 200]$ depending on which of the 100 groups are active.
- 3. A random overlap case where the first 50 groups are non overlapping and the remaining 50 are assigned uniformly at random from the existing p = 2000 indices. $s \le 200$ in this case.

The scenarios we consider are depicted in Fig. 7. In each of the cases, we compute the bound to be ≈ 630 . The bound becomes looser as the complexity of the groupings increases. This, as argued before, is a result of the bound for the signal sparsity becoming looser.



Figure 7: Types of groupings considered. Each set of coefficients encompassed by one color belongs to one group.

We can see from Fig. 8(a) that our group lasso bound (≈ 630) holds for all cases. For the sake of comparison, we considered the lasso performance on the signals in cases (1) - (3) as well, and these are plotted in Fig. 8(b). From the values of p and s computed for the three cases, we have the corresponding bounds for the lasso [6] to be 3305 for the no overlap case (1), [1819, 3010] for the partial overlap case (2) and 3000 for case (3).



(a) performance of the group lasso on cases con- (b) performance of the lasso on cases considered sidered in Figure 7. Note that our bound evaluates in Figure 7.to 630, clearly sufficient measurements to recover the signal in all cases.



We consider exact recovery of the wavelet transform coefficients of the "blocks" signal (Fig. 9(a)). We group the wavelet transform coefficients into parent child pairs as outlined in Section 6. In this case, for a p = 16384 length signal, we have M = 16382 groups, and the maximum group size is B = 2. We use the Haar wavelet bases to decompose the image. Fig. 9(b) shows the reconstruction obtained from 1690 measurements, corresponding to the bound computed for k = 47. Of course, we can compute k since we have the original signal with us. We see that our bound yields a sufficient number of measurements for exact recovery.

Our final experiment outlines the relationship between the number of measurements taken and the size of the problem. We generated test signals that were group sparse, with each active group having coefficients selected randomly from a uniform $\mathcal{U}[-1,1]$ distribution. We fix the group size B to be 6. We consider two cases:

- The non overlapping case (1), and
- The partial overlapping case (2)

Fig. 10 shows the probability of error as the number of measurements increases. In the figure, note that we show the total number of groups (M) in the signal. For each M, we fix the group sparsity level k to be M/10. The results are averaged over 100 tests, and the probability of error is computed empirically. It can be seen in Fig. 10 that, regardless of the groups overlapping or not, we need roughly the same number of measurements to achieve a low probability of error.



Figure 9: Exact reconstruction of a length 16384 signal from 1690 measurements in the wavelet domain



Figure 10: Number of groups vs Number of measurements for exact signal reconstruction. The color bar indicates the probability of error, computed empirically. Notice how the number of measurements needed to achieve low probability of error is nearly the same in both cases, highlighting that we indeed do not pay a penalty for complicated grouping strategies

7.2 Modeling DWT coefficients into groups

In the spirit of [2,36], we considered a 128×128 section of the cameraman image, and obtained 5000 iid gaussian measurements from it. No noise was added to the image. We compare our methods with the ones displayed in [36]. Fig 11(a) has been taken directly from [36], and our result is shown in Fig. 11(b).

For the basis of comparison, we zoom into similar regions from the best performing methods in Fig. 11(a) and Fig. 11(b), in Fig. 12. It can be seen that our method performs comparably to the turbo-BG and turbo-GM methods.

Along similar lines, we tested our methods for noiseless image recovery using the Microsoft Research Object Class Recognition database¹. The dataset consists of images categorized into 20

¹http://research.microsoft.com/en-us/projects/ObjectClassRecognition



(a) Reconstruction Performance of Various Methods



(b) Reconstruction using Glasso

Figure 11: Reconstruction of a section of the cameraman image using various methods.



(a) MCMC

(b) Turbo-BG

(c) Turbo-GM

(d) Glasso

Figure 12: comparison of our method against the state of the art. The figures shown are zoomed in versions of those in Fig. 11

types, with roughly 30 images of each type. We used the first 10 images of each type to generate a training set of 200 images, which was used to learn the regularization parameters. To compare and contrast our results with other methods tested in [36] (Fig. 6)², we compute the Normalized Mean Square Error (NMSE) of our methods over the same dataset for comparison. The normalized mean square error (in dB) for the true image \boldsymbol{x} is given by $10 \times \log \left(\frac{\|\hat{\boldsymbol{x}}-\boldsymbol{x}\|^2}{\|\boldsymbol{x}\|^2}\right)$. We resized the images to size 128×128 , and obtained 5000 measurements for each case.



Figure 13: comparison of various methods, and Glasso. Note that, the lower the value of NMSE, the better the performance.



Figure 14: Comparison of the two methods in the presence of noise.

Fig. 14 shows the results we obtain as a function of the noise standard deviation. For the purpose of the experiment, we consider piecewise constant signals of length 1024, having 5 jumps. The location of the jumps is chosen at random, and the magnitude of each "piece" is chosen uniformly between [-1,1]. We take 256 measurements for both the lasso and group lasso. From the figure, it is clear that by modeling the wavelet coefficients into parent-child pairs, we can better reconstruct signals in the presence of noise. The results are averaged over 1000 randomly generated signals.

²The authors thank Subhojit Som and Phil Schniter for sharing data for Fig. 13

8 Conclusions and Discussion

In this paper, we showed that one can recover a signal known to lie in a sparse union of subspaces exactly using atomic norm minimization algorithms. The number of measurements needed depend on the number of active subspaces, their dimensions and the relative angles between subspaces. We also showed that the measurement bound we derived is universal, in that it holds regardless of the nature and specific structure of the subspaces, in terms of overlaps. Indeed, the bounds can be specialized to cases where one is interested in a specific type of grouping. We subsequently extended these results for signals lying in a generic union of groups.

We also proposed a novel modeling strategy for DWT coefficients of signals. By modeling the coefficients into parent-child groups, we were able to take advantage of convex recovery methods that provably guarantee signal recovery. We showed that our method is at least as good as the current state of the art in non adaptive compressed sensing.

We note here that we do not claim the optimality of the particular grouping method that we have used , *viz.* grouping the parent child pairs together. How best to group wavelet coefficients is still an open question, and is an avenue for further research. We prefer the parent-child pairs for its simplicity, and due to the fact that the groupings yield acceptable results, as seen in Section 7.

9 Appendices

9.1 Proof of Lemma 3.2

Proof Let $M_L := \max_{1 \le i \le L} q_i$. For t > 0, we have that

$$\mathbb{E}[M_L] = \frac{\log[\exp(t \cdot \mathbb{E}[M_L])]}{t}$$

$$\stackrel{(a)}{\leq} \frac{\log[\mathbb{E}[\exp(t \cdot M_L)]]}{t}$$

$$\stackrel{(b)}{=} \frac{\log[\mathbb{E}[\max_{1 \le j \le L} \exp(t \cdot q_j)]]}{t}$$

$$\stackrel{(c)}{\leq} \frac{\log[L\mathbb{E}[\exp(t \cdot q_1)]]}{t}$$

$$= \frac{\log(L) - \frac{d}{2}\log(1 - 2t)}{t}$$

Where (a) follows from Jensen's inequality, (b) follows from the monotonicity of the exponential function, and (c) merely bounds the maximum by the sum over all the elements. Now, setting $t = [(2+2\epsilon)]^{-1}$ with $\epsilon = \sqrt{\frac{d}{2\log(L)}}$ yields

$$\mathbb{E}[M_L] \le (\sqrt{2\log(L)} + \sqrt{d})^2$$

Note that t can be optimized depending on the application. We use this particular choice because it makes no assumptions about the relative magnitudes of (M - k) and B.

9.2 Proof of Lemma 3.3

Proof Let $K^* = [K_j]_{j \in J^*}$. By duality, it suffices to show that $||z||_{\mathcal{A}} \leq \sqrt{|\mathcal{K}^*|} \sigma_p(\mathbf{K}^*) ||z||$ for all z with $supp(z) \subset \mathcal{K}^*$. For any such z, there exists a representation $z = \mathcal{K}^* b$ such that

$$z = [K_{j1}K_{j2}\dots K_{j|J|}][b^1b^2\dots b^{|\mathcal{K}^*|}]^T$$
$$= \sum_{K\in\mathcal{K}^*} K_i b^i$$

so that none of the supports of b^i overlap. It then follows that

$$\begin{aligned} \|z\|_{\mathcal{A}} &= \|\sum_{K \in \mathcal{K}^{\star}} K_{i} b^{i}\| \\ &\stackrel{(i)}{\leq} \sum_{K \in \mathcal{K}^{\star}} \|b^{i}\| \\ &\stackrel{(ii)}{\leq} \sqrt{|\mathcal{K}^{\star}|} \left(\sum_{K \in \mathcal{K}^{\star}} \|b_{K}\|^{2}\right)^{1/2} \\ &= \sqrt{|\mathcal{K}^{\star}|} \|b\| \\ &\stackrel{(iii)}{\leq} \sqrt{|\mathcal{K}^{\star}|} \ \sigma_{p}(\mathbf{K}^{\star}) \ \|\mathbf{z}\| \end{aligned}$$

Where (i) follows from the definition of the norm $\|\cdot\|_{\mathcal{A}}$, (ii) is a consequence of the relation $\|\beta\|_1 \leq \sqrt{k}\|\beta\|_2$ for k dimensional vectors β and (iii) follows from minimizing $\|b\|$ subject to $z = K^*b$

9.3 Bounding the last term in the proof of Theorem 3.1

Proof First, note that we can write

$$\begin{split} \mathbb{E}\left[\|\boldsymbol{K}^{S}\bar{\boldsymbol{w}}^{S}\|^{2}\right] &= \mathbb{E}\left[\|\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\bar{\boldsymbol{w}}\|^{2}\right] \\ &= \mathbb{E}\left[\|\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{w}\|^{2}\right] \end{split}$$

Where $P_{S}(\cdot)$ is the operator that projects (·) onto the space spanned by the indices in S. Now,

$$\begin{split} & \mathbb{E}\left[\|\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{w}\|^{2}\right] \\ &= \mathbb{E}\left[\|\boldsymbol{w}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}^{T}\boldsymbol{K}^{T}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{w}\|^{2}\right] \\ &= \mathbb{E}\left[tr\left(\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{w}\boldsymbol{w}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}^{T}\boldsymbol{K}^{T}\right)\right] \\ &= tr\left(\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\mathbb{E}[\boldsymbol{w}\boldsymbol{w}^{T}](\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}^{T}\boldsymbol{K}^{T}\right) \\ &\stackrel{(a)}{=}tr\left(\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}^{T}\boldsymbol{K}^{T}\right) \\ &= \|(\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T}\|_{F}^{2} \end{split}$$

where (a) is a consequence w being a standard Gaussian vector. Now, let the singular value decomposition of K be $K = U\Sigma V^T$. Then,

$$\begin{aligned} \| (\boldsymbol{K}\boldsymbol{K}^{T})^{-1}\boldsymbol{K}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{K}^{T} \|_{F}^{2} \\ &= \| (\boldsymbol{U}\boldsymbol{\Sigma}^{2}\boldsymbol{U}^{T})^{-1}\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{T}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{U}^{T} \|_{F}^{2} \\ &= \| \boldsymbol{U}^{-T}\boldsymbol{\Sigma}^{-2}\boldsymbol{\Sigma}\boldsymbol{V}^{T}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{U}^{T} \|_{F}^{2} \\ &= \| \boldsymbol{\Sigma}^{-1}\boldsymbol{V}^{T}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{V}\boldsymbol{\Sigma} \|_{F}^{2} \\ &= \| \boldsymbol{\Sigma}^{-1} \|^{2} \| \boldsymbol{V}^{T}\boldsymbol{P}_{\boldsymbol{S}}\boldsymbol{V} \|_{F}^{2} \| \boldsymbol{\Sigma} \|^{2} \\ &= \kappa(\boldsymbol{K})^{2} \| \boldsymbol{P}_{\boldsymbol{S}} \|_{F}^{2} \\ &= \kappa(\boldsymbol{K})^{2} |S| \end{aligned}$$

Acknowledgments

This work was partially supported by AFOSR grant FA9550-09-1-0140 and the DARPA KECOM Program, and by ONR Award N00014-11-1-0723.

References

- F. Bach. Consistency of the group lasso and multiple kernel learning. Journal of Machine Learning Research, 9:1179–1225, June 2008.
- [2] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Transactions on Information Theory*, 2010.
- [3] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of the lasso and dantzig selector. Annals of Statistics, 37(4):1705–1732, 2009.
- [4] T. Blumensath and M. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE transactions on Information Theory*, 55(4):1872–1882, Apr 2009.
- [5] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Information Theory*, 52:489–509, 2006.
- [6] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. preprint arXiv:1012.0621v1, 2010.
- [7] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. SIAM Review, 43(1):129– 159, 2009.
- [8] M. S. Crouse, R. D. Nowak, and R. G. Baraniuk. Wavelet based statistical signal processing using hidden markov models. *Transactions on Signal Processing*, 46(4):886–902, 1998.
- [9] D. Donoho, A. Maleki, and A. Montanari. Message passing algorithms for compressed sensing. *National Academy of Sciences*, 2009.
- [10] D. L. Donoho. Compressed sensing. IEEE Trans. Information Theory, 52:1289–1306, 2006.
- [11] P. Dragotti, M. Vetterli, and T. Blu. Sampling moments and reconstructing signals of nite rate of innovation: Shannon meets strangfix. *IEEE transactions on Signal Processing*, 55(5):1741–1757, May 2007.

- [12] M. F. Duarte, V. Cevher, and R. G. Baraniuk. Model-based compressive sensing for signal ensembles. Allerton, 2009.
- [13] M. F. Duarte, M. B. Wakin, and R. G. Baraniuk. Wavelet-domain compressive signal reconstruction using a hidden markov tree model. *Proceedings of the IEEE International Conference on Acoustics*, Speech, and Signal Processing (ICASSP), pages 5137–5140, Mar 2008.
- [14] Y. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE transactions on Information Theory*, 55(11):5302–5316, Nov 2009.
- [15] Y. C. Eldar. Compressed sensing of analog signals in shift-invariant spaces. IEEE transactions on Signal Processing, 57(8):2986–2997, Aug 2009.
- [16] Y. C. Eldar. Uncertainty relations for shift-invariant analog signals. IEEE transactions on Information Theory, 55(12):5742–5757, Dec 2009.
- [17] Y. Gordon. On Milman's inequality and random subspaces which escape through a mesh in \mathbb{R}^n . Geometric aspects of functional analysis, Isr. Semin., 1317:84–106, 1986 87.
- [18] J. Huang and T. Zhang. The benefit of group sparsity. Technical report, arXiv:0901.2962. Preprint available at http://arxiv.org/pdf/0903.2962v2, May 2009.
- [19] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. Technical report, arXiv:0903.3002. Preprint available at http://arxiv.org/pdf/0903.3002v2, May 2009.
- [20] L. Jacob, G. Obozinski, and J. P. Vert. Group lasso with overlap and graph lasso. *Proceedings of the* 26th International Conference on machine Learning, 2009.
- [21] R. Jenatton, J. Audibert, and F. Bach. Structured variable selection with sparsity inducing norms. Technical report, arXiv:0904.3523. Preprint available at http://arxiv.org/pdf/0904.3523v3, Sep 2009.
- [22] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *Technical report, arXiv:1009.3139. submitted*, 2010.
- [23] C. La and M. N. Do. Tree based orthogonal matching pursuit algorithm for signal reconstruction. IEEE International Conference on Image Processing, Atlanta, GA., pages 1277 – 1280, Oct 2006.
- [24] Y. M. LU and M. N. Do. A theory for sampling signals from a union of subspaces, *IEEE transactions on Signal Processing*, 56(6):2334–2345, 2008.
- [25] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis*, 17(4):1248 – 1282, 2006.
- [26] M. Mishali and Y. Eldar. Blind multi-band signal reconstruction: compressed sensing for analog signals. IEEE Trans. Signal Processing, 57(30):993–1009, March 2009.
- [27] S. Mosci, S. Villa, A. Verri, and L. Rosasco. A primal-dual algorithm for group sparse regularization with overlapping groups. *Neural Information Processing Systems*, 2010.
- [28] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. Appl. Comput. Harmon. Anal., 26:301–321, 2008.
- [29] S. Negahban, P. Ravikumar, M. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Preprint ArXiv :1010:2731v1*, October 2010.
- [30] G. Obozinski, L. Jacob, and J. Vert. Group lasso with overlaps: The latent group lasso approach. Preprint arXiv:1110.0413v1 [stat.ML], Oct 2011.
- [31] F. Parvaresh, H. Vikalo, S. Misra, and B. Hassibi. Recovering sparse signals using sparse measurement matrices in compressed dna microarrays. *IEEE Journal of Selected Topics in Signal Processing*, 2(3):275– 285, June 2008.

- [32] D. Percival. Theoretical properties of the overlapping groups lasso. Preprint arXiv:1103.4614v2 [stat.ML], Nov 2011.
- [33] T. Rockafellar and J. B. Wets. Variational analysis. Springer Series of Comprehensive Studies in Mathematics, 317, 1997.
- [34] J. K. Romberg, H. Choi, and R. G. Baraniuk. Bayesian tree structured image modeling using wavelet domain hidden markov models. *Transactions on Image Processing*, March 2000.
- [35] P. Schniter. Turbo reconstruction of structured sparse signals. Proc. Conference on Information Sciences and Systems, Mar 2010.
- [36] S. Som and P. Schniter. Compressive imaging using approximate message passing and a markov-tree prior. *IEEE transactions on signal processing*, 2011.
- [37] A. Subramanian et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression proles. *National Academy of Sciences*, 102:1554515550, 2005.
- [38] G. Tang and A. Nehorai. Performance analysis of sparse recovery based on constrained minimal singular values. *IEEE Transactions of Signal Processing*, to appear.
- [39] R. Tibshirani. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B, pages 267–288, 1996.
- [40] M. V. Vetterli, P. Marziliano, and T. Blu. Sampling signals with nite rate of innovation. *IEEE trans-actions on Signal Processing*, 50:1417–1428, Jun 2002.
- [41] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. *Transactions on Signal Processing*, 57:2479–2493, 2009.
- [42] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. Journal of the royal statistical society. Series B, 68:49–67, 2006.