

Structured Sparsity: Discrete and Convex approaches

Anastasios Kyrillidis, Luca Baldassarre, Marwa El-Halabi, Quoc Tran-Dinh and Volkan Cevher
EPFL, Lausanne, Switzerland.

February 12, 2016

Abstract

Compressive sensing (CS) exploits sparsity to recover sparse or compressible signals from dimensionality reducing, non-adaptive sensing mechanisms. Sparsity is also used to enhance interpretability in machine learning and statistics applications: While the ambient dimension is vast in modern data analysis problems, the relevant information therein typically resides in a much lower dimensional space. However, many solutions proposed nowadays do not leverage the true underlying structure. Recent results in CS extend the simple sparsity idea to more sophisticated *structured* sparsity models, which describe the interdependency between the nonzero components of a signal, allowing to increase the interpretability of the results and lead to better recovery performance. In order to better understand the impact of structured sparsity, in this chapter we analyze the connections between the discrete models and their convex relaxations, highlighting their relative advantages. We start with the general group sparse model and then elaborate on two important special cases: the dispersive and the hierarchical models. For each, we present the models in their discrete nature, discuss how to solve the ensuing discrete problems and then describe convex relaxations. We also consider more general structures as defined by set functions and present their convex proxies. Further, we discuss efficient optimization solutions for structured sparsity problems and illustrate structured sparsity in action via three applications.

Contents

1	Introduction	1
1.1	Why structured sparsity?	1
2	Preliminaries	2
3	Sparse group models	3
3.1	The discrete model	5
3.2	Convex approaches	5
3.3	Extensions	6
4	Sparse dispersive models	7
4.1	The discrete model	7
4.2	Convex approaches	8
5	Hierarchical sparse models	9
5.1	The discrete model	9
5.2	Convex approaches	10
6	Submodular models	11
6.1	The discrete model	12
6.2	Convex approaches	12
6.3	Examples	13
7	Optimization	14
7.1	Greedy and discrete approaches	16
7.1.1	Projected gradient descent and matching pursuit variants	16
7.1.2	Submodular function minimization	16
7.2	Convex approaches	16
7.2.1	Proximity methods	16
7.2.2	Primal-dual and alternating minimization approaches	17

8 Applications	17
8.1 Compressive Imaging	17
8.1.1 Results	19
8.2 Neuronal spike detection from compressed data	19
8.3 Digital Confocal Imaging	22
8.3.1 Results	24

1 Introduction

Information in many natural and man-made signals can be exactly represented or well approximated by a sparse set of nonzero coefficients in an appropriate basis [90]. Compressive sensing (CS) [24,40,120] exploits this fact to recover signals from their compressive samples through dimensionality reducing, non-adaptive sensing mechanisms. From a different perspective, sparsity is used to enhance *interpretability* in machine learning [53,58] and statistics [123] applications: While the ambient dimension is vast in modern data analysis problems, the relevant information therein typically resides in a much lower dimensional space. This conclusion has led to several new theoretical and algorithmic developments in different communities, including theoretical computer science [52], neuronal imaging [57,74], gene expression inference [108,121], bioinformatics [112,141] and X-ray crystallography [23,70,94]. In all these disciplines, given a data set to analyze, one is usually interested in the *simplest* model that well explains the observations.

Apart from the rather mature theory describing the sparse signal approximation problem, its success lies also in the computational tractability of such task. In the CS regime, Donoho [40] and Candes et al. [24] utilize simple convex optimization—i.e., linear programming solvers or quadratic programming methods—and matrix isometry assumptions over sparsely restricted sets to compute a sparse approximation from a limited set of linear measurements in polynomial time. Alternatively and in contrast to the conventional convex relaxation approaches, iterative greedy algorithms [46,80,99] rely on the discrete structure of the sparse signal approximation problem to find a solution and are characterized by identical approximation guarantees as in [24,40]. Moreover, their overall complexity matches, if not improves, the computational cost of convex approaches.

1.1 Why structured sparsity?

While many of the optimization solutions proposed nowadays result in *sparse* model selections, in many cases they do not capture the true underlying structure to best explain the observations [8]. In fact, this *uninformed* selection process has been the center of attention in sparse approximation theory [81,138], because it not only prevents interpretability of results in many problems, but also fails to exploit key prior information that could radically improve estimation performance.

Recent results in CS extend the simple sparsity idea to consider more sophisticated *structured* sparsity models, which describe the interdependency between the nonzero coefficients, increase the interpretability of the results and lead to better recovery performance [8,15,44,111]. Nowadays, we have witnessed aplenty elaborate approaches that guide the selection process: (overlapping) group Lasso, fused Lasso, greedy approaches for signal approximation under tree-structure assumptions just to name a few; see [47,73,124,137]. To show the merits of such approaches, consider the problem of image recovery from a *limited* set of measurements using the tree-structured group sparse model [6]. Figure 1 shows the performance of structured sparsity models, compared to simple ones.

Moreover, such a priori model-based assumptions result into more robust solutions and allow recovery with far fewer samples, e.g. $O(k)$ samples for k sparse signals whose non-zero coefficients are arranged into few blocks or form a rooted connected subtree of a given tree over the coefficients [8]. To highlight the importance of this property, in the case of Magnetic Resonance Imaging (MRI), reducing the total number of measurements is highly desirable for both capturing functional activities within small time periods and rendering the whole procedure less “painful” for the patient [89].

2 Preliminaries

Plain lowercase and uppercase letters represent scalars while boldface lowercase letters represent vectors. We reserve boldface uppercase letters for matrices. Calligraphic uppercase letters are mostly used to denote sets. The i -th entry of a vector \mathbf{w} is denoted as w_i . We use superscripts such as \mathbf{w}^i to denote the estimate at the i -th iteration of an algorithm. The sign of a scalar $\alpha \in \mathbb{R}$ is given by $\text{sign}(\alpha)$.

Given a set $\mathcal{S} \subseteq \mathcal{N} := \{1, \dots, n\}$, the complement \mathcal{S}^c is defined with respect to \mathcal{N} , and its cardinality as $|\mathcal{S}|$. The support set of \mathbf{w} is $\text{supp}(\mathbf{w}) = \{i \in \mathcal{N} : w_i \neq 0\}$. Given a vector $\mathbf{w} \in \mathbb{R}^n$, $\mathbf{w}_{\mathcal{S}}$ is the projection (in \mathbb{R}^n) of \mathbf{w} onto \mathcal{S} , i.e. $(\mathbf{w}_{\mathcal{S}})_{\mathcal{S}^c} = \mathbf{0}$, whereas $\mathbf{w}|_{\mathcal{S}} \in \mathbb{R}^{|\mathcal{S}|}$ is \mathbf{w} limited to \mathcal{S} entries. We define $\Sigma_k := \{\mathbf{w} : \mathbf{w} \in \mathbb{R}^n, |\text{supp}(\mathbf{w})| \leq k\}$ as the set of all k -sparse vectors in n -dimensions; with a slight abuse of notation, we use Σ_k to further denote the set of k -sparse supports in \mathcal{N} . We sometimes write $\mathbf{x} \in \Sigma_k$ to mean $\text{supp}(\mathbf{x}) \in \Sigma_k$.




Figure 1: Empirical performance of simple and structured sparsity recovery on natural images. In all cases, the number of linear measurements are 5% (top row) and 10% (bottom row) of the actual image dimensions. **Left panel:** Original images of dimension: (Top row) 2048×2048 , (Bottom row) 512×512 . **Middle panel:** Conventional recovery using simple sparsity model. **Right panel:** Tree-structured sparse recovery.

We use \mathbb{B}^n to represent the space of n -dimensional binary vectors and define $\iota : \mathbb{R}^n \rightarrow \mathbb{B}^n$ to be the indicator function of the nonzero components of a vector in \mathbb{R}^n , i.e., $\iota(\mathbf{x})_i = 1$ if $x_i \neq 0$ and $\iota(\mathbf{x})_i = 0$, otherwise. We let $\mathbb{1}_n$ to be the n -dimensional vector of all ones, $\mathbb{1}_{n,S}$ the n -dimensional vector of all ones projected onto S and \mathbf{I}_n the $n \times n$ identity matrix; we often use \mathbf{I} when dimension is clear from the context.

Norms: For completeness, we repeat some of the norm definitions presented in the Introduction chapter of this book. We define the ℓ_p^n -norm in n -dimensions as:

$$\|\mathbf{x}\|_p = \begin{cases} (\sum_{i=1}^n |x_i|^p)^{1/p} & \text{if } p \in (0, \infty), \\ \max_i |x_i| & \text{if } p = \infty. \end{cases}$$

The ℓ_0 pseudo-norm is defined as: $\|\mathbf{x}\|_0 := |\text{supp}(\mathbf{x})|$.

Projections operations: Given a discrete set \mathcal{M}_k with sparsity level k and an anchor point $\mathbf{x} \in \mathbb{R}^n$, a key operation in our subsequent discussions is the following projection problem:

$$\mathcal{P}_{\mathcal{M}_k}(\mathbf{x}) \in \underset{\mathbf{w} \in \mathbb{R}^n}{\text{argmin}} \{ \|\mathbf{w} - \mathbf{x}\|_2^2 \mid \text{supp}(\mathbf{w}) \in \mathcal{M}_k \}. \quad (1)$$

Proximity operations: Consider $g(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a regularizer function. We define the proximity operator of $g(\cdot)$ as [32, eq. (2.13)]:

$$\text{prox}_\lambda^g(\mathbf{x}) := \underset{\mathbf{w} \in \mathbb{R}^n}{\text{argmin}} \left\{ \frac{1}{2} \|\mathbf{w} - \mathbf{x}\|_2^2 + \lambda \cdot g(\mathbf{w}) \right\}, \quad (2)$$

where $\lambda > 0$ is the regularizer weight.

Optimization preliminaries: Existing algorithmic solutions invariably rely on two structural assumptions on the objective function that particularly stand out among many others: the *Lipschitz continuous gradient* assumption and the *strong regularity* condition.




Figure 4: Consider the group structure \mathfrak{G}^1 defined by the following groups, $\mathcal{G}_1 = \{1\}$, $\mathcal{G}_2 = \{2\}$, $\mathcal{G}_3 = \{1, 2, 3, 4, 5\}$, $\mathcal{G}_4 = \{4, 6\}$, $\mathcal{G}_5 = \{3, 5, 7\}$ and $\mathcal{G}_6 = \{6, 7, 8\}$.




Figure 5: Bipartite group graph with loops induced by the group structure \mathfrak{G}^1 described in Figure 4, where on each edge we report the elements of the intersection.

Given the above and for a user-defined group budget $G \in \mathbb{Z}_+$, we define the group model \mathcal{M}_G as $\mathcal{M}_G := \{\bigcup_{\mathcal{G}_\ell \in \mathcal{J}} \mathcal{G}_\ell, \mathcal{J} \subseteq \mathfrak{G}, |\mathcal{J}| \leq G\}$, that is all sets of indexes that are the union of at most G groups from the collection \mathfrak{G} . Then, the projection problem (1) becomes:

$$\hat{\mathbf{x}} =: \mathcal{P}_{\mathcal{M}_G}(\mathbf{x}) \in \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \{ \|\mathbf{x} - \mathbf{z}\|_2^2 \mid \operatorname{supp}(\mathbf{z}) \in \mathcal{M}_G \}. \quad (3)$$

Moreover, one might be only interested in identifying the *group-support* of the approximation $\hat{\mathbf{x}}$, that is the G groups that constitute its support. We call this the *group-sparse model selection* problem.

3.1 The discrete model

According to (3), we search for $\hat{\mathbf{x}} \in \mathbb{R}^n$ such that $\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2$ is minimized, while $\hat{\mathbf{x}}$ does not exceed a given group budget G . A useful notion in the group sparse model is that of the **group ℓ_0 -“norm”**:

$$\|\mathbf{w}\|_{\mathfrak{G}, 0} := \min_{\boldsymbol{\omega} \in \mathbb{B}^M} \left\{ \sum_{j=1}^M \omega_j \mid \mathbf{A}\boldsymbol{\omega} \geq \iota(\mathbf{w}) \right\}; \quad (4)$$

here, \mathbf{A} denotes the adjacency matrix as defined in the previous subsection, the binary vector $\boldsymbol{\omega}$ indicates which groups are active and the constraint $\mathbf{A}\boldsymbol{\omega} \geq \iota(\mathbf{w})$ makes sure that, for every non-zero component of \mathbf{w} , there is at least one active group that covers it. Given the above definitions, the group-based signal approximation problem (3) can be reformulated as

$$\hat{\mathbf{x}} \in \underset{\mathbf{w} \in \mathbb{R}^N}{\operatorname{argmin}} \{ \|\mathbf{w} - \mathbf{x}\|_2^2 \mid \|\mathbf{w}\|_{\mathfrak{G}, 0} \leq G \}. \quad (5)$$

One can easily observe that, in the case where we already know the group cover of the approximation $\hat{\mathbf{x}}$, we can obtain $\hat{\mathbf{x}}$ as $\hat{\mathbf{x}}_{\mathcal{J}} = \mathbf{x}_{\mathcal{J}}$ and $\hat{\mathbf{x}}_{\mathcal{J}^c} = 0$, where $\mathcal{J} = \bigcup_{\mathcal{G}_\ell \in \mathcal{S}^G(\hat{\mathbf{x}})} \mathcal{G}_\ell$, with $\mathcal{S}^G(\hat{\mathbf{x}})$ denoting the group support of $\hat{\mathbf{x}}$ and $\mathcal{J}^c = \mathcal{N} \setminus \mathcal{J}$. I.e., if we know the group support of the solution, the entries' values are naturally given by the




Figure 6: Consider $\mathcal{G}_1 = \{1, 2, 3\}$, $\mathcal{G}_2 = \{3, 4, 5\}$, $\mathcal{G}_3 = \{5, 6, 7\}$, which can be represented by the graph in the left plot. If \mathcal{G}_3 were to include an element from \mathcal{G}_1 , for example $\{2\}$, we would have the loopy graph of the right plot. Note that \mathfrak{G}^1 is pairwise overlapping, but not loopless, since $\mathcal{G}_3, \mathcal{G}_4, \mathcal{G}_5$ and \mathcal{G}_6 form a loop.

anchor point \mathbf{x} . The authors in [5] prove that the group support can be obtained by solving a discrete problem, according to the next lemma.

Lemma 3.1 [5] *Given $\mathbf{x} \in \mathbb{R}^N$ and a group structure \mathfrak{G} , the group support of the solution $\hat{\mathbf{x}}$ —denoted as $\mathcal{S}^G(\hat{\mathbf{x}}) = \{\mathcal{G}_j \in \mathfrak{G} : \omega_j^G = 1\}$ —is given by the solution $(\boldsymbol{\omega}^G, \mathbf{y}^G)$ of the following binary maximization problem:*

$$\max_{\boldsymbol{\omega} \in \mathbb{B}^M, \mathbf{y} \in \mathbb{B}^n} \left\{ \sum_{i=1}^N y_i x_i^2 : \mathbf{A}\boldsymbol{\omega} \geq \mathbf{y}, \sum_{j=1}^M \omega_j \leq G \right\}. \quad (6)$$

Here, \mathbf{y} denotes the selected variables, while $\boldsymbol{\omega}$ the active groups. Thus, the constraint $\mathbf{A}\boldsymbol{\omega} \geq \mathbf{y}$ guarantees that, for every selected variable, there is at least one group that covers it.

The problem in (6) can produce all the instances of the weighted maximum coverage problem (WMC) [100], where the weights for each element are given by x_i^2 ($1 \leq i \leq n$) and the index sets are given by the groups $\mathcal{G}_j \in \mathfrak{G}$ ($1 \leq j \leq M$). Since WMC is in general NP-hard [100] and given Lemma 1, finding the groups that cover the solution $\hat{\mathbf{x}}$ is in general NP-hard.

However, it is possible to approximate the solution of (6) using the greedy WMC algorithm [101]. At each iteration, the algorithm selects the group that covers new variables with maximum combined weight until G groups have been selected. To the best of our knowledge, only the work in [5] presents a polynomial time algorithm for solving (6) for loopless pairwise overlapping groups structures, using dynamic programming. The algorithm gradually explores the group graph, which in this case is a tree or a forest of trees, from the leaves towards the root, storing the best solutions found so far and dynamically updating them until the entire tree is examined.

3.2 Convex approaches

Recent works in compressive sensing and machine learning with group sparsity have mainly focused on leveraging the group structures for lowering the number of samples required for recovering signals [8, 15, 44, 67, 69, 108, 111, 119].

For the special case of non-overlapping groups, dubbed as the block-sparsity model, the problem of model selection does not present computational difficulties and features a well-understood theory [119]. The first convex relaxations for group-sparse approximation [137] considered only non-overlapping groups: the authors proposed the `Group LARS` (**L**east **A**nge **R**egre**S**sion) algorithm to solve this problem, a natural extension of simple sparsity LARS algorithm [42]. Using the same algorithmic principles, its extension to overlapping groups [140] has the drawback of selecting supports defined as the complement of a union of groups, even though it is possible to engineer the groups in order to favor certain sparsity patterns over others [73]. Eldar et al. [44] consider the union of subspaces framework and cast the model selection problem as a block-sparse model selection one by duplicating the variables that belong to overlaps between the groups, which is the optimization approach proposed also in [69]. Moreover, [44] considers a model-based pursuit approach [31] as potential solver for this problem, based on a predefined model \mathcal{M}_k . For these cases, one uses the group lasso norm

$$\sum_{\mathcal{G} \in \mathfrak{G}} \|\mathbf{x}_{|\mathcal{G}}\|_p, \quad (7)$$

where $\mathbf{x}_{|\mathcal{G}}$ is the restriction of \mathbf{x} to only the components indexed by \mathcal{G} .

In addition, convex proxies to the group ℓ_0 -norm (4) have been proposed (e.g., [69]) for finding group-sparse approximations of signals. Given a group structure \mathfrak{G} , an example generalization is defined as

$$\|\mathbf{x}\|_{\mathfrak{G}, \{1, p\}} := \inf_{\substack{\mathbf{v}_1, \dots, \mathbf{v}_M \in \mathbb{R}^n \\ \forall i, \text{supp}(\mathbf{v}_i) = \mathcal{G}_i}} \left\{ \sum_{i=1}^M d_i \|\mathbf{v}_i\|_p \mid \sum_{i=1}^M \mathbf{v}_i = \mathbf{x} \right\}, \quad (8)$$

where $\|\mathbf{x}\|_p = \left(\sum_{i=1}^N x_i^p \right)^{1/p}$ is the ℓ_p -norm, and d_j are positive weights that can be designed to favor certain groups over others [108]. This norm can be seen as a weighted instance of the atomic norm described in [30, 111], based on the seminal work of [31] on the sparse synthesis model. There, the authors leverage convex optimization for signal recovery, but not for model selection: Let $\mathcal{A} := \{\mathbf{a}_1, \mathbf{a}_2, \dots \mid \mathbf{a}_i \in \mathbb{R}^n, \forall i\}$ be the *atomic set of group sparse signals* \mathbf{x} that can be synthesized as a k -sparse combination of atoms \mathbf{a}_i in \mathcal{A} , i.e.,

$$\mathbf{x} = \sum_{i=1}^k c_i \mathbf{a}_i, \quad c_i \geq 0, \mathbf{a}_i \in \mathcal{A} \text{ and } \|\mathbf{c}\|_0 \leq k. \quad (9)$$

Here, $\mathbf{a}_i \in \mathbb{R}^n$, $\text{supp}(\mathbf{a}_i) = \mathcal{G}_i$ and $\|\mathbf{a}_i\|_2 = 1$. We can then define the atomic norm:

$$\|\mathbf{x}\|_{\mathcal{A}} = \inf \left\{ \sum_{i=1}^{|\mathcal{A}|} c_i \mid \mathbf{x} = \sum_{i=1}^{|\mathcal{A}|} c_i \mathbf{a}_i, c_i \geq 0, \forall \mathbf{a}_i \in \mathcal{A} \right\}. \quad (10)$$

Lemma 3.2 [30, 111] *If in (8) the weights are all equal to 1 ($d_i = 1, \forall i$), we have*

$$\|\mathbf{x}\|_{\mathcal{A}} = \|\mathbf{x}\|_{\mathfrak{G},\{1,p\}}.$$

The group-norm (8) can also be viewed as the tightest convex relaxation of a particular set function related to the *weighted set-cover* (see Section 6 and [107]).

One can in general use (8) to find a group-sparse approximation under the chosen group norm

$$\hat{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^N} \left\{ \|\mathbf{w} - \mathbf{x}\|_2^2 : \|\mathbf{w}\|_{\mathfrak{G},\{1,p\}} \leq \lambda \right\}, \quad (11)$$

where $\lambda > 0$ controls the trade-off between approximation accuracy and group-sparsity. However, solving (11) does not necessarily yield a group-support for $\hat{\mathbf{x}}$: even though we can recover one through the decomposition $\{\mathbf{v}^j\}$ used to compute $\|\hat{\mathbf{x}}\|_{\mathfrak{G},\{1,p\}}$, it may not be unique and when it is unique it may not capture the minimal group-cover of \mathbf{x} [108]. Therefore, the equivalence of ℓ_0 and ℓ_1 minimization [24, 40] in the standard compressive sensing setting does not generally hold in the overlapping group-based setting.

The regularized version of problem (11) is equivalent to the proximity operator of $\|\mathbf{x}\|_{\mathfrak{G},\{1,p\}}$. Recently, [96, 131] proposed an efficient algorithm for this proximity operator in large scale settings with extended overlap among groups. In this case, the proximity operator involves: (i) an active set preprocessing step [135] that restricts the prox operations on a subset of the model—i.e., “active” groups and, (ii) a dual optimization step based on Bertsekas’ projected Newton method [12]; however, its convergence requires the strong regularity of the Hessian of the objective near the optimal solution. The authors in [1] propose a fixed point method to compute the proximity operator for a wide range of group sparse model variants, given the model-structured mapping of the fixed-point Picard iterations is non-expansive. [136] develops an efficient primal-dual criterion for the overlapping group sparse proximity operator: the solution of its dual formulation is used to compute the duality gap and tweak the accuracy of the solution and, thus, the convergence of the algorithm.

3.3 Extensions

In many applications, such as genome-wide association studies [141] and multinomial classification [133], it is desirable to find approximations that are not only group-sparse, but also sparse in the usual sense (see [118] for an extension of the group lasso). A classic illustrative example is the sparse group lasso approach [47], a regularization method that combines the lasso [123] and the group lasso [93].

From a discrete perspective, the original problem (6) can be generalized by introducing a sparsity constraint k and allowing to individually select variables within a group. The generalized integer problem then becomes

$$\max_{\omega \in \mathbb{B}^M, \mathbf{y} \in \mathbb{B}^n} \left\{ \sum_{i=1}^n y_i x_i^2 \mid \mathbf{A}\omega \geq \mathbf{y}, \sum_{i=1}^n y_i \leq k, \sum_{j=1}^M \omega_j \leq G \right\}. \quad (12)$$

This problem is in general NP-hard too, but it turns out that it can be solved in polynomial time for the same group structures that allow to solve (6).

Proposition 1 [5] *Given a loopless group structure \mathfrak{G} , there exists dynamic programming algorithm that solves (12) with complexity linear in n and k .*

From a convex point of view, Simon et al. [118] are the first to propose the sparse group model regularization in the context of linear regression. Given a group model \mathfrak{G} and constants $\lambda_1, \lambda_2 > 0$, they consider the problem

$$\operatorname{minimize}_{\mathbf{x} \in \mathbb{R}^n} \left\{ f(\mathbf{x}) + \lambda_1 \sum_{g \in \mathfrak{G}} \|\mathbf{x}_{|g}\|_2 + \lambda_2 \|\mathbf{x}\|_1 \right\}.$$

For this purpose, a generalized gradient descent algorithm is proposed. [133] provides an efficient and robust sparse group lasso algorithm: in this case, a penalized quadratic approximation of the loss function is optimized via a Newton type algorithm in a coordinate descent framework [136].

4 Sparse dispersive models

To describe the *dispersive* structure, we motivate our discussion with an application from neurobiology. Living beings behave and function via transmission of electrical signals between electrically excitable neuronal brain cells. Such chemical “information” causes a swift change in the electrical potential of a possibly discharged neuron cell, which results in its electrical excitation. Currently, we are far from understanding the grid of neurons in its *entirety*: large-scale brain models are difficult to handle while complex neuronal signal models lead to non-interpretable results.




Figure 7: Neuronal spike train example.

Inspired by the statistical analysis in [51], the authors in [63] consider a simple one-dimensional model, where the neuronal signal behaves as a train of spike signals with some *refractoriness* period $\Delta > 0$: there is a minimum nonzero time period Δ where a neuron remains inactive between two consecutive electrical excitations. In statistical terms, neuronal signals are defined by a inter-spike interval distribution that characterizes the probability a new spike to be generated as a function of the inter-arrival time. Figure 7 illustrates how a collection of noisy neuronal spike signals with $\Delta > 0$ might appear in practice.

4.1 The discrete model

We provide next a formal definition of the sparse dispersive model:

Definition 4.1 (Dispersive model) We define the dispersive model \mathcal{D}_k in n -dimensions with sparsity level k and refractory parameter $\Delta \in \mathbb{Z}_+$ as:

$$\mathcal{D}_k = \{ \mathcal{S}_q \mid \forall q, \mathcal{S}_q \subseteq \mathcal{N}, |\mathcal{S}_q| \leq k \text{ and } |i - j| > \Delta, \forall i \neq j \in \mathcal{S}_q \}, \quad (13)$$

i.e., \mathcal{D}_k is a collection of index subsets \mathcal{S}_q in \mathcal{N} with number of elements no greater than k and with distance between the indices in \mathcal{S}_q greater than the interval Δ .

We note that if there are no constraints on the interval of consecutive spikes, the dispersive model naturally boils down to the simple sparsity model Σ_k .

Given the definition above, the projection (1) is:

$$\mathcal{P}_{\mathcal{D}_k}(\mathbf{x}) \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^n} \{ \|\mathbf{w} - \mathbf{x}\|_2^2 \mid \operatorname{supp}(\mathbf{w}) \in \mathcal{D}_k \}. \quad (14)$$

Let $\boldsymbol{\omega} \in \mathbb{B}^n$ be a *support indicator* binary vector, i.e., $\boldsymbol{\omega}$ represents the support set of a sparse vector \mathbf{x} such that $\operatorname{supp}(\boldsymbol{\omega}) = \operatorname{supp}(\mathbf{x})$. Moreover, let $\mathbf{D} \in \mathbb{B}^{(n-\Delta+1) \times n}$ such that:

$$\mathbf{D} = \begin{bmatrix} 1 & 1 & \cdots & 1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \cdots & 1 & 1 & 0 & \cdots & 0 \\ & & & & & \ddots & & & \\ 0 & \cdots & 0 & 0 & 1 & 1 & \cdots & 1 & 1 \end{bmatrix}_{(n-\Delta+1) \times n} \quad (15)$$

Here, per row, there are Δ consecutive ones that denote the time interval between two potential consecutive spikes. Finally, let $\mathbf{b} \in \mathbb{R}^{n-\Delta+2}$ such that $\mathbf{b} := [k \ 1 \ 1 \ \cdots \ 1 \ 1]^T$.

According to [63], the following linear support constraints encodes the definition of the dispersive model \mathcal{D}_k :

$$\mathbf{A} := \begin{bmatrix} \mathbf{1} \\ \mathbf{D} \end{bmatrix} \boldsymbol{\omega} \leq \mathbf{b}. \quad (16)$$

One can observe that $\mathcal{D}_k \equiv \{ \bigcup_{\boldsymbol{\omega} \in \mathcal{Z}} \operatorname{supp}(\boldsymbol{\omega}) \mid \mathcal{Z} := \{ \boldsymbol{\omega} \in \mathbb{B}^n : \mathbf{A}\boldsymbol{\omega} \leq \mathbf{b} \} \}$. To this end, (14) becomes:

$$\mathcal{P}_{\mathcal{D}_k}(\mathbf{x}) \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^n} \{ \|\mathbf{w} - \mathbf{x}\|_2^2 \mid \mathbf{A} \cdot \operatorname{supp}(\mathbf{w}) \leq \mathbf{b} \}. \quad (17)$$

A key observation is given in the next lemma.

Lemma 4.2 [63] Given the problem setting above, it is easy to observe that (17) has solution $\mathcal{P}_{\mathcal{D}_k}(\mathbf{x})$ such that $\mathcal{S} := \text{supp}(\mathcal{P}_{\mathcal{D}_k}(\mathbf{x}))$ and $(\mathcal{P}_{\mathcal{D}_k}(\mathbf{x}))_{\mathcal{S}} = \mathbf{x}_{\mathcal{S}}$ where:

$$\mathcal{S} \in \text{supp} \left(\underset{\omega \in \mathbb{B}^n: \mathbf{A}\omega \leq \mathbf{b}}{\text{argmax}} \{ \mathbf{c}^T \omega \} \right), \quad \text{where } \mathbf{c} := [x_1^2 \quad x_2^2 \quad \dots \quad x_n^2]^T, \quad (18)$$

i.e., we target to capture most of the signal's \mathbf{x} energy, given structure \mathcal{D}_k . To solve (18), the authors in [63] identify that the binary integer program (18) is identical to the solution of the linear program, obtained by relaxing the integer constraints into continuous constraints.

To this end, Lemma 4.2 indicates that (17) can be efficiently performed using linear programming tools [20]. Once (17) is relaxed to a convex problem, decades of knowledge on convex analysis and optimization can be leveraged. Interior point methods find a solution with fixed precision in polynomial time but their complexity might be prohibitive even for moderate-sized problems.

4.2 Convex approaches

The constraint matrix \mathbf{D} describes a *collection of groups*, where each group is assumed to have at most one non-zero entry to model the refractoriness property.¹ Moreover, these groups are *overlapping* which aggrandizes the “clash” between neighboring groups: a non-zero entry in a group discourages every other overlapping group to have a distinct non-zero entry.

In mathematical terms, each row i of \mathbf{D} defines a group \mathcal{G}_i such that $\mathcal{G}_i = \text{supp}(\mathbf{d}_i) \subseteq \mathcal{N}$ where \mathbf{d}_i denotes the i -th row of \mathbf{D} , $\forall i \in \{1, \dots, M := n - \Delta + 1\}$:

$$\mathbf{D} = \begin{bmatrix} \overbrace{\quad \quad \quad}^{\Delta} \\ \mathcal{G}_1 \\ \mathcal{G}_2 \\ \vdots \\ \mathcal{G}_M \end{bmatrix}$$

Given such group structure, the dispersive model is characterized both by *inter-group* and *intra-group* properties:

- *Intra-group sparsity*: we desire $\|\mathbf{D}\omega\|_{\infty} \leq 1$, i.e., per refractoriness period of length Δ , we require only one “active” spike.
- *Inter-group exclusion*: due to the refractoriness property, the activation of a group implies the deactivation of its closely neighboring groups.

While the sparsity level within a group can be easily “convexified” using standard ℓ_1 -norm regularization, the dispersive model further introduces the notion of *inter-group exclusion*, which is highly *combinatorial*. However, one can relax it by introducing *competitions* among variables in overlapping groups: variables that have a “large” neighbor should be penalized more than variables with “smaller” neighbors.

In this premise and based on [142], we identify the following family of norms²:

$$\Omega_{\text{exclusive}}(\mathbf{x}) = \sum_{\mathcal{G}_i} \left(\sum_{j \in \mathcal{G}_i} |x_j| \right)^p, \quad p = 2, 3, \dots, \quad (20)$$

as convex regularizers that imitate the dispersive model. In (20), $\left(\sum_{j \in \mathcal{G}_i} |x_j| \right) := \|\mathbf{x}_{\mathcal{G}_i}\|_1$ promotes sparsity within each group \mathcal{G}_i , while the outer sum over groups $\sum_{\mathcal{G}_i} \|\mathbf{x}_{\mathcal{G}_i}\|_1^p$ imposes sparsity over the number of groups that are activated. Observe that for $p = 1$, (20) becomes the standard ℓ_1 -norm over \mathcal{N} . Notice that the definition of the overlapping groups (instead of non-overlapping) is a key property for capturing the discrete structure: variables belonging to overlapping groups are weighted differently when considered parts of different groups. This leads to variable “suppression” (i.e., thresholding) of elements, depending on the “weight” of their neighborhood within the groups they belong to.

¹Other convex structured models that can be described as the composition of a simple function over a linear transformation \mathbf{D} can be found in [1].

²The proposed norm originates from the composite absolute penalties (CAP) convex norm, proposed in [140], according to which:

$$g(\mathbf{x}) = \sum_{\mathcal{G}_i} \left(\sum_{j \in \mathcal{G}_i} |x_j|^\gamma \right)^p, \quad (19)$$

for various values of γ and p . Observe that this model also includes the famous group sparse model where $g(\mathbf{x}) = \sum_{\mathcal{G}_i} \|\mathbf{x}_{\mathcal{G}_i}\|_2$, described in Section 3, for $p = 1/2$ and $\gamma = 2$.




Figure 8: Wavelet coefficients naturally cluster along a rooted connected subtree of a regular tree and tend to decay towards the leaves. (Left) Example of wavelet tree for a 32×32 image. The root of the tree is the top-left corner and there are three regular subtrees related to horizontal, vertical and diagonal details. Each node is connected to four children representing detail at a finer scale. (Centre) Grayscale 512×512 image. (Right) Wavelet coefficients for the image at center. Best viewed in color, dark blue represents values closer to zero.




Figure 9: Hierarchical constraints. Each node represent a variable. (Left) A valid selection of nodes. (Right) An *invalid* selection of nodes.

5 Hierarchical sparse models

Hierarchical structures are found in many signals and applications. For example, the wavelet coefficients of images are naturally organized on regular quad-trees to reflect their multi-scale structure, Figure 8 and [6–8,33, 61,67,90,116,140]; gene networks are described by a hierarchical structure that can be leveraged for multi-task regression [77]; hierarchies of latent variables are typically used for deep learning [11].

In essence, a hierarchical structure defines an ordering of importance among the elements (either individual variables or groups of them) of a signal with the rule that an element can be selected only after its ancestors. Such structured models result into more robust solutions and allow recovery with far fewer samples. In compressive sensing, assuming that the signal possesses a hierarchical structure with sparsity k leads to improved sample complexity bounds of the order of $O(k)$ for dense measurement matrices [8], compared to the bound of $O(k \log(n/k))$ for standard sparsity. Also in the case of sparse measurement matrices, e.g. expanders, hierarchical structures yield improved sample complexity bounds [4,68].

5.1 The discrete model

The discrete model underlying the hierarchical structure is given by the next definition.

Definition 5.1 (Hierarchical model) Let \mathcal{T} denote an arbitrary tree or forest representation over the variables in a set \mathcal{N} . We define a k rooted connected (RC) subtree \mathcal{S} with respect to \mathcal{T} as a collection of k variables in \mathcal{N} such that $v \in \mathcal{S}$ implies $\mathcal{A}(v) \in \mathcal{S}$, where $\mathcal{A}(v)$ is the set that contains all the ancestors of the node v .

The hierarchical model of budget k , \mathcal{T}_k is the set of all k rooted-connected subtrees of \mathcal{T} .

An example of rooted connected subtree over $|\mathcal{N}| = 9$ variables is given in Figure 9.

Given a tree \mathcal{T} , the rooted connected approximation is the solution of the following discrete problem

$$\mathcal{P}_{\mathcal{T}_k}(\mathbf{x}) \in \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} \{ \|\mathbf{x} - \mathbf{z}\|_2^2 \mid \operatorname{supp}(\mathbf{z}) \in \mathcal{T}_k \}. \quad (21)$$




Figure 10: Examples of parent-child and family models. Active groups are indicated by dotted ellipses. The support (black nodes) is given by the union of the active groups. (Left) Parent-child model. (Right) Family model.

which can be reformulated as follows

$$\hat{\mathbf{y}} \in \operatorname{argmax}_{\mathbf{y} \in \mathbb{B}^n} \left\{ \sum_{i=1}^N y_i x_i^2 : \mathbf{y} \in \mathcal{T}_k \right\}, \quad (22)$$

where \mathbf{y} is a binary vector with k non-zero components that indicates which components of \mathbf{x} are selected. Given a solution $\hat{\mathbf{y}}$ of the above problem, a solution $\hat{\mathbf{z}}$ of (21) is then obtained as $\hat{\mathbf{z}}_{|\mathcal{S}} = \mathbf{x}_{|\mathcal{S}}$ and $\hat{\mathbf{z}}_{|\mathcal{N} \setminus \mathcal{S}} = 0$, where $\mathcal{S} = \operatorname{supp}(\hat{\mathbf{y}})$.

This type of constraint can be represented by a group structure with an overall sparsity constraint k , where for each node in the tree we define a group consisting of that node and all its ancestors. When a group is selected, we require that all its elements are selected as well. Problem (22) can then be cast as a special case of the Weighted Maximum Coverage problem (6). Fortunately, this particular group structure leads to tractable solutions.

Indeed (22) can be solved exactly via a dynamic program that runs in polynomial time [5,25]. For d -regular trees, that is trees for which each node has d children, the algorithm in [5] has complexity $\mathcal{O}(nkd)$.

5.2 Convex approaches

The hierarchical structure can also be enforced by convex penalties, based on groups of variables. Given a tree structure \mathcal{T} , define groups consisting of a node and all its descendants and let $\mathcal{G}_{\mathcal{T}}$ represent the set of all these groups. Based on this construction, the hierarchical group lasso penalty [75,77,140] imitates the hierarchical sparse model and is defined as follows

$$\Omega(\mathbf{x})_{\text{HGL}} = \sum_{\mathcal{G} \in \mathcal{G}_{\mathcal{T}}} w_{\mathcal{G}} \|\mathbf{x}_{|\mathcal{G}}\|_p \quad (23)$$

where $p \geq 1$, $w_{\mathcal{G}}$ are positive weights and $\mathbf{x}_{|\mathcal{G}}$ is the restriction of \mathbf{x} to the elements contained in \mathcal{G} . Since the nodes lower down in the tree appear in more groups than their ancestors, they will contribute more to $\Omega(\mathbf{x})_{\text{HGL}}$ and therefore will be more easily encouraged to be zero. The proximity operator of Ω_{HGL} can be computed exactly for $p = 2$ and $p = \infty$ via an active set algorithm [75].

Other convex penalties have been recently proposed in order to favor hierarchical structures, but also allowing for a certain degree of flexibility in deviations from the discrete model. One approach considers groups consisting of all parent-child pairs and uses the latent group lasso penalty (see Section 3.2) in order to obtain solutions whose support is the union of few such pairs [110], see Figure 10 (left).

An interesting extension is given by the *family* model [13,139], where the groups consist of a node and all its children, see Figure 10 (right). Again the latent groups lasso penalty is used. This model is better suited for wavelet decomposition of images because it better reflects the fact that a large coefficient value implies large coefficients values for all its children at a finer scale.

For both these cases, one can use the duplication strategy to transform the overlapping proximity problem into a block one, which can be efficiently solved in closed-form [69].

6 Submodular models

Most of the structures described so far are naturally combinatorial, but the models presented in the previous sections are incapable of capturing more complex structures. A more general approach is to use a set function

R to quantify how much a given support set deviates from the desired structure. For instance, we can describe all previous structures by using an indicator function that assigns an infinite value to sets that do not belong to the chosen structure \mathcal{M}_k . Next, we provide a definition of such set functions and the sets that they describe:

Definition 6.1 Given a set function $R : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, we can define a model \mathcal{R}_τ consisting of all sets \mathcal{S} for which $R(\mathcal{S}) \leq \tau$:

$$\mathcal{R}_\tau := \{\mathcal{S} \mid R(\mathcal{S}) \leq \tau\}. \quad (24)$$

Unfortunately, computing the projection of a given signal onto this set is not generally feasible for any set function, due to their intractable combinatorial nature.

Therefore, it is necessary to restrict our attention to set functions with specific properties that lead to tractable problems. It turns out that in many applications (See for e.g. [2, 3, 6, 26, 62, 78, 97, 115]) the combinatorial penalties have a convenient “diminishing returns” property that we can exploit. These are the so-called *submodular* set functions.

Definition 6.2 A set function $R : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is *submodular* iff $R(\mathcal{S} \cup \{v\}) - R(\mathcal{S}) \geq R(\mathcal{U} \cup \{v\}) - R(\mathcal{U})$, $\forall \mathcal{S} \subseteq \mathcal{U} \subseteq \mathcal{N}$, and $v \in \mathcal{N} \setminus \mathcal{U}$.

Given a combinatorial penalty that encodes the desired structure, we want to solve an inverse problem of the form:

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} f(\mathbf{x}) \quad \text{subject to} \quad \text{supp}(\mathbf{x}) \in \mathcal{R}_\tau. \quad (25)$$

To simplify things, one can relax the hard constraint required in (25), by using the combinatorial penalty as a regularizer to discourage unwanted patterns of the support:

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} f(\mathbf{x}) + \lambda \cdot R(\text{supp}(\mathbf{x})), \quad (26)$$

for some regularization parameter $\lambda > 0$.

Unfortunately, the composite minimization of a continuous and discrete function required in (26) leads to an NP-hard problem, even for submodular set functions. In fact, even in the simple sparsity case in CS setting, where $f(\mathbf{x}) := \frac{1}{2} \|\mathbf{u} - \Phi \mathbf{x}\|_2^2$, and $R(\mathcal{S}) = |\mathcal{S}|$, the problem (26) is NP-hard.

When submodular functions are paired with continuous functions like in (26), the minimization becomes very difficult. However when considered alone, submodular functions can be efficiently minimized, see Section 7.1.2. Submodular function minimization constitutes a key component in both the convex and discrete approaches to solving (26).

6.1 The discrete model

In the discrete setting, (26) is preserved in an attempt to faithfully encode the discrete model at hand. While such criterion seems cumbersome to solve, it has favorable properties that lead to polynomial solvability, irrespective of its combinatorial nature: there are efficient combinatorial algorithms that provide practical solutions to overcome this bottleneck and are guaranteed to converge in polynomial time [43]; see Section 7.1.

Within this context, our intention here is to present fundamental algorithmic steps that reveal the underlying structure of (26). Here, we assume that f has an L -Lipschitz continuous gradient and, thus, it admits the following upper bound:

$$f(\mathbf{x}) \leq f(\mathbf{x}^i) + \langle \nabla f(\mathbf{x}^i), \mathbf{x} - \mathbf{x}^i \rangle + \frac{L}{2} \|\mathbf{x} - \mathbf{x}^i\|_2^2 := Q(\mathbf{x}, \mathbf{x}^i) \quad \forall \mathbf{x}, \mathbf{x}^i \in \text{dom}(f).$$

Then, we perform the following iterative Majorization-Minimization (MM) scheme:

$$\mathbf{x}^{i+1} \in \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \{Q(\mathbf{x}, \mathbf{x}^i) + \lambda R(\text{supp}(\mathbf{x}))\} = \underset{\mathcal{S} \subseteq \mathcal{N}}{\text{argmin}} \left\{ \min_{\mathbf{x} : \text{supp}(\mathbf{x}) = \mathcal{S}} Q(\mathbf{x}, \mathbf{x}^i) + \lambda R(\mathcal{S}) \right\} \quad (27)$$

If we focus on in the inner minimization above and for a given support \mathcal{S} , one can compute the optimal minimizer $\hat{\mathbf{x}}^i$ as: $\hat{\mathbf{x}}_{\mathcal{S}^c}^i = 0$ and $\hat{\mathbf{x}}_{\mathcal{S}}^i = \mathbf{x}_{\mathcal{S}}^i - (1/L)\nabla f(\mathbf{x}^i)_{\mathcal{S}}$. By substituting $\hat{\mathbf{x}}_{\mathcal{S}}^i$ into the upper bound $Q(\cdot, \cdot)$, we obtain a modular majorizer M^i :

$$\begin{aligned} f(\mathbf{x}) + \lambda R(\text{supp}(\mathbf{x})) &\leq C - \frac{L}{2} \sum_{j \in \text{supp}(\mathbf{x})} \left(x_j^i - \frac{1}{L} \nabla f(\mathbf{x}^i)_j \right)^2 + \lambda R(\text{supp}(\mathbf{x})) \\ &:= M^i(\text{supp}(\mathbf{x})) + \lambda R(\text{supp}(\mathbf{x})), \end{aligned}$$

where C is a constant. Finally, the update (27) then becomes:

$$\mathbf{x}^{i+1} = \hat{\mathbf{x}}_{\mathcal{S}}^i \quad \text{subject to} \quad \mathcal{S} \in \arg \min_{\mathcal{S} \subseteq \mathcal{N}} M^i(\mathcal{S}) + R(\mathcal{S}) \quad (28)$$

The minimization in (28) is a *submodular function minimization* (SFM) problem, since the majorizer $M^i + R$ is submodular, and thus can be done efficiently.

Proposition 2 ([43]) *At each iteration, the new estimate \mathbf{x}^{i+1} produced by the Majorization-Minimization algorithm satisfies*

$$f(\mathbf{x}^{i+1}) + R(\text{supp}(\mathbf{x}^{i+1})) \leq f(\mathbf{x}^i) + R(\text{supp}(\mathbf{x}^i)),$$

which implies convergence.

Note that this scheme is a type of proximal gradient method (cf., Sect. 7.2.1), applied to submodular regularizers, where the proximal operation in (41) corresponds to the SFM step in (28). We refer the reader to Section 7.2 for more information.

6.2 Convex approaches

The non-convex problem (26) can be converted to a closely related convex problem by replacing the discrete regularizer $g(\mathbf{x}) = R(\text{supp}(\mathbf{x}))$ by its largest convex lower bound. This is also called convex envelope and is given by the biconjugate g^{**} , i.e. by applying the Fenchel-Legendre conjugate twice [18]. We call this approach a convex relaxation of (26). However, for some functions g , there is no meaningful convex envelope. For these cases, one can use the convex closure $R^- : [0, 1]^n \rightarrow \mathbb{R}$, of $R(\mathcal{S})$, which is the point-wise largest lower bound of R [41]. Note that both the convex envelope and the convex closure are “tight” relaxations, but one in the continuous domain, the other in the discrete domain.

Both notions of convex relaxation for submodular functions use the Lovász extension (LE), introduced in [88]. We give here only one of the equivalent definitions:

Definition 6.3 *Given a set function R such that $R(\emptyset) = 0$, we define its Lovász extension as follows: $\forall \mathbf{x} \in \mathbb{R}^N$,*

$$r(\mathbf{x}) = \sum_{k=1}^N x_{j_k} \left(R(\{j_1, \dots, j_k\}) - R(\{j_1, \dots, j_{k-1}\}) \right) \quad (29)$$

where $x_{j_1} \geq \dots \geq x_{j_N}$.

The Lovász extension is the convex closure of its corresponding submodular function on the unit hypercube, i.e $R^- = r$ [41]. In this sense, this extension already gives a convex relaxation for any submodular function. However, it turns out that using the Lovász extension as a convexification might not always fully capture the structure of the discrete penalty. We elaborate more on this in Section 6.3.

For a monotone³ submodular function its convex envelope is also obtained through the LE.

Proposition 3 ([2]) *Given a monotone submodular function R , s.t $R(\emptyset) = 0$, and its Lovász extension r , the convex envelope of $g(\mathbf{x}) = R(\text{supp}(\mathbf{x}))$ on the unit ℓ_∞ -ball is given by the norm $\Omega_{sub}(\mathbf{x}) := g^{**}(\mathbf{x}) = r(|\mathbf{x}|)$.*

The proximity operator (2) of Ω_{sub} can be efficiently computed. In fact, computing $\text{prox}_\lambda^{\Omega_{sub}}(\mathbf{x})$ is shown in [2] to be equivalent to minimizing the submodular function $\lambda R(\mathcal{S}) - \sum_{i \in \mathcal{S}} |x_i|$ using the minimum-norm point algorithm. Other SFM algorithms can also be used, but then solving a sequence of SFMs is needed. Note that simpler subcases of submodular functions may admit more efficient proximity operator than this general one.

As a result, the convex relaxation of (26), where $R(\text{supp}(\mathbf{x}))$ is replaced by $\Omega_{sub}(\mathbf{x})$, can then be efficiently solved by a proximal gradient method (cf., Sect. 7.2.1).

The monotonicity requirement in (3) is necessary to guarantee the convexity of $r(|\mathbf{x}|)$. Unfortunately, all symmetric⁴ submodular functions, among which undirected cut functions (cf., Sect. 6.3), are not monotone. In fact, the convex envelope of $R(\text{supp}(\mathbf{x}))$ on the unit ℓ_∞ -ball, for R any submodular function, is given by [2]:

$$g^{**}(\mathbf{x}) = \min_{\delta \in [0, 1]^n, \delta \geq |\mathbf{x}|} r(\delta). \quad (30)$$

Thus, when R is non-decreasing, the convex envelope is $r(|\mathbf{x}|)$ as stated in Proposition 3. When R is symmetric and $R(\mathcal{N}) = R(\emptyset) = 0$ (which is assumed w.l.o.g since addition of constants doesn’t affect the regularization), $g^{**}(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathbb{R}^n$ and the minimum in (30) is attained at any constant value vector δ .

So the convexification of symmetric submodular functions consists of the Lovász extension alone. However the LE is a poor convexification that can significantly modify the intended structure. This can already be seen by the fact that the LE is tight only on the unit hypercube, while most penalties $R(\text{supp}(\mathbf{x}))$ are not constrained there, and that $R(\text{supp}(\mathbf{x}))$ is symmetric around the origin which is not the case for $r(\mathbf{x})$. Some artifacts of using the LE as a convexification for symmetric functions are illustrated in [43].

For some problems, the submodularity requirement can be relaxed. Convex relaxations of general set functions, when paired with the ℓ_p -norm ($p > 1$) as a continuous prior are studied in [107].

³A monotone function is a function that satisfies: $\forall \mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{N}, R(\mathcal{S}) \leq R(\mathcal{T})$

⁴A symmetric function is a function that satisfies: $\forall \mathcal{S} \subseteq \mathcal{N}, R(\mathcal{S}) = R(\mathcal{N} \setminus \mathcal{S})$

Proposition 4 ([107]) Define the norm

$$\Omega_p(\mathbf{x}) := \min_{\mathbf{v} \in \mathcal{V}} \left\{ \sum_{\mathcal{S} \subseteq \mathcal{N}} R(\mathcal{S})^{\frac{1}{q}} \|\mathbf{v}^{\mathcal{S}}\|_p \mid \sum_{\mathcal{S} \subseteq \mathcal{N}} \mathbf{v}^{\mathcal{S}} = \mathbf{x} \right\}, \quad (31)$$

where $\frac{1}{p} + \frac{1}{q} = 1$ and $\mathcal{V} = \{\mathbf{v} = (\mathbf{v}^{\mathcal{S}})_{\mathcal{S} \subseteq \mathcal{N}} \in (\mathbb{R}^n)^{2^{\mathcal{N}}} \text{ s.t. } \text{supp}(\mathbf{v}^{\mathcal{S}}) \subseteq \mathcal{S}\}$.

Ω_p is the convex positively homogeneous envelope of the function $\lambda R(\text{supp}(\mathbf{x})) + \mu \|\mathbf{x}\|_p^p$ (up to a constant factor), and the convex envelope of $R(\text{supp}(\mathbf{x}))^{\frac{1}{q}} \|\mathbf{x}\|_p$.

Note that this is the same norm as (8), with weights d_j given by a set function R , such that $d_j = R(\mathcal{G}_j)^{1/q}$, and the group structure \mathfrak{G} is the entire power set of \mathcal{N} .

The convex relaxation proposed in Proposition 4, when applied to a submodular function, generalizes the relaxation of Proposition 3 to unit balls with respect to the ℓ_p -norm ($p > 1$) and not only to the ℓ_∞ one.

6.3 Examples

We now offer some examples of structures that can be described via submodular functions.




Figure 11: Wavelet coefficients and underlying hierarchical group model.




Figure 12: Background subtraction and underlying Ising model.

Simple sparsity: The most ubiquitous structure in practice is simple sparsity, which corresponds to the set function $R(\mathcal{S}) = |\mathcal{S}|$ that measures the cardinality of \mathcal{S} . The convexification obtained by Proposition 3 is the usual ℓ_1 -norm.

Group sparsity: Given a group structure \mathfrak{G} , a group sparse model can be enforced by penalizing the number of groups that intersects with the support: $R_\cap(\mathcal{S}) = \sum_{\mathcal{S} \cap \mathcal{G}_i \neq \emptyset} d_i$ where $d_i \geq 0$ is the weight associated with the group $\mathcal{G}_i \in \mathfrak{G}$. For example, defining the groups to be each node and all of its descendants on a tree (see Figure 11), enforces the support to form a rooted connected tree, which is characteristic of wavelet coefficients (cf., Sect. 5). The way the groups and their weights are defined leads to different sparsity patterns [3].

The convexification of $R_\cap(\mathcal{S})$ is the ℓ_1/ℓ_∞ -norm⁵ over groups $\sum_{\mathcal{G}_i \in \mathfrak{G}} d_i \|\mathbf{x}\|_\infty$ (even for overlapping groups) [2, 140]. Note that for groups that form a partition of \mathcal{N} , $R_\cap(\mathcal{S})$ is equivalent to the minimum weight set cover, defined as:

$$R_{sc}(\mathcal{S}) = \min_{s \in \mathbb{B}^M} \sum_{i=1}^M d_i s_i \text{ s.t. } \sum_{i=1}^M s_i \mathbb{1}_{n, \mathcal{G}_i} \geq \mathbb{1}_{n, \mathcal{S}} \quad (32)$$

with the same set of groups \mathfrak{G} and weights d . R_{sc} computes the total weight of the lightest cover for \mathcal{S} using the groups in \mathfrak{G} . Note that $R_{sc}(\mathcal{S})$ is not a submodular function⁶, unless the groups form a partition.

Ising model: The Ising model [92] is a model that associates with each coefficient x_i of a signal $\mathbf{x} \in \mathbb{R}^n$ a discrete variable $s_i \in \{-1, 1\}$ to represent the state (zero/non-zero) of the coefficient. The Ising penalty enforces clustering over the coefficients, which is a desired structure for example in background subtraction in images or videos [26]. It can be naturally encoded on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where the vertices are the coefficients of the signal, i.e. $\mathcal{V} = \mathcal{N}$, and the edges connects neighboring coefficients. For example, for images, the coefficients of the signal are the pixels of the image, and edges connect pixels next to each other, forming a lattice (see Figure 12). This is called the two dimensional lattice Ising model. The Ising penalty is then expressed via the following symmetric submodular function:

$$R_{\text{ISING}}(\mathcal{S}) = \frac{1}{2} \left(|\mathcal{E}| - \sum_{(i,j) \in \mathcal{E}} s_i s_j \right), \quad (33)$$

⁵Actually, it is a norm iff $\mathcal{N} = \cup_{d_i > 0} \mathcal{G}_i$

⁶Consider this example: let $\mathcal{N} = \{1, 2, 3, 4\}$, and $\mathfrak{G} = \{\mathcal{G}_1 = \{1\}, \mathcal{G}_2 = \{2, 3\}, \mathcal{G}_3 = \{1, 2, 4\}\}$, with weights defined as $d_i = |\mathcal{G}_i|$. Then the inequality in definition 6.2 is not satisfied for the sets $\mathcal{S} = \{1, 2\}$ for which $R_{sc}(\mathcal{S}) = 3$, and $\mathcal{T} = \{1, 2, 4\}$ for which $R_{sc}(\mathcal{T}) = 4$, with the addition of the element $\{3\}$.

where $\mathbf{s} \in \mathbb{R}^n$ is an indicator vector for the set \mathcal{S} such that $s_i = 1$ if $i \in \mathcal{S}$ and $s_i = -1$, otherwise. When $\mathcal{S} = \text{supp}(\mathbf{x})$ for a signal $\mathbf{x} \in \mathbb{R}^n$, s_i encodes the state of the corresponding coefficient x_i . We can also view the Ising penalty as a cut function: $R_{\text{ISING}}(\mathcal{S})$ indeed just counts the number of edges that are cut by the set \mathcal{S} . Cut functions with appropriate graphs and weights actually capture a large subset of submodular functions [72,78]. Since R_{ISING} is symmetric, its convexification is just its Lovász extension, as discussed above, which is shown [43] to be the anisotropic discrete Total Variation semi-norm

$$\|\mathbf{x}\|_{TV} = \sum_{(i,j) \in \mathcal{E}} |x_i - x_j| \quad (34)$$

While the Ising model enforces the clustering of non-zero coefficients, its convexification encodes a different structure by favoring piece-wise constant signals. Furthermore, this penalty cannot be described by any of the structures introduced in the previous section.

Entropy: Given n random variables X_1, \dots, X_n , the joint entropy of $X_{\mathcal{S}} = (X_i)_{i \in \mathcal{S}}$ defines a submodular function, $R(\mathcal{S}) = H(X_{\mathcal{S}})$. Moreover, the mutual information also defines a symmetric submodular function, $R(\mathcal{S}) = I(X_{\mathcal{S}}, X_{\mathcal{N} \setminus \mathcal{S}})$. Such functions occur, for example, in experimental design [115], in learning of Bayesian networks [62], in semi-supervised clustering [97], and in diverse feature selection [36]. Both of these set functions encourage structures that cannot be defined by the models presented in the previous sections.

Supermodular functions: Another variant of set function regularizers used in practice, for example in diverse feature selection, are supermodular functions, i.e. the negative of a submodular function. The problem (26) with supermodular regularizers can then be cast as a maximization problem with a submodular regularizer. The notion of “approximate” submodularity introduced in [37,79] is used to reduce the problem to an “approximate” submodular maximization problem, where several efficient algorithms for constrained/unconstrained submodular maximization [21,22] with provable approximation guarantees can be employed. For more information, we refer the reader to [36].

7 Optimization

In order to use the previous structures in practice, one needs efficient optimization solutions for structured sparsity problems that scale up in high-dimensional settings. From our discussions above, it is apparent that the key actors for this purpose are projection and proximity operations over structured sets that go beyond simple selection heuristics and towards provable solution quality as well as runtime/space bounds.

Projection operations faithfully follow the underlying combinatorial model but, in most cases, result in hard-to-solve or even combinatorial optimization problems. Furthermore, model misspecification often results in wildly inaccurate solutions.

Proximity operators of convex sparsity-inducing norms often can only partially describe the underlying discrete model and might lead to “rules-of-thumb” in problem solving (e.g., how to set up the regularization parameter). However, such approaches work quite well in practice and are more robust to deviations from the model, leading to satisfactory solutions.

Here, our intention is to present an overview of the dominant approaches followed in practice. We consider the following three general optimization formulations:⁷

- *Discrete projection formulation:* Given a signal model \mathcal{M}_k , let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a closed convex data fidelity/loss function. Here, we focus on the *projected* non-convex minimization problem:

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{x} \in \mathcal{M}_k. \quad (35)$$

- *Convex proximity formulation:* Given a signal model \mathcal{M}_k , let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a closed convex data fidelity/loss function, $g : \mathbb{R}^n \rightarrow \mathbb{R}$ a closed convex regularization term, possibly non-smooth, that faithfully models \mathcal{M}_k and $\lambda > 0$. In this chapter, we focus on the convex composite minimization problem:

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\mathbf{x}) + \lambda \cdot g(\mathbf{x}). \quad (36)$$

- *Convex structured-norm minimization:* Given a signal model \mathcal{M}_k , let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a closed convex regularization term, possibly non-smooth, that faithfully models \mathcal{M}_k . Moreover, let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a closed convex data fidelity/loss function and $\sigma > 0$. We consider the following minimization problem:

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad g(\mathbf{x}) \quad \text{subject to} \quad f(\mathbf{x}) \leq \sigma. \quad (37)$$

⁷We acknowledge that there are other criteria that can be considered in practice; for completeness, in the simple sparsity case, we refer the reader to the ℓ_1 -norm constrained linear regression (a.k.a. Lasso [123])—similarly, there are alternative optimization approaches for the discrete case [134]. However, our intention in this chapter is to use the most prevalent formulations used in practice.

In addition, we also briefly mention submodular function minimization (SFM)

$$\underset{\mathcal{S} \in 2^{\mathcal{N}}}{\text{minimize}} \quad R(\mathcal{S})$$

where $R : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is a submodular function; see Definition 6.2. SFM is at the core of both discrete and convex approaches for solving estimation problems, where the structure is imposed via a submodular set function; see Section 6.

Some representative problem instances are given next.

Compressive sensing (CS): In classic CS, we are interested in recovering a high-dimensional signal $\mathbf{x}^* \in \mathcal{M}_k$ from an underdetermined set of linear observations $\mathbf{u} \in \mathbb{R}^m$ ($m < n$), generated via a known $\Phi \in \mathbb{R}^{m \times n}$:

$$\mathbf{u} = \Phi \mathbf{x}^* + \varepsilon.$$

Here, $\varepsilon \in \mathbb{R}^m$ denotes an additive noise term with $\|\varepsilon\|_2 \leq \sigma$.

For $f(\mathbf{x}) := \frac{1}{2} \|\mathbf{u} - \Phi \mathbf{x}\|_2^2$ and using the fact that $\mathbf{x}^* \in \mathcal{M}_k$, one can formulate the problem at hand as in (35); observe that, in the case where $\mathcal{M}_k \equiv \Sigma_k$, (35) corresponds to the ℓ_0 pseudo-norm optimization problem [98]. Alternatively, given a function g that encourages solutions to be in \mathcal{M}_k , one can solve the CS problem with convex machinery as in (36)-(37); when $\mathcal{M}_k \equiv \Sigma_k$, (36)-(37) corresponds to the Basis Pursuit DeNoising (BPDN) criterion [31].

Sparse graph modeling: In Gaussian graph model selection [34, 39, 64, 65, 82, 83, 127], the sparse-regularized maximum likelihood estimation yields the criterion:

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \in \mathbb{R}^{n^2}}{\text{argmin}} \quad \left\{ -\log \det(\text{mat}(\mathbf{x})) + \text{trace}(\text{mat}(\mathbf{x}) \cdot \hat{\mathbf{C}}) + \lambda \cdot g(\mathbf{x}) \right\}, \quad (38)$$

where $\mathbf{x} \in \mathbb{R}^{n^2}$ is the vectorized version of the inverse covariance estimate, $\text{mat}(\cdot)$ is the linear matricization operation and $\hat{\mathbf{C}}$ is the sample covariance. Developments in random matrix theory [76] divulge the poor performance of solving (38) without regularization: in that case, the solution is usually fully dense and no inference about the dependence graph structure is possible. Choosing a suitable regularizer g [71, 85, 87], that models the connections between the variables, reduces the degrees of freedom, leading to robust and more interpretable results.

Other applications: Similar sparsity regularization in discrete or convex settings is found many diverse disciplines: low-light imaging problem in signal processing [59], heteroschedastic LASSO [35], dictionary learning, [122], etc.

7.1 Greedy and discrete approaches

7.1.1 Projected gradient descent and matching pursuit variants

Iterative greedy algorithms maintain the combinatorial nature of (35), but instead of tackling it directly, which would be intractable, the algorithms of this class greedily refine a k -model-sparse solution using only “local” information available at the current iteration. Within this context, matching pursuit approaches [91, 129] gradually construct the sparse estimate by greedily choosing the non-zero coefficients that best explain the current residual, e.g. $\mathbf{y} - \mathbf{A}\mathbf{x}^i$, where \mathbf{x}^i is the current iterate. Extensions of such greedy approaches have been recently proposed to accommodate structured models in the selection process [67].

Most of the algorithmic solutions so far concentrate on the non-convex projected gradient descent algorithm: a popular method, known for its simplicity and ease of implementation. Per iteration, the total computational complexity is determined by the calculation of the gradient and the projection operation on \mathcal{M}_k as in (1). Most algorithms in the discrete model can be described or easily deduced by the following simple recursion:

$$\mathbf{x}^{i+1} = \mathcal{P}_{\mathcal{M}_k} \left(\mathbf{x}^i - \frac{\mu}{2} \nabla f(\mathbf{x}^i) \right), \quad (39)$$

where μ is a step size and $\mathcal{P}_{\mathcal{M}_k}(\cdot)$ is the projection onto k -model-sparse signals.

Representative examples for the CS application are hard thresholding methods over simple sparse sets [14, 27, 46, 80, 84, 99]. [8] further extends these ideas to *model-based CS*, where non-overlapping group structures and tree structures are used to perform the projection (1). From a theoretical computer science perspective, [4, 68] address the CS problem using *sparse sensing matrices*. Exploiting model-based sparsity in recovery provably reduces the number of measurements m , without sacrificing accuracy. The resulting algorithm reduces the main computational cost of the proposed scheme on the difficulty of projecting onto the model-sparse set, which, as mentioned in Section 3, in most relevant cases, can be computed in linear time using dynamic programming.⁸

⁸In the case of CS, an important modification of (39) to achieve linear computational time per iteration is the substitution of the gradient with the *median* operator, which is nonlinear and defined component-wise on a vector; for more information, we refer to [46, 52].

7.1.2 Submodular function minimization

Submodularity is considered the discrete equivalent of convexity in the sense that it allows efficient minimization. The best combinatorial algorithm for submodular function minimization (SFM) has a proven complexity of $O(n^5T + n^6)$, where T is the function evaluation complexity [109]. For practical purposes however, another algorithm called minimum-norm point algorithm [48] is usually used. This algorithm has no known worst case complexity but in practice it usually runs in $O(n^2)$. Furthermore, for certain functions which are “graph representable” [49, 72], submodular minimization becomes equivalent to computing the minimum s-t cut on the corresponding graph $G(\mathcal{V}, \mathcal{E})$, which has time complexity⁹ $\tilde{O}(|\mathcal{E}| \min\{|\mathcal{V}|^{2/3}, |\mathcal{E}|^{1/2}\})$ [54].

7.2 Convex approaches

7.2.1 Proximity methods

Proximity gradient methods are iterative processes that rely on two key structural assumptions: i) f has Lipschitz continuous gradient¹⁰ (see Definition 2.1) and ii) the regularizing term g is endowed with a *tractable* proximity operator.

By the Lipschitz gradient continuity and given a putative solution $\mathbf{x}^i \in \text{dom}(f + g)$, one can locally approximate f around \mathbf{x}^i using a quadratic function as:

$$f(\mathbf{x}) \leq Q(\mathbf{x}, \mathbf{x}^i) := f(\mathbf{x}^i) + \nabla f(\mathbf{x}^i)^T(\mathbf{x} - \mathbf{x}^i) + \frac{L}{2}\|\mathbf{x} - \mathbf{x}^i\|_2^2, \quad \forall \mathbf{x} \in \text{dom}(F).$$

The special structure of this upper-bound allows us to consider a majorization-minimization approach: instead of solving (36) directly, we solve a sequence of simpler composite quadratic problems:

$$\mathbf{x}^{i+1} \in \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \{Q(\mathbf{x}, \mathbf{x}^i) + g(\mathbf{x})\}. \quad (40)$$

In particular, we observe that (40) is equivalent to the following iterative *proximity* operation, similar to (2):

$$\mathbf{x}^{i+1} \in \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \left\{ \frac{1}{2} \left\| \mathbf{x} - \left(\mathbf{x}^i - \frac{1}{L} \nabla f(\mathbf{x}^i) \right) \right\|_2^2 + \frac{1}{2L} g(\mathbf{x}) \right\}. \quad (41)$$

Here, the anchor point \mathbf{w} in (2) is the gradient descent step: $\mathbf{w} := \mathbf{x}^i - \frac{1}{L} \nabla f(\mathbf{x}^i)$.

In the case where $g(\mathbf{x}) := \|\mathbf{x}\|_1$, the proximity algorithm in (41) is known as the Iterative Soft-Thresholding Algorithm (ISTA) [28, 32, 38]. Iterative algorithms can use memory to provide momentum in convergence. Based on Nesterov’s optimal gradient methods [103], [10] proves the universality of such acceleration in the composite convex minimization case of (36), where $g(\mathbf{x})$ can be any convex norm with tractable proximity operator. However, the resulting optimization criterion in (41) is more challenging when g stands for more elaborate sparsity structures.

Within this context, [114, 132] present a new convergence analysis for proximity (accelerated) gradient problems, under the assumption of *inexact proximity evaluations* and study how these errors propagate into the convergence rate.

An emerging direction for solving composite minimization problems of the form (36) is based on the proximity-Newton method [50]. The origins of this method can be traced back to the work of [16, 50], which relies on the concept of *strong regularity* introduced by [113] for generalized equations. In this case, we identify that the basic optimization framework above can be easily adjusted to second-order Newton gradient and quasi-Newton approaches:

$$\mathbf{x}^{i+1} \in \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \left\{ \frac{1}{2} \left\| \mathbf{x} - \left(\mathbf{x}^i - \mathbf{H}_i^{-1} \nabla f(\mathbf{x}^i) \right) \right\|_{\mathbf{H}_i}^2 + \frac{1}{2} g(\mathbf{x}) \right\}. \quad (42)$$

where \mathbf{H}_i represents either the actual Hessian of f at \mathbf{x}^i (i.e., $\nabla^2 f(\mathbf{x}_i)$) or a symmetric positive definite matrix approximating $\nabla^2 f(\mathbf{x}^i)$. Given a computationally efficient Newton direction, one can re-use the model-based proximity solutions presented in the previous subsection along with a second order *variable metric* gradient descent scheme, as presented in (42) [128].

7.2.2 Primal-dual and alternating minimization approaches

Several model-based problems can be solved using the convex structured-norm minimization in (37). As a stylized example, consider the noiseless CS problem formulation

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad g(\mathbf{x}) \quad \text{subject to} \quad \mathbf{u} = \Phi \mathbf{x}, \quad (43)$$

⁹the notation $\tilde{O}(\cdot)$ ignores log terms

¹⁰In [128], the authors consider a more general class of functions with no *global* Lipschitz constant L over their domain. The description of this material is out of the scope of this chapter and is left to the reader who is interested in deeper convex analysis and optimization.

where g is a convex structured norm, mimicking the model at hand.

Within this context, primal-dual convex optimization methods provide attractive approaches, by exploiting both the primal and the dual formulations at the same time (via Lagrangian dualization). More precisely, primal-dual methods solve the following minimax problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \max_{\mathbf{y} \in \mathbb{R}^m} \{g(\mathbf{x}) + \langle \Phi \mathbf{x} - \mathbf{u}, \mathbf{y} \rangle\}, \quad (44)$$

where \mathbf{y} is the Lagrange multiplier associated to the equality constraint $\mathbf{u} = \Phi \mathbf{x}$. The minimax formulation (44) can be considered as a special case of the general formulation in [29, 104].

If the function g is nonsmooth and possesses a closed form proximity operator, then general primal-dual methods such as [29, 60] have been known as powerful tools to tackle problem (44). Since g is nonsmooth, one can also use the primal-dual subgradient method [106] as well as the prox-method proposed in [102] for solving (43). An alternative approach is to combine primal-dual optimization methods and smoothing technique as done in [104, 105] for solving (44).

When (44) possesses a separable structure $g(\mathbf{x}) := \sum_{i=1}^p g_i(\mathbf{x}_i)$, distributed approaches can be applied. Recently, the authors in [126] developed a unified primal-dual decomposition framework for solving (43), using also ideas from the alternating direction methods of multipliers [19, 32, 130]. Such alternating optimization methods offer a unifying computational perspective on a broad set of large-scale convex estimation problems in machine learning [73] and signal processing [32], including sparse recovery, deconvolution, and de-mixing [55]. They have numerous universality and scalability benefits and, in most cases, they are well-suited to distributed convex optimization. In this case, one can borrow ideas and tools from gradient descent and Newton schemes in solving the subproblems of the alternating minimization. While their theoretical convergence guarantees are not optimal, they usually work well in practice.

8 Applications

8.1 Compressive Imaging

Natural images are usually sparse in wavelet basis. In this experiment, we study the image reconstruction problem from compressive measurements, where structured sparsity ideas are applied in practice.

For this purpose and given a $p \times p$ natural grayscale image $\mathbf{x} \in \mathbb{R}^{p^2}$, we use the Discrete Wavelet Transform (DWT) with $\log_2(p)$ levels, based on the Daubechies 4 wavelet, to represent \mathbf{x} ; see the Wavelet representation of two images in Figures 13-14. In math terms, the DWT can be represented by an operator matrix \mathbf{W}^\top , so that \mathbf{x} can be sparsely represented (or well-approximated) as $\mathbf{x} = \mathbf{W}\mathbf{c}$, where $\mathbf{c} \in \mathbb{R}^n$, $n := p^2$ are the wavelet coefficients for \mathbf{x} .

To exploit this fact in practice, we consider the problem of recovering $\mathbf{x} \in \mathbb{R}^n$ from m compressive measurements $\mathbf{u} \in \mathbb{R}^m$. The measurements are obtained by applying a sparse matrix $\mathbf{A} \in \mathbb{R}^{m \times p^2}$ to the vectorized image such that:

$$\mathbf{u} = \mathbf{A}\mathbf{x}.$$

Here, \mathbf{A} is the adjacency matrix of an expander graph of degree $d = 8$, so that $\|\mathbf{A}\|_0 = dn$. Thus, the overall measurement operator on the wavelet coefficients is then given by the concatenation of the expander matrix with the DWT: $\mathbf{u} = \mathbf{A}\mathbf{W}\mathbf{c}$, with $\mathbf{c} \approx \hat{\mathbf{c}}$ with $\|\hat{\mathbf{c}}\|_0 \ll m^2$, i.e., \mathbf{x} can be well-approximated by using only a limited number of wavelet coefficients.

We use the following methods for recovering \mathbf{c} from the measurements \mathbf{u} :

$$\begin{aligned} & \underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} && \|\mathbf{u} - \mathbf{A}\mathbf{W}\mathbf{c}\|_2^2 \\ & \text{subject to} && \text{supp}(\mathbf{c}) \in \mathcal{T}_k. \end{aligned} \quad (\text{Rooted Connected Tree model (RC)})$$

$$\begin{aligned} & \underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} && \|\mathbf{c}\|_1 \\ & \text{subject to} && \mathbf{u} = \mathbf{A}\mathbf{W}\mathbf{c}. \end{aligned} \quad (\text{Basis Pursuit (BP)})$$

$$\begin{aligned} & \underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} && \|\mathbf{c}\|_{\text{HGL}} \\ & \text{subject to} && \mathbf{u} = \mathbf{A}\mathbf{W}\mathbf{c}. \end{aligned} \quad (\text{Hierarchical Group Lasso (HGL) pursuit})$$

$$\begin{aligned} & \underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} && \|\mathbf{c}\|_{\text{PC}} \\ & \text{subject to} && \mathbf{u} = \mathbf{A}\mathbf{W}\mathbf{c}. \end{aligned} \quad (\text{Parent-Child Latent Group Lasso (PC) pursuit})$$

$$\begin{aligned} & \underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} && \|\mathbf{c}\|_{\text{FAM}} \\ & \text{subject to} && \mathbf{u} = \mathbf{A}\mathbf{W}\mathbf{c}. \end{aligned} \quad (\text{Family Latent Group Lasso (FAM) pursuit})$$

The RC model is solved via the improved projected gradient descent given in [80] with the projections computed via the dynamic program proposed in [5]. All the remaining methods are solved using the primal-dual method described in [125] which relies on the proximity operator of the associated structure-sparsity inducing penalties. For BP the proximity operator is given by the standard soft-thresholding function. For HGL, we use the algorithm and code given by [75]. For the latent group lasso approaches, PC and FAM, we adopt the duplication strategy proposed in [69, 108], for which the proximity operator reduces to the standard block-wise soft-thresholding on the duplicated variables. All algorithms are written in `Matlab`, except for the HGL proximity operator and the RC projection that are in `C`.

The duplication approach consists in creating a latent vector that contains copies of the original variables. The number of copies is determined by the number of groups that a given variable belongs to. For the Parent-Child model, each node belongs to the four groups that contain each of its children and the group that contains its father. The root has only three children, corresponding to the roots of the horizontal, vertical and diagonal wavelet trees. The leaves belong only to the group that contains their fathers. A simple calculation shows that the latent vector for the PC model contains $2(n - 1)$ variables.

For the Family model, instead, each node belongs to only two groups: the group containing its children and the group containing its siblings and its father. Each leaf belongs to only the group that contains its siblings and its father. Overall, the number of variables in the latent vector for the Family models is equal to $\frac{5n}{4} - 1$.

The duplication approach does not significantly increase the problem size, while it allows an efficient implementation of the proximity operator. Indeed, given that the proximity operator can be computed in closed form over the duplicated variables, this approach is as fast as the hierarchical group lasso one, where the proximity operator is computed via `C` code.

In order to obtain a good performance, both the parent-child and the family model require a proper weighting scheme to penalize groups lower down in the tree, where smaller wavelet coefficients are expected, compared to nodes closer to the root, which normally carry most of the energy of the signals and should be penalized less. We have observed that setting the group weights proportional to the level L of the node of the group closest to the root gives good results. In particular, we set the weights equal to L^2 , with 0 being the root level.

8.1.1 Results

We performed the compressive imaging experiments on both a 256×256 portrait of a woman and a 2048×2048 mountain landscape. Apart from conversion to grayscale and resizing through the `Matlab` function `imresize`, no preprocessing has been carried out. The primal-dual algorithm of [125] has been run up to precision 10^{-5} . We measure the recovery performance in terms of Power Signal to Noise Ratio (PSNR) and relative recovery error ℓ_2 norm as $\frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2}{\|\mathbf{x}\|}$, where $\hat{\mathbf{x}}$ is the estimated image and \mathbf{x} is the true image.

Figures 13-14 report the recovery results, using $m = \frac{n}{8}$, that is using only 12.4% samples compared to the ambient dimension. The estimated images are on the top two rows, while the third and fourth rows show the estimated wavelet coefficients.

The effect of imposing structured sparsity can be clearly seen for the HGL, PC and FAM models, where the high values of the coefficients tend to cluster around the root of the wavelet tree (i.e., top-left corner of the image) and their intensity decreases descending the tree. The family model shows the grouping among the siblings, where four leaves are either all zero or all non-zero. For the 256×256 image, despite being coded in `C`, the discrete model is approximately 160 times slower than the other methods, which are computationally equivalent: e.g., in our tests, the family model took around 60 seconds, while the RC one required almost 2 hours. We therefore did not use the RC model on the larger 2048×2048 mountain image, but we compared also against Total Variation (TV) pursuit, which obtains the best performance on this image.

8.2 Neuronal spike detection from compressed data

In the experiments that follow, we compare the performance of the following three optimization criteria, assuming the dispersive model \mathcal{D}_k .

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{x} \in \mathcal{D}_k. \quad (\text{Discrete dispersive})$$

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\mathbf{x}) + \lambda \cdot \Omega_{\text{exclusive}}(\mathbf{x}). \quad (\text{Exclusive norm regularization})$$

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad \Omega_{\text{exclusive}}(\mathbf{x}) \quad \text{subject to} \quad f(\mathbf{x}) \leq \sigma^2. \quad (\text{Exclusive norm pursuit})$$

Empirical performance on synthetic data: Figures 15-16 illustrate the utility of each approach in the compressed sensing setting where $f(\mathbf{x}) := \frac{1}{2} \|\mathbf{u} - \Phi \mathbf{x}\|_2^2$. That is, we observe $\mathbf{x}^* \in \mathbb{R}^n$ through a limited set of linear sketches $\mathbf{u} = \Phi \mathbf{x}^* + \varepsilon \in \mathbb{R}^m$ where $\Phi \in \mathbb{R}^{m \times n}$ is a known linear sketch matrix. Here, we assume $n = 500$ and $m = 70$ for $\|\mathbf{x}^*\|_0 = 25$. Without loss of generality, we assume $(\mathbf{x}^*)_i \geq 0, \forall i$ and $\Delta^* = 20$.

In the discrete case, we relax the refractory period Δ to model signal structure deviations; here, we assume $\Delta = 15$. The discrete exclusive model [8, 99] clearly outperforms the rest of the approaches under comparison;

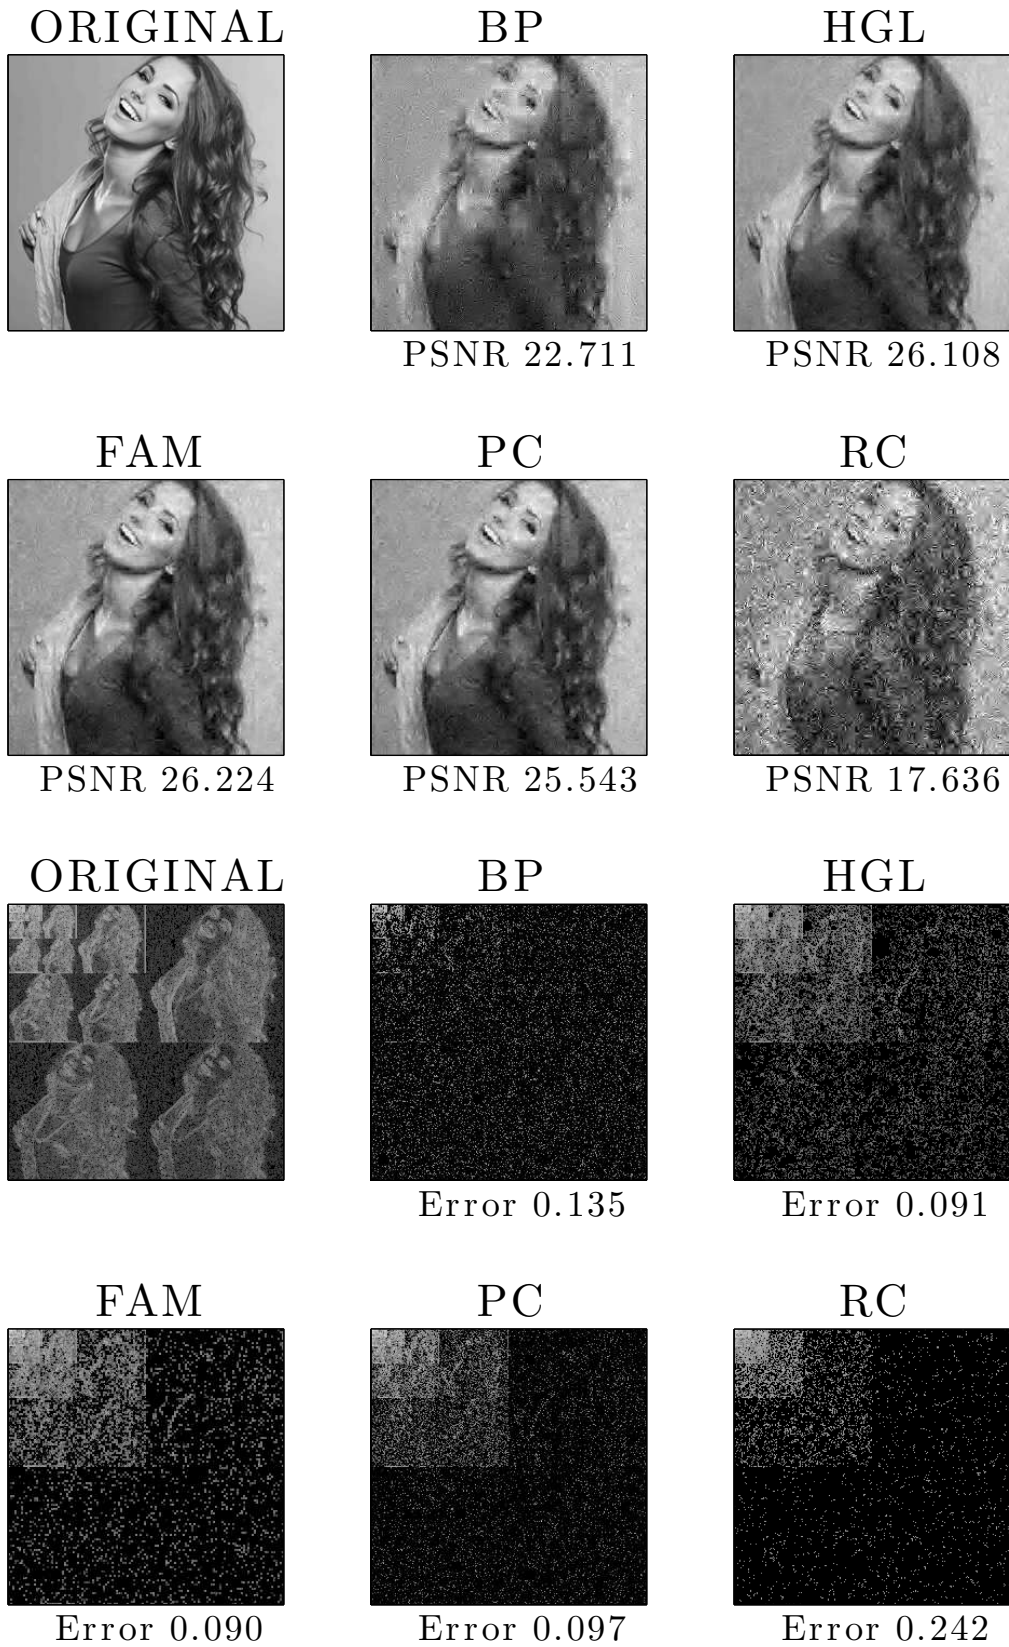


Figure 13: Woman image recovery performance from compressive measurements. Here, $p = 256$. The top two rows show the reconstruction performance in the original domain, along with the PSNR levels achieved. The bottom two rows show the corresponding representations into the wavelet domain.

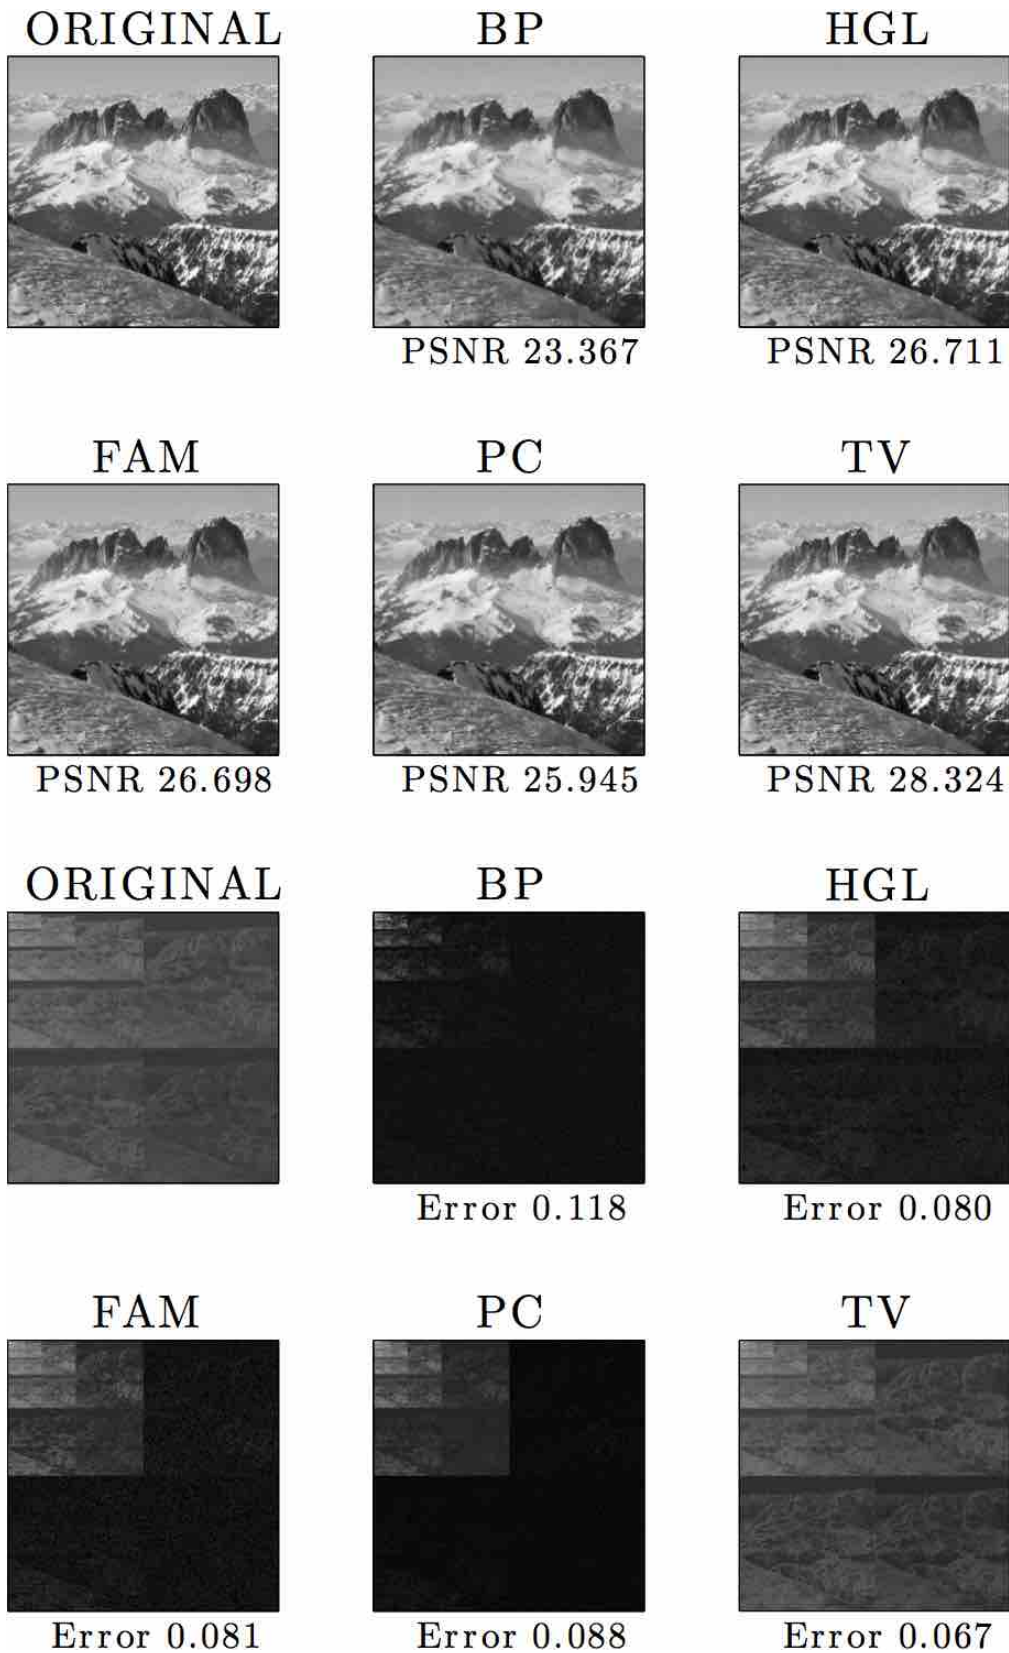


Figure 14: Mountains image recovery performance from compressive measurements. Here, $p = 2048$. The top two rows show the reconstruction performance in the original domain, along with the PSNR levels achieved. The bottom two rows show the corresponding representations into the wavelet domain.




Figure 15: Performance of the discrete dispersive approach for the problem spike train recovery from a limited set of linear measurements.

such behavior is also observed on average over the set of experiments conducted (Figure 15). This also implies that the discrete model usually requires fewer measurements for accurate recovery compared to conventional sparse approximation, as long as the underlying signal approximately follows \mathcal{D}_k .

On the other hand, due to convex relaxations, convex approaches introduce unnecessary nonzero coefficients that do not comply with the underlying model. However, both approaches show good performance in recovering \mathbf{x}^* from limited measurements; see Figure 16.

Real neuronal spike data: In order to understand the functioning of the human brain, it is necessary to identify and study the behavior of neuronal cell membranes under rapid change in the electric potential. However, to observe such phenomena, electrical activities on neurons need to be recorded using specialized equipment. In this experiment, we perform somatic spike detection of a tufted L5 pyramidal cell responding to in-vivo-like current injected in the apical dendrites and the soma simultaneously (see [45] for the experimental details).

A snapshot of the neuronal spikes waveforms is shown in Figure 17. In order to accurately detect the neuronal spikes, a high-frequency sample acquisition equipment is required. Within this context, we apply CS ideas to decrease the number of samples needed to approximately detect the *positions* of the spike train. Let $\mathbf{x}^* \in \mathbb{R}^n$ with $n = 832$ represent the signal in Figure 17a; furthermore, let $\Phi \in \mathbb{R}^{m \times n}$ be the sensing matrix where $m = 0.25 \cdot n$, i.e., we perform a 75% compression.

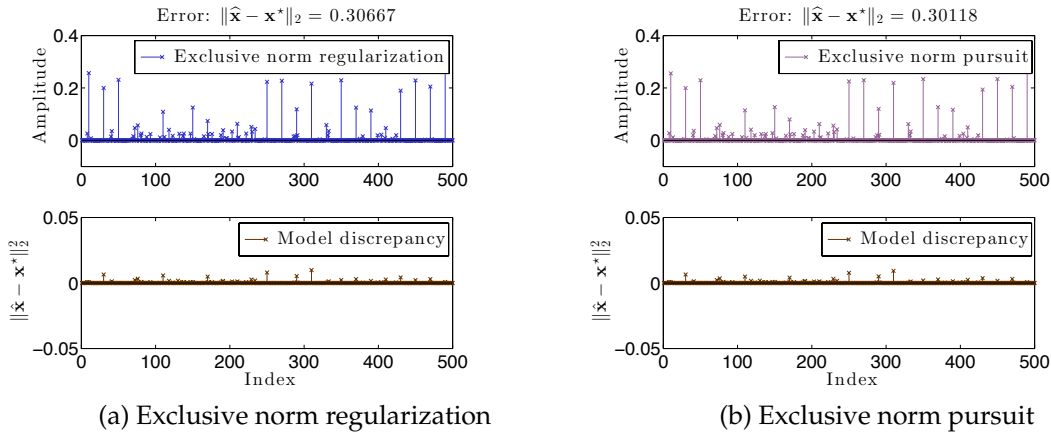


Figure 16: Performance of dispersive convex relaxations for the problem of spike train recovery from a limited set of linear measurements. **(Left panel:)** Exclusive norm regularization approach. **(Right panel:)** Exclusive norm pursuit approach.

We use the proposed schemes to recover the locations of the neuronal spikes under the assumption of the dispersive model with refractory period Δ . Here, Δ is set equal to the *average* period between two consecutive spikes.

Figure 17b represents the recovered k -sparse approximation using the discrete dispersive model \mathcal{D}_k . Here, k is set to the number of spikes expected to appear for a given time period—such number can be easily deduced by observing the behavior of a specific neuron type. From Figure 17b, we observe that the discrete model approximates the locations of the spikes quite accurately: most of the spike locations are exactly recovered.

However, due to the “strictness” of the discrete model, we observe that small deviations from \mathcal{D}_k lead to imprecise estimations; e.g., between the 12th and 13th spike of the sequence, a larger (than usual) refractory period is observed that leads to mis-location of the next spike estimation.

Figures 17c-17d depict the performance of convex solvers using the exclusive norm as (i) regularizer and (ii) objective function. Tweaking the λ parameter in the (i) case, one can achieve *sparse* solutions that approximate the underlying model (Figure 17c); however, one can observe multiple detected spikes with separation less than Δ , violating the assumed model. In the model-based Basis pursuit case, the solver tries to *fit* the solution to the data, which usually leads to less sparse solutions (Figure 17d). One can further sparsify the convex solutions to obtain a k -sparse answer as in Figures 17e-17f: however, in most cases, further processing of the returned signal is required to maintain a \mathcal{D}_k -modeled solution. E.g., in this case, due to the fact that convex norms force the solution to fully *explain* the observations, the sparsified solution includes more than one spike per true spike location.

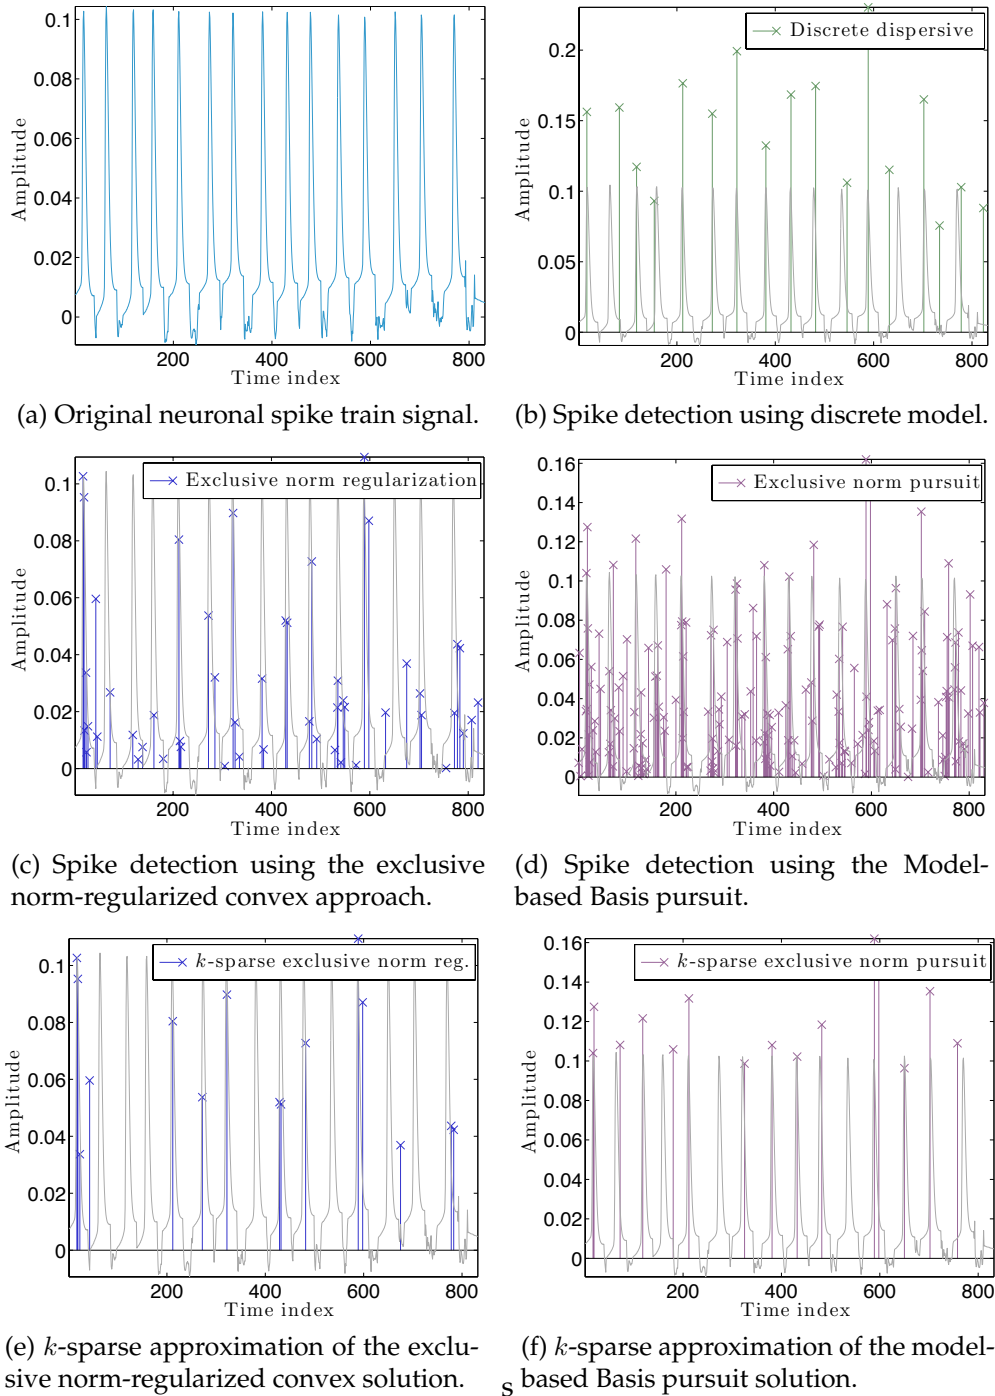


Figure 17: Spike detection in real neuronal data using the dispersive model. Figure 17a depicts the original signal \mathbf{x}^* . We observe $\mathbf{u} = \Phi\mathbf{x}^*$ using only 25% measurements through a linear sketch Φ . Figures 17b-17d illustrate the performance of the three approaches under comparison. Figures 17e-17f show the convex solutions, sparsified to be k -sparse.

8.3 Digital Confocal Imaging

Confocal imaging [95, 117] has become one of the best techniques for 3D imaging, due to its ability to reduce the blur caused by out-of-focus scattering by means of focus light and a physical pinhole. The recent work of Goy et al. [56] combines this technique with digital holography which allows access to the complex field at the recording CCD sensor and enables a novel and flexible way of eliminating the out-of-focus contribution via a so-called “virtual pinhole” in the digital domain. Under the assumption that the scattered light is not distorted by the matter in the light cones around the focus, this makes it possible to obtain very sharp images of the sample at all depths.

We are interested in modeling the scattering process so that we can instead leverage the information carried by the out-of-focus scattering for reducing the number of focusing locations that are needed for a faithful imaging of the sample. We create a linear model by adopting the classical first Born approximation for the scattering [17]. Our model takes into account the specific form of the Gaussian beam used to focus the light inside the sample, the numerical aperture of the transmission lens that determines which fraction of the scattered light is captured and the resolution and size of the CCD that records the hologram. For simplicity, we model the scattered field as it would arrive on the CCD and not its interaction with the reference field normally used to obtain the hologram. We use a 3D discrete model for representing the imaged sample, where one 3D pixel, or voxel, contains the average value of the scattering potential in the volume contained therein.

For a specific focus location i , the measurement model \mathbf{A}_i maps a 3D object represented by its scattering potential \mathbf{F} to a 2D complex field \mathbf{S} . The complete measurement model \mathbf{A} is then given by the concatenation of \mathbf{A}_i for all focus points

$$\mathbf{A}(\mathbf{F}) = \begin{bmatrix} \mathbf{A}_1(\mathbf{F}) \\ \mathbf{A}_2(\mathbf{F}) \\ \vdots \\ \mathbf{A}_P(\mathbf{F}) \end{bmatrix},$$

where P is the total number of focus points.

For the experiments in this section, we consider a $64 \times 8 \times 64$ volume filled with oil with refractive index 1.52, containing two concentric cylinders with axis aligned with the y axis. Their refractive indexes are 1.54 for the outer one and 1.56 for the inner one. A x - z cross section of their scattering potential is shown in Figure 18 (Left). Note that \mathbf{F} is equal to zero outside the cylinder and has constant value inside each. These properties could be captured via sparsity and clusteredness, for example via ℓ_1 norm or Total Variation minimization or via the Ising discrete model; see Section 6. We choose 16 focus locations taken from a uniform grid over the central x - z plane. The focusing lens and the transmission lens have numerical aperture 1 and 0.65 respectively, while the coherent light of the Gaussian beam has wavelength $\lambda = 405 \cdot 10^{-9}$. We consider a 91×91 CCD that covers the entire field of view of the transmission lens. We simulate the measurements by applying the measurement model to the synthetic scattering potential \mathbf{F} , that is $\mathbf{u} = \mathbf{A}(\mathbf{F})$, where we have $\mathbf{F} \in \mathbb{R}^{64 \times 8 \times 64}$ and $\mathbf{u} \in \mathbb{C}^{16 \times 91 \times 91}$. Despite the apparent oversampling ratio, the problem is still ill-posed, because each recorded field carries information only about a small neighborhood around each focus point.

We compare four methods for estimating \mathbf{F} from the measurements \mathbf{u} . Let the data-fit term be the square loss $L(\mathbf{F}) = \frac{1}{2} \|\mathbf{u} - \mathbf{A}(\mathbf{F})\|_2^2$.

$$\underset{\mathbf{F}}{\text{minimize}} \quad L(\mathbf{F}) . \quad (\text{Conjugate Gradient (CG)})$$

$$\underset{\mathbf{F}}{\text{minimize}} \quad \|\mathbf{F}\|_1 \quad \text{subject to} \quad \mathbf{u} = \mathbf{A}(\mathbf{F}) . \quad (\text{Basis Pursuit (BP)})$$

$$\underset{\mathbf{F}}{\text{minimize}} \quad \|\mathbf{F}\|_{TV} \quad \text{subject to} \quad \mathbf{u} = \mathbf{A}(\mathbf{F}) . \quad (\text{Total Variation (TV) pursuit})$$

$$\underset{\mathbf{F}}{\text{minimize}} \quad L(\mathbf{F}) + \lambda R_{\text{ISING}}(\text{supp}(\mathbf{F})) + \tau |\text{supp}(\mathbf{F})| , \quad (\text{Ising plus Cardinality (IC)})$$

In the last approach, one has two regularization parameters $\lambda, \tau \geq 0$ to be set.

8.3.1 Results

We evaluate the recovery performance of the methods using the relative recovery error $\|\widehat{\mathbf{F}} - \mathbf{F}^*\|_2 / \|\mathbf{F}^*\|_2$, where \mathbf{F}^* is the synthetic scattering potential used to simulate the output \mathbf{y} . We explored several pairs of regularization parameters λ and τ for the discrete IC model, selecting the pairs that gave the best performance. Of course, in practice one cannot resort to this type of parameter exploration and needs to choose them according to knowledge regarding the noise level and the desired amount of clusteredness. Furthermore, this method is very sensitive to changes in the regularization parameters, making it difficult to find the best value for them.

Figure 18 (right) reports the results on this experimental setup. Total Variation pursuit achieves the best performance without having to tune any parameter, closely followed by the discrete model. Despite the scattering potential is also sparse, Basis Pursuit enforces too much sparsity, especially in the regions of the sample that lie farthest away from the focus points and hence do not contribute to the the output \mathbf{u} . It is interesting to note, that in this case, enforcing the incorrect structure leads to poorer performance than unstructured recovery.

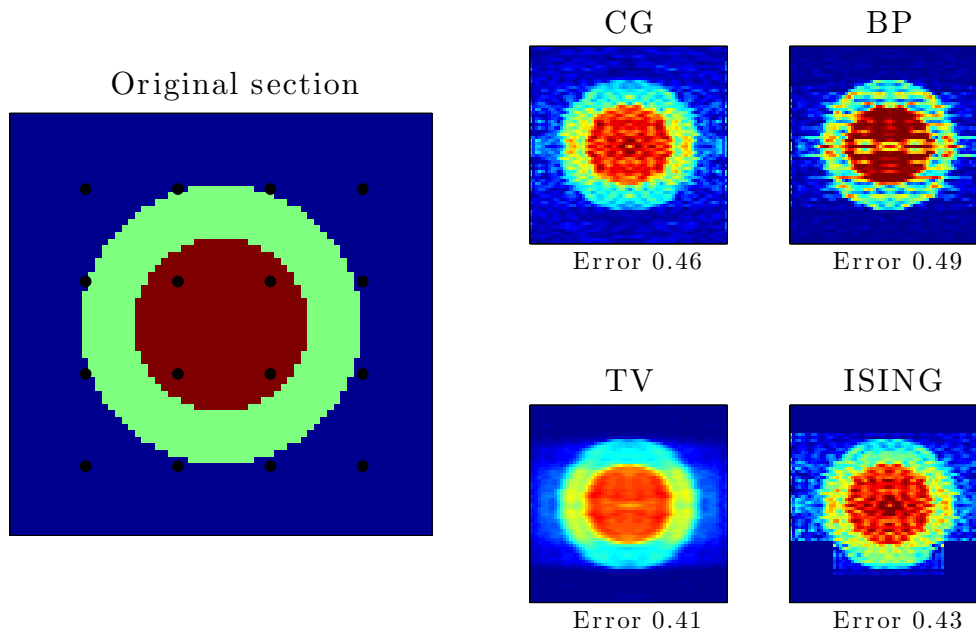


Figure 18: (Left) Central section of the sample to be reconstructed. The black dots represent the 16 focus locations. (Right) Reconstruction results for the four methods.

References

- [1] A. Argyriou, C. Micchelli, M. Pontil, L. Shen, and Y. Xu. Efficient first order methods for linear composite regularizers. *arXiv preprint arXiv:1104.1436*, 2011.
- [2] F. Bach. Structured sparsity-inducing norms through submodular functions. In *NIPS*, pages 118–126, 2010.
- [3] F. Bach. Learning with submodular functions: A convex optimization perspective. *arXiv preprint arXiv:1111.6453*, 2011.
- [4] B. Bah, L. Baldassarre, and V. Cevher. Model-based sketching and recovery with expanders. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, number EPFL-CONF-187484, 2014.
- [5] L. Baldassarre, N. Bhan, V. Cevher, and A. Kyrillidis. Group-sparse model selection: Hardness and relaxations. *arXiv preprint arXiv:1303.3207*, 2013.
- [6] R. Baraniuk. Optimal tree approximation with wavelets. In *SPIE's International Symposium on Optical Science, Engineering, and Instrumentation*, pages 196–207. International Society for Optics and Photonics, 1999.
- [7] R. Baraniuk, R. DeVore, G. Kyriazis, and X. Yu. Near best tree approximation. *Advances in Computational Mathematics*, 16(4):357–373, 2002.
- [8] R.G. Baraniuk, V. Cevher, M.F. Duarte, and C. Hegde. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- [9] R.G. Baraniuk, V. Cevher, and M.B. Wakin. Low-dimensional models for dimensionality reduction and signal recovery: A geometric perspective. *Proceedings of the IEEE*, 98(6):959–971, 2010.
- [10] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [11] Y. Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.
- [12] D. Bertsekas. Projected newton methods for optimization problems with simple constraints. *SIAM Journal on control and Optimization*, 20(2):221–246, 1982.
- [13] N. Bhan, L. Baldassarre, and V. Cevher. Tractability of interpretability via selection of group-sparse models. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, 2013.
- [14] T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- [15] T. Blumensath and M. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *Information Theory, IEEE Transactions on*, 55(4):1872–1882, 2009.

- [16] J.F. Bonnans. Local Analysis of Newton-Type Methods for Variational Inequalities and Nonlinear Programming. *Appl. Math. Optim*, 29:161–186, 1994.
- [17] M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. CUP Archive, 1999.
- [18] J. Borwein and A. Lewis. *Convex analysis and nonlinear optimization: theory and examples*. Springer, 2006.
- [19] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [20] S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [21] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 649–658. IEEE, 2012.
- [22] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. Submodular maximization with cardinality constraints. SIAM, 2014.
- [23] E. Candes, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, 2013.
- [24] E. J. Candes. Compressive sampling. In *Proceedings of the International Congress of Mathematicians: Madrid, August 22-30, 2006: invited lectures*, pages 1433–1452, 2006.
- [25] C. Cartis and A. Thompson. An exact tree projection algorithm for wavelets. *arXiv preprint arXiv:1304.4570*, 2013.
- [26] V. Cevher, C. Hegde, M.F. Duarte, and R.G. Baraniuk. Sparse signal recovery using markov random fields. In *NIPS*, 2009.
- [27] V. Cevher, S. Jafarpour, and A. Kyrillidis. Linear inverse problems with norm and sparsity constraints. *Practical Applications of Sparse Modeling*, page 179, 2014.
- [28] A. Chambolle, R. De Vore, N. Lee, and B. Lucier. Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage. *Image Processing, IEEE Transactions on*, 7(3):319–335, 1998.
- [29] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [30] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12:805–849, 2012.
- [31] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1998.
- [32] P. Combettes and V. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
- [33] M. Crouse, R. Nowak, and R. Baraniuk. Wavelet-based statistical signal processing using hidden markov models. *Signal Processing, IEEE Transactions on*, 46(4):886–902, 1998.
- [34] J. Dahl, L. Vandenberghe, and V. Roychowdhury. Covariance selection for nonchordal graphs via chordal embedding. *Optimization Methods & Software*, 23(4):501–520, 2008.
- [35] A. Dalalyan, M. Hebiri, K. Meziani, and J. Salmon. Learning heteroscedastic models by convex programming under group sparsity. In *Proceedings of The 30th International Conference on Machine Learning*, pages 379–387, 2013.
- [36] A. Das, A. Dasgupta, and R. Kumar. Selecting diverse features via spectral regularization. In *NIPS*, pages 1592–1600, 2012.
- [37] A. Das and D. Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *arXiv preprint arXiv:1102.3975*, 2011.
- [38] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on pure and applied mathematics*, 57(11):1413–1457, 2004.
- [39] Q. Tran Dinh, A. Kyrillidis, and V. Cevher. An inexact proximal path-following algorithm for constrained convex minimization. *SIAM Journal on Optimization*, 2013.
- [40] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.
- [41] S. Dughmi. Submodular functions: Extensions, distributions, and algorithms. a survey. *arXiv preprint arXiv:0912.0322*, 2009.
- [42] B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, et al. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

- [43] M. El Halabi, L. Baldassarre, and V. Cevher. To convexify or not? regression with clustering penalties on graphs. In *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013 IEEE 5th International Workshop on*, pages 21–24. IEEE, 2013.
- [44] Y.C. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *Information Theory, IEEE Transactions on*, 55(11):5302–5316, 2009.
- [45] International Neuroinformatics Coordinating Faculty. Spike time prediction - challenge c, 2009.
- [46] S. Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49(6):2543–2563, 2011.
- [47] J. Friedman, T. Hastie, and R. Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- [48] S. Fujishige and S. Isotani. A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, 7(1):3–17, 2011.
- [49] S. Fujishige and S. Patkar. Realization of set functions as cut functions of graphs and hypergraphs. *Discrete Mathematics*, 226(1):199–210, 2001.
- [50] M. Fukushima and H. Mine. A generalized proximal point algorithm for certain non-convex minimization problems. *International Journal of Systems Science*, 12(8):989–1000, 1981.
- [51] W Gerstner and W. Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.
- [52] A. Gilbert and P. Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010.
- [53] F. Girosi. An equivalence between sparse approximation and support vector machines. *Neural computation*, 10(6):1455–1480, 1998.
- [54] A. Goldberg and S. Rao. Beyond the flow decomposition barrier. *J. ACM*, 45(5):783–797, September 1998.
- [55] T. Goldstein, B. Donoghue, and S. Setzer. Fast alternating direction optimization methods. *CAM report*, pages 12–35, 2012.
- [56] A. Goy and D. Psaltis. Digital confocal microscope. *Opt. Exp*, 20(20):22720, 2012.
- [57] A. Gramfort and M. Kowalski. Improving m/eeg source localization with an inter-condition sparse prior. In *IEEE International Symposium on Biomedical Imaging*, 2009.
- [58] V. Guigue, A. Rakotomamonjy, and S. Canu. Kernel basis pursuit. In *Machine Learning: ECML 2005*, pages 146–157. Springer, 2005.
- [59] Z. Harmany, R. Marcia, and R. Willett. This is spiral-tap: sparse poisson intensity reconstruction algorithm theory and practice. *Image Processing, IEEE Transactions on*, 21(3):1084–1096, 2012.
- [60] B.S. He and X.M. Yuan. On the $O(1/n)$ convergence rate of the Douglas-Rachford alternating direction method. *SIAM J. Numer. Anal.*, 50:700–709, 2012.
- [61] L. He and L. Carin. Exploiting structure in wavelet-based bayesian compressive sensing. *Signal Processing, IEEE Transactions on*, 57(9):3488–3497, 2009.
- [62] D. Heckerman, D. Geiger, and D. Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.
- [63] C. Hegde, M. Duarte, and V. Cevher. Compressive sensing recovery of spike trains using a structured sparsity model. In *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- [64] C. Hsieh, M. Sustik, I. Dhillon, and P. Ravikumar. Sparse inverse covariance matrix estimation using quadratic approximation. In *NIPS*, pages 2330–2338, 2011.
- [65] C. Hsieh, M. Sustik, I. Dhillon, P. Ravikumar, and R. Poldrack. Big & quic: Sparse inverse covariance estimation for a million variables. In *Advances in Neural Information Processing Systems*, pages 3165–3173, 2013.
- [66] J. Huang and T. Zhang. The benefit of group sparsity. *The Annals of Statistics*, 38(4):1978–2004, 2010.
- [67] J. Huang, T. Zhang, and D. Metaxas. Learning with structured sparsity. *The Journal of Machine Learning Research*, 12:3371–3412, 2011.
- [68] P. Indyk and I. Razenshteyn. On model-based rip-1 matrices. In *Automata, Languages, and Programming*, pages 564–575. Springer, 2013.
- [69] L. Jacob, G. Obozinski, and J.P. Vert. Group lasso with overlap and graph lasso. In *International Conference on Machine Learning*, 2009.
- [70] K. Jaganathan, S. Oymak, and B. Hassibi. Sparse phase retrieval: Uniqueness guarantees and recovery algorithms. *arXiv preprint arXiv:1311.2745*, 2013.

- [71] A. Jalali, P. Ravikumar, V. Vasuki, and S. Sanghavi. On learning discrete graphical models using group-sparse regularization. In *International Conference on Artificial Intelligence and Statistics*, pages 378–387, 2011.
- [72] S. Jegelka, H. Lin, and J. Bilmes. On fast approximate submodular minimization. In *NIPS*, pages 460–468, 2011.
- [73] R. Jenatton, J.-Y. Audibert, and F. Bach. Structured variable selection with sparsity-inducing norms. *Journal of Machine Learning Research*, 12:2777–2824, 2011.
- [74] R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, F. Bach, and B. Thirion. Multi-scale mining of fmri data with hierarchical structured sparsity. In *Pattern Recognition in NeuroImaging (PRNI)*, 2011.
- [75] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12:2297–2334, 2011.
- [76] I. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of statistics*, pages 295–327, 2001.
- [77] S. Kim and E. Xing. Tree-guided group lasso for multi-task regression with structured sparsity. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 543–550, 2010.
- [78] V. Kolmogorov and R. Zabini. What energy functions can be minimized via graph cuts? *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(2):147–159, 2004.
- [79] A. Krause and V. Cevher. Submodular dictionary selection for sparse representation. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 567–574, 2010.
- [80] A. Kyrillidis and V. Cevher. Recipes on hard thresholding methods. In *4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, number EPFL-CONF-183059, 2011.
- [81] A. Kyrillidis and V. Cevher. Combinatorial selection and least absolute shrinkage via the clash algorithm. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pages 2216–2220. Ieee, 2012.
- [82] A. Kyrillidis and V. Cevher. Fast proximal algorithms for self-concordant function minimization with application to sparse graph selection. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6585–6589. IEEE, 2013.
- [83] A. Kyrillidis, R. Karimi Mahabadi, Q. Tran Dinh, and V. Cevher. Scalable sparse covariance estimation via self-concordance. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.
- [84] A. Kyrillidis, G. Puy, and V. Cevher. Hard thresholding with norm constraints. In *2012 Ieee International Conference On Acoustics, Speech And Signal Processing (Icassp)*, number EPFL-CONF-183061, pages 3645–3648. Ieee, 2012.
- [85] J. Lee and T. Hastie. Structure learning of mixed graphical models. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pages 388–396, 2013.
- [86] E. Lein, M. Hawrylycz, N. Ao, M. Ayres, A. Bensinger, A. Bernard, A. Boe, M. Boguski, K. Brockway, E. Byrnes, et al. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, 445(7124):168–176, 2007.
- [87] P. Loh, M. Wainwright, et al. Structure estimation for discrete graphical models: Generalized covariance matrices and their inverses. *The Annals of Statistics*, 41(6):3022–3049, 2013.
- [88] L. Lovász. Submodular functions and convexity. In *Mathematical Programming The State of the Art*, pages 235–257. Springer, 1983.
- [89] M. Lustig, D. Donoho, and J. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine*, 58(6):1182–1195, 2007.
- [90] S. Mallat. *A wavelet tour of signal processing*. Academic press, 1999.
- [91] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41(12):3397–3415, 1993.
- [92] B. McCoy and T. Wu. *{The Two-Dimensional Ising Model}*. Harvard University Press, 1973.
- [93] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- [94] RP Millane. Phase retrieval in crystallography and optics. *JOSA A*, 7(3):394–411, 1990.
- [95] M. Minsky. Microscopy apparatus, December 19 1961. US Patent 3,013,467.
- [96] S. Mosci, S. Villa, A. Verri, and L. Rosasco. A primal-dual algorithm for group ℓ_1 regularization with overlapping groups. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [97] M. Narasimhan, N. Jojic, and J. Bilmes. Q-clustering. In *NIPS*, volume 5, page 5, 2005.

- [98] B. Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- [99] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [100] G. Nemhauser and L. Wolsey. *Integer and combinatorial optimization*, volume 18. Wiley New York, 1988.
- [101] G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher. An analysis of approximations for maximizing submodular set functions — I. *Mathematical Programming*, 14(1):265–294, 1978.
- [102] A. Nemirovskii. Prox-method with rate of convergence $\mathcal{O}(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. Op*, 15(1):229–251, 2004.
- [103] Y. Nesterov. A method of solving a convex programming problem with convergence rate $\mathcal{O}(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
- [104] Y. Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM J. Optimization*, 16(1):235–249, 2005.
- [105] Y. Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, 2005.
- [106] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Math. Program.*, 120(1, Ser. B):221–259, 2009.
- [107] G. Obozinski and F. Bach. Convex relaxation for combinatorial penalties. *arXiv preprint arXiv:1205.1240*, 2012.
- [108] G. Obozinski, L. Jacob, and J.P. Vert. Group lasso with overlaps: The latent group lasso approach. *arXiv preprint arXiv:1110.0413*, 2011.
- [109] J.B. Orlin. A faster strongly polynomial time algorithm for submodular function minimization. *Mathematical Programming*, 118(2), 2009.
- [110] N.S. Rao, R.D. Nowak, S.J. Wright, and N.G. Kingsbury. Convex approaches to model wavelet sparsity patterns. In *Image Processing (ICIP), 2011 18th IEEE International Conference on*, pages 1917–1920, 2011.
- [111] N.S. Rao, B. Recht, and R.D. Nowak. Signal recovery in unions of subspaces with applications to compressive imaging. *arXiv preprint arXiv:1209.3079*, 2012.
- [112] F. Rapaport, E. Barillot, and J.P. Vert. Classification of arraycgh data using fused svm. *Bioinformatics*, 24(13):i375–i382, 2008.
- [113] S. M. Robinson. Strongly Regular Generalized Equations. *Mathematics of Operations Research*, Vol. 5, No. 1 (Feb., 1980), pp. 43–62, 5:43–62, 1980.
- [114] M. Schmidt, N.L. Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. *NIPS, Granada, Spain*, 2011.
- [115] M. Seeger. On the submodularity of linear experimental design. Technical report, 2009.
- [116] J. Shapiro. Embedded image coding using zerotrees of wavelet coefficients. *Signal Processing, IEEE Transactions on*, 41(12):3445–3462, 1993.
- [117] C. Sheppard and D. Shotton. *Confocal laser scanning microscopy*. BIOS Scientific Publishers, 1997.
- [118] N. Simon, J. Friedman, T. Hastie, and R Tibshirani. A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231–245, 2013.
- [119] M. Stojnic, F. Parvaresh, and B. Hassibi. On the reconstruction of block-sparse signals with an optimal number of measurements. *Signal Processing, IEEE Transactions on*, 57(8):3075–3085, 2009.
- [120] M. Stojnic, W. Xu, and B. Hassibi. Compressed sensing-probabilistic analysis of a null-space characterization. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 3377–3380. IEEE, 2008.
- [121] A. Subramanian, P. Tamayo, V.K. Mootha, S. Mukherjee, B.L. Ebert, M.A. Gillette, A. Paulovich, S.L. Pomeroy, T.R. Golub, E.S. Lander, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 102(43):15545–15550, 2005.
- [122] Z. Szabó, B. Póczos, and A. Lorincz. Online group-structured dictionary learning. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2865–2872. IEEE, 2011.
- [123] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [124] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.

- [125] Q. Tran-Dinh and V. Cevher. An optimal primal-dual decomposition framework. Technical report, LIONS - EPFL, 2014.
- [126] Q. Tran-Dinh and V. Cevher. A unified optimal primal-dual framework for constrained convex minimization. *Tech. Report., LIONS*, pages 1–32, 2014.
- [127] Q. Tran Dinh, A. Kyrillidis, and V. Cevher. A proximal Newton framework for composite minimization: Graph learning without Cholesky decompositions and matrix inversions. In *Proceedings of The 30th International Conference on Machine Learning*, pages 271–279, 2013.
- [128] Q. Tran-Dinh, A. Kyrillidis, and V. Cevher. Composite self-concordant minimization. *Journal of Machine Learning Research*, 16:371–416, 2015.
- [129] J. Tropp and A. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 53(12):4655–4666, 2007.
- [130] P. Tseng. Applications of splitting algorithm to decomposition in convex programming and variational inequalities. *SIAM J. Control Optim.*, 29:119–138, 1991.
- [131] S. Villa, L. Rosasco, S. Mosci, and A. Verri. Proximal methods for the latent group lasso penalty. *Computational Optimization and Applications*, pages 1–27, 2012.
- [132] S. Villa, S. Salzo, L. Baldassarre, and A. Verri. Accelerated and inexact forward-backward algorithms. *SIAM Journal on Optimization*, 23(3):1607–1633, 2013.
- [133] M. Vincent and N. Hansen. Sparse group lasso and high dimensional multinomial classification. *Computational Statistics & Data Analysis*, 71:771–786, 2014.
- [134] S. Wright, R. Nowak, and M. Figueiredo. Sparse reconstruction by separable approximation. *Signal Processing, IEEE Transactions on*, 57(7):2479–2493, 2009.
- [135] SJ Wright and J Nocedal. *Numerical optimization*, volume 2. Springer New York, 1999.
- [136] L. Yuan, J. Liu, and J. Ye. Efficient methods for overlapping group lasso. In *NIPS*, volume 35, pages 352–360, 2011.
- [137] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [138] X. Zeng and M. Figueiredo. A novel sparsity and clustering regularization. *arXiv preprint arXiv:1310.4945*, 2013.
- [139] Z. Zhang, Y. Shi, and B. Yin. Mr images reconstruction based on tv-group sparse model. In *Multimedia and Expo (ICME), 2013 IEEE International Conference on*, pages 1–6. IEEE, 2013.
- [140] P. Zhao, G. Rocha, and B. Yu. The composite absolute penalties family for grouped and hierarchical variable selection. *The Annals of Statistics*, 37(6A):3468–3497, 2009.
- [141] H. Zhou, M.E. Sehl, J.S. Sinsheimer, and K. Lange. Association screening of common and rare genetic variants by penalized regression. *Bioinformatics*, 26(19):2375, 2010.
- [142] Y. Zhou, R. Jin, and S. Hoi. Exclusive lasso for multi-task feature selection. In *International Conference on Artificial Intelligence and Statistics*, pages 988–995, 2010.