Low Complexity Regularization of Linear Inverse Problems

Samuel Vaïter, Gabriel Peyré, Jalal M. Fadili

► To cite this version:

Samuel Vaïter, Gabriel Peyré, Jalal M. Fadili. Low Complexity Regularization of Linear Inverse Problems. Sampling Theory, a Renaissance, Pfander, Götz E. (Ed.), 50 p., 2015, 978-3-319-19748-7. 10.1007/978-3-319-19749-4 . hal-01018927v3

HAL Id: hal-01018927

<https://hal.science/hal-01018927v3>

Submitted on 20 May 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Low Complexity Regularization of Linear Inverse Problems

Samuel Vaïter, Gabriel Peyré, Jalal M. Fadili

► To cite this version:

Samuel Vaïter, Gabriel Peyré, Jalal M. Fadili. Low Complexity Regularization of Linear Inverse Problems. [Research Report] CNRS. 2014. <hal-01018927v2>

HAL Id: hal-01018927

<https://hal.archives-ouvertes.fr/hal-01018927v2>

Submitted on 7 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Low Complexity Regularization of Linear Inverse Problems

Samuel Vaiter, Gabriel Peyré and Jalal Fadili

Abstract Inverse problems and regularization theory is a central theme in imaging sciences, statistics and machine learning. The goal is to reconstruct an unknown vector from partial indirect, and possibly noisy, measurements of it. A now standard method for recovering the unknown vector is to solve a convex optimization problem that enforces some prior knowledge about its structure. This chapter delivers a review of recent advances in the field where the regularization prior promotes solutions conforming to some notion of simplicity/low-complexity. These priors encompass as popular examples sparsity and group sparsity (to capture the compressibility of natural signals and images), total variation and analysis sparsity (to promote piecewise regularity), and low-rank (as natural extension of sparsity to matrix-valued data). Our aim is to provide a unified treatment of all these regularizations under a single umbrella, namely the theory of partial smoothness. This framework is very general and accommodates all low-complexity regularizers just mentioned, as well as many others. Partial smoothness turns out to be the canonical way to encode low-dimensional models that can be linear spaces or more general smooth manifolds. This review is intended to serve as a one stop shop toward the understanding of the theoretical properties of the so-regularized solutions. It covers a large spectrum including: (i) recovery guarantees and stability to noise, both in terms of ℓ^2 -stability and model (manifold) identification; (ii) sensitivity analysis to perturbations of the parameters involved (in particular the observations), with applications to unbiased risk estimation ; (iii) convergence properties of the forward-backward proximal splitting scheme, that is particularly well suited to solve the corresponding large-scale regularized optimization problem.

Samuel Vaiter
Ceremade, Université Paris-Dauphine, e-mail: samuel.vaiter@ceremade.dauphine.fr

Gabriel Peyré
CNRS and Ceremade, Université Paris-Dauphine, e-mail: gabriel.peyre@ceremade.dauphine.fr

Jalal Fadili
GREYC, CNRS-ENSICAEN-Université de Caen, e-mail: Jalal.Fadili@greyc.ensicaen.fr

1 Inverse Problems and Regularization

In this chapter, we deal with finite-dimensional linear inverse problems.

1.1 Forward Model

Let $x_0 \in \mathbb{R}^N$ be the unknown vector of interest. Suppose that we observe a vector $y \in \mathbb{R}^P$ of P linear measurements according to

$$y = \Phi x_0 + w, \quad (1)$$

where $w \in \mathbb{R}^P$ is a vector of unknown errors contaminating the observations. The forward model (1) offers a model for data acquisition that describes a wide range of problems in data processing, including signal and image processing, statistics, and machine learning. The linear operator $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^P$, assumed to be known, is typically an idealization of the acquisition hardware in imaging science applications, or the design matrix in a parametric statistical regression problem. The noise w can be either deterministic (in this case, one typically assumes to know some bound on its ℓ^2 norm $\|w\|$), or random (in which case its distribution is assumed to be known). Except in Sections 4.4 and 5.3 where the noise is explicitly assumed random, w is deterministic throughout the rest of the chapter. We refer to [191] and [22] for a comprehensive account on noise models in imaging systems.

Solving an inverse problem amounts to recovering x_0 , to a good approximation, knowing y and Φ according to (1). Unfortunately, the number of measurements P can be much smaller than the ambient dimension N of the signal. Even when $P = N$, the mapping Φ is in general ill-conditioned or even singular. This entails that the inverse problem is in general ill-posed. In signal or image processing, one might for instance think of Φ as a convolution with the camera point-spread function, or a subsampling accounting for low-resolution or damaged sensors. In medical imaging, typical operators represent a (possibly subsampled) Radon transform (for computerized tomography), a partial Fourier transform (for magnetic resonance imaging), a propagation of the voltage/magnetic field from the dipoles to the sensors (for electro- or magneto-encephalography). In seismic imaging, the action of Φ amounts to a convolution with a wavelet-like impulse response that approximates the solution of a wave propagation equation in media with discontinuities. For regression problems in statistics and machine learning, Φ is the design matrix whose columns are P covariate vectors.

1.2 Variational Regularization

As argued above, solving an inverse problem from the observations (1) is in general ill-posed. In order to reach the land of well-posedness, it is necessary to restrict the inversion process to a well-chosen subset of $\mathbb{R}^N$ containing the plausible solutions including x_0 ; e.g. a linear space or a union of subspaces. A closely related procedure, that we describe next, amounts to adopting a variational framework where the sought-after solutions are those where a prior penalty/regularization function is the smallest. Though this approach may have a maximum a posteriori Bayesian interpretation, where a random prior is placed on x_0 , this is not the only interpretation. In fact, we put no randomness whatsoever on the class of signals we look for. We will not elaborate more on these differences in this chapter, but the reader may refer to [119] for an insightful discussion.

The foundations of regularization theory can be traced back to the pioneering work of the Russian school, and in particular of Tikhonov in 1943 when he proposed the notion of conditional well-posedness. In 1963, Tikhonov [217, 218] introduced what is now commonly referred to as Tikhonov (or also Tikhonov-Phillips) regularization, see also the book [219]. This corresponds, for $\lambda > 0$, to solving an optimization problem of the form

$$x^* \in \underset{x \in \mathbb{R}^N}{\operatorname{Argmin}} \frac{1}{2\lambda} \|\Phi x - y\|^2 + J(x). \quad (\mathcal{P}_{y,\lambda})$$

1.2.1 Data fidelity

In $(\mathcal{P}_{y,\lambda})$, $\|\Phi x - y\|^2$ stands for the data fidelity term. If the noise happens to be random, then using a likelihood argument, an appropriate fidelity term conforming to the noise distribution can be used instead of the quadratic data fidelity. Clearly, it is sufficient then to replace the latter by the negative log-likelihood of the distribution underlying the noise. Think for instance of the Csiszár's I-divergence for Poisson noise. We would also like to stress that many of the results provided in this chapter extend readily when the quadratic loss in the fidelity term, i.e. $\mu \mapsto \|y - \mu\|^2$, is replaced by any smooth and strongly convex function, see in particular Remark 13. To make our exposition concrete and digestible, we focus in the sequel on the quadratic loss.

1.2.2 Regularization

The function $J : \mathbb{R}^N \rightarrow \mathbb{R}$ is the regularization term which is intended to promote some prior on the vector to recover. We will consider throughout this chapter that J is a convex finite-valued function. Convexity plays an important role at many locations, both on the recovery guarantees and the algorithmic part. See for instance Section 6 which gives a brief overview of recent algorithms that are able to tackle

this class of convex optimization problems. It is however important to realize that non-convex regularizing penalties, as well as non-variational methods (e.g. greedy algorithms), are routinely used for many problems such as sparse or low-rank recovery. They may even outperform in practice their convex counterparts/relaxation. It is however beyond the scope of this chapter to describe these algorithms and the associated theoretical performance guarantees. We refer to Section 2.1 for a brief account on non-convex model selection approaches.

The scalar $\lambda > 0$ is the regularization parameter. It balances the trade-off between fidelity and regularization. Intuitively, and anticipating on our theoretical results hereafter, this parameter should be adapted to the noise level $\|w\|$ and the known properties of the vector x_0 to recover. Selecting optimally and automatically λ for a given problem is however difficult in general. This is at the heart of Section 5, where unbiased risk estimation strategies are shown to offer a versatile solution.

Note that since Φ is generally not injective and J is not coercive, the objective function of $(\mathcal{P}_{y,\lambda})$ is neither coercive nor strictly convex. In turn, there might be existence (of minimizers) issues, and even if minimizers exist, there are not unique in general.

Under mild assumptions, problem $(\mathcal{P}_{y,\lambda})$ is formally equivalent to the constrained formulations

$$\begin{aligned} \min \{J(x) ; \|y - \Phi x\| \leq \varepsilon\}, & \quad (\mathcal{P}_{y,\varepsilon}^1) \\ \min \{\|y - \Phi x\| ; J(x) \leq \gamma\}, & \quad (\mathcal{P}_{y,\gamma}^2) \end{aligned}$$

in the sense that there exists a bijection between each pair of parameters among $(\lambda, \varepsilon, \gamma)$ so that the corresponding problems share the same set of solutions. However, this bijection is not explicit and depends on y , so that both from an algorithmic point of view and a theoretical one, each problem may need to be addressed separately. See the recent paper [60] and references therein for a detailed discussion, and [154, Theorem 2.3] valid also in the non-convex case. We focus in this chapter on the penalized/Tikhonov formulation $(\mathcal{P}_{y,\lambda})$, though most of the results stated can be extended to deal with the constrained ones $(\mathcal{P}_{y,\varepsilon}^1)$ and $(\mathcal{P}_{y,\gamma}^2)$ (the former is known as the residual method or Morozov regularization and the latter as Ivanov regularization in the inverse problems literature).

The value of λ should typically be an increasing function of $\|w\|$. In the special case where there is no noise, i.e. $w = 0$, the fidelity to data should be perfect, which corresponds to considering the limit of $(\mathcal{P}_{y,\lambda})$ as $\lambda \rightarrow 0^+$. Thus, assuming that $y \in \text{Im}(\Phi)$, as is the case when $w = 0$, it can be proved that the solutions of $(\mathcal{P}_{y,\lambda})$ converge to the solutions of the following constrained problem [217, 198]

$$x^* \in \underset{x \in \mathbb{R}^N}{\text{Argmin}} J(x) \quad \text{subject to} \quad \Phi x = y. \quad (\mathcal{P}_{y,0})$$

1.3 Notations

For any subspace T of $\mathbb{R}^N$, we denote P_T the orthogonal projection onto T , $x_T = P_T(x)$ and $\Phi_T = \Phi P_T$. For a matrix A , we denote A^* its transpose, and A^+ its Moore-Penrose pseudo-inverse. For a convex set E , $\text{aff}(E)$ denotes its affine hull (i.e. the smallest affine space containing it), and $\text{lin}(E)$ its linear hull (i.e. the linear space parallel to $\text{aff}(E)$). Its relative interior $\text{ri}(E)$ is the interior for the topology of $\text{aff}(E)$ and $\text{rbd}(E)$ is its relative boundary. For a manifold $\mathcal{M}$, we denote $\mathcal{T}_{\mathcal{M}}(x)$ the tangent space of $\mathcal{M}$ at $x \in \mathcal{M}$. A good source on smooth manifold theory is [144].

A function $J : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is said to be proper if it is not identically $+\infty$. It is said to be finite-valued if $J(x) \in \mathbb{R}$ for all $x \in \mathbb{R}^N$. We denote $\text{dom}(J)$ the set of points x where $J(x) \in \mathbb{R}$ is finite. J is said to be closed if its epigraph $\{(x, y) ; J(x) \leq y\}$ is closed. For a set $C \subset \mathbb{R}^N$, the indicator function ι_C is defined as $\iota_C(x) = 0$ if $x \in C$ and $\iota_C(x) = +\infty$ otherwise.

We recall that the subdifferential at x of a proper and closed convex function $J : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is the set

$$\partial J(x) = \{\eta \in \mathbb{R}^N ; \forall \delta \in \mathbb{R}^N, J(x + \delta) \geq J(x) + \langle \eta, \delta \rangle\}.$$

Geometrically, when J is finite at x , $\partial J(x)$ is the set of normals to the hyper-planes supporting the graph of J and tangent to it at x . Thus, $\partial J(x)$ is a closed convex set. It is moreover bounded, hence compact, if and only if $x \in \text{int}(\text{dom}(J))$. The size of the subdifferential at $x \in \text{dom}(J)$ reflects in some sense the degree of non-smoothness of J at x . The larger the subdifferential at x , the larger the “kink” of the graph of J at x . In particular, if J is differentiable at x , then $\partial J(x)$ is a singleton and $\partial J(x) = \{\nabla J(x)\}$.

As an illustrative example, the subdifferential of the absolute value is

$$\forall x \in \mathbb{R}, \quad \partial |\cdot|(x) = \begin{cases} \text{sign}(x) & \text{if } x \neq 0, \\ [-1, 1] & \text{otherwise.} \end{cases} \quad (2)$$

The ℓ^1 norm

$$\forall x \in \mathbb{R}^N, \quad \|x\|_1 = \sum_{i=1}^N |x_i|$$

is a popular low-complexity prior (see Section 2.3.1 for more details). Formula (2) is extended by separability to obtain the subdifferential of the ℓ^1 norm

$$\partial \|\cdot\|_1(x) = \{\eta \in \mathbb{R}^N ; \|\eta\|_\infty \leq 1 \quad \text{and} \quad \forall i \in I, \text{sign}(\eta_i) = \text{sign}(x_i)\} \quad (3)$$

where $I = \text{supp}(x) = \{i ; x_i \neq 0\}$. Note that at a point $x \in \mathbb{R}^N$ such that $x_i \neq 0$ for all i , $\|\cdot\|_1$ is differentiable, and $\partial \|\cdot\|_1(x) = \{\text{sign}(x)\}$.

2 Low Complexity Priors

A recent trend in signal and image processing, statistics and machine learning is to make use of large collections of so-called “models” to account for the complicated structures of the data to handle. Generally speaking, these are manifolds $\mathcal{M}$ (most of the time linear subspaces), and hopefully of low complexity (to be detailed later), that capture the properties of the sought after signal, image or higher dimensional data. In order to tractably manipulate these collections, the key idea underlying this approach is to encode these manifolds in the non-smooth parts of the regularizer J . As we detail here, the theory of partial smoothness turns out to be natural to provide a mathematically grounded and unified description of these regularizing functions.

2.1 Model Selection

The general idea is thus to describe the data to recover using a large collection of models $\mathbb{M} = \{\mathcal{M}\}_{\mathcal{M} \in \mathbb{M}}$, which are manifolds. The “complexity” of elements in such a manifold $\mathcal{M}$ is measured through a penalty $\text{pen}(\mathcal{M})$. A typical example is simply the dimensionality of $\mathcal{M}$, and it should reflect the intuitive notion of the number of parameters underlying the description of the vector $x_0 \in \mathcal{M}$ that one aims at recovering from the noisy measurements of the form (1). As popular examples of such low complexity, one thinks of sparsity, piecewise regularity, or low rank. Penalizing in accordance to some notion of complexity is a key idea, whose roots can be traced back to the statistical and information theory literature, see for instance [162, 2].

Within this setting, the inverse problem associated to the measurements (1) is solved by restricting the inversion to an optimal manifold as selected by $\text{pen}(\mathcal{M})$. Formally, this would correspond to solving $(\mathcal{P}_{y,\lambda})$ with the combinatorial regularizer

$$J(x) = \inf \{ \text{pen}(\mathcal{M}) ; \mathcal{M} \in \mathbb{M} \text{ and } x \in \mathcal{M} \}. \quad (4)$$

A typical example of such a model selection framework is that with sparse signals, where the collection $\mathbb{M}$ corresponds to a union of subspaces, each of the form

$$\mathcal{M} = \{x \in \mathbb{R}^N ; \text{supp}(x) \subseteq I\}.$$

Here $I \subseteq \{1, \dots, N\}$ indexes the supports of signals in $\mathcal{M}$, and can be arbitrary. In this case, one uses $\text{pen}(\mathcal{M}) = \dim(\mathcal{M}) = |I|$, so that the associated combinatorial penalty is the so-called ℓ^0 pseudo-norm

$$J(x) = \|x\|_0 = |\text{supp}(x)| = |\{i \in \{1, \dots, N\} ; x_i \neq 0\}|. \quad (5)$$

Thus, solving $(\mathcal{P}_{y,\lambda})$ is intended to select a few active variables (corresponding to non-zero coefficients) in the recovered vector.

These sparse models can be extended in many ways. For instance, piecewise regular signals or images can be modeled using manifolds $\mathcal{M}$ that are parameterized by the locations of the singularities and some low-order polynomial between these singularities. The dimension of $\mathcal{M}$ thus grows with the number of singularities, hence the complexity of the model.

Literature review. The model selection literature [17, 11, 18] proposes many theoretical results to quantify the performance of these approaches. However, a major bottleneck of this class of methods is that the corresponding J function defined in (4) is non-convex, and even not necessarily closed, thus typically leading to highly intractable combinatorial optimization problems. For instance, in the case of the ℓ^0 penalty (5) and for an arbitrary operator Φ , $(\mathcal{P}_{y,\lambda})$ is known to be NP-hard, see e.g. [168].

It then appears crucial to propose alternative strategies which allow to deploy fast computational algorithms. A first line of work consists in finding stationary points of $(\mathcal{P}_{y,\lambda})$ using descent-like schemes. For instance, in the case of the ℓ^0 pseudo-norm, this can be achieved using iterative hard thresholding [20, 124], or iterative reweighting schemes which consist of solving a sequence of weighted ℓ^1 - or ℓ^2 -minimization problems where the weights used for the next iteration are computed from the values of the current solution, see for instance [189, 49, 72] and references therein. Another class of approaches is that of greedy algorithms. These are algorithms which explore the set of possible manifolds $\mathcal{M}$ by progressively, actually in a greedy fashion, increasing the value of $\text{pen}(\mathcal{M})$. The most popular schemes are matching pursuit [161] and its orthogonal variant [180, 73], see also the comprehensive review [169] and references therein. The last line of research, which is the backbone of this chapter, consists in considering convex regularizers which are built in such away that they promote the same set of low-complexity manifolds $\mathbb{M}$. In some cases, the convex regularizer proves to be the convex hull of the initial (restricted) non-convex combinatorial penalty (4). But these convex penalties can also be designed without being necessarily convexified surrogates of the original non-convex ones.

In the remainder of this section, we describe in detail a general framework that allows model selection through the general class of convex partly smooth functions.

2.2 Encoding Models into Partly Smooth Functions

Before giving the precise definition of our class of convex priors, we define formally the subspace T_x .

Definition 1 (Model tangent subspace). For any vector $x \in \mathbb{R}^N$, we define the *model tangent subspace* of x associated to J

$$T_x = \text{lin}(\partial J(x))^\perp,$$

In fact, the terminology “tangent” originates from the sharpness property of Definition 2(ii) below, when x belongs to the manifold $\mathcal{M}$.

When J is differentiable at x , i.e. $\partial J(x) = \{\nabla J(x)\}$, one has $T_x = \mathbb{R}^N$. On the contrary, when J is not smooth at x , the dimension of T_x is of a strictly smaller dimension, and J essentially promotes elements living on or close to the affine space $x + T_x$.

We can illustrate this using the ℓ^1 norm $J = \|\cdot\|_1$ defined in (2). Using formula (3) for the subdifferential, one obtains that

$$T_x = \{u \in \mathbb{R}^N ; \text{supp}(u) \subseteq \text{supp}(x)\},$$

which is the set of vector having the same sparsity pattern as x .

Toward the goal of studying the recovery guarantees of problem (4), our central assumption is that J is a partly smooth function relative to some manifold $\mathcal{M}$. Partial smoothness of functions was originally defined [146]. Loosely speaking, a partly smooth function behaves smoothly as we move on the manifold $\mathcal{M}$, and sharply if we move normal to it. Our definition hereafter specializes that of [146] to the case of finite-valued convex functions.

Definition 2. Let J be a finite-valued convex function. J is *partly smooth at x relative to a set $\mathcal{M}$* containing x if

- (i) (Smoothness) $\mathcal{M}$ is a C^2 -manifold around x and J restricted to $\mathcal{M}$ is C^2 around x .
- (ii) (Sharpness) The tangent space $\mathcal{T}_{\mathcal{M}}(x)$ is T_x .
- (iii) (Continuity) The set-valued mapping ∂J is continuous at x relative to $\mathcal{M}$.

J is said to be *partly smooth relative to a set $\mathcal{M}$* if $\mathcal{M}$ is a manifold and J is partly smooth at each point $x \in \mathcal{M}$ relative to $\mathcal{M}$. J is said to be *locally partly smooth at x relative to a set $\mathcal{M}$* if $\mathcal{M}$ is a manifold and there exists a neighbourhood U of x such that J is partly smooth at each point of $\mathcal{M} \cap U$ relative to $\mathcal{M}$.

Remark 1 (Uniqueness of $\mathcal{M}$). In the previous definition, $\mathcal{M}$ needs only to be defined locally around x , and it can be shown to be locally unique, see [130, Corollary 4.2]. In the following we will thus often denote $\mathcal{M}_x$ any such a manifold for which J is partly smooth at x .

Taking once again the example of $J = \|\cdot\|_1$, one sees that in this case, $\mathcal{M}_x = T_x$ because this function is polyhedral. Section 2.3.6 below defines functions J for which $\mathcal{M}_x$ differs in general from T_x .

2.3 Examples of Partly Smooth Regularizers

We describe below some popular examples of partly smooth regularizers that are widely used in signal and image processing, statistics and machine learning. We

first expose basic building blocks (sparsity, group sparsity, anti-sparsity) and then show how the machinery of partial smoothness enables a powerful calculus to create new priors (using pre- and post-composition, spectral lifting, and positive linear combinations).

2.3.1 ℓ^1 Sparsity

One of the most popular non-quadratic convex regularization is the ℓ^1 norm

$$J(x) = \|x\|_1 = \sum_{i=1}^N |x_i|,$$

which promotes sparsity. Indeed, it is easy to check that J is partly smooth at x relative to the subspace

$$\mathcal{M}_x = T_x = \{u \in \mathbb{R}^N ; \text{supp}(u) \subseteq \text{supp}(x)\}.$$

Another equivalent way to interpret this ℓ^1 prior is that it is the convex envelope (restricted to the ℓ^2 -ball) of the ℓ^0 pseudo-norm (5), in the sense that the ℓ^1 -unit ball is the convex hull of the restriction of the unit ball of the ℓ^0 -pseudo norm to the ℓ^2 -unit ball.

Literature review. The use of the ℓ^1 norm as a sparsity-promoting regularizer traces back several decades. An early application was deconvolution in seismology [61, 197, 212]. Rigorous recovery results began to appear in the late 1980's [81, 80]. In the mid-1990's, ℓ^1 regularization of least-square problems has been popularized in the signal processing literature under the name basis pursuit [58] and in the statistics literature under the name Lasso [213]. Since then, the applications and understanding of ℓ^1 minimization have continued to increase dramatically.

2.3.2 $\ell^1 - \ell^2$ Group Sparsity

To better capture the sparsity pattern of natural signals and images, it is useful to structure the sparsity into non-overlapping groups $\mathcal{B}$ such that $\bigcup_{b \in \mathcal{B}} b = \{1, \dots, N\}$. This group structure is enforced by using typically the mixed $\ell^1 - \ell^2$ norm

$$J(x) = \|x\|_{1, \mathcal{B}} = \sum_{b \in \mathcal{B}} \|x_b\|, \quad (6)$$

where $x_b = (x_i)_{i \in b} \in \mathbb{R}^{|b|}$. Unlike the ℓ^1 norm, and except the case $|b| = 1$ for all $b \in \mathcal{B}$, the $\ell^1 - \ell^2$ norm is not polyhedral, but is still partly smooth at x relative to the linear manifold

$$\mathcal{M}_x = T_x = \{u ; \text{supp}_{\mathcal{B}}(u) \subseteq \text{supp}_{\mathcal{B}}(x)\} \quad \text{where} \quad \text{supp}_{\mathcal{B}}(x) = \bigcup \{b ; x_b \neq 0\}.$$

Literature review. The idea of group/block sparsity has been first proposed by [127, 126, 31] for wavelet block shrinkage, i.e. when $\Phi = \text{Id}$. For over-determined regression problems of the form (1), it has been introduced by [9, 244]. Group sparsity has also been extensively used in machine learning in e.g. [7] (regression and multiple kernel learning) and [175] (for multi-task learning). The wavelet coefficients of a natural image typically exhibit some group structure, see [160] and references therein on natural image modeling. Indeed, edges and textures induce strong dependencies between coefficients. In audio processing, it has proved useful to structure sparsity in multi-channel data [121]. Group sparsity is also at the heart of the so-called multiple measurements vector (MMV) model, see for instance [69, 57]. It is possible to replace the ℓ^2 norm with more general functionals, such as ℓ^p norms for $p > 1$, see for instance [225, 171, 237].

2.3.3 ℓ^∞ Anti-sparsity

In some cases, the vector to be reconstructed is expected to be flat. Such a prior can be captured using the ℓ^∞ norm

$$J(x) = \|x\|_\infty = \max_{i \in \{1, \dots, n\}} |x_i|.$$

It can be readily checked that this regularizer is partly smooth (in fact polyhedral) relative to the subspace

$$\mathcal{M}_x = T_x = \{u; u_I = \rho x_I \text{ for some } \rho \in \mathbb{R}\}, \quad \text{where } I = \{i; x_i = \|x\|_\infty\}.$$

Literature review.

The ℓ^∞ regularization has found applications in computer vision, such as for database image retrieval [137]. For this application, it is indeed useful to have a compact signature of a signal x , ideally with only two values $\pm \|x\|_\infty$ (thus achieving optimal anti-sparsity since $\dim(T_x) = 1$ in such a case). An approach proposed in [138] for realizing this binary quantification is to compute these vectors as solutions of $(\mathcal{P}_{y,\lambda})$ for $J = \|\cdot\|_\infty$ and a random Φ . A study of this regularization is done in [108], where an homotopy-like algorithm is provided. The use of this ℓ^∞ regularization is also connected to Kashin's representation [157], which is known to be useful in stabilizing the quantization error for instance. Other applications such as wireless network optimization [211] also rely on the ℓ^∞ prior.

2.3.4 Synthesis Regularizers

Sparsity or more general low-complexity regularizations are often used to model coefficients $\alpha \in \mathbb{R}^Q$ describing the data $x = D\alpha$ in a dictionary $D \in \mathbb{R}^{N \times Q}$ of Q atoms in $\mathbb{R}^N$. Given a partly smooth function $J_0 : \mathbb{R}^Q \rightarrow \mathbb{R}$, we define the following synthesis-type prior $J : \mathbb{R}^N \rightarrow \mathbb{R}$ as the pre-image of J_0 under the linear mapping D

$$J(x) = \min_{\alpha \in \mathbb{R}^Q} J_0(\alpha) \quad \text{s.t.} \quad D\alpha = x$$

Since J_0 is bounded below and convex, J is convex. If D is surjective (as in most cases with redundant dictionaries), then J is also finite-valued. The initial optimization $(\mathcal{P}_{y,\lambda})$ can equivalently be solved directly over the coefficients domain to obtain $x^* = D\alpha^*$ where

$$\alpha^* \in \operatorname{Argmin}_{\alpha \in \mathbb{R}^Q} \frac{1}{2\lambda} \|y - \Phi D\alpha\|^2 + J_0(\alpha) \quad (7)$$

which can be interpreted as a regularized inversion of the operator ΦD using the prior J_0 .

It is possible to study directly the properties of the solutions α^* to (7), which involves directly partial-smoothness of J_0 . A slightly different question is to understand the behavior of the solutions $x^* = D\alpha^*$ of $(\mathcal{P}_{y,\lambda})$, which requires to study partial smoothness of J itself. In the case where D is invertible, both problems are completely equivalent.

Literature review. Sparse synthesis regularization using $J_0 = \|\cdot\|_1$ is popular in signal and image processing to model natural signals and images, see for instance [160, 207] for a comprehensive account. The key problem to achieve good performance in these applications is to design a dictionary to capture sparse representations of the data to process. Multiscale dictionaries built from wavelet pyramids are popular to sparsely represent transient signals with isolated singularities and natural images [159]. The curvelet transform is known to provide non-adaptive near-optimal sparse representation of piecewise smooth images away from smooth edges (so-called cartoon images) [34]. Gabor dictionaries (made of localized and translated Fourier atoms) are popular to capture locally stationary oscillating signals for audio processing [3]. To cope with richer and diverse contents, researchers have advocated to concatenate several dictionaries to solve difficult problems in signal and image processing, such as component separation or inpainting, see for instance [99]. A line of current active research is to learn and optimize the dictionary from exemplars or even from the available data themselves. We refer to [97, Chapter 12] for a recent overview of the relevant literature.

2.3.5 Analysis Regularizers

Analysis-type regularizers (following the terminology introduced in [98]) are of the form

$$J(x) = J_0(D^*x),$$

where $D \in \mathbb{R}^{N \times Q}$ is a linear operator. Such a prior controls the low complexity (as measured by J_0) of the correlations between the columns of D and the signal x . If J_0 is partly smooth at $z = D^*x$ for the manifold $\mathcal{M}_z^0$, then it is shown in [146, Theorem 4.2] that J is partly smooth at x relative to the manifold

$$\mathcal{M}_x = \{u \in \mathbb{R}^N ; D^*u \in \mathcal{M}_z^0\}$$

provided that the following transversality condition holds [144, Theorem 6.30(a)]

$$\text{Ker}(D) \cap \mathcal{T}_{\mathcal{M}_z^0}(z)^\perp = \{0\} \iff \text{Im}(D^*) + \mathcal{T}_{\mathcal{M}_z^0}(z) = \mathbb{R}^N.$$

Literature review. A popular example is when $J_0 = \|\cdot\|_1$ and D^* a finite-difference discretization of the derivative of a 1-D signal or a 2-D image. This defines the anisotropic total variation semi-norm, which promotes piecewise constant signals or images [196]. The 2-D isotropic total variation semi-norm can be interpreted as taking $J_0 = \|\cdot\|_{1,2}$ with blocks of size two. A comprehensive review of total variation regularization can be found in [52]. TV regularization has been extended in several ways to model piecewise polynomial functions, see in particular the Total Generalized Variation prior [26].

One can also use a wavelet dictionary D which is shift-invariant, such that the corresponding regularization J can be seen as a kind of multi-scale total variation. This is typically the case of the Haar wavelet dictionary [208]. When using higher order wavelets, the corresponding priors favors models $\mathcal{M}$ composed of discrete piecewise polynomials.

The Fused Lasso [214] corresponds to J_0 being the ℓ^1 -norm and D is the concatenation of the identity and the adjoint of a finite-difference operator. The corresponding models $\mathcal{M}$ are composed of disjoint blocks over which the signals are constant.

Defining a block extracting operator $D^*x = (x_b)_{b \in \mathcal{B}}$ allows to re-write the group $\ell^1 - \ell^2$ norm (6), even with overlapping blocks (i.e. $\exists(b, b') \in \mathcal{B}^2$ with $b \cap b' \neq \emptyset$), as $J = J_0 \circ D^*$ where $J_0 = \|\cdot\|_{1,2}$ without overlap, see [139, 182, 245, 32]. To cope with correlated covariates in linear regression, analysis-type sparsity-enforcing prior were proposed in [118, 193] using $J_0 = \|\cdot\|_*$ the nuclear norm (as defined in Section 2.3.6).

For unitary D , the solutions of $(\mathcal{P}_{y,\lambda})$ with synthesis and analysis regularizations are obviously the same. In the general case (e.g. D overcomplete), however, these two regularizations are different. Some authors have reported results comparing these two priors for the case where J_0 is the ℓ^1 norm [98, 199]. A first discussion on the relation and distinction between analysis and synthesis ℓ^1 -sparse regularizations can be found in [98]. But only very recently, some theoretical recovery results and algorithmic developments on ℓ^1 -analysis regularization (so-called cosparsity model) have began to be developed, see e.g. [167, 230].

2.3.6 Spectral Functions

The natural extension of low-complexity priors to matrix-valued data $x \in \mathbb{R}^{N_0 \times N_0}$ (where $N = N_0^2$) is to impose the low-complexity on the singular values of the matrix. We denote $x = U_x \text{diag}(\Lambda_x) V_x^*$ an SVD decomposition of x , where $\Lambda_x \in \mathbb{R}_+^{N_0}$. If $j : \mathbb{R}^{N_0} \rightarrow \mathbb{R}$ is a permutation-invariant closed convex function, then one can con-

sider the function

$$J(x) = j(\Lambda_x)$$

which can be shown to be a convex function as well [147]. When restricted to the linear space of symmetric matrices, j is partly smooth at Λ_x for a manifold m_{Λ_x} , if and only if J is partly smooth at x relative to the manifold

$$\mathcal{M}_x = \{U \operatorname{diag}(\Lambda) U^* ; \Lambda \in m_{\Lambda_x}, U \in \mathcal{O}_{N_0}\},$$

where $\mathcal{O}_{N_0} \subset \mathbb{R}^{N_0 \times N_0}$ is the orthogonal group. The proof of this assertion can be found in [70, Theorem 3.19], which builds upon the work of [71] on manifold smoothness transfer under spectral lifting. This result can be extended to non-symmetric matrices by requiring that j is an absolutely permutation-invariant closed convex function, see [70, Theorem 5.3].

Literature review. The most popular spectral prior is obtained for $j = \|\cdot\|_1$. This defines the nuclear norm, or 1-Schatten norm, as

$$J(x) = \|x\|_* = \|\Lambda_x\|_1. \quad (8)$$

It can be shown that the nuclear norm is the convex hull of the rank function with respect to the spectral norm ball, see [102, 133]. It then corresponds to promoting a low-rank prior. Moreover, the nuclear norm can be shown to be partly smooth at x relative to the set [148, Example 2]

$$\mathcal{M}_x = \{u ; \operatorname{rank}(u) = \operatorname{rank}(x)\}$$

which is a manifold around x .

The nuclear norm has been used in signal and image processing, statistics and machine learning for various applications, including low rank matrix completion [205, 190, 40], principal component pursuit [36], model reduction [103], and phase retrieval [45]. It is also used for some imaging applications, see for instance [152].

2.3.7 Mixed Regularizations

Starting from a collection of convex functions $\{J_\ell\}_{\ell \in \mathcal{L}}$, $\mathcal{L} = \{1, \dots, L\}$, it is possible to design a convex function as

$$J_\ell(x) = \sum_{\ell \in \mathcal{L}} \rho_\ell J_\ell(x),$$

where $\rho_\ell > 0$ are weights. If each J_ℓ is partly smooth at x relative to a manifold $\mathcal{M}_x^\ell$, then it is shown in [146, Corollary 4.8] that J is also partly smooth at x for

$$\mathcal{M}_x = \bigcap_{\ell \in \mathcal{L}} \mathcal{M}_x^\ell,$$

with the proviso that the manifolds $\mathcal{M}_x^\ell$ intersect transversally [144, Theorem 6.30(b)], i.e. the sum of their respective tangent spaces $\mathcal{T}_{\mathcal{M}_x^\ell}(x)$ spans the whole ambient space $\mathbb{R}^N$.

Literature review. A popular example is to impose both sparsity and low rank of a matrix, when using $J_1 = \|\cdot\|_1$ and $J_2 = \|\cdot\|_*$, see for instance [114, 177].

2.3.8 Separable Regularization

Let $\{J_\ell\}_{\ell \in \mathcal{L}}$, $\mathcal{L} = \{1, \dots, L\}$, be a family of convex functions. If J_ℓ is partly smooth at x_ℓ relative to a manifold $\mathcal{M}_{x_\ell}^\ell$, then the separable function

$$J(\{x_\ell\}_{\ell \in \mathcal{L}}) = \sum_{\ell \in \mathcal{L}} J_\ell(x_\ell)$$

is partly smooth at $(x_1, \dots, x_L)$ relative to $\mathcal{M}_{x_1}^1 \times \dots \times \mathcal{M}_{x_L}^L$ [146, Proposition 4.5].

Literature review. One fundamental problem that has attracted a lot of interest in the recent years in data processing involves decomposing an observed object into a linear combination of components/constituents x_ℓ , $\ell \in \mathcal{L} = \{1, \dots, L\}$. One instance of such a problem is image decomposition into texture and piece-wise-smooth (cartoon) parts. The corresponding forward model can be cast in the form (1), where $x_0 = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$, x_1 and x_2 are the texture and cartoon components, and $\Phi = [\text{Id} \quad \text{Id}]$. The decomposition is then achieved by solving the variational problem $(\mathcal{P}_{y,\lambda})$, where J_1 is designed to promote the discontinuities in the image, and J_2 to favor textures; see [206, 6, 184] and references therein. Another example of decomposition is principal component pursuit, proposed in [36], to decompose a matrix which is the superposition of a low-rank component and a sparse component. In this case $J_1 = \|\cdot\|_1$ and $J_2 = \|\cdot\|_*$.

3 ℓ^2 Stability

In this section, we assume that J is a finite-valued convex function, but it is not assumed to be partly smooth.

The observations y are in general contaminated by noise, as described by the forward model (1). It is thus important to study the ability of $(\mathcal{P}_{y,\lambda})$ to recover x_0 to a good approximation in presence of such a noise w , and to assess how the reconstruction error decays as a function of the noise level. In this section, we present a generic result ensuring a so-called “linear convergence rate” in terms of ℓ^2 -error between a recovered vector and x_0 (see Theorem 1), which encompasses a large body of literature from the inverse problems community.

3.1 Dual Certificates

It is intuitively expected that if $(\mathcal{P}_{y,\lambda})$ is good at recovering an approximation of x_0 in presence of noise, then $(\mathcal{P}_{y,0})$ should be able to identify x_0 uniquely when the noise vanishes, i.e. $y = \Phi x_0$. For this to happen, the solution to $(\mathcal{P}_{y,0})$ has to satisfy some non-degeneracy condition. To formalize this, we first introduce the notion of dual certificate.

Definition 3 (Dual certificates). For any vector $x \in \mathbb{R}^N$, the set of *dual certificates* at x is defined as

$$\mathcal{D}(x) = \text{Im}(\Phi^*) \cap \partial J(x) .$$

The terminology “dual certificate” was introduced in [40]. One can show that the image by Φ^* of the set of solutions of the Fenchel-Rockafellar dual to $(\mathcal{P}_{y,0})$ is precisely $\mathcal{D}(x)$.

It is also worth noting that x_0 being a solution of $(\mathcal{P}_{y,0})$ for $y = \Phi x_0$ is equivalent to $\mathcal{D}(x_0) \neq \emptyset$. Indeed, this is simply a convenient re-writing of the first order optimality condition for $(\mathcal{P}_{y,0})$.

To ensure stability of the set of minimizers $(\mathcal{P}_{y,\lambda})$ to noise perturbing the observations Φx_0 , one needs to introduce the additional requirement that the dual certificates should be strictly inside the subdifferential of J at x_0 . This is precisely the non-degeneracy condition mentioned previously.

Definition 4 (Non-degenerate dual certificates). For any vector $x \in \mathbb{R}^N$, we define the set of *non-degenerate dual certificates* of x

$$\tilde{\mathcal{D}}(x) = \text{Im}(\Phi^*) \cap \text{ri}(\partial J(x)) .$$

3.2 Stability in ℓ^2 Norm

The following theorem, proved in [101], establishes a linear convergence rate valid for any regularizer J , without any particular assumption beside being a proper closed convex function. In particular, it does not assume partial smoothness of J . This generic result encompasses many previous works, as discussed in Section 3.3.

Theorem 1. Assume that

$$\text{Ker}(\Phi) \cap T_{x_0} \cap = \{0\} \quad \text{and} \quad \tilde{\mathcal{D}}(x_0) \neq \emptyset \quad (9)$$

and consider the choice $\lambda = c \|w\|$, for some $c > 0$. Then we have for all minimizers x^* of $(\mathcal{P}_{y,\lambda})$

$$\|x^* - x_0\|_2 \leq C \|w\| , \quad (10)$$

where $C > 0$ is a constant (see Remark 4 for details).

In plain words, this bound tells us that the distance of x_0 to the set of minimizers of $(\mathcal{P}_{y,\lambda})$ is within a factor of the noise level, which justifies the terminology “linear convergence rate”.

Remark 2 (The role of non-smoothness). The injectivity of Φ when restricted to T_{x_0} is intimately related to the fact that J is non-smooth at x_0 . The higher the degree of non-smoothness, the lower the dimension of the subspace T_{x_0} , and hence the more likely the restricted injectivity. If J is smooth around x_0 (e.g. quadratic regularizers), however, the restricted injectivity condition cannot be fulfilled, unless Φ is itself injective. The reason is that T_{x_0} is the whole $\mathbb{R}^N$ at the smoothness points. For smooth regularizations, it can be shown that the convergence rate is slower than linear, we refer to [198] for more details.

Remark 3 (Uniqueness). One can show that condition (9) implies that x_0 is the unique solution of $(\mathcal{P}_{y,0})$ for $y = \Phi x_0$. This condition however does not imply in general that $(\mathcal{P}_{y,\lambda})$ has a unique minimizer for $\lambda > 0$.

Remark 4 (Stability constant). The result (10) ensures that the mapping $y \mapsto x^*$ (that might be set-valued) is C -Lipschitz-continuous at $y = \Phi x_0$. Condition $\tilde{\mathcal{D}}(x_0) \neq \emptyset$ is equivalent to the existence of some $\eta \in \tilde{\mathcal{D}}(x_0)$. The value of C (in fact an upper-bound) can be found in [101]. It depends on Φ , T_{x_0} , c and the chosen non-degenerate dual certificate η . In particular, the constant degrades critically as η gets closer to the relative boundary of $\tilde{\mathcal{D}}(x_0)$, which reflects the intuition of how far is η from being a non-degenerate certificate.

Remark 5 (Source condition). The condition $\mathcal{D}(x_0) \neq \emptyset$ is often called “source condition” or “range condition” in the literature of inverse problems. We refer to the monograph [198] for a general overview of this condition and its implications. It is an abstract condition, which is not easy to check in practice, since exhibiting a valid non-degenerate certificate is not trivial. We give in Section 4.1 further insights about this in the context of compressed sensing. Section 4.1 describes a particular construction of a good candidate (the so-called linearized pre-certificate) for being such an $\eta \in \tilde{\mathcal{D}}(x_0)$, and it is shown to govern stability of the manifold $\mathcal{M}_{x_0}$ for partly smooth regularizers.

Remark 6 (Infinite dimension). It is important to remind that, in its full general form, Theorem 1 only holds in finite dimension. The constant C indeed may depend on the ambient dimension N , in which case the constant can blow-up as the discretization grid of the underlying continuous problem is made finer (i.e. as N grows). We detail below some relevant literature where similar results are shown in infinite dimension.

3.3 Related Works

3.3.1 Convergence Rates

For quadratic regularizations of the form $J = \|D^* \cdot\|^2$ for some linear operator D^* , the ℓ^2 -error decay can be proved to be $O(\sqrt{\|w\|})$, which is not linear, see [198, Chapter 3] for more details and extensions to infinite dimensional Hilbert spaces. For non-smooth priors, in [30], the authors show the Bregman distance between x^* and x_0 exhibits a linear convergence rate for both the Lagrangian $(\mathcal{P}_{y,\lambda})$ and the constrained $(\mathcal{P}_{y,\varepsilon}^1)$ problems under the source condition $\mathcal{D}(x_0) \neq 0$. These results hold more generally over infinite dimensional Banach spaces. They have been subsequently generalized to ill-posed non-linear inverse problems by [192] and [134]. It is important to observe that in order to prove convergence rates in terms of ℓ^2 -error, as done in (10), it is necessary to strengthen the source condition to its non-degenerate version, i.e. $\widehat{\mathcal{D}}(x_0) \neq 0$.

In [155], the authors consider the case where J is a ℓ^p norm with $1 \leq p \leq 2$ and establish convergence rates of $\|\Phi x_0 - \Phi x^*\|$ in $O(\|w\|)$ and of $\|x^* - x_0\|$ in $O(\sqrt{\|w\|})$. [117] prove Theorem 1 for $J = \|\cdot\|_1$. They show that the non-degeneracy condition is also necessary for linear convergence, and draw some connections with the restricted isometry property (RIP), see below. Under a condition that bears similarities with (9), linear convergence with respect to J , i.e. $J(x^* - x_0) = O(\|w\|)$, is proved in [116] for positively homogeneous regularizers. This result is equivalent to Theorem 1 but only when J is coercive, which precludes many important regularizers, such as for instance analysis-type regularizers including total variation.

3.3.2 RIP-based Compressed Sensing

The recovery performance of compressed sensing (i.e. when Φ is drawn from suitable random ensembles) for $J = \|\cdot\|_1$ has been widely analyzed under the so-called restricted isometry property (RIP) introduced in [42, 43, 47]. For any integer $k \geq 0$, the k -th order restricted isometry constant of a matrix Φ is defined as the smallest $\delta_k \geq 0$ such that

$$(1 - \delta_k) \|x\|^2 \leq \|\Phi x\|^2 \leq (1 + \delta_k) \|x\|^2,$$

for all vectors x such that $\|x\|_0 \leq k$. It is shown [42] that if $\delta_{2k} + \delta_{3k} < 1$, then for every vector x_0 with $\|x_0\|_0 \leq k$, there exists a non-degenerate certificate [46, Lemma 2.2], see also the discussion in [117]. In turn, this implies linear convergence rate, and is applied in [43] to show ℓ^2 -stability to noise of compressed sensing. This was generalized in [35] to analysis sparsity $J = \|D^* \cdot\|_1$, where D is assumed to be a tight frame, structured sparsity in [35] and matrix completion in [190, 39] using $J = \|\cdot\|_*$. The goal is then to design RIP matrices Φ with constants such that $\delta_{2k} + \delta_{3k}$ (or a related quantity) is small enough. This is possible if Φ is drawn from an appropriate random ensemble for some (hopefully optimal) scaling of (N, P, k) . For instance, if Φ is drawn from the standard Gaussian ensemble (i.e. with i.i.d. zero-mean standard

Gaussian entries), there exists a constant C such that the RIP constants of $\Phi/\sqrt{P}$ obey $\delta_{2k} + \delta_{3k} < 1$ with overwhelming probability provided that

$$P \geq Ck \log(N/k), \quad (11)$$

see for instance [47]. This result remains true when the entries of Φ are drawn independently from a subgaussian distribution. When Φ is a structured random matrix, e.g. random partial Fourier matrix, the RIP constants of $\Phi/\sqrt{P}$ can also satisfy the desired bound, but at the expense of polylog terms in the scaling (11), see [105] for a comprehensive treatment. Note that in general, computing the RIP constants for a given matrix is an NP-hard problem [10, 220].

3.3.3 RIP-less Compressed Sensing

RIP-based guarantees are uniform, in the sense that the recovery holds with high probability for *all* sparse signals. There is a recent wave of work in RIP-less analysis of the recovery guarantees for compressed sensing. The claims are non-uniform, meaning that they hold for a fixed signal with high probability on the random matrix Φ . This line of approaches improves on RIP-based bounds providing typically sharper constants. When Φ is drawn from the Gaussian ensemble, it is proved in [195] for $J = \|\cdot\|_1$ that if the number of measurements P obeys $P \geq Ck \log(N/k)$ for some constant $C > 0$, where $k = \|x_0\|_0$, then condition (9) holds with high probability on Φ . This result is based on Gordon's comparison principle for Gaussian processes and depends on a summary parameter for convex cones called the Gaussian width. Equivalent lower bounds on the number of measurements for matrix completion from random measurements by minimizing the nuclear norm were provided in [48] to ensure that (9) holds with high probability. This was used to prove ℓ^2 -stable matrix completion in [37].

The authors in [54] have recently showed that the Gaussian-width based approach leads to sharp lower bounds on P required to solve regularized inverse problems from Gaussian random measurements. For instance, they showed for $J = \|\cdot\|_1$ that

$$P > 2k \log(N/k) \quad (12)$$

guarantees exact recovery from noiseless measurements by solving $(\mathcal{P}_{y,0})$. An overhead in the number of measurements is necessary to get linear convergence of the ℓ^2 -error in presence of noise by solving $(\mathcal{P}_{y,\varepsilon}^1)$ with $\varepsilon = \|w\|$, i.e. x_0 is feasible. Their results handle for instance the case of group sparsity (6) and the nuclear norm (8). In the polyhedral case, it can be shown that (12) implies the existence of a non-degenerate dual certificate, i.e. (9), with overwhelming probability. The Gaussian width is closely related to another geometric quantity called the statistical dimension in conic integral geometry. The statistical dimension canonically extends the linear dimension to convex cones, and has been proposed in [4] to deliver reliable predictions about the quantitative aspects of the phase transition for exact noiseless recovery from Gaussian measurements.

To deal with non-Gaussian matrix measurements (such as for instance partial Fourier matrices), [123] introduced the “golfing scheme” for noiseless low-rank matrix recovery guarantees using $J = \|\cdot\|_*$. The golfing scheme is an iterative procedure to construct an (approximate) non-degenerate certificate. This construction is also studied in [38] for noiseless and noisy sparse recovery with $J = \|\cdot\|_1$. In another chapter of this volume [221], the author develops a technique, called the “bowling scheme”, which is able to deliver bounds on the number of measurements that are similar with the Gaussian width-based bounds for standard Gaussian measurements, but the argument applies to a much wider class of measurement ensembles.

4 Model Stability

In the remainder of this chapter, we assume that J is finite-valued convex and locally partly smooth around x_0 , as defined in Section 2.2. This means in particular that the prior J promotes locally solution which belong to the manifold $\mathcal{M} = \mathcal{M}_{x_0}$. In the previous section, we were only concerned with ℓ^2 stability guarantees and partial smoothness was not necessary then. Owing to the additional structure conveyed by partial smoothness, we will be able to provide guarantees on the identification of the correct $\mathcal{M} = \mathcal{M}_{x_0}$ by solving $(\mathcal{P}_{y,\lambda})$, i.e. whether the (unique) solution x^* of $(\mathcal{P}_{y,\lambda})$ satisfies $x^* \in \mathcal{M}$. Such guarantees are of paramount importance for many applications. For instance, consider the case where ℓ^1 regularization is used to localize some (sparse) sources. Then $x^* \in \mathcal{M}$ means that one perfectly identifies the correct source locations. Another example is that of the nuclear norm for low-rank matrix recovery. The correct model identification implies that x^* has the correct rank, and consequently that the eigenspaces of x^* have the correct dimensions and are close to those of x_0 .

4.1 Linearized Pre-certificate

We saw in Section 3.2 that ℓ^2 -stability of the solutions to $(\mathcal{P}_{y,\lambda})$ is governed by the existence of a non-degenerate dual certificate $p \in \tilde{\mathcal{D}}(x_0)$. It turns out that not all dual certificates are equally good for stable model identification, and toward the latter, one actually needs to focus on a particular dual certificate, that we call “minimal norm” certificate.

Definition 5 (Minimal norm certificate). Assume that x_0 is a solution of $(\mathcal{P}_{y,0})$. We define the “minimal-norm certificate” as

$$\eta_0 = \Phi^* \underset{p \in \partial J(x_0)}{\operatorname{argmin}} \|p\| . \quad (13)$$

A remarkable property, stated in Proposition 1 below, is that, as long as one is concerned with checking whether η_0 is non-degenerate, i.e. $\eta_0 \in \text{ri}(\partial J(x_0))$, one can instead use the vector η_F defined below, which can be computed in closed form.

Definition 6 (Linearized pre-certificate). Assume that

$$\text{Ker}(\Phi) \cap T_{x_0} = \{0\}. \quad (14)$$

We define the “linearized pre-certificate” as

$$\eta_F = \Phi^* \underset{\Phi^* p \in \text{aff}(\partial J(x_0))}{\text{argmin}} \|p\|. \quad (15)$$

Remark 7 (Well-posedness of the definitions). Note that the hypothesis that x_0 is a solution of $(\mathcal{P}_{y,0})$ is equivalent to saying that $\mathcal{D}(x_0)$ is a non-empty convex compact set. Hence in (13), the optimal p is the orthogonal projection of 0 on a non-empty closed convex set, and thus η_0 is uniquely defined. Similarly, the hypothesis (14) implies that the constraint set involved in (15) is a non-empty affine space, and thus η_F is also uniquely defined.

Remark 8 (Certificate vs. pre-certificate). Note that the only difference between (13) and (15) is that the convex constraint set $\partial J(x_0)$ is replaced by a simpler affine constraint. This means that η_F does not always qualify as a valid certificate, i.e. $\eta_F \in \partial J(x_0)$, hence the terminology “pre-certificate” used. This condition is actually at the heart of the model identification result exposed in Theorem 2.

For now on, let us remark that η_F is actually simple to compute, since it amounts to solving a linear system in the least-squares sense.

Proposition 1. Under condition (14), one has

$$\eta_F = \Phi^* \Phi_{T_{x_0}}^{+,*} e_{x_0} \quad \text{where} \quad e_{x_0} = P_{T_{x_0}}(\partial J(x_0)) \in \mathbb{R}^N. \quad (16)$$

Remark 9 (Computating e_x). The vector e_x appearing in (16) can be computed in closed form for most of the regularizers discussed in Section 2.2. For instance, for $J = \|\cdot\|_1$, $e_x = \text{sign}(x)$. For $J = \|\cdot\|_{1,\mathcal{B}}$, it reads $e_x = (e_b)_{b \in \mathcal{B}}$, where $e_b = x_b / \|x_b\|$ if $x_b \neq 0$, and $e_b = 0$ otherwise. For $J = \|\cdot\|_*$ and a SVD decomposition $x = U_x \text{diag}(\Lambda_x) V_x^*$, one has $e_x = U_x V_x^*$.

The following proposition, whose proof can be found in [231], exhibits a precise relationship between η_0 and η_F . In particular, it implies that η_F can be used in place of η_0 to check whether η_0 is non-degenerate, i.e. $\eta_0 \in \text{ri}(\partial J(x_0))$.

Proposition 2. Under condition (14), one has

$$\begin{aligned} \eta_F \in \text{ri}(\partial J(x_0)) &\implies \eta_F = \eta_0, \\ \eta_0 \in \text{ri}(\partial J(x_0)) &\implies \eta_F = \eta_0. \end{aligned}$$

4.2 Model Identification

The following theorem provides a sharp sufficient condition to establish model selection. It is proved in [231]. It encompasses as special cases many previous works in the signal processing, statistics and machine learning literatures, as we discuss in Section 4.5.1.

Theorem 2. *Let J be locally partly smooth at x_0 relative to $\mathcal{M} = \mathcal{M}_{x_0}$. Assume that*

$$\text{Ker}(\Phi) \cap T_{x_0} = \{0\} \quad \text{and} \quad \eta_F \in \text{ri}(\partial J(x_0)). \quad (17)$$

Then there exists C such that if

$$\max(\lambda, \|w\|/\lambda) \leq C, \quad (18)$$

the solution x^ of $(\mathcal{P}_{y,\lambda})$ from the measurements (1) is unique and satisfies*

$$x^* \in \mathcal{M} \quad \text{and} \quad \|x_0 - x^*\| = O(\max(\lambda, \|w\|)). \quad (19)$$

Remark 10 (Linear convergence rate vs. model identification). Obviously, the assumptions (17) of Theorem 2 imply those of Theorem 1. They are of course stronger, but imply a stronger result, since uniqueness of x^* and model identification (i.e. $x^* \in \mathcal{M}$) are not guaranteed by Theorem 1 (which does not even need J to be partly smooth). A chief advantage of Theorem 2 is that its hypotheses can be easily checked and analyzed for a particular operator Φ . Indeed, computing η_F only requires solving a linear system, as clearly seen from formula (16).

Remark 11 (Minimal signal-to-noise ratio). Another important distinction between Theorems 1 and 2 is the second assumption (18). In plain words, it requires that the noise level is small enough and that the regularization parameter is wisely chosen. Such an assumption is not needed in Theorem 2 to ensure linear convergence of the ℓ^2 -error. In fact, this condition is quite natural. To see this, consider for instance the case of sparse recovery where $J = \|\cdot\|_1$. If the minimal signal-to-noise ratio is low, the noise will clearly dominate the amplitude of the smallest entries, so that one cannot hope to recover the exact support, but it is still possible to achieve a low ℓ^2 -error by forcing those small entries to zero.

Remark 12 (Identification of the manifold). For all the regularizations considered in Section 2.3, the conclusion of Theorem 2 is even stronger as it guarantees that $\mathcal{M}_{x^*} = \mathcal{M}$. The reason is that for any x and nearby points x' with $x' \in \mathcal{M}_x$, one has $\mathcal{M}_{x'} = \mathcal{M}_x$.

Remark 13 (General loss/data fidelity). It is possible to extend Theorem 2 to account for general loss/data fidelity terms beyond the quadratic one, i.e. $\frac{1}{2}\|y - \Phi x\|^2$. More precisely, this result holds true for loss functions of the form $F(\Phi x, y)$, where $F : \mathbb{R}^P \times \mathbb{R}^P \rightarrow \mathbb{R}$ is a C^2 strictly convex function in its first argument, ∇F is C^1 in the second argument, with $\nabla F(y, y) = 0$, where ∇F is the gradient with respect to the first variable. In this case, the expression (16) of η_F becomes simply

$$\eta_F = \Gamma(\mathbf{P}_T \Gamma \mathbf{P}_T)^+ e_{x_0} \quad \text{where} \quad \begin{cases} T = T_{x_0} \\ \Gamma = \Phi^* \partial^2 F(\Phi x_0, \Phi x_0) \Phi \end{cases},$$

and where $\partial^2 F$ is the Hessian with respect to the first variable (which is a positive definite operator). We refer to [231] for more details.

4.3 Sharpness of the Model Identification Criterion

The following proposition, proved in [231], shows that Theorem 2 is in some sense sharp, since the hypothesis $\eta_F \in \text{ri}(\partial J(x_0))$ (almost) characterizes the stability of $\mathcal{M}$.

Proposition 3. *We suppose that x_0 is the unique solution of $(\mathcal{P}_{y,0})$ for $y = \Phi x_0$ and that*

$$\text{Ker}(\Phi) \cap T_{x_0} = \{0\}, \quad \text{and} \quad \eta_F \notin \partial J(x_0). \quad (20)$$

Then there exists $C > 0$ such that if (18) holds, then any solution x^ of $(\mathcal{P}_{y,\lambda})$ for $\lambda > 0$ obeys $x^* \notin \mathcal{M}$.*

In the particular case where $w = 0$ (no noise), this result shows that the manifold $\mathcal{M}$ is not correctly identified when solving $(\mathcal{P}_{y,\lambda})$ for $y = \Phi x_0$ and for any $\lambda > 0$ small enough.

Remark 14 (Critical case). The only case not covered by neither Theorem 2 nor Proposition 3 is when $\eta_F \in \text{rbd}(\partial J(x_0))$, where rbd stands for the boundary relative to the affine hull. In this case, one cannot conclude, since depending on the noise w , one can have either stability or non-stability of $\mathcal{M}$. We refer to [230] where an example illustrates this situation for the 1-D total variation $J = \|D_{\text{DIF}}^* \cdot\|_1$, where D_{DIF}^* is a finite-difference discretization of the 1-D derivative operator.

4.4 Probabilistic Model Consistency

Theorem 2 assumes a deterministic noise w , and the operator Φ is fixed. For applications in statistics and machine learning, it makes sense to rather assume a random model for both Φ and w . The natural question is then to assert that the estimator defined by solving $(\mathcal{P}_{y,\lambda})$ is consistent in the sense that it correctly estimates x_0 and possibly the model $\mathcal{M}_{x_0}$ as the number of observations $P \rightarrow +\infty$. This requires to handle operators Φ with an increasing number of rows, and thus to also assess sensitivity of the optimization problem $(\mathcal{P}_{y,\lambda})$ to perturbations of Φ (and not only to (w, λ) as done previously).

To be more concrete, in this section, we work under the classical setting where N and x_0 are fixed as the number of observations $P \rightarrow +\infty$. The data (φ_i, w_i) are assumed to be random vectors in $\mathbb{R}^N \times \mathbb{R}$, where φ_i is the i -th row of Φ for $i =$

$1, \dots, P$. These vectors are supposed independent and identically distributed (i.i.d.) samples from a joint probability distribution such that $\mathbb{E}(w_i|\varphi_i) = 0$, finite fourth-order moments, i.e. $\mathbb{E}(w_i^4) < +\infty$ and $\mathbb{E}(\|\varphi_i\|^4) < +\infty$. Note that in general, w_i and φ_i are not necessarily independent. It is possible to consider other distribution models by weakening some of the assumptions and strengthening others, see e.g. [143, 246, 7]. Let us denote $\Gamma = \mathbb{E}(\varphi_i^* \varphi_i) \in \mathbb{R}^{N \times N}$, where φ_i is any row of Φ . We do not make any assumption on the invertibility of Γ .

In this setting, a natural extension of η_F defined by (16) in the deterministic case is

$$\tilde{\eta}_F = \Gamma \Gamma_{T_{x_0}}^+ e_{x_0}$$

where $\Gamma_{T_{x_0}} = P_{T_{x_0}} \Gamma P_{T_{x_0}}$, and we use the fact that $\Gamma_{T_{x_0}}$ is symmetric and $\text{Im}(\Gamma_{T_{x_0}}^+) \subset T_{x_0}$. It is also implicitly assumed that $\text{Ker}(\Gamma) \cap T_{x_0} = \{0\}$ which is the equivalent adaptation of the restricted injectivity condition in (17) to this setting.

To make the discussion clearer, the parameters $(\lambda = \lambda_P, \Phi = \Phi_P, w = w_P)$ are now indexed by P . The estimator x_P^* obtained by solving $(\mathcal{P}_{\lambda_P, y_P})$ for $y_P = \Phi_P x_0 + w_P$ is said to be consistent for x_0 if,

$$\lim_{P \rightarrow +\infty} \Pr(x_P^* \text{ is unique}) = 1$$

and $x_P^* \rightarrow x_0$ in probability. The estimator is said to be model consistent if

$$\lim_{P \rightarrow +\infty} \Pr(x_P^* \in \mathcal{M}) = 1,$$

where $\mathcal{M} = \mathcal{M}_{x_0}$ is the manifold associated to x_0 .

The following result, whose proof can be found in [231], guarantees model consistency for an appropriate scaling of μ_P . It generalizes several previous works in the statistical and machine learning literature as we review in Section 4.5.1.

Theorem 3. *If*

$$\text{Ker}(\Gamma) \cap T_{x_0} = \{0\} \quad \text{and} \quad \tilde{\eta}_F \in \text{ri}(\partial J(x_0)), \quad (21)$$

and

$$\lambda_P = o(P) \quad \text{and} \quad \lambda_P^{-1} = o(P^{-1/2}). \quad (22)$$

Then the estimator x_P^ of x_0 is model consistent.*

4.5 Related Works

4.5.1 Model Consistency

Theorem 2 is a generalization of a large body of results in the literature. For the Lasso, i.e. $J = \|\cdot\|_1$, to the best of our knowledge, this result was initially stated

in [107]. In this setting, the result (19) corresponds to the correct identification of the support, i.e. $\text{supp}(x^*) = \text{supp}(x_0)$. Condition (21) for $J = \|\cdot\|_1$ is known in the statistics literature under the name “irrepresentable condition” (generally stated in a non-geometrical form), see e.g. [246]. [143] have shown estimation consistency for Lasso for fixed N and x_0 and asymptotic normality of the estimates. The authors in [246] prove Theorem 3 for $J = \|\cdot\|_1$, though under slightly different assumptions on the covariance and noise distribution. A similar result is established in [141] for the elastic net, i.e. $J = \|\cdot\|_1 + \rho \|\cdot\|_2^2$ for $\rho > 0$. In [7] and [8], the author proves Theorem 3 for two special cases, namely the group Lasso and nuclear norm minimization. Note that these previous works assume that the asymptotic covariance Γ is invertible. We do not impose such an assumption, and only require the weaker restricted injectivity condition $\text{Ker}(\Gamma) \cap T = \{0\}$. In a previous work [230], we have proved an instance of Theorem 2 when $J(x) = \|D^*x\|_1$, where $D \in \mathbb{R}^{N \times Q}$ is an arbitrary linear operator. This covers as special cases the discrete anisotropic total variation or the fused Lasso. This result was further generalized in [229] when J belongs to the class of partly smooth functions relative to linear manifolds $\mathcal{M}$, i.e. $\mathcal{M} = T_x$. Typical instances encompassed in this class are the $\ell^1 - \ell^2$ norm, or its analysis version, as well as polyhedral gauges including the ℓ^∞ norm. Note that the nuclear norm (and composition of it with linear operators as proposed for instance in [118, 193]), whose manifold is not linear, does not fit into the framework of [229], while it is covered by Theorem 2. Lastly, a similar result is proved in [90] for a continuous (infinite dimensional) sparse recovery problem over the space of Radon measures normed by J the total variation of a measure (not to be confused with the total variation of functions). In this continuous setting, an interesting finding is that, when $\eta_0 \in \text{ri}(\partial J(x_0))$, η_0 is not equal to η_F but to a different certificate (called “vanishing derivative” certificate in [90]) that can also be computed by solving a linear system.

4.5.2 Stronger Criteria for ℓ^1

Many sufficient conditions have been proposed in the literature to ensure that η_F is a non-degenerate certificate, and hence to guarantee stable identification of the support (i.e. model). We illustrate this here for $J = \|\cdot\|_1$, but similar reasoning can be carried out for $\|\cdot\|_{1,\mathcal{B}}$ or $\|\cdot\|_*$.

The strongest criterion makes use of mutual coherence, first considered in [78]

$$\mu(\Phi) = \max_{i \neq j} |\langle \varphi_i, \varphi_j \rangle|$$

where each column φ_i of Φ are assumed normalized to a unit ℓ^2 norm. Mutual coherence measures the degree of ill-conditioning of Φ through the correlation of its columns $(\varphi_i)_{1 \leq i \leq N}$. Mutual coherence is always lower-bounded by $\sqrt{\frac{N-P}{P(N-1)}}$, and equality holds if and only if $(\varphi_i)_{1 \leq i \leq N}$ is an equiangular tight frame, see [210]. Finer variants based on cumulative coherences have been proposed in [120, 24]. To take into account the influence of the support $I = \text{supp}(x_0)$ of the vector x_0 to

recover, Tropp introduced in [222] the Exact Recovery Condition (ERC), defined as

$$\text{ERC}(I) = \left\| \Phi_{I^c}^* \Phi_I^{+,*} \right\|_{\infty, \infty} = \max_{j \notin I} \left\| \Phi_I^+ \varphi_j \right\|_1$$

where $\|\cdot\|_{\infty, \infty}$ is the matrix operator norm induced by the ℓ^∞ vector norm, $\Phi_I = (\varphi_i)_{i \in I}$ and I^c is the complement of the set I . Φ_I is assumed injective which, in view of Section 2.3.1, is nothing but a specialization to ℓ^1 of the restricted injectivity condition in (17). A weak ERC criterion, which does not involve matrix inversion, is derived in [85]

$$\text{wERC}(I) = \frac{\max_{j \in I^c} \sum_{i \in I} |\langle \varphi_i, \varphi_j \rangle|}{1 - \max_{j \in I} \sum_{i \neq j \in I} |\langle \varphi_i, \varphi_j \rangle|}.$$

Given the structure of the subdifferential of the ℓ^1 norm, it is easy to check that

$$\eta_F \in \text{ri}(\partial J(x_0)) \iff \text{IC}(x_0) = \left\| \Phi_{I^c}^* \Phi_I^{+,*} \text{sign}(x_{0,I}) \right\|_\infty < 1.$$

The right hand side in the equivalence is precisely what is called the irrepresentable condition in statistics and machine learning. Clearly, $\text{IC}(x_0)$ involves both the sign vector and the support of x_0 . The following proposition gives ordered upper bounds of $\text{IC}(x_0)$ in terms of the cruder criteria ERC, wERC and mutual coherence. A more elaborate discussion of them can be found in [160].

Proposition 4. *Assume that Φ_I is injective and denote $k = |I| = \|x_0\|_0$. Then,*

$$\text{IC}(x_0) \leq \text{ERC}(I) \leq \text{wERC}(I) \leq \frac{k\mu(\Phi)}{1 - (k-1)\mu(\Phi)}.$$

4.5.3 Linearized Pre-certificate for Compressed Sensing Recovery

Stable support identification has been established in [240, 83] for the Lasso problem when Φ is drawn from the Gaussian ensemble. These works show that for $k = \|x_0\|_0$, if

$$P > 2k \log(N)$$

then indeed $\eta_F \in \text{ri}(\partial J(x_0))$, and this scaling can be shown to be sharp. This scaling should be compared with (12) ensuring that there exists a non-degenerate certificate. The gap in the log term indicates that there exists vectors that can be stably recovered by ℓ^1 minimization in ℓ^2 -error sense, but whose support cannot be stably identified. Equivalently, for these vectors, there exists a non-degenerate certificate but it is not η_F .

The pre-certificate η_F is also used to ensure exact recovery of a low-rank matrix from incomplete noiseless measurements by minimizing the nuclear norm [40, 48]. This idea is further generalized by [41] for a family of decomposable norms (including in particular ℓ^1 - ℓ^2 norm and the nuclear norm), which turns to be a subset of

partly smooth regularizers. In these works, lower bounds on the number of random measurements needed for η_F to be a non-degenerate certificate are developed. In fact, these measurement lower bounds combined with Theorem 2 allow to conclude that matrix completion by solving $(\mathcal{P}_{y,\lambda})$ with $J = \|\cdot\|_*$ identifies the correct rank at high signal-to-noise levels.

4.5.4 Sensitivity Analysis

Sensitivity analysis is a central theme in variational analysis. Comprehensive monographs on the subject are [23, 166]. The function to be analyzed underlying problems $(\mathcal{P}_{y,\lambda})$ and $(\mathcal{P}_{y,0})$ is

$$f(x, \theta) = \begin{cases} \frac{1}{2\lambda} \|y - \Phi x\|^2 + J(x) & \text{if } \lambda > 0, \\ \iota_{\mathcal{H}_y}(x) + J(x) & \text{if } \lambda = 0, \end{cases}, \quad (23)$$

where $\mathcal{H}_y = \{y; \Phi x = y\}$ and where the parameters are $\theta = (\lambda, y, \Phi)$ for $\lambda \geq 0$. Theorems 2 and 3 can be understood as a sensitivity analysis of the minimizers of f at a point $(x = x_0, \theta = \theta_0 = (0, \Phi x_0, \Phi))$.

Classical sensitivity analysis of non-smooth optimization problems seeks conditions to ensure smoothness of the mapping $\theta \mapsto x_\theta$ where x_θ is a minimizer of $f(\cdot, \theta)$, see for instance [194, 23]. This is usually guaranteed by the non-degenerate source condition and restricted injectivity condition (9), which, as already exposed in Section 3.2, ensure linear convergence rate, and hence Lipschitz behavior of this mapping. The analysis proposed by Theorem 2 goes one step further, by assessing that $\mathcal{M}_{x_0}$ is a stable manifold (in the sense of [241]), since the minimizer x_θ is unique and remains in $\mathcal{M}_{x_0}$ for θ close to θ_0 . Our starting point for establishing Theorem 2 is the inspiring work of Lewis [146] who first introduced the notion of partial smoothness and showed that this broad class of functions enjoys a powerful calculus and sensitivity theory. For convex functions (which is the setting considered in our work), partial smoothness is closely related to $\mathcal{U} - \mathcal{V}$ -decompositions developed in [145]. In fact, the behavior of a partly smooth function and of its minimizers (or critical points) depend essentially on its restriction to the manifold, hence offering a powerful framework for sensitivity analysis theory. In particular, critical points of partly smooth functions move stably on the manifold as the function undergoes small perturbations [149]. A important and distinctive feature of Theorem 2 is that, partial smoothness of J at x_0 relative to $\mathcal{M}$ transfers to $f(\cdot, \theta)$ for $\lambda > 0$, but not when $\lambda = 0$ in general. In particular, [146, Theorem 5.7] does not apply to prove our claim.

5 Sensitivity Analysis and Parameter Selection

In this section, we study local variations of the solutions of $(\mathcal{P}_{y,\lambda})$ considered as functions of the observations y . In a variational-analytic language, this corresponds to analyzing the sensitivity of the optimal values of $(\mathcal{P}_{y,\lambda})$ to small perturbations of y seen as a parameter. This analysis will have important implications, and we exemplify one of them by constructing unbiased estimators of the quadratic risk, which in turn will allow to have an objectively-guided way to select the optimal value of the regularization parameter λ .

As argued in Section 4.5.4, assessing the recovery performance by solving $(\mathcal{P}_{y,\lambda})$ for w and λ small amounts to a sensitivity analysis of the minimizers of f in (23) at $(x = x_0, \theta = \theta_0 = (0, \Phi x_0, \Phi))$. This section involves again sensitivity analysis of (23) to perturbations of y but for $\lambda > 0$. Though we focus our attention on sensitivity to y , our arguments extend to any parameters, for instance λ or Φ .

Similarly to the previous section, we suppose here that J is a finite-valued convex and partly smooth function. For technical reasons, we furthermore assume that the partial smoothness manifold is linear, i.e. $\mathcal{M}_x = T_x$. We additionally suppose that the set of all possible models $\mathcal{T} = \{T_x\}_{x \in \mathbb{R}^N}$ is finite. All these assumptions hold true for the regularizers considered in Section 2.3, with the notable exception of the nuclear norm, whose manifolds of partial smoothness are non-linear.

5.1 Differentiability of Minimizers

Let us denote $x^*(y)$ a minimizer of $(\mathcal{P}_{y,\lambda})$ for a fixed value of $\lambda > 0$. Our main goal is to study differentiability of $x^*(y)$ and find a closed-form formula of the derivative of $x^*(y)$ with respect to the observations y . Since $x^*(y)$ is not necessarily a unique minimizer, such a result means actually that we have to single out one solution $x^*(y)$, which hopefully should be a locally smooth function of y . However, as J is non-smooth, one cannot hope for such a result to hold for any observation $y \in \mathbb{R}^P$. For applications to risk estimation (see Section 5.3), it is important to characterize precisely the smallest set $\mathcal{H}$ outside of which $x^*(y)$ is indeed locally smooth. It turns out that one can actually write down an analytical expression of such a set $\mathcal{H}$, containing points where one cannot find locally a smooth parameterization of the minimizers. This motivates our definition of what we coin a “transition space”.

Definition 7 (Transition space). We define the *transition space* $\mathcal{H}$ as

$$\mathcal{H} = \bigcup_{T \in \mathcal{T}} \text{bd}(\mathcal{H}_T),$$

where $\text{bd}(C)$ is the boundary of a set C , and

$$\mathcal{H}_T = \left\{ y \in \mathbb{R}^P ; \exists x \in \tilde{T}, \lambda^{-1} \Phi_T^* (\Phi x - y) \in \text{rbd}(\partial J(x)) \right\}.$$

where $\tilde{T} = \{x \in \mathbb{R}^N ; T_x = T\}$.

The set $\mathcal{H}$ contains the observations $y \in \mathbb{R}^P$ such that the model subspace $T_{\tilde{x}(y)}$ associated to a well chosen solution $\tilde{x}(y)$ of $(\mathcal{P}_{y,\lambda})$ is not stable with respect to small perturbations of y . In particular, when $J = \|\cdot\|_1$, it can be checked that $\mathcal{H}$ is a finite union of hyperplanes and when $J = \|\cdot\|_{1,2}$ it is a semi-algebraic set (see Definition 8). This stability is not only crucial to prove smoothness of $\tilde{x}(y)$, it is also important to be able to write down an explicit formula for the derivative, as detailed in the following theorem whose proof is given in [227].

Theorem 4. *Let $y \notin \mathcal{H}$ and x^* a solution of $(\mathcal{P}_{y,\lambda})$ such that*

$$\text{Ker } \Phi_T \cap \text{Ker } D^2 J_T(x^*) = \{0\} \quad (\mathcal{I}_{x^*})$$

where $T = T_{x^}$. Then, there exists an open neighborhood $\mathcal{V} \subset \mathbb{R}^N$ of y , and a mapping $\tilde{x} : \mathcal{V} \rightarrow T$ such that*

1. *for every $\bar{y} \in \mathcal{V}$, $\tilde{x}(\bar{y})$ is a solution of $(\mathcal{P}_{\lambda,\bar{y}})$, and $\tilde{x}(y) = x^*$;*
2. *the mapping $\tilde{x}$ is $C^1(\mathcal{V})$ and*

$$\forall \bar{y} \in \mathcal{V}, \quad D\tilde{x}(\bar{y}) = (\Phi_T^* \Phi_T + \lambda D^2 J_T(x^*))^{-1} \Phi_T.$$

Here $D^2 J_T$ is the Hessian (second order derivative) of J restricted to T . This Hessian is surely well-defined owing to partial smoothness, see Definition 2(i).

5.2 Semi-algebraic Geometry

Our goal now is to show that the set $\mathcal{H}$ is in some sense “small” (in particular to show that it has zero Lebesgue measure), which will entail differentiability of $y \mapsto x^*$ Lebesgue almost everywhere. For this, additional geometrical structure on J is needed. Such a rich class of functions is provided by the notion of a semi-algebraic subset of $\mathbb{R}^N$ to be defined shortly. Semi-algebraic sets and functions have been broadly applied to various areas of optimization. The wide applicability of semi-algebraic functions follows largely from their stability under many mathematical operations. In particular, the celebrated Tarski-Seidenberg theorem states, loosely, that the projection of a semi-algebraic set is semi-algebraic. These stability properties are crucial to obtain the following result, proved in [227].

Definition 8 (Semi-algebraic set and function). A set E is semi-algebraic if it is a finite union of sets defined by polynomial equations and (possibly strict) inequalities. A function $f : E \rightarrow F$ is semi-algebraic if E and its graph $\{(u, f(u)) ; u \in E\}$ are semi-algebraic sets.

Remark 15 (From semi-algebraic to o-minimal geometry). The class of semi-algebraic functions is large, and subsumes, for instance, all the regularizers J described

in Section 2.3. The qualitative properties of semi-algebraic functions are shared by a much bigger class called functions definable in an o-minimal structure over $\mathbb{R}$, or simply definable functions. O-minimal structures over $\mathbb{R}$ correspond in some sense to an axiomatization of some of the prominent geometrical properties of semi-algebraic geometry [68] and particularly of the stability under projection. For example, the function $J(x) = \sum_i |x_i|^s$, for an arbitrary $s \geq 0$, is semi-algebraic only for rational $s \in \mathbb{Q}$, while it is always definable in an o-minimal structure [236]. Due to the variety of regularizations J that can be formulated within the framework of o-minimal structures, all our results stated in this section apply to definable functions, see [227] for a detailed treatment.

Semi-algebraic functions are stable for instance under (sub)differentiation and projection. These stability properties are crucial to obtain the following result, proved in [227].

Proposition 5. *If J is semi-algebraic, the transition space $\mathcal{H}$ is semi-algebraic and has zero Lebesgue measure.*

5.3 Unbiased Risk Estimation

A problem of fundamental practical importance is to automatically adjust the parameter λ to reach the best recovery performance when solving $(\mathcal{P}_{y,\lambda})$. Parameter selection is a central theme in statistics, and is intimately related to the question of model selection, as introduced in Section 2.1.

We then adopt a statistical framework in which the observation model (1) becomes

$$Y = \Phi x_0 + W \quad (24)$$

where W is random noise having an everywhere strictly positive probability density function, assumed to be known. Though the forthcoming results can be stated for a large family of distributions, for the sake of concreteness, we only consider the white Gaussian model where $W \sim \mathcal{N}(0, \sigma^2 \text{Id}_{P \times P})$, with known variance σ^2 .

Under the observation model (24), the ideal choice of λ should be the one which minimizes the quadratic estimation risk $\mathbb{E}_W(\|x^*(Y) - x_0\|^2)$. This is obviously not realistic as x_0 is not available, and in practice, only one realization of Y is observed. To overcome these obstacles, the traditional approach is to replace the quadratic risk with an some estimator that solely depends on Y . The risk estimator is also expected to enjoy nice statistical properties among which unbiasedness is highly desirable.

However, it can be shown, see e.g. [100, Section IV], that the quadratic risk $\mathbb{E}_W(\|x^*(Y) - x_0\|^2)$ cannot be reliably estimated on $\text{Ker}(\Phi)$. Nonetheless, we may still obtain a reliable assessment of the part that lies in $\text{Im}(\Phi^*) = \text{Ker}(\Phi)^\perp$ or any linear image of it. For instance, the most straightforward surrogate of the above risk is the so-called prediction risk $\mathbb{E}_W(\|\mu(Y) - \mu_0\|^2)$, where

$$\mu_0 = \Phi x_0 \quad \text{and} \quad \mu(y) = \Phi x^*(y),$$

where $x^*(y)$ is any solution of $(\mathcal{P}_{y,\lambda})$. One can easily show that $\mu(y) \in \mathbb{R}^P$ is well-defined as a single-valued mapping and thus does not depend on the particular choice of $x^*(y)$, see [227]. Consequently, Theorem 4 shows that $y \mapsto \mu(y)$ is a C^1 mapping on $\mathbb{R}^P \setminus \mathcal{H}$.

5.4 Degrees of Freedom

The degrees of freedom (DOF) quantifies the model “complexity” of a statistical modeling procedure [95]. It is at the heart of several risk estimation procedures. Therefore, in order to design estimators of the prediction risk, an important step is to get an estimator of the corresponding DOF.

Definition 9 (Empirical DOF). Suppose that $y \mapsto \mu(y)$ is differentiable Lebesgue almost everywhere, as is the case when it is Lipschitz-continuous (Rademacher’s theorem). The empirical number of degrees of freedom is defined as

$$\text{df}(y) = \text{div}(\mu)(y) = \text{tr}(\text{D}\mu(y)),$$

where the derivative is to be understood in the weak sense, i.e. to hold Lebesgue almost everywhere (a.e.).

An instructive example to get the gist of this formula is the case where μ is the orthogonal projection onto some linear subspace V . We then get easily that $\text{df}(y) = \dim(V)$, which is in agreement with the intuitive notion of the number of DOF.

The following result delivers the closed-form expression of $\text{df}(y)$, valid on a full Lebesgue measure set, for $\mu(y) = \Phi x^*(y)$ and $x^*(y)$ an appropriate solution of $(\mathcal{P}_{y,\lambda})$. At this stage, it is important to realize that the main difficulty does not lie in showing almost everywhere differentiability of $\mu(y)$; this mapping is in fact Lipschitz-continuous by classical arguments of sensitivity analysis applied to $(\mathcal{P}_{y,\lambda})$. Rather, it is the existence of such a formula and its validity Lebesgue a. e. that requires more subtle arguments obtained owing to partial smoothness of J . For this, we need also to rule out the points y where $(\mathcal{I}_{x^*})$ does not hold. This is the rationale behind the following set.

Definition 10 (Non-injectivity set). We define the *Non-injectivity set* $\mathcal{G}$ as

$$\mathcal{G} = \{y \notin \mathcal{H} ; (\mathcal{I}_{x^*}) \text{ does not hold for any minimizer } x^* \text{ of } (\mathcal{P}_{y,\lambda})\}.$$

Theorem 5. For every $y \notin \mathcal{H} \cup \mathcal{G}$, there is x^* such that $(\mathcal{P}_{y,\lambda})$ holds and

$$\text{df}(y) = \text{tr}(\Delta_{x^*}(y)) \quad \text{where} \quad \Delta_{x^*}(y) = \Phi_T \circ (\Phi_T^* \Phi_T + \lambda \text{D}^2 J_T(x^*))^{-1} \circ \Phi_T^*, \quad (25)$$

where $T = T_{x^*}$.

Remark 16 (Non-injectivity set). It turns out that $\mathcal{G}$ is in fact empty for many regularizers. This is typically the case for $J = \|\cdot\|_1$ [84], $J = \|D^* \cdot\|_1$ [228], and the

underlying reasoning can be more generally extended to polyhedral regularizers. The same result was also shown for $J = \|\cdot\|_{1,2}$ in [232]. More precisely, in all these works, it was shown that for each $y \notin \mathcal{H}$, there exists a solution x^* of $(\mathcal{P}_{y,\lambda})$ that fulfills $(\mathcal{I}_{x^*})$. The proof is moreover constructive allowing to build such a solution starting from any other one.

5.5 Stein Unbiased Risk Estimator (SURE)

We now have all necessary ingredients at hand to design an estimator of the prediction risk.

Definition 11. Suppose that $y \mapsto \mu(y)$ is differentiable Lebesgue almost everywhere, as is the case when it is Lipschitz-continuous. The SURE is defined as

$$\text{SURE}(y) = \|y - \mu(y)\|^2 + 2\sigma^2 \text{df}(y) - P\sigma^2. \quad (26)$$

In this definition, we have anticipated on unbiasedness of this estimator. In fact, this turns out to be a fundamental property owing to the celebrated lemma of Stein [209], which indeed asserts that the SURE (26) is an unbiased estimator of the prediction risk. Therefore, putting together Theorem 5, Proposition 5 and Stein's lemma, we get the following.

Theorem 6. Suppose that J is semi-algebraic and $\mathcal{G}$ is of zero Lebesgue measure. Then,

$$\mathbb{E}_W(\text{SURE}(Y)) = \mathbb{E}_W(\|\mu(Y) - \mu_0\|^2)$$

where (25) is plugged into (26), and $\mu(Y) = \Phi x^*(Y)$.

Remark 17 (Parameter selection). A practical usefulness of the SURE is its ability to provide an objectively guided way to select a good λ from a single observation y by minimizing $\text{SURE}(y)$. While unbiasedness of the SURE is guaranteed, it is hard to control its variance and hence its consistency. This is an open problem in general, and thus little can be said about the actual theoretical efficiency of such an empirical parameter selection method. It works however remarkably well in practice, see the discussion in Section 5.6.5 and references therein.

Remark 18 (Projection risk). The SURE can be extended to unbiasedly estimate other risks than the prediction one. For instance, as argued in Section 5.3, one can estimate the so-called projection risk defined as $\mathbb{E}_W(\|P_{\text{Ker}(\Phi)^\perp}(x^*(Y) - x_0)\|^2)$. This is obviously better than the prediction risk as a surrogate for the estimation risk.

5.6 Related Works

5.6.1 Sensitivity Analysis

In Section 4.5.4, we reviewed the relevant literature pertaining to sensitivity analysis for partly smooth functions, which is obviously very connected to Theorem 4. See also [21] for the case of linear optimization over a convex semi-algebraic partly smooth feasible set, where the authors prove a sensitivity result with a zero-measure transition space. A distinctive feature of our analysis toward proving unbiasedness of the SURE is the need to ensure that sensitivity analysis can be carried out on a full Lebesgue measure set. In particular, it necessitates local stability of the manifold $\mathcal{M}_{x^*}$ associated to an appropriate solution x^* , and this has to hold Lebesgue almost everywhere. Thus the combination of partial smoothness and semi-algebraicity is the key.

5.6.2 Risk Estimators

In this section, we put emphasis on the SURE as an unbiased estimator of the prediction risk. There are other alternatives in the literature which similarly rely on estimator of the DOF. One can think for instance of the generalized cross-validation (GCV) [115]. Thus our results apply equally well to such risk estimators. Extensions of the SURE to independent variables from a continuous exponential family are considered in [135]. [100] generalizes the SURE principle to continuous multivariate exponential families, see also [181, 228] for the multivariate Gaussian case. The results described here can be extended to these setting as well, see [227].

5.6.3 Applications of SURE in Statistics and Imaging

Applications of SURE emerged for choosing the parameters of linear estimators such ridge regression or smoothing splines [150]. After its introduction in the wavelet community through the SURE-Shrink estimator [79], it has been extensively used for various image restoration problems, e.g. with sparse regularization [19, 238, 186, 55, 181, 33, 156, 187, 188] or with non-local means [234, 89, 76, 235].

5.6.4 Closed-form Expressions for SURE

For the Lasso problem, i.e. $J = \|\cdot\|_1$, the divergence formula (25) reads

$$\text{df}(y) = |\text{supp}(x^*)|,$$

where x^* is a solution of $(\mathcal{P}_{y,\lambda})$ such that $(\mathcal{I}_{x^*})$ holds, i.e. $\Phi_{\text{supp}(x^*)}$ has full rank. This result is proved in [247] for injective Φ and in [84] for arbitrary Φ . This result is extended to analysis ℓ^1 -sparsity, i.e. $J = \|D^* \cdot\|_1$, in [216, 228]. A formula for the DOF in the case where $x^*(y)$ is the orthogonal projection onto a partly smooth convex set C is proved in [142]. This work extends that of [164] which treats the case where C is a convex polyhedral cone. These two works allow one to compute the degrees of freedom of estimators defined by solving $(\mathcal{P}_{y,\gamma}^2)$ in the case where Φ is injective. [128] studied the DOF of the metric projection onto a closed set (non-necessarily convex), and gave a precise representation of the bias when the projection is not sufficiently differentiable.

A formula of an estimate of the DOF for the group Lasso, i.e. $J = \|\cdot\|_{1,2}$ when Φ is orthogonal within each group was conjectured in [244]. An estimate is also given by [202] using heuristic derivations that are valid only when Φ is injective, though its unbiasedness is not proved. [226] derived an estimator of the DOF of the group Lasso and proved its unbiasedness when Φ is injective. Closed-form expression of the DOF estimate for denoising with the nuclear norm, i.e. $\Phi = \text{Id}$ and $J = \|\cdot\|_*$, were concurrently provided in [77, 44].

5.6.5 Numerical Methods for SURE

Deriving the closed-form expression of the DOF is in general challenging and has to be addressed on a case by case basis. The implementation of the divergence formula such as (25) can be computationally expensive in high dimension. But since only the trace of the Jacobian is needed, it is possible to speed up these computations through Monte-Carlo sampling, but at the price of mild approximations. If the Jacobian is not known in closed-form or prohibitive to compute, one may appeal to finite-difference approximations along Monte Carlo sampled directions [243, 201], see [111, 186] for applications to imaging problems.

In practice, the analytical formula (25) might be subject to serious numerical instabilities, and thus cannot always be applied safely when the solution x^* is only known approximately. Think for instance of the case where x^* is approximated by an iterate computed after finitely-many iterations of an algorithm as detailed in Section 6. A better practice is then to directly compute the DOF, hence the SURE, recursively from the iterates themselves, as proposed by [238, 112, 75].

6 Proximal Splitting for Structured Optimization

Though problems $(\mathcal{P}_{y,\lambda})$, $(\mathcal{P}_{y,0})$, $(\mathcal{P}_{y,\epsilon}^1)$ or $(\mathcal{P}_{y,\gamma}^2)$ are non-smooth, they enjoy enough structure to be solved by efficient algorithms. The type of algorithm to be used depends in particular on the properties of J . We first briefly mention some popular non-smooth optimization schemes in Section 6.1, and focus our attention on proximal splitting schemes afterwards.

6.1 Convex Optimization for Regularized Inverse Problems

6.1.1 (Sub)-gradient Descent

Consider for example problem $(\mathcal{P}_{y,\lambda})$. This is a convex composite optimization problem where one of the functions is smooth with a Lipschitz-continuous gradient. If J were smooth enough, then a simple gradient (or possibly (quasi-)Newton) descent method could be used. However, as detailed in Section 2.2, low-complexity regularizers J are intended to be non-smooth in order to promote models $\mathcal{M}$ of low intrinsic dimension, and J is precisely non-smooth transverse to $\mathcal{M}$. One can think of replacing gradients by subgradients (elements of the subdifferential), since J is assumed finite-valued (hence closed) convex, which are bounded. This results in a subgradient descent algorithm which is guaranteed to converge but under stringent assumptions on the descent step-sizes, which in turn makes their global convergence rate quite slow, see [172].

6.1.2 Interior Point Methods

Clearly, the key to getting efficient algorithms is to exploit the structure of the optimization problems at hand while handling non-smoothness properly. For a large class of regularizers J , such as those introduced in Section 2.3, the corresponding optimization problems can be cast as conic programs. The cone constraint can be enforced using a self-concordant barrier function, and the optimization problem can hence be solved using interior point methods, as pioneered by [174], see also the monograph [25]. This class of methods enjoys fast convergence rate. Each iteration however is typically quite costly and can become prohibitive as the dimension increases.

6.1.3 Conditional Gradient

This algorithm is historically one of the first method for smooth constrained convex optimization (a typical example being $(\mathcal{P}_{y,\gamma}^2)$), and was extensively studied in the 70's. It is also known as Frank-Wolfe algorithm, since it was introduced by [106] for quadratic programming and extended in [88]. The conditional gradient algorithm is premised on being able to easily solve (at each iteration) linear optimization problems over the feasible region of interest. This is in contrast to other first-order methods, such as forward-backward splitting and its variants (see Section 6.3), which are premised on being able to easily solve (at each iteration) a projection problem. Moreover, in many applications the solutions to the linear optimization subproblem are highly structured and exhibit particular sparsity and/or low-rank properties. These properties have renewed interest in the conditional gradient method to solve sparse recovery (ℓ_1 and total variation), low-rank matrix recovery (nuclear norm

minimization), anti-sparsity recovery, and various other problems in signal processing and machine learning; see e.g. [62, 136, 200, 87, 129].

6.1.4 Homotopy/Path-following

Homotopy and path-following-type methods have been introduced in the case of ℓ^1 -minimization to solve $(\mathcal{P}_{y,\lambda})$ by [176]. They were then adapted to analysis ℓ^1 , i.e. $J = \|D^* \cdot\|_1$, in [215], and ℓ^∞ regularization, $\|\cdot\|_\infty$, in [108]. One can in fact show that these methods can be applied to any polyhedral regularization (see [233]), because these methods only rely on the crucial fact that the solution path $\lambda \mapsto x_\lambda^*$, where x_λ^* is a solution of $(\mathcal{P}_{y,\lambda})$, is piecewise affine. The LARS algorithm [96] is an accelerated version of homotopy which computes an approximate homotopy path for $J = \|\cdot\|_1$ along which the support increases monotonically along the course of iterations. In the noiseless compressed sensing case, with Φ drawn from the Gaussian ensemble, it is shown in [82] that if x_0 is k -sparse with $P > 2k \log(N)$, the homotopy method reaches x_0 in only k iterations. This k -solution property was empirically observed for other random matrix ensembles, but at different thresholds for P . In [158], the authors proved that in the worst case, the number of segments in the solution path is exponential in the number of variables, and thus the homotopy method can then take as many iterations to converge.

As for interior points, the cost per iteration of homotopy-like methods, without particular ad hoc optimization, scales badly with the dimension, thus preventing them to be used for large-scale problems such as those encountered in imaging. This class of solvers is thus a wise choice for problems of medium size, and when high accuracy (or even exact computation up to machine precision for the homotopy algorithm) is needed. Extensions of these homotopy methods can deal with progressive changes in the operator Φ or the observations y , and are thus efficient for these settings, see [5].

6.1.5 Approximate Message Passing

In the last five years, ideas from graphical models and message passing and approximate message passing algorithms have been proposed to solve large-scale problems of the form $(\mathcal{P}_{y,\lambda})$ for various regularizers J , in particular ℓ^1 , $\ell^1 - \ell^2$ and the nuclear norm. A comprehensive review is given in [165]. However, rigorous convergence results have been proved so far only in the case in which Φ is standard Gaussian, though numerical results show that the same behavior should apply for broader random matrix ensembles.

6.2 Proximal Splitting Algorithms

Proximal splitting methods are first-order iterative algorithms that are tailored to solve structured non-smooth (essentially convex) optimization problems. The first operator splitting method has been developed from the 70's. Since then, the class of splitting methods have been regularly enriched with increasingly sophisticated algorithms as the structure of problems to handle becomes more complex.

To make our discussion more concrete, consider the general problem of minimizing the proper closed convex function

$$f = h + \sum_{k=1}^K g_k \circ A_k$$

where $h : \mathbb{R}^N \rightarrow \mathbb{R}$ is convex and smooth, the $A_k : \mathbb{R}^N \rightarrow \mathbb{R}^{N_k}$ are linear operators and $g_k : \mathbb{R}^{N_k} \rightarrow \mathbb{R}$ are proper closed convex functions for which the so-called proximity operator (to be defined shortly) can be computed easily (typically in closed form). We call such a function g_k “simple”.

Definition 12. The proximity operator of a proper closed convex function g is defined as, for $\gamma > 0$

$$\text{prox}_{\gamma g}(x) = \underset{u \in \mathbb{R}^N}{\operatorname{argmin}} \frac{1}{2} \|x - u\|^2 + \gamma g(u).$$

The proximal operator generalizes the notion of orthogonal projection onto a non-empty closed convex set C that one recovers by taking $g = \iota_C$.

Proximal splitting algorithms may evaluate (possibly approximately) the individual operators (e.g. gradient of h), the proximity operators of the g_k 's, the linear operators A_k , all separately at various points in the course of iteration, but never those of sums of functions nor composition by a linear operator. Therefore, each iteration is cheap to compute for large-scale problems. They also enjoy rigorous convergence guarantees, stability to errors, with possibly quantified convergence rates and iteration complexity bounds on various quantities. This justifies their popularity in contemporary signal and image processing or machine learning, despite that their convergence is either sublinear or at best linear.

It is beyond the scope of this Chapter to describe thoroughly the huge literature on proximal splitting schemes, as it is a large and extremely active research field in optimization theory. Good resources and reviews on the subject are [16, 13, 64, 178]. We instead give a brief classification of the most popular algorithms according to the class of structured objective functions they are able to handle:

- *Forward-Backward (FB)* algorithm [163, 179, 66]. It is designed to minimize (6.2) when h has a Lipschitz-continuous gradient, $K = 1$, $A_1 = \text{Id}$, and g_1 is simple. There are accelerated (optimal) variants of FB, such as the popular Nesterov [173] or Fista [15], but the convergence of the iterates is not longer guar-

teed for these schemes. FB and its variants are good candidates to solve $(\mathcal{P}_{y,\lambda})$. We will further elaborate on FB in Section 6.3.

- *Douglas-Rachford (DR)* algorithm [86, 153]. It is designed to minimize (6.2) for $h = 0$, $K = 2$, $A_k = \text{Id}$ and g_k is simple for $k = 1, 2$. It can be easily extended to the case of $K > 2$ by either lifting to a product space, see e.g. [63], or through projective splitting [94]. DR can be used to solve $(\mathcal{P}_{y,0})$, $(\mathcal{P}_{y,\varepsilon}^1)$ or $(\mathcal{P}_{y,\gamma}^2)$ for certain operators Φ .
- *Generalized Forward-Backward (GFB)* algorithm [185]. It can handle the case of an arbitrary K with $A_k = \text{Id}$, g_k simple and h has a Lipschitz-continuous gradient. It can be interpreted as hybridization of FB scheme and the DR scheme on a product space.
- *Alternate Direction Method of Multipliers (ADMM)* algorithm [104, 109, 110, 113]. It is adapted to minimize (6.2) for $h = 0$, $K = 2$ with $A_1 = \text{Id}$ and A_2 is injective. It can be shown [110, 93] that ADMM is equivalent to DR applied to the Fenchel-Rockafellar dual problem $\min_u g_1^* \circ -A_2^*(u) + g_2^*(u)$, where g_k^* is the Legendre-Fenchel conjugate of g_k . While DR applies when g_1 and $g_2 \circ A_2$ are simple, ADMM is a better alternative whereas both $g_1 \circ -A_2^*$ and g_2^* are simple. Extension to the case $K > 2$ was proposed for instance in [92].
- *Dykstra* algorithm [91]. It is able to solve the case where $h(x) = \|x - y\|^2$, $A_k = \text{Id}$ and the g_k are simple functions. It was initially introduced by [91] in the case where the g_k are indicator functions of closed convex sets, and is generalized in [12] to arbitrary convex functions. It is also extended in [51, 14] to the case where h is a Bregman divergence.
- *Primal-Dual schemes*. Recently, primal-dual splitting algorithms have been proposed to minimize (6.2) in its full generality, and even more complex objectives, see for instance [56, 223, 203, 29, 53, 65, 67, 239]. Primal-dual schemes can be used to solve $(\mathcal{P}_{y,\lambda})$, $(\mathcal{P}_{y,0})$, $(\mathcal{P}_{y,\varepsilon}^1)$ or $(\mathcal{P}_{y,\gamma}^2)$.

6.3 Finite Model Identification with Forward Backward

The FB algorithm is a good candidate to solve $(\mathcal{P}_{y,\lambda})$ when J is simple. Starting from some $x^{(0)} \in \mathbb{R}^N$, the FB iteration applied to $(\mathcal{P}_{y,\lambda})$ reads

$$x^{(n+1)} = \text{Prox}_{\tau_n \lambda J} \left(x^{(n)} + \tau_n \Phi^*(y - \Phi x^{(n)}) \right),$$

where the step-size sequence should satisfy $0 < \underline{\tau} \leq \tau_n \leq \bar{\tau} < 2/\|\Phi\|^2$ to ensure convergence of the sequence $x^{(n)}$ to a minimizer of $(\mathcal{P}_{y,\lambda})$.

In fact, owing to partial smoothness of J , much more can be said about the iterates of the FB algorithm. More precisely, after a finite number of iterations, Forward-Backward algorithm correctly identifies the manifold $\mathcal{M}$. This is made formal in the following theorem whose proof can be found in [151].

Theorem 7. *Under the assumptions of Theorem 2, $x^{(n)} \in \mathcal{M}$ for n large enough.*

This result sheds some light on the convergence behavior of this algorithm in the favorable case where condition (14) holds and $(\|w\|/\lambda, \lambda)$ are sufficiently small. In fact, it is shown in [151] that FB identifies in finite time the manifold of any non-degenerate minimizer x^* . As a corollary, if condition (14) holds at x_0 and $(\|w\|/\lambda, \lambda)$ are sufficiently small, then we recover Theorem 7. These results shed light on the typical convergence behavior of FB observed in such circumstances (e.g. in compressed sensing problems).

Remark 19 (Local linear convergence). The FB generally exhibits a global sublinear $O(1/n)$ convergence rate in terms of the objective function. However, under partial smoothness of J , it is shown in [151] that once the active manifold is identified, the FB algorithm enters a local linear convergence regime (Q -linear in general and R -linear if $\mathcal{M}$ is a linear manifold), whose rate can be characterized precisely in terms of the condition number of $\Phi_{T_{x_0}}$.

6.4 Related Works

Finite support identification and local R -linear convergence of FB to solve $(\mathcal{P}_{y,\lambda})$ is established in [27] under either a very restrictive injectivity assumption, or a non-degeneracy assumption that is a specialization of ours to the ℓ_1 norm. A similar result is proved in [125]. The ℓ_1 norm is a partly smooth function and is therefore covered by Theorem 7. [170] proved Q -linear convergence of FB to solve $(\mathcal{P}_{y,\lambda})$ with a data fidelity satisfying restricted smoothness and strong convexity assumptions, and J a so-called convex decomposable regularizer. Again, the latter falls within the class of partly smooth functions, and their result is then subsumed by our analysis.

For general programs, a variety of algorithms, such as proximal and projected-gradient schemes were observed to have the finite identification property of the active manifold. In [132, 130], the authors have shown finite identification of manifolds associated to partly smooth functions via the (sub)gradient projection method, Newton-like methods, and the proximal point algorithm. Their work extends that of e.g. [242] on identifiable surfaces from the smooth constrained convex case to a general non-smooth setting. Using these results, [131] considered the algorithm [224] to solve (6.2) when h is C^2 , $K = 1$, $A_1 = \text{Id}$, and g_1 is simple and partly smooth, but not necessarily convex, and proved finite identification of the active manifold. However, the convergence rates remain an open problem in all these works.

7 Summary and Perspectives

In this chapter, we have reviewed work covering a large body of literature on the regularization of linear inverse problems. We also showed how these previous works can be all seen as particular instances of a unified framework, namely sensitivity

analysis for minimization of convex partly smooth functions. We believe this general framework is the one that should be adopted as long as one is interested in studying fine properties and guarantees of these regularizers, and in particular when the stability of the low-complexity manifold associated to the data to recover is at stake.

This analysis is however only the tip of the iceberg, and there is actually a flurry of open problems to go beyond the theoretical results presented in this chapter. We list here a few ones that we believe are important avenues for future works:

- *Non-convexity and/or non-finiteness*: in this chapter, for the sake of simplicity, we focused on smooth convex fidelity terms and finite-valued convex regularizers. All the results stated in this chapter extend readily to proper lower semicontinuous convex regularizers, since any such a function is subdifferentially regular. Generalizations of some of the results to non-convex regularizers is possible as well, though some regularity assumptions are needed. This is of practical importance to deal with settings where Φ is not a linear operator, or to impose more aggressive regularization (for instance when using ℓ^p functional with $0 \leq p < 1$ instead of the ℓ^1 norm). There are however many difficulties to tackle in this case. For instance, regularity properties that hold automatically for the convex case have to be either imposed or proved. Another major bottleneck is that some of the results presented here, if extended verbatim, will only assess the recovery of a stationary/critical point. The latter is not a local minimum in general, and even less global.
- *Dictionary learning*: a related non-convex sensitivity analysis problem is to understand the recovery of the dictionary D in synthesis regularization (as defined in Section 2.3.4) when solving problems of the form

$$\min_{\{\alpha_k\}_k, D \in \mathcal{D}} \sum_k \frac{1}{2} \|y - \Phi D \alpha_k\|^2 + \lambda J_0(\alpha_k)$$

where the $(y_k)_k$ are a set of input exemplars and $\mathcal{D}$ stands for the set of constraints imposed on the dictionary to avoid trivial solutions. Such a non-convex variational problem is popular to compute adapted dictionaries, in particular when $J_0 = \|\cdot\|_1$, see [97] and references therein. Although the dictionary learning problem has been extensively studied when $J_0 = \|\cdot\|_1$, most of the methods lack of theoretical guarantees. The theory of dictionary learning is only beginning to develop, see e.g. [122, 140, 204, 1]. Tackling other regularizers, including analysis ℓ^1 of the form $J = J_0 \circ D^*$ is even more difficult, see e.g. [183, 59] for some computational schemes.

- *Infinite dimensional problems*: we dealt in this chapter with finite-dimensional vector spaces. It is not straightforward to extend these results to infinite-dimensional cases. As far as ℓ_2 -stability is concerned, the constants involved in the upper bounds depend on the dimension N , and the scaling might diverge as $N \rightarrow +\infty$. We refer to Section 3.3 for previous works on convergence rates of Tikhonov regularization in infinite-dimensional Hilbert or Banach spaces. Extending Theorem 2 for possibly non-reflexive Banach spaces is however still

out of reach (non-reflexivity is a typical degeneracy when considering low-complexity regularization). There exists however some extensions of classical stability results over spaces of measures, such as weak convergence [28], exact recovery [50, 74] and stable support recovery [90].

- *Compressed sensing*: as highlighted in Sections 3.3.3 and 4.5.3, the general machinery of partly smooth regularizers (and the associated dual certificates) is well adapted to derive optimal recovery bounds for compressed sensing. Unfortunately, this analysis has been for now only applied to norms ($\|\cdot\|_1$, $\|\cdot\|_{1,\mathcal{B}}$, $\|\cdot\|_*$ and $\|\cdot\|_\infty$). Extending this framework for synthesis and analysis regularizers (see Sections 2.3.4 and 2.3.5) is a difficult open problem.
- *Convergence and acceleration of the optimization schemes*: Section 6.3 showed how partial smoothness can be used to achieve exact manifold identification after a finite number of iterations using the FB algorithm. This in turn implies a local linear convergence of the iterates, and raises the hope of acceleration using either first-order or second-order information for the function along the identified manifold (in which we recall it is C^2). Studying such accelerations and their guarantees as well as extending this idea other proximal splitting schemes is thus of practical importance to tackle more complicated problems such as e.g. $(\mathcal{P}_{y,0})$, $(\mathcal{P}_{y,\epsilon}^1)$ or $(\mathcal{P}_{y,\gamma}^2)$.

Acknowledgements

This work has been supported by the European Research Council (ERC project SIGMA-Vision). We would like to thank our collaborators Charles Deledalle, Charles Dossal, Mohammad Golbabaee and Vincent Duval who have helped to build this unified view of the field.

References

1. A. Agarwal, A. Anandkumar, and P. Netrapalli. Exact recovery of sparsely used overcomplete dictionaries. arxiv, 2013.
2. H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Second international symposium on information theory*, pages 267–281. Springer Verlag, 1973.
3. J. Allen. Short-term spectral analysis, and modification by discrete Fourier transform. *IEEE Transactions on Acoustics Speech and Signal Processing*, 25(3):235–238, 1977.
4. D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: A geometric theory of phase transitions in convex optimization. *CoRR*, abs/1303.6672, 2013.
5. M. S. Asif and J. Romberg. Sparse recovery of streaming signals using L1-homotopy. Technical report, Preprint arxiv 1306.3331, 2013.
6. J.-F. Aujol, G. Aubert, L. Blanc-Féraud, and A. Chambolle. Image decomposition into a bounded variation component and an oscillating component. *Journal of Mathematical Imaging and Vision*, 22:71–88, 2005.

7. F. Bach. Consistency of the group Lasso and multiple kernel learning. *The Journal of Machine Learning Research*, 9(Jun):1179–1225, 2008.
8. F. Bach. Consistency of trace norm minimization. *The Journal of Machine Learning Research*, 9(Jun):1019–1048, 2008.
9. S. Bakin. *Adaptive regression and model selection in data mining problems*. PhD thesis, Australian National University, 1999.
10. A. S. Bandeira, E. Dobriban, D. G. Mixon, and W. F. Sawin. Certifying the restricted isometry property is hard. *IEEE transactions on information theory*, 59(6):3448–3450, 2013.
11. A. Barron, L. Birgé, and P. Massart. Risk bounds for model selection via penalization. *Probability theory and related fields*, 113(3):301–413, 1999.
12. H. H. Bauschke and P. L. Combettes. A dykstra-like algorithm for two monotone operators. *Pacific Journal of Optimization*, 4(3):383–391, 2008.
13. H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011.
14. H. H. Bauschke and A. S. Lewis. Dykstras algorithm with bregman projections: A convergence proof. *Optimization*, 48(4):409–427, 2000.
15. 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.
16. A. Beck and M. Teboulle. Gradient-based algorithms with applications to signal recovery. *Convex Optimization in Signal Processing and Communications*, 2009.
17. L. Birgé and P. Massart. From model selection to adaptive estimation. In D. Pollard, E. Torgersen, and L. Y. Grace, editors, *Festschrift for Lucien Le Cam*, chapter 4, pages 55–87. Springer New York, 1997.
18. L. Birgé and P. Massart. Minimal penalties for Gaussian model selection. *Probability theory and related fields*, 138(1-2):33–73, 2007.
19. T. Blu and F. Luisier. The SURE-LET approach to image denoising. *IEEE Trans. Image Process.*, 16(11):2778–2786, 2007.
20. T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265 – 274, 2009.
21. J. Bolte, A. Daniilidis, and A. S. Lewis. Generic optimality conditions for semialgebraic convex programs. *Mathematics of Operations Research*, 36(1):55–70, 2011.
22. C. Bonchelet. *Handbook of Image and Video Processing*, chapter Image Noise Models. Academic Press, Inc., 2005.
23. J. F. Bonnans and A. Shapiro. *Perturbation analysis of optimization problems*. Springer Series in Operations Research and Financial Engineering. Springer Verlag, 2000.
24. L. Borup, R. Gribonval, and M. Nielsen. Beyond coherence : recovering structured time-frequency representations. *Applied and Computational Harmonic Analysis*, 24(1):120–128, 2008.
25. S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
26. K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010.
27. K. Bredies and D. A. Lorenz. Linear convergence of iterative soft-thresholding. *Journal of Fourier Analysis and Applications*, 14(5-6):813–837, 2008.
28. K. Bredies and H.K. Pikkariainen. Inverse problems in spaces of measures. *ESAIM: Control, Optimisation and Calculus of Variations*, 19:190–218, 2013.
29. L. M. Briceño Arias and P. L. Combettes. A monotone+skew splitting model for composite monotone inclusions in duality. *SIAM Journal on Optimization*, 21(4):1230–1250, 2011.
30. M. Burger and S. Osher. Convergence rates of convex variational regularization. *Inverse Problems*, 20(5):1411, 2004.
31. T. T. Cai. Adaptive wavelet estimation: a block thresholding and oracle inequality approach. *Annals of statistics*, 27(3):898–924, 1999.
32. T.T. Cai and B.W. Silverman. Incorporating information on neighbouring coefficients into wavelet estimation. *Sankhya: The Indian Journal of Statistics, Series B*, pages 127–148, 2001.

33. T.T. Cai and H.H. Zhou. A data-driven block thresholding approach to wavelet estimation. *The Annals of Statistics*, 37(2):569–595, 2009.
34. E. J. Candès and D. L. Donoho. Curvelets: A surprisingly effective nonadaptive representation for objects with edges. Technical report, DTIC Document, 2000.
35. E. J. Candès, Y. C. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2011.
36. E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, June 2011.
37. E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
38. E. J. Candès and Y. Plan. A probabilistic and RIPless theory of compressed sensing. *Information Theory, IEEE Transactions on*, 57(11):7235–7254, 2011.
39. E. J. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *Information Theory, IEEE Transactions on*, 57(4):2342–2359, 2011.
40. E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
41. E. J. Candès and B. Recht. Simple bounds for recovering low-complexity models. *Mathematical Programming*, 141(1–2):577–589, 2013.
42. E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
43. E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.
44. E. J. Candès, C. A. Sing-Long, and J. D. Trzasko. Unbiased risk estimates for singular value thresholding and spectral estimators. *IEEE Transactions on Signal Processing*, 61(19):4643–4657, 2012.
45. E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
46. E. J. Candès and T. Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.
47. E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *Information Theory, IEEE Transactions on*, 52(12):5406–5425, 2006.
48. E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.
49. E. J. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2007.
50. E.J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.
51. Y. Censor and S. Reich. The dykstra algorithm with bregman projections. *Communications in Applied Analysis*, 2:407–419, 1998.
52. A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock. An introduction to total variation for image analysis. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*. De Gruyter, 2010.
53. 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.
54. V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.
55. C. Chaux, L. Duval, A. Benazza-Benyahia, and J-C. Pesquet. A nonlinear stein-based estimator for multichannel image denoising. *IEEE Transactions on Signal Processing*, 56(8):3855–3870, 2008.
56. G. Chen and M. Teboulle. A proximal-based decomposition method for convex minimization problems. *Mathematical Programming*, 64(1-3):81–101, 1994.

57. J. Chen and X. Huo. Theoretical results on sparse representations of multiple-measurement vectors. *Signal Processing, IEEE Transactions on*, 54(12):4634–4643, 2006.
58. S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1999.
59. Yunjin Chen, Thomas Pock, and Horst Bischof. Learning ℓ_1 -based analysis and synthesis sparsity priors using bi-level optimization. In *NIPS*, 2012.
60. R. Ciak, B. Shafei, and G. Steidl. Homogeneous penalizers and constraints in convex image restoration. *Journal of Mathematical Imaging and Vision*, 47:210–230, 2013.
61. J. F. Claerbout and F. Muir. Robust modeling with erratic data. *Geophysics*, 38(5):826–844, 1973.
62. K.L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. In *19th ACM-SIAM Symposium on Discrete Algorithms*, pages 922–931, 2008.
63. P. L. Combettes and J.-C. Pesquet. A proximal decomposition method for solving convex variational inverse problems. *Inverse Problems*, 24(6):065014, 2008.
64. P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In H. H. Bauschke, Burachik R. S., P. L. Combettes, Elser. V., D. R. Luke, and H. Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212. Springer, 2011.
65. P. L. Combettes and J. C. Pesquet. Primal–dual splitting algorithm for solving inclusions with mixtures of composite, lipschitzian, and parallel-sum type monotone operators. *Set-Valued and variational analysis*, 20(2):307–330, 2012.
66. P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
67. L. Condat. A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, pages 1–20, 2012.
68. M. Coste. An introduction to o-minimal geometry. Lecture notes, November 1999.
69. S. F. Cotter, B. D. Rao, J. Engan, and K. Kreutz-Delgado. Sparse solutions to linear inverse problems with multiple measurement vectors. *Signal Processing, IEEE Transactions on*, 53(7):2477–2488, 2005.
70. A. Daniilidis, D. Drusvyatskiy, and A. S. Lewis. Orthogonal invariance and identifiability. Technical report, arXiv 1304.1198, 2013.
71. A. Daniilidis, J. Malick, and H.S. Sendov. Spectral (isotropic) manifolds and their dimension. *Journal d’Analyse Mathématique*, 2014. to appear.
72. I. Daubechies, R. DeVore, M. Fornasier, and C. S. Gunturk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010.
73. G. Davis, S.G. Mallat, and Z. Zhang. Adaptive time-frequency approximations with matching pursuits. Technical report, Courant Institute of Mathematical Sciences, 1994.
74. Y. de Castro and F. Gamboa. Exact reconstruction using beurling minimal extrapolation. *Journal of Mathematical Analysis and Applications*, 395(1):336–354, 2012.
75. C. Deledalle, S. Vaite, G. Peyré, M. J. Fadili, and C. Dossal. Proximal Splitting Derivatives for Risk Estimation. In *2nd International Workshop on New Computational Methods for Inverse Problems (NCMIP)*, Paris, 2012.
76. C.-A. Deledalle, V. Duval, and J. Salmon. Non-local Methods with Shape-Adaptive Patches (NLM-SAP). *Journal of Mathematical Imaging and Vision*, pages 1–18, 2011.
77. C.-A. Deledalle, S. Vaite, G. Peyré, M.J. Fadili, and C. Dossal. Risk estimation for matrix recovery with spectral regularization. In *ICML’12 Workshops*, 2012. (arXiv:1205.1482v1).
78. D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *Information Theory, IEEE Transactions on*, 47(7):2845–2862, 2001.
79. D. L. Donoho and I. M. Johnstone. Adapting to Unknown Smoothness Via Wavelet Shrinkage. *Journal of the American Statistical Association*, 90(432):1200–1224, 1995.
80. D. L. Donoho and B. F. Logan. Signal recovery and the large sieve. *SIAM J. Appl. Math.*, 52(2):577–591, 1992.

81. D. L. Donoho and P. B. Stark. Uncertainty principles and signal recovery. *SIAM J. Appl. Math.*, 49(3):906–931, 1989.
82. D. L. Donoho and Y. Tsaig. Fast solution of ℓ^1 -norm minimization problems when the solution may be sparse. *Information Theory, IEEE Transactions on*, 54(11):4789–4812, 2008.
83. C. Dossal, M.-L. Chabanol, G. Peyré, and J. M. Fadili. Sharp support recovery from noisy random measurements by ℓ^1 -minimization. *Applied and Computational Harmonic Analysis*, 33(1):24–43, 2012.
84. C. Dossal, M. Kachour, M. J. Fadili, G. Peyré, and C. Chesneau. The degrees of freedom of the Lasso for general design matrix. *Statistica Sinica*, 23:809–828, 2013.
85. C. Dossal and S. Mallat. Sparse spike deconvolution with minimum scale. In *Proc. SPARS 2005*, 2005.
86. J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American mathematical Society*, 82(2):421–439, 1956.
87. M. Dudík, Z. Harchaoui, and J. Malick. Lifted coordinate descent for learning with trace-norm regularization. In Neil D. Lawrence and Mark Girolami, editors, *Proc. AISTATS*, volume 22 of *JMLR Proceedings*, pages 327–336. JMLR.org, 2012.
88. J. C. Dunn and S. Harshbarger. Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications*, 62(2):432–444, 1978.
89. V. Duval, J.-F. Aujol, and Y. Gousseau. A bias-variance approach for the non-local means. *SIAM Journal on Imaging Sciences*, 4(2):760–788, 2011.
90. V. Duval and G. Peyré. Exact support recovery for sparse spikes deconvolution. Technical report, Preprint hal-00839635, 2013.
91. R. L. Dykstra. An algorithm for restricted least squares regression. *J. Amer. Stat.*, 78:839–842, 1983.
92. J. Eckstein. Parallel alternating direction multiplier decomposition of convex programs. *Journal of Optimization Theory and Applications*, 80(1):39–62, 1994.
93. J. Eckstein and D. P. Bertsekas. On the douglas–rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
94. J. Eckstein and B. F. Svaiter. General projective splitting methods for sums of maximal monotone operators. *SIAM J. Control Optim.*, 48(2):787–811, 2009.
95. B. Efron. How biased is the apparent error rate of a prediction rule? *Journal of the American Statistical Association*, 81(394):461–470, 1986.
96. B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–451, 2004.
97. M. Elad. *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer, 2010.
98. M. Elad, P. Milanfar, and R. Rubinstein. Analysis versus synthesis in signal priors. *Inverse problems*, 23(3):947, 2007.
99. M. Elad, J.-L. Starck, P. Querre, and D. L. Donoho. Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA). *Applied and Computational Harmonic Analysis*, 19(3):340 – 358, 2005.
100. Y. C. Eldar. Generalized SURE for exponential families: Applications to regularization. *Signal Processing, IEEE Transactions on*, 57(2):471–481, 2009.
101. M. J. Fadili, G. Peyré, S. Vaïter, C.-A. Deledalle, and J. Salmon. Stable recovery with analysis decomposable priors. In *Proc. SampTA*, 2013.
102. M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford University, 2002.
103. M. Fazel, H. Hindi, and S. P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *American Control Conference, 2001. Proceedings of the 2001*, volume 6, pages 4734–4739. IEEE, 2001.
104. M. Fortin and R. Glowinski. *Augmented Lagrangian methods: applications to the numerical solution of boundary-value problems*. Access Online via Elsevier, 2000.

105. S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhäuser series in applied and numerical harmonic analysis. Birkhäuser, 2013.
106. M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
107. J.-J. Fuchs. On sparse representations in arbitrary redundant bases. *Information Theory, IEEE Transactions on*, 50(6):1341–1344, 2004.
108. J.-J. Fuchs. Spread representations. In *Signals, Systems and Computers (ASILOMAR)*, pages 814–817. IEEE, 2011.
109. D. Gabay. Applications of the method of multipliers to variational inequalities. In M. Fortin and R. Glowinski, editors, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-value Problems*, Amsterdam, 1983. North-Holland Publishing Company.
110. D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.
111. A. Girard. A fast Monte-Carlo cross-validation procedure for large least squares problems with noisy data. *Numerische Mathematik*, 56(1):1–23, 1989.
112. R. Giryas, M. Elad, and Y.C. Eldar. The projected GSURE for automatic parameter tuning in iterative shrinkage methods. *Applied and Computational Harmonic Analysis*, 30(3):407–422, 2011.
113. R. Glowinski and P. Le Tallec. *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*, volume 9. SIAM, 1989.
114. M. Golbabaee and P. Vandergheynst. Hyperspectral Image Compressed Sensing Via Low-Rank And Joint-Sparse Matrix Recovery. In *2012 IEEE International Conference On Acoustics, Speech And Signal Processing (ICASSP)*, pages 2741–2744. IEEE, 2012.
115. G.H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.
116. M. Grasmair. Linear convergence rates for Tikhonov regularization with positively homogeneous functionals. *Inverse Problems*, 27(7):075014, 2011.
117. M. Grasmair, O. Scherzer, and M. Haltmeier. Necessary and sufficient conditions for linear convergence of ℓ_1 -regularization. *Communications on Pure and Applied Mathematics*, 64(2):161–182, 2011.
118. E. Grave, G. Obozinski, and F. Bach. Trace Lasso: a trace norm regularization for correlated designs. In *Neural Information Processing Systems (NIPS)*, Spain, 2012.
119. R. Gribonval. Should penalized least squares regression be interpreted as maximum a posteriori estimation? *Signal Processing, IEEE Transactions on*, 59(5):2405–2410, 2011.
120. R. Gribonval and M. Nielsen. Beyond sparsity : recovering structured representations by ℓ^1 -minimization and greedy algorithms. *Advances in Computational Mathematics*, 28(1):23–41, 2008.
121. R. Gribonval, H. Rauhut, K. Schnass, and P. Vandergheynst. Atoms of all channels, unite! average case analysis of multi-channel sparse recovery using greedy algorithms. *Journal of Fourier analysis and Applications*, 14(5-6):655–687, 2008.
122. R. Gribonval and K. Schnass. Dictionary identification - sparse matrix factorization via ℓ_1 -minimization. *IEEE Transactions on Information Theory*, 56(7):3523–3539, 2010.
123. D. Gross. Recovering low-rank matrices from few coefficients in any basis. *Information Theory, IEEE Transactions on*, 57(3):1548–1566, 2011.
124. B.F. Svaiter H. Attouch, J. Bolte. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized gauss-seidel methods. *Mathematical Programming, Series A*, 137(1-2):91–129, 2013.
125. E. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence. *SIAM Journal on Optimization*, 19(3):1107–1130, 2008.
126. P. Hall, G. Kerkycharian, and D. Picard. On the minimax optimality of block thresholded wavelet estimators. *Statistica Sinica*, 9(1):33–49, 1999.
127. P. Hall, S. Penev, G. Kerkycharian, and D. Picard. Numerical performance of block thresholded wavelet estimators. *Statistics and Computing*, 7(2):115–124, 1997.

128. N. R. Hansen and A. Sokol. Degrees of freedom for nonlinear least squares estimation. Technical report, arXiv 1402.2997, 2014.
129. E. Harchaoui, A. Juditsky, and A. Nemirovski. Conditional gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*, pages 1–38, 2014.
130. W. Hare and A. S. Lewis. Identifying active manifolds. *Algorithmic Operations Research*, 2(2), 2007.
131. W. L. Hare. Identifying active manifolds in regularization problems. In H. H. Bauschke, R. S., Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, volume 49 of *Springer Optimization and Its Applications*, chapter 13. Springer, 2011.
132. W. L. Hare and A. S. Lewis. Identifying active constraints via partial smoothness and prox-regularity. *J. Convex Anal.*, 11(2):251–266, 2004.
133. J.-B. Hiriart-Urruty and H. Y. Le. Convexifying the set of matrices of bounded rank: applications to the quasiconvexification and convexification of the rank function. *Optimization Letters*, 6(5):841–849, 2012.
134. B. Hofmann, B. Kaltenbacher, C. Poeschl, and O. Scherzer. A convergence rates result for Tikhonov regularization in Banach spaces with non-smooth operators. *Inverse Problems*, 23(3):987, 2007.
135. H. M. Hudson. A natural identity for exponential families with applications in multiparameter estimation. *Annals of Statistics*, 6(3):473–484, 1978.
136. M. Jaggi and M. Sulovsky. A simple algorithm for nuclear norm regularized problems. In *ICML*, 2010.
137. H. Jégou, M. Douze, and C. Schmid. Improving bag-of-features for large scale image search. *International Journal of Computer Vision*, 87(3):316–336, 2010.
138. H. Jégou, T. Furon, and J.-J. Fuchs. Anti-sparse coding for approximate nearest neighbor search. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 2029–2032. IEEE, 2012.
139. R. Jenatton, J.Y. Audibert, and F. Bach. Structured variable selection with sparsity-inducing norms. *The Journal of Machine Learning Research*, 12:2777–2824, 2011.
140. R. Jenatton, R. Gribonval, and F. Bach. Local stability and robustness of sparse dictionary learning in the presence of noise. arxiv:1210.0685, 2012.
141. J. Jia and B. Yu. On model selection consistency of the elastic net when $p \gg n$. *Statistica Sinica*, 20:595–611, 2010.
142. K. Kato. On the degrees of freedom in shrinkage estimation. *Journal of Multivariate Analysis*, 100(7):1338–1352, 2009.
143. K. Knight and W. Fu. Asymptotics for Lasso-Type Estimators. *The Annals of Statistics*, 28(5):1356–1378, 2000.
144. J. M. Lee. *Smooth manifolds*. Springer, 2003.
145. C. Lemaréchal, F. Oustry, and C. Sagastizábal. The $\mathcal{U}$ -lagrangian of a convex function. *Trans. Amer. Math. Soc.*, 352(2):711–729, 2000.
146. A. S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM Journal on Optimization*, 13(3):702–725, 2002.
147. A. S. Lewis. The mathematics of eigenvalue optimization. *Mathematical Programming*, 97(1–2):155–176, 2003.
148. A. S. Lewis and J. Malick. Alternating projections on manifolds. *Mathematics of Operations Research*, 33(1):216–234, 2008.
149. A. S. Lewis and S. Zhang. Partial smoothness, tilt stability, and generalized Hessians. *SIAM Journal on Optimization*, 23(1):74–94, 2013.
150. K.-C. Li. From Stein’s unbiased risk estimates to the method of generalized cross validation. *The Annals of Statistics*, 13(4):1352–1377, 1985.
151. J. Liang, M.J. Fadili, and G. Peyré. Local linear convergence of forward-backward under partial smoothness. Technical report, arxiv preprint arXiv:1407.5611, 2014.
152. S. G. Lingala, Y. Hu, E. V. R. Di Bella, and M. Jacob. Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR. *IEEE Trans. Med. Imaging*, 30(5):1042–1054, 2011.

153. P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
154. D. Lorenz and N. Worliczek. Necessary conditions for variational regularization schemes. *Inverse Problems*, 29(7):075016, 2013.
155. D.A. Lorenz. Convergence rates and source conditions for Tikhonov regularization with sparsity constraints. *Journal of Inverse and Ill-Posed Problems*, 16(5):463–478, 2008.
156. F. Luisier, T. Blu, and M. Unser. Sure-let for orthonormal wavelet-domain video denoising. *Circuits and Systems for Video Technology, IEEE Transactions on*, 20(6):913–919, 2010.
157. Y. Lyubarskii and R. Vershynin. Uncertainty principles and vector quantization. *Information Theory, IEEE Transactions on*, 56(7):3491–3501, 2010.
158. J. Mairal and B. Yu. Complexity analysis of the lasso regularization path. In *ICML'12*, 2012.
159. S. G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 11(7):674–693, 1989.
160. S. G. Mallat. *A wavelet tour of signal processing*. Elsevier/Academic Press, Amsterdam, third edition, 2009.
161. S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, 41(12):3397–3415, 1993.
162. C. L. Mallows. Some comments on C_p . *Technometrics*, 15(4):661–675, 1973.
163. B. Mercier. Topics in finite element solution of elliptic problems. *Lectures on Mathematics*, 63, 1979.
164. M. Meyer and M. Woodroffe. On the degrees of freedom in shape-restricted regression. *Annals of Statistics*, 28(4):1083–1104, 2000.
165. A. Montanari. Graphical models concepts in compressed sensing. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing*. Cambridge University Press, 2012.
166. B. S. Mordukhovich. Sensitivity analysis in nonsmooth optimization. In D. A. Field and V. Komkov, editors, *Theoretical Aspects of Industrial Design*, volume 58, pages 32–46. SIAM Volumes in Applied Mathematics, 1992.
167. S. Nam, M. E. Davies, M. Elad, and R. Gribonval. The cosparsity analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013.
168. B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
169. D. Needell, J. Tropp, and R. Vershynin. Greedy signal recovery review. In *Signals, Systems and Computers, Conference on*, pages 1048–1050. IEEE, 2008.
170. S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, December 2012.
171. S. N. Negahban and M. J. Wainwright. Simultaneous support recovery in high dimensions: Benefits and perils of block-regularization. *Information Theory, IEEE Transactions on*, 57(6):3841–3863, 2011.
172. Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
173. Y. Nesterov. Gradient methods for minimizing composite objective function. CORE Discussion Papers 2007076, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), Sep 2007.
174. Y. Nesterov, A. Nemirovskii, and Y. Ye. *Interior-point polynomial algorithms in convex programming*, volume 13. SIAM, 1994.
175. G. Obozinski, B. Taskar, and M. I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20(2):231–252, 2010.
176. M. R. Osborne, B. Presnell, and B. A. Turlach. A new approach to variable selection in least squares problems. *IMA journal of numerical analysis*, 20(3):389–403, 2000.
177. S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *arXiv preprint arXiv:1212.3753*, 2012.
178. N. Parikh and S. P. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.

179. G. B. Passty. Ergodic convergence to a zero of the sum of monotone operators in hilbert space. *Journal of Mathematical Analysis and Applications*, 72(2):383–390, 1979.
180. Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Signals, Systems and Computers, Conference on*, pages 40–44. IEEE, 1993.
181. J.-C. Pesquet, A. Benazza-Benyahia, and C. Chaux. A SURE approach for digital signal/image deconvolution problems. *IEEE Transactions on Signal Processing*, 57(12):4616–4632, 2009.
182. G. Peyré, J. Fadili, and C. Chesneau. Adaptive structured block sparsity via dyadic partitioning. In *Proc. EUSIPCO 2011*, pages 1455–1459, 2011.
183. G. Peyré and M.J. Fadili. Learning analysis-sparsity priors. In *International Conference on Sampling Theory and Applications (SampTA)*, Singapour, 2011.
184. G. Peyré, M.J. Fadili, and J.-L. Starck. Learning the morphological diversity. *SIAM Journal on Imaging Sciences*, 3(3):646–669, 2010.
185. H. Raguet, J. Fadili, and G. Peyré. Generalized forward-backward splitting. *SIAM Journal on Imaging Sciences*, 6(3):1199–1226, 2013.
186. S. Ramani, T. Blu, and M. Unser. Monte-Carlo SURE: a black-box optimization of regularization parameters for general denoising algorithms. *Image Processing, IEEE Transactions on*, 17(9):1540–1554, 2008.
187. S. Ramani, Zhihao Liu, J. Rosen, J.-F. Nielsen, and J.A. Fessler. Regularization parameter selection for nonlinear iterative image restoration and mri reconstruction using GCV and SURE-based methods. *Image Processing, IEEE Transactions on*, 21(8):3659–3672, Aug 2012.
188. S. Ramani, J. Rosen, Z. Liu, and J. A. Fessler. Iterative weighted risk estimation for nonlinear image restoration with analysis priors. In *Computational Imaging X*, volume 8296, pages 82960N–82960N–12, 2012.
189. B.D. Rao and K. Kreutz-Delgado. An affine scaling methodology for best basis selection. *Signal Processing, IEEE Transactions on*, 47(1):187–200, Jan 1999.
190. B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
191. R. Refregier and F. Goudail. *Statistical Image Processing techniques for Noisy Images - An application Oriented Approach*. Kluwer, 2004.
192. E. Resmerita. Regularization of ill-posed problems in Banach spaces: convergence rates. *Inverse Problems*, 21(4):1303, 2005.
193. E. Richard, F. Bach, and J.-P. Vert. Intersecting singularities for multi-structured estimation. In *International Conference on Machine Learning*, Atlanta, États-Unis, 2013.
194. R. T. Rockafellar and R. Wets. *Variational analysis*, volume 317. Springer, Berlin, 1998.
195. M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008.
196. L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
197. F. Santosa and W.W. Symes. Linear inversion of band-limited reflection seismograms. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1307–1330, 1986.
198. O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*, volume 167. Springer, 2009.
199. I. W. Selesnick and M. A. T. Figueiredo. Signal restoration with overcomplete wavelet transforms: comparison of analysis and synthesis priors. In *Proceedings of SPIE*, volume 7446, page 74460D, 2009.
200. S. Shalev-Shwartz, A. Gonen, and O. Shamir. Large-scale convex minimization with a low-rank constraint. In *ICML*, 2011.
201. X. Shen and J. Ye. Adaptive model selection. *Journal of the American Statistical Association*, 97(457):210–221, 2002.
202. V. Solo and M. Ulfarsson. Threshold selection for group sparsity. In *Acoustics Speech and Signal Processing (ICASSP), IEEE International Conference on*, pages 3754–3757. IEEE, 2010.

203. M. V. Solodov. A class of decomposition methods for convex optimization and monotone variational inclusions via the hybrid inexact proximal point framework. *Optimization Methods and Software*, 19(5):557–575, 2004.
204. D. A. Spielman, H. Wang, and J. Wright. Exact recovery of sparsely-used dictionaries. *Journal of Machine Learning Research*, pages 1–35, 2012.
205. N. Srebro. *Learning with matrix factorizations*. PhD thesis, MIT, 2004.
206. J.-L. Starck, M. Elad, and D.L. Donoho. Image decomposition via the combination of sparse representations and variational approach. *IEEE Trans. Image Processing*, 14(10):1570–1582, 2005.
207. J.-L. Starck, F. Murtagh, and J. M. Fadili. *Sparse image and signal processing: wavelets, curvelets, morphological diversity*. Cambridge University Press, 2010.
208. G. Steidl, J. Weickert, T. Brox, P. Mrázek, and M. Welk. On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and sides. *SIAM Journal on Numerical Analysis*, 42(2):686–713, 2004.
209. C.M. Stein. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6):1135–1151, 1981.
210. T. Strohmer and R.W. Heath Jr. Grassmannian frames with applications to coding and communication. *Applied and Computational Harmonic Analysis*, 14(3):257–275, 2003.
211. C. Studer, W. Yin, and R. G. Baraniuk. Signal representations with minimum ℓ_∞ -norm. In *Communication, Control, and Computing, Proc. 50th Ann. Allerton Conf. on*, 2012.
212. H. L. Taylor, S. C. Banks, and J. F. McCoy. Deconvolution with the ℓ_1 norm. *Geophysics*, 44(1):39–52, 1979.
213. R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B. Methodological*, 58(1):267–288, 1996.
214. 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.
215. R. J. Tibshirani and J. Taylor. The solution path of the generalized Lasso. *The Annals of Statistics*, 39(3):1335–1371, 2011.
216. R. J. Tibshirani and J. Taylor. Degrees of freedom in Lasso problems. *The Annals of Statistics*, 40(2):1198–1232, 2012.
217. A. N. Tikhonov. Regularization of incorrectly posed problems. *Soviet Math. Dokl.*, 4:1624–1627, 1963.
218. A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization methods. *Soviet Math. Dokl.*, 4:1035–1038, 1963.
219. A. N. Tikhonov and V. Arsenin. *Solutions of ill-posed problems*. V. H. Winston and Sons, 1977.
220. A. M. Tillman and M. E. Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *Information Theory, IEEE Transactions on*, 60(2):1248–1259, 2014.
221. J. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*. Birkhäuser, 2014.
222. J.A. Tropp. Just relax: Convex programming methods for identifying sparse signals in noise. *Information Theory, IEEE Transactions on*, 52(3):1030–1051, 2006.
223. P. Tseng. Alternating projection-proximal methods for convex programming and variational inequalities. *SIAM Journal on Optimization*, 7(4):951–965, April 1997.
224. P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Math. Prog. (Ser. B)*, 117, 2009.
225. B. A. Turlach, W. N. Venables, and S. J. Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.
226. S. Vaiter, C. Deledalle, G. Peyré, J. Fadili, and C. Dossal. Degrees of freedom of the group Lasso. In *ICML’12 Workshops*, pages 89–92, 2012.
227. S. Vaiter, C. Deledalle, G. Peyré, J. Fadili, and C. Dossal. The degrees of freedom of partly smooth regularizers. Technical report, Preprint Hal-00768896, 2013.

228. S. Vaïter, C.-A. Deledalle, G. Peyré, C. Dossal, and J. Fadili. Local behavior of sparse analysis regularization: Applications to risk estimation. *Applied and Computational Harmonic Analysis*, 35(3):433–451, 2013.
229. S. Vaïter, M. Golbabaee, M. J. Fadili, and G. Peyré. Model selection with low complexity priors. Technical report, arXiv preprint arXiv:1307.2342, 2013.
230. S. Vaïter, G. Peyré, C. Dossal, and M. J. Fadili. Robust sparse analysis regularization. *IEEE Transactions on Information Theory*, 59(4):2001–2016, 2013.
231. S. Vaïter, G. Peyré, and J. Fadili. Model consistency of partly smooth regularizers. Technical report, Preprint Hal-00987293, 2014.
232. S. Vaïter, G. Peyré, J. M. Fadili, C.-A. Deledalle, and C. Dossal. The degrees of freedom of the group Lasso for a general design. In *Proc. SPARS'13*, 2013.
233. S. Vaïter, G. Peyré, and M. J. Fadili. Robust polyhedral regularization. In *Proc. SampTA*, 2013.
234. D. Van De Ville and M. Kocher. SURE-based Non-Local Means. *Signal Processing Letters, IEEE*, 16(11):973–976, 2009.
235. D. Van De Ville and M. Kocher. Non-local means with dimensionality reduction and SURE-based parameter selection. *IEEE Trans. Image Process.*, 9(20):2683–2690, 2011.
236. L. van den Dries and C. Miller. Geometric categories and o-minimal structures. *Duke Mathematical Journal*, 84(2):497–540, 08 1996.
237. J. E. Vogt and V. Roth. A complete analysis of the $\ell_{1,p}$ group-Lasso. In *Machine Learning, International Conference on*, 2012.
238. C. Vonesch, S. Ramani, and M. Unser. Recursive risk estimation for non-linear image deconvolution with a wavelet-domain sparsity constraint. In *Image Processing, International Conference on*, pages 665–668. IEEE, 2008.
239. B. C. Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, pages 1–15, 2011.
240. M. J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ^1 -constrained quadratic programming (Lasso). *IEEE Transactions on Information Theory*, 55(5):2183–2202, 2009.
241. S. J. Wright. Identifiable surfaces in constrained optimization. *SIAM Journal Control Optimisation*, 31(4):1063–1079, 1993.
242. S. J. Wright. Identifiable surfaces in constrained optimization. *SIAM Journal on Control and Optimization*, 31(4):1063–1079, 1993.
243. J. Ye. On measuring and correcting the effects of data mining and model selection. *Journal of the American Statistical Association*, pages 120–131, 1998.
244. 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, 2005.
245. 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.
246. P. Zhao and B. Yu. On model selection consistency of Lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.
247. H. Zou, T. Hastie, and R. Tibshirani. On the “degrees of freedom” of the Lasso. *The Annals of Statistics*, 35(5):2173–2192, 2007.