Total Variation Minimization in Compressed Sensing

Felix Krahmer, Christian Kruschel, Michael Sandbichler

APPLIEDMATHEMATICS

Total Variation Minimization in Compressed Sensing

Felix Krahmer, Christian Kruschel, and Michael Sandbichler

April 10, 2017

Abstract

This chapter gives an overview over recovery guarantees for total variation minimization in compressed sensing for different measurement scenarios. In addition to summarizing the results in the area, we illustrate why an approach that is common for synthesis sparse signals fails and different techniques are necessary. Lastly, we discuss a generalizations of recent results for Gaussian measurements to the subgaussian case.

1 Introduction

The central aim of Compressed Sensing (CS) [CRT06, Don06] is the recovery of an unknown vector from very few linear measurements. Put formally, we would like to recover $x \in \mathbb{R}^n$ from $y = Ax + e \in \mathbb{R}^m$ with $m \ll n$, where e denotes additive noise.

For general x, recovery is certainly not possible, hence additional structural assumptions are necessary in order to be able to guarantee recovery. A common assumption used in CS is that the signal is *sparse*. Here for x we assume

$$\|x\|_0 := |\{k \in [n]\colon x_k \neq 0\}| \leq s,$$

that is, there are only very few nonzero entries of x. And say that x is s-sparse for some given sparsity level $s \ll n$. We call a vector *compressible*, if it can be approximated well by a sparse vector. To quantify the quality of approximation, we let

$$\sigma_s(x)_q := \inf_{\|z\|_0 \leq s} \|z - x\|_q$$

denote the error of the best s-sparse approximation of x.

In most cases, the vector x is not sparse in the standard basis, but there is a basis Ψ, such that $x = \Psi z$ and z is sparse. This is also known as *synthesis sparsity* of x. To find an (approximately) synthesis sparse vector, we can instead solve the problem of recovering z from $y = A\Psi z$. A common strategy in CS is to solve a basis pursuit program in order to recover the original vector. For a fixed noise level ε, it is given by

$$\text{minimize } \|z\|_1 \text{ such that } \|Az - y\|_2 \leq \varepsilon. \tag{1}$$

While this and related approaches of convex regularization have been studied in the inverse problems and statistics literature long before the field of compressed sensing developed, these works typically assumed the measurement setup was given. The new paradigm arising in the context of

1

Figure 1: The original Fabio image (left) and the absolute values after application of a discrete gradient operator(right).

compressed sensing was to attempt to use the remaining degrees of freedom of the measurement system to reduce the ill-posedness of the system as much as possible. In many measurement systems, the most powerful known strategies will be based on randomization, i.e., the free parameters are chosen at random.

Given an appropriate amount of randomness (i.e., for various classes of random matrices A, including some with structure imposed by underlying applications), one can show that the minimizer $\hat{x}$ of (1) recovers the original vector x with error

$$\|x - \hat{x}\|_2 \le c \left(\frac{\sigma_s(x)_1}{\sqrt{s}} + \varepsilon \right), \tag{2}$$

see, e.g., [BDDW08] for an elementary proof in the case of subgaussian matrices without structure, and [KR14] for an overview, including many references, of corresponding results for random measurement systems with additional structure imposed by applications. Note that (2) entails that if x is s-sparse and the measurements are noiseless, the recovery is exact.

For many applications, however, the signal model of sparsity in an orthonormal basis has proven somewhat restrictive. Two main lines of generalization have been proposed. The first line of work, initiated by [RSV08] is the study of sparsity in redundant representation systems, at first under incoherence assumptions. More recently, also systems without such assumptions have been analyzed [CENR10, KNW15], the main idea being that even when due to conditioning problems one may not be able to recover the coefficients correctly, one can still hope for a good approximation of the signal.

The second line of work focuses on signals that are sparse after the application of some transform, one speaks of *cosparsity* or *analysis sparsity* [NDEG13], see, e.g., [KR15] for an analysis of the Gaussian measurement setup in this framework. A special case of particular importance, especially for imaging applications, is that of sparse gradients. Namely, as it turns out, natural images often admit very sparse approximations in the gradient domain, see, e.g., Figure 1. Here the discrete gradient at location $i = (i_1, \ldots, i_n)$ is defined as the vector with its n entries given by $\left((\nabla z)_i\right)_j = z_{i+e_j} - z_i$, $j = 1, \ldots, n$, where e_j is the j-th standard basis vector.

2

A first attempt to recover a gradient sparse signal is to formulate a compressed sensing problem in terms of the sparse gradient. When this is possible (for instance in the example of Fourier measurements [CRT06]), applying (1) will correspond to minimizing $\|\nabla z\|_1 =: \|z\|_{TV}$, the *total variation seminorm*. Then (under some additional assumptions) compressed sensing recovery guarantees of the form (2) can apply. This approach, however, only yields that the gradient can be approximately recovered, not the signal. When no noise is present and the gradient is exactly sparse (which is not very realistic), this allows for signal recovery via integrating the gradient, but in case of noisy measurements, this is highly unstable.

Nevertheless, this motivates to minimize the total variation seminorm if one attempts to recover the signal directly, not the gradient. In analogy with (1), this yields the following minimization problem.

$$\text{minimize } \|z\|_{TV} = \|\nabla z\|_1 \text{ such that } \|Az - y\|_2 \leq \varepsilon.$$

For A the identity (i.e., not reducing the dimension), this relates to the famous Rudin-Osher-Fatemi functional, a classical approach for signal and image denoising [ROF92]. Due to its high relevance for image processing, this special case of analysis sparsity has received a lot of attention recently also in the compressed sensing framework where A is dimension reducing. The purpose of this chapter is to give an overview of recovery results for total variation minimization in this context of compressed sensing (Section 2) and to provide some geometric intuition by discussing the one-dimensional case under Gaussian or subgaussian measurements (to our knowledge, a generalization to the latter case does not appear yet in the literature) with a focus on the interaction between the high-dimensional geometry and spectral properties of the gradient operator (Section 3).

2 An overview over TV recovery results

In this section, we will give an overview of the state of the art recovery guarantees for recovery of gradient sparse signals via total variation minimization. We start by discussing in Section 2.1 sufficient conditions for the success of TV minimization.

Subsequently, we focus on recovery results for random measurements. Interestingly, the results in one dimension differ severely from the ones in higher dimensions. Instead of obtaining a required number of measurements roughly on the order of the sparsity level s, we need $\sqrt{sn}$ measurements for recovery. We will see this already in Subsection 2.2, where we present the results of Cai and Xu [CX15] for recovery from Gaussian measurements. In Section 3, we will use their results to obtain refined results for noisy measurements as well as guarantees for Subgaussian measurements using an argument of Tropp [Tro15]. In Subsection 2.3 we will present results by Ward and Needell for dimensions larger or equal than two showing that recovery can be achieved from Haar incoherent measurements.

2.1 Sufficient Recovery Conditions

Suppose we are given linear measurements $Ax = y$ for an arbitrary $A \in \mathbb{R}^{m \times n}$ and a signal x, which is s-sparse after application of the discrete gradient defined as $(\nabla x)_i = x_{i+1} - x_i$ for $i = 1, \ldots, n-1$. Comparing this to the setup in Section 1, we see that a natural way to recover the signal is solving

$$\text{minimize } \|\nabla z\|_1 \text{ such that } Az = y. \tag{3}$$

3

For $I \subset [n]$ we denote A_I as the columns of A indexed by $I \subset [n]$, and for a consecutive notation we denote $\mathcal{I}_I^T \nabla$ as the rows of ∇ indexed by I and $\mathcal{I}$ as the identity matrix. The following results can be easily applied to any arbitrary matrix $D \in \mathbb{R}^{p \times n}$ which replaces ∇ in (3). This more general case is called *analysis ℓ_1-minimization*, we can highlight other information from the solutions of $Ax = y$ than few gradients. Moreover, most of the following results can be generalized to a real Hilbert space setting as well as to isotropic total variation [Kru15].

A crucial question for applying TV-minimization in applications is to verify whether there is exactly one solution of (3). As we consider ∇ not being injective, we cannot easily use the well-known recovery results in compressed sensing [FR13] for the matrix $A\nabla^{\dagger}$. However, we can give one crucial condition for unique solutions since x can only satisfy $Ax = y$ and $(\nabla x)_{I^c} = 0$ if

$$\ker(\mathcal{I}_{I^c}^T \nabla) \cap \ker(A) = \{0\}.$$

If ∇ is replaced by the identity, this is equivalent to A_I being injective. Since this injectivity condition is unavoidable, we assume for the rest of this section that it is satisfied.

The strongest condition is the following ∇-*null space property*, which gives sufficient and necessary conditions for unique solutions up to a certain sparsity of the transformed solutions.

Theorem 2.1. [NDEG13] Any $x \in \mathbb{R}^n$ with $s := \|\nabla x\|_0$ solves (3) uniquely if and only if for all $I \subset \{1, ..., n\}$ with $|I| \leq s$ it holds that

$$\forall w \in \ker(A) \backslash \{0\}: \quad \sup_{x \in \ker(\mathcal{I}_I^T \nabla)} |\langle (\nabla w)_I, \operatorname{sign}(\nabla x)_I \rangle| < \|(\nabla w)_{I^c}\|_1.$$

If ∇ is replaced by the identity, this is similar to the well-known *null space property* [DH01].

Corollary 2.1. For all $x \in \mathbb{R}^n$ with $s := \|\nabla x\|_0$, the solution of (3) with $y = Ax$ is unique and equal to x if for all $I \subset \{1, ..., n\}$ with $|I| \leq s$ it holds that

$$\forall w \in \ker(A) \backslash \{0\}: \quad \|(\nabla w)_I\|_1 < \|(\nabla w)_{I^c}\|_1.$$

However, it is hard to verify whether the latter condition is satisfied, even in the classical compressed sensing setup [TP14]. A common strategy to tackle this uniform recovery condition will be considered in the subsequent sections by assuming random measurements of different kind.

To consider measurements for specific applications, where it is difficult to prove when the *null space property* is satisfied, one can empirically examine whether specific elements x solve (3) uniquely. By repeatedly checking different x with $\|\nabla x\|_0 = s$, for being a unique solution of (3), one can approximate among all gradient s-sparse vectors the fraction of those who can uniquely be recovered via TV minimization. Such a *Monte Carlo Experiment* is done in [JKL15] where the TV-minimization is applied to computed tomography measurements. Here, different images $X \in \mathbb{R}^{n \times n}$ with gradient sparsity s were constructed via a random model and tested whether they are unique solutions of (3). The results prompt that there is a sharp transition between the case that every vector with a certain gradient sparsity is uniquely recoverable and the case that TV-minimization will find a different solution than the desired vector. This behavior empirically agrees with the phase transition that can be proved to appear for ∇ replaced by the identity and A a Gaussian measurement matrix [Don04].

To efficiently check whether many specific vectors x can be uniquely recovered via (3), one needs to establish characteristics of x which must be easily verifiable. Such a non-uniform recovery condition is given in the following theorem.

4

Theorem 2.2. [JKL15] It holds that $x \in \mathbb{R}^n$ is a unique solution of (3) if and only if there exists $w \in \mathbb{R}^m$ and $v \in \mathbb{R}^{n-1}$ such that

$$\nabla^T v = A^T w, v_I = \text{sign}(\nabla x)_I, \|v_{I^c}\|_\infty < 1. \tag{4}$$

The basic idea of the proof is using the optimality condition for convex optimization problems [Roc72]. Equivalent formulations of the latter theorem can be found in [ZMY16, KR15] where the problem is considered from a geometric perspective. However, verifying the conditions in Theorem 2.2 still requires solving the linear program

$$\text{minimize } \|v\|_\infty \text{ such that } v_I = \text{sign}(\nabla x)_I, \nabla^T v \in \text{rg}(A^T). \tag{5}$$

When ∇ is replaced by the identity, the *Fuchs Condition* [Fuc04] is known as a weaker result as it combines (4) to one short formula considering a particular w and avoids solving this linear program. Obviously, a weaker condition is left because the most of the null space of A is not considered. A similar result can be formulated for the analysis ℓ_1-minimization.

Corollary 2.2. If $x \in \mathbb{R}^n$ satisfies

$$\|(\mathcal{I}_{I^c}^T \nabla (\nabla^T \mathcal{I}_{I^c} \mathcal{I}_{I^c}^T \nabla + A^T A)^{-1} \nabla \text{sign}(\nabla x))_I\|_\infty < 1$$

then x is the unique solution of (3).

2.2 Recovery from Gaussian measurements

As discussed above, to date no deterministic constructions of compressed sensing matrices are known that get anywhere near an optimal number of measurements. Also for the variation of aiming to recover approximately gradient sparse measurements, the only near-optimal recovery guarantees have been established for random measurement models. Both under (approximate) sparsity and gradient sparsity assumptions, an important benchmark is that of a measurement matrix with independent standard Gaussian entries. Even though such measurements are hard to realize in practice, they can be interpreted as the scenario with maximal randomness, which often has particularly good recovery properties. For this reason, the recovery properties of total variation minimization have been analyzed in detail for such measurements. Interestingly, as shown by the following theorem, recovery properties in the one-dimensional case are significantly worse than for synthesis sparse signals and also for higher dimensional cases. That is why we focus on this case in Section 3, providing a geometric viewpoint and generalizing the results to subgaussian measurements.

Theorem 2.3 ([CX15]). Let the entries of $A \in \mathbb{R}^{m \times n}$ be i.i.d. standard Gaussian random variables and let $\hat{x}$ be a solution of (3) with input data $y = Ax_0$. Then

1. There exist constants $c_1, c_2, c_3, c_4 > 0$, such that for $m \geq c_1 \sqrt{sn}(\log n + c_2)$

$$\mathbb{P}(\forall x_0 \colon \|\nabla x_0\|_0 \leq s \colon \hat{x} = x_0) \geq 1 - c_3 e^{-c_4 \sqrt{m}}.$$

2. For any $\eta \in (0, 1)$, there are constants $\tilde{c}_1, \tilde{c}_2 > 0$ and a universal constant $c_2 > 0$, such that for $s \geq \tilde{c}_0$ and $(s + 1) < \frac{n}{4}$. If $m \leq \tilde{c}_1 \sqrt{sn} - \tilde{c}_2$, there exist infinitely many $x_0 \in \mathbb{R}^n$ with $\|\nabla x_0\|_0 \leq s$, such that $\mathbb{P}(\hat{x} \neq x_0) \geq 1 - \eta$.

Indeed, one would have expected by analogy with basis pursuit results a number of measurements on the order of s up to log factors. As we can for example see in Theorem 2.6 below, this is also the scaling one obtains for dimensions larger than 1. Thus the scaling of $\sqrt{sn}$ is a unique feature of the 1-dimensional case. Also note that the logarithmic factor in the upper bound makes the result meaningless for a sparsity level on the order of the dimension. This has been addressed in [KRZ15], showing that a dimension reduction is also possible if the sparsity level is a (small) constant multiple of the dimension.

The proof of Theorem 2.3 uses Gordon's escape through the mesh Theorem [Gor88]. We will elaborate on this topic in Section 3.

In case we are given noisy measurements $y = Ax_0 + e$ with $\|e\|_2 \leq \varepsilon$, we can instead of solving (3) consider

$$\text{minimize } \|\nabla z\|_1 \text{ such that } \|Az - y\|_2 \leq \varepsilon. \tag{6}$$

If ∇x_0 is not exactly, but approximately sparse, and our measurements are corrupted with noise, the following result can be established.

Theorem 2.4 ([CX15]). Let the entries of $A \in \mathbb{R}^{m \times n}$ be i.i.d. standard Gaussian random variables and let $\hat{x}$ be a solution of (6) with input data y satisfying $\|Ax_0 - y\|_2 \leq \varepsilon$. Then for any $\alpha \in (0, 1)$, there are positive constants $\delta, c_0, c_1, c_2, c_3$, such that for $m = \alpha n$ and $s = \delta n$

$$\mathbb{P}\left(\|x_0 - \hat{x}\|_2 \leq c_2 \frac{\min_{|S| \leq s} \|(\nabla x_0)_{S^c}\|_1}{\sqrt{n}} + c_3 \frac{\varepsilon}{\sqrt{n}} \right) \geq 1 - c_0 \mathrm{e}^{-c_1 n}.$$

This looks remarkably similar to the recovery guarantees obtained for compressed sensing, note however that the number of measurements needs to be proportional to n, which is not desirable. We will present a similar result with improved number of measurements in Section 3.5.

Theorem 2.5. (Corollary of Theorem 3.4) Let $x_0 \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m \times n}$ with $m \geq C\sqrt{ns}\log(2n)$ be a standard Gaussian matrix. Furthermore, set $y = Ax_0 + e$, where $\|e\| \leq \varepsilon$ denotes the (bounded) error of the measurement and for some absolute constants $c, \tilde{c} > 0$ the solution $\hat{x}$ of (13) satisfies

$$\mathbb{P}\left(\|\hat{x} - x_0\| > \frac{2\varepsilon}{c\sqrt[4]{ns}(\sqrt{\log(2n)} - 1)} \right) \leq \mathrm{e}^{-\tilde{c}\sqrt{ns}}.$$

Note, however that this theorem does not incorporate the case of compressible vectors, but on the other hand Theorem 3.4 also incorporates the case of special subgaussian measurement ensembles.

2.3 Recovery from Haar-incoherent measurements

For dimensions $d \geq 2$, Needell and Ward [NW13a, NW13b] derived recovery results for measurement matrices having the restricted isometry property (RIP) when composed with the Haar wavelet transform. Here we say that a matrix Φ has the RIP of order k and level δ if for every k-sparse vector x it holds that

$$(1 - \delta)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1 - \delta)\|x\|_2^2.$$

The results of [NW13a, NW13b] build upon a connection between a signal's wavelet representation and its total variation seminorm first noted by Cohen, Dahmen, Daubechies and De-Vore [CDDD03].

Their theorems yield stable recovery via TV minimization for N^d dimensional signals. For $d = 2$, notably these recovery results concern images of size $N \times N$.

Several definitions are necessary in order to be able to state the theorem. The d dimensional discrete gradient is defined via $\nabla \colon \mathbb{R}^{C^d} \to \mathbb{C}^{N^d \times d}$ and maps $x \in \mathbb{C}^{N^d}$ to its discrete derivative which, for each $\alpha \in [N]^d$ is a vector $(\nabla x)_\alpha \in \mathbb{C}^d$ composed of the derivatives in all d directions. Up to now, we have always used the anisotropic version of the TV seminorm, which can be seen as taking the ℓ_1 norm of the discrete gradient. The isotropic TV seminorm is defined via a combination of ℓ_2 and ℓ_1 norms. It is given by $\|z\|_{TV_2} := \sum_{\alpha \in [N]^d} \|(\nabla z)_\alpha\|_2$. The result in [NW13a] is given in terms of the isotropic TV seminorm but can also be formulated for the anisotropic version.

Furthermore, we will need to concatenate several measurement matrices in order to be able to state the theorem. This will be done via the concatenation operator $\oplus \colon \mathrm{Lin}(\mathbb{C}^n, \mathbb{C}^{k_1}) \times \mathrm{Lin}(\mathbb{C}^n, \mathbb{C}^{k_2}) \to \mathrm{Lin}(\mathbb{C}^n, \mathbb{C}^{k_1+k_2})$, which 'stacks' two linear maps.

Finally, we need the notion of shifted operators. For an operator $\mathcal{B} \colon \mathbb{C}^{N^{l-1} \times (N-1) \times N^{d-l}} \to \mathbb{C}^q$, these are defined as the operators $\mathcal{B}_{0_l} \colon \mathbb{C}^{N^d} \to \mathbb{C}^q$ and $\mathcal{B}^{0_l} \colon \mathbb{C}^{N^d} \to \mathbb{C}^q$ concatenating a column of zeros to the end or beginning of the l-th component, respectively.

Theorem 2.6 ([NW13a])**.** Let $N = 2^n$ and fix integers p and q. Let $\mathcal{A} \colon \mathbb{C}^{N^d} \to \mathbb{C}^p$ be a map that has the restricted isometry property of order $2ds$ and level $\delta < 1$ if it is composed with the orthonormal Haar wavelet transform. Furthermore let $\mathcal{B}_1, \ldots, \mathcal{B}_d$ with $\mathcal{B}_j \colon \mathbb{C}^{(N-1)N^{d-1}} \to \mathbb{C}^q$ be such that $\mathcal{B} = \mathcal{B}_1 \oplus \mathcal{B}_2 \oplus \cdots \oplus \mathcal{B}_d$ has the restricted isometry property of order $5ds$ and level $\delta < \frac{1}{3}$. Consider the linear operator $\mathcal{M} = \mathcal{A} \oplus [\mathcal{B}_1]_{0_1} \oplus [\mathcal{B}_1]^{0_1} \oplus \cdots \oplus [\mathcal{B}_d]_{0_d} \oplus [\mathcal{B}_d]^{0_d}$. Then $\mathcal{M} \colon \mathbb{C}^{N^d} \to \mathbb{C}^m$ with $m = 2dq + p$ and for all $x \in \mathbb{C}^{N^d}$ we have the following. Suppose we have noisy measurements $y = \mathcal{M}(x) + e$ with $\|e\|_2 \leq \varepsilon$, then the solution to

$$\hat{x} = \mathrm{argmin}_z \|z\|_{TV_2} \text{ such that } \|\mathcal{M}(z) - y\|_2 \leq \varepsilon$$

satisfies

1. $\|\nabla(x - \hat{x})\|_2 \leq c_1 \left(\frac{\|\nabla x - (\nabla x)_S\|_{1,2}}{\sqrt{s}} + \sqrt{d}\varepsilon \right)$,

2. $\|x - \hat{x}\|_{TV_2} \leq c_2 \left(\|\nabla x - (\nabla x)_S\|_{1,2} + \sqrt{sd}\varepsilon \right)$,

3. $\|x - \hat{x}\|_2 \leq c_3 d \log N \left(\frac{\|\nabla x - (\nabla x)_S\|_{1,2}}{\sqrt{s}} + \sqrt{d}\varepsilon \right)$,

for some absolute constants c_1, c_2, c_3.

From the last point of the previous theorem, we see that for noiseless measurements and gradient sparse vectors x, perfect recovery can be achieved provided the RIP assumption holds. Subgaussian measurement matrices, for example, will have the RIP, also when composed with the Haar wavelet transform H (this is a direct consequence of rotation invariance). Moreover, as shown in [KW11], randomizing the column signs of an RIP matrix will, with high probability, also yield a matrix that has the RIP when composed with H. An important example is a subsampled Fourier matrix with random column signs, which relates to spread spectrum MRI (cf. [PMG$^+$12]).

2.4 Recovery from subsampled Fourier measurements

Fourier measurements are widely used in many applications. Especially in medical applications as parallel-beam tomography and magnetic resonance imaging it is desirable to reduce the number of

samples to spare patients burden. In Section 2.1, this is a motivation for introducing algorithmic checks for unique solutions of (3). In this section, we consider a probabilistic approach where an incomplete measurement matrix $A \in \mathbb{C}^{m \times n}$ chosen from the discrete Fourier transform on $\mathbb{C}^N$ is considered. Therefore we consider a subset Ω of the index set $\{-\lfloor n/2 \rfloor + 1, ..., \lceil n/2 \rceil\}$, where Ω consists of m integers chosen uniformly at random and, additionally, $0 \in \Omega$. Hence, we want to recover a signal, sparse in the gradient domain, with a measurement matrix $A = (e^{2\pi i k j/n})_{k \in \Omega, j \in [n]}$. In [CRT06] the optimal sampling cardinality for s-sparse signals in the gradient domain was given and enables to recover one-dimensional signals signals from $\mathcal{O}(k \log(n))$ Fourier samples. It naturally extends to two dimensions.

Theorem 2.7. [CRT06] With probability exceeding $1 - \eta$, a signal z, which is k-sparse in the gradient domain is the unique solution of (3) if

$$m \gtrsim k(\log(n) + \log(\eta^{-1})).$$

As already discussed in the introduction, the proof of this result proceeds via recovering the gradient and then using that the discrete gradient (with periodic boundary conditions) is injective. Due to the poor conditioning of the gradient, however, this injectivity results do not directly generalize to recovery guarantees for noisy measurements. For two (and more) dimensions, such results can be obtained via the techniques discussed in the previous subsection.

These techniques, however, do not apply directly. Namely, the Fourier (measurement) basis is not incoherent to the Haar wavelet basis; in fact, the constant vector is contained in both, which makes them maximally coherent. As observed in [PVW11], this incoherence phenomenon only occurs for low frequencies, the high frequency Fourier basis vectors exhibit small inner products to the Haar wavelet basis. This can be taken into account using a *variable density* sampling scheme with sampling density that is larger for low frequencies and smaller for high frequencies. For such a sampling density, one can establish the restricted isometry for the corresponding randomly subsampled discrete Fourier matrix combined with the Haar wavelet transform with appropriately rescaled rows [KW14]. This yields the following recovery guarantee.

Theorem 2.8. [KW14] Fix integers $N = 2^p, m$, and s such that $s \gtrsim \log(N)$ and

$$m \gtrsim s \log^3(s) \log^5(N). \tag{7}$$

Select m frequencies $\{(\omega_1^j, \omega_2^j)\}_{j=1}^m \subset \{-N/2 + 1, \ldots, N/2\}^2$ i.i.d. according to

$$\mathbb{P}\big[(\omega_1^j, \omega_2^j) = (k_1, k_2)\big] = C_N \min\left(C, \frac{1}{k_1^2 + k_2^2}\right) =: \eta(k_1, k_2), \quad -N/2 + 1 \leq k_1, k_2 \leq N/2, \tag{8}$$

where C is an absolute constant and C_N is chosen such that η is a probability distribution. Consider the weight vector $\rho = (\rho_j)_{j=1}^m$ with $\rho_j = (1/\eta(\omega_1^j, \omega_2^j))^{1/2}$, and assume that the noise vector $\xi = (\xi_j)_{j=1}^m$ satisfies $\|\rho \circ \xi\|_2 \leq \varepsilon\sqrt{m}$, for some $\epsilon > 0$. Then with probability exceeding $1 - N^{-C \log^3(s)}$, the following holds for all images $f \in \mathbb{C}^{N \times N}$:
Given noisy partial Fourier measurements $y = \mathcal{F}_\Omega f + \xi$, the estimation

$$f^\# = \text{argmin}_{g \in \mathbb{C}^{N \times N}} \|g\|_{TV} \quad \text{such that} \quad \|\rho \circ (\mathcal{F}_\Omega g - y)\|_2 \leq \varepsilon\sqrt{m}, \tag{9}$$

where $\circ$ denotes the Hadamard product, approximates f up to the noise level and best s-term approximation error of its gradient:

$$\|f - f^\#\|_2 \lesssim \frac{\|\nabla f - (\nabla f)_s\|_1}{\sqrt{s}} + \varepsilon. \tag{10}$$

8

A similar optimality result is given in [Poo15], also for noisy data and inexact sparsity. In contrast to the previous result, this result includes the one-dimensional case. The key to obtaining such a result is showing that the stable gradient recover implies the stable signal recovery, i.e.,

$$\|z\|_2 \lesssim \gamma + \|z\|_{TV} \text{ with } \|Az\|_2 \le \gamma. \tag{11}$$

Again the sampling distribution is chosen as a combination of the uniform distribution and a decaying distribution. The main idea is to use this sampling to establish (11) via the RIP. We skip technicalities for achieving the optimality in the following theorem and refer to the original article for more details.

Theorem 2.9. [Poo15] Let $z \in \mathbb{C}^n$ be fixed and x be a minimizer of (6) with $\varepsilon = \sqrt{m}\delta$ for some $\delta > 0$, $m \gtrsim k \log(n)(1+\log(\eta^{-1}))$, and an appropriate sampling distribution. Then with probability exceeding $1 - \eta$, it holds that

$$\|\nabla z - \nabla x\|_2 \lesssim \left(\delta\sqrt{k} + C_1 \frac{\|P\nabla z\|_1}{\sqrt{k}} \right), \frac{\|z - x\|_2}{\sqrt{n}} \lesssim C_2 \left(\frac{\delta}{\sqrt{s}} + C_1 \frac{\|P\nabla z\|_1}{k} \right),$$

where P is the orthogonal projection onto a k-dimensional subspace,

$$C_1 = \log(k) \log^{1/2}(m), \text{ and } C_2 = \log^2(k) \log(n) \log(m).$$

In the two-dimensional setting the result changes to

$$\|\nabla z - \nabla x\|_2 \lesssim \left(\delta\sqrt{k} + C_3 \frac{\|P\nabla z\|_1}{\sqrt{k}} \right), \|z - x\|_2 \lesssim C_2 \left(\delta + C_3 \frac{\|P\nabla z\|_1}{k} \right),$$

with remaining C_2 and

$$C_3 = \log(k) \log(n^2/k) \log^{1/2}(n) \log^{1/2}(m).$$

These results are optimal since the best error one can archive [NW13b] is $\|z - x\|_2 \lesssim \|P\nabla z\|_1 k^{-1/2}$.

The optimality in the latter theorems is achieved by considering a combination of uniform random samling and variable density sampling. Uniform sampling on its own can achieve robust and stable recovery. However, the following theorem shows that the signal error is no longer optimal but the bound on the gradient error is still optimal up to log factors. Here (11) is obtained by using the Poincaré inequality.

Theorem 2.10. [Poo15] Let $z \in \mathbb{C}^n$ be fix and x be a minimizer of (6) with $\varepsilon = \sqrt{m}\delta$ for some $\delta > 0$ and $m \gtrsim k \log(n)(1 + \log(\eta^{-1}))$ with random uniform sampling. Then with probability exceeding $1 - \eta$, it holds that

$$\|\nabla z - \nabla x\|_2 \lesssim \left(\delta\sqrt{k} + C \frac{\|P\nabla z\|_1}{\sqrt{k}} \right), \frac{\|z - x\|_2}{\sqrt{n}} \lesssim (\delta\sqrt{s} + C\|P\nabla z\|_1),$$

where P is the orthogonal projection onto a k-dimensional subspace and $C = \log(k) \log^{1/2}(m)$.

3 TV-recovery from subgaussian measurements in 1D

In this section, we will apply the geometric viewpoint discussed in [Ver15] to the problem, which will eventually allow us to show the TV recovery results for noisy subgaussian measurements mentioned in Section 2.2.

As in the original proof of the 1D recovery guarantees for Gaussian measurements [CX15], the *Gaussian mean width* will play an important role in our considerations.

Definition 3.1. The (Gaussian) mean width of a bounded subset K of $\mathbb{R}^n$ is defined as

$$w(K) := \mathbb{E} \sup_{x \in K-K} \langle g, x \rangle,$$

where $g \in \mathbb{R}^n$ is a vector of i.i.d. $\mathcal{N}(0,1)$ random variables.

In [CX15], the mean width appears in the context of the *Gordon's escape through the mesh* approach [Gor88] (see Section 3.4 below), but as we will see, it will also be a crucial ingredient in applying the Mendelson small ball method [KM15, Men14].

The mean width has some nice (and important) properties, it is for example invariant under taking the convex hull, i.e.

$$w(\mathrm{ch}(K)) = w(K).$$

Furthermore, it is also invariant under translations of K, as $(K - x_0) - (K - x_0) = K - K$. Due to the rotational invariance of Gaussian random variables, i.e. $Ug \sim g$, we also have that $w(UK) = w(K)$. Also, it satisfies the inequalities

$$w(K) = \mathbb{E} \sup_{x \in K-K} \langle g, x \rangle \leq 2\mathbb{E} \sup_{x \in K} \langle g, x \rangle \leq 2\mathbb{E} \sup_{x \in K} |\langle g, x \rangle|,$$

which are equalities if K is symmetric about 0, because then $K = -K$ and hence $K - K = 2K$.

3.1 M^* bounds and recovery

In order to highlight the importance of the Gaussian mean width in signal recovery, we present some arguments from [Ver15]. Thus in this section we present a classical result, the M^* bound, which connects the mean width to recovery problems, cf. [Ver15]. Namely, recall that due to rotational invariance, the kernel of a Gaussian random matrix $A \in \mathbb{R}^{m \times n}$ is a random subspace distributed according to the uniform distribution (the Haar measure) on the Grassmannian

$$G_{n,n-m} := \{V \leq \mathbb{R}^n \colon \dim(V) = n - m\}.$$

Consequently, the set of all vectors that yield the same measurements directly correspond to such a random subspace.

The average size of the intersection of this subspace with a set reflecting the minimization objective now gives us an average bound on the worst case error.

Theorem 3.1 (M^* bound, Theorem 3.12 in [Ver15])**.** Let K be a bounded subset of $\mathbb{R}^n$ and E be a random subspace of $\mathbb{R}^n$ of drawn from the Grassmanian $G_{n,n-m}$ according to the Haar measure. Then

$$\mathbb{E} \operatorname{diam}(K \cap E) \leq C \frac{w(K)}{\sqrt{m}}, \tag{12}$$

10

where C is absolute constant.

Given the M^*-bound it is now straightforward to derive bounds on reconstructions from linear observations. We first look at feasibility programs - which in turn can be used to obtain recovery results for optimization problems. For that, let $K \subset \mathbb{R}^n$ be bounded and $x \in K$ be the vector we seek to reconstruct from measurements $Ax = y$ with a Gaussian matrix $A \in \mathbb{R}^{m \times n}$.

Corollary 3.1. [MPTJ07] Choose $\hat{x} \in \mathbb{R}^n$, such that

$$\hat{x} \in K \text{ and } A\hat{x} = y,$$

then one has, for an absolute constant C',

$$\mathbb{E} \sup_{x \in K} \|\hat{x} - x\|_2 \leq C' \frac{w(K)}{\sqrt{m}}.$$

This corollary directly follows by choosing $C' = 2C$, observing that $\hat{x} - x \in K - K$, and that the side constraint enforces $A(\hat{x} - x) = 0$.

Via a standard construction in functional analysis, the so called *Minkowski functional*, one can now cast an optimization problem as a feasiblity program so that Corollary 3.1 applies.

Definition 3.2. The Minkowski functional of a bounded convex set $K \subset \mathbb{R}^n$ is given by

$$\| \cdot \|_K \colon \mathbb{R}^n \to \mathbb{R} \colon x \mapsto \inf\{t > 0 \colon x \in tK\}.$$

So the Minkowski functional tells us, how much we have to 'inflate' our given set K in order to capture the vector x. Clearly, from the definition we have that if K is closed

$$K = \{x \colon \|x\|_K \leq 1\}.$$

If a convex set K is closed and symmetric, then $\| \cdot \|_K$ defines a norm on $\mathbb{R}^n$.

For bounded, star shaped K, the notion of $\|\cdot\|_K$ now allows to establish a direct correspondence between norm minimization problems and feasibility problems. With this observation, Corollary 3.1 translates to the following result.

Corollary 3.2. For K bounded and star-shaped, let $x \in K$ and $y = Ax$. Choose $\hat{x} \in \mathbb{R}^n$, such that it solves

$$\min \|z\|_K \text{ with } Az = y,$$

then

$$\mathbb{E} \sup_{x \in K} \|\hat{x} - x\|_2 \leq C' \frac{w(K)}{\sqrt{m}}.$$

Here $\hat{x} \in K$ is due to the fact that the minimum satisfies $\|\hat{x}\|_K \leq \|x\|_K \leq 1$, as $x \in K$ by assumption.

This result directly relates recovery guarantees to the mean width, it thus remains to calculate the mean width for the sets under consideration. In the following subsections, we will discuss two cases. The first one directly corresponds to the desired signal model, namely gradient sparse vectors. These considerations are mainly of theoretical interest, as the associated minimization problem closely relates to support size minimization, which is known to be NP hard in general. The second case considers the TV minimization problem introduced above, which then also yields guarantees for the (larger) set of vectors with bounded total variation.

11

3.2 The mean width of gradient sparse vectors in 1d

Here [PV13] served as an inspiration, as the computation is very similar for the set of sparse vectors.

Definition 3.3. The jump support of a vector x is given via

$$\mathrm{Jsupp}(x) := \{i \in [n-1] \colon x_{i+1} - x_i \neq 0\}.$$

The jump support captures the positions, in which a vector x changes its values. With this, we now define the set

$$K_0^s := \{x \in \mathbb{R}^n \colon \|x\|_2 \leq 1, |\mathrm{Jsupp}(x)| \leq s\}.$$

The set K_0^s consists of all s-gradient sparse vectors, which have 2-norm smaller than one. We will now calculate the mean width of K_0^s in order to apply Corrolary 3.1 or 3.2.

Note that we can decompose the set K_0^s into smaller sets $K_J \cap B_2^n$ with $K_J = \{x \colon \mathrm{Jsupp}(x) \subset J\}$, $|J| = s$ and $B_2^n = \{x \in \mathbb{R}^n \colon \|x\|_2 \leq 1\}$. As we can't add any jumps within the set K_J, it is a subspace of $\mathbb{R}^n$. We can even quite easily find an orthonormal basis for it, if we define

$$(e_{[i,j]})_k := \frac{1}{\sqrt{j-i+1}} \begin{cases} 1, & \text{if } k \in [i,j] \\ 0, & \text{else} \end{cases} \quad .$$

As we can align all elements of $J = \{j_1, j_2, \ldots, j_s\}$ with $1 \leq j_1 < j_2 < \ldots < j_s = n$, we see that $\{e_{[1,j_1]}, e_{[j_1+1,j_2]}, e_{[j_2+1,j_3]}, \ldots, e_{[j_{s-1}+1,j_s]}\}$ forms an ONB of K_J. Now, we can write all elements $x \in K_J \cap B_2^n$ as $x = \sum_{i=1}^s \alpha_i e_{[j_{i-1}+1,j_i]}$ by setting $j_0 := 0$. The property that $x \in B_2^n$ now enforces (ONB) that $\|\alpha\|_2 \leq 1$. Now, note that $K_0^s = -K_0^s$, so we have

$$w(K_0^s) = \mathbb{E} \sup_{x \in K_0^s - K_0^s} \langle g, x \rangle = 2\mathbb{E} \sup_{x \in K_0^s} \langle g, x \rangle.$$

Using the decomposition $K_0^s = \bigcup_{|J|=s} (K_J \cap B_2^n)$, we get

$$w(K_0^s) = 2\mathbb{E} \sup_{|J|=s} \sup_{x \in K_J \cap B_2^n} \langle g, x \rangle.$$

Now

$$\sup_{x \in K_J \cap B_2^n} \langle g, x \rangle \leq \sup_{\alpha \in B_2^s} \sum_{i=1}^s \alpha_i \langle g, e_{[j_{i-1}+1,j_i]} \rangle = \sup_{\alpha \in B_2^s} \sum_{i=1}^s \alpha_i \underbrace{\sum_{k=j_{i-1}+1}^{j_i} \frac{g_k}{\sqrt{j_i - j_{i-1}}}}_{=:G_i^J} .$$

Note that G_i^J is again a Gaussian random variable with mean 0 and variance 1. Furthermore, the supremum over α is attained, if $\alpha \parallel G^J$, so we have $\sup_{x \in K_J \cap B_2^n} \langle g, x \rangle = \|G^J\|_2$. Also note that G^J has i.i.d. entries, but for different J_1, J_2, the random vectors G^{J_1} and G^{J_2} may be dependent. Our task is now to calculate $\mathbb{E} \sup_{|J|=s} \|G^J\|_2$. As it has been shown for example in [FR13], we have that

$$\sqrt{\frac{2}{\pi}} \sqrt{s} \leq \mathbb{E}\|G^J\|_2 \leq \sqrt{s}$$

and from standard results for Gaussian concentration (cf. [PV13]), we get

$$\mathbb{P}(\|G^J\|_2 \geq \sqrt{s} + t) \leq \mathbb{P}(\|G^J\|_2 \geq \mathbb{E}\|G_J\|_2 + t) \leq e^{-t^2/2}.$$

By noting that $|\{J \subset [n] : |J| = s\}| = \binom{n}{s}$, we see by a union bound that

$$\mathbb{P}(\sup_{|J|=s} \|G^J\|_2 \geq \sqrt{s} + t) \leq \binom{n}{s} \mathbb{P}(\|G^J\|_2 \geq \sqrt{s} + t) \leq \binom{n}{s} e^{-t^2/2}.$$

For the following calculation, set $X := \sup_{|J|=s} \|G^J\|_2$. By Jensen's inequality and rewriting the expectation, we have that

$$e^{\lambda \mathbb{E} X} \leq \mathbb{E} e^{\lambda X} = \int_0^\infty \mathbb{P}(e^{\lambda X} \geq \tau) d\tau.$$

Now, the previous consideration showed, that

$$\mathbb{P}(e^{\lambda X} \geq \underbrace{e^{\lambda(\sqrt{s}+t)}}_{=:\tau}) = \mathbb{P}(X \geq \sqrt{s} + t) \leq \binom{n}{s} e^{-t^2/2} = \binom{n}{s} e^{-(\log(\tau)/\lambda - \sqrt{s})^2/2},$$

Computing the resulting integrals yields

$$e^{\lambda \mathbb{E} X} \leq \binom{n}{s} e^{-s/2} \lambda \sqrt{2\pi} e^{(\sqrt{s}+\lambda)^2/2}.$$

Using a standard bound for the binomial coefficients, namely $\binom{n}{s} \leq e^{s \log(en/s)}$, we see

$$e^{\lambda \mathbb{E} X} \leq e^{s \log(en/s) - s/2 + (\sqrt{s}+\lambda)^2/2 + \log(\lambda) + \log(\sqrt{2\pi})},$$

or equivalently

$$\lambda \mathbb{E} X \leq s \log(en/s) - s/2 + (\sqrt{s} + \lambda)^2/2 + \log(\lambda) + \log(\sqrt{2\pi})$$

By setting $\lambda = \sqrt{s \log(en/s)}$ and assuming (reasonably) large n, we thus get

$$\mathbb{E} X \leq 5\sqrt{s \log(en/s)}.$$

From this, we see that

$$w(K_0^s) \leq 10\sqrt{s \log(en/s)}.$$

It follows that the Gaussian mean width of the set of gradient sparse vectors is the same as the mean width of sparse vectors due to the similar structure. If we want to obtain accuracy δ for our reconstruction, according to Theorem 3.1, we need to take

$$m = \mathcal{O}\left(\frac{s \log(en/s)}{\delta^2}\right)$$

measurements.

In Compressed Sensing, the squared mean width of the set of s-sparse vectors (its so called *statistical dimension*) already determines the number of required measurements in order to recover a sparse signal with basis pursuit. This is the case because the convex hull of the set of sparse vectors can be embedded into the ℓ_1-ball inflated by a constant factor. In the case of TV minimization, as we will see in the following section, this embedding yields a (rather large) constant depending on the dimension.

13

3.3   Extensions to the TV-norm

In the previous subsection, we considered exactly gradient sparse vectors. However searching all such vectors x that satisfy $Ax = y$ is certainly not a feasible task. Instead, we want to solve the convex program

$$\min \|z\|_{TV} \text{ with } Az = y,$$

with $\|z\|_{TV} = \|\nabla z\|_1$ the total variation seminorm. Now if we have that $x \in K_0^s$, we get that

$$\|x\|_{TV} \le 2\|\alpha\|_1 \le 2\sqrt{s}\|\alpha\|_2 = 2\sqrt{s},$$

with α as in section 3.2, so $K_0^s \subset K_{TV}^{2\sqrt{s}} := \{x \in B_2^n : \|x\|_{TV} \le 2\sqrt{s}\}$. As $K_{TV}^{2\sqrt{s}}$ is convex, we even have $\mathrm{ch}(K_0^s) \subset K_{TV}^{2\sqrt{s}}$. We can think of the set $K_{TV}^{2\sqrt{s}}$ as 'gradient- compressible' vectors.

In the proof of Theorem 3.3 in [CX15], the Gaussian width of the set $K_{TV}^{4\sqrt{s}}$ has been calculated via a wavelet-based argument. One obtains that $w(K_{TV}^{2\sqrt{s}}) \le C\sqrt{\sqrt{ns}\log(2n)}$ with $C \le 20$ being an absolute constant.

The proof strategy fundamentally differs from the analogous results for synthesis sparsity. To illustrate the necessity of the new approach, we will illustrate in this subsection how the tried and tested strategy leads highly suboptimal bounds when applied for the mean width of $K_{TV}^{2\sqrt{s}}$. To find a constant, such that it is contained in the 'inflated' set $c_{n,s}\mathrm{ch}(K_0^s)$. Then $w(K_{TV}^{2\sqrt{s}}) \le c_{n,s}w(K_0^s)$. This works well for sparse recovery, where $c_{n,s} = 2$, but pityably fails in the case of TV recovery as we will see below.

Let us start with $x \in K_{TV}^{2\sqrt{s}}$. Now we can decompose $J := \mathrm{Jsupp}(x) = J_1 \uplus J_2 \uplus \dots J_p$ with $|J_k| \le s$ in an ascending manner, i.e. for all $k \in J_i, l \in J_{i+1}$, we have that $\alpha_k < \alpha_l$. Note that the number p of such sets satisfies $p \le \frac{n}{s}$. Similarly as above, we now write $x = \sum_{i=1}^{|J|} \alpha_i e_{[j_{i-1}+1,j_i]} = \sum_{k=1}^{p} \sum_{i \in J_k} \alpha_i e_{[j_{i-1}+1,j_i]}$. From this, we see that

$$x = \sum_{k=1}^{p} \|\alpha_{J_k}\|_2 \underbrace{\sum_{i \in J_k} \frac{\alpha_i}{\|\alpha_{J_k}\|_2} e_{[j_{i-1}+1,j_i]}}_{\in K_0^s} .$$

The necessary factor $c_{n,s}$ can be found by bounding the size of $\|\alpha_{J_k}\|_2$, i.e.

$$\max(\|\alpha_{J_k}\|_2) \le \sum_{k=1}^{p} \|\alpha_{J_k}\|_2 \overset{C-S}{\le} \underbrace{\|\alpha\|_2}_{\le 1} \sqrt{p} \le \sqrt{\frac{n}{s}}.$$

From this, we see that $K_{TV}^{2\sqrt{s}} \subset \sqrt{\frac{n}{s}}\mathrm{ch}(K_0^s)$. To see that this embedding constant is optimal, we construct a vector, for which it is needed.

To simplify the discussion, suppose that n and s are even and $s|n$. For even n, the vector $x_1 = (\sqrt{\frac{1-(-1)^k\varepsilon}{n}})_k$ has unity norm, lies in $K_{TV}^{2\sqrt{s}}$ for $\varepsilon < \frac{2\sqrt{s}}{n}$ and has jump support on all of $[n]$!

Splitting $\mathrm{Jsupp}(x_1)$ into sets $J_1, \dots, J_{n/s}$. Setting $a_k = \sqrt{\frac{n}{s}}x_1|_{J_k} \in K_0^s$, we see that $x_1 = \sum_{k=1}^{n/s} \sqrt{\frac{s}{n}}a_k$ and in order for this to be elements of $c_{n,s}\mathrm{ch}(K_0^s)$, we have to set $c_{n,s} = \sqrt{\frac{n}{s}}$. This

follows from

$$x_1 = \sum_{k=1}^{n/s} x_1|_{J_k} = \sum_{k=1}^{n/s} \sqrt{\frac{s}{n}}\frac{p}{p} a_k = \sum_{k=1}^{n/s} \frac{1}{p} \underbrace{\left(\sqrt{\frac{n}{s}} a_k\right)}_{\in \sqrt{\frac{n}{s}} K_0^s} \in \sqrt{\frac{n}{s}} \mathrm{ch}(K_0^s)$$

and no smaller inflation factor than $\sqrt{\frac{n}{s}}$ can suffice.

So from the previous discussion, we get

Lemma 3.1. We have the series of inclusions

$$\mathrm{ch}(K_0^s) \subset K_{TV}^{2\sqrt{s}} \subset \sqrt{\frac{n}{s}}\mathrm{ch}(K_0^s).$$

In view of the results obtainable for sparse vectors and the ℓ_1-ball, this is very disappointing, because Lemma 3.1 now implies that the width of $K_{TV}^{2\sqrt{s}}$ satisfies

$$w(K_{TV}^{2\sqrt{s}}) \leq w(\sqrt{\frac{n}{s}}\mathrm{ch}(K_0^s)) = \sqrt{\frac{n}{s}}w(K_0^s) \leq 10\sqrt{n\log(\mathrm{e}(n-1)/s)},$$

which is highly suboptimal.

Luckily, the results in [CX15] suggest, that the factor n in the previous equation can be replaced by $\sqrt{sn}$. However, they have to resort to a direct calculation of the Gaussian width of $K_{TV}^{2\sqrt{s}}$. The intuition why the Gaussian mean width can be significantly smaller than the bound given in Lemma 3.1 stems from the fact, that in order to obtain an inclusion we need to capure all 'outliers' of the set - no matter how small their measure is.

3.4 Exact recovery

For exact recovery, the M^*-bound is not suitable anymore and, as suggested in [Ver15], we will use 'Gordon's escape through the mesh' in order to find conditions on exact recovery. Exact recovery for TV minimization via this approach has first been considered in [CX15].

Suppose, we want to recover $x \in K_0^s$ from Gaussian measurements $Ax = y$. Given, that we want our estimator $\hat{x}$ to lie in a set K, exact recovery is achieved, if $K \cap \{z \colon Az = y\} = \{x\}$. This is equivalent to requiring

$$(K - x) \cap \underbrace{\{z - x \colon Az = y\}}_{=\ker(A)} = \{0\}.$$

With the descent cone $D(K, x) = \{t(z - x) \colon t \geq 0, z \in K\}$, we can rewrite this condition as

$$D(K, x) \cap \ker(A) = \{0\},$$

by introducing the set $S(K, x) = D(K, x) \cap B_2^n$, we see that if

$$S(K, x) \cap \ker(A) = \emptyset,$$

we get exact recovery. The question, when a section of a subset of the sphere with a random hyperplane is empty is answered by Gordon's escape through a mesh.

15

Theorem 3.2 ([Gor88])**.** Let $S \subset \mathbb{S}^{n-1}$ be fixed and $E \in G_{n,n-m}$ be drawn at random according to the Haar measure. Assume that $\hat{w}(S) = \mathbb{E} \sup_{u \in S} \langle g, u \rangle < \sqrt{m}$, then $S \cap E = \emptyset$ with probability exceeding

$$1 - 2.5 \exp\left(-\frac{(m/\sqrt{m+1} - \hat{w}(S))^2}{18}\right).$$

So we get exact recovery with high probability from a program given in Theorem 3.1 or 3.2, provided that $m > \hat{w}(S(K, x_0))^2$.

Let's see how this applies to TV minimization. Suppose, we are given $x \in K_0^s$ and Gaussian measurements $Ax = y$. Solving

$$\min \|z\|_{TV} \text{ with } Az = y,$$

amounts to using the Minkowski functional of the set $K = \{z \in \mathbb{R}^n : \|z\|_{TV} \le \|x\|_{TV}\}$, which is a scaled TV-Ball.

In [CX15], the null space property for TV minimization given in Corollary 2.1 has been used in order to obtain recovery guarantees.

They consider the set, where this condition is not met, i.e.

$$\mathcal{S} := \{x' \in B_2^n : \exists J \subset [n], |J| \le s, \|(\nabla x')_J\|_1 \ge \|(\nabla x')_{J^c}\|_1\},$$

and apply Gordon's escape through the mesh to see that with high probability, its intersection with the kernel of A is empty, thus proving exact recovery with high probability. Their estimate to the mean width of the set $\mathcal{S}$,

$$\hat{w}(\mathcal{S}) \le c \sqrt[4]{ns} \sqrt{\log(2n)}$$

with $c < 19$ is essentially optimal (up to logarithmic factors), as they also show that $w(\mathcal{S}) \ge C\sqrt[4]{ns}$. So uniform exact recovery can only be expected for $m = \mathcal{O}(\sqrt{sn} \log n)$ measurements.

Let us examine some connections to the previous discussion abot the descent cone.

Lemma 3.2. We have that for $K = \{z \in \mathbb{R}^n : \|z\|_{TV} \le \|x\|_{TV}\}$ defined as above and $x \in K_0^s$, it holds that $S(K, x) \subset \mathcal{S}$.

Proof. Let $y \in S(K, x)$. Then there exists a $x \ne z \in K$, such that $y = \frac{z-x}{\|z-x\|_2}$. Set $J = \text{Jsupp}(x)$, then, as $z \in K$, we have that $\|z\|_{TV} \le \|x\|_{TV}$, or

$$\sum_{i \in J} |(\nabla x)_i| \ge \sum_{i \in J} |(\nabla z)_i| + \sum_{i \notin J} |(\nabla z)_i|$$

Now, by the triangle inequality and this observation, we have

$$\sum_{i \in J} |(\nabla x)_i - (\nabla z)_i| \ge \sum_{i \in J} |(\nabla x)_i| - |(\nabla z)_i| \ge \sum_{i \notin J} |(\nabla z)_i| = \sum_{i \notin J} |(\nabla x)_i - (\nabla z)_i|.$$

The last equality follows from the fact that ∇x is zero outside of the gradient support of x. Multiplying both sides with $\frac{1}{\|z-x\|_2}$ gives the desired result

$$\|(\nabla y)_J\|_1 = \frac{1}{\|z-x\|_2} \sum_{i \in J} |(\nabla x)_i - (\nabla z)_i| \ge \ge \frac{1}{\|z-x\|_2} \sum_{i \notin J} |(\nabla x)_i - (\nabla z)_i| = \|(\nabla y)_{J^c}\|_1.$$

$\square$

The previous lemma shows that the recovery guarantees derived from the null space property and via the descent cone are actually connected in a very simple way. Clearly, now if we do not intersect the set $\mathcal{S}$, we also do not intersect the set $S(K, x)$, which yields exact recovery for example with the same upper bounds on m as for $\mathcal{S}$. Even more specifically, in the calculation of $\hat{w}(\mathcal{S})$ given in [CX15], an embedding into a slightly larger set $\tilde{\mathcal{S}} = \{x \in B_2^n : \|x\|_{TV} \leq 4\sqrt{s}\}$ is made. This embedding can also quite easily be done if we note that $\|x\|_{TV} \leq 2\sqrt{s}$, as we showed above and $\|z\|_{TV} \leq \|x\|_{TV}$.

3.5   Subgaussian measurements

Up to this point, all our measurement matrices have been assumed to consist of i.i.d. Gaussian random variables. We will reduce this requirement in this section to be able to incorporate also subgaussian measurement matrices into our framework. In order to do this, we rely on results given by Tropp in [Tro15] using the results of Mendelson [KM15, Men14]. We will consider problems of the form

$$\min \|z\|_{TV} \text{ such that } \|Az - y\| \leq \varepsilon, \tag{13}$$

where A is supposed to be a matrix with independent subgaussian rows. Furthermore, we denote the exact solution by x_0, i.e. $Ax_0 = y$. We pose the following assumptions on the distribution of the rows of A.

(M1) $\mathbb{E}A_i = 0$,

(M2) There exists $\alpha > 0$, such that for all $u \in \mathbb{S}^{n-1}$ it holds that $\mathbb{E}|\langle A_i, u\rangle| \geq \alpha$,

(M3) There is a $\sigma > 0$, such that for all $u \in \mathbb{S}^{n-1}$ it holds that $\mathbb{P}(|\langle A_i, u\rangle| \geq t) \leq 2\exp(-t^2/(2\sigma^2))$,

(M4) The constant $\rho := \frac{\sigma}{\alpha}$ is small.

Then the small ball methods yields the following recovery guarantee (we present the version of [Tro15]).

Theorem 3.3. Let $x_0 \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m \times n}$ be a subgaussian matrix satisfying (M1)-(M4) above. Furthermore, set $y = Ax_0 + e$, where $\|e\| \leq \varepsilon$ denotes the (bounded) error of the measurement. Then the solution $\hat{x}$ of (13) satisfies

$$\|\hat{x} - x_0\| \leq \frac{2\varepsilon}{\max\{c\alpha\rho^{-2}\sqrt{m} - C\sigma w(S(K, x_0)) - \alpha t, 0\}}$$

with probability exceeding $1 - e^{-ct^2}$. $D(K, x_0)$ denotes the descent cone of the set K at x_0, as defined in the previous section.

From this we see that, provided

$$m \geq \tilde{C}\rho^6 w^2(S(K, x_0)),$$

we obtain stable reconstruction of our original vector from (13).

In the previous section, we have shown the inclusion $S(K, x_0) \subset \mathcal{S}$ for $x_0 \in K_s^0$ and hence we have that

$$w(S(K, x_0) \leq w(\mathcal{S}) \leq c\sqrt[4]{ns}\sqrt{\log(2n)}.$$

17

So we see that for $m \geq \tilde{C}\rho^6 \sqrt{ns}\log(2n)$, we obtain the bound

$$\|\hat{x} - x_0\| \leq \frac{2\varepsilon}{\max\{c\alpha\rho^{-2}\sqrt{\tilde{C}}\rho^3 \sqrt[4]{ns}\sqrt{\log(2n)} - C\sigma \sqrt[4]{ns}\sqrt{\log(2n)} - \alpha t, 0\}}$$

$$= \frac{2\varepsilon}{\max\{\sigma(c\sqrt{\tilde{C}} - C)\sqrt[4]{ns}\sqrt{\log(2n)} - \alpha t, 0\}}$$

with high probability. We conclude that, given the absolute constants c, C, we need to set $\tilde{C} \geq \frac{C^2}{c^2}$ in order to obtain a meaningful result. Combining all our previous discussions with Theorem 3.3, we get

Theorem 3.4. Let $x_0 \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m \times n}$ be a subgaussian matrix satisfying (M1)-(M4). Furthermore, set $y = Ax_0 + e$, where $\|e\| \leq \varepsilon$ denotes the (bounded) error of the measurement, constants $c, C, \tilde{C} > 0$ as above and $t \leq \frac{\sigma(c\sqrt{\tilde{C}}-C)\sqrt[4]{ns}\sqrt{\log(2n)}}{\alpha}$. Then the solution $\hat{x}$ of (13) satisfies

$$\mathbb{P}\left(\|\hat{x} - x_0\| > \frac{2\varepsilon}{\sigma(c\sqrt{\tilde{C}} - C)\sqrt[4]{ns}\sqrt{\log(2n)} - \alpha t}\right) \leq \mathrm{e}^{-ct^2}.$$

We can for example set $t = \rho(c\sqrt{\tilde{C}} - C)\sqrt[4]{ns}$ (for $n \geq 2$) to obtain the bound

$$\mathbb{P}\left(\|\hat{x} - x_0\| > \frac{2\varepsilon}{\sigma(c\sqrt{\tilde{C}} - C)\sqrt[4]{ns}(\sqrt{\log(2n)} - 1)}\right) \leq \mathrm{e}^{-\tilde{c}\rho\sqrt{ns}}.$$

For example for i.i.d. standard Gaussian measurements, the constant $\rho = \sqrt{\frac{2}{\pi}}$.

Note that in the case of noisefree measurements $\varepsilon = 0$, Theorem 3.4 gives an exact recovery result for a wider class of measurement ensembles with high probability. Furthermore with a detailed computation of $w(S(K, x_0))$ one might be able to improve the number of measurements for nonuniform recovery.

4 Discussion and open problems

As the considerations in the previous sections illustrate, the mathematical properties of total variation minimization differ significantly from algorithms based on synthesis sparsity, especially in one dimension. For this reason, there are a number of questions that have been answered for synthesis sparsity, but which are still open for the framework of total variation minimization. For example, the analysis provided in [RRT12, KMR14] for deterministically subsampled partial random circulant matrices, as they are used to model measurement setups appearing in remote sensing or coded aperture imaging, could not be generalized to total variation minimization. The difficulty in this setup is that the randomness is encoded by the convolution filter, so it is not clear what the analogy of variable density sampling would be.

Another case of practical interest is that of sparse 0/1 measurement matrices. Recently it has been suggested that such meausurements increase efficiency in photoacoustic tomography, while at the same time, the signals to be recovered (after a suitable temporal transform) are approximately

18

gradient sparse. This suggests the use of total variation minimization for recovery, and indeed empirically, this approaches yields good recovery results [SKB$^+$15]. Theoretical guarantees, however, (as they are known for synthesis sparse signals via an expander graph construction [BGI$^+$08]) are not available to date for this setup.

Acknowledgements

References

[BDDW08] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. Wakin. A simple proof of the Restricted Isometry Property for random matrices. *Constr. Approx.*, 28(3):253–263, 2008.

[BGI$^+$08] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 798–805. IEEE, 2008.

[CDDD03] A. Cohen, W. Dahmen, I. Daubechies, and R DeVore. Harmonic analysis of the space bv. *Revista Matematica Iberoamericana*, 19(1):235–263, 2003.

[CENR10] E. J. Candès, Y. C. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Appl. Comput. Harmon. Anal.*, 31(1):59–73, 2010.

[CRT06] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(3):489–509, 2006.

[CX15] J.-F. Cai and W. Xu. Guarantees of total variation minimization for signal recovery. *Information and Inference*, 4(4):328–353, 2015.

[DH01] D. L. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7):2845–2862, 2001.

[Don04] D. Donoho. High-dimensional centrally-symmetric polytopes with neighborliness proportional to dimension. Technical report, Department of Statistics, Stanford University, 2004.

[Don06] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[FR13]   S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.

[Fuc04]   J. J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Transactions on Information Theory,*, 50(6), 2004.

[Gor88]   Y. Gordon. *On Milman's inequality and random subspaces which escape through a mesh in $\mathbb{R}^n$*. Springer, 1988.

[JKL15]   J. Jørgensen, C. Kruschel, and D. Lorenz. Testable uniqueness conditions for empirical assessment of undersampling levels in total variation-regularized x-ray ct. *Inverse Problems in Science and Engineering*, 23(8):1283–1305, 2015.

[KM15]   V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *International Mathematics Research Notices*, 2015(23):12991–13008, 2015.

[KMR14]   F. Krahmer, S. Mendelson, and H. Rauhut. Suprema of chaos processes and the restricted isometry property. *Comm. Pure Appl. Math.*, 67(11):1877–1904, 2014.

[KNW15]   F. Krahmer, D. Needell, and R Ward. Compressive sensing with redundant dictionaries and structured measurements. *SIAM J. Math. Anal.*, 47(6):4606–4629, 2015.

[KR14]   F. Krahmer and H. Rauhut. Structured random measurements in signal processing. *GAMM-Mitteilungen*, 37(2):217–238, 2014.

[KR15]   M. Kabanava and H. Rauhut. Analysis ℓ_1-recovery with frames and gaussian measurements. *Acta Applicandae Mathematicae*, 140(1):173–195, 2015.

[Kru15]   C. Kruschel. *Geometrical Interpretations and Algorithmic Verification of Exact Solutions in Compressed Sensing*. PhD thesis, TU Braunschweig, 2015.

[KRZ15]   M. Kabanava, H. Rauhut, and H. Zhang. Robust analysis ℓ_1-recovery from gaussian measurements and total variation minimization. *European J. Appl. Math.*, 26(06):917–929, 2015.

[KW11]   F. Krahmer and R. Ward. New and improved johnson-lindenstrauss embeddings via the restricted isometry property. *SIAM J. Math. Anal.*, 43(3):1269–1281, 2011.

[KW14]   F. Krahmer and R. Ward. Stable and robust sampling strategies for compressive imaging. *IEEE Trans. Image Proc.*, 23(2):612–622, 2014.

[Men14]   S. Mendelson. Learning without concentration. In *COLT*, pages 25–39, 2014.

[MPTJ07]   S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis*, 17(4):1248–1282, 2007.

[NDEG13]   S. Nam, M. Davies, M. Elad, and R. Gribonval. The cosparse analysis model and algorithms. *Appl. Comp. Harmon. Anal.*, 34(1):30–56, 2013.

[NW13a]   D. Needell and R. Ward. Near-optimal compressed sensing guarantees for total variation minimization. *IEEE Transactions on Image Processing*, 22(10):3941–3949, 2013.

[NW13b]   D. Needell and R. Ward. Stable image reconstruction using total variation minimization. *SIAM J. Imaging Sci.*, 6(2):1035–1058, 2013.

[PMG+12]   G. Puy, J. Marques, R. Gruetter, J.-P. Thiran, D. Van De Ville, P. Vandergheynst, and Y. Wiaux. Spread spectrum magnetic resonance imaging. *IEEE Trans. Medical Imaging*, 31(3):586–598, 2012.

[Poo15]   C. Poon. On the role of total variation in compressed sensing. *SIAM J. Imag. Sci*, 8(1):682–720, 2015.

[PV13]   Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *Information Theory, IEEE Transactions on*, 59(1):482–494, 2013.

[PVW11]   G. Puy, P. Vandergheynst, and Y. Wiaux. On variable density compressive sampling. *Signal Processing Letters*, 18:595–598, 2011.

[Roc72]   R.T. Rockafellar. *Convex Analysis*. Princeton University Press, 1972.

[ROF92]   Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.

[RRT12]   H. Rauhut, J. Romberg, and J.A. Tropp. Restricted isometries for partial random circulant matrices. *Appl. Comp. Harmon. Anal.*, 32(2):242–254, 2012.

[RSV08]   H. Rauhut, K. Schnass, and P. Vandergheynst. Compressed sensing and redundant dictionaries. *IEEE Trans. Inform. Theory*, 54(5):2210–2219, 2008.

[SKB+15]   M. Sandbichler, F. Krahmer, T. Berer, P. Burgholzer, and M. Haltmeier. A novel compressed sensing scheme for photoacoustic tomography. *SIAM J. Appl. Math.*, 75(6):2475–2494, 2015.

[TP14]   A. Tillmann and M. Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *IEEE Transactions on Information Theory*, 60(2):1248–1259, 2014.

[Tro15]   J. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*, pages 67–101. Springer, 2015.

[Ver15]   R. Vershynin. Estimation in high dimensions: a geometric perspective. In *Sampling theory, a renaissance*, pages 3–66. Springer, 2015.

[ZMY16]   H. Zhang, Y. Ming, and W. Yin. One condition for solution uniqueness and robustness of both ℓ_1-synthesis and ℓ_1-analysis minimizations. *Advances in Computational Mathematics*, 42(6):1381–1399, 2016.