On the Effective Measure of Dimension in Total Variation Minimization

Raja Giryes
Department of Electrical and Computer Engineering
Duke University
Durham, NC, 27708 USA
Email: raja.giryes@duke.edu

Yaniv Plan
Department of Mathematics
University of British Columbia
Vancouver, BC V6T 1Z2, Canada
Email: yaniv@math.ubc.ca

Roman Vershynin
Department of Mathematics
University of Michigan
Ann Arbor, MI 48109, U.S.A
Email: romanv@umich.edu

***Abstract*—Total variation (TV) is a widely used technique in many signal and image processing applications. One of the famous TV based algorithms is TV denoising that performs well with piecewise constant images. The same prior has been used also in the context of compressed sensing for recovering a signal from a small number of measurements. Recently, it has been shown that the number of measurements needed for such a recovery is proportional to the size of the edges in the sampled image and not the number of connected components in the image. In this work we show that this is not a coincidence and that the number of connected components in a piecewise constant image cannot serve alone as a measure for the complexity of the image. Our result is not limited only to images but holds also for higher dimensional signals. We believe that the results in this work provide a better insight into the TV prior.**

I. Introduction

Consider the problem of recovering an unknown signal $\boldsymbol{x} \in \mathbb{R}^d$ from a set of linear measurements

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{z}, \tag{I.1}$$

where $\boldsymbol{A} \in \mathbb{R}^{m\times d}$ is the measurements matrix and $\boldsymbol{z}$ is an additive white Gaussian noise with variance σ^2. It is clear that when $m < d$ the problem is underdetermined and it is impossible to stably recover $\boldsymbol{x}$ from $\boldsymbol{y}$ without having any prior knowledge on $\boldsymbol{x}$.

In this paper we consider the case of piecewise constant b-dimensional signals, where $b \geq 2$[1]. A popular prior for this type of signals is the total variation one that penalizes the magnitudes of the gradients in $\boldsymbol{x}$ [1], [2]. For recovering $\boldsymbol{x}$ from $\boldsymbol{y}$, it reads as

$$\min_{\tilde{\boldsymbol{x}}} \left\| |\nabla \tilde{\boldsymbol{x}}| \right\|_1 \quad s.t. \quad \|\boldsymbol{y} - \boldsymbol{A}\tilde{\boldsymbol{x}}\|_2^2 \leq \epsilon, \tag{I.2}$$

and in its unconstrained form of as

$$\min_{\tilde{\boldsymbol{x}}} \|\boldsymbol{y} - \boldsymbol{A}\tilde{\boldsymbol{x}}\|_2^2 + \gamma \left\| |\nabla \tilde{\boldsymbol{x}}| \right\|_1, \tag{I.3}$$

where ϵ is a proxy for the energy of the noise $\boldsymbol{z}$, γ is a regularization parameter dependent on $\boldsymbol{z}$, and $|\nabla \boldsymbol{x}|$ is the vector of magnitudes of the gradients of $\boldsymbol{x}$. For example, in the two dimensional case ($b = 2$)

$$|\nabla \boldsymbol{x}| = \sum_{i,j} \sqrt{\left(\frac{\partial \boldsymbol{x}(i,j)}{\partial i}\right)^2 + \left(\frac{\partial \boldsymbol{x}(i,j)}{\partial j}\right)^2}, \tag{I.4}$$

where we abuse notation and refer by $\boldsymbol{x}(i,j)$ to the location (i,j) in the image $\boldsymbol{x}$.

The TV formulation is very popular for image denoising (the case $\boldsymbol{A} = \boldsymbol{I}$) [1], [2] and has been proven to be edge preserving [3]. An extension of TV, which uses overparameterization, has been applied with sparsity-based techniques for curve fitting, and image denoising and segmentation [4]. It has been also successfully applied in the context of compressed sensing, for recovering a signal from a small number of linear Gaussian measurements [5], [6], [7], [8], [9], [10]. In [11], this technique has been utilized for reducing the number of needed MRI measurements.

The work in [7], [8] analyzed theoretically the performance of TV. Given the number of edges, k, in $\boldsymbol{x}$, it showed that for a certain $\boldsymbol{A}$ it is sufficient to use only $m = O(bk\log(d))$ random measurements[2] to recover $\boldsymbol{x}$ stably from $\boldsymbol{y}$, implying a perfect recovery in the noiseless case $\boldsymbol{z} = 0$.

[1]Throughout the paper we assume that $d = n^b$ for a certain integer n and $\boldsymbol{x}$ is a column stacked representation of the b-dimensional signal

[2]Note that (i) the dependence on b is likely to be only due to a flaw in the proof technique as for the anisotropic version of TV it is sufficient to have $m = O(k\log(d))$ [8]; and (ii) the conditions in [8] are in terms of the sparsity of the gradient magnitude, k', and not the number of edges k, i.e., $m = O(bk'\log(d))$. However, both measures are related to each other via the inequalities $\frac{1}{b}k \leq k' \leq k$.

This result resembles the compressed sensing guarantees with the standard sparsity prior that requires $m = O(s\log(d))$ [5], [12], [13]. However, while in the latter requirement (which was shown to be optimal [14]) s is the intrinsic dimension of the signal, this is not the case in the bound for TV that relies on the number of edges. Each group of edges creates a connected component in $\boldsymbol{x}$ and the number of these components defines the manifold dimension of this signal. It is clear that unless we are in the one-dimensional case ($b = 1$), which has been analyzed in [9], the number of edges in the signal does not set the number of connected components. The length of the boundaries within an image might be much larger than the number of connected components within it. Notice that this might happen even in the case of a signal with a manifold dimension equals to two as demonstrated in Fig. 1a.

A. Our Contribution

Therefore, it is natural to ask whether this gap is an artificial outcome of the proof technique or that there is a more fundamental barrier that prevents sampling in the manifold dimension. In [15], we have examined this question for the case of $b = 2$ in the context of the analysis cosparse model, treating also other types of manifolds, showing that the manifold dimension is not a sufficient proxy for the needed number of measurements in this model. In this work we focus on the case of TV and extend our results to $b \geq 2$.

We show that both in images and higher dimensional signals, it is not possible to have a stable recovery from a number of measurements proportional to the number of connected components in the signals. Our main result is stated in the following theorem.

Theorem I.1. *Let $K \subset \mathbb{R}^d$ be the set of all b-dimensional signals with only two connected components and an ambient dimension $d = n^b$. Suppose $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{z}$ for some $\boldsymbol{x} \in K$, $\boldsymbol{A} \in \mathbb{R}^{m\times d}$ with $\|\boldsymbol{A}\| \leq 1$, and $\boldsymbol{z} \sim \mathcal{N}(0, \sigma^2 \cdot \mathrm{Id})$. Then for any estimator $\hat{\boldsymbol{x}}(\boldsymbol{y})$ we have*

$$\max_{\boldsymbol{x}\in K_2} \mathbb{E}\, \|\hat{\boldsymbol{x}} - \boldsymbol{x}\|_2 \geq C\sigma \exp(cd/m).$$

The theorem shows that though we may have a signal that has only two connected components in it, i.e., its manifold dimension is two, we will still need $O(d)$ measurements to recover it without incurring a huge error. Therefore, we may conclude that the fact that the needed number of measurements for recovery is of the order of the borders between the connected componenets is not a result of a flaw in the existing reconstruction guarantees or due to a bad selection of the measurement matrix $\boldsymbol{A}$.

The theorem is proven in Section II by providing packing for the set K, and combining it with a hypothesis testing argument. Then we perform several experiments in Section III to demonstrate this fact and show that indeed, the size of the borders in a signal is a better measure for its compressibility than its manifold dimension. We believe that our results may help to understand better the types of images and signals with which it is better to use the TV prior.

II. Proof of Theorem I.1

In our proof we take a similar path to the one we used in [15] for the two-dimensional case. The core of our proof is a construction of a random packing for K. With that packing we utilize several Lemmas proved in [15] to finish our proof.

Recall that a packing $\mathcal{X} \subset K$ with ℓ_2-balls of radius δ is a set satisfying $\|x - y\|_2 \geq \delta$ for any $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$ with $\boldsymbol{x} \neq \boldsymbol{y}$. We denote by $P(K, \delta)$ the maximal cardinality of such a set, i.e., the *packing number*.

A. Packing Construction

Lemma II.1. *Suppose that $n = d^{1/b} \geq 8$ is an integer. Then*

$$P(K, 1/2) \geq \exp(d/16).$$

Proof. Note that each element $\boldsymbol{x} \in \mathbb{R}^{n^b}$ in the set K can be viewed as a b-dimensional grid composed of two connected components, $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, with $\boldsymbol{x}$ constant on each component (See Figure 1a for the two-dimensional case). Our question reduces to constructing a packing for pairs of blobs in the b-dimensional cube. Note that each entry of our blobs has a magnitude $1/n$.

Our goal is to give a lower bound on the packing number $P(K, 1/2)$. By abuse of notation we treat $\boldsymbol{x}$ as a grid and refer by $\boldsymbol{x}(i_1, i_2, \dots, i_b)$ to the location $(i_1, i_2, \dots, i_b)$ on this grid, where $0 \leq i_l < n, \forall 1 \leq l \leq b$. We restrict ourselves only to blobs of the form

$$\boldsymbol{x}(i_1,\dots,i_b) = \begin{cases} \frac{1}{n} & \begin{matrix} i_1 = 0 \textit{ or} \\ (i_1 < n \textit{ and } i_2 \bmod 3 = 1), \end{matrix} \\ -\frac{1}{n} & \begin{matrix} i_1 = n-1 \textit{ or} \\ (i_1 \neq 0 \textit{ and } i_2 \bmod 3 = 0), \end{matrix} \\ \zeta & \text{else}, \end{cases}$$

where $\zeta = \pm\frac{1}{n}$ w.p. 0.5 and mod is the modulo operation. This type of blob is illustrated in Fig. 1b for the two-dimensional case. Note that our problem turns to be constructing a packing for the Hamming cube

$$\tilde{K} := \frac{1}{n}\{+1, -1\}^q,$$

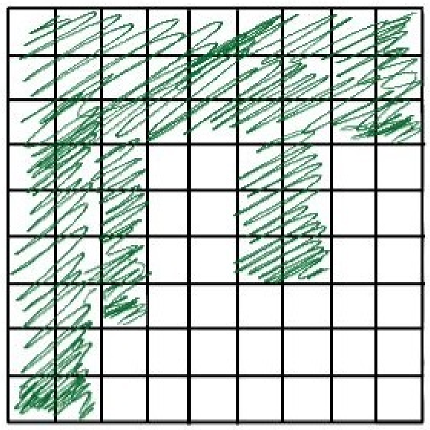

(a) A point in K_2. All green squares have one value and all white squares have another.

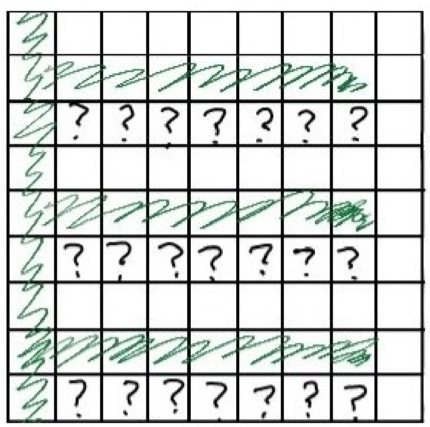

(b) Visualization of packing patterns. Green corresponds with $1/n$, white corresponds with $-1/n$, and ? can be either.

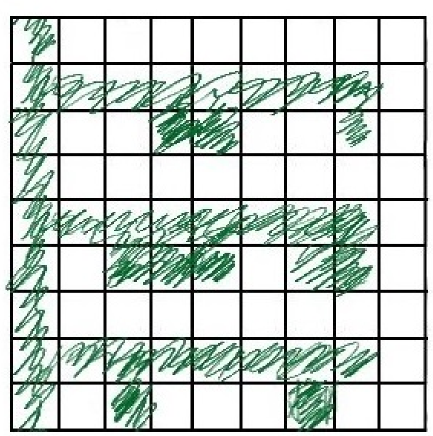

(c) One possible point in the packing.

where $q = \frac{dn(n-2)}{3n^2}$ is the number of questions marks.

We use a random construction, noticing that $\|n \cdot \boldsymbol{x} - n \cdot \boldsymbol{x}'\|_2^2 \sim 4 \cdot \text{Binomial}(q, 1/2)$ for $\boldsymbol{x}$ and $\boldsymbol{x}'$ being two points picked uniformly at random. From Hoeffding's inequality we have

$$\begin{aligned} &\mathbb{P}\left\{\left\|\boldsymbol{x}_{t,l} - \boldsymbol{x}'_{t,l'}\right\|_2^2 < \frac{q}{d}\right\} = \\ &\mathbb{P}\left\{\left\|n \cdot \boldsymbol{x} - n \cdot \boldsymbol{x}'\right\|_2^2 < q\right\} \\ &\leq \exp(-q/8). \end{aligned} \tag{II.1}$$

To get a packing number for $\tilde{K}$ using the above inequality, we introduce the following Lemma (proved in [15])

Lemma II.2 (Random packing)**.** *Let F be a distribution supported on some set $K \subset \mathbb{R}^d$. Let $\boldsymbol{x}, \boldsymbol{x}'$ be independently chosen from F. Suppose that*

$$\mathbb{P}\left\{\|\boldsymbol{x} - \boldsymbol{x}'\|_2 < \delta\right\} \leq \eta$$

for some $\eta, \delta > 0$. Then,

$$P(K, \delta) \geq \eta^{-1/2}.$$

Since we assumed $n \geq 8$ then $q = dn(n-2)/3n^2 \geq d/4$. Therefore with this Lemma, we have $P(\tilde{K}_l, 1/2) \geq \exp(d/16)$. ■

B. Lower Bound by Packing Number

Having the packing number we prove our theorem by showing that subsampling causes the distance between some pair of points in the packing to reduce greatly and therefore these points will become indistinguishable due to the noise. This is done using the following proposition from [15]

Proposition II.3. *Let $K \subset \mathbb{R}^d$ be a cone. Let $\mathcal{X}$ be a δ-packing of $K \cap B^d$. Suppose $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{z}$ for $\boldsymbol{x} \in K$, $\boldsymbol{A} \in \mathbb{R}^{m \times d}$ with $\|\boldsymbol{A}\| \leq 1$, and $\boldsymbol{z} \sim \mathcal{N}(0, \sigma^2 \cdot \text{Id})$. Then for any estimator $\hat{\boldsymbol{x}} = \hat{\boldsymbol{x}}(\boldsymbol{y})$, we have*

$$\sup_{\boldsymbol{x} \in K} \mathbb{E}\,\|\hat{\boldsymbol{x}} - \boldsymbol{x}\|_2 \geq \frac{\delta\sigma\,|\mathcal{X}|^{1/m}}{32}.$$

By combining this proposition with the packing number we have from Lemma II.1 we get Theorem I.1..

III. Experiments

To demonstrate the results of the theorem we test the performance of TV in recovering signals with only two connected components but different number of edges. In all the experiments, the measurement matrix $\boldsymbol{A}$ is a random Gaussian matrix with normalized columns.

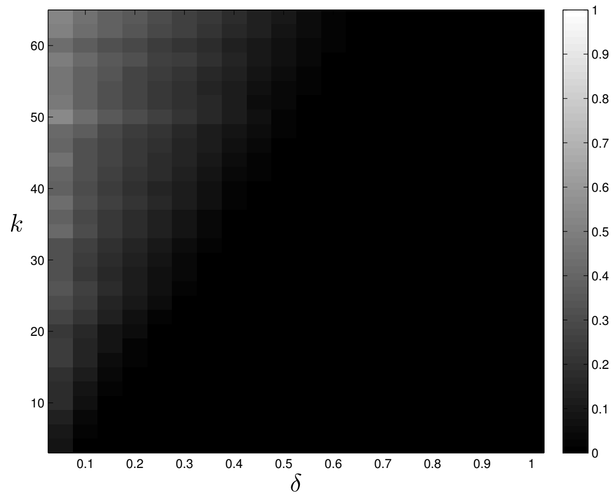

Fig. 2: Recovery error of TV in the noiseless case. Experiment setup: $d = 12^2$, $\delta = \frac{m}{d}$ and $\rho = \frac{p}{d}$. For each configuration we average over 200 realizations. Color attribute: Mean Squared Error.

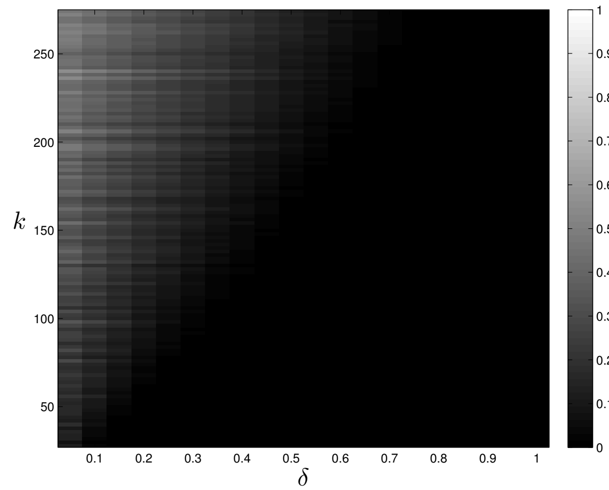

Fig. 4: Recovery error of TV in the noiseless case. Experiment setup: $d = 7^3$, $\delta = \frac{m}{d}$ and k is the edge size. For each configuration we average over 200 realizations. Color attribute: Mean Squared Error.

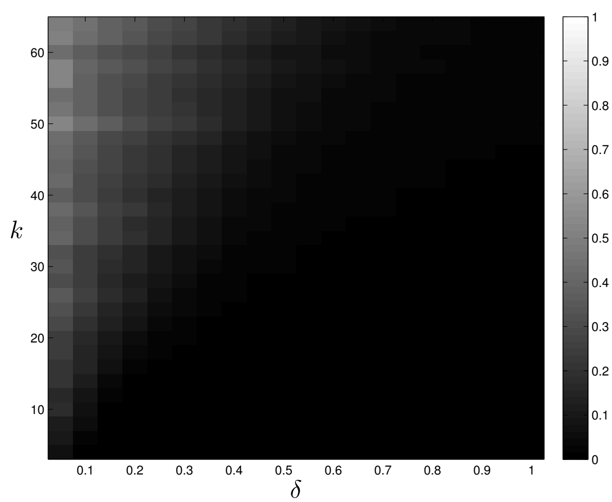

Fig. 3: Recovery error of TV in the noisy case with $\sigma = 0.01$. Experiment setup: $d = 12^2$, $\delta = \frac{m}{d}$ and k is the edge size. For each configuration we average over 200 realizations. Color attribute: Mean Squared Error.

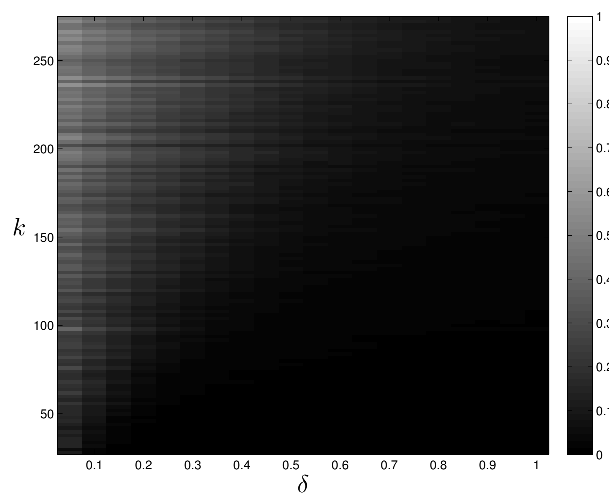

Fig. 5: Recovery error of TV in the noisy case with $\sigma = 0.01$. Experiment setup: $d = 7^3$, $\delta = \frac{m}{d}$ and k is the edge size. For each configuration we average over 200 realizations. Color attribute: Mean Squared Error.

In the first experiment, we consider images of size 12×12 with two connected components, each having a different gray value. We generate these images randomly using the same setup as in [15]. The values in the first and second connected components are selected randomly from the ranges $[0, 1]$ and $[-1, 0]$ respectively. The images are generated by setting all the pixels in the image to a value from the range $[0, 1]$ and then picking one pixel at random and starting a random walk (from its location) that assigns to all the pixels in its path a value from the range $[-1, 0]$ (the same value). The random walk stops once it gets to a pixel that it has visited before. With high probability the resulted image will be with only two connected components. We generate many images like that and sort them according to the length of the edges in them (eliminating images that have more than two connected components).

We test the reconstruction performance for different edges' length and sampling rates. The recovery error in the noiseless and noisy cases are presented in Figs. 4 and 5. It can be clearly seen that the estimation quality in both cases is determined by the size of the boundary within the image and not by the manifold dimension of the signal (that equals two), which is fixed in all the experiments. As predicted by Theorem I.1, if we rely only on the manifold dimension we get a very unstable recovery. However, if we take into account also the edge size, we can have a better prediction of our success rate. As we have seen in the previous experiment, also here we can see that instability in the noisy case also implies instability to relaxation.

We repeat the experiment for three dimensional grids of size $7 \times 7 \times 7$, also with only two connected components. These are generated in a very similar way to the images in the previous experiments. The only difference is that the path of the random walk is three dimensional instead of two.

Observe that also in this case the number of edges in the image is a better proxy for the reconstruction quality. However, it may happen that for longer edges the recovery would be easier than shorter edges. Therefore, we suggest that an exploration for a better measure for the complexity of piecewise constant high dimensional signals.

IV. Conclusion

In this work we have inquired whether it is possible to provide recovery guarantees for compressed sensing with total variation as a prior, using only information about the manifold dimension of the signal. Though the answer for this question is positive for standard compressed sensing (with the standard sparsity model) and the matrix completion problem [16], [17], we have shown that this is not true for images and higher dimensional signals. This has been demonstrated both in theory and simulations. Our conclusion is that the number of edges in the signal is a better proxy, than the signal's manifold dimension, for the needed number of measurements that provides a successful recovery using TV. However, we believe that there is still a need for an exploration of a better measure for the complexity of the image.

Acknowledgment

RG is partially supported by AFOSR. YP is partially supported by an NSF Postdoctoral Research Fellowship under award No. 1103909. RV is partially supported by NSF under grant DMS 1265782, by USAF under grant FA9550-14-1-0009, and by Humboldt Research Foundation.

References

[1] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

[2] A. Chambolle, "An algorithm for total variation minimization and applications," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 89–97, 2004.

[3] D. Strong and T. Chan, "Edge-preserving and scale-dependent properties of total variation regularization," *Inverse Problems*, vol. 19, no. 6, p. S165, 2003.

[4] R. Giryes, M. Elad, and A. Bruckstein, "Sparsity based methods for overparameterized variational problems," 2014, http://arxiv.org/abs/1405.4969.

[5] E. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

[6] S. Nam, M. Davies, M. Elad, and R. Gribonval, "The cosparse analysis model and algorithms," *Appl. Comput. Harmon. Anal.*, vol. 34, no. 1, pp. 30 – 56, 2013.

[7] D. Needell and R. Ward, "Stable image reconstruction using total variation minimization," *SIAM Journal on Imaging Sciences*, vol. 6, no. 2, pp. 1035–1058, 2013.

[8] ——, "Near-optimal compressed sensing guarantees for total variation minimization," *IEEE Transactions on Image Processing*, vol. 22, no. 10, pp. 3941–3949, 2013.

[9] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M. Davies, "Greedy-like algorithms for the cosparse analysis model," *Linear Algebra and its Applications*, vol. 441, no. 0, pp. 22 – 60, 2014, special Issue on Sparse Approximate Solution of Linear Systems.

[10] R. Giryes, "Sampling in the analysis transform domain," 2014, http://arxiv.org/abs/1410.6558.

[11] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.

[12] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203 – 4215, dec. 2005.

[13] M. Rudelson and R. Vershynin, "Sparse reconstruction by convex relaxation: Fourier and gaussian measurements," in *Information Sciences and Systems, 2006 40th Annual Conference on*, 22-24 2006, pp. 207 –212.

[14] E. Candès, "Modern statistical estimation via oracle inequalities," *Acta Numerica*, vol. 15, pp. 257–325, 2006.

[15] R. Giryes, Y. Plan, and R. Vershynin, "On the effective measure of dimension in analysis cosparse models," 2014, http://arxiv.org/abs/1410.0989.

[16] E. Candes and Y. Plan, "Matrix completion with noise," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 925–936, June 2010.

[17] Y. Eldar, D. Needell, and Y. Plan, "Uniqueness conditions for low-rank matrix recovery," *Applied and Computational Harmonic Analysis*, vol. 33, no. 2, pp. 309 – 314, 2012.