

NUMERICAL EVALUATION OF SAMPLING BOUNDS FOR NEAR-OPTIMAL RECONSTRUCTION IN COMPRESSED SENSING

Yoann Le Montagner^{1,2}, Marcio Marim^{1,2}, Elsa Angelini², Jean-Christophe Olivo-Marin¹

¹Institut Pasteur, Unité d'Analyse d'Images Quantitative CNRS URA 2582, F-75015 Paris

²Institut Télécom, Télécom ParisTech CNRS LTCl, F-75013 Paris

ABSTRACT

In this paper, we propose an empirical review of the conditions under which the *compressed sensing* framework allows to achieve exact image reconstruction. After a short presentation of the theoretical results related to this subject, we investigate the relevance and the limits of these theoretical results through several numerical reconstructions of some benchmark images. In particular, we discuss quantitative and qualitative artifacts that affect the reconstructed image when reducing the number of measurements in the Fourier domain. Finally, we conclude our study by extending our results to some real microscopic images.

Index Terms— Compressed sensing, image reconstruction, sampling rate, total variation minimization, Fourier transform.

1. THEORETICAL APPROACH OF COMPRESSED SENSING

1.1. Generalities on compressed sensing

The recent sampling theory of compressed sensing (CS) predicts that sparse signals and images can be reconstructed from what was previously believed to be incomplete information. CS was introduced by Candès *et. al* in [1] and Donoho in [2]. It relies on the fact that many types of signals or images can be well-approximated by a sparse decomposition on a suitable basis.

CS provides a formal reconstruction framework exploiting a compressed version of the original signal by taking only a small amount of linear and non-adaptive measurements. Now considering the noiseless incomplete measurements $y = \Phi x$ the recovery of a signal $x \in \mathbb{R}^N$ is achieved by solving the convex program:

$$\hat{x} = \arg \min_{x \in \mathbb{R}^N} \|\Psi x\|_{l_1} \text{ s.t. } \Phi x = y \quad (1)$$

where Φ is a M by N sampling matrix selecting only M coefficients of x along a sampling basis, with typically $M \ll N$, and Ψ is a basis transformation matrix chosen such that Ψx is a sparse representation of the original signal. In (1), the term $\|\Psi x\|_{l_1}$ is a convex surrogate to $\|\Psi x\|_{l_0}$: it aims at enforcing the sparsity of Ψx during the reconstruction process.

When CS is applied to image signals, it is quite common to use to following recovery program instead of (1):

$$\hat{x} = \arg \min_{x \in \mathbb{R}^N} \|x\|_{\text{TV}} \text{ s.t. } \Phi x = y \quad (2)$$

where $\|\cdot\|_{\text{TV}}$ stands for the total variation semi-norm. When proceeding this way, the underlying hypothesis that is assumed on the

original signal is that it has a sparse gradient¹.

In what follows, we will focus on the sparsity prior is enforced through TV regularization, and where the measurement basis is the Fourier basis. In this framework, the matrix Φ stands for a subsampling version of the Fourier transform.

1.2. Sampling bound and sparsity

Suppose that $\|x\|_{\text{TV}}$ is S -sparse or has S non-zero coefficients; in [3], the authors proposed to compute the required number of measurements in the following manner:

$$M \geq C \cdot \mu^2 \cdot S \cdot \log N \quad (3)$$

where C is a positive constant, and μ stands for the coherence between the sensing basis A and the representation basis $\|\cdot\|_{\text{TV}}$, as defined in [3]. Then the solution of (2) is exact with overwhelming probability and the success probability exceeds $1 - \delta$ if $M \geq C \cdot \mu^2 \cdot S \cdot \log N / \delta$.

Following (3), there exists a sampling bound such that the reconstruction is possible when sampling up to this value and fails when sampling down to the value. Theoretically, the transition is sharp but in practice the transition is smooth and there is an intermediate domain where reconstructions are not exact but do not fail completely.

However, for real signals the straightforward applicability of the measure (3) is limited. The reason is that the sparsity S has a restricted sense when dealing with real images which are not perfectly sparse, and possibly noisy.

Here we show that the sharp transition is perfectly verified for simulated piece-wise smooth images but only partially for real images with smooth edges and noise.

2. TOWARDS A NUMERIC EVALUATION OF AN OPTIMAL BOUND FOR THE SAMPLING RATE

2.1. Measuring the sampling threshold

Our first experiment consists in reconstructing elementary images composed of a single circular white object on a black background (see fig. 1), for which we carried out several CS reconstructions with various sampling rates. For each reconstruction, we sampled a few Fourier coefficients and computed the corresponding solution to the problem (2). The Fourier coefficients were selected with a uniform random scheme: the probability of selecting each coefficient followed an independent Bernoulli law whose parameter was tuned ac-

¹However, this is not the only way to deal with image signals: one could also use the usual recovery scheme (1), with for exemple a wavelet transform for Ψ .




Fig. 1. Examples of three CS reconstructions of the same one-disk image with radius $\rho = 22$, for three different sampling rates τ . On the very left: original image (top) and the logarithmic amplitude of its Fourier transform (bottom). Then, from left to right: sampling mask in the Fourier space (top) and the corresponding reconstructed images (bottom), with $\tau = 0.2\%$, $\tau = 0.7\%$, $\tau = 4\%$. The reconstruction obtained with 4% of the Fourier samples is identical to the original image.

cording to the targeted sampling rate². We studied the evolution of the reconstruction error between the CS reconstructed image \hat{x} and the original image x_0 as a function of the sampling rate τ ; the reconstruction error was defined as:

$$\text{RecErr} = \frac{\|\hat{x} - x_0\|_{l_2}}{\|\mu_0 - x_0\|_{l_2}} \quad (4)$$

where μ_0 is a constant signal, whose value is equal to the mean value of x_0 ³. Results are presented on fig. 2⁴.

For each curve, we observe three distinct domains:

- for small values of τ , the reconstruction error is constant at a high level: in this domain, the number of Fourier samples is too low to achieve a correct CS reconstruction, and the solution computed from (2) is roughly unstructured; this case corresponds to the first reconstruction presented in fig. 1;
- for high values of τ , the reconstruction error is also almost constant at a value close to zero: in this domain, the sampling rate is sufficient to perform an exact reconstruction of the original image from the subset of Fourier coefficients that are actually sampled⁵; this case corresponds to the third reconstruction presented in fig. 1;
- between these two constant domains, there is a narrow area around a transition sampling rate τ^* where the reconstruction error decreases from high value to almost zero; this case corresponds to the second reconstruction presented in fig. 1.

²Except for the central Fourier coefficient, which was always selected in order to ensure the uniqueness of the solution to the problem (2).

³With the normalization factor $\|\mu_0 - x_0\|_{l_2}$, we ensure $\text{RecErr}(\tau) \rightarrow 1$ when $\tau \rightarrow 0$, as we always sample the DC component

⁴The value of the reconstruction error obtained for each simulation depends on the sampling mask that is randomly generated at the beginning of each reconstruction. To get rid of this dependency, each simulation was run ten to twenty times, and median error values are reported. The variability encountered among these aggregated values will be discussed further.

⁵The fact that there still remains a small gap between the image obtained out of (2) and the original image might be due to the numerical algorithm that is used to solve this optimization problem. This algorithm (NESTA) is described in [4].




Fig. 2. Normalized quadratic reconstruction error between the original images and the CS reconstructions as a function of the sampling rate τ , for four one-disk bench images with different radius ρ .

The value τ^* of the sampling rate, as it somehow separates the domain where reconstruction is possible from the domain where it is not, is a measure of the sampling threshold that is defined in a theoretical manner in (3). There are several ways to define the actual value τ^* from the curve $\text{RecErr} = f(\tau)$: one could define τ^* such that $\text{RecErr}(\tau^*) = 0.5$ ⁶, or decide that τ^* is the point where the first derivative of the function takes its maximal absolute value⁷; however, as long as the transition domain is sufficiently narrow, all these definitions are likely to be equivalent. For the sake of simplicity, we have defined τ^* such that $\text{RecErr}(\tau^*) = 0.5$ in our simulations; this definition does not depend on the spread of the transition area.

The fact that this threshold is drastically modified depending on the input image reflects on the variation of the underlying sparsity coefficient S .

2.2. Optimal sampling rate and sparsity

When performing a CS reconstruction using the optimization problem (2), we know that the underlying *a priori* hypothesis that is made on the input image is that it has a sparse gradient. In the case of our simple binary images, the number of non-zero gradient coefficients is roughly equal to the perimeter of the object. Then, together with (3), we can make the assumption that the transitional sampling rate τ^* is an increasing linear function of the perimeter of the object. Therefore, in the case of our one-disk images, τ^* may increase linearly with ρ .

In order to check this hypothesis, we computed the transition sampling rate τ^* for eight values of the disk radius ρ . Results illustrated in fig. 3, confirmed that τ^* obeys a linear increasing with respect to ρ , hence somehow confirming empirically the theoretical relation (3).

2.3. Optimal sampling rate and shape factor

We have also investigate the dependency of the transitional sampling rate τ^* with respect to the shape factor of the imaged object. Equation (3) suggests that τ^* depends only on the number S of non-zero gradient coefficients, that is related to the perimeter of the object, but not to its shape factor. Therefore, two objects with the same perimeter should have the same transitional sampling rate, even if

⁶0.5 being the mean value between the two constant domains of the curve
⁷as the function seems to have an inflexion point in the transition area




Fig. 3. Transition sampling rate τ^* for eight values of ρ . Linear regression confirms that, for the single-disk images, τ^* obeys a linear increasing law with respect to ρ .




Fig. 4. Reconstruction error between the original images and the CS reconstructions as a function of the sampling rate τ , for the four images presented above (from left to right: ellipses with ratio between the half minor axis and the half major axis $\gamma = 1$, $\gamma = 0.4$, $\gamma = 0.15$, and $\gamma = 0.06$).

one is isotropic (for example: a disk) and one has a spatial dimension much larger than the others (for example: a flat ellipse).

To validate this hypothesis, we followed a similar approach than for the single-disk bench images, but we replaced the disks with ellipses of constant perimeter and various eccentricities. By varying the ratio γ between the half minor axis and the half major axis from 1 (circle) to almost 0 (flat shape), we tested shapes with different spatial and frequency characteristics; on the other hand, by setting a constant perimeter, we have maintained a constant TV-based sparsity measure for all the test images: then the observed τ^* should be the same for all the inputs. The results are presented on fig. 4.

Although the four curves do not perfectly overlap, the associated transitional sample rates are distributed in a narrow domain, approximately $[10^{-2}, 2 \times 10^{-2}]$. Moreover, the issue of the CS reconstruction process when operated with τ in the neighborhood of τ^* is highly dependent on the actual sampling mask that is randomly generated: this induces some uncertainty about the actual value of τ^* measured from these curves.




Fig. 5. Reconstruction error as a function of the sampling rate τ , for four one-disk test images with different radius ρ , and using a Gaussian sampling strategy.

2.4. Sampling strategy

The notion of transitional sampling rate is also related to the Fourier sampling strategy.

In all previous simulations, we used a random uniform sampling strategy, which means that the probability of selecting the Fourier coefficient corresponding to the spatial frequency ω is given by $p(\omega) = \tau$, with τ the targeted sampling rate. In this case, no spatial frequency is privileged as $p(\omega)$ does not actually depend on ω . However, it could be profitable to get more measurements within some sub-domains of the Fourier space, in order to take into account any *a priori* knowledge about the input image. For instance, in order to favor low frequencies, one can choose to allocate the random measurement according to a Gaussian sampling strategy:

$$p(\omega) = e^{-\frac{\|\omega\|_2^2}{R^2}} \quad (5)$$

where the parameter R is tuned according to the overall targeted sampling rate.

This kind of Fourier sampling strategy modifies the profile of the reconstruction error as a function of τ , as seen in fig. 5. For these simulations, we tried to reconstruct the single-disk test images for different sampling rates, using a Gaussian sampling strategy. As such strategy induces somehow a low-pass filtering effect on the reconstructed image, small objects become harder to retrieve than large ones⁸, as the first have a broader Fourier spectra than the second. On the contrary, when using a uniform sampling strategy, the small objects are retrieved first (see fig. 2).

Moreover, with the Gaussian sampling strategy, it becomes difficult to define what the critical sampling threshold τ^* is, as the evolution of the reconstruction error as a function of τ is far more complex than with a uniform sampling strategy. Indeed, for the single-disk images with $\rho = 3$ and $\rho = 8$, we can identify at least three domains where the reconstruction error is quite stable, and two transitional domains in between. Fig. 6 illustrates the CS reconstructed images typically obtained with sampling rates set in the three stable domains.

2.5. Realistic image reconstructions

So far, simulations were carried out on simple test images, whose structure is very poor compared to what is encountered in realis-

⁸meaning that a higher sampling rate is required to reconstruct the small-disk images




Fig. 6. Examples of three CS reconstructions of the same image (one disk with $\rho = 8$) using a Gaussian sampling strategy. From left to right: original image, reconstruction using $\tau = 0.05\%$, using $\tau = 1\%$, and using $\tau = 20\%$ of the Fourier coefficients. With $\tau = 1\%$, the reconstructed image is very close to the original one, but the sharp edges appear to be blurred due to the low-pass filtering effect of the Gaussian sampling strategy. With $\tau = 20\%$, the reconstructed image is identical to the original one.

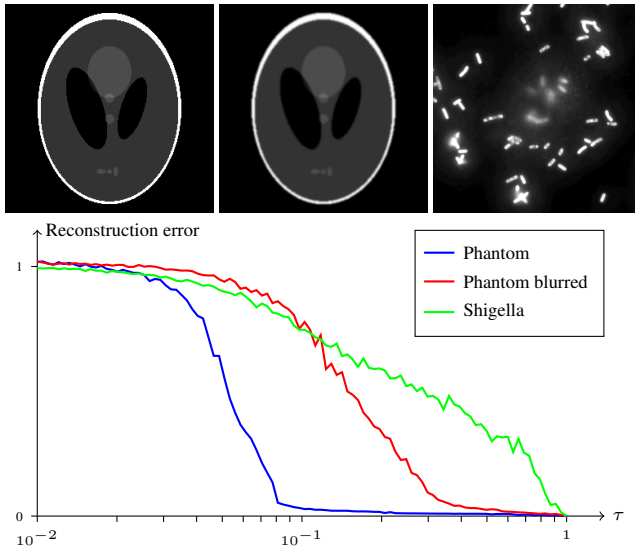


Fig. 7. Reconstruction error as a function of the sampling rate τ , for the three benchmark images presented above (from left to right: *Phantom*, *Phantom blurred*, *Shigella*).

tic ones. For real images, the notion of sparsity might be difficult to quantify, due either to the complexity of the textural content of the image, or to the degradation (blur, noise) induce by the acquisition process. Therefore, the practical definition of a sharp sampling threshold as defined in (3) might not be so obvious, even if we restrict ourselves to a uniform random sampling strategy, which does not add any extra filtering effect that would complicate the analysis.

As an empirical study, we measured the evolution of the reconstruction error as a function of τ using a uniform random sampling strategy for three test images (see fig. 7):

1. the well-known Shepp-Logan phantom image, with strictly piecewise constant regions and sharp edges;
2. the Shepp-Logan image degraded with a small Gaussian blur;
3. a true fluorescence microscopy image of *Shigella* bacteria.

For both Shepp-Logan images, the error curves look very similar to the one obtained for the disks and the ellipses in figs. 2 and 4. This means that, for these input signals, the notion of sampling threshold as defined in (3) holds, and therefore that an underlying sparsity level S can be defined. Going further, we can verify that blurring

the *Phantom* image has indeed induced an increase of its associated sampling threshold.

However, in the case of the *Shigella* image, the reconstruction error remains significantly high up until the sampling rate τ reaches 100%: for this image, it is not possible to define a critical sampling threshold τ^* . The underlying reason might be that the sparse model does not hold for this image, as it is not piecewise constant⁹; a more realistic assumption would be to consider this image as compressible in some basis Ψ , meaning that the error between the image and its approximation by at most S basis vectors of Ψ decreases proportionally to $S^{-\alpha}$ for some $\alpha > 0$ when S increases. It has been proved in [5] that CS can handle compressive signals through a relaxation of the optimization problem (2) into:

$$\hat{x} = \arg \min_{x \in \mathbb{R}^N} \|x\|_{\text{TV}} \text{ s.t. } \|\Phi x - y\|_2 \leq \epsilon \quad (6)$$

where ϵ is a positive parameter. Unfortunately, with this minimization scheme, we must give up the hope of achieving an exact reconstruction of the input image.

3. CONCLUSION

In this paper, we have presented an experimental exploration of the conditions under which the exact recovery theorems stated in [1, 3] hold, in a context of image reconstruction through TV minimization and Fourier undersampling. In the case of some simple test images, we studied the evolution of CS reconstructed images when modifying either the sampling pattern or the sampling rate of the Fourier coefficients. Finally, in the case of realistic images, we have discussed the practical relevance of such exact recovery theorems relating critical sampling rate to sparsity measures.

4. REFERENCES

- [1] Emmanuel Candès, Justin Romberg, and Terence Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, pp. 489–509, 2006.
- [2] David L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, 2006.
- [3] Emmanuel Candès and J. Romberg, “Sparsity and incoherence in compressive sampling,” *Inverse Problems*, vol. 23, pp. 969–985, 2006.
- [4] Stephen Becker, Jérôme Bobin, and Emmanuel Candès, “NESTA: A fast and accurate first-order method for sparse recovery,” Tech. Rep., Caltech, 2009.
- [5] Emmanuel Candès and Terence Tao, “Near-optimal signal recovery from random projections: universal encoding strategies?,” *IEEE Transactions on Information Theory*, vol. 52, pp. 5406–5425, 2006.

⁹Changing the sparse gradient model for another sparse regularization scheme (such as a sparse wavelet decomposition) would not drastically modify the result, as this real-world signal is not strictly sparse whatever the decomposition basis is, assuming the basis is built in a reasonable way.