ON THE EMPIRICAL RATE-DISTORTION PERFORMANCE OF COMPRESSIVE SENSING

Adriana Schulz[1,2], *Luiz Velho*[1], *Eduardo A. B. da Silva*[2]

[1] IMPA - Instituto de Matemática Pura e Aplicada, Brazil;
[2]PEE/COPPE/DEL/Poli, Univ. Fed. do Rio de Janeiro, Brazil.
e-mails: aschulz@impa.br, lvelho@impa.br, eduardo@lps.ufrj.br

ABSTRACT

Compressive Sensing (CS) is a new paradigm in signal acquisition and compression. In compressive sensing, a compressible signal is acquired using much less measurements than the ones required by Nyquist theorem, provided that it is sparse in some transform domain. Recovery of the signal from the measurements is achieved by convexoptimization techniques (e.g. l1-norm minimization). CS has been attracting the interest of the signal compression community, and a vast body of knowledge about it has already been developed. When it comes to image compression applications, one is ultimately interested in how many bits one spends for a given image quality. Although several theoretical results regarding the rate-distortion performance of CS have been published recently, there are not too many practical image compression results available. The main goal of this paper is to carry out an empirical analysis of the rate distortion performance of CS in image compression. We analyze issues such as the minimization algorithm used and the transform employed, as well as the trade-off between number of measurements and quantization error. From the experimental results obtained we highlight the potential and limitations of CS when compared to traditional image compression methods.

***Index Terms*—** Compressive Sensing, Rate-Distortion Analysis, Quantization, Image Compression, Sparsity.

1. INTRODUCTION

Ordinary images, as well as most natural and manmade signals, tend to be compressible and, therefore, can be well approximated by a sparse representation. Standard image acquisition techniques follow the sample-then-compress framework. It involves sampling at a large rate only to discard most of the acquired information using a compression scheme that explores the domain in which the image is sparse.

In this context, CS comes out as a new paradigm for data acquisition; it gives a stable and robust algorithm that allows sensing at rates much smaller than the Nyquist limit, while recovering the signals with little corruption [1, 2].

To compress sensed information, it is necessary to make use of quantization schemes that add distortion to the acquired data. Therefore, a relevant contribution to CS theory consists in verifying how it performs in the presence of quantization errors and in a rate-distortion sense.

Theoretical results have been established guaranteeing stability of CS to the addition of quantization errors. In [3] CS encoding of approximately sparse signals with quantized measurements is studied and performance is demonstrated to be within a logarithmic factor to the one of the optimal encoder.

Related works consider strictly sparse signals and evaluate CS when quantization errors are added. In [4], the results of [3] are extended to the scenario where exact sparsity is guaranteed and inefficiencies in terms of rate and performance are verified, suggesting modifications in the uniform scalar quantization method and the reconstruction algorithm. In [5], both of these changes are explored and extensive computer simulations are made confirming their advantages.

The rate-distortion function is used in [6] to compare CS to the ideal compression scheme (where an oracle informs the sparsity pattern) and the loss in performance is evaluated as relatively small (an additive logarithmic rate penalty is observed). In [7], a lower bound on the number of measurements needed to reconstruct a signal is set as a function of the measurements' SNR and rate-distortion function.

However, fundamental questions regarding performance in practical applications still remain unanswered.

In this work we aim at contributing to answering some of these questions from an empirical point of view. Though innumerable applications have been suggested, we will concentrate our study on the scenario in which CS was first presented and is mostly discussed: image sensing and compression.

It is important to emphasize that it is not the intention of this work neither to elaborate further theoretical results nor to analyze CS fundamental limitations in idealized scenarios. Instead we aim at evaluating applications in image acquisition by means of empirical analysis.

1.1. Overview of CS and Basic Notation

Let $x \in C^{N\times 1}$ be a vector representation of an image and $\Psi \in C^{N\times N}$ a transform that makes x sparse, i.e., $\Psi x = s$, where s has only S nonzero coefficients. Since we are acquiring only $M \ll N$ measurements, sensing can be denoted by $y = \Phi x = \Theta s$, where $\Phi, \Theta \in C^{M\times N}$ and $\Theta = \Phi\Psi^*$.

CS theory states that it is possible, through the use of a convex optimization algorithm, to recover x from y with over-

(a) *Phantom* (b) *Lena* (c) *Camera man* (d) *Text*

Fig. 1. Test images.

whelm probability if Θ satisfies a *Restricted Isometry Property (RIP)* [8]. It has also been shown that this property can by assumed if the entries of Θ belong to certain random ensembles and M is on the same order as $S\log(N/M)$.

Moreover, this acquisition technique is robust to sparsity approximations and measurements errors [9]. Let $y \in C^{M\times 1}$ be the acquired data corrupted by noise, i.e., $y = \Phi x + n$, where $\|n\| \leq \epsilon_q$. If we reconstruct the signal by solving the convex optimization problem

$$\min_x \|\Psi x\|_{l_1} \quad \text{subject to} \quad \|y - \Phi x\|_{l_2} \leq \epsilon_q. \tag{1}$$

then the recovery error is bounded by the sum of the measurement (quantization) error and the error due to the fact that the signal is not strictly sparse, i.e.

$$\|y - \Phi x\|_{l_2} \leq C \cdot \left(\epsilon_q + \underbrace{S^{-1/2}\|x_S - x\|_{l_1}}_{\epsilon_s} \right), \tag{2}$$

where C is relatively small and x_S is an approximation of x where the S largest coefficients in the Ψ domain are observed. This implies that the reconstruction error in CS is of the order of the maximum of the quantization and measurement errors. For details see [9].

2. EXPERIMENTAL SETUP

CS investigations were made on four different images of size $N = 256^2$, which differ in terms of both sparsity and high energy coefficient distribution in the frequency domain (see Figure 1). Since the images are stored in the computer as a matrix of pixels, we simulate acquisition by means of measurements that involve linear combinations of these pixels.

Measurements were taken by choosing at random M waveforms of an $N \times N$ Noiselet transform [10]. Such measurements were chosen because they are highly *incoherent* [1] with the considered sparse domains and the RIP tends to hold for reasonable values of M. In addition, the matrix created is orthogonal and self-adjoint, thus being easy to manipulate.

The following recovery strategies were considered:

A. Minimization of the l_1-norm of the image's DCT;

B. Minimization of the l_1-norm of the image's DWT;

C. Minimization of the image's TV-norm; and

D. Minimization of the l_1-norm of the image's SVD

The efficiency of each strategy is related to how sparse the images are at the considered domain. The DCT and the Wavelet domains were chosen due to their large use in image compression standards. In addition, since most published theorems relate to orthogonal rather than to the more efficient biorthogonal basis, we used an orthonormal Wavelet basis (Coiflet of 2 vanishing moments).

In many recent publications [1, 11], CS researchers have used the total variation (TV) norm, which can be interpreted as the l_1-norm of the (appropriately discretized) gradient. Applied to images, the TV-norm minimization favors a certain smoothness that is usually found in natural and manmade pictures and is, therefore, very effective.

Finally, the SVD was calculated for each image and used to determine sparsity domains because it gives a very accurate sparse representation. This technique requires the knowledge of the SVD basis, that is calculated from the whole image information (not available in CS) and requires a large data rate for transmission (which is not taken into account). Nevertheless we used such results as upper bounds that, although loose, give interesting insights on performance limitations.

While strategies A, B and D solve Equation 1, in strategy C the image is reconstructed by solving the following convex optimization problem:

$$\hat{x} = \min_x \|x\|_{TV} \quad \text{subject to} \quad \|y - \Phi x\|_{l_2} \leq \epsilon. \tag{3}$$

2.1. The Rate-Distortion Curve

Scalar uniform quantization was considered and tests were made for different quantization steps in order to select the best for each compression rate and plot rate-distortion curves.

Rate was calculated as $(M/N)H_y$, where H_y is the entropy (in bits per pixel) of the quantized measured data y estimated from its histogram. The case of unused quantization values is dealt with by considering each of them to have occurred once.

2.2. Implementation Aspects

The experiments were implemented in `MATLAB` and the l_1`-Magic` [12] toolbox was used to solve both optimization problems (Equation 1 and the minimization of the TV norm). For each image, recovery strategy and quantization step, the parameter ϵ_q (see Equation 1) was chosen experimentally in order to maximize the PSNR.

The Wavelet basis was generated using the `WAVELAB` [13] package and the Noiselet basis using an algorithm made available by Romberg [11].

3. RESULTS

In Figure 2 the rate-distortion curve was plotted for all tested images and considered strategies. We observe that CS recovery schemes that perform the l_1-norm minimization in the

Wavelet domain are far less efficient than the JPEG2000 standard. However, by analyzing the results for strategy D and for the test image *Phantom* on strategy C, we see that there is room for improvements; in both cases one gets better results than JEPG2000. The *Phantom* image in the frequency domain and the SVD transform are both very sparse. This fact gives an indication that by choosing representations that strengthen sparsity one can reduce not only the number of measurements needed to reconstruct the signal but also the approximation error.

It is important to mention that, though strategy D presents an upper bound to CS performance, it is not a practical because it requires an *a priori* knowledge of the image's SVD. Figure 5 highlights this argument by contrasting recovery of the image *Camera man* using as basis *Camera man*'s SVD and *Lena*'s SVD.

In Figure 3 the Rate $\times$ PSNR curve was plotted for strategies B and C using varying quantization steps. We observe that for a particular compression rate, each image and recovery strategy has an optimal quantization step that produces the highest PSNR. If the image is not sparse in the considered domain, the curves show that it is more efficient to take a large amount of measurements and compensate for the potential rate increase by enlarging the quantization step.

One can also observe that for a fixed PSNR, the ideal quantization step is approximately the same in all evaluated scenarios. This observation is closely related to the result in Equation 2, which indicates that the recovery error is on the same order as the largest of the approximation and measurement errors [3]. The PSNR determines the acceptable distortion and, therefore, the values of ϵ_q and ϵ_s. While ϵ_q only depends on the quantization step, ϵ_s depends on the sparsity distribution and, hence, on the number of measurements.

The same comment can be made after observing Figure 4, that shows the results in terms of number of measurements $\times$ PSNR. For each strategy the number of measurements determines ϵ_s; in addition, all quantization steps that make ϵ_q on the order of ϵ_s (or smaller) result in the same PSNR (see Equation 2). Therefore, all curves overlap until the number of measurements is large enough so that ϵ_s exceeds ϵ_q (see Figure 4.(b)). In Figure 4.(a), it is noteworthy that for quantization steps smaller then 3, the curves overlap completely; this is so because once the errors due to sparsity are very large, reducing the quantization step is ineffective in increasing PSNR. In contrast, in Figure 4.(c), where the image is strongly sparse in the considered domain (SVD), ϵ_s tends to be much smaller, and therefore such behavior is not observed.

4. DISCUSSION AND CONCLUSIONS

The results obtained during this study suggest contexts in which improvements in the CS acquisition strategy could lead to better rate-distortion performance.

It has already been emphasized that sparsity plays a very important role in the recovery performance (e.g. SVD). Therefore, to make CS applications in imaging practical, domains that enhance sparsity, such as biorthogonal Wavelets and gradient-based models (e.g. TV), should be investigated.

Another significant aspect in CS development is the recovery algorithm. Though only the l_1 and TV minimization were evaluated in this work, there are recent algorithms that not only speed up, but also improve reconstruction [14].

Different quantization models have also been proposed as a way of improving CS performance [5]. Moreover, alternitive sensing matrices are also worth investigating.

A critical aspect to discuss is the universality of CS. Most of the referenced publications point out that one of the greatest advantage in CS is that it does not need to be adaptive. By this, we mean that encoding can be done without the knowledge of the sparsity distribution. In truth, as long as there exists a domain in which the signal is sparse, the same random measurements can be taken to reconstruct it. However, in the performed experiments, the best trade-off between number of measurements and quantization step varies according to the signal's sparsity distribution. Devising strategies that mitigate this effect is another topic worth investigating.

5. REFERENCES

[1] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. on Inf. Theory*, 2006.

[2] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, 2006.

[3] E. Candès and J. Romberg, "Encoding the l_p ball from limited measurements," *DCC*, 2006.

[4] P. Boufounos and R. Baraniuk, "Quantization of sparse representations," *DCC*, 2007.

[5] W. Dai, H. V. Pham, and O. Milenkovic, "Distortion-Rate Functions for Quantized CS," *ArXiv e-prints*, Jan. 2009.

[6] A. Fletcher, S. Rangan, and V. Goyal, "On the rate-distortion performance of compressed sensing.," *ICASSP*, 2007.

[7] S. Sarvotham, D. Baron, and R. Baraniuk, "Measurements vs. bits: Compressed sensing meets information theory," *Allerton Conf. on Communication, Control, and Computing*, 2006.

[8] E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. on Information Theory*, vol. 51, no. 12, 2005.

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

[10] R. Coifman, F. Geshwind, and Y. Meyer, "Noiselets," *Appl. Comp. Harmon.Anal.*, vol. 10, no. 1, 2001.

[11] J. Romberg, "Imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, 2008.

[12] E. Candès and J. Romberg, "L1–magic," www.l1-magic.org.

[13] D. Donoho, A. Maleki, and M. Shahram, "Wavelab," http://www-stat.stanford.edu/ wavelab.

[14] E. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted l_1 minimization," *J. Fourier Anal. Appl.*, 2007.

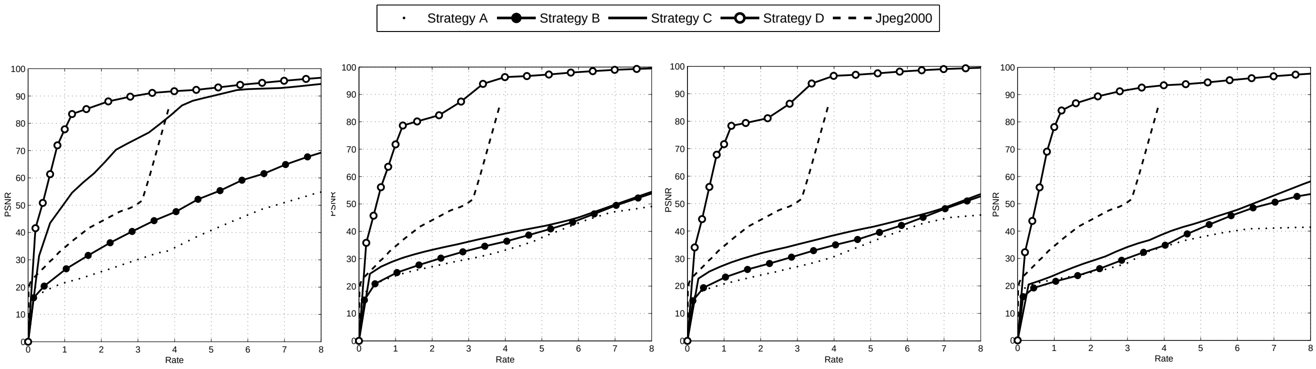

(a) *Phantom* (b) *Lena* (c) *Camera man* (d) *Text*

Fig. 2. Rate-Distortion curves for all considered strategies and JPEG2000.

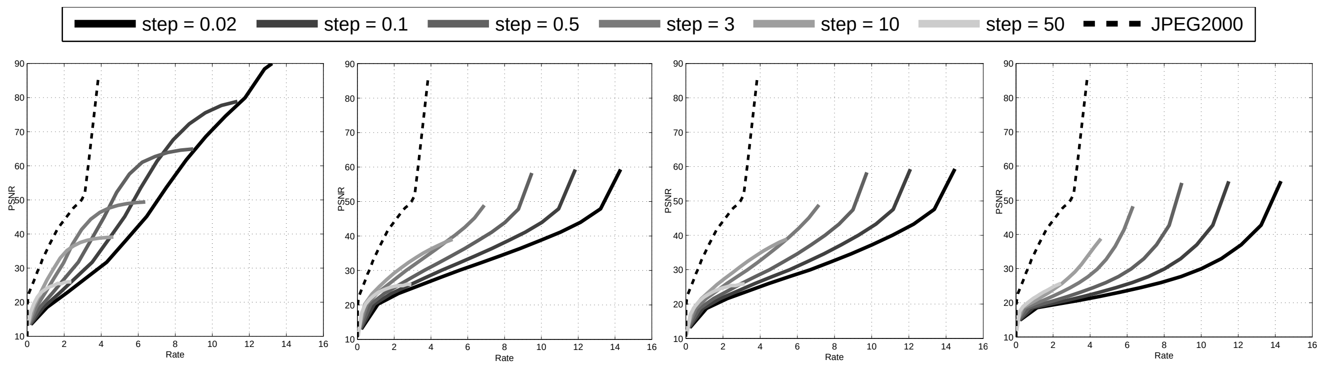

(a) *Phantom* (b) *Lena* (c) *Camera man* (d) *Text*

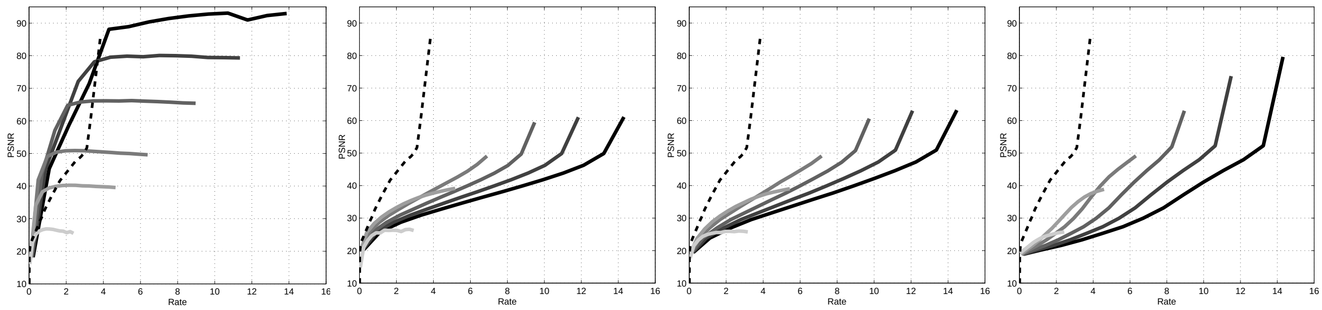

(e) *Phantom* (f) *Lena* (g) *Camera man* (h) *Text*

Fig. 3. Rate × PSNR for varying quantization steps: (a-d) strategy B and (e-h) strategy C (see Section 2 for strategies definition).

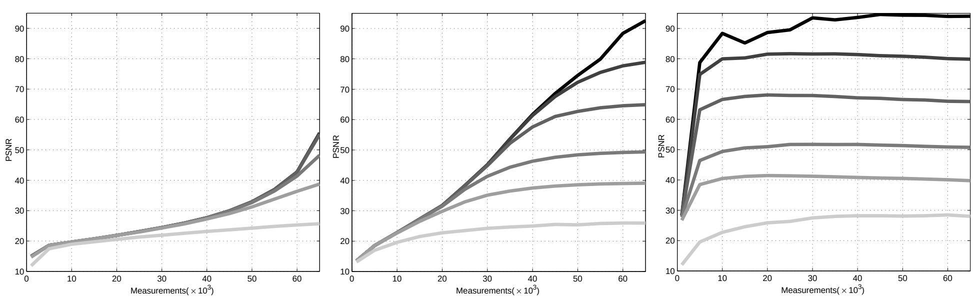

(a) *Text*, Strategy B (b) *Phantom*, Strategy B (c) *Text*, Strategy D

Fig. 4. Number of Measurements × PSNR (see Section 2 for strategies definition).

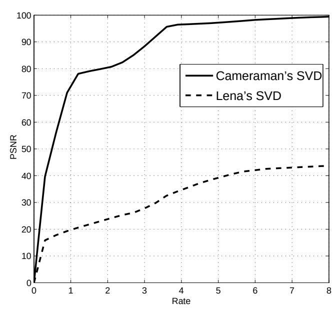

Fig. 5. Rate × PSNR *Camera man*, Strategy D.