To be published in Applied Optics:

Title:	An Active Illumination Single-Pixel Camera Based on Compressive Sensing
Authors:	Filipe Magalhaes, Francisco Araujo, Miguel Correia, Mehrdad Abolbashari, and Faramarz Farahi
Accepted:	20 December 2010
Posted:	21 December 2010
Doc. ID:	136601

An Active Illumination Single-Pixel Camera Based on Compressive Sensing

Filipe Magalhães,[1,2*] Francisco M. Araújo,[1] and Miguel V. Correia[1,2]

[1] *INESC Porto, Rua do Campo Alegre, 687, 4169-007 Porto, Portugal*

[2] *Universidade do Porto, Faculdade de Engenharia, Departamento de Engenharia Electrotécnica e de Computadores, Rua Dr. Roberto Frias, s/n 4200-465 Porto, Portugal*

Mehrdad Abolbashari,[3] Faramarz Farahi,[3,4]

[3] *Center for Optoelectronics and Optical Communications, University of North Carolina at Charlotte, Charlotte, NC 28223, USA*

[4] *Department of Physics and Optical Science, University of North Carolina at Charlotte, Charlotte, NC 28223, USA*

[*] *Corresponding author: filipe.magalhaes@inescporto.pt*

In this paper an optical imaging system based on compressive sensing (CS) is presented along with its principal mathematical aspects. Although CS is undergoing significant advances and empowering many discussions and applications throughout various fields, this article focus on the analysis of a single-pixel camera. This work was the core for the development of a single-pixel camera approach based on active illumination. Therefore, the active illumination concept is described along with experimental results, which were

very encouraging towards the development of compressive sensing based cameras for various applications, such as pixel level programmable gain imaging.

OCIS codes: (110.0110) Imaging systems; (110.3010) Image reconstruction techniques; (110.1758) Computational imaging.

Introduction

It is clear that the Nyquist-Shannon sampling theorem has been a fundamental rule of signal processing for many years and can be found in nearly all signal acquisition protocols, being extensively used from consumer video and audio electronics to medical imaging devices or communication systems. Basically, it states that a band-limited input signal can be recovered without distortion if it is sampled at a rate of at least twice the highest frequency component of interest within the signal. For some signals, such as images that are not naturally band limited, the sampling rate is dictated not by the Nyquist-Shannon theorem but by the desired temporal or spatial resolution. However, it is common in such systems to use an anti-aliasing low-pass filter to band limit the signal before sampling it, and so the Nyquist Shannon theorem plays an implicit role [1].

In the last few years, an alternative theory has emerged, showing that super-resolved signals and images can be reconstructed from far fewer data or measurements than what is usually considered necessary. This is the main concept of compressive sensing (CS), also known as compressed sensing, compressive sampling and sparse sampling. In fact, "the theory was so revolutionary when it was created a few years ago that an early paper outlining it was initially rejected on the basis that its claims appeared impossible to substantiate [2]."

CS relies on the empirical observation that many types of signals or images can be well approximated by a sparse expansion in terms of a suitable basis, that is, by only a small number of non-zero coefficients. This is the key aspect of many lossy compression techniques such as JPEG and MP3, where compression is achieved by simply storing only the largest basis coefficients.

In CS, since the number of samples taken is smaller than the number of coefficients in the full image or signal, converting the information back to the intended domain would involve solving an underdetermined matrix equation. Thus, there would be a huge number of candidate solutions and, as a result, we must find a strategy to select the "best" candidate.

Different approaches to recover information from incomplete data sets have existed for several decades. One of its earliest applications was related with reflection seismology, in which a sparse reflection function (indicating meaningful changes between surface layers) was sought from band limited data [1, 3, 4]. It was, however, very recently, that the field has gained increasing attention, when Emmanuel J. Candès, Justin Romberg and Terence Tao [5], discovered that it was possible to reconstruct Magnetic Resonance Imaging (MRI) data from what appeared to be highly incomplete data sets in face of the Nyquist-Shannon criterion (see Fig. 1). Following, Candès et al. work , this decoding or reconstruction problem can be seen as an optimization problem and be efficiently solved using the ℓ_1-norm [6].

As a result, CS has become a kind of revolutionary research topic that draws from diverse fields, such as mathematics, engineering, signal processing, probability and statistics, convex optimization, random matrix theory and computer science.

Undergoing significant advances, CS has proved to be far reaching and has enabled several applications in many fields, such as: distributed source coding in sensor networks [7, 8], coding, analog–digital (A/D) conversion, remote wireless sensing [1, 9] and inverse problems, such as those presented by MRI [10].

One application with particular interest within the aim of the work presented here, is the ground-breaking single-pixel imaging setup developed by D. Takhar, et al. at the Rice University [11].

This paper is organized as follows. After a brief introduction (section I), some mathematical background essential to the understanding of CS is shown (section II). Then, on section III, CS is presented along with some of its principal properties. Section IV explains why ℓ_1-norm is such a good option for compressive sensing. Some insights about the robustness of CS in the presence of noise are given on section V. Next, on section VI the single-pixel camera developed at Rice is discussed. Subsequently, the innovative active illumination single-pixel camera developed in the scope of the current work is described. Following, experimental results from the single–pixel cameras are presented. In the end, the main conclusions of this work are exposed.

Compressive Sensing Background

In order to become possible, CS is built upon two principles: sparsity, related with the signals of interest, and incoherence, related with the sensing modality.

K-sparse and compressible signals

Let's consider a real-valued, finite-length, one dimensional, discrete-time signal x, which can be viewed as a $\mathrm{N}\times 1$ column vector in $\Re^{\mathrm{N}}$ with elements $x[n]$, with $n = 1,2,\ldots,\mathrm{N}$. Any signal in $\Re^{\mathrm{N}}$ can be represented in terms of a basis of $\mathrm{N}\times 1$ vectors $\{\psi_i\}_{i=1}^{\mathrm{N}}$. For simplicity, let's assume that the basis is orthonormal. Using the $\mathrm{N}\times\mathrm{N}$ basis matrix $\Psi = \{\psi_1, \psi_2, \ldots, \psi_{\mathrm{N}}\}$ with the vectors $\{\psi_i\}$ as columns, a signal x can be expressed as:

$$x = \sum_{i=1}^{\mathrm{N}} s_i \psi_i \text{ or } x = \Psi s, \tag{1}$$

where s is the $\mathrm{N}\times 1$ column vector of weighting coefficients $s_i = \langle x, \psi_i \rangle = \psi_i^T x$. s and x are equivalent representations of the signal with x in time or space domain and s in Ψ domain.

The signal x is K-sparse if it is a linear combination of only K basis vectors, which means that only K of the s_i coefficients in Eq. (1) are nonzero, while the remaining $(N-K)$ coefficients are zero. In addition, the signal x is compressible if the representation in Eq. (1) has just a few large coefficients and many small coefficients, setting the basis of transform coding. Therefore, we can say that a signal x is sparse in the Ψ domain if the coefficient sequence is supported on a small set, and compressible if the sequence is concentrated near a small set.

In face of the typical data acquisition paradigm, huge amounts of data are collected only to be in large part discarded at the compression stage to facilitate storage and transmission. Imagine, for example, a digital camera that captures images with millions of sensors (pixels) but eventually encodes the image in just a few hundred kilobytes. Clearly, this is a tremendously wasteful process and suffers from three principal drawbacks. First, the initial number of samples N may be large, even if the desired K is small. Second, the set of all N transform coefficients

$\{s_i\}$ must be computed even though all but K of them will be discarded. Third, there is an overhead that is introduced by the encoding of the large coefficients locations.

Recovering K-sparse signals

Following the work presented in [12], Candès and Tao developed a refined version of the *Uniform Uncertainty Principle* (UUP) [13], which has proved to be essential to the study of the general robustness of CS. This key notion was then named *Restricted Isometry Property* (RIP) and can be defined as follows:

For each integer $K = 1,2,\ldots$, define the isometry constant δ_K of a measurement matrix A as the smallest number such that

$$(1-\delta_K)\|x\|_{\ell_2}^2 \le \|Ax\|_{\ell_2}^2 \le (1+\delta_K)\|x\|_{\ell_2}^2 \tag{2}$$

holds for all K-sparse vectors x. Therefore, we can say that a matrix A obeys the RIP of order K if δ_K differs enough from one. When this condition is verified, A approximately preserves the Euclidean length of K-sparse signals, which in turn implies that K-sparse vectors cannot be in the null space of A. An alternative description of this property is to say that all subsets of K columns taken from A are in fact nearly orthogonal (they cannot be exactly orthogonal since we have more columns than rows).

Let's imagine we want to acquire K-sparse signals making use of matrix A. Suppose that δ_{2K} is not close to one. This indicates that all pair-wise distances between K-sparse signals must be well preserved in the measurement space, which means that

$$(1-\delta_{2s})\|x_1 - x_2\|_{\ell_2}^2 \le \|Ax_1 - Ax_2\|_{\ell_2}^2 \le (1+\delta_{2s})\|x_1 - x_2\|_{\ell_2}^2 \tag{3}$$

is true for all K-sparse vectors x_1, x_2 [1, 14].

Incoherence

Let's now consider $M < N$ linear measurements of x and a collection of test functions $\{\varphi_m\}_{m=1}^{M}$ such that $y[m] = \langle x, \varphi_m \rangle$. By stacking the measurements $y[m]$ into the $M \times 1$ vector y and the test functions φ_m^T as rows into an $M \times N$ sensing matrix Φ we can write

$$y = \Phi x = \Phi \Psi s = \Theta s \,. \tag{4}$$

A condition related with RIP is *incoherence*, which requires that the rows of Φ (the measurement or sensing matrix) cannot represent the columns of Ψ in a sparse way (and vice-versa).

Incoherence extends the duality between time and frequency and expresses the idea that an object having a sparse representation in Ψ must be spread out in the domain in which it was acquired. This incoherence property is verified for many pairs of bases, including, for instance, delta spikes and sine waves of the Fourier basis, or the Fourier basis and wavelets.

The coherence between the sensing basis Φ and the representation basis Ψ can be given by the following equation:

$$\mu(\Phi, \Psi) = \sqrt{n} \cdot \max_{1 \le k, j \le n} \left| \langle \varphi_k, \psi_j \rangle \right| , \tag{5}$$

which, in simple words, is measuring the largest correlation between any two elements of Φ and Ψ. CS is essentially interested in low coherence pairs. For instance, for the previously referred delta spikes and sine waves (time-frequency) pair, $\mu(\Phi, \Psi) = 1$, therefore, indicating maximal incoherence [1, 15].

A particular aspect of interest is that random matrices are largely incoherent with any fixed basis Ψ. This empowers the use of known fast transforms such as a Walsh, Hadamard, or Noiselet transform [16].

Furthermore, what is most remarkable about this concept is that it allows capturing information contained in a sparse signal in a very efficient way without trying to understand that signal.

How Compressive Sensing Works

Compressive sensing addresses the inefficiencies presented by the *sample-then-compress* framework by directly acquiring a compressed signal representation, avoiding the intermediate stage of acquiring N samples [5]. CS bypasses the sampling process and directly acquires a condensed representation y consisting of M linear measurements. Furthermore, The measurement process is nonadaptive in that Φ does not depend in any way on the signal x.

The transformation from x to y is a dimensionality reduction and so loses information in general. In particular, since $M < N$, for a given y, there is an infinite number of x' such that $\Phi x' = y$. The overwhelming capacity of CS is that Φ can be designed such that sparse/compressible x can be recovered exactly/approximately from measurements of y.

To recover the signal x from the random measurements y, the traditional favorite method of least squares has been shown to fail with high probability. Instead, it has been demonstrated that using the ℓ_1 optimization [12]

$$\hat{s} = \arg\min \left\| s' \right\|_{\ell_1} \text{ such that } \Theta s' = y \tag{6}$$

it is possible to exactly reconstruct K-*sparse* vectors and closely approximate compressible vectors stably with high probability using just $M \geq O(K \cdot \log(N/K))$ random measurements [5, 6]. Minimizing the ℓ_1-norm subject to linear equality constraints can easily be recast as a linear program, also known as *basis pursuit*, which can find several alternative reconstruction techniques based on greedy, stochastic and variational algorithms [5, 9, 17, 18].

In addition to enabling sub-Nyquist sampling, CS exhibits a number of attractive properties.

- **Universality:** Φ can be considered a universal encoding strategy, as it does not need to be designed with regards to the structure of Ψ. This allows exactly the same encoding strategy to be applied in a variety of different sensing environments; no knowledge is required about the subtleties of the data being acquired. Random measurements are also future proof – i.e., if new research yields a better sparsity inducing basis, then the same set of random measurements can be used to reconstruct data with even better quality –.
- **Encryption:** A pseudorandom basis can be generated using a simple algorithm according to a random seed. Such encoding effectively implements a form of encryption: randomized measurements will themselves resemble noise and be meaningless to an observer who does not know the associated seed.
- **Robustness and progressivity:** Random coding is robust in that the randomized measurements have equal priority, unlike the Fourier or wavelet coefficients in current transform coders. Thus, this enables a progressively better data reconstruction as more measurements are obtained. Besides this, one or more measurements can also be lost without corrupting the entire reconstruction.

Oppositely, all the bits in JPEG 2000 do not have the same value and if important bits are missing – e.g., because of packet loss –, then it is impossible to retrieve the information accurately.

In addition, the robustness of CS does not become impaired by the quantization process.

- **Scalability:** the number of measurements to compute can be adaptively selected in order to trade off the amount of compression of the acquired image versus acquisition time. In contrast, conventional cameras trade off resolution versus the number of pixel sensors.
- **Computational asymmetry:** CS places most of its computational complexity in the recovery system (decoder), which will often have more substantial computational resources than the encoder. The encoder is very simple since it merely computes incoherent projections and makes no decisions [11].

The Geometry of ℓ_1-norm

The geometry of CS problems in $\Re^N$ helps to visualize why ℓ_2 reconstruction fails to find the sparse solution that can be identified by ℓ_1 reconstruction. Figure 2 presents significant information to this subject. Part (a) illustrates the ℓ_2 ball in $\Re^3$ with a certain radius. It must be emphasized that this ball is isotropic. Part (b) represents the ℓ_1 ball in $\Re^3$, which is anisotropic ("pointy" along the axes).

The ℓ_2 minimizer $\hat{s}$ is the point from H closest to the origin. This point can be found by blowing the ℓ_2 ball until it bumps into H. Due to the random orientation of H (imposed by the

randomness in matrix Θ), the closest point $\hat{s}$ will be away from the coordinate axes with high probability and, therefore, will not be sparse and will be far from the correct answer s. In higher dimensions, this difference becomes even more significant. Paying attention to the part (b) of Figure 2, it can be seen that the point of intersection $\hat{s}$ is now defined by the vector that solves equation (6).

Robustness of Compressive Sensing

In any realistic application, we cannot expect to measure Φx without any error. Therefore, now, it is important to analyze the robustness of compressive sampling in face of measurement errors. This is a critical topic since any real-world sensor is subject to at least a small amount of noise. For that reason, one immediately understands that to be widely applicable, the methodology needs to be stable. Small perturbations in the observed data should, then, induce small perturbations in the reconstruction. Fortunately, the recovery procedures may be adapted to be surprisingly stable and robust in the presence of arbitrary perturbations.

Let's suppose the measurements are affected by noise and define the following model:

$$y = \Phi x + e, \tag{7}$$

where e is a stochastic or deterministic error term with bounded energy $\|e\|_{\ell_2} \leq \varepsilon$, being ε an upper bound on the noise magnitude.

Because of the measurement inaccuracies, a modification has been introduced to equation (6) to make it noise-aware. In this way, the reconstruction proposal has the following form:

$$\hat{s} = \arg\min \|s'\|_{\ell_1} \text{ such that } \|y - \Phi\Psi s'\|_{\ell_2} < \varepsilon, \tag{8}$$

which satisfies $\|\hat{s}-s\|_{\ell_2} < C_N\varepsilon + C_K\sigma_K(x)$ with overwhelming probability. C_N and C_K are the noise and approximation error amplification constants, respectively, and $\sigma_K(x)$ is the ℓ_2 error incurred by approximating s using its largest K terms. Once again, this problem is convex and can be solved using standard convex programming algorithms [14, 19].

The Single-Pixel Camera

The single-pixel camera, developed originally at the Rice University [11], is one of the most paramount examples of CS. It can be seen as an optical computer comprising a digital micro-mirror device (DMD) with 1024 × 768 micromirrors, two lenses, a single photodetector and an A/D converter. Basically, this configuration computes random linear measurements of the scene under view. The image is then reconstructed from these measurements by a digital computer. A block-diagram depicting the single-pixel camera setup can be seen on Figure 3.

This time-multiplexing technique enables the use of a single and yet more sensitive photon detector. This is particularly important when the detector is expensive, making an N-pixel array/matrix prohibitive. A single-pixel camera can also be adapted to image at wavelengths that are currently impossible with conventional digital cameras.

Figure 4 presents the experimental setup comprising the optical hardware of the single-pixel camera previously described [11]. Following the red arrows in Figure 4, it can be seen that a light source is used to illuminate the object (in this case, a black and white printout of an "R" character). Then, the object's image is formed by the means of Lens 1 on the DMD that adequately reflects or not the light incident on each of its pixels, depending on the imposed modulation pattern. The light reflected towards Lens 2 will finally be concentrated on the single

light detector that will integrate it, thus, yielding an output voltage that depends on the used DMD modulation pattern. This voltage is amplified through an operational amplifier circuit to be finally digitized by an A/D converter. This process is repeated until M values are acquired so that we can, finally, use them to reconstruct the imaged object. Each of these values (output voltage of the photodiode) can be interpreted as the inner product of the desired image x with a measurement basis $\Phi(m), m = 1,2,\ldots,M$.

With this setup the resolution of the reconstructed image is limited by the pixel arrangement of the DMD.

Active Illumination Single-Pixel Camera

Following the work of the Rice group [11, 19], we have developed an active illumination single-pixel camera. Instead of the DMD, our setup makes use of a video projector to incorporate the random measurement matrix into the system. The proposed experimental setup is presented in Figure 5.

Figure 6 shows an integrated version of the proposed single-pixel camera in a much smaller assembly comprising the same configuration of Figure 5. This integrated setup comprises the lens and the photodiode circuit in a single integrated module (see Figure 6 (b)), therefore, its compactness.

A large area (1 cm^2) silicon photodiode (RS 303-674) was used in order to facilitate the optical alignment. The experimental setup includes a video projector with 1280×800 maximum resolution (Epson® model EB-W7), a 12 bits A/D data acquisition board with up-to 10kS/s

(National Instruments™ USB-6008), and a 8 mm focal length Computar® lens with focus distance ranging from 30 cm to infinity.

Regarding our configuration, the video projector was used to project the result of the product between the image to be reconstructed and the random measurement patterns (see Figure 7). Therefore, each of the output voltages of the photodiode amplifier circuit is representative of the inner product between the used pattern for that measurement and the image to be reconstructed.

Results and Discussion

Original Single-Pixel Camera

The Rice group set the single-pixel camera to acquire a 64 × 64 pixels image (hence, N = 4096). This size was chosen to ensure quick reconstruction during tests so that focusing and other adjustments could be made.

Since the test image was piecewise constant (with sharp edges) it could be sparsely represented in the wavelet domain. Figures 9(a) and 9(b) show the best K-term Haar wavelet approximation of the idealized image in Figure 8, with K = 400 and 675, respectively. Using M = 820 and 1600 pseudorandom projections, it is possible to obtain the reconstructed images shown in Figures 10(a) and 10(b), respectively, using Basis Pursuit.

From these results, it is clear that the recognizable features of the “R” can be recovered, even with a number of measurements below half of the total number of pixels of the reconstructed image. It can also be seen that the reconstruction quality turns progressively better

as M increases, as well as more robust to noisy measurements, enhancing the reconstruction of the singularities (sharp edges). As stated by the authors, the sources of noise included subtle nonlinearities in the photodiode, nonuniform reflectance of the mirrors towards the lens that focused onto the photodiode, and nonuniform mirror positions. The robustness of the CS reconstruction tends to suppress quantization noise from the A/D converter and photodiode circuit noise during detection[11].

Some other results, obtained with the same setup, are presented in Figures 11 and 12.

Comparatively to imaging by raster scanning, the Rice single-pixel camera provided results far superior. Figure 13 shows comparative results of the two approaches.

From Figure 13, it is clear that the results from CS are much better. This is due to the fact that raster scanning measures only 1/N of total light per measurement, while CS measures ½ of total light per measurement [20]. Operation in a CS mode also significantly reduces dark noise.

Active Illumination Single-Pixel Camera

Comparatively to the work presented by the Rice group, some preliminary results were obtained using the active illumination configuration for the single-pixel camera. Initial results obtained with the active illumination approach for an image containing sharp edges are presented in Figure 14. The reconstructed images have 32 × 32 pixels, in order to speed up the reconstruction and testing procedures during the experimental phase. The reconstructions were obtained with the ℓ_1-Magic software package[1]. ℓ_1-Magic is a collection of MATLAB® routines for solving the convex optimization programs central to compressive sampling. The results were obtained with

MATLAB® on Windows Vista™ with an Intel® Core™2 Duo CPU @ 2.50GHz and 3GB of RAM. The average processing time to reconstruct a 32×32 pixels image from 410 measurements was approximately 10 seconds and the average processing time to reconstruct a 64×64 pixels image from 1640 measurements was approximately 400 seconds. In addition, it should be taken into account the time consumed to project all the codes (300 ms for each code). Therefore, bigger images can be reconstructed at the expense of more time.

It is clear that with an increasing number of measurements, the overall quality of the reconstructed images is also increased. The size of the "F" character on the wall was 4cm × 3cm, thus resulting on a resolution of approximately 2 mm per pixel.

Figure 15 depicts the results obtained with the active illumination system for imaging of a more complex scene. In this case, only the random patterns were projected as the scene was composed by real objects.

Despite of the video projector resolution limit (even though it was not fully exploited), super-resolution images could have been obtained either by using multiple sub-pixel shifted images of the same scene [21] or by exploring the fact that patches in a natural image tend to redundantly recur many times inside the image, both within the same scale, as well as across different scales [22].

As stated before, random matrices are largely incoherent with any fixed basis Ψ and, therefore, for the measurement or sensing matrices, we have used pseudorandomly generated (with uniform distribution) and Hadamard based ones, having both given similar results.

[1] available at http://www.acm.caltech.edu/l1magic/

We also performed some experiments in order to test the robustness of the system in the presence of noise. We concluded that the system was able to reconstruct images of similar quality when the amplitude of the added noise was below 20% of the maximum amplitude of the signal. The results of the conducted experiments are depicted on Figure 16.

Some experiments regarding the reconstruction of colored images were also conducted. In this case, the object to be imaged is a red squared contour on a green background painted on normal paper (see Figure 17a). For this experiment, the object was fixed and only the random patterns were projected. For each color channel (RGB), 410 values were measured and a single image was reconstructed. In this experiment, the random patterns projected to obtain the measurements to reconstruct each of the RGB images were not black and white but red and black, for the case of the red channel reconstruction, green and black for the green case and blue and black for the remaining case. This procedure was adopted to boost the independent influence of the RGB colors in the scene on their respective measurements.

The final image corresponds to the combination of the three RGB images in order to create the colored image (see Figure 17b).

Even though it is possible to recognize the red squared contour and the green background, the result is not very defined. Besides the reduced size of the painted area, we believe that the low reflectivity of the paper surface and the difference between the projected colors and those present in the imaged object strongly conditioned the results. It must also be emphasized that the contours are not sharp.

In the same manner, we tried to reconstruct a color image of the real scene depicted on Figure 15 (a). The obtained results are presented in Figure 18. In this case the quality of the reconstruction is significantly better and this is mainly due to the object's materials, which are more reflective than paper.

Conclusions

In this paper an analysis of the theory of compressive sensing was presented towards the implementation of an innovative active single-pixel camera. We have presented experimental results for a flexible single-pixel CS imaging architecture based on an active illumination concept that yielded good results.

For certain applications, an active illumination single-pixel camera can represent a good way for reducing deployment complexity and costs, while increasing the performance and capabilities of data acquisition and processing systems.

In the light of the work we have already developed, our intent is to extend our work to the development of other configurations that will disregard the need of an active illumination source and operate in a transmissive mode rather than reflective. It is also within our agenda the development of a single-pixel camera with local programmable gain, which could be used to detect weak signals in the presence of very strong optical signals in an image. Such solutions can find applications in various domains such as biology [23], non-destructive inspection [22, 24], security (terahertz imaging for drug detection [22]) and astronomy [25-27].

CS imaging can also be further extended to statistical inference related tasks, such as detection, classification and recognition, since the image reconstruction is not explicitly required, but only the relevant statistics for the problem at hand.

Acknowledgements

This work was supported by a PhD grant (SFRH/BD/45380/2008) from Fundação para Ciência e Tecnologia. The author would also like to acknowledge Fundação Calouste Gulbenkian for the financial support provided.

References

1. E. J. Candès and M. B. Wakin, "An Introduction To Compressive Sampling," Signal Processing Magazine, IEEE **25**, 21-30 (2008).
2. L. K. Kirk, "Rethinking signal processing," Commun. ACM **52**, 13-15 (2009).
3. J. F. Claerbout and F. Muir, "Robust modeling with erratic data," Geophysics **38**, 826-844 (1973).
4. S. Fadil and W. S. William, "Linear Inversion of Band-Limited Reflection Seismograms," (SIAM, 1986), pp. 1307-1330.
5. E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," Information Theory, IEEE Transactions on **52**, 489-509 (2006).
6. E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," Communications on Pure and Applied Mathematics **59**, 1207-1223 (2006).
7. D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed Compressive Sensing [preprint]," (2005).
8. M. Golbabaee and P. Vandergheynst, "Distributed Compressed Sensing for Sensor Networks, Using p-thresholding," (HAL - CCSD, 2009).
9. J. Haupt and R. Nowak, "Signal Reconstruction From Noisy Random Projections," Information Theory, IEEE Transactions on **52**, 4036-4048 (2006).
10. M. Lustig, D. L. Donoho, and J. M. Pauly, "Rapid MR imaging with compressed sensing and randomly under-sampled 3DFT trajectories," in *14th Ann. Meeting ISMRM*, 2006),
11. D. Takhar, J. N. Laska, M. B. Wakin, M. F. Duarte, D. Baron, S. Sarvotham, K. F. Kelly, and R. G. Baraniuk, "A new compressive imaging camera architecture using optical-domain compression," in *Computational Imaging IV*, (SPIE, 2006), 606509-606510.
12. E. J. Candès and T. Tao, "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?," Information Theory, IEEE Transactions on **52**, 5406-5425 (2006).
13. E. J. Candès and T. Tao, "Decoding by Linear Programming," IEEE Transactions on Information Theory **51**, 4203-4215 (2005).
14. E. J. Candès, "Compressive sampling," in *Int. Cong. Mathematicians*, (Madrid, Spain, 2006), pp. 1433-1452.
15. R. G. Baraniuk, "Compressive Sensing [Lecture Notes]," Signal Processing Magazine, IEEE **24**, 118-121 (2007).
16. E. J. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," Inverse Problems **23**, 969-985 (2007).

17. S. Chen, D. Donoho, and M. Saunders, "Atomic Decomposition by Basis Pursuit," SIAM Review **43**, 129-159 (2001).
18. J. A. Tropp and A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," Information Theory, IEEE Transactions on **53**, 4655-4666 (2007).
19. M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, S. Ting, K. F. Kelly, and R. G. Baraniuk, "Single-Pixel Imaging via Compressive Sampling," Signal Processing Magazine, IEEE **25**, 83-91 (2008).
20. R. Baraniuk and K. Kelly, "Lecture 5: A Single-Pixel Compressive Camera ", retrieved http://ima.umn.edu/2006-2007/ND6.4-15.07/activities/Baraniuk-Richard/baraniuk-IMA-CScamera-june07.pdf.
21. S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," Signal Processing Magazine, IEEE **20**, 21-36 (2003).
22. K. Kawase, Y. Ogawa, Y. Watanabe, and H. Inoue, "Non-destructive terahertz imaging of illicit drugs using spectral fingerprints," Opt. Express **11**, 2549-2554 (2003).
23. R. A. Hoebe, C. H. Van Oven, T. W. J. Gadella, P. B. Dhonukshe, C. J. F. Van Noorden, and E. M. M. Manders, "Controlled light-exposure microscopy reduces photobleaching and phototoxicity in fluorescence live-cell imaging," Nat Biotech **25**, 249-253 (2007).
24. W. L. Chan, K. Charan, D. Takhar, K. F. Kelly, R. G. Baraniuk, and D. M. Mittleman, "A single-pixel terahertz imaging system based on compressed sensing," Applied Physics Letters **93**, 121105-121105-121103 (2008).
25. Bobin, J, Starck, and L. J, *Compressed Sensing in astronomy and remote sensing: a data fusion perspective* (Society of Photo-Optical Instrumentation Engineers, Bellingham, WA, ETATS-UNIS, 2009), Vol. 7446, p. 1 vol.
26. J. Bobin, J. L. Starck, and R. Ottensamer, "Compressed Sensing in Astronomy," Selected Topics in Signal Processing, IEEE Journal of **2**, 718-726 (2008).
27. H. G. E. Joachim, "On compressive sensing applied to radar," Signal Process. **90**, 1402-1414 (2010).

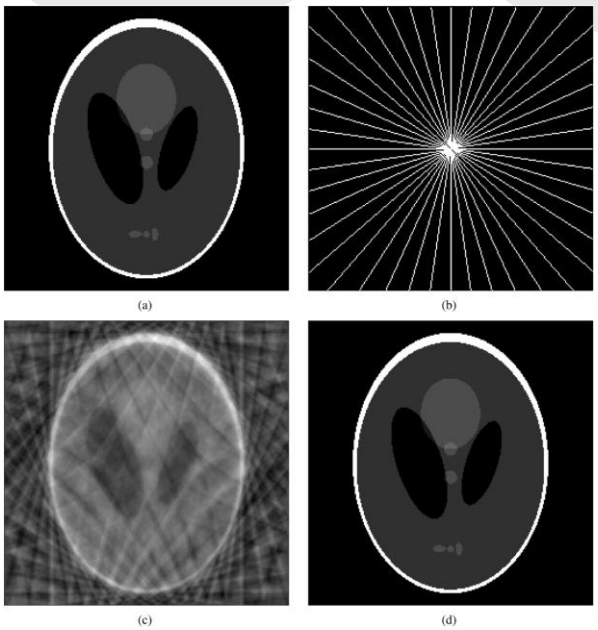

Fig. 1. Example of a simple recovery problem. (a) The Logan–Shepp phantom test image. (b) Sampling domain in the frequency plane; Fourier coefficients are sampled along 22 approximately radial lines. (c) Minimum energy reconstruction obtained by setting unobserved Fourier coefficients to zero. (d) Compressive sensing based reconstruction. This reconstruction is an exact replica of the image in (a) [5].

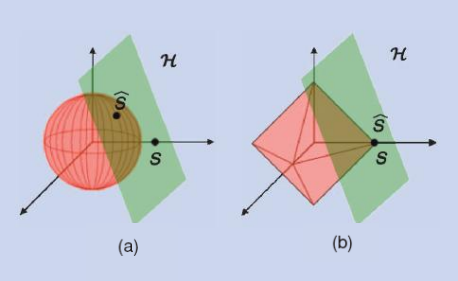

Fig. 2. (Color online) Geometry of ℓ_1 recovery. (a) Visualization of the ℓ_2 minimization that does not find the sparse point of contact $\hat{s}$ between the ℓ_2 ball (hypersphere, in red) and the translated measurement matrix null space (in green). (b) Visualization of the ℓ_1 minimization solution that finds the sparse point of contact $\hat{s}$ with high probability thanks to the pointiness of the ℓ_1 ball [15].

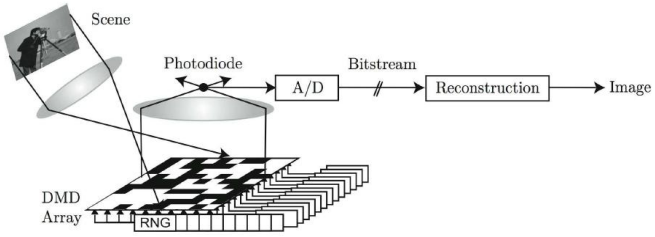

Fig. 3. Single-Pixel Camera block-diagram. Incident light field (corresponding to the desired image x) is reflected off a DMD array whose mirror orientations are modulated by a pseudorandom pattern. Each different mirror pattern produces a voltage at the single photodiode

that corresponds to one measurement $y[m]$. From M measurements a sparse approximation to the desired image x using CS techniques can be obtained [11].

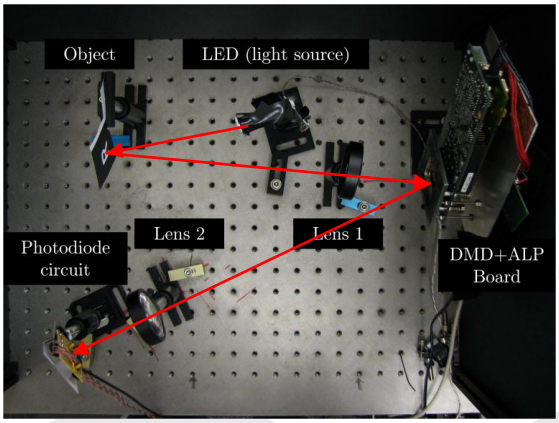

Fig. 4. (Color online) Optical setup of the single-pixel camera [11].

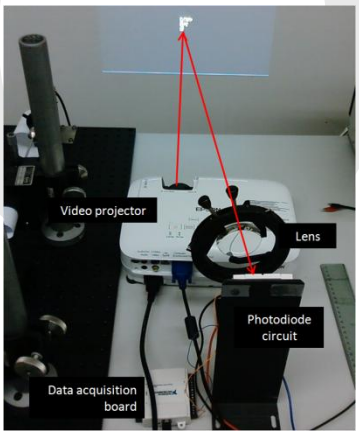

Fig. 5. (Color online) Active illumination single-pixel-camera experimental setup. Following the red arrows, it can be seen that the image projected by the video projector is reflected on the wall and by the means of a lens is focused on the photodiode active area. The output of the photodiode amplifier circuit is connected to a data acquisition board.

(a) 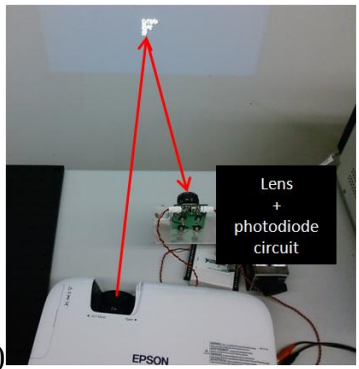

(b) 

Fig. 6. (Color online) (a) Compact active illumination single-pixel camera setup. (b) Detailed photo of the assembly comprising the lens and the photodiode circuit.

Fig. 7. Example of one of the projected images, representing the product between a random measurement pattern and the image to be reconstructed.

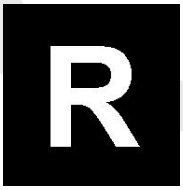

Fig. 8. Ideal image with 64 × 64 pixels (N = 4096).

(a)  (b) 

Fig. 9. Images (64 × 64 pixels) reconstructed using the best K-term Haar wavelet approximation: (a) K = 400; (b) K = 675 [11].

(a) 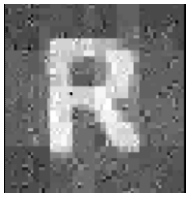 (b) 

Fig. 10. CS reconstruction from: (a) M=820 measurements ⇒ 20%; (b) M = 1600 measurements ⇒ 39% [11].

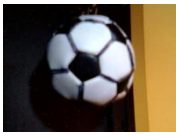  

Fig. 11. (Color online) (left) Original object (ball); (center) 4096 pixels (800 measurements ⇒ 20%); (right) 4096 pixels (1600 measurements ⇒ 40%) (pictures taken from http://dsp.rice.edu/cscamera).

  

Fig. 12. (Color online) (left) Original object (Mandril); (center) 4096 pixels (800 measurements ⇒ 20%); (right) 4096 pixels (1600 measurements ⇒ 40%) (pictures taken from http://dsp.rice.edu/cscamera). This image was reconstructed using RGB color filters to separately acquire each channel and then combine them.

(a) 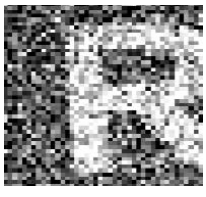 (b) 

Fig. 13. (a) 64 × 64 pixels (N=4096) raster scan image obtained with 4096 measurements.

(b) 64 × 64 pixels image reconstructed via CS from 2700 measurements [20].

Fig. 14. First results obtained (32×32 pixels ⇒ N = 1024) with the active illumination single-pixel camera. (left) 205 measurements ⇒ 20% (PSNR = 11.08 dB); (center) 410 measurements ⇒ 40% (PSNR = 12.30 dB); (right) 717 measurements ⇒ 70% (PSNR = 13.21 dB).

(a) 

(b) 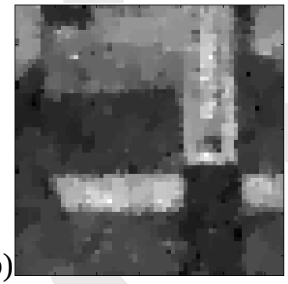

(c) 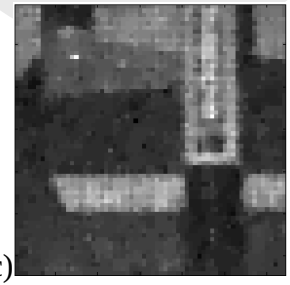

(d) 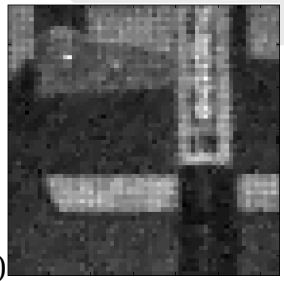

Fig. 15. (Color online) (a) Original scene; Image reconstruction using: (b) 20% of the measurements (PSNR = 69.74 dB); (c) 40% of the measurements (PSNR = 75.60 dB); (d) 60% of the measurements. All the reconstructions are images with 64 × 64 pixels (N=4096). All the PSNR values were obtained comparing the respective image with the image reconstructed using 60% of the measurements.

(a) 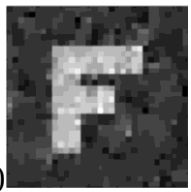 (b) 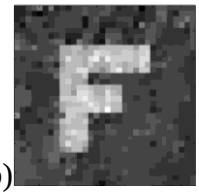

Fig. 16. Reconstruction of the central image of Fig. 14. after the addition of uniformly distributed noise with maximum amplitude of: (a) 10% of the maximum amplitude of the measured signal (SNR = 20.63 dB); (b) 20% of the maximum amplitude of the measured signal (SNR = 14.54 dB).

(a) 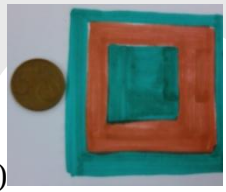 (b) 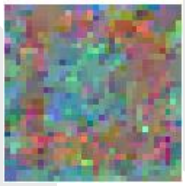

Fig. 17. (Color online) (a) Piece of paper with the painted red contour and green background (the coin is there only for size comparison). (b) Reconstructed color image of the painted area in (a).

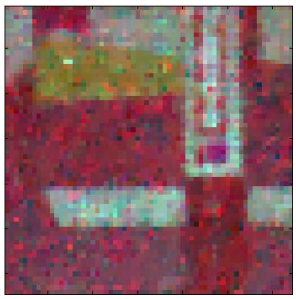

Fig. 18. (Color online) 64 × 64 pixels color reconstruction of the real scene depicted in Figure 15 (a). 40% of the measurements were used to reconstruct the image for each of the RGB channels.