Optical Engineering

[SPIDigitalLibrary.org/oe](https://spiedigitallibrary.org/oe)

Compressed sensing for practical optical imaging systems: a tutorial

Rebecca M. Willett
Roummel F. Marcia
Jonathan M. Nichols

Compressed sensing for practical optical imaging systems: a tutorial

Rebecca M. Willett
Duke University
Durham, North Carolina 27708
E-mail: willett@duke.edu

Roummel F. Marcia
University of California, Merced
Merced, California 95343

Jonathan M. Nichols
Naval Research Laboratory
Washington, DC 20375

Abstract. The emerging field of compressed sensing has potentially powerful implications for the design of optical imaging devices. In particular, compressed sensing theory suggests that one can recover a scene at a higher resolution than is dictated by the pitch of the focal plane array. This rather remarkable result comes with some important caveats however, especially when practical issues associated with physical implementation are taken into account. This tutorial discusses compressed sensing in the context of optical imaging devices, emphasizing the practical hurdles related to building such devices, and offering suggestions for overcoming these hurdles. Examples and analysis specifically related to infrared imaging highlight the challenges associated with large format focal plane arrays and how these challenges can be mitigated using compressed sensing ideas. © 2011 Society of Photo-Optical Instrumentation Engineers (SPIE). [DOI: 10.1117/1.3596602]

Subject terms: compressed sensing; sampling; image reconstruction; inverse problems; computational imaging; infrared; coded aperture imaging; optimization.

Paper 100978TR received Nov. 29, 2010; revised manuscript received May 8, 2011; accepted for publication May 12, 2011; published online Jul. 6, 2011.

1 Introduction

This tutorial describes new methods and computational imagers for increasing system resolution based on recently developed compressed sensing (CS, also referred to as compressive sampling)^{1,2} techniques. CS is a mathematical framework with several powerful theorems that provide insight into how a high resolution image can be inferred from a relatively small number of measurements using sophisticated computational methods. For example, in theory a 1 mega-pixel array could potentially be used to reconstruct a 4 mega-pixel image by projecting the desired high resolution image onto a set of low resolution measurements (via spatial light modulators, for instance) and then recovering the 4 mega-pixel scene through sparse signal reconstruction software. However, it is not immediately clear how to build a practical system that incorporates these theoretical concepts. This paper provides a tutorial on CS for optical engineers which focuses on 1. a brief overview of the main theoretical tenets of CS, 2. physical systems designed with CS theory in mind and the various tradeoffs associated with these systems, and 3. an overview of the state-of-the-art in sparse reconstruction algorithms used for CS image formation. There are several other tutorials on CS available in the literature which we highly recommend;³⁻⁷ however, these papers do not address important technical issues related to optical systems, including a discussion of the tradeoffs associated with non-negativity, photon noise, and the practicality of implementation in real imaging systems.

Although the CS theory and methods we describe in this paper can be applied to many general imaging systems, we concentrate on infrared (IR) technology as a specific example to highlight the challenges associated with applying CS to practical optical systems and to illustrate its potential benefits for improving system resolution. Much of the research and

development in IR imaging is driven by a continued desire for high resolution, large-format focal plane arrays (FPAs). The push for high quality, wide field-of-view IR imagery is particularly strong in military applications. Modern maritime combat, for example, requires consideration of small, difficult-to-detect watercraft operating in low-light environments or in areas where obscurants such as smoke or marine-layer haze make conventional imaging difficult. While high resolution IR imaging can meet these challenges, there is a considerable cost in terms of power, form factor, and finance, in part due to higher costs of IR detector materials and material processing challenges when constructing small pitch FPAs. Despite recent advances,⁸⁻¹⁰ copious research funds continue to be spent on developing larger format arrays with small pitch. CS addresses the question of whether it is necessary to physically produce smaller sensors in order to achieve higher resolution. Conventional practice would require one sensor (e.g., focal plane array element) per image pixel; however, the theory, algorithms, and architectures described in this paper may alleviate this constraint.

1.1 Paper Structure and Contribution

In Sec. 2, we introduce the key theoretical concepts of CS with an intuitive explanation of how high resolution imagery can be achieved with a small number of measurements. Practical architectures which have been developed to exploit CS theory and associated challenges of photon efficiency and noise, non-negativity, and dynamic range are described in Sec. 3. Section 4 describes computational methods designed for inferring high resolution imagery from a small number of compressive measurements. This is an active research area, and we provide a brief overview of the broad classes of techniques used and their tradeoffs. These concepts are then brought together in Sec. 5 with the description of a physical system in development for IR imaging and how CS impacts the contrast of bar target images commonly used to assess

the resolution of IR cameras. We offer some brief concluding remarks in Sec. 6.

2 Compressed Sensing

The basic idea of CS theory is that when the image of interest is very sparse or highly compressible in some basis (i.e., most basis coefficients are small or zero-valued), relatively few well-chosen observations suffice to reconstruct the most significant nonzero components. In particular, judicious selection of the type of image transformation introduced by measurement systems may dramatically improve our ability to extract high quality images from a limited number of measurements. In this section we review the intuition and theory underlying these ideas. By designing optical sensors to collect measurements of a scene according to CS theory, we can use sophisticated computational methods to infer critical scene structure and content.

2.1 Underlying Intuition

One interpretation for why CS is possible is based on the inherent compressibility of most images of interest. Since images may be stored with modern compression methods using fewer than one bit per pixel, we infer that the critical information content of the image is much less than the number of pixels times the bit depth of each pixel. Thus rather than measuring each pixel and then computing a compressed representation, CS suggests that we can measure a “compressed” representation directly.

More specifically, consider what would happen if we knew ahead of time that our scene consisted of one bright point source against a black background. Conventional measurement of this scene would require an FPA with N elements to localize this bright spot with accuracy $1/N$, as depicted in Fig. 1(a). Here, I_j is an indicator image for the j 'th pixel on a $\sqrt{N} \times \sqrt{N}$ grid, so that x_j is a direct measurement of the j 'th pixel's intensity. In other words, the point spread function being modeled is a simple Dirac delta function at the resolution of the detector. If the one bright spot in our image is in pixel k , then x_k will be proportional to its brightness, and the remaining x_j 's will be zero-valued.

However, our experience with binary search strategies and group testing suggests that, armed with our prior knowledge about the sparsity of the signal, we should be able to determine the location of the bright pixel with significantly fewer than N measurements. For instance, consider the binary sensing strategy depicted in Fig. 1(b). The first measurement immediately narrows down the set of possible locations for our bright pixel to half its original size, and the second measurement reduces the size of this set by half again. Thus with only $M = \log_2 N$ measurements of this form, we may accurately localize our bright source.

It is easiest to see that this approach should work in the setting where there is only one nonzero pixel in the original scene and measurements are noise-free. Compressed sensing provides a mechanism for transforming this intuition into settings with noisy measurements where (i) the image contains a small, unknown number ($K \ll N$) of nonzero pixels or (ii) the image contains significantly more structure, such as texture, edges, boundaries, and smoothly varying surfaces, but can be approximated accurately with $K \ll N$ nonzero coefficients in some basis (e.g., a wavelet basis). Case (i) is depicted in Fig. 1(c), and case (ii) is depicted

Fig. 1 Potential imaging modes. (a) Each detector in a focal plane array measures the intensity of a single pixel. This corresponds to a conventional imaging setup. For an N -pixel image, we would require N elements in the FPA. (b) Binary system for noise-free sensing of an image known to have only one nonzero pixel. The first measurement, x_1 , would indicate which half of the imaging plane contains the nonzero pixel. The second measurement, x_2 , combined with x_1 , narrows down the location of the nonzero pixel to one of the four quadrants. To localize the nonzero element to one of N possible locations, only $M = \log_2 N$ binary measurements of this form are required. (c) An extension of the binary sensing system to images which may contain more than one non-zero pixel. Each measurement is the inner product between the image and a binary, possibly pseudorandom, array. Compressed sensing says that if M , the number of measurements, is a small multiple of $\log_2 N$, the image can be accurately reconstructed using appropriate computational tools. (d) Similar concepts hold when the image is sparse or compressible in some orthonormal basis, such as a wavelet basis.

in Fig. 1(d). The CS approach requires that the binary projections be constructed slightly differently than those considered in Fig. 1(b), e.g., using random matrices, but each measurement is nevertheless the (potentially weighted) sum of a subset of the original pixels.

2.2 Underlying Theory

The above concepts are formalized in this section. We only provide a brief overview of a few main theoretical results in this burgeoning field and refer readers to other tutorials³⁻⁷ for additional details. Consider an N -pixel image (which we represent as a length- N column vector) f^* , represented in terms of a basis expansion with N coefficients:

$$f^* = W\theta^* = \sum_{i=1}^N \theta_i^* w_i,$$

where w_i is the i 'th basis vector and θ_i^* is the corresponding coefficient. In many settings, the basis $W \triangleq [w_1, \dots, w_N]$ can be chosen so that only $K \ll N$ coefficients have significant magnitude, i.e., many of the θ_i^* 's are zero or very small for large classes of images; we then say that $\theta^* \triangleq [\theta_1^*, \dots, \theta_N^*]^T$ is sparse or compressible. Sparsity (or, more generally, low-dimensional structure) has long been recognized as a highly useful metric in a variety of inverse problems, but much of the underlying theoretical support was lacking. More recent theoretical studies have provided strong justification for the use of sparsity constraints and

quantified the accuracy of sparse solutions to underdetermined systems.^{11,12}

In such cases, it is clear that if we knew which K of the θ_i^* 's were significant, we would ideally just measure these K coefficients directly, resulting in fewer measurements to obtain an accurate representation of f^* . Of course, in general we do not know *a priori* which coefficients are significant. The key insight of CS is that, with slightly more than K well-chosen measurements, we can determine which θ_i^* 's are significant and accurately estimate their values. Furthermore, fast algorithms which exploit the sparsity of θ^* make this recovery computationally feasible.

The data collected by an imaging or measurement system are represented as

$$x = Rf^* + n = RW\theta^* + n, \quad (1)$$

where $R \in \mathbb{R}^{M \times N}$ linearly projects the scene onto an M -dimensional set of observations, $n \in \mathbb{R}^M$ is noise associated with the physics of the sensor, and $x \in \mathbb{R}_+^M$ is the observed data. (Typically n is assumed bounded or bounded with high probability to ensure that x , which is proportional to photon intensities, is non-negative.) Sparse recovery algorithms address the problem of solving for f^* when the number of unknowns, N , is much larger than the number of observations, M . In general, this is an ill-posed problem as there are a possibly infinite number of candidate solutions for f^* ; nevertheless, CS theory provides a set of conditions on R , W , and f^* that, if satisfied, assure an accurate estimation of f^* .

Much of the CS literature revolves around determining when a sensing matrix $A \triangleq RW$ allows accurate reconstruction using an appropriate algorithm. One widely used property used in such discussions is the restricted isometry property (RIP):

Definition 2.1 (Restricted Isometry Property¹¹): The matrix A satisfies the restricted isometry property of order K with parameter $\delta_K \in [0, 1)$ if

$$(1 - \delta_K)\|\theta\|_2^2 \leq \|A\theta\|_2^2 \leq (1 + \delta_K)\|\theta\|_2^2$$

holds simultaneously for all sparse vectors θ having no more than K nonzero entries. Matrices with this property are denoted $\text{RIP}(K, \delta_K)$.

[In the above, $\|\theta\|_2 \triangleq (\sum_{i=1}^N \theta_i^2)^{1/2}$.] For example, if the entries of A are independent and identically distributed according to

$$A_{i,j} \sim \mathcal{N}\left(0, \frac{1}{M}\right)$$

or

$$A_{i,j} = \begin{cases} M^{-1/2} & \text{with probability } 1/2 \\ -M^{-1/2} & \text{with probability } 1/2 \end{cases},$$

then A satisfies $\text{RIP}(K, \delta_K)$ with high probability for any integer $K = O(M/\log N)$.^{11,13,14} Matrices which satisfy the RIP combined with sparse recovery algorithms are guaranteed to yield accurate estimates of the underlying function f^* , as specified by the following theorem. [As shown below, $\|\theta\|_1 \triangleq \sum_{i=1}^N |\theta_i|$]

Theorem 2.2 (Noisy Sparse Recovery with RIP Matrices^{3,15}): Let A be a matrix satisfying $\text{RIP}(2K, \delta_{2K})$

with $\delta_{2K} < \sqrt{2} - 1$, and let $x = A\theta^* + n$ be a vector of noisy observations of any signal $\theta^* \in \mathbb{R}^N$, where n is a noise or error term with $\|n\|_2 \leq \epsilon$. Let θ_K^* be the best K -sparse approximation of θ^* ; that is, θ_K^* is the approximation obtained by keeping the K largest entries of θ^* and setting the others to zero. Then the estimate

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^N} \|\theta\|_1 \quad \text{subject to } \|x - A\theta\|_2 \leq \epsilon, \quad (2)$$

obeys

$$\|\theta^* - \hat{\theta}\|_2 \leq C_{1,K}\epsilon + C_{2,K} \frac{\|\theta^* - \theta_K^*\|_1}{\sqrt{K}},$$

where $C_{1,K}$ and $C_{2,K}$ are constants which depend on K but not on N or M .

In other words, the accuracy of the reconstruction of a general image f^* from measurements collected using a system which satisfies the RIP depends on a. the amount of noise present and b. how well f^* may be approximated by an image sparse in W . In the special case of *noiseless acquisition of K -sparse signals*, we have $\epsilon = 0$ and $\theta^* \equiv \theta_K^*$, and Theorem 2.2 implies exact recovery so that $\theta^* = \hat{\theta}$.

Note that if the noise were Gaussian white noise with variance σ^2 , then $\|n\|_2 \approx \sigma\sqrt{N}$, so the first term in the error bound in Theorem 2.2 scales like $\sigma\sqrt{N}$. If the noise is Poisson, corresponding to a low-light or infrared setting, then $\|n\|_2$ may be arbitrarily large; theoretical analysis in this setting must include several physical constraints not considered above but described in Sec. 3.2. Additional theory on reconstruction accuracy, optimality, stability with respect to noise, and the uniqueness of solutions is prevalent in the literature.^{3,12,16} Finally, note that the reconstruction (2) in Theorem 2.2 is equivalent to

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^N} \frac{1}{2} \|x - A\theta\|_2^2 + \tau \|\theta\|_1 \quad \text{and} \quad \hat{f} = W\hat{\theta}, \quad (3)$$

where $\tau > 0$ is a regularization parameter which depends on ϵ . Methods for computing this and related formulations are described in Sec. 4.

A related criteria for determining the quality of the measurement matrix for CS is the worst-case *coherence* of $A \equiv RW$.¹⁷⁻¹⁹ Formally, denote the Gram matrix $G \triangleq A^T A$ when the columns of A have unit norm and let

$$\mu(A) \triangleq \max_{\substack{1 \leq i, j \leq N \\ i \neq j}} |G_{i,j}| \quad (4)$$

be the largest off-diagonal element of the Gram matrix. A good goal in designing a sensing matrix is to therefore choose R and W so that μ is as close as possible to $M^{-1/2}$. There also exist matrices A associated with highly overcomplete dictionaries where $N \approx M^2$ and $\mu(A) \approx 1/\sqrt{M}$.²⁰ Several recent works have examined the recovery guarantees of CS as a function of $\mu(A)$,^{17,21-26} such as the following:

Theorem 2.3 (Noisy Sparse Recovery with Incoherent Matrices²¹): Let $x = A\theta^* + n$ be a vector of noisy observations of any K -sparse signal $\theta^* \in \mathbb{R}^N$, where $K \leq [\mu(A)^{-1} + 1]/4$ and n is a noise or error term with $\|n\|_2 \leq \epsilon$. Then the estimate in Eq. (2) obeys

$$\|\theta^* - \hat{\theta}\|_2 \leq \frac{4\epsilon^2}{1 - \mu(A)(4K - 1)}.$$

One of the main practical advantages of coherence-based theory is that it is possible to compute μ for a given CS system. Furthermore, Gersgorin's circle theorem²⁷ states that if a matrix has coherence μ then it satisfies RIP of order (k, ϵ) with $k \sim \epsilon/\mu$; this fact was used in several papers analyzing the performance of CS.^{19,28,29}

Many of the ideas connected to CS, such as using the ℓ_1 norm during reconstruction to achieve sparse solutions, existed in the literature long before the central theory of CS was developed. However, we had limited fundamental insight into why these methods worked, including necessary conditions which would guarantee successful image recovery. CS theory provides significant inroads in this direction. In particular, we now understand how to assess when a sensing matrix will facilitate accurate sparse recovery, and we can use this insight to guide the development of new imaging hardware. Given a particular CS system, governed by $A = RW$, let the ratios $\rho = K/M$ and $\delta = M/N$, respectively, quantify the degree of sparsity and the degree to which the problem is underdetermined. It has been shown that for many CS matrices, there exist sharp boundaries in ρ, δ space that clearly divide the "solvable" from "unsolvable" problems in the noiseless case.^{30,31} This boundary is shown in Fig. 2 for the case of a random Gaussian CS matrix (entries of A are drawn independently from a Gaussian distribution); however it has been shown by Donoho and Tanner³⁰ that this same boundary holds for many other CS matrices as well. Above the boundary, the system lacks sparsity and/or is too underdetermined to solve; below the boundary, solutions are readily obtainable by solving Eq. (2) with $\epsilon = 0$.

Finally, it should be noted that the matrix R is a model for the propagation of light through an optical system. The reconstruction performance is thus going to depend not only on the RIP or incoherence properties of R , but also on modeling inaccuracies due to misalignments, calibration, diffraction, sensor efficiency and bit depth, and other practical challenges. These model inaccuracies can often be incorporated in the noise term for the theoretical analysis above, and typically play a significant role in determining overall system performance.

Fig. 2 For a CS matrix A with Gaussian distributed entries, the boundary separates regions in the problem space where (2) can and cannot be solved (with $\epsilon = 0$). Below this curve solutions can be readily obtained, above the curve they cannot.

3 CS Imaging Systems and Practical Challenges

To date there have been several CS imaging devices built and tested in the laboratory. In general, the central challenge addressed by these methods is to find an architecture which effectively balances between a. physical considerations such as size and cost, b. reconstruction accuracy, and c. reconstruction speed. We review several recently proposed architectures in this section and outline some of the challenges associated with non-negativity and photon noise in this section.

3.1 Imaging Systems

Developing practical optical systems to exploit CS theory is a significant challenge being explored by investigators in the signal processing, optics, optimization, astronomy, and coding theory communities. In addition to implicitly placing hard constraints on the nature of the measurements which can be collected, such as non-negativity of both the projection vectors and the measurements, practical CS imaging systems must also be robust and reasonably sized. Neifeld and Ke³² describe three general optical architectures for compressive imaging: 1. sequential, where measurements are taken one at a time, 2. parallel, where multiple measurements are taken simultaneously using a fixed mask, and 3. photon-sharing, where beam-splitters and micromirror arrays are used to collect measurements. Here, we describe some optical hardware with these architectures that have been recently considered in literature.

3.1.1 Rice single-pixel camera

Perhaps the most well-known example of a CS imager is the rice single-pixel camera developed by Duarte et al.^{33,34} and the more recent single-pixel microscope.³⁵ This architecture uses only a single detector element to image a scene. A digital micromirror array is used to represent a pseudorandom binary array, and the scene of interest is then projected onto that array before the aggregate intensity of the projection is measured with a single detector. Since the individual orientations of the mirrors in the micromirror array can be altered very rapidly, a series of different pseudorandom projections can be measured successively in relatively little time. The original image is then reconstructed from the resulting observations using CS reconstruction techniques such as those described in Section 4. One of the chief benefits of this architecture is that any binary projection matrix can readily be implemented in this system, so that existing CS theory can be directly applied to the measurements. The drawback to this architecture is that one is required to keep the camera focused on the object of interest until enough samples have been collected for reconstruction. The time required may be prohibitive in some applications. Although we can rapidly collect many projections sequentially at lower exposure, this increases the amount of noise per measurement, thus diminishing its potential for video imaging applications.

3.1.2 Coded aperture imagers

Marcia and Willett,³⁶ Marcia et al.³⁷ Romberg,³⁸ and Stern and Javidi³⁹ propose practical implementations of CS ideas using coded apertures, demonstrating that if the coded apertures are designed using a pseudorandom construction, then the resulting observation model satisfies the RIP. Furthermore, the resulting sensing matrix R has a Toeplitz structure that allows for very fast computation within reconstruc-

tion algorithms, providing a significant speedup over random matrix constructions.⁴⁰ The “random lens imaging” optical system⁴¹ is another parallel architecture that is highly suitable for practical and implementable compressive imaging since it provides a snapshot image (i.e., all M measurements are collected simultaneously) and does not require complex, and potentially large, imaging apparatuses.

3.1.3 CMOS CS imagers

Both Robucci et al.⁴² and Majidzadeh et al.⁴³ have proposed performing an analog, random convolution step in complementary, metal-oxide-semiconductor (CMOS) electronics. A clear advantage to this architecture is that the additional optics required for spatial light modulation are removed in favor of additional circuitry. In general, this seems to be a wise trade-off to make, considering the immediate reduction in imager size. The device described by Robucci et al.⁴² also leverages the ability of CMOS electronics to perform fast, block-wise inner products between the incoming data and a pre-defined random sequence. In the cited work the authors used a noiselet basis with binary $\{1, -1\}$ coefficients for the projection operation, however, the architecture is extremely flexible and could admit many other choices (e.g., discrete cosine transform). The CMOS implementation is also likely to be low cost, relative to other approaches requiring expensive optical components.

3.1.4 Spectral imagers

Tradeoffs between spectral and spatial resolution limit the performance of modern spectral imagers, especially in photon-limited settings where the small number of photons must be apportioned between the voxels in the data cube, resulting in low signal-to-noise ratio (SNR) per voxel. Gehm et al.⁴⁴ and Wagadarikar et al.⁴⁵ proposed innovative, real-time spectral imagers, where each pixel measurement is the coded projection of the spectrum in the corresponding spatial location in the data cube. This was first implemented using two dispersive elements separated by binary-coded masks; later, simpler designs omitted one dispersive element. In related work, objects are illuminated by light sources with tunable spectra using spatial light modulators to facilitate compressive spectral image acquisition.⁴⁶ Finally, we mention more recent examples of compressive spectral imagers, such as compressive structured light codes where each camera pixel measures light from points along the line of sight within a volume density,⁴⁷ and cameras that use dispersers for imaging piecewise “macropixel” objects (e.g., biochip microarrays in biochemistry).⁴⁸

3.1.5 Application-specific architectures

Compressive imaging can afford other possible advantages besides simply reducing the number of pixels. Shankar et al. used CS to develop an infrared camera that significantly reduced the thickness of the optics (by an order of magnitude).⁴⁹ Coskun et al.⁵⁰ also used CS principles to eliminate optical components (a lens) in a fluorescent imaging application of naturally sparse signals. Other recent works include the application of CS theory to radar imaging⁵¹ and the recovery of volumetric densities associated with translucent media (e.g., smoke, clouds, etc.).⁵² Still, other CS optical systems have also been proposed in a variety of applications including DNA microarrays,^{53,54} magnetic resonance imaging

(MRI),⁵⁵ ground penetrating radar,⁵⁶ confocal microscopy,⁵⁷ and astronomical imaging.⁵⁸

3.2 Non-negativity and Photon Noise

The theory in Sec. 2 described pseudorandom sensing matrices that satisfied the RIP, and hence led to theoretical guarantees on reconstruction accuracy in the presence of Gaussian or bounded noise. However, these sensing matrices were zero mean, so that approximately half of the elements were negative. Such a system is impossible to construct with linear optical elements. In addition, Gaussian or bounded noise models are not appropriate for all optical systems. Finally, we have the (typically not modeled) physical constraint that the total light intensity incident upon our detector cannot exceed the total light intensity entering our aperture; i.e., $\|Rf\|_1 \leq \|f\|_1$. These practical considerations lead to an active area of ongoing research. Several of the physical architectures described above were designed based on zero-mean sensing models, and then subjected to a mean shift to make every element of the sensing matrix non-negative. In high SNR settings, reconstruction algorithms can compensate for this offset, rendering it negligible. This adjustment is critical to the success of many sparse reconstruction algorithms discussed in the literature, which perform best when $R^T R \approx I$. In particular, assume we measure $x_p = R_p f^* + n$, where R_p is defined as $R - \mu_R$ for a zero-mean CS matrix R and $\mu_R \triangleq (\min_{i,j} R_{i,j}) \mathbf{1}_{M \times N}$; every element of R_p is non-negative by definition. Note $x_p = Rf^* + \mu_R f^* + n$. Since $\mu_R f^*$ is a constant vector proportional to the total scene intensity we can easily estimate $z \triangleq \mu_R f^*$ from data and apply sparse reconstruction algorithms to $\hat{x} \triangleq x_p - z \approx Rf^* + n$.³⁷

In low SNR photon-limited settings, however, the combination of non-negative sensing matrices and light intensity preservation present a significant challenge to compressive optical systems. In previous work, we evaluated CS approaches for generating high resolution images and video from photon-limited data.^{59,60} In particular, we showed how a feasible positivity- and flux-preserving sensing matrix can be constructed, and analyzed the performance of a CS reconstruction approach for Poisson data that minimizes an objective function consisting of a negative Poisson log likelihood term and a penalty term which measures image sparsity. We showed that for a fixed image intensity, the error bound actually grows with the number of measurements or sensors. This surprising fact can be understood intuitively by noting that dense positive sensing matrices will result in measurements proportional to the average intensity of the scene plus small fluctuations about that average. Accurate measurement of these fluctuations is critical to CS reconstruction, but in photon-limited settings the noise variance is proportional to the mean background intensity and overwhelms the desired signal.

3.3 Dynamic Range

An important consideration in practical implementations of CS hardware architectures is the quantization of the measurements, which involves encoding the values of x in Eq. (1) in finite-length bit strings. The representation of real-valued measurements as bit strings introduces error in addition to the noise discussed above. Moreover, if the dynamic range of the sensor is limited, very high and low intensity values that

are outside this range will be truncated and simply be given the maximum and minimum values of the quantizer. These measurement inaccuracies can be mitigated by incorporating the quantization distortion within the observation model^{61,62} and by either judiciously rejecting “saturated” measurements or factoring them in as inequality constraints within the reconstruction method.⁶³

4 CS Reconstruction Methods

The ℓ_2 - ℓ_1 CS problem (3) can be solved in a variety of ways. However, many off-the-shelf optimization software packages are often unsuitable because the size of imaging problems is generally prohibitively too large. For instance, second-derivative methods require solving an $N \times N$ linear equation at each iteration of the underlying Newton’s method. In our settings, N corresponds to the number of pixels. Thus, to reconstruct a 1024×1024 pixel image requires solving a $1024^2 \times 1024^2$ linear system at each iteration, which is computationally too expensive and memory intensive. In addition, the ℓ_1 term in the objective function in Eq. (3) is not differentiable; thus, the ℓ_2 - ℓ_1 CS reconstruction problem must be reformulated so that gradient-based methods can be applied. Finally, the ℓ_1 regularization common in CS can be implemented efficiently using simple thresholding schemes within the optimization algorithms; computational shortcuts like this are not exploited in most off-the-shelf software. Current sparse reconstruction methods are designed so that these computational considerations are taken into account.

For example, gradient projection methods^{64,65} introduce additional variables and recast Eq. (3) as a constrained optimization problem with a differentiable objective function. Gradient descent directions, which are generally easy to compute, are used at each iteration, and are then projected onto the constraint set so that each step is feasible. The projection involves only simple thresholding and can be done very quickly, which leads to fast computation at each iteration. Alternatively, methods called iterative shrinkage/thresholding algorithms^{66–69} map the objective function onto a sequence of simpler optimization problems which can be solved efficiently by shrinking or thresholding small values in the current estimate of θ . Another family of methods based on matching pursuits (MP)^{70–72} starts with $\hat{\theta} = 0$ and greedily chooses elements of $\hat{\theta}$ to have nonzero magnitude by iteratively processing residual errors between y and $A\hat{\theta}$. MP approaches are well-suited to settings with little or no noise; in contrast, gradient-based methods are more robust to noisier problems. Additionally, they are generally faster than MP methods, as A is not formed explicitly and is used only for computing matrix-vector products. Finally, while these methods do not have the fast quadratic convergence properties that are theoretically possible with some black-box optimization methods, they scale better with the size of the problem, which is perhaps the most important issue in practical CS imaging.

There is a wide variety of algorithms in the literature and online available for solving Eq. (3) and its variants. Many of these algorithms aim to solve the exact same convex optimization problem, and hence will all yield the same reconstruction result. However, because of specific implementation aspects and design decisions, some algorithms require fewer iterations or less computation time per iteration depending on the problem structure. For instance, some algorithms will exhibit faster convergence with wavelet sparsity regularization, while others may converge more quickly with

a total variation regularizer. Similarly, if the A matrix has a helpful structure (e.g., Toeplitz), then multiplications by A and A^T can be computed very quickly and algorithms which exploit this structure will converge most quickly, while if A is a pseudorandom matrix then algorithms which attempt to limit the number of multiplications by A or A^T may perform better.

4.1 Alternative Sparsity Measures

While most of the early work in CS theory and methods focused on measuring and reconstructing signals which are sparse in some basis, current thinking in the community is more broadly focused on high-dimensional data with underlying low-dimensional structure. Examples of this include conventional sparsity in an orthonormal basis,¹¹ small total variation,^{73–77} a low-dimensional submanifold of possible scenes,^{78–82} a union of low-rank subspaces,^{83,84} or a low-rank matrix in which each column is a vector representation of a small image patch, video frame, or small localized collection of pixels.^{85,86} Each of the above examples has been successfully used to model structure in “natural” images and hence facilitate accurate image formation from compressive measurements (although some models do not admit reconstruction methods based on convex optimization, and hence may be computationally complex or only yield locally optimal reconstructions). A variety of penalties based on similar intuition have been developed from a Bayesian perspective. While these methods can require significant burn-in and computation time, they allow for complex, nonparametric models of structure and sparsity within images, producing compelling empirical results in denoising, interpolating and compressive sampling that are on par with, if not better than, established methods.⁸⁷

4.2 Non-negativity and Photon Noise

In optical imaging, we often estimate light intensity, which *a priori* be non-negative. Thus, it is necessary that the reconstruction $\hat{f} = W\hat{\theta}$ is non-negative, which involves adding constraints to the CS optimization problem Eq. (3), i.e.,

$$\begin{aligned} \hat{\theta} &= \arg \min_{\theta \in \mathbb{R}^N} \frac{1}{2} \|x - RW\theta\|_2^2 + \tau \text{pen}(\theta) \\ &\text{subject to } W\theta \geq 0 \\ \hat{f} &= W\hat{\theta} \end{aligned} \quad (5)$$

where $\text{pen}(\theta)$ is a general sparsity-promoting penalty term. The addition of the nonnegativity constraint in Eq. (5) makes the problem more challenging than the conventional CS minimization problem, and it has been shown that simply thresholding the unconstrained solution so that the constraints in Eq. (5) are satisfied leads to suboptimal estimates.⁸⁸

In the context of low-light settings, where measurements are inherently noisy due to low count levels, the inhomogeneous Poisson process model⁸⁹ has been used in place of the $\|x - RW\theta\|_2^2$ term in Eq. (5). The adaptation of the reconstruction methods described previously can be very challenging in these settings. In addition to enforcing non-negativity, the negative Poisson log likelihood used in the formulation of an objective function often requires the application of relatively sophisticated optimization theory principles.^{90–96}

4.3 Video

Compressive image reconstruction naturally extends to sparse video recovery since videos can be viewed as sequences of correlated images. Video compression techniques can be incorporated into reconstruction methods to improve speed and accuracy. For example, motion compensation and estimation can be used to predict changes within a scene to achieve a sparse representation.⁹⁷ Here, we describe a less computationally intensive approach based on exploiting interframe differences.⁹⁸

Let $f_1^*, f_2^*, \dots$ be a sequence of images comprising a video, and let W be an orthonormal basis in which each f_t^* is sparse, i.e., $f_t^* = W\theta_t^*$, where θ_t^* is mostly zeros. To recover the video sequence $\{f_t^*\}$, we need to solve for each f_t^* . Simply applying Eq. (5) at each time frame works well, but this approach can be improved upon by solving for multiple frames simultaneously. In particular, rather than treating each frame separately, we can exploit interframe correlation and solve for θ_t^* and the difference $\Delta\theta_t^* \triangleq \theta_{t+1}^* - \theta_t^*$ instead of θ_t^* and θ_{t+1}^* . This results in the coupled optimization problem:

$$\begin{aligned} \begin{bmatrix} \hat{\theta}_t \\ \Delta\hat{\theta}_t \end{bmatrix} &= \arg \min_{\theta_t, \Delta\theta_t} \frac{1}{2} \left\| \begin{bmatrix} x_t \\ x_{t+1} \end{bmatrix} - \begin{bmatrix} R_t & 0 \\ 0 & R_{t+1} \end{bmatrix} \begin{bmatrix} W & 0 \\ W & W \end{bmatrix} \begin{bmatrix} \theta_t \\ \Delta\theta_t \end{bmatrix} \right\|_2^2 + \tau_1 \|\theta_t\|_1 + \tau_2 \|\Delta\theta_t\|_1 \\ &\text{subject to } W\theta_t \geq 0, W(\theta_t + \Delta\theta_t) \geq 0, \end{aligned} \quad (6)$$

where $\tau_1, \tau_2 > 0$ and R_t and R_{t+1} are the observation matrices at times t and $t+1$, respectively. When $\theta_{t+1}^* \approx \theta_t^*$, then $\Delta\theta_t^* = \theta_{t+1}^* - \theta_t^*$ is very sparse compared to θ_{t+1}^* , which makes it even better suited to the sparsity-inducing ℓ^1 -penalty terms in Eq. (6). We use different regularization parameters τ_1 and τ_2 to promote greater sparsity in $\Delta\theta_t$. This approach can easily be extended to solve for many frames simultaneously.^{98–100}

4.4 Connection to Super-Resolution

The above compressive video reconstruction method has several similarities with super-resolution reconstruction. With super-resolution, observations are typically not compressive in the CS sense, but rather downsampled observations of a high resolution scene. Reconstruction is performed by: a. estimating any shift or dithering between successive frames, b. estimating the motion of any dynamic scene elements, c. using these motion estimates to construct a sensing matrix R , and d. solving the resulting inverse problem, often with a sparsity-promoting regularization term.

There are two key differences between super-resolution image/video reconstruction and compressive video reconstruction. First, CS recovery guarantees to place specific demands on the size and structure of the sensing matrix; these requirements may not be satisfied in many super-resolution contexts, making it difficult to predict performance or assess optimality. Second, super-resolution reconstruction as described above explicitly assumes that the observed frames consist of a small number of moving elements superimposed upon slightly dithered versions of the same background; thus estimating motion and dithering is essential to accurate reconstruction. CS video reconstruction, however, does not require dithering or motion modeling as long as each frame is sufficiently sparse. Good motion models can improve on reconstruction performance, as shown in Sec. 6, by increasing the amount of sparsity in the variable (i.e., sequence of frames) to be estimated, but even if successive frames are very different, accurate reconstruction is still feasible. This is not the case with conventional super-resolution estimation.

5 Infrared Camera Examples

IR cameras are a particularly promising target for compressive sampling owing in large part to the manufacturing costs associated with the FPA. Demands for large format, high resolution FPAs have historically meant even smaller pixels grown on even larger substrates. For typical materials used in infrared imagers (e.g., HgCdTe, InSb), the manufacturing process is extremely challenging and requires (among other things) controlling material geometry at the pixel-level. Recent developments in FPA technology have resulted in $20\text{ }\mu\text{m}$ pitch FPAs across the infrared spectrum. Yuan et al. (Teledyne Judson Technologies) have recently demonstrated a 1280×1024 , $20\text{ }\mu\text{m}$ pitch InGaAs FPA for the short-wave infrared (SWIR) waveband, 1.0 to $1.7\text{ }\mu\text{m}$.⁸ In the mid-wave infrared, 3 to $5\text{ }\mu\text{m}$, Nichols et al.⁹ used a 2048×2048 , $20\text{ }\mu\text{m}$ pitch InSb FPA (Cincinnati Electronics), while Carmody et al. report on a 640×480 , $20\text{ }\mu\text{m}$ pitch HgCdTe FPA.¹⁰ Despite these advances, numerous research funds continue to be spent on developing even larger format arrays while attempting to decrease the pixel pitch to $\leq 15\text{ }\mu\text{m}$. The ability to improve the resolution of existing FPA technology without physically reducing pixel size, therefore, has some potentially powerful implications for the IR imaging community.

In this section, we explore taking existing focal plane array technology and using spatial light modulation (SLM) to produce what is effectively a higher resolution camera. The specific architecture used will depend on whether one is imaging coherent or incoherent light and on the physical quantity being measured at the detector. For incoherent light, typical imaging systems are linear in intensity and the focal plane array is directly measuring this intensity. Thus, one way to design a CS imager is to use a fixed, coded aperture as part of the camera architecture.^{36,38,42,101} The imager shown schematically at the top of Fig. 3(a) illustrates one such approach, whereby the aperture modulates the light in the image's Fourier plane. Mathematically, this architecture is convolving the desired image intensity f with the magnitude squared of the point spread function $|h|^2 = |\mathcal{F}^{-1}(H)|^2$

Fig. 3 Infrared camera examples. (a) Two possible IR camera architectures. In the first, a coded aperture is placed in the lens Fourier plane, while in the second, two different phase masks are used to create the measurement matrix. (b) Panoramic midwave camera assembly.

where H is the mask pattern and $\mathcal{F}$ denotes the Fourier transform. Low-resolution observations are then obtained by downsampling the impinging light intensity either by integration downsampling at the detector or by subsampling using SLMs. Each downsampling approach has its advantages and disadvantages: integration downsampling allows for more light and hence higher SNR observations, but results in loss of high frequency components and adversely impacts the incoherence of A , whereas subsampling results in much stronger theoretical properties but allows for less light. Thus, depending on the light conditions, one approach might be more suitable than the other. For example, in high photon count regimes, subsampling will more likely yield a more accurate reconstruction. Finally, if the mask is designed so that the resulting observation model satisfies the RIP property and the image is sparse in some basis (e.g., wavelet), then the underdetermined problem can be solved for the high-resolution image f .³⁶ Using this approach, the spatial resolution of the modulation becomes the resolution of the recovered image. That is to say, if we desire a $2\times$ improvement in resolution, our coded aperture would possess $1/2$ the pixel width of the FPA. Heuristically, one can think of the convolution as a mechanism for spreading localized (spatial) image information over all pixels at the detector. One can additionally think of the fixed, known, mask pattern as the unique “key” that allows the reconstruction algorithm to then extract the true image, despite the apparent ambiguity introduced by the detector downsampling. This is precisely what the RIP property means physically for the system architecture, namely that it (a) modulates the image in a manner that allows all of the needed information to be encoded in the compressed samples, yet (b) ensures a unique solution to the true image can be found. Although all the information needed for reconstruction is present in the compressed samples, the light blocked by the aperture decreases the signal-to-noise ratio. One possible solution is to increase the camera integration time though this can result in motion blur artifacts.

Another possibility is to use the architecture shown at the bottom of Fig. 3(a) as proposed by Rivenson et al.¹⁰² (see also Romberg³⁸). This architecture is more appropriate for

imaging coherent light. In this approach, one makes use of Fourier optics to convolve the electromagnetic field associated with the image and a random phase pattern. A second phase modulator, located at the image plane, is necessary in order for the resulting sensing matrix satisfies the RIP.¹⁰² Note that the detector in this architecture needs to be capable of measuring both real and imaginary parts of the complex image (as opposed to image intensity). The potential advantage of this architecture is that it modulates phase only, thus light is not blocked as it is in the coded aperture architecture. However, this camera architecture requires two SLMs, each with its own inherent losses. Furthermore, detectors that can measure the complex image field are far less common than those that measure image intensity.

Regardless of the specific architecture used, one could either consider the coded apertures to be fixed patterns that do not change in time or, given the state of SLM technology, one could dynamically modify the pattern for video data. This could potentially provide a benefit (i.e., help guarantee the RIP) to the multiframe reconstruction methods described in Section 4. Finally, we should point out that in a typical camera the image is not located exactly one focal length away. Rather, the system possesses a finite aperture D , designed in conjunction with the focal length FL to collect imagery at some prescribed range from the camera. In this more general case there will still be a lens Fourier plane at which we can apply our modulation, however, the location of the mask will need to be specified to within $\delta \propto \lambda(FL/D)^2 m$. Diffraction effects may also need to be considered in developing the signal model to be used in solving the $\ell_2 - \ell_1$ reconstruction problem; potential solutions to this problem for conventional coded apertures have been described by Stayman et al.¹⁰³

Central to each of the above-described architectures is the availability of fixed apertures, digital micromirror devices, or SLMs capable of spatially modulating the light at the desired spatial resolution. A fixed coded aperture is clearly the simplest design and can be manufactured easily to have micrometer resolution using chrome on quartz. SLMs allow codes to be changed over time, but current SLM technology restricts the masks to have lower resolution than chrome on

Fig. 4 (a) Original 128×128 bar target image for hot target $T = 45$ C, (b) Same imagery as (a) but downsampled by a factor of 2 to simulate a $30 \mu\text{m}$ pitch camera, (c) Compressively sampled image using a simulated $30 \mu\text{m}$ pitch, (d) Bar target at full resolution recovered from the measurements (c) using Eq. (3), (e) Image slice of the medium resolution bar target image [the lower right set of bars in (a)], (f) Contrast as a function of bar target temperature.

quartz solutions.¹⁰⁴ For example, Boulder Nonlinear Systems manufactures a 512×512 , $15 \mu\text{m}$ pitch device that can operate in the visible, short-wave IR, midwave IR, or long-wave IR wavelengths and can be used to spatially modulate the amplitude, phase, or both.¹⁰⁵ Sections 5.1 and 5.2 illustrate how coded-aperture IR imaging can be used to enhance the resolution of existing IR cameras.

5.1 Midwave Camera Example

As an illustrative example, we consider bar target imagery collected from a midwave infrared camera. The camera, shown in Fig. 3(b), was designed for maritime search and tracking applications. Optically, the camera is a catadioptric, $f/2.5$ aperture system that provides a 360° horizontal and a -10° to $+30^\circ$ elevation field of view. The core of the imaging system is a 2048×2048 pixel, $15 \mu\text{m}$ pitch, cryogenically cooled indium antimonide focal plane array. The sensor is designed to operate in the 3.5 to $5 \mu\text{m}$ spectral band with a CO_2 notch filter near $4.2 \mu\text{m}$. Additional details for the camera are described by Nichols et al.⁹

In order to evaluate the compressive imaging reconstruction algorithm, we collected bar target imagery generated using a CI System extended blackbody source held at fixed temperatures ranging from 5 to 45 C. Four bar patterns of varying spatial frequencies occupied a 128×128 pixel area. Figure. 4(a) shows the image acquired for a bar target temperature of $T = 45$ C. This image was then downsampled by a factor of 2 in order to simulate observations from a

lower resolution, $30 \mu\text{m}$ focal plane array. The downsampled 64×64 image is shown in Fig. 4(b). Low resolution, compressed imagery was also generated by numerically applying a fixed, coded aperture with random, Gaussian distributed entries in the lens Fourier plane prior to downsampling [e.g., simulate the architecture of Fig. 3(a)]. That is to say the degree of transparency for each of the elements in the mask H is chosen to be random draws from a truncated Gaussian distribution spanning the range 0 (no light blocked) to 1 (all light blocked). Each sample collected by a $30 \mu\text{m}$ pixel is therefore modeled as a summation over four, $15 \mu\text{m}$ blocks of the image intensity convolved with the magnitude squared of the point spread function $h = \mathcal{F}^{-1}(H)$. This process is shown at the single pixel level in Fig. 5.

There are certainly multiple images that could give rise to the same compressed sample value. What is remarkable about CS theory is that it tells us that if the image is sparse in the chosen basis and if we modulate the image in accordance with the RIP property, then this ambiguity is effectively removed and the true image recovered. The practical advantage to using this system, as was pointed out in Sec. 3, is that the information needed to recover the high resolution image is encoded in a single low-resolution image. Thus there is no need for multiple shots for the recovery to work. The cost, however, as was already mentioned, is a decrease in signal to noise ratio.

The compressed measurements are shown in Fig. 4(c) and clearly illustrate the influence of the randomly coded aper-

Fig. 5 Pixel-level description of the process by which data at the 15 μm scale are converted to compressed samples at the 30 μm scale. The desired image intensity is convolved with a point spread function dictated by the magnitude of the inverse Fourier Transform of the mask pattern.

ture on the image. Finally, to recover the image, we solved Eq. (3) using the gradient projection for sparse reconstruction (GPSR) algorithm,⁶⁴ which is a gradient-based optimization method that is very fast, accurate, and efficient. The basis used in the recovery was the length-12 Coiflet. We have no reason to suspect this is the optimal basis (in terms of promoting sparsity) for recovery; however among the various wavelet basis tried by the authors, little difference in the results was observed. The image in Fig. 4(d) is the full resolution (128×128) reconstruction.

Qualitatively it can be seen that the reconstruction algorithm correctly captures some of the details lost in the down-sampling. This comparison can be made quantitative by considering a “slice” through the 0.36 (cycle/mrad) resolution bar target (lower right set of bars). Figure 4(e) compares the measured number of photons in the original 45 C bar target image to that achieved by interpolating the downsampled image and the recovered image. Clearly, some of the contrast is lost in the downsampled imagery, whereas the compressively sampled, reconstructed image suffers little loss. Figure 4(f) shows how the estimated image contrast behaves as a function of bar target temperature. Contrast was determined for the bar target pixels as $|\max(f) - \min(f)|/\bar{f}$, where $\bar{f}$ is the mean number of photons collected by the background and

$\max(f)$, $\min(f)$ are, respectively, the largest and smallest number of photons collected across the bar target.

Regardless of sampling strategy, the contrast is greater near the extremes ($T = 5\text{ C}$, $T = 45\text{ C}$) and is lowest near the ambient (room) temperature of $T = 27\text{ C}$ as expected. However, for areas of high contrast we see a factor of 2 improvement by using the compressive sampling strategy over the more conventional bi-cubic interpolation approach. We have also tried linear interpolation of the low-resolution imagery and found no discernible difference in the results. Based on these results we see that a CS camera with half the resolution of a standard camera would be capable of producing the same resolution imagery as the standard camera without any significant reduction in image contrast.

5.2 SWIR Video Example

We now consider the application of compressive coded apertures on video collected by a short-wave IR (0.9 to 1.7 μm) camera. The camera is based on a 1024×1280 InGaAs (indium, gallium arsenide) focal plane array with 20 μm pixel pitch. Optically, the camera was built around a fixed, $f/2$ aperture and provides a 6° field of view along the diagonal with a focus range of $50\text{ m} \rightarrow \infty$. Imagery were output at the standard 30 Hz frame rate with a 14 bit dynamic range. The video used in the following example consists of 25 256×256 pixel frames cropped from the full imagery. All video data were collected at the Naval Research Laboratory, Chesapeake Bay Detachment, in November 2009. We simulate the performance of a low-resolution noncoded video system by simply downsampling the original video by a factor of 2 in each direction. We also simulate the performance of a low-resolution noncoded video system and perform two methods for reconstruction: the first solves for each frame individually while the second solves for two frames simultaneously. For this experiment, we again used the GPSR algorithm. The maximum number of iterations for each method are chosen so that the aggregate time per frame for each method is approximately the same. The 25th frames for the original signal, the reconstruction using the single-frame noncoded method, and the reconstruction using the two-frame coded method are presented in Fig. 6. The root-mean squared error (RMSE) values for this frame for the single-frame noncoded and the two-frame coded methods are 3.78% and 2.03%, respectively.

Fig. 6 SWIR compressive video simulation. (a) Original scene, (b) Reconstruction of the 25th frame without coded apertures (RMSE = 3.78%). (c) Reconstruction using coded apertures (RMSE = 2.03%). Note the increased resolution, resulting in smoother edges and less pixelated reconstruction.

6 Conclusions

This tutorial is aimed at introducing optical engineers to several of the theoretical breakthroughs and practical challenges associated with compressed sensing in optical systems. While many of the theoretical results are promising, in that a relatively small focal plane array can be used to collect high resolution imagery, translating this theory to practice requires careful attention to the tradeoffs between focal plane array size; optical component size, weight, and expense; admissibility of theory in practical systems; and choice of reconstruction method. Through proof-of-concept experiments with a bar target and with video in SWIR systems, we demonstrate how compressed sensing concepts can be used to improve contrast and resolution in practical optical imaging settings.

Acknowledgments

The authors would like to thank Zachary Harmany for sharing his sparse reconstruction algorithms for compressive coded apertures. This work was supported by NSF CAREER Awards CCF-06-43947, DMS-08-11062, DARPA Grant No. HR0011-09-1-0036, NGA Award HM1582-10-1-0002, and AFRL Grant FA8650-07-D-1221.

References

1. E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory* **51**(12), 4203–4215 (2005).
2. D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006).
3. E. Candès, "Compressive sampling," in *Proc. Int. Congress of Mathematicians*, Madrid, Spain Vol. 3, pp. 1433–1452 (2006).
4. R. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.* **24**(4), 118–121 (2007).
5. E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 21–30 (2008).
6. J. Romberg, "Imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 14–20 (2008).
7. M. Fornasier and H. Rauhut, "Compressive sensing," in *Handbook of Mathematical Methods in Imaging*, Springer, Heidelberg, Germany, (2011).
8. H. Yuan, G. Apgar, J. Kim, J. Laquindanum, V. Nalavade, P. Beer, J. Kimchi, and T. Wong, "FPA development: from InGaAs, InSb, to HgCdTe," in *Infrared Technology and Applications XXXIV*, B. F. Andresen, G. F. Fulop, and P. R. Norton, Eds., *Proc. SPIE Defense, Security, and Sensing Symposium* Vol. 6940 (2008).
9. J. M. Nichols, J. R. Waterman, R. Menon, and J. Devitt, "Modeling and analysis of a high performance midwave infrared panoramic periscope," *Opt. Eng.* **49**(11) (2010).
10. M. Carmody, J. G. Pasko, D. Edwall, E. Piquette, M. Kangas, S. Freeman, J. Arias, R. Jacobs, W. Mason, A. Stoltz, Y. Chen, and N. K. Dhar, "Status of LWIR HgCdTe-on-Silicon FPA Technology," *J. Electron. Mater.* **37**(9), 1184–1188 (2008).
11. E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory* **15**(12), 4203–4215 (2005).
12. J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inf. Theory* **52**(9), 4036–4048 (2006).
13. E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory* **52**(2), 489–509 (2006).
14. R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approx.* **28**(3), 253–263 (2008).
15. E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.* **59**(8), 1207–1223 (2006).
16. E. J. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," *Ann. Stat.* **35**, 2313–2351 (2007).
17. J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory* **50**(10), 2231–2242 (2004).
18. E. J. Candès and Y. Plan, "Near-ideal model selection by ℓ_1 minimization," *Ann. Stat.* **37**, 2145–2177 (2009).
19. D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization," *Proc. Natl. Acad. Sci. U.S.A.* **100**(5), 2191–2202 (2003).
20. T. Strohmer and R. W. Heath Jr., "Grassmannian frames with applications to coding and communications," *Appl. Comput. Harmon. Anal.* **14**(3), 257–275 (2003).
21. D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory* **52**(1), 6–18 (2006).
22. D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory* **47**(7), 2845–2862 (2001).
23. D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization," *Proc. Natl. Acad. Sci. U.S.A.* **100**(5), 2197–2202 (2003).
24. R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inf. Theory* **49**(12), 3320–3325 (2003).
25. J. A. Tropp, "Just relax: convex programming methods for identifying sparse signals in noise," *IEEE Trans. Inf. Theory* **52**(3), 1030–1051 (2006).
26. A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss, "Approximation of functions over redundant dictionaries using coherence," in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Baltimore, MD (2003).
27. R. Varga, *Geršgorin and His Circles*, Springer-Verlag, Berlin, Germany (2004).
28. J. A. Tropp, A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss, "Improved sparse approximation over quasincoherent dictionaries," in *IEEE International Conference on Image Processing*, Barcelona, Spain (2003).
29. J. A. Tropp, "On the conditioning of random subdictionaries," *Appl. Comput. Harmon. Anal.* **25**, 1–24 (2008).
30. D. L. Donoho and J. Tanner, "Precise undersampling theorems," *Proc. IEEE* **98**(6), 913–924 (2010).
31. D. L. Donoho and J. Tanner, "Exponential bounds implying construction of compressed sensing matrices, error-correcting codes, and neighborly polytopes by random sampling," *IEEE Trans. Inf. Theory*, **56**(4), 2002–2016 (2010).
32. M. Neifeld and J. Ke, "Optical architectures for compressive imaging," *Appl. Opt.* **46**, 5293–5303 (2007).
33. M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 83–91 (2008).
34. 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," *Appl. Phys. Lett.* **93**, 121105 (2008).
35. Y. Wu, C. Chen, P. Ye, G. R. Arce, and D. W. Prather, "Development of a compressive programmable array microscope," *Proceedings of the Conference on Lasers and Electro-optics/Quantum Electronics and Laser Science Conference*, Baltimore, MD, Vol. 1–5, pp. 3135–3136 (2009).
36. R. F. Marcia and R. M. Willett, "Compressive coded aperture super-resolution image reconstruction," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 833–836, Las Vegas, NV (2008).
37. R. F. Marcia, Z. T. Harmany, and R. M. Willett, "Compressive coded aperture imaging," in *Proceedings of the 2009 IS&T/SPIE Electronic Imaging: Computational Imaging VII*, San Jose, CA (2009).
38. J. Romberg, "Compressive sampling by random convolution," *SIAM J. Imaging Sci.* **2**(4), 1098–1128 (2009).
39. A. Stern and B. Javid, "Random projections imaging with extended space-bandwidth product," *J. Disp. Tech.* **3**(3), 315–320 (2007).
40. W. Yin, S. P. Morgan, J. Yang, and Y. Zhang, "Practical compressive sensing with Toeplitz and circulant matrices," in *Proceedings of Visual Communications and Image Processing (VCIP)*, SPIE, San Jose, CA (2010).
41. R. Fergus, A. Torralba, and W. T. Freeman, "Random lens imaging," Tech. Rep. MIT-CSAIL-TR-2006-058, MIT Computer Science and Artificial Intelligence Laboratory (2006).
42. R. Robucci, J. D. Gray, L. K. Chiu, J. Romberg, and P. Hasler, "Compressive sensing on a CMOS separable-transform image sensor," *Proc. IEEE* **98**(6), 1089–1101 (2010).
43. V. Majidzadeh, L. Jacques, A. Schmid, P. Vanderghenst, and Y. Leblebici, "A (256*256) pixel 76.7mW CMOS imager/compressor based on real-time in-pixel compressive sensing," in *IEEE International Symposium on Circuits and Systems*, Paris, France, pp. 2956–2959 (2010).
44. M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schultz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express* **15**(21), 14013–14027 (2007).
45. A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Appl. Opt.* **47**(10), B44–B51 (2008).
46. M. Maggioni, G. L. Davis, F. J. Warner, F. B. Geshwind, A. C. Coppi, R. A. DeVerse, and R. R. Coifman, "Hyperspectral microscopic analysis of normal, benign and carcinoma microarray tissue sections," *Proc. SPIE* **6091** (2006).
47. J. Gu, S. K. Nayar, E. Grinspun, P. N. Belhumeur, and R. Ramamoorthi, "Compressive Structured Light for Recovering Inho-

- homogeneous Participating Media," in *European Conference on Computer Vision (ECCV)*, Marseille-France (2008).
48. M. A. Golub, M. Nathan, A. Averbuch, E. Lavi, V. A. Zheludev, and A. Schlar, "Spectral multiplexing method for digital snapshot spectral imaging," *Appl. Opt.* **48**, 1520–1526 (2009).
49. M. Shankar, R. Willett, N. Pitsianis, T. Schulz, R. Gibbons, R. T. Kolste, J. Carriere, C. Chen, D. Prather, and D. Brady, "Thin infrared imaging systems through multichannel sampling," *Appl. Opt.* **47**(10), B1–B10 (2008).
50. A. F. Coskun, I. Sencan, T.-W. Su, and A. Ozcan, "Lensless wide-field fluorescence imaging on a chip using compressive decoding of sparse objects," *Opt. Express* **18**(10), 10510–10523 (2010).
51. L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin, "Sparsity and compressed sensing in radar imaging," *Proc. IEEE* **98**(6), 1006–1020 (2010).
52. J. Gu, S. K. Nayar, E. Grinspun, P. N. Belhumeur, and R. Ramamoorthi, "Compressive structured light for recovering inhomogeneous participating media," in *Proceedings of the European Conference on Computer Vision*, Marseille, France (2008).
53. M. Mohtashemi, H. Smith, D. Walburger, F. Sutton, and J. Diggins, "Sparse sensing DNA microarray-based biosensor: Is it feasible?," in *2010 IEEE Sensors Applications Symposium*, Limerick, Ireland, pp. 127–130 (2010).
54. M. Sheikh, O. Milenkovic, and R. Baraniuk, "Designing compressive sensing DNA microarrays," in *2nd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, St. Thomas, VI, USA 2007, pp. 141–144 (2007).
55. M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.* **58**, 1182–1195 (2007).
56. A. C. Gurbuz, J. H. McClellan, and W. R. Scott, "A compressive sensing data acquisition and imaging method for stepped frequency GPRs," *IEEE Trans. Signal Process.* **57**, 2640–2650 (2009).
57. P. Ye, J. L. Paredes, G. Arce, Y. Wu, C. Chen, and D. Prather, "Compressive confocal microscopy," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, Taipei, Taiwan, pp. 429–432 (2009).
58. J. Bobin, J.-L. Starck, and R. Ottensmeyer, "Compressed sensing in astronomy," *IEEE J. Sel. Top. Signal Process.* **2**, 718–726 (2008).
59. M. Raginsky, R. M. Willett, Z. T. Harmany, and R. F. Marcia, "Compressed sensing performance bounds under Poisson noise," *IEEE Trans. Signal Process.* **58**(8), 3990–4002 (2010).
60. M. Raginsky, S. Jafarpour, Z. Harmany, R. Marcia, R. Willett, and R. Calderbank, "Performance bounds for expander-based compressed sensing in Poisson noise," accepted to *IEEE Transactions on Signal Processing*, 2011.
61. W. Dai, H. V. Pham, and O. Milenkovic, "Distortion-rate functions for quantized compressive sensing," in *IEEE Information Theory Workshop on Networking and Information Theory, 2009, ITW 2009*, pp. 171–175 (2009).
62. L. Jacques, D. K. Hammond, and M. J. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Trans. Inf. Theory* **57**(1), 559–571 (2011).
63. J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing," in *Applied and Computational Harmonic Analysis*, (2011) (in press).
64. M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Top. Signal Process.* **1**(4), 586–597 (2007).
65. E. van den Berg and M. P. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comput.* **31**(2), 890–912 (2008).
66. S. Wright, R. Nowak, and M. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Process.* **57**, 2479–2493 (2009).
67. P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Model. Simul.* **4**(4), 1168–1200 (2005).
68. I. Daubechies, M. Deffrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.* **LVII**, 1413–1457 (2004).
69. J. M. Bioucas-Dias and M. A. Figueiredo, "A new TwIST: two-step iterative shrinkage/thresholding algorithms for image restoration," *IEEE Trans. Image Process.* **16**(12), 2992–3004 (2007).
70. G. Davis, S. Mallat, and M. Avellaneda, "Greedy adaptive approximation," *J. Constructive Approximation* **13**, 57–98 (1997).
71. D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory* **52**(1), 6–18 (2006).
72. J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory* **50**, 2231–2242 (2004).
73. T. Chan and J. Shen, *Image Processing And Analysis: Variational, PDE, Wavelet, and Stochastic Methods*, Society for Industrial and Applied Mathematics, Philadelphia, PA (2005).
74. W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing," *SIAM J. Imag. Sci.* **1**(1), 143–168 (2008).
75. J. P. Oliveira, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Review: Adaptive total variation image deblurring: A majorization-minimization approach," *Signal Process.* **89**(9), 1683–1693 (2009).
76. A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imaging Vision* **20**(1–2), 89–97 (2004).
77. J. Darbon and M. Sigelle, "Image restoration with discrete constrained total variation part i: Fast and exact optimization," *J. Math. Imaging Vision* **26**(3), 261–276 (2006).
78. J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science* **290**(5500), 2319–2323 (2000).
79. S. Roweis and L. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science* **22**(5500), 2323–2326 (2000).
80. M. Belkin, P. Niyogi, and V. Singhwani, "Manifold regularization: A geometric framework for learning from labeled and unlabeled examples," *J. Mach. Learn. Res.* **7**, 2399–2434 (2006).
81. M. B. Wakin, "Manifold-based signal recovery and parameter estimation from compressive measurements," (2009), (submitted) <http://arxiv.org/abs/1002.1247>.
82. M. Davenport, M. Duarte, M. Wakin, J. Laska, D. Takhar, K. Kelly, and R. Baraniuk, "The smashed filter for compressive classification and target recognition," in *Proc. of SPIE Computational Imaging V*, San Jose, CA (2007).
83. Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Info. Theory* **55**(11), 5302–5316 (2009).
84. R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Inf. Theory* **56**(4), 1982–2001 (2010).
85. E. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.* **9**(6), 717–772 (2009).
86. E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" [arXiv:0912.3599v1](http://arxiv.org/abs/0912.3599v1) (2009).
87. M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. Dunson, G. Sapiro, and L. Carin, "Nonparametric bayesian dictionary learning for analysis of noisy and incomplete images," *IEEE Trans. Image Process.* (2010) (submitted).
88. Z. T. Harmany, D. O. Thompson, R. M. Willett, and R. F. Marcia, "Gradient projection for linearly constrained convex optimization in sparse signal recovery," in *IEEE International Conference on Image Processing*, Hong Kong, China (2010).
89. D. Snyder, *Random Point Processes*, Wiley-Interscience, New York, (1975).
90. S. Ahn and J. Fessler, "Globally convergent image reconstruction for emission tomography using relaxed ordered subsets algorithms," *IEEE Trans. Med. Imaging* **22**, 613–626 (2003).
91. R. Willett and R. Nowak, "Multiscale Poisson intensity and density estimation," *IEEE Trans. Inf. Theory* **53**(9), 3171–3187 (2007).
92. Z. T. Harmany, R. F. Marcia, and R. M. Willett, "SPIRAL out of convexity: Sparsity-regularized algorithms for photon-limited imaging," in *Proc. of SPIE Computational Imaging VIII*, San Jose, CA (2010).
93. R. M. Willett, Z. T. Harmany, and R. F. Marcia, "Poisson image reconstruction with total variation regularization," in *Proceedings of IEEE International Conference on Image Processing*, Hong Kong, China (2010).
94. Z. T. Harmany, R. F. Marcia, and R. M. Willett, "Sparse Poisson intensity reconstruction algorithms," in *IEEE Workshop on Statistical Signal Processing*, Cardiff, Wales, UK (2009).
95. F.-X. Dupe, J. M. Fadili, and J.-L. Starck, "A proximal iteration for deconvolving Poisson noisy images using sparse representations," *IEEE Trans. Image Process.* **18**(2), 310–321 (2009).
96. M. Figueiredo and J. Bioucas-Dias, "Deconvolution of Poissonian images using variable splitting and augmented Lagrangian optimization," in *IEEE Workshop on Statistical Signal Processing*, Cardiff, United Kingdom (2009). Available at <http://arxiv.org/abs/0904.4868>.
97. J. Y. Park and M. B. Wakin, "A multiscale framework for compressive sensing of video," in *Picture Coding Symposium (PCS)*, Chicago, IL (2009).
98. R. F. Marcia and R. M. Willett, "Compressive coded aperture video reconstruction," in *Proceedings 16th European Signal Processing Conference*, 2008, Lausanne, Switzerland (2008).
99. N. Jacobs, S. Schuh, and R. Pless, "Compressive sensing and differential image motion," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Dallas, TX (2010).
100. D. O. Thompson, Z. T. Harmany, and R. F. Marcia, "Sparse video recovery using linearly constrained gradient projection," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Prague, Czech Republic (2011).
101. R. F. Marcia, Z. T. Harmany, and R. M. Willett, "Compressive coded apertures for high-resolution imaging," in *Proceedings of the 2010 SPIE Photonics Europe*, Brussels, Belgium (2010).

102. Y. Rivenson, A. Stern, and B. Javidi, "Single exposure super-resolution compressive imaging by double phase encoding," *Opt. Express* **18**(14), 15094–15103 (2010).
103. J. W. Stayman, N. Subotic, and W. Buller, "An analysis of coded aperture acquisition and reconstruction using multiframe code sequences for relaxed optical design constraints," in *Proc. SPIE* **7468** (2009).
104. N. Savage, "Digital spatial light modulators," *Nature Photonics* **3**, 170–172 (2009).
105. <http://www.bnonlinear.com/products/xyslm/xy.htm>.

Rebecca M. Willett is an assistant professor in the Electrical and Computer Engineering Department at Duke University. She completed her PhD in Electrical and Computer Engineering at Rice University in 2005. Prof. Willett received the National Science Foundation CAREER Award in 2007, is a member of the DARPA Computer Science Study Group, and received an Air Force Office of Scientific Research Young Investigator Program award in

2010. Prof. Willett has also held visiting researcher positions at the Institute for Pure and Applied Mathematics at UCLA in 2004, the University of Wisconsin-Madison 2003-2005, the French National Institute for Research in Computer Science and Control (INRIA) in 2003, and the Applied Science Research and Development Laboratory at GE Healthcare in 2002. Her research interests include network and imaging science with applications in medical imaging, wireless sensor networks, astronomy, and social networks. Additional information, including publications and software, are available online at <http://www.ee.duke.edu/~willett/>.

Roummel F. Marcia received a BA in Mathematics from Columbia University in 1995 and his PhD in Mathematics from the University of California, San Diego in 2002. He was a Computation and Informatics in Biology and Medicine Postdoctoral Fellow in the Biochemistry Department at the University of Wisconsin-Madison and a Research Scientist in the Electrical and Computer Engineering at Duke University. He is currently an Assistant Professor of Applied Mathematics at the University of California, Merced. His research interests include nonlinear optimization, numerical linear algebra, signal and image processing, and computational biology.

Jonathan M. Nichols received a BSc degree from the University of Delaware in 1997 and MSc and PhD degrees from Duke University in 1999 and 2002 respectively, all in Mechanical Engineering. He currently works for the Naval Research Laboratory in Washington, DC as a member of the Maritime Sensing Section in the Optical Sciences Division. His research interests include infrared imaging, signal processing, and inverse problems.