

A New Compressive Imaging Camera Architecture using Optical-Domain Compression

Dharmpal Takhar, Jason N. Laska, Michael B. Wakin, Marco F. Duarte, Dror Baron
Shriram Sarvotham, Kevin F. Kelly, Richard G. Baraniuk

Department of Electrical and Computer Engineering
Rice University

ABSTRACT

Compressive Sensing is an emerging field based on the revelation that a small number of linear projections of a compressible signal contain enough information for reconstruction and processing. It has many promising implications and enables the design of new kinds of *Compressive Imaging* systems and cameras. In this paper, we develop a new camera architecture that employs a digital micromirror array to perform optical calculations of linear projections of an image onto pseudorandom binary patterns. Its hallmarks include the ability to obtain an image with a single detection element while sampling the image fewer times than the number of pixels. Other attractive properties include its universality, robustness, scalability, progressivity, and computational asymmetry. The most intriguing feature of the system is that, since it relies on a single photon detector, it can be adapted to image at wavelengths that are currently impossible with conventional CCD and CMOS imagers.

Keywords: Compressed sensing, sparsity, incoherent projections, random matrices, linear programming, imaging, camera

1. INTRODUCTION

Imaging sensors, hardware, and algorithms are under increasing pressure to accommodate ever larger and higher-dimensional data sets; ever faster capture, sampling, and processing rates; ever lower power consumption; communication over ever more difficult channels; and radically new sensing modalities. Fortunately, over the past few decades, there has been an enormous increase in computational power and data storage capacity, which provides a new angle to tackle these challenges. We could be on the verge of moving from a “digital signal processing” (DSP) paradigm, where analog signals (including light fields) are sampled periodically to create their digital counterparts for processing, to a “computational signal processing” (CSP) paradigm, where analog signals are converted directly to any of a number of intermediate, “condensed” representations for processing using various nonlinear techniques.

1.1. Compressive sensing

CSP builds upon a core tenet of signal processing and information theory: that signals, images, and other data often contain some type of *structure* that enables intelligent representation and processing. The notion of structure has been characterized and exploited in a variety of ways for a variety of purposes. Current state-of-the-art compression algorithms employ a decorrelating transform to compact a correlated signal’s energy into just a few essential coefficients.^{1–3} Such *transform coders* exploit the fact that many signals have a *sparse* representation in terms of some basis, meaning that a small number K of adaptively chosen transform coefficients can be transmitted or stored rather than $N \gg K$ signal samples. For example, smooth images are sparse in the Fourier basis, and piecewise smooth images are sparse in a wavelet basis;⁴ the commercial coding standards JPEG⁵ and JPEG2000⁶ directly exploit this sparsity.

The standard procedure for transform coding of sparse signals is to (i) acquire the full N -sample signal x ; (ii) compute the complete set of transform coefficients; (iii) locate the K largest, significant coefficients and discard the (many) small coefficients; and (iv) encode the *values and locations* of the largest coefficients. In cases where N is large and K is small, this procedure can be quite inefficient. Much of the output of the analog-to-digital conversion process ends up being discarded (though it is not known a priori which pieces are needed). Arguably, this “sample first, ask questions later”

Supported by NSF, ONR, AFOSR, and the Texas Instruments Leadership University Program.
Email: {kaka, laska, wakin, duarte, drorb, shri, kkelly, richb}@rice.edu; Web: dsp.rice.edu/cs.

process places unnecessary demands on DSP systems, particularly imaging systems where each digital sample requires its own imaging sensor (pixel).

This raises a simple question: For a given signal, is it possible to directly estimate the set of large coefficients that will not be discarded by the transform coder? While this seems improbable, the recent theory of *Compressive Sensing* (also known as Compressed Sensing, or CS) introduced by Candès, Romberg, and Tao^{7,8} and Donoho⁹ demonstrates that a signal that is K -sparse in one basis (call it the *sparsity basis*) can be recovered from cK *nonadaptive* linear projections onto a second basis (call it the *measurement basis*) that is incoherent with the first, where c is a small oversampling constant. While the measurement process is linear, the reconstruction process is decidedly *nonlinear*.

1.2. Compressive imaging

A critical aspect of CS measurements is *multiplexing*: each measurement is a function of several of the signal samples or image pixels. From this reduced set of measurements, it can still be possible (using CS techniques) to extract the salient signal information. This principle of “sample less, compute later” shifts the technological burden from the sensor to the processing. Thus, CS is an enabling framework for the CSP paradigm.

In this paper, we develop a new system to support what can be called Compressive Imaging (CI). Our system incorporates a microcontrolled mirror array driven by pseudorandom and other measurement bases and a single or multiple photodiode optical sensor. This hardware optically computes incoherent image measurements as dictated by the CS theory; we then apply CS reconstruction algorithms — described below — to obtain the acquired images.

Our imaging system enjoys a number of desirable features:

- **Single detector:** By time multiplexing a single detector, we can use a less expensive and yet more sensitive photon detector. This is particularly important when the detector is expensive, making an N -pixel array prohibitive. A single detector camera can also be adapted to image at wavelengths that are currently impossible with conventional CCD and CMOS imagers.
- **Universality:** Random and pseudorandom measurement bases are *universal* in the sense that they can be paired with any sparse basis. This allows exactly the same encoding strategy to be applied in a variety of different sensing environments; knowledge of the nuances of the environment are needed only at the decoder. Random measurements are also *future-proof*: if future research in image processing yields a better sparsity-inducing basis, then the same set of random measurements can be used to reconstruct an even better quality image.
- **Encryption:** A pseudorandom basis can be generated using a simple algorithm according to a random seed. Such encoding effectively implements a form of *encryption*: the 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 they allow a *progressively better reconstruction* of the data as more measurements are obtained; one or more measurements can also be lost without corrupting the entire reconstruction.
- **Scalability:** We can adaptively select how many measurements to compute 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:** Finally, CI places most of its computational complexity in the decoder, which will often have more substantial computational resources than the encoder/imager. The encoder is very simple; it merely computes incoherent projections and makes no decisions.

This paper is organized as follows. Section 2 provides a brief overview of the CS theory. Section 3 outlines our proposed CI system. Section 4 presents preliminary experimental results, and Section 5 concludes.

2. COMPRESSIVE SENSING BACKGROUND

2.1. Sparse representations

Consider a length- N , real-valued signal x of any dimension indexed as $x(n)$, $n \in \{1, 2, \dots, N\}$. For two-dimensional (2D) images we simply choose a 1D ordering of the N pixels. We use the terms “signal” and “image” interchangeably below. Suppose that the basis $\Psi = [\psi_1, \dots, \psi_N]$ provides a K -sparse representation of x ; that is

$$x = \sum_{n=1}^N \theta(n) \psi_n = \sum_{\ell=1}^K \theta(n_\ell) \psi_{n_\ell}. \quad (1)$$

Here x is a linear combination of K vectors chosen from Ψ , $\{n_\ell\}$ are the indices of those vectors, and $\{\theta(n)\}$ are the coefficients; the concept is extendable to tight frames.⁹ Alternatively, we can write in matrix notation

$$x = \Psi\theta, \quad (2)$$

where x is an $N \times 1$ column vector, the *sparse basis* matrix Ψ is $N \times N$ with the basis vectors ψ_n as columns, and θ is an $N \times 1$ column vector with K nonzero elements. Using $\|\cdot\|_p$ to denote the ℓ_p norm,* we can write that $\|\theta\|_0 = K$. Various expansions, including wavelets,⁴ Gabor bases,⁴ and curvelets¹⁰ are widely used for representation and compression of natural images and other data.

For the moment, we will focus on exactly K -sparse signals and defer discussion of the more general situation where the coefficients decay rapidly but not to zero. Similar principles hold for “compressible” signals that are well-approximated using K terms, and this robustness is critical to the success of our imaging system.

2.2. Incoherent projections

In CS, we do not measure or encode the K significant $\theta(n)$ directly. Rather, we measure and encode $M < N$ projections $y(m) = \langle x, \phi_m^T \rangle$ of the signal onto a *second set* of basis functions $\{\phi_m\}$, $m \in \{1, 2, \dots, M\}$, where ϕ_m^T denotes the transpose of ϕ_m and $\langle \cdot, \cdot \rangle$ denotes the inner product. In matrix notation, we measure

$$y = \Phi x, \quad (3)$$

where y is an $M \times 1$ column vector, and the *measurement basis* matrix Φ is $M \times N$ with each row a basis vector ϕ_m . Since $M < N$, recovery of the signal x from the measurements y is ill-posed in general; however the additional assumption of signal *sparsity* makes recovery possible and practical.

The CS theory tells us that when certain conditions hold, namely that the basis $\{\phi_m\}$ cannot sparsely represent the elements of the basis $\{\psi_n\}$ (a condition known as *incoherence* of the two bases^{7-9, 11}) and the number of measurements M is large enough, then it is indeed possible to recover the set of large $\{\theta(n)\}$ (and thus the signal x) from a similarly sized set of measurements $\{y(m)\}$. This incoherence property holds for many pairs of bases, including for example, delta spikes and the sine waves of the Fourier basis, or the Fourier basis and wavelets. Significantly, this incoherence also holds with high probability between an arbitrary fixed basis and a randomly generated one (consisting of i.i.d. Gaussian or Bernoulli/Rademacher ± 1 vectors). Signals that are sparsely represented in frames or unions of bases can be recovered from incoherent measurements in the same fashion.

2.3. Signal recovery via ℓ_0 optimization

The recovery of the sparse set of significant coefficients $\{\theta(n)\}$ can be achieved using *optimization* by searching for the signal with ℓ_0 -sparsest coefficients $\{\theta(n)\}$ that agrees with the M observed measurements in y (recall that $M < N$). Reconstruction relies on the key observation that, given some technical conditions on Φ and Ψ , the coefficient vector θ is the solution to the ℓ_0 minimization

$$\hat{\theta} = \arg \min \|\theta\|_0 \quad \text{s.t. } y = \Phi\Psi\theta \quad (4)$$

with overwhelming probability. (Thanks to the incoherence between the two bases, if the original signal is sparse in the θ coefficients, then no other set of sparse signal coefficients θ' can yield the same projections y .)

*The ℓ_0 “norm” $\|\theta\|_0$ merely counts the number of nonzero entries in the vector θ .

In principle, remarkably few incoherent measurements are required to recover a K -sparse signal via ℓ_0 minimization. Clearly, more than K measurements must be taken to avoid ambiguity; it has been established that $K + 1$ random measurements will suffice.^{12, 13}

Unfortunately, solving this ℓ_0 optimization problem is prohibitively complex, requiring a combinatorial enumeration of the $\binom{N}{K}$ possible sparse subspaces. In fact, the ℓ_0 -recovery problem is known to be NP-complete.¹⁴ Yet another challenge is robustness, since the recovery may be very poorly conditioned. In fact, *both* of these considerations (computational complexity and robustness) can be addressed, but at the expense of slightly more measurements.

2.4. Signal recovery via ℓ_1 optimization

The practical revelation that supports the new CS theory is that it is not necessary to solve the ℓ_0 -minimization problem to recover the set of significant $\{\theta(n)\}$. In fact, a much easier problem yields an equivalent solution (thanks again to the incoherency of the bases); we need only solve for the ℓ_1 -sparsest coefficients θ that agree with the measurements y ^{7-9, 15-19}

$$\hat{\theta} = \arg \min \|\theta\|_1 \quad \text{s.t. } y = \Phi\Psi\theta. \quad (5)$$

This optimization problem, also known as *Basis Pursuit*,²⁰ is significantly more approachable and can be solved with traditional linear programming techniques whose computational complexities are polynomial in N .

There is no free lunch, however; according to the theory, more than $K + 1$ measurements are required to recover sparse signals via Basis Pursuit. Instead, one typically requires $M \geq cK$ measurements, where $c > 1$ is an *oversampling factor*. As an example, for Gaussian random matrices, the rule of thumb $c \approx \log_2(\frac{N}{K} + 1)$ provides a convenient approximation to the required oversampling factor.¹²

2.5. Signal recovery via greedy pursuit

At the expense of slightly more measurements, iterative greedy algorithms have also been developed to recover the signal x from the measurements y . Examples include the iterative Orthogonal Matching Pursuit (OMP),¹¹ matching pursuit (MP),²¹ and tree matching pursuit (TMP)^{21, 22} algorithms. OMP, for example, iteratively selects the vectors from $\Phi\Psi$ that contain most of the energy of the measurement vector y . The selection at each iteration is made based on inner products between the columns of $\Phi\Psi$ and a residual; the residual reflects the component of y that is orthogonal to the previously selected columns. OMP is guaranteed to converge within a finite number of iterations. In CS applications, OMP requires $c \approx 2 \ln(N)$ to succeed with high probability.¹¹

TMP algorithms are especially well suited for reconstructing natural images. This class of algorithms exploits the structure of images that are sparse in a wavelet (or curvelet) basis; the wavelet basis functions can be grouped by scale and sorted in each group by offset. Two wavelet functions in consecutive scales are said to be linked by a *father-child* relationship when the support of the coarser-scale wavelet completely contains that of the finer-scale wavelet. For 2D signals such as images, each parent wavelet is linked to four children wavelets; the graph described by these links is a *quad-tree*.

Thanks to the singularity analysis properties of wavelets, wavelet coefficient values tend to propagate through scales. A large wavelet coefficient (in magnitude) generally indicates the presence of a singularity inside its support; a small wavelet coefficient generally indicates a smooth region. Thanks to the nesting of child wavelets inside their parents, edges manifest themselves in the wavelet domain as chains of large coefficients propagating through the scales of the wavelet quad-tree. Wavelet coefficients also have decaying magnitudes at finer scales.⁴ This induces the large wavelet coefficients of piecewise smooth images to form a connected subtree within the wavelet quad-tree.

Since greedy algorithms select the dictionary vector that explains the most energy from the image, it turns out to be unnecessary to check *all* possible coefficients at each iteration. Rather, the next most significant coefficient at each stage is likely to be among the children of the currently selected coefficients. By limiting the greedy algorithm to search only among these elements of the dictionary, the computational complexity of the TMP algorithm is significantly reduced compared to the original MP algorithm.²¹ Additionally, by considering the contribution of all the elements of the chain of coefficients anchored at the root of the quad-tree, the salient edge features of the signal under consideration can be reconstructed with higher priority and accuracy by the TMP algorithm.²²

While requiring a similar order of computation to linear programming-based reconstruction algorithms, greedy algorithms are often significantly faster in practice.




Figure 1. Compressive Imaging (CI) camera block diagram. Incident lightfield (corresponding to the desired image x) is reflected off a digital micro-mirror device (DMD) array whose mirror orientations are modulated in the pseudorandom pattern supplied by the random number generators (RNG). Each different mirror pattern produces a voltage at the single photodiode (PD) that corresponds to one measurement $y(m)$. From M measurements y we can reconstruct a sparse approximation to the desired image x using CS techniques.

2.6. Other signal recovery algorithms

The introduction of CS has inspired new measurement schemes and fast reconstruction algorithms. For instance, Ref. 23 employs group testing and random subset selection to estimate the locations and values of the nonzero coefficients in strictly sparse signals. The scheme provides a compact representation of the measurement basis as well as a fast $O(N \text{polylog} N)$ reconstruction algorithm that succeeds with high probability. Additionally, the complexity of the measurement calculation (3) is very low for signals that are sparse in the space domain. The algorithm can be extended to signals sparse in other bases and to compressible signals; however its main drawback is the large oversampling factor c required for high probability of successful reconstruction. Similar algorithms have been described elsewhere.²⁴

3. A COMPRESSIVE IMAGING TESTBED

CS/CI principles enable the design of flexible new imaging devices and techniques. Our hardware realization of the CI concept is a *single pixel camera*; it combines a micro-controlled mirror array displaying a time sequence of M pseudo-random basis functions with a single optical sensor to compute incoherent image measurements y as in (3) (see Figure 1). By time multiplexing a single detector,²⁵ we can employ a less expensive and yet more sensitive photon sensor. We can also adaptively select how many measurements to compute 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.

3.1. Camera hardware

Micro-actuated mirrors have proven to be a commercially viable MEMS technology for the video/projector display market as well as laser systems and telescope optics.²⁶ In our system, we employ a Texas Instruments (TI) digital micromirror device (DMD). The combination of a TI DMD developer's kit and accessory light modulator package (ALP) allows us to operate in the variety of modes necessary to test various CS approaches.

The DMD consists of an array of electrostatically actuated micro-mirrors where each mirror the array is suspended above an individual SRAM cell. The DMD micro-mirrors form a pixel array of size 1024×768 . Each mirror rotates about a hinge and can be positioned in one of two states (+12 degrees and -12 degrees from horizontal); thus light falling on the DMD may be reflected in two directions depending on the orientation of the mirrors. While the DMD was originally created for displaying images in televisions and digital projectors, other groups have recently begun to use it for enhanced image acquisition.²⁷⁻³⁰

In our CI setup (see Figures 1 and 2), the desired image is formed on the DMD plane with the help of a biconvex lens; this image acts as an object for the second biconvex lens which focuses the image onto the photodiode. The light is collected from one out of the two directions in which it is reflected (e.g., the light reflected by mirrors in the +12 degree state). The light from a given configuration of the DMD mirrors is summed at the photodiode to yield an absolute voltage




Figure 2. Compressive Imaging (CI) camera hardware setup.

that yields a coefficient $y(m)$ for that configuration. The output is amplified through an op-amp circuit and then digitized by a 12-bit analog-to-digital converter.

3.2. Optical multiplexing

The output voltage of the photodiode can be interpreted as the inner product of the desired image x with a measurement basis vector $\phi(m)$. In particular, letting $\rho(m)$ denote the mirror positions of the m -th measurement pattern, the voltage reading from the photodiode v can be written as

$$v(m) \propto \langle x, \phi(m) \rangle + \text{DC offset}, \quad (6)$$

where

$$\phi(m) = \mathbf{1}_{\{\rho(m)=+12 \text{ degrees}\}} \quad (7)$$

and $\mathbf{1}$ is the indicator function. (The DC offset can be measured by setting all mirrors to -12 degrees; it can then be subtracted off.)

Equation (6) holds the key for implementing a CI system. For a given image x , we take M measurements $\{y(1), y(2), \dots, y(M)\}$ corresponding to mirror configurations $\{\rho(1)\rho(2) \dots, \rho(M)\}$.[†] Since the patterns $\rho(m)$ are programmable, we can select them to be incoherent with the sparsity-inducing basis (e.g., wavelets or curvelets). As mentioned previously, random or pseudorandom measurement patterns enjoy a useful universal incoherence property with any fixed basis, and so we employ pseudorandom ± 12 degree patterns on the mirrors. These correspond to pseudorandom 0/1 Bernoulli measurement vectors $\phi_m = \mathbf{1}_{\{\rho(m)=+12 \text{ degrees}\}}$. (The measurements may easily be converted to ± 1 Rademacher patterns by setting all mirrors in $\rho(1)$ to $+12$ degrees and then letting $y(m) \leftarrow 2y(m) - y(1)$ for $m > 1$.) Other options for incoherent CI mirror patterns include $-1/0/1$ group-testing patterns.^{23,24} Mirrors can also be duty-cycled to give the elements of ϕ finer precision, for example to approximate Gaussian measurement vectors.^{8,9}

This system directly acquires a reduced set of M incoherent projections of an N -pixel image x *without* first acquiring the N pixel values.[‡] The key idea is that each measurement multiplexes several pixel values, and CS methodologies can be used to tease them out.

Since the camera is “progressive,” better quality images (larger K) can be obtained by taking more measurements M . Also, since the data measured by the camera is “future-proof,” new reconstruction algorithms based on better sparsifying image transforms can be applied at a later date to obtain even better quality images.

[†]We assume that the image x is stationary during the acquisition process. Currently, this is possible in controlled environments; ultimately the fast switching rates of the DMD will allow for applications in practical settings.

[‡]In our setup, the number of reconstructed pixels (and thus the resolution of the final image) corresponds to the number of micromirrors in the DMD array $N = 1024 \times 768$.

3.3. Related work

Two notable previous DMD-driven applications involve confocal microscopy^{31,32} and micro-optoelectromechanical (MOEM) systems.^{27–30}

The three primary differences between our CI/DMD camera and MOEM systems are the placement of the DMD between the image and the detector (as opposed to placement between the light source and the image); the replacement of the CCD detector with a single photodiode; and the large size of MOEM data sets. In a MOEM system, a DMD is positioned between the image source and the sensing element; its function is to limit the number of image pixel columns being sensed at a given time by reflecting off light from unwanted pixels and allowing light from the desired pixels to pass through. A MOEM system obtains the sum of the intensities of sets of columns from the sensed image; the purpose of multiplexing is to increase the signal-to-noise ratio of the measurements, which is lower when each column is sensed separately due to the low illumination caused by DMD modulation. In Refs. 27–29 the authors propose sets of N Hadamard patterns, which allows for simple demultiplexing, and randomized Hadamard patterns, which yield a uniform signal-to-noise ratio among the measurements. MOEM data sets are very large, requiring a large amount of storage and processing. In contrast, in our CI/DMD camera, compression in the optical domain greatly reduces the storage and processing of the image.

Other efforts on CI include Refs. 33, 34, which employ optical elements to perform transform coding of multispectral images. These designs obtain sampled outputs that correspond to coded information of interest, such as the wavelength of a given light signal or the transform coefficients in a basis of interest. The elegant hardware designed for these purposes uses concepts that include optical projections, group testing,²³ and signal inference.

4. EXPERIMENTAL RESULTS

4.1. Hardware configuration

In the optical setup depicted in Figure 2, the object is front illuminated by an LED light source. To minimize noise, the initial measurements have been performed by lock-in detection with the LED driven at 1kHz.

For simplicity, the initial images are square in shape and therefore use only a 768×768 array of pixels on the DMD. This array can be further sub-divided into blocks depending on the desired resolution for the reconstructed image, e.g., for a 64×64 pixel image we actuate 12×12 blocks of micro-mirrors in unison. Additionally, since the DMD array is programmed on the fly, it can be adaptively sectioned to highlight (zoom in on) part of the image (as opposed to acquiring the entire image and then digitally zooming through post-processing) or to adaptively modify the projection basis functions during acquisition for enhanced reconstruction performance.

4.2. Imaging results

For our imaging experiment, we displayed a printout of the letter “R” in front of the camera; Figure 3(a) shows the printout. We set the camera to acquire a 64×64 pixel image (hence, $N = 4096$). This size was chosen to ensure quick reconstruction during tests so that focusing and other adjustments could be made. The image quality benefits from the reduced resolution since activating the DMD in blocks of pixels reflects more light, thus reducing the sensitivity requirement on the photodetector and thus noise. As we discuss later, better hardware exists for some of the components, and higher resolution images will result.

Since our test image is piecewise constant (with sharp edges) it can be sparsely represented in the wavelet domain. Figures 3(b) and 3(c) show the best K -term Haar wavelet approximation of the idealized image in Figure 3(a) with $K = 400$ and 675 , respectively. Figure 3(d) shows the projection of the original test image onto the surface of the DMD; this image is a still frame from a video taken with a conventional digital camera, and the poor quality is due to the low resolution of the obtained NTSC signal with the video capture card and the low light level. Using $M = 1600$ and 2700 pseudorandom projections (roughly $4 \times$ the K used in (b) and (c)), we reconstructed the images shown in Figures 3(e) and 3(f) using Basis Pursuit.

From these results, it is clear that the recognizable features of the “R” can be recovered, even with half as many measurements as the number of pixels acquired. The reconstruction quality is also progressively better with higher M as well as more robust to noisy measurements, enhancing the reconstruction of the singularities (sharp edges). The sources of noise include subtle nonlinearities in the photodiode, nonuniform reflectance of the mirrors through the lens that focuses

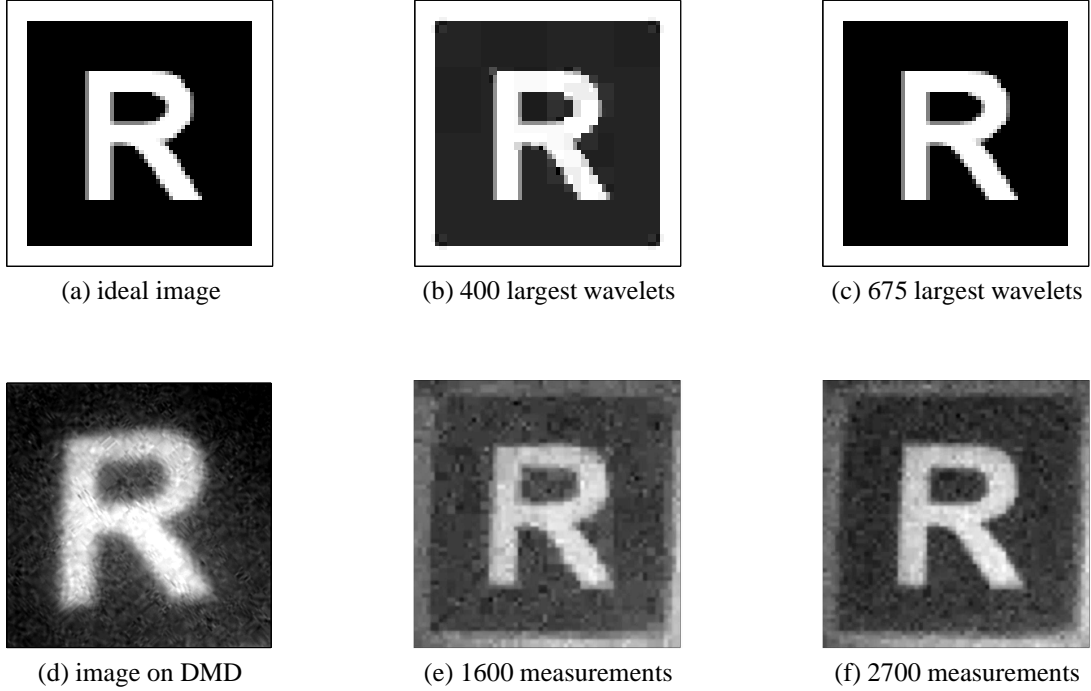


Figure 3. CI/DMD imaging of a 64×64 ($N = 4096$ pixel) image. Ideal image (a) of full resolution and approximated by its (b) largest 400 wavelet coefficients and (c) largest 675 wavelet coefficients. (d) Conventional 320×240 camera image acquired at the DMD plane. CS reconstruction from (e) 1600 random measurements and (f) 2700 random measurements. In all cases, Haar wavelets were used for approximation or reconstruction.

onto the photodiode (thus changing the weighting of the pattern blocks), and nonuniform mirror positions; work is ongoing to remove these sources of noise at the measurement setup. The robustness of the CS reconstruction will tend to suppress quantization noise from the A/D converter and photodiode circuit noise during detection.

5. DISCUSSION AND CONCLUSIONS

In this paper, we have presented a first prototype imaging system that successfully employs compressive sensing (CS) principles. The camera has many attractive features, including simplicity, universality, robustness, and scalability, that should enable it to impact a variety of different applications. Another interesting and potentially useful practical feature of our system is that it off-loads processing from data collection into data reconstruction. Not only will this lower the complexity and power consumption of the device, but it will enable adaptive new measurements schemes. The most intriguing feature of the system is that, since it relies on a single photon detector, it can be adapted to image at wavelengths that are currently impossible with conventional CCD and CMOS imagers.

In current research, we are developing sequences of projection functions that reduce the mechanical movement of the mirrors, designing measurement matrices with lower oversampling factors c , analyzing and characterizing the noise and resolution effects on reconstruction, and designing and implementing fast reconstruction algorithms.

There are many possibilities for extensions to our system, including:

- Plenty of measurement bases can be implemented with the DMD beyond Rademacher and Gaussian vectors.
- Reconstruction in a curvelet frame¹⁰ or joint wavelet/curvelet frame could provide higher-quality images due to their pumped-up sparsification capabilities.
- Color, multispectral, and hyperspectral imaging can be achieved using a more complex photon sensing element such as a dual sandwich photodetector or by combining multiple photodiodes.

- We should be able to greatly reduce the size and complexity of the setup by replacing the DMD chip with a MEMS-based shutter array placed directly over the photodiode.³⁵ This will also enable imaging of wavelengths that poorly reflect from the aluminum micromirror array.
- We can implement super-resolution where multiple shifted offsets are incorporated to reconstruct images with sub-pixel resolution.³⁶ In our setup, the DMD can be mounted onto a micro/nano-positioning device to make measurements with different lateral translations.
- Video encoding using CI can be achieved in several ways. Spatiotemporal random patterns can be used to encode a clip of video that is sparse in a 3D wavelet/curvelet domain. Also, separate frames of video can be encoded independently and decoded jointly by exploiting the correlation between the acquired frames.¹² Preliminary experiments with 3D incoherent measurements appear promising.
- CS principles can be extended to multiple imagers acquiring correlated images as in a camera network^{37,38} or lightfield/lumigraph capture system.³⁹

ACKNOWLEDGEMENTS

Thanks to Texas Instruments for providing the TI DMD developer's kit and accessory light modulator package (ALP). Thanks also to Dave Brady and Dennis Healy for enlightening discussions.

REFERENCES

1. R. A. DeVore, B. Jawerth, and B. J. Lucier, "Image compression through wavelet transform coding," *IEEE Trans. Inform. Theory*, vol. 38, pp. 719–746, Mar. 1992.
2. J. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Signal Processing*, vol. 41, pp. 3445–3462, Dec. 1993.
3. Z. Xiong, K. Ramchandran, and M. T. Orchard, "Space-frequency quantization for wavelet image coding," *IEEE Trans. Image Processing*, vol. 6, no. 5, pp. 677–693, 1997.
4. S. Mallat, *A Wavelet Tour of Signal Processing*. San Diego: Academic Press, 1999.
5. W. Pennebaker and J. Mitchell, "JPEG: Still image data compression standard," *Van Nostrand Reinhold*, 1993.
6. D. S. Taubman and M. W. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards and Practice*. Kluwer, 2001.
7. E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, 2004. Submitted.
8. E. Candès and T. Tao, "Near optimal signal recovery from random projections and universal encoding strategies," *IEEE Trans. Inform. Theory*, 2004. Submitted.
9. D. Donoho, "Compressed sensing," 2004. Preprint.
10. E. Candès and D. Donoho, "Curvelets — A surprisingly effective nonadaptive representation for objects with edges," *Curves and Surfaces*, 1999.
11. J. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," Apr. 2005. Preprint.
12. D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed compressed sensing," 2005. Preprint.
13. R. Venkataramani and Y. Bresler, "Further results on spectrum blind sampling of 2D signals," in *Proc. IEEE Int. Conf. Image Proc. (ICIP)*, vol. 2, (Chicago), Oct. 1998.
14. E. Candès and T. Tao, "Error correction via linear programming," *Found. of Comp. Math.*, 2005. Submitted.
15. E. Candès and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Found. of Comp. Math.*, 2004. Submitted.
16. E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. on Pure and Applied Math.*, 2005. Submitted.
17. E. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inform. Theory*, vol. 51, pp. 4203–4215, Dec. 2005.
18. E. Candès and J. Romberg, "Practical signal recovery from random projections," *IEEE Trans. Signal Processing*, 2005. Submitted.

19. D. Donoho and Y. Tsaig, "Extensions of compressed sensing," 2004. Preprint.
20. S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. on Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.
21. M. F. Duarte, M. B. Wakin, and R. G. Baraniuk, "Fast reconstruction of piecewise smooth signals from random projections," in *Proc. SPARS05*, (Rennes, France), Nov. 2005.
22. C. La and M. N. Do, "Signal reconstruction using sparse tree representation," in *Proc. Wavelets XI at SPIE Optics and Photonics*, (San Diego), Aug. 2005.
23. G. Cormode and S. Muthukrishnan, "Towards an algorithmic theory of compressed sensing," *DIMACS Tech. Report 2005-25*, September 2005.
24. A. C. Gilbert, M. J. Strauss, J. Tropp, and R. Vershynin, "Sublinear, small space approximation of compressible signals and uniform algorithmic embeddings," *Preprint*, November 2004.
25. D. Brady, "Multiplex sensors and the constant radiance theorem," *Optics Letters*, vol. 27, no. 1, pp. 16–18, 2002.
26. J. Sampsell, "An overview of the digital micromirror device (DMD) and its application to projection displays," in *1993 SID International Symposium Digest of Technical Papers*, vol. 24, p. 1012, 1993.
27. R. A. DeVerse, R. R. Coifman, A. C. Coppi, W. G. Fateley, F. Geshwind, R. M. Hammaker, S. Valenti, and F. J. Warner, "Application of spatial light modulators for new modalities in spectrometry and imaging," in *Proc. SPIE*, vol. 4959, pp. 12–22, 2003. <http://www.math.yale.edu/mmm82/hyperspectral.html>.
28. G. L. Davis, M. Maggioni, F. J. Warner, F. B. Geshwind, A. C. Coppi, R. A. DeVerse, and R. Coifman, "Hyperspectral analysis of normal and malignant colon tissue microarray sections using a novel DMDsystem," in *NIH Optical Imaging Workshop (poster)*, 2004. <http://www.math.yale.edu/mmm82/hyperspectral.html>.
29. G. L. Davis, M. Maggioni, F. J. Warner, F. B. Geshwind, A. C. Coppi, R. A. DeVerse, and R. R. Coifman, "Hyperspectral analysis of normal and malignant microarray tissue sections using a novel micro-optoelectrical-mechanical system," *Mod. Pathol.*, vol. 17 Suppl. 1:358A, 2004.
30. R. Muise and A. Mahalanobis, "Target detection using integrated hyper spectral sensing and processing." IMA Workshop on Integration of Sensing and Processing, December 2005.
31. P. M. Lane, R. P. Elliott, and C. E. MacAulay, "Confocal microendoscopy with chromatic sectioning," in *Proc. SPIE*, vol. 4959, pp. 23–26, 2003.
32. V. Bansal, S. Patel, and P. Saggau, "High-speed confocal laser scanning microscopy using acousto-optic deflectors and a digital micromirror device," in *Proc. SPIE*, vol. 5324, pp. 47–54, 2004.
33. N. P. Pitsianis, D. J. Brady, and X. Sun, "Sensor-layer image compression based on the quantized cosine transform," in *Visual Information Processing XIV*, vol. 5817, (Orlando, FL, USA), p. 250, SPIE, 2005.
34. D. J. Brady, M. Feldman, N. Pitsianis, J. P. Guo, A. Portnoy, and M. Fiddy, "Compressive optical montage photography," in *Photonic Devices and Algorithms for Computing VII*, vol. 5907, (San Diego, CA, USA), p. 590708, SPIE, 2005.
35. M. Pizzi, V. Koniachkine, N. Nieri, S. Sinesi, and P. Perlo, "Electrostatically driven film light modulators for display applications," *Microsystem Technologies - Micro- and Nanosystems-Information Storage and Processing Systems*, vol. 10, pp. 17–21, 2003.
36. J. Tanida, T. Kumagai, K. Yamada, S. Miyatake, K. Ishida, T. Morimoto, N. Kondou, D. Miyazaki, and Y. Ichioka, "Thin observation module by bound optics (TOMBO): Concept and experimental verification," *Applied Optics*, vol. 40, no. 11, pp. 1806–1813, 2001.
37. R. Wagner, R. G. Baraniuk, and R. Nowak, "Distributed image compression for sensor networks using correspondence analysis and super-resolution," *Proceedings of IEEE ICIP*, vol. 1, pp. 597–600, 2003.
38. W. E. Mantzel, H. Choi, and R. G. Baraniuk, "Distributed camera network localization," *Proceedings of the 38th Asilomar Conference on Signals, Systems and Computers*, vol. 2, pp. 1381–1386, 2004.
39. S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in *SIGGRAPH '96: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, (New York, NY, USA), pp. 43–54, ACM Press, 1996.