

Compressive imaging: hybrid measurement basis design

Amit Ashok^{1,*} and Mark A. Neifeld^{1,2}

¹1230 East Speedway Boulevard, Department of Electrical and Computer Engineering,
University of Arizona, Tucson, Arizona 85721, USA

²1630 East University Boulevard, College of Optics, University of Arizona, Tucson, Arizona 85721, USA

*Corresponding author: ashoka@ece.arizona.edu

Received February 2, 2011; accepted March 9, 2011;
posted March 25, 2011 (Doc. ID 142151); published May 17, 2011

The inherent redundancy in natural scenes forms the basis of compressive imaging where the number of measurements is less than the dimensionality of the scene. The compressed sensing theory has shown that a purely random measurement basis can yield good reconstructions of sparse objects with relatively few measurements. However, additional prior knowledge about object statistics that is typically available is not exploited in the design of the random basis. In this work, we describe a hybrid measurement basis design that exploits the power spectral density statistics of natural scenes to minimize the reconstruction error by employing an optimal combination of a nonrandom basis and a purely random basis. Using simulation studies, we quantify the reconstruction error improvement achievable with the hybrid basis for a diverse set of natural images. We find that the hybrid basis can reduce the reconstruction error up to 77% or equivalently requires fewer measurements to achieve a desired reconstruction error compared to the purely random basis. It is also robust to varying levels of object sparsity and yields as much as 40% lower reconstruction error compared to the random basis in the presence of measurement noise. © 2011 Optical Society of America

OCIS codes: 110.0110, 100.0100, 110.1758, 100.6640.

1. INTRODUCTION

It is well known that band-limited signals can be sampled at the Nyquist rate [1–3] without loss of information. However, many band-limited signals possess additional internal structure that reduces the inherent degrees of freedom. For example, it has been observed that natural scenes tend to be redundant and therefore sparse/compressible in some transform domains. As a result, such signals can be sampled at sub-Nyquist rates with minimal or no loss of information. The notion of signal structure, which, for example, is the basis of sparse/compressible signal representations, forms the foundation of the compressed sensing theory, which suggests parsimonious sampling of signals commensurate with their sparsity [4–7]. It should be noted that compressed sensing theory exploits the prior knowledge of signal sparsity to derive an efficient measurement basis without assuming any specific knowledge of the sparsity domain. Perhaps surprisingly, it has been shown that a random measurement basis is efficient for sampling sparse/compressible signals [4]. The theoretical underpinnings for the choice of random basis has a strong relation to the Johnson–Lindenstrauss lemma in approximation theory, which is described in [8]. The evolution of the sampling theory, beginning with the Nyquist sampling theory to the recent emergence of the compressed sensing theory, can be viewed as injecting an increasing amount of prior knowledge about the signals of interest to improve sampling efficiency. The compressed sensing theory also applies to imaging [9] and shares common elements with the independent and parallel development of compressive imaging, sometimes also referred to as feature-specific imaging, in the optics community [10–12]. In this work, we describe an application of additional prior information about natural scene statistics to develop an alternate measurement basis with improved effi-

ciency that departs from the celebrated random measurement basis. Specifically, we consider a measurement basis design approach that utilizes both nonrandom projections (from the signal sparsity basis) and random projections to construct a hybrid measurement basis. The motivation for this approach arises from the well-known power-law power spectral model of the second-order statistics of natural scenes [13–15].

2. THEORY: HYBRID MEASUREMENT BASIS DESIGN

Natural scenes tend to be sparse/compressible in transform domains, such as discrete cosine transform (DCT) or discrete wavelet transform (DWT). In addition to being sparse in some transform domains, it has been also observed that the power spectral density of such scenes reduces exponentially with increasing spatial/angular frequency [13–15]. Figure 1 shows a plot of the average radial power spectral density obtained from a set of diverse natural images. A power-law model provides a good approximation as shown by the solid red line in Fig. 1 [13–15]. Note that a power-law power spectral model implies that the expected magnitude (i.e., on average) of the transform coefficients (e.g., DCT coefficients) decreases with increasing spatial/angular frequencies. This is illustrated in Fig. 2, which shows the typical distribution of DCT coefficients for a sample image block (64 × 64 pixels) as a function of spatial/angular frequency. Given this distribution, one may infer that transform coefficients corresponding to lower spatial/angular frequencies tend to have larger magnitude and therefore are less sparse compared to coefficients associated with higher spatial frequencies. This structure of coefficient sparsity represents additional prior information available in the case of natural scenes. Recall that the compressed sensing theory does not make any specific assumption about the

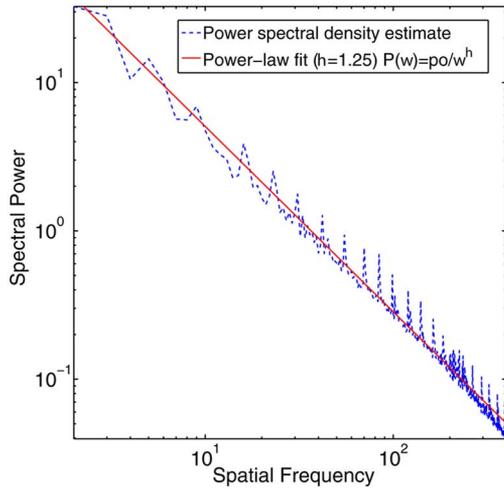


Fig. 1. (Color online) Radial power spectral density function and a power-law fit for natural images.

nature of sparse domain or the distribution of transform domain coefficients. In the absence of such knowledge, the compressed sensing theory shows that an efficient nonadaptive measurement basis is the random basis, which is maximally incoherent [16] with any unknown (at measurement time), sparse representation basis. In the case of natural scenes, we can and should exploit this additional prior knowledge to design a more efficient measurement basis that could potentially improve upon the purely random measurement basis. In fact, Donoho and Candes [4,5,16] have suggested such an approach: use DCT/DWT coefficients directly to capture low-frequency information in images and use random projections to estimate the high-frequency information. However, to the best of our knowledge, no work has been done to formalize this approach toward maximizing the reconstruction performance for natural images. Here, we describe a design approach that employs both nonrandom and random projections to construct a hybrid measurement basis by exploiting this additional prior information.

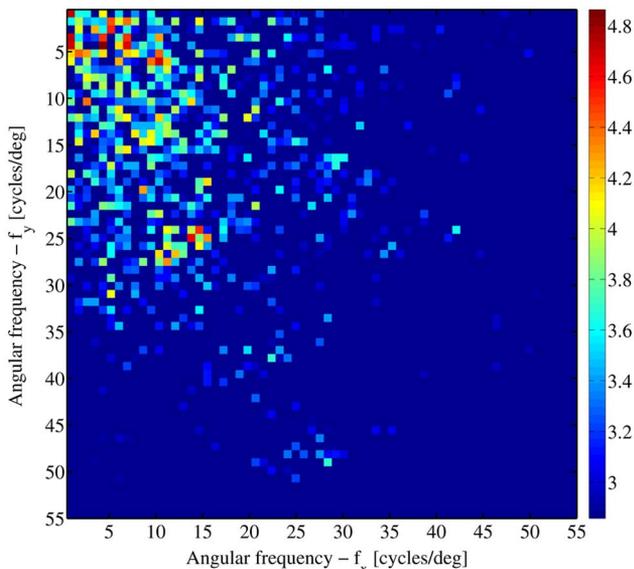


Fig. 2. (Color online) Distribution of energy in DCT coefficients for a 64 × 64 image patch of a natural image.

It is well known that the DCT transform is a good choice of compressible representation basis for natural images motivating its widespread use in the JPEG image compression standard [17,18]. We also know that the DCT coefficients corresponding to low frequencies are significant and nonsparse; therefore, we should use DCT projections to measure such coefficients directly. On the other hand, DCT coefficients corresponding to relatively higher frequencies become more sparse, so it would be more efficient to use random projections to estimate these coefficients indirectly. This is indeed the main idea of the hybrid measurement basis design, which combines M_1 DCT projections and M_2 random projections such that $M_1 + M_2 = M$. Now, we describe the hybrid measurement basis design approach formally. Let us denote the DCT coefficients of an image block with N pixels represented by the column vector $\vec{\theta} = [\theta_1, \theta_2, \dots, \theta_N]^T$. Note that the DCT coefficients are naturally ordered according to increasing spatial frequencies. If we denote the image block by a column vector $\vec{f} = [f_1, f_2, \dots, f_N]^T$, obtained by lexicographically arranging the pixels within the block, we can express

$$\vec{\theta} = \mathbf{D}\vec{f}, \tag{1}$$

where \mathbf{D} is a DCT transform matrix of size $N \times N$. We can rearrange the DCT coefficients $\vec{\theta}$ in descending magnitude order and express them in a new column vector $\hat{\theta} = [\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_N]^T$, where $|\hat{\theta}_1| \geq |\hat{\theta}_2| \geq \dots \geq |\hat{\theta}_N|$. For natural scenes, these coefficients can be described by a power-law distribution; therefore, we can state

$$|\hat{\theta}_n| \leq R \cdot n^{-\frac{1}{p}}, \tag{2}$$

where p is the power-law exponent, n is the index of the DCT coefficient, and R is the DCT coefficient value at zero frequency (DC). A K -term approximation $\hat{\theta}_K$, obtained by retaining only the K largest $\hat{\theta}$ coefficients (i.e., setting the remaining coefficients to zero), incurs the following ℓ_2 error [5,7]:

$$\|\vec{\theta} - \hat{\theta}_K\|_{\ell_2} \leq C_p \cdot R \cdot K^{-\left(\frac{1}{p-1}\right)}, \tag{3}$$

and because \mathbf{D} is a unitary transform, we can also state

$$\|\vec{f} - \hat{f}_K\|_{\ell_2} \leq C_p \cdot R \cdot K^{-\left(\frac{1}{p-1}\right)}, \tag{4}$$

where C_p is a function of power-law exponent p . Note that above relationship is upper bound on the ℓ_2 error [5,7] and the exact error depends on the particular values of p and K .

According to compressed sensing theory [7], a K -term approximation of a sparse signal uses M random measurements, where M is given by

$$M = C_k(\mu) \cdot K \cdot \log(N), \tag{5}$$

where C_k is a function of the coherence μ between measurement and sparsity bases and N is the dimensionality of the data. For real signal/images, it has been observed that M is typically in the range of 3 to 5 K . Given the relationships in Eqs. (4) and (5), we can upper bound ℓ_2 reconstruction error achievable with M purely random compressive measurements as

$$\|\vec{f} - \hat{f}_K\|_{\ell_2} \leq C_p \cdot R \cdot \left(\frac{M}{C_k \cdot \log(N)} \right)^{-\left(\frac{1}{p-1}\right)}. \tag{6}$$

Now, let us define a random variable X_M that denotes the number of the most significant DCT coefficients (i.e., the first θ s) that appear among the first M DCT coefficients (i.e., the first θ s). Equivalently, we can define X_M as the cardinality of the intersection set: $\{\theta_i : i \in 1 \dots M\} \cap \{\theta_i : i \in 1 \dots M\}$ and, as a result, $X_M \leq M$. For a power-law DCT coefficient model, we expect that, for small M , X_M grows as $X_M = M$. However, as M increases and DCT coefficients become increasingly sparse, the rate of change of X_M decreases and X_M approaches the number of nonzero DCT coefficients P as $M \rightarrow N$.

The hybrid measurement basis design can be stated as follows: given M compressive measurements, use M_1 DCT projections associated with low frequencies to measure the corresponding DCT coefficients directly and use M_2 random projections to measure the remaining sparser DCT coefficients. Now, the problem is to find the optimal M_1 and M_2 distribution so as to minimize the ℓ_2 reconstruction error such that $M_1 + M_2 = M$. It should be noted that here the ℓ_2 norm is a measure of reconstruction error and it does not preclude or substitute the use of the ℓ_1 norm in the reconstruction algorithm to yield sparse reconstructions.

Recall that the ℓ_2 error with a K -term approximation of a signal is given by Eq. (4). Therefore, we must now determine the equivalent K that results from the sum of K_1 (associated with M_1 DCT projections) and K_2 (resulting from M_2 random projections) within the hybrid basis. Because DCT projections are completely coherent with the sparsity basis (i.e., the DCT itself), measuring M_1 DCT coefficients results in $K_1 = X_{M_1}$. For M_2 random projections, which are incoherent with the sparsity basis, K_2 is obtained by applying Eq. (5), which yields

$$K_2 = \frac{M_2}{C_K \cdot \log(N)}. \tag{7}$$

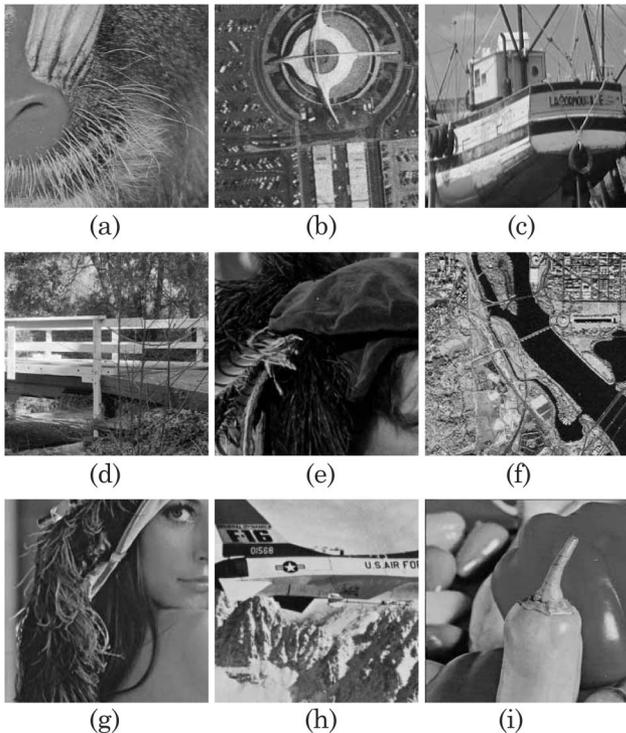


Fig. 3. Samples from the image set used in the simulation study.

Therefore, we can now express K as

$$K = X_{M_1} + \frac{M_2}{C_K \cdot \log(N)}, \tag{8}$$

and, substituting this expression into Eq. (4), we obtain the following expression for ℓ_2 reconstruction error using M measurements with the hybrid basis:

$$\|\vec{f} - \hat{f}_K\|_{\ell_2} \leq C_p \cdot R \cdot \left(\mathbb{E}(X_{M_1}) + \frac{M - M_1}{C_K \cdot \log(N)} \right)^{-\left(\frac{1}{p} - \frac{1}{2}\right)}. \tag{9}$$

Note the expectation in the Eq. (9) is over the ensemble of natural images that leads to the expected reconstruction error for the hybrid measurement basis. For each value of M , there is an optimal M_1^{opt} DCT projections and $M_2^{\text{opt}} = M - M_1^{\text{opt}}$ random projections comprising the hybrid measurement basis that minimizes the expected reconstruction error. In Section 3, we describe a simulation study that employs a diverse set of natural images to quantify the reconstruction performance of the hybrid, random, and DCT measurement bases for different object sparsity levels, noiseless and noisy measurement models.

3. SIMULATION STUDY

We expected that incorporating additional prior information about natural scenes (i.e., beyond sparsity) into the hybrid measurement basis should yield improved reconstruction compared to the random measurement basis. Here, we use a simulation study to quantify the reconstruction performance for the three different measurement bases for a variety of natural images. Specifically, a set of 17 diverse images (256×256 pixels) drawn from the University of Southern California Signal and Image Processing Institute image database [19] is used to form a representative ensemble of natural images. Figure 3 shows some samples images from this image set. These images are analyzed in a blockwise scheme where each image block of size $\sqrt{N} \times \sqrt{N}$ is considered a sample. Here, DCT is used as the sparsity basis to represent these image blocks. Sparsity, denoted by α , is defined as $\alpha = 100 \times \frac{N-P}{N}\%$ (i.e., percentage of zero coefficients). Figure 4 shows an image and the corresponding sparse image with $\alpha = 70\%$ for a block size of 32×32 . Note that the original and sparse images are nearly indistinguishable, illustrating the effectiveness of DCT as a sparsity basis. We begin by analyzing the reconstruction performance of the hybrid basis in the absence of measurement noise.

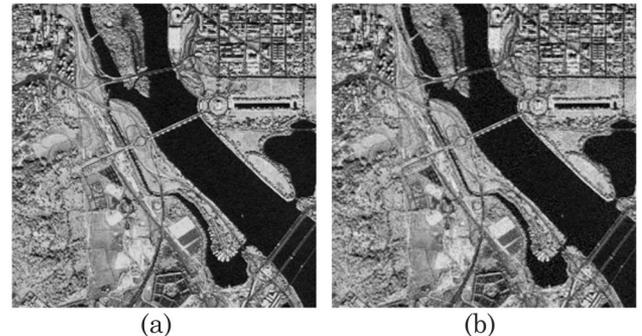


Fig. 4. Reference image and corresponding sparse image ($\alpha = 70\%$).

A. Noiseless Compressive Reconstruction

Mathematically, a noiseless compressive measurement can be written as

$$\vec{g} = \mathbf{H}\vec{f}, \tag{10}$$

where \vec{f} is the column vector of length N representing an object, \vec{g} is the column vector of length M representing the compressive measurements, and \mathbf{H} is the measurement matrix of size $M \times N$ whose row vectors are the projection vectors. As noted previously, an object can be expressed in the sparse DCT basis (represented by transform matrix \mathbf{D}) as

$$\vec{\theta} = \mathbf{D}\vec{f} \quad \vec{f} = \mathbf{D}^T\vec{\theta}. \tag{11}$$

Recall that $\vec{\theta}$ is a column vector of length N whose elements are the sparse DCT coefficients.

To construct an optimal hybrid measurement basis, $I - 1$ images from the image data set ($I = 17$) are used to estimate the power-law parameters, such as R, p, C_p , and constant C_k , from Eq. (5) using a least-squares fit. Here, we choose an image block size of 32×32 (i.e., $N = 1024$) with a sparsity of $\alpha = 85\%$. We estimate the expected value of the random variable X_M , which is upper bounded by the P , using overlapping image blocks from the $I - 1$ image set. Note that, for $N = 1024$ and $\alpha = 85\%$, we obtain $P = 154$. We find that following relation provides a good empirical fit to the estimated $\mathbb{E}(X_M)$:

$$\mathbb{E}(X_M) = K \left(1 - \exp\left(-\frac{M}{\beta}\right) \right), \tag{12}$$

where β is estimated using a least-squares fit to the data. Given M_1 DCT projections and M_2 random projections, the compressive measurement vector can be expressed as

$$\vec{g} = \begin{bmatrix} \vec{g}_d \\ \vec{g}_c \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{\text{dct}} \\ \mathbf{H}_{\text{rand}} \end{bmatrix} \mathbf{D}^T \begin{bmatrix} \vec{\theta}_d \\ \vec{\theta}_c \end{bmatrix}, \tag{13}$$

where $\mathbf{H}_{\text{dct}} = \mathbf{D}_{M_1}$ and $\mathbf{H}_{\text{rand}} = \Phi_{\text{rand}} \mathbf{D}_{N-M_1}^T \mathbf{D}_{N-M_1}$, \mathbf{D}_{M_1} contains the first M_1 rows of \mathbf{D} , \mathbf{D}_{N-M_1} is composed of the last $N - M_1$ rows of \mathbf{D} , and Φ_{rand} matrix's elements are independent random numbers generated from a standard normal distribution.

Here, we employ a nonlinear ℓ_1 norm-based reconstruction method, described in [4], for the random subset of measurements. This nonlinear reconstruction method attempts to minimize the ℓ_1 norm of the estimated DCT coefficients subject to the data measurement constraint. The algorithm can be mathematically described as

$$\min |\vec{\theta}_c|_{\ell_1} \quad \text{s.t.} \quad \mathbf{H}_{\text{rand}} \mathbf{D}^T \vec{\theta}_c = \vec{g}_c. \tag{14}$$

We use the ℓ_1 magic algorithm implementation here [20]. The DCT subset of measurements is reconstructed using a linear minimum mean square error method where the reconstruction can be expressed as

$$\hat{\vec{\theta}}_d = \mathbf{D} \mathbf{H}_{\text{dct}}^T (\mathbf{D} \mathbf{H}_{\text{dct}}^T \mathbf{H}_{\text{dct}} \mathbf{D}^T)^{-1} \vec{g}_d. \tag{15}$$

Given the reconstruction error model defined in Eq. (9), the optimal hybrid basis is constructed by finding the optimal M_1 for each candidate value of M . Figure 5 shows a plot of the

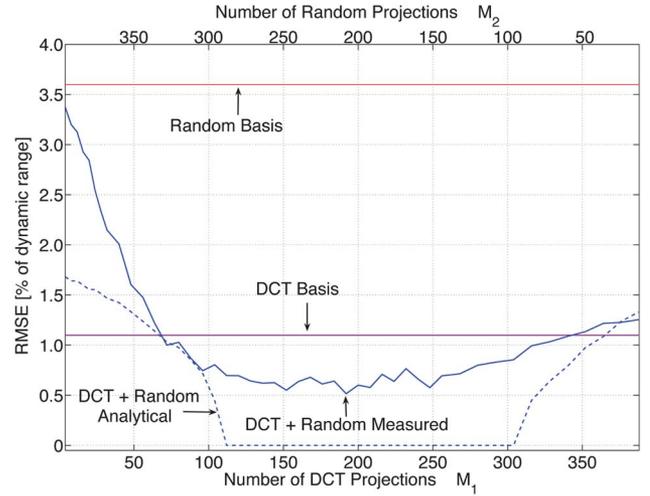


Fig. 5. (Color online) Reconstruction error versus the number of DCT projections (M_1) comprising the hybrid basis and DCT/random projections with $M = 400$.

reconstruction error as a function of M_1 for $M = 400$, as predicted by the reconstruction error model along with the actual observed error for the hybrid basis. The reconstruction error associated with DCT and random bases is also plotted for reference. Note that the reconstruction error model data is obtained with a training set of 16 images and the actual error is computed using a test image that is distinct from the training set. In this particular case, there is no unique value of M_1 that minimizes the reconstruction error. Therefore, we use the mean value of the optimal set of M_1 to define M_1^{opt} . Note that, while the model error is different from the actual observed error, they both show similar qualitative trends with M_1 . For example, the model predicts $M_1^{\text{opt}} = 200$, whereas the actual is $M_1^{\text{opt}} = 190$.

The optimal mixture of DCT (M_1) and random (M_2) projections comprising the hybrid basis is obtained by averaging all possible combinations of training sets and test images for each value of M . The resulting average reconstruction error performance of the optimal hybrid basis is plotted in Fig. 6 as a function of M along with that of the random and DCT bases. Here, we use a normalized rms error (RMSE) metric for reconstruction error, which is defined

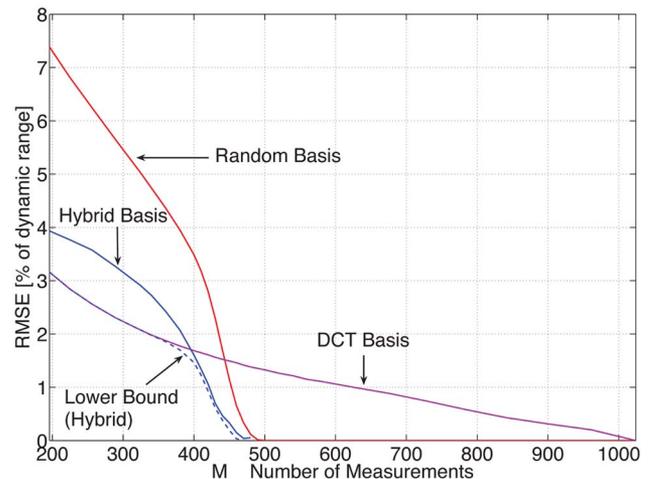


Fig. 6. (Color online) Reconstruction error versus the number of measurements for hybrid, DCT, and random bases.

as $\text{RMSE} = \sqrt{\frac{\mathbb{E}(\|\tilde{f} - \hat{f}\|_{\ell_2})}{DR}} \times 100\%$, where DR is the object dynamic range. Note that the reconstruction error is also averaged over multiple realizations of the random projections comprising the hybrid and random bases. In Fig. 6, note that, for relatively small values of M , the DCT basis actually outperforms the random basis due to the nonsparse nature of DCT coefficients at low spatial frequencies. For example, at $M = 256$, the DCT basis achieves a $\text{RMSE} = 2.56\%$, whereas the random basis incurs a $\text{RMSE} = 6.23\%$, which is nearly three times worse. As the sparsity of the DCT coefficients increases with M , the random basis becomes progressively more efficient, achieving perfect reconstruction at $M = 500$ measurements while the DCT basis yields a nonzero error of 1.33%. As expected, the DCT basis requires the maximum number of measurements $M = N = 1024$ to achieve zero error. We observe that the reconstruction error for the hybrid basis is significantly lower than that of the random basis for all values of M . It is somewhat surprising to note that, for a relatively small number of measurements (e.g., $M = 256$), the hybrid basis yields reconstructions that have a worse error ($\text{RMSE} = 3.57\%$) compared to the DCT basis ($\text{RMSE} = 2.56\%$). To understand this behavior, recall that the error model [Eq. (9)] for the hybrid basis is upper bound on the reconstruction error. Therefore, it is to be expected that the resulting mixture of DCT and random projections based on this model would be near-optimal but not optimal and, as a result, may not achieve the minimum reconstruction error. In the simulation study, we can determine the minimum reconstruction error for the hybrid basis precisely by conducting an exhaustive search for the M_1 (for a given M) that minimizes the actual observed error. This minimum reconstruction error estimate serves to lower bound the performance of the hybrid basis and is shown here by the dashed blue curve in Fig. 6. For small values of M , we note that the lower bound on the reconstruction error for the hybrid basis actually coincides with the DCT basis, suggesting that the DCT is the optimal basis in this region. However, as M increases, we observe that the hybrid basis outperforms both the DCT and random bases. Furthermore, we observe that the designed hybrid basis performance approaches its lower bound as M increases. Table 1 summarizes the reconstruction performance of the three measurement bases for various values of M . Note that the hybrid basis consistently outperforms the random basis for all values of M . For example, at $M = 400$, the hybrid basis provides a reconstruction improvement of 54% compared to the random basis, while, at $M = 460$, this improvement increases to nearly 77%. Equivalently, we can also quantify the performance of the hybrid basis by observing the number of measurements required to achieve a particular reconstruction error. For an RMSE of 1%, the hybrid basis requires $M = 420$, while the random and DCT bases require $M = 451$ and $M = 625$, respectively. Similarly, for an RMSE of 3%,

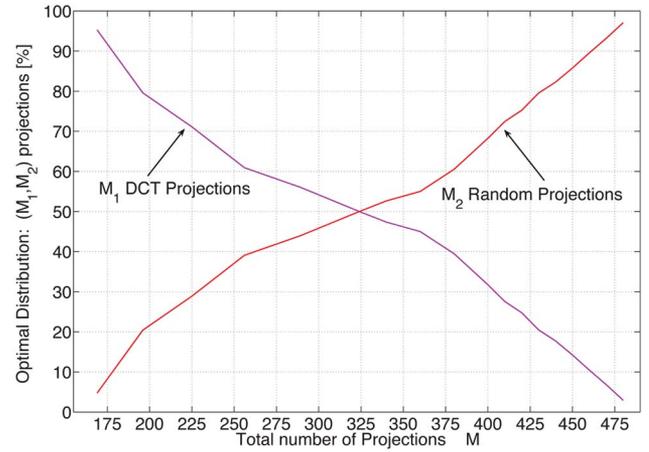


Fig. 7. (Color online) Optimal number of DCT (M_1) and random (M_2) projections comprising the hybrid basis.

the random basis needs $M = 414$ compared to $M = 315$ for the hybrid basis.

It is also instructive to inspect the mixture of DCT and random projections comprising the hybrid basis to gain further insights into its performance. Figure 7 shows a plot of the optimal number of DCT projections M_1 and random projections M_2 in the hybrid basis as a function of M . As expected, for small M , the DCT generates the majority of the projections within the hybrid basis. As M increases, DCT coefficients associated with increasingly higher spatial frequencies become more sparse. Therefore, one would expect that an optimal mixture of DCT and random projections would include an increasing number of random projections. This is indeed the case, as shown by the solid red curve in Fig. 7, the proportion of random projections increases with M . Further, it should also be noted that the hybrid basis achieves perfect reconstruction at slightly smaller $M = 470$ compared to $M = 500$ for a purely random basis. This is due to the fact that the hybrid basis exploits the nearly perfect nonsparsity of the first few DCT coefficients near DC to offset the number of random projections required to estimate the remaining DCT nonzero coefficients.

So far, we have only discussed the reconstruction error performance of the three compressive measurement bases quantitatively. The visual quality of the reconstructed images is also an important element of the imaging performance of a compressive imaging system. Figure 8 shows the reference image and example reconstruction images obtained with $M = 400$ measurements for the hybrid, DCT, and random bases. Observe that the DCT image reconstruction is more blurred relative to the other two images and, of course, the reference image. Recall that the DCT reconstruction is essentially a low-frequency representation of the object (high-frequency coefficients are truncated). This is clearly illustrated by the

Table 1. RMSE Performance of Three Compressive Measurement Bases for Various Values of M^a

$M_1^{\text{opt}}/M_2^{\text{opt}} \rightarrow$	174/82	174/166	174/226	97/333	59/401	41/439	25/475
$\downarrow \text{RMSE } M \rightarrow$	256	340	400	430	460	480	500
Random Basis	6.25%	4.74%	3.49%	2.30%	0.65%	0.1%	0%
DCT Basis	2.56% (59%)	1.98% (58%)	1.69% (52%)	1.56% (37%)	1.46% (-124%)	1.38% (-1280%)	1.33%
Hybrid Basis	3.59% (43%)	2.73% (43%)	1.61% (54%)	0.70% (70%)	0.15% (77%)	0% (100%)	0%

^aNumber in brackets indicates improvement relative to random basis.

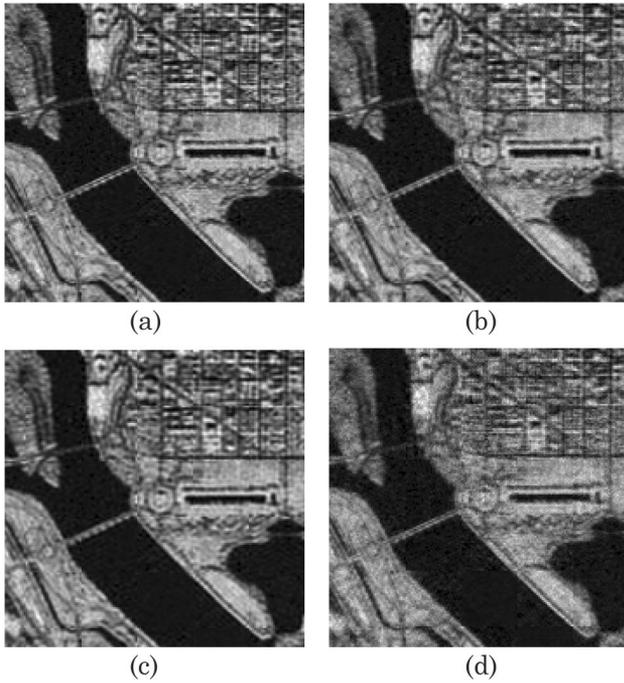


Fig. 8. Example reconstructions obtained with $M = 400$ compressive measurements. Clockwise from top left: reference image and reconstructions with hybrid basis, random basis, and DCT basis.

bridge across the river (middle section) in the image, which is resolved into two distinct lanes in the hybrid and random reconstructions but not in the DCT reconstruction. Further, note that, in the top half of the reconstructed images, the DCT and hybrid reconstructions offer higher fidelity than the random reconstruction. Example reconstructions in Fig. 9 for a different objects also show similar visual quality trends

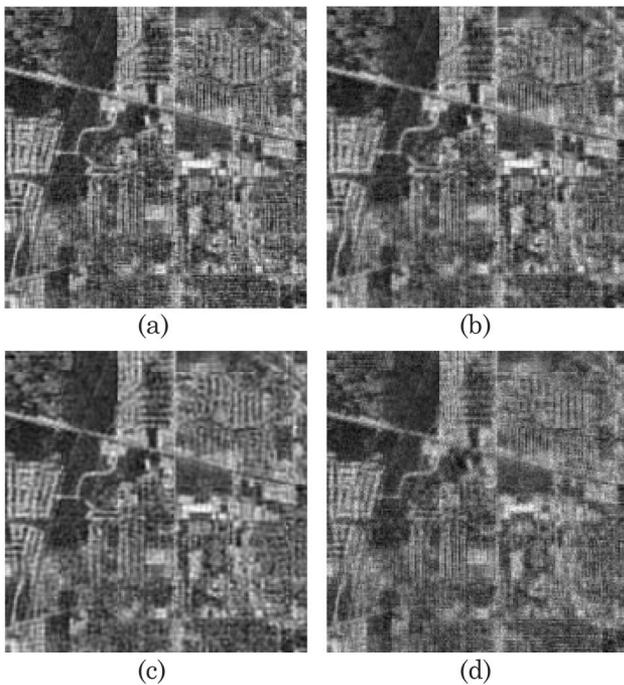


Fig. 9. Example reconstructions obtained with $M = 400$ compressive measurements. Clockwise from top left: reference image and reconstructions with hybrid basis, random basis, and DCT basis.

and reaffirm our conclusions regarding the performance of the hybrid basis.

B. Different Object Sparsity

So far, we have considered objects with a sparsity level of $\alpha = 85\%$ in the simulation study. Because objects can have significantly different sparsity, it is important to understand the performance of the hybrid basis as a function of sparsity. Here, we consider four other levels of sparsity to emulate objects with varying levels of sparsity in the DCT domain. The reconstruction error for these four different sparsity level objects is plotted in Fig. 10 as a function of number of measurements. Observe that, as the sparsity level α decreases, more measurements are required for perfect reconstruction for both the hybrid and random bases. Note that the hybrid basis maintains a superior reconstruction performance compared to the random basis for all sparsity levels. The difference between the hybrid and random bases reconstruction error, however, reduces with decreasing sparsity level. For example, at $\alpha = 85\%$ and $M = 2P = 307$, the random RMSE = 5.34% compared to RMSE = 3.1% for the hybrid basis, a relative improvement of 42%; while at $\alpha = 70\%$ and $M = 2P = 614$, the random RMSE = 2.73% and the hybrid RMSE = 1.82%, which is a smaller 33% relative improvement. This reduced difference between the hybrid and random bases performance can be understood by considering the decreasing sparsity of DCT coefficients with increasing spatial frequencies. As mentioned earlier, the main idea behind the design of the hybrid basis is to exploit the sparsity structure of DCT coefficients. Therefore, for objects with higher sparsity, the relative proportion of nonzero DCT coefficients corresponding to lower frequencies is higher compared to objects with lower sparsity. Thus, the resulting hybrid basis design tends to be more efficient in the error per measurement sense for objects with higher sparsity. Therefore, the hybrid basis performance improves with increasing sparsity relative to the random basis. This fact is confirmed by the data tabulated in Table 2 summarizing the performance of the three measurement bases for the five different object sparsity levels. For example, at $\alpha = 95\%$ the hybrid basis requires 18% fewer measurements to achieve an RMSE of 2.0% compared to the random basis, while, at $\alpha = 70\%$, the hybrid basis needs only 4% fewer measurements than random basis to achieve the same error.

C. Noisy Compressive Reconstruction

In any physical compressive imager, there are various sources of noise that can affect a compressive measurement, such as photon noise, thermal noise, readout noise, quantization noise, etc. Here, we consider a detector-noise-limited compressive imager model. The noise is composed of the readout (thermal) and quantization noise that can be effectively modeled as additive white Gaussian noise (AWGN) with zero mean and variance σ_n^2 . The noise standard deviation σ_n is defined as a percentage of the dynamic range of a conventional imager. Mathematically, a noisy compressive measurement can be expressed as

$$\vec{g} = \mathbf{H}\vec{f} + \vec{n}, \quad (16)$$

where \vec{n} is the measurement noise column vector of length M representing the independently and identically distributed

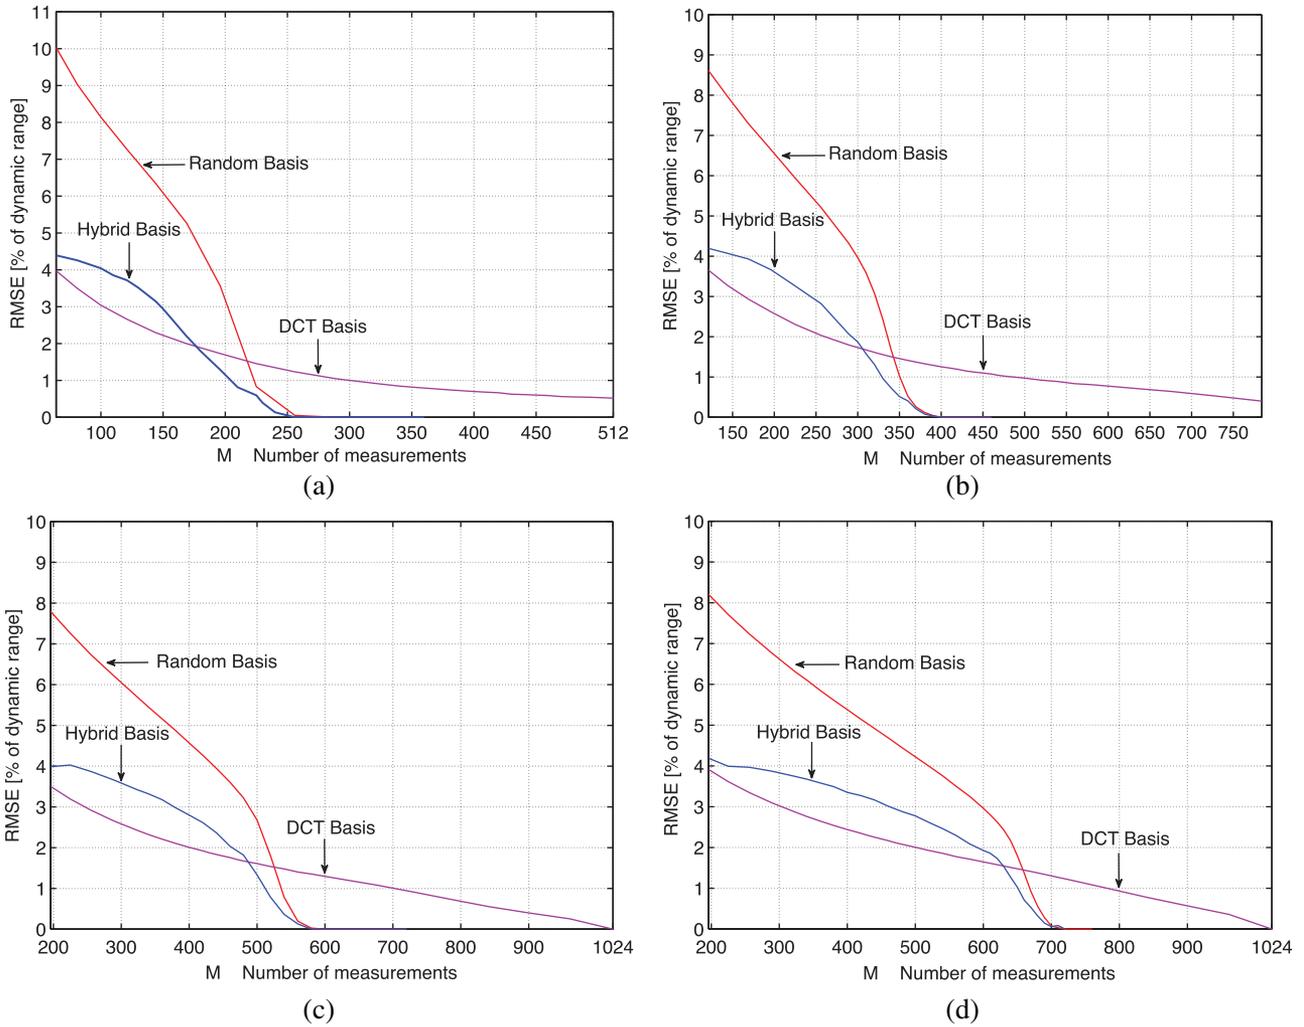


Fig. 10. (Color online) Reconstruction error versus the number of measurements for hybrid, DCT, and random bases for four different sparsity levels: (a) $\alpha = 95\%$, (b) $\alpha = 90\%$, (c) $\alpha = 80\%$, and (d) $\alpha = 70\%$.

detector noise in each compressive measurement. It is important to note that, in a passive imaging system, the total photon count incident on the imaging aperture remains constant for a fixed exposure time as the number of measurements vary. This fixed photon constraint can be mathematically expressed as

$$\mathbf{H} = \frac{\mathbf{H}^*}{C_M}, \tag{17}$$

where $C_M = \max_j \sum_{i=1}^M |H^*|_{ij}$ and \mathbf{H}^* is the measurement matrix before scaling. Note that the scaling constant C_M is a function of the number of measurements. A compressive imaging system can be physically implemented in different ways as

Table 2. Number of Measurements M ($N = 1024$) Required to Achieve a Desired RMSE Performance for Objects with Five Different Sparsity Levels ^a

↓ RMSE Basis →	Sparsity Level														
	$\alpha = 95\%$			$\alpha = 90\%$			$\alpha = 85\%$			$\alpha = 80\%$			$\alpha = 70\%$		
	R	D	H	R	D	H	R	D	H	R	D	H	R	D	H
RMSE = 3.0%	202	102	148	321	169	244	414	209	315	490	245	376	600	304	460
RMSE = 2.5%	206	130	161	329	206	271	426	263	354	503	309	428	626	387	534
RMSE = 2.0%	211	169	173	335	260	293	434	337	381	514	401	462	645	501	589
RMSE = 1.5%	216	220	189	341	341	312	444	448	404	525	534	492	656	644	631
RMSE = 1.0%	222	299	203	350	480	328	451	625	420	534	700	512	666	783	650
RMSE = 0.5%	234	534	227	360	741	350	464	813	437	547	861	532	681	918	670
RMSE = 0.0%	256	1024	246	400	1024	390	500	1024	470	600	1024	580	720	1024	700

^aR, random; D, DCT; H, hybrid.

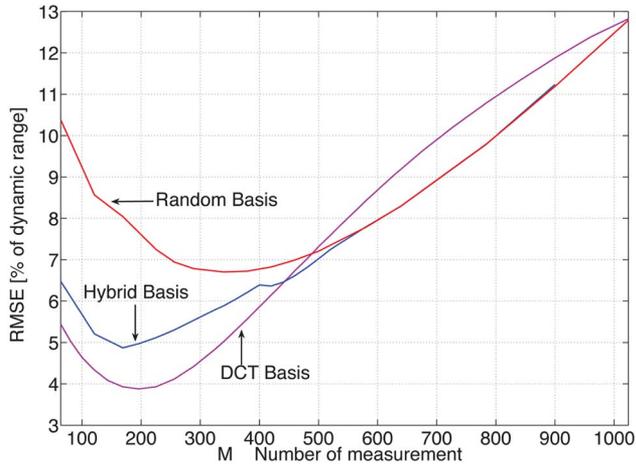


Fig. 11. (Color online) Reconstruction RMSE versus the number of noisy measurements with a noise strength $\sigma_n = 0.5\%$.

described in [12]. Here, we assume a sequential implementation where each measurement is acquired sequentially.

The object reconstruction technique for the hybrid and random measurements is based on minimizing the total variation (TV) subject to the quadratic data measurement error constraint, which can be mathematically expressed as

$$\min \text{TV}(\hat{\theta}) \quad \text{s.t.} \quad \|\mathbf{H}\mathbf{D}^T\hat{\theta} - \vec{g}\|_{\ell_2} \leq \epsilon, \quad (18)$$

where TV denotes the TV of a reconstruction and ϵ is the error in the measurement process ($\epsilon = \sigma_n^2$). TV is defined in terms of the first-order horizontal and vertical gradients of a candidate reconstruction [21]. We employ the TV reconstruction algorithm from the ℓ_1 magic toolbox here [20]. For DCT measurements, the linear minimum square error operator is used to obtain the following object estimate:

$$\hat{\theta} = \sigma_s^2 \mathbf{D}\mathbf{H}^T (\sigma_s^2 \mathbf{D}\mathbf{H}^T \mathbf{H}\mathbf{D}^T + \sigma_n^2 \mathbf{I})^{-1} \vec{g}, \quad (19)$$

where σ_s^2 is the second moment of the object estimated from the image data set and \mathbf{I} is a $M \times M$ identity matrix.

The reconstruction performance of each candidate measurement basis is analyzed in presence of measurement noise. The object sparsity level is set to $\alpha = 85\%$ and $\sigma_n = 0.5\%$ in this study. Figure 11 shows the reconstruction RMSE (averaged over 20 noise realizations and 12 random projection realizations) with the number of measurements. Observe that, for all bases, the reconstruction error decreases initially, reaching a minima, and then increases monotonically with the number of measurements. This behavior results from two underlying mechanisms. First, the truncation/approximation error decreases monotonically as more measurements are acquired. However, a fixed photon budget (i.e., a fixed exposure time)

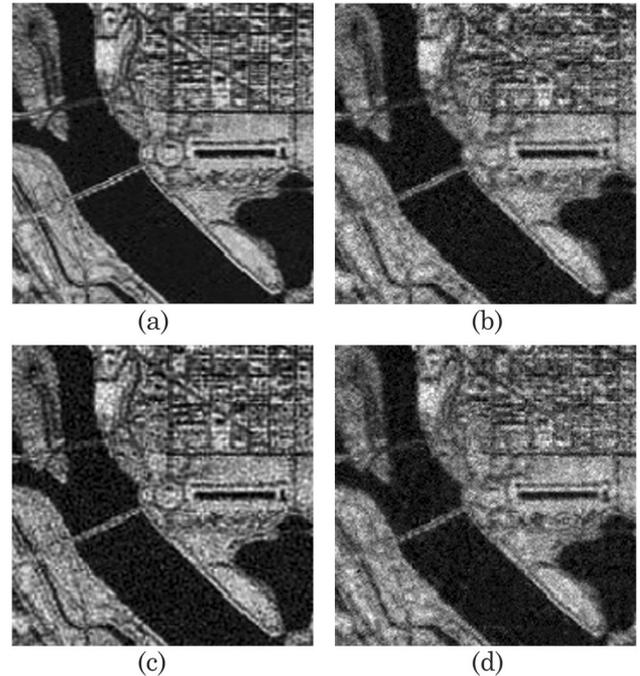


Fig. 12. Example reconstructions obtained with $M = 460$ noisy compressive measurements. Clockwise from top left: reference image and reconstructions with hybrid basis, random basis, and DCT basis.

implies that the photon allocation per measurement decreases with increasing measurement dimensionality, while the detector noise remains fixed. As a result, the measurement fidelity (i.e., measurement signal-to-noise ratio) decreases monotonically with M . The tradeoff between these two error mechanisms leads to the nonmonotonic RMSE behavior [10]. Overall, the hybrid basis provides a lower reconstruction RMSE relative to the random basis. For example, at $M = 169$, the random basis yields an RMSE of 8.0% relative to an RMSE of 4.9% for the hybrid basis, an improvement of nearly 40%. Furthermore, the minimum reconstruction error for the hybrid basis (RMSE = 4.9%) is significantly lower than the error achievable with the random basis (RMSE = 6.7%). Similar to the noiseless case, the DCT basis provides the best reconstruction performance for small values of M for two reasons. First, the linear reconstruction technique used for DCT reconstruction is more robust to noise compared to the nonlinear reconstruction method employed for hybrid and random bases [12]. Second, the hybrid design for small value of M is suboptimal as described in Subsection 3.A. However, as the number of measurements increases, the hybrid basis yields a lower reconstruction error (RMSE = 7.0%) than that of the DCT basis at $M = 500$ (RMSE = 7.32%). Table 3 summarizes the RMSE performance of the three bases. Figures 12 and 13 show example reconstructions for two objects in the

Table 3. RMSE Performance of the Three Measurement Bases for Noisy Measurements^a

↓RMSE M →	169	256	340	420	460	500	580
Random Basis	8.04%	6.95%	6.7%	6.83%	6.99%	7.21%	7.8%
DCT Basis	3.93% (51.1%)	4.21% (39.4%)	5.03% (24.9%)	6.15% (9.9%)	6.74% (3.6%)	7.32% (-1.5%)	8.4% (-7.7%)
Hybrid Basis	4.87% (39.4%)	5.3% (23.7%)	5.89% (12.1%)	6.37% (7.3%)	6.62% (5.3%)	7.02% (2.6%)	7.8% (0%)

^aNumber in brackets indicates improvement relative to random basis.

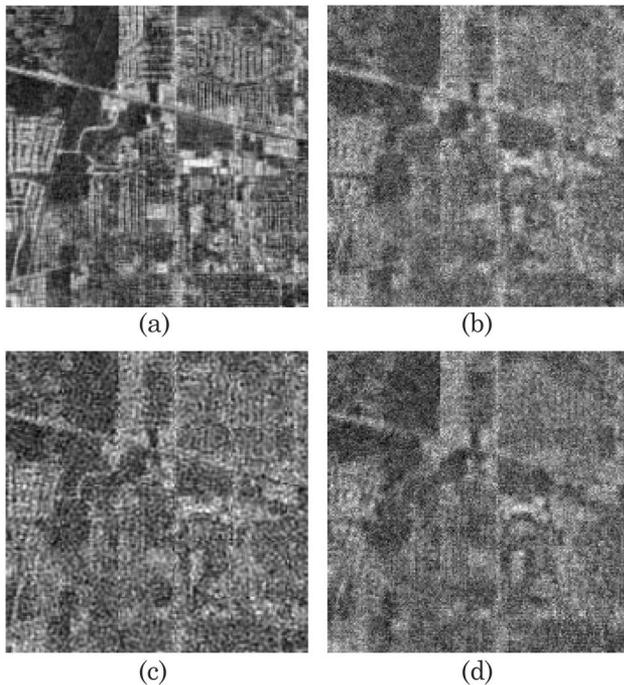


Fig. 13. Example reconstructions obtained with $M = 460$ noisy compressive measurements. Clockwise from top left: reference image and reconstructions with hybrid basis, random basis, and DCT basis.

image data set. We observe that the hybrid reconstructions are consistently superior to random reconstructions. This affirms the improved reconstruction performance of the hybrid basis relative to the random basis in the presence of noise.

4. CONCLUSIONS

The sparse nature of natural scenes has been exploited in compressive imaging via the use of random measurement basis. However, prior knowledge about the sparsity structure of natural scenes, such as the well-known power-law spectral density models, has not been utilized for measurement basis design yet to our knowledge. In this work, we described the design of a hybrid measurement basis that employs a combination of a known sparsity basis, such as DCT, and the random projection basis to improve reconstruction performance for natural scenes compared to the random basis. Specifically, the optimal combination of DCT and random projections is derived from the ensemble object power spectral density estimated from a diverse set of images. The reconstruction performance is quantified for each image in the image set (using the leave-one-out testing method) using a simulation study. The simulation results demonstrate that the reconstruction performance is significantly improved with the hybrid basis relative to DCT and random bases. For example, the hybrid basis achieves a reconstruction error of 0.15% compared to an RMSE of 0.65% for the random basis with $M = 460$ measurements ($N = 1024$, $\alpha = 85\%$), which is an improvement of nearly 77%. Furthermore, we also observe that the hybrid basis design is robust to varying levels of object sparsity and measurement noise. The hybrid basis maintained superior reconstruction performance relative to the random and DCT bases in the presence of measurement noise. In fact, with 50% ($M = 500$) measurements, the random basis yielded a RMSE of 8%, while the hybrid basis achieved a nearly 40% low-

er RMSE of 4.9%. It should be emphasized that the hybrid basis design method described here is also applicable to other non-random sparsity bases besides DCT, such as wavelets or any other sparsity basis, where prior knowledge about the sparsity structure is available. Furthermore, a structured measurement basis, such as the hybrid basis, may also be attractive from the practical implementation perspective in terms of an efficient measurement basis representation within the compressive imager hardware. A compressive imager with a hybrid measurement basis can be implemented using one of the several architectures described in [12,22–25] that employ active or passive optical modulation mechanisms to acquire compressive measurements. As an extension to this work, we are actively pursuing the development of more efficient adaptive compressive measurement bases that incorporate prior knowledge and information from past measurements within a sequential measurement framework.

ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge support for this research under the Defense Advanced Research Projects Agency (DARPA) grant N66001-10-C-4051.

REFERENCES

1. E. T. Whittaker, "On the functions which are represented by the expansions of the interpolation theory," *Proc. R. Soc. Edinb., Sect. A, Math.* **35**, 181–194 (1915).
2. H. Nyquist, "Certain topics in telegraph transmission theory," *Trans. AIEE* **47**, 617–644 (1928).
3. C. E. Shannon, "Communication in the presence of noise," *Proc. IRE* **37**, 10–21 (1949).
4. Y. Tsaig and D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006).
5. D. Donoho and Y. Tsaig, "Extensions of compressed sensing," *Signal Process.* **86**, 549–571 (2006).
6. E. Candes and J. Romberg, "Signal recovery from random projections," *Proc. SPIE* **5674**, 76–86 (2005).
7. E. Candes and T. Tal, "Near-optimal signal recovery from random projections: universal encoding strategies?" *IEEE Trans. Inf. Theory* **52**, 5406–5424 (2006).
8. F. Krahermer and R. Ward, "New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property," submitted for publication (2010), arXiv:1009.0744v4.
9. J. Romberg, "Imaging via compressing sampling," *IEEE Signal Process. Mag.* **25**, 14–20 (2008).
10. M. A. Neifeld and P. Shankar, "Feature-specific imaging," *Appl. Opt.* **42**, 3379–3389 (2003).
11. H. Pal and M. A. Neifeld, "Multispectral principal component imaging," *Opt. Express* **11**, 2118–2125 (2003).
12. M. A. Neifeld and J. Ke, "Optical architectures for compressive imaging," *Appl. Opt.* **46**, 5293–5303 (2007).
13. D. J. Tolhurst, Y. Tadmor, and T. Chao, "Amplitude spectra of natural images," *Ophthalmic. Physiol. Opt.* **12**, 229–232 (1992).
14. D. L. Ruderman, "Origins of scaling in natural images," *Vision Res.* **37**, 3385–3398 (1997).
15. E. P. Simoncelli and B. A. Olshausen, "Natural image statistics and neural representation," *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
16. E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Probl.* **23**, 969–986 (2007).
17. W. Chen and W. Pratt, "Scene adaptive coder," *IEEE Trans. Commun.* **32**, 225–232 (1984).
18. G. Wallace, "The JPEG still picture compression standard," *Commun. ACM* **34**, 30–44 (1991).
19. University of Southern California Signal and Image Processing Institute, "The USC-SIPI image database," <http://sipi.usc.edu/database>.
20. California Institute of Technology, " ℓ_1 -MAGIC," <http://www.acm.caltech.edu/l1magic/>.

21. S. D. Babacan, R. Molina, and A. K. Katsaggelos, "Variational Bayesian blind deconvolution using a total variation prior," *IEEE Trans. Image Process.* **18**, 12–26 (2009).
22. P. K. Baheti and M. A. Neifeld, "Feature-specific structured imaging," *Appl. Opt.* **45**, 7382–7391 (2006).
23. N. P. Pitsianis, D. J. Brady, and X. Sun, "Sensor-layer image compression based on the quantized cosine transform," *Proc. SPIE* **5817**, 250–257 (2005).
24. D. J. Brady, N. P. Pitsianis, X. Sun, and P. Potuluri, "Compressive sampling and signal inference," U.S. patents 7283231 (16 Oct. 2007), 7432843 (7 Oct. 2008), 7463174 (9 Dec. 2008), 7463179 (9 Dec. 2008), 7616306 (10 Nov. 2009).
25. M. E. Gehm, S. T. McCain, N. P. Pitsianis, D. J. Brady, P. Potuluri, and M. E. Sullivan, "Static two-dimensional aperture coding for multimodal multiplex spectroscopy," *Appl. Opt.* **45**, 2965–2974 (2006).