

The quest for optimal sampling: Computationally efficient, structure-exploiting measurements for compressed sensing

Ben Adcock, Anders C. Hansen and Bogdan Roman

Abstract An intriguing phenomenon in many instances of compressed sensing is that the reconstruction quality is governed not just by the overall sparsity of the signal, but also on its structure. This paper is about understanding this phenomenon, and demonstrating how it can be fruitfully exploited by the design of suitable sampling strategies in order to outperform more standard compressed sensing techniques based on random matrices.

1 Introduction

Compressed sensing concerns the recovery of signals and images from a small collection of linear measurements. It is now a substantial area of research, accompanied by a mathematical theory that is rapidly reaching a mature state. Applications of compressed sensing can roughly be divided into two areas. First, *type I* problems, where the physical device imposes a particular type of measurements. This is the case in numerous real-world problems, including medical imaging (e.g. Magnetic Resonance Imaging (MRI) and Computerized Tomography (CT)), electron microscopy, seismic tomography and radar. Second, *type II* problems, where the sensing mechanism allows substantial freedom to design the measurements so as to

Ben Adcock
Purdue University, Department of Mathematics, 150 N. University St, West Lafayette, IN 47906,
USA, e-mail: adcock@purdue.edu

Anders C. Hansen
DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Rd, Cambridge
CB3 0WA, United Kingdom e-mail: a.hansen@maths.cam.ac.uk

Bogdan Roman
DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Rd, Cambridge
CB3 0WA, United Kingdom e-mail: b.roman@maths.cam.ac.uk

improve the reconstructed image or signal. Applications include compressive imaging and fluorescence microscopy.

This paper is devoted to the role of structured sparsity in both classes of problems. It is well known that the standard sparsifying transforms of compressed sensing, i.e. wavelets and their various generalizations, not only give sparse coefficients, but that there is a distinct structure to this sparsity: wavelet coefficients of natural signals and images are far more sparse at fine scales than at coarse scales. For type I problems, recent developments [2, 4] have shown that this structure plays a key role in the observed reconstruction quality. Moreover, the optimal subsampling strategy depends crucially on this structure. We recap this work in this paper.

Since structure is vitally important in type I problems, it is natural to ask whether or not it can be exploited to gain improvements in type II problems. In this paper we answer this question in the affirmative. We show that appropriately-designed structured sensing matrices can successfully exploit structure. Doing so leads to substantial gains over classical approaches in type II problems, based on convex optimization with universal, random Gaussian or Bernoulli measurements, as well as more recent structure-exploiting algorithms – such as model-based compressed sensing [5], TurboAMP [46], and Bayesian compressed sensing [29, 30] – which incorporate structure using bespoke recovery algorithms. The proposed matrices, based on appropriately subsampled Fourier/Hadamard transforms, are also computationally efficient

We also review the theory of compressed sensing introduced in [2, 6] with such structured sensing matrices. The corresponding sensing matrices are highly non-universal and do not satisfy a meaningful Restricted Isometry Property (RIP). Yet, recovery is still possible, and is vastly superior to that obtained from standard RIP matrices. It transpires that the RIP is highly undesirable if one seeks to exploit structure by designing appropriate measurements. Thus we consider an alternative that takes such structure into account, known as the *RIP in levels* [6].

2 Recovery of wavelet coefficients

In many applications of compressed sensing, we are faced with the problem of recovering an image or signal x , considered as a vector in \mathbb{C}^n or a function in $L^2(\mathbb{R}^d)$, that is sparse or compressible in an orthonormal basis of wavelets. If $\Phi \in \mathbb{C}^{n \times n}$ or $\Phi \in \mathcal{B}(L^2(\mathbb{R}^d), \ell_2(\mathbb{N}))$ (the set of bounded linear operators) is the corresponding sparsifying transformation, then we write $x = \Phi c$, where $c \in \mathbb{C}^n$ or $c \in \ell_2(\mathbb{N})$ is the corresponding sparse or compressible vector of coefficients. Given a sensing operator $A \in \mathbb{C}^{m \times n}$ or $A \in \mathcal{B}(L^2(\mathbb{R}^d), \mathbb{C}^m)$ and noisy measurements $y = Ax + e$ with $\|e\|_2 \leq \eta$, the usual approach is to solve the ℓ_1 -minimization problem:

$$\min_{z \in \mathbb{C}^n} \|\Phi z\|_1 \quad \text{s.t.} \quad \|y - Az\|_2 \leq \eta. \quad (1)$$

or

$$\inf_{z \in L^2(\mathbb{R}^d)} \|\Phi z\|_1 \quad \text{s.t.} \quad \|y - Az\|_2 \leq \eta. \quad (2)$$

Throughout this paper, we shall denote a minimizer of (1) or (2) as \hat{x} . Note that (2) must be discretized in order to be solved numerically, and this can be done by restricting the minimization to be taken over a finite-dimensional space spanned by the first n wavelets, where n is taken sufficiently large [1].

As mentioned, compressed sensing problems arising in applications can be divided into two classes:

- I. *Imposed sensing operators.* The operator A is specified by the practical device and is therefore considered fixed. This is the case in MRI – where A arises by subsampling the Fourier transform [38, 39] – as well as other examples, including X-ray CT (see [14] and references therein), radar [31], electron microscopy [7, 35], seismic tomography [37] and radio interferometry [53].
- II. *Designed sensing operators.* The sensing mechanism allows substantial freedom to design A so as to improve the compressed sensing reconstruction. Some applications belonging to this class are *compressive imaging*, e.g. the single-pixel camera [21] and the more recent lensless camera [32], and compressive fluorescence microscopy [47]. In these applications A is assumed to take binary values (typically $\{-1, 1\}$), yet, as we will see later, this is not a significant practical restriction.

As stated, the purpose of this paper is to show that insight gained from understanding the application of compressed sensing to type I problems leads to more effective strategies for type II problems.

2.1 Universal sensing matrices

Let us consider type II problems. In finite dimensions, the traditional compressed sensing approach has been to construct matrices A possessing the following two properties. First, they should satisfy the *Restricted Isometry Property (RIP)*. Second, they should be *universal*. That is, if $\Phi \in \mathbb{C}^{n \times n}$ is an arbitrary isometry, then $A\Phi$ also satisfies the RIP of the same order as A . Subject to these conditions, a typical result in compressed sensing is as follows (see [24], for example): if A satisfies the RIP of order $2k$ with constant $\delta_{2k} < 4/\sqrt{41}$ then, for any $x \in \mathbb{C}^n$, we have

$$\|x - \hat{x}\|_2 \leq C \frac{\sigma_k(\Phi^* x)_1}{\sqrt{k}} + D\eta, \quad (3)$$

where \hat{x} is any minimizer of (1), C and D are positive constants depending only on δ_{2k} and, for $c \in \mathbb{C}^n$,

$$\sigma_k(c)_1 = \inf_{z \in \Sigma_k} \|c - z\|_1, \quad \Sigma_k = \{z \in \mathbb{C}^n : \|z\|_0 \leq k\}.$$

Hence, x is recovered exactly up to the noise level η and the error $\sigma_k(c)_1$ of the best approximation of $c = \Phi^*x$ with a k -sparse vector. Since A is universal, one has complete freedom to choose the sparsifying transformation Φ so as to minimize the term $\sigma_k(\Phi^*x)_1$ for the particular signal x under consideration.

Typical examples of universal sensing matrices A arise from random ensembles. In particular, Gaussian or Bernoulli random matrices (with the latter having the advantage of being binary) both have this property with high probability whenever m is proportional to k times by a log factor. For this reason, such matrices are often thought of as ‘optimal’ matrices for compressed sensing.

Remark 1. One significant drawback of random ensembles, however, is that the corresponding matrices are dense and unstructured. Storage and the lack of fast transforms render them impractical for all but small problem sizes. To overcome this, various structured random matrices have also been developed and studied e.g. pseudo-random permutations of columns of Hadamard or DFT (Discrete Fourier Transform) matrices [25, 32]. Often these admit fast, $\mathcal{O}(n \log n)$ transforms. However, the best known theoretical RIP guarantees are usually larger than for (sub)Gaussian random matrices [24].

2.2 Sparsity structure dependence and the flip test

Since it will become important later, we now describe a quick and simple test, which we call the *flip test*, to investigate the presence or absence of an RIP. Success of this test suggests the existence of an RIP and failure demonstrates its lack.

Let $A \in \mathbb{C}^{m \times n}$ be a sensing matrix, $x \in \mathbb{C}^n$ an image and $\Phi \in \mathbb{C}^{n \times n}$ a sparsifying transformation. Recall that sparsity of the vector $c = \Phi^*x$ is unaffected by permutations. Thus, let us define the flipped vector

$$P(c) = c' \in \mathbb{C}^n, \quad c'_i = c_{n+1-i}, \quad i = 1, \dots, n,$$

and using this, we construct the flipped image $x' = \Phi c'$. Note that, by construction, we have $\sigma_k(c)_1 = \sigma_k(c')_1$. Now suppose we perform the usual compressed sensing reconstruction (1) on both x and x' , giving approximations $\hat{x} \approx x$ and $\hat{x}' \approx x'$. We now wish to reverse the flipping operation. Thus, we compute $\check{x} = \Phi P(\Phi^* \hat{x}')$, which gives a second approximation to the original image x .

This test provides a simple way to investigate whether or not the RIP holds. To see why, suppose that A satisfies the RIP. Then by construction, we have that

$$\|x - \hat{x}\|_2, \|x - \check{x}\|_2 \leq C \frac{\sigma_k(\Phi^*x)_1}{\sqrt{k}} + D\eta.$$

Hence both \hat{x} and \check{x} should recover x equally well. In the top row of Figure 1 we present the result of the flip test for a Gaussian random matrix. As is evident, the reconstructions \hat{x} and \check{x} are comparable, thus indicating the RIP.




Fig. 1: Recovery of an MRI image of a passion fruit using (1) from $m = 8192$ samples at $n = 256 \times 256$ resolution (i.e. a 12.5% subsampling rate) with Daubechies-4 wavelets as the sparsifying transform Φ . Top row: Flip test for Gaussian random measurements. Bottom row: Flip test for subsampled DFT measurements taken according to the subsampling pattern shown in the bottom left panel. The flip test suggests that the RIP holds for random sensing matrices (top row), but that there is no RIP for structured sensing matrices with structured sampling (bottom row).

Having considered type II problems, let us now examine the flip test for a type I problem. As discussed, in applications such as MRI, X-ray CT, radio interferometry, etc, the matrix A is imposed by the physical sensing device and arises from subsampling the rows of the DFT matrix $F \in \mathbb{C}^{n \times n}$.¹ Whilst one often has some freedom to choose which rows to sample (corresponding to selecting particular frequencies at which to take measurements), one cannot change the matrix F .

It is well known that in order to ensure a good reconstruction, one cannot subsample the DFT uniformly at random (recall that the sparsifying transform is a wavelet basis), but rather one must sample randomly according to an appropriate nonuni-

¹ In actual fact, the sensing device takes measurements of the *continuous* Fourier transform of a function $x \in L^2(\mathbb{R}^d)$. As discussed in [1, 4], modelling continuous Fourier measurements as discrete Fourier measurements can lead to inferior reconstructions, and worse, inverse crimes. To avoid this, one must consider an infinite-dimensional compressed sensing approach, as in (2). See [2, 4] for details, as well as [26] for implementation in MRI. However, for simplicity, we shall continue to work with the finite-dimensional model in the remainder of this paper.

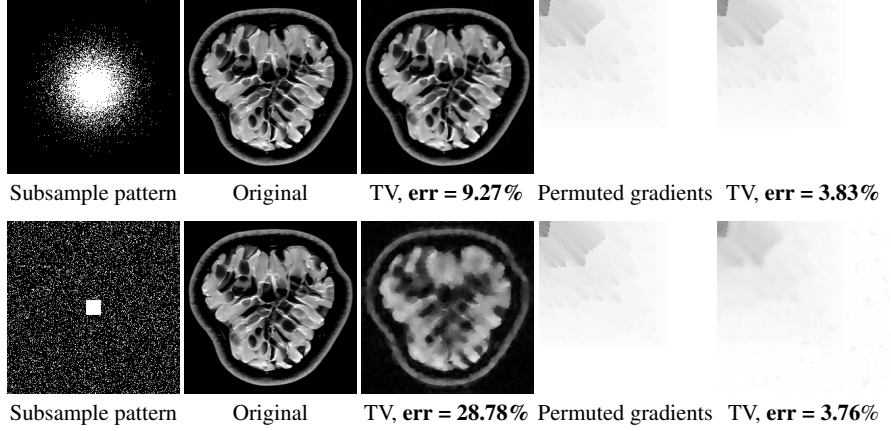


Fig. 2: TV recovery (4) from $m = 16384$ DFT samples at $n = 512 \times 512$ (6.25% subsampling rate). The *Permuted gradients* image was built from the same image gradient vectors as the *Original* image, having the same TV norm and gradient sparsity structure, differing only in the ordering and sign of the gradient vectors. The large error differences confirm that, much like the flip test for wavelet coefficients, the sparsity structure matters for TV reconstructions as well.

form density [2, 12, 39, 52]. See the bottom left panel of Figure 1 for an example of a typical density. As can be seen in the next panel, by doing so one achieves a great recovery. However, the result of the flip test in the bottom right panel clearly demonstrates that the matrix $F\Phi$ does not satisfy an RIP. In particular, the ordering of the wavelet coefficients plays a crucial role in the reconstruction quality. To explain this, and in particular, the high-quality reconstruction seen in the unflipped case, one evidently requires a new analytical framework.

Note that the flip test in Figure 1 also highlights another important phenomenon: namely, the effectiveness of the subsampling strategy depends on the sparsity structure of the image. In particular, two images with the same total sparsity (the original x and the flipped x') result in wildly different errors when the same sampling pattern is used. Thus we conclude that there is no one optimal sampling strategy for all sparse vectors of wavelet coefficients.

Let us also note that the same conclusions of the flip test hold when (1) with wavelets is replaced by TV-norm minimization:

$$\min_{z \in \mathbb{C}^n} \|z\|_{TV} \quad \text{s.t.} \quad \|y - Az\|_2 \leq \eta. \quad (4)$$

Recall that $\|x\|_{TV} = \sum_{i,j} \|\nabla x(i,j)\|_2$, where we have $\nabla x(i,j) = \{D_1 x(i,j), D_2 x(i,j)\}$, $D_1 x(i,j) = x(i+1,j) - x(i,j)$, $D_2 x(i,j) = x(i,j+1) - x(i,j)$. In the experiment leading to Figure 2, we chose an image $x \in [0, 1]^{N \times N}$, and then built a different image x' from the gradient of x so that $\{\|\nabla x'(i,j)\|_2\}$ is a permutation of $\{\|\nabla x(i,j)\|_2\}$ for which $x' \in [0, 1]^{N \times N}$. Thus, the two images have the same “TV sparsity” and the same TV norm. In Figure 2 we demonstrate how the errors differ substantially for

the two images when using the same sampling pattern. Note also how the improvement depends both on the TV sparsity structure and on the subsampling pattern. Analysis of this phenomenon is work in progress. In the remainder of this paper we will focus on structured sparsity for wavelets.

2.3 Structured sparsity

One of the foundational results of nonlinear approximation is that, for natural images and signals, the best k -term approximation error in a wavelet basis decays rapidly in k [18, 40]. In other words, wavelet coefficients are approximately k -sparse. However, wavelet coefficients possess far more structure than mere sparsity. Recall that a wavelet basis for $L^2(\mathbb{R}^d)$ is naturally partitioned into dyadic scales. Let $0 = M_0 < M_1 < \dots < \infty$ be such a partition, and note that $M_{l+1} - M_l = \mathcal{O}(2^l)$ in one dimension and $M_{l+1} - M_l = \mathcal{O}(4^l)$ in two dimensions. If $x = \Phi c$, let $c^{(l)} \in \mathbb{C}^{M_l - M_{l-1}}$ denote the wavelet coefficients of x at scale $l = 1, 2, \dots$, so that $c = (c^{(1)} | c^{(2)} | \dots)^\top$. Suppose that $\varepsilon > 0$ and define

$$k_l = k_l(\varepsilon) = \min \left\{ K : \sum_{i=1}^K |c_{\pi(i)}^{(l)}|^2 \geq \varepsilon^2 \|c^{(l)}\|_2^2 \right\}, \quad l = 1, 2, \dots, \quad (5)$$

where π is a bijection that gives a nonincreasing rearrangement of the entries of $c^{(l)}$, i.e. $|c_{\pi(i)}^{(l)}| \geq |c_{\pi(i+1)}^{(l)}|$ for $i = 1, \dots, M_l - M_{l-1} - 1$. Sparsity of the whole vector c means that for large r we have $k/n \ll 1$, where $k = k(\varepsilon) = \sum_{l=1}^r k_l$, is the total effective sparsity up to finest scale r . However, Figure 3 reveals that not only is this the case in practice, but we also have so-called *asymptotic sparsity*. That is

$$k_l / (M_l - M_{l-1}) \rightarrow 0, \quad l \rightarrow \infty. \quad (6)$$

Put simply, wavelet coefficients are much more sparse at fine scales than they are at coarse scales.

Note that this observation is by no means new: it is a simple consequence of the dyadic scaling of wavelets, and is a crucial step towards establishing the nonlinear approximation result mentioned above. However, given that wavelet coefficients always exhibit such structure, one may ask the following question. For type II problems, are the traditional sensing matrices of compressed sensing – which, as shown by the flip test (top row of Figure 1), recover all sparse vectors of coefficients equally well, regardless of ordering – optimal when wavelets are used as the sparsifying transform? It has been demonstrated in Figure 1 (bottom row) that structure plays a key role in type I compressed sensing problems. Leveraging this insight, in the next section we show that significant gains are possible for type II problems in terms of both the reconstruction quality and computational efficiency when the sensing matrix A is designed specifically to take advantage of such inherent structure.




Fig. 3: Left: GLPU phantom [27]. Right: Relative sparsity of Daubechies-4 coefficients. Here the levels correspond to wavelet scales and $k_l(\epsilon)$ is given by (5). Each curve shows the relative sparsity at level l as a function of ϵ . The decreasing nature of the curves for increasing l confirms (6).

Remark 2. Asymptotic sparsity is by no means limited to wavelets. A similar property holds for other types of -lets, including curvelets [9, 10], contourlets [19, 42] or shearlets [16, 17, 34] (see [4] for examples of Figure 3 based on these transforms). More generally, *any* sparsifying transform that arises (possibly via discretization) from a countable basis or frame will typically exhibit asymptotic sparsity.

Remark 3. Wavelet coefficients (and their various generalizations) actually possess far more structure than the asymptotic sparsity (6). Specifically, they tend to live on rooted, connected trees [15]. There are a number of existing algorithms which seek to exploit such structure within a compressed sensing context. We shall discuss these further in Section 5.

3 Efficient sensing matrices for structured sparsity

For simplicity, we now consider the finite-dimensional setting, although the arguments extend to the infinite-dimensional case [2]. Suppose that $\Phi \in \mathbb{C}^{n \times n}$ corresponds to a wavelet basis so that $c = \Phi^* x$ is not just k -sparse, but also asymptotically sparse with the sparsities k_1, \dots, k_r within the wavelet scales being known. As before, write $c^{(l)}$ for set of coefficients at scale l . Considering type II problems, we now seek a sensing matrix A with as few rows m as possible that exploits this local sparsity information.

3.1 Block-diagonality and structured Fourier/Hadamard sampling

For the l^{th} level, suppose that we assign a total of $m_l \in \mathbb{N}$ rows of A in order to recover the k_l nonzero coefficients of $c^{(l)}$. Note that $m = \sum_{l=1}^r m_l$. Consider the product $B = A\Phi$ of the sensing matrix A and the sparsifying transform Φ . Then there is a

natural decomposition of B into blocks $\{B_{jl}\}_{j,l=1}^r$ of size $m_j \times (M_l - M_{l-1})$, where each block corresponds to the m_j measurements of the $M_l - M_{l-1}$ wavelet functions at the l^{th} scale.

Suppose it were possible to construct a sensing matrix A such that (i) B was block diagonal, i.e. $B_{jl} = 0$ for $j \neq l$, and (ii) the diagonal blocks B_{ll} satisfied an RIP of order $2k_l$ whenever m_l was proportional to k_l times by the usual log factor. In this case, one recovers the coefficients $c^{(l)}$ at scale l from near-optimal numbers of measurements using the usual reconstruction (1).

This approach, originally proposed by Donoho [20] and Tsaig & Donoho [48] under the name of ‘multiscale compressed sensing’, allows for structure to be exploited within compressed sensing. Similar ideas were also pursued by Romberg [44] within the context of compressive imaging. Unfortunately, it is normally impossible to design an $m \times N$ matrix A such that $B = A\Phi$ is exactly block diagonal. Nevertheless, the notion of block-diagonality provides insight into better designs for A than purely random ensembles. To proceed, we relax the requirement of strict block-diagonality, and instead ask whether there exist practical sensing matrices A for which B is approximately block-diagonal whenever the sparsifying transform Φ corresponds to wavelets. Fortunately, the answer to this question is affirmative: as we shall explain next, and later confirm with a theorem, approximate block-diagonality can be ensured whenever A arises by appropriately subsampling the rows of the Fourier or Hadamard transform. Recalling that the former arises naturally in type I problems, this points towards the previously-claimed conclusion that new insight brought about by studying imposed sensing matrices leads to better approaches for the type II problem.

Let $F \in \mathbb{C}^{n \times n}$ be either the discrete Fourier or discrete Hadamard transform. Let $\Omega \subseteq \{1, \dots, n\}$ be an index set of size $|\Omega| = m$. We now consider choices for A of the form $A = P_\Omega F$, where $P_\Omega \in \mathbb{C}^{m \times n}$ is the restriction operator that selects rows of F whose indices lie in Ω . We seek an Ω that gives the desired block-diagonality. To do this, it is natural to divide up Ω itself into r disjoint blocks

$$\Omega = \Omega_1 \cup \dots \cup \Omega_r, \quad |\Omega_l| = m_l,$$

where the l^{th} block $\Omega_l \subseteq \{N_{l-1}, \dots, N_l\}$ corresponds to the m_l samples required to recover the k_l nonzero coefficients at scale l . Here the parameters $0 = N_0 < N_1 < \dots < N_r = n$ are appropriately chosen and delineate frequency bands from which the m_l samples are taken.

In Section 3.3, we explain why this choice of A works, and in particular, how to choose the sampling blocks Ω_l . In order to do this, it is first necessary to recall the notion of incoherent bases.

3.2 Incoherent bases and compressed sensing

Besides random ensembles, a common approach in compressed sensing is to design sensing matrices using orthonormal systems that are incoherent with the particular choice of sparsifying basis Φ [11, 12, 24]. Let $\Psi \in \mathbb{C}^{n \times n}$ be an orthonormal basis of \mathbb{C}^n . The (mutual) coherence of Φ and Ψ is the quantity

$$\mu = \mu(\Psi^* \Phi) = \max_{i,j=1,\dots,n} |(\Psi^* \Phi)_{i,j}|^2.$$

We say Ψ and Φ are *incoherent* if $\mu(\Psi, \Phi) \leq a/n$ for some $a \geq 1$ independent of n . Given such a Ψ , one constructs the sensing matrix $A = P_\Omega \Psi$, where $\Omega \subseteq \{1, \dots, N\}$, $|\Omega| = m$ is chosen uniformly at random. A standard result gives that a k -sparse signal x in the basis Φ is recovered exactly with probability at least $1 - p$, provided

$$m \geq Cka \log(1 + p^{-1}) \log(n),$$

for some universal constant $C > 0$ [24]. As an example, consider the Fourier basis $\Psi = F$. This is incoherent with the canonical basis $\Phi = I$ with optimally-small constant $a = 1$. Fourier matrices subsampled uniformly at random are efficient sensing matrices for signals that are themselves sparse.

However, the Fourier matrix is not incoherent with a wavelet basis: $\mu(F, \Phi) = \mathcal{O}(1)$ as $n \rightarrow \infty$ for any orthonormal wavelet basis [2]. Nevertheless, Fourier samples taken within appropriate frequency bands (i.e. not uniformly at random) are *locally* incoherent with wavelets in the corresponding scales. This observation, which we demonstrate next, explains the success of the sensing matrix A constructed in the previous subsection for an appropriate choice of $\Omega_1, \dots, \Omega_r$.

3.3 Local incoherence and near block-diagonality of Fourier measurements with wavelets

For expository purposes, we consider the case of one-dimensional Haar wavelets. We note however that the arguments generalize to arbitrary compactly-supported orthonormal wavelets, and to the infinite-dimensional setting where the unknown image x is a function. See Section 4.3.

Let $j = 0, \dots, r-1$ (for convenience we now index from 0 to $r-1$, as opposed to 1 to r) be the scale and $p = 0, \dots, 2^j - 1$ the translation. The Haar basis consists of the functions $\{\psi\} \cup \{\phi_{j,p}\}$, where $\psi \in \mathbb{C}^n$ is the normalized scaling function and $\phi_{j,p}$ are the scaled and shifted versions of the mother wavelet $\phi \in \mathbb{C}^n$. It is a straightforward, albeit tedious, exercise to show that

$$|\mathcal{F}\phi_{j,p}(\omega)| = 2^{l/2-r+1} \frac{|\sin(\pi\omega/2^{l+1})|^2}{|\sin(\pi\omega/2^r)|} \lesssim 2^{l/2} \frac{|\sin(\pi\omega/2^{l+1})|^2}{|\omega|}, \quad |\omega| < 2^r,$$

where \mathcal{F} denotes the DFT [3]. This suggests that the Fourier transform $\mathcal{F}\phi_{l,p}(\omega)$ is large when $\omega \approx 2^l$, yet smaller when $\omega \approx 2^j$ with $j \neq l$. Hence we should separate frequency space into bands of size roughly 2^j .

Let $F \in \mathbb{C}^{n \times n}$ be the DFT matrix with rows indexed from $-n/2 + 1$ to $n/2$. Following an approach of [12], we now divide these rows into the following disjoint frequency bands

$$W_0 = \{0, 1\}, \quad W_j = \{-2^j + 1, \dots, -2^{j-1}\} \cup \{2^{j-1} + 1, \dots, 2^j\}, \quad l = 0, \dots, r-1.$$

With this to hand, we now define Ω_j to be a subset of W_j of size $|\Omega_j| = m_j$ chosen uniformly at random. Thus, the overall sensing matrix $A = P_\Omega F$ takes measurements of the signal x by randomly drawing m_j samples its Fourier transform within the frequency bands W_j .

Having specified Ω , let us note that the matrix $U = F\Phi$ naturally divides into blocks, with the rows corresponding to the frequency bands W_l and the columns corresponding to the wavelet scales. Write $U = \{U_{jl}\}_{j,l=0}^{r-1}$ where $U_{jl} \in \mathbb{C}^{2^j \times 2^l}$. For compressed sensing to succeed in this setting, we require two properties. First, the diagonal blocks U_{jj} should be incoherent, i.e. $\mu(U_{jj}) = \mathcal{O}(2^{-j})$. Second, the coherences $\mu(U_{jl})$ of the off-diagonal blocks U_{jl} should be appropriately small in comparison to $\mu(U_{jj})$. These two properties are demonstrated in the following lemma:

Lemma 1. *We have $\mu(U_{jj}) \lesssim 2^{-j}$ and, in general,*

$$\mu(U_{jl}) \lesssim \mu(U_{jj})2^{-|j-l|}, \quad j, l = 0, \dots, r-1$$

Hence U is approximately block diagonal, with exponential decay away from the diagonal blocks. Fourier measurements subsampled according to the above strategy are therefore ideally suited to recover structured sparse wavelet coefficients.²

The left panel of Figure 4 exhibits this decay by plotting the absolute values of the matrix U . In the right panel, we also show a similar result when the Fourier transform is replaced by the Hadamard transform. This is an important case, since the measurement matrix is binary. The middle panel of the figure shows the U matrix when Legendre polynomials are used as the sparsifying transform, as is sometimes the case for smooth signals. It demonstrates that diagonally-dominated coherence is not just a phenomenon associated with wavelets.

Having identified a measurement matrix to exploit structured sparsity, let us demonstrate its effectiveness. In Figure 5 we compare these measurements with the case of random Bernoulli measurements (this choice was made over random Gaussian measurements because of storage issues). As is evident, at all resolutions we see a significant advantage, since the former strategy exploits the structured sparsity. Note that for both approaches, the reconstruction quality is *resolution dependent*: the error decreases as the resolution increases, due to the increasing sparsity of wavelet coefficients at higher resolutions. However, because the Fourier/wavelets matrix U

² For brevity, we do not give the proof of this lemma or the later recovery results for Haar wavelets, Theorem 2. Details of the proof can be found in the short note [3].




Fig. 4: The absolute values of the matrix $U = F\Phi$, where F is the discrete Fourier transform (left and middle) or the Hadamard transform (right) and the sparsifying transform Φ corresponds to Haar wavelets (left and right) or Legendre polynomials (middle). Light colours correspond to large values and dark colours to small values.




Fig. 5: Recovery from 12.5% measurements using (1) with Daubechies-4 wavelets as the sparsifying transform. Top row: Random Bernoulli sensing matrix. Bottom row: Fourier sensing matrix with multilevel subsampling (see Definition 3). All images are 256×256 crops of the original full resolution versions in order to aid the visual comparison.

is *asymptotically incoherent* (see also Section 4.1), it exploits the inherent asymptotic sparsity structure (6) of the wavelet coefficients as the resolution increases, and thus gives successively greater improvements over random Bernoulli measurements.

Remark 4. Besides improved reconstructions, an important feature of this approach is storage and computational time. Since F and Φ have fast, $\mathcal{O}(n \log n)$, transforms, the matrix $A = P_\Omega F \Phi$ does not need to be stored, and the reconstruction (1) can be performed efficiently with standard ℓ_1 solvers (we use SPGL1 [49, 50] throughout).

Recall that in type I problems such as MRI, we are constrained by the physics of the device to take Fourier measurements. A rather strange conclusion of Figure 5 is the following: compressed sensing actually works better for MRI with the intrinsic measurements, than if one were able to take optimal (in the sense of the standard sparsity-based theory) random (sub)Gaussian measurements. This has practical consequences. In MRI there is actually a little flexibility to design measurements, based on specifying appropriate pulses. By doing this, a number of approaches [28, 36, 43, 45, 54] have been proposed to make MRI measurements closer to uniformly incoherent with wavelets (i.e. similar to random Gaussians). On the other hand, Figure 5 suggests that one can obtain great results in practice by appropriately subsampling the unmodified Fourier operator.

4 A general framework for compressed sensing based on structured sparsity

Having argued for the particular case of Fourier samples with Haar wavelets, we now describe a general mathematical framework for structured sparsity. This is based on work in [2].

4.1 Concepts

We shall work in both the finite- and infinite-dimensional settings, where $U \in \mathbb{C}^{n \times n}$ or $U \in \mathcal{B}(\ell_2(\mathbb{N}))$ respectively. We assume throughout that U is an isometry. This occurs for example when $U = \Psi^* \Phi$ for an orthonormal basis Ψ and an orthonormal system Φ , as is the case for the example studied in the previous section: namely, Fourier sampling and a wavelet sparsifying basis. However, framework we present now is valid for an arbitrary isometry U , not just this particular example. We discuss this case further in Section 4.3.

We first require the following definitions. In the previous section it was suggested to divide both the sampling strategy and the sparse vector of coefficients into disjoint blocks. We now formalize these notions:

Definition 1 (Sparsity in levels). Let c be an element of either \mathbb{C}^N or $\ell_2(\mathbb{N})$. For $r \in \mathbb{N}$ let $\mathbf{M} = (M_1, \dots, M_r) \in \mathbb{N}^r$ with $1 \leq M_1 < \dots < M_r$ and $\mathbf{k} = (k_1, \dots, k_r) \in \mathbb{N}^r$, with $k_j \leq M_j - M_{j-1}$, $k = 1, \dots, r$, where $M_0 = 0$. We say that c is (\mathbf{k}, \mathbf{M}) -sparse if, for each $j = 1, \dots, r$, the set

$$\Delta_j := \text{supp}(c) \cap \{M_{j-1} + 1, \dots, M_j\},$$

satisfies $|\Delta_j| \leq k_j$. We denote the set of (\mathbf{k}, \mathbf{M}) -sparse vectors by $\Sigma_{\mathbf{k}, \mathbf{M}}$.

This definition allows for differing amounts of sparsity of the vector c in different levels. Note that the levels \mathbf{M} do not necessarily correspond to wavelet scales – for now, we consider a general setting. We also need a notion of best approximation:

Definition 2 ((\mathbf{k}, \mathbf{M})-term approximation). Let c be an element of either \mathbb{C}^N or $\ell_2(\mathbb{N})$. We say that c is (\mathbf{k}, \mathbf{M}) -compressible if $\sigma_{\mathbf{k}, \mathbf{M}}(c)$ is small, where

$$\sigma_{\mathbf{k}, \mathbf{M}}(c)_1 = \inf_{z \in \Sigma_{\mathbf{k}, \mathbf{M}}} \|c - z\|_1. \quad (7)$$

As we have already seen for wavelet coefficients, it is often the case that $k_j/(M_j - M_{j-1}) \rightarrow 0$ as $j \rightarrow \infty$. In this case, we say that c is *asymptotically sparse in levels*. However, we stress that this framework does not explicitly require such decay.

We now consider the level-based sampling strategy:

Definition 3 (Multilevel random sampling). Let $r \in \mathbb{N}$, $\mathbf{N} = (N_1, \dots, N_r) \in \mathbb{N}^r$ with $1 \leq N_1 < \dots < N_r$, $\mathbf{m} = (m_1, \dots, m_r) \in \mathbb{N}^r$, with $m_j \leq N_j - N_{j-1}$, $j = 1, \dots, r$, and suppose that

$$\Omega_j \subseteq \{N_{j-1} + 1, \dots, N_j\}, \quad |\Omega_j| = m_j, \quad j = 1, \dots, r,$$

are chosen uniformly at random, where $N_0 = 0$. We refer to the set

$$\Omega = \Omega_{\mathbf{N}, \mathbf{m}} = \Omega_1 \cup \dots \cup \Omega_r.$$

as an (\mathbf{N}, \mathbf{m}) -multilevel sampling scheme.

As discussed in Section 3.2, the (infinite) Fourier/wavelets matrix $U = F\Phi$ is globally coherent. However, as shown in Lemma 1, the coherence of its $(j, l)^{\text{th}}$ block is much smaller. We therefore require a notion of local coherence:

Definition 4 (Local coherence). Let U be an isometry of either \mathbb{C}^N or $\ell_2(\mathbb{N})$. If $\mathbf{N} = (N_1, \dots, N_r) \in \mathbb{N}^r$ and $\mathbf{M} = (M_1, \dots, M_r) \in \mathbb{N}^r$ with $1 \leq N_1 < \dots < N_r$ and $1 \leq M_1 < \dots < M_r$ the $(j, l)^{\text{th}}$ local coherence of U with respect to \mathbf{N} and \mathbf{M} is given by

$$\mu_{\mathbf{N}, \mathbf{M}}(j, l) = \sqrt{\mu(P_{N_j}^{N_{j-1}} U P_{M_l}^{M_{l-1}}) \mu(P_{N_j}^{N_{j-1}} U)}, \quad k, l = 1, \dots, r,$$

where $N_0 = M_0 = 0$ and P_b^a denotes the projection matrix corresponding to indices $\{a+1, \dots, b\}$. In the case where U is an operator on $\ell_2(\mathbb{N})$, we also define

$$\mu_{\mathbf{N}, \mathbf{M}}(j, \infty) = \sqrt{\mu(P_{N_j}^{N_{j-1}} U P_{M_{r-1}}^\perp) \mu(P_{N_j}^{N_{j-1}} U)}, \quad j = 1, \dots, r,$$

Note that the local coherence $\mu_{\mathbf{N}, \mathbf{M}}(j, l)$ is not just the coherence $\mu(P_{N_j}^{N_{j-1}} U P_{M_l}^{M_{l-1}})$ in the $(j, l)^{\text{th}}$ block. For technical reasons, one requires the product of this and the coherence $\mu(P_{N_j}^{N_{j-1}} U)$ in the whole j^{th} row block.

Remark 5. In [33], the authors define a local coherence of a matrix U in the j^{th} row (as opposed to row block) to be the maximum of its entries in that row. Using this, they prove recovery guarantees for the Fourier/Haar wavelets problem based on the RIP and the global sparsity k . Unfortunately, these results do not explain the importance of structured sparsity as shown by the flip test. Conversely, our notion of local coherence also takes into account the sparsity levels. As we will see in Section 4.2, this allows one to establish recovery guarantees that are consistent with the flip test and properly explain the role of structure in the reconstruction.

Recall that in practice (see Figure 4), the local coherence often decays along the diagonal blocks and in the off-diagonal blocks. Loosely speaking, we say that the matrix U is *asymptotically incoherent* in this case.

In Section 3.3 we argued that the Fourier/wavelets matrix was nearly block-diagonal. In our theorems, we need to account for the off-diagonal terms. To do this in the general setup, we require a notion of a relative sparsity:

Definition 5 (Relative sparsity). Let U be an isometry of either \mathbb{C}^N or $\ell_2(\mathbb{N})$. For $\mathbf{N} = (N_1, \dots, N_r) \in \mathbb{N}^r$, $\mathbf{M} = (M_1, \dots, M_r) \in \mathbb{N}^r$ with $1 \leq N_1 < \dots < N_r$ and $1 \leq M_1 < \dots < M_r$, $\mathbf{k} = (k_1, \dots, k_r) \in \mathbb{N}^r$ and $j = 1, \dots, r$, the j^{th} relative sparsity is given by

$$K_j = K_j(\mathbf{N}, \mathbf{M}, \mathbf{k}) = \max_{z \in \Sigma_{\mathbf{k}, \mathbf{M}}, \|\eta\|_\infty \leq 1} \|P_{N_j}^{N_{j-1}} U z\|^2,$$

where $N_0 = M_0 = 0$.

The relative sparsities K_j take into account *interferences* between different sparsity level caused by the non-block diagonality of U .

4.2 Main theorem

Given the matrix/operator U and a multilevel sampling scheme Ω , we now consider the solution of the convex optimization problem

$$\min_{z \in \mathbb{C}^n} \|z\|_1 \quad \text{s.t.} \quad \|P_\Omega y - P_\Omega U z\|_2 \leq \eta, \quad (8)$$

where $y = Uc + e$, $\|e\|_2 \leq \eta$. Note that if $U = \Psi^* \Phi$, $x = \Phi c$ is the signal we wish to recover and \hat{c} is a minimizer of (8) then this gives the approximation $\hat{x} = \Phi \hat{c}$ to x .

Theorem 1. Let $U \in \mathbb{C}^{N \times N}$ be an isometry and $c \in \mathbb{C}^N$. Suppose that $\Omega = \Omega_{\mathbf{N}, \mathbf{m}}$ is a multilevel sampling scheme, where $\mathbf{N} = (N_1, \dots, N_r) \in \mathbb{N}^r$, $N_r = n$, and $\mathbf{m} = (m_1, \dots, m_r) \in \mathbb{N}^r$. Let $\varepsilon \in (0, e^{-1}]$ and suppose that (\mathbf{k}, \mathbf{M}) , where $\mathbf{M} = (M_1, \dots, M_r) \in \mathbb{N}^r$, $M_r = n$, and $\mathbf{k} = (k_1, \dots, k_r) \in \mathbb{N}^r$, are any pair such that the following holds:

(i) We have

$$1 \gtrsim \frac{N_j - N_{j-1}}{m_j} \log(\varepsilon^{-1}) \left(\sum_{l=1}^r \mu_{\mathbf{N}, \mathbf{M}}(j, l) k_l \right) \log(n), \quad j = 1, \dots, r. \quad (9)$$

(ii) We have $m_j \gtrsim \hat{m}_j \log(\varepsilon^{-1}) \log(n)$, where \hat{m}_j is such that

$$1 \gtrsim \sum_{j=1}^r \left(\frac{N_j - N_{j-1}}{\hat{m}_j} - 1 \right) \mu_{\mathbf{N}, \mathbf{M}}(j, l) \tilde{k}_j, \quad l = 1, \dots, r, \quad (10)$$

for all $\tilde{k}_1, \dots, \tilde{k}_r \in (0, \infty)$ satisfying

$$\tilde{k}_1 + \dots + \tilde{k}_r \leq k_1 + \dots + k_r, \quad \tilde{k}_j \leq K_j(\mathbf{N}, \mathbf{M}, \mathbf{k}).$$

Suppose that $\hat{c} \in \mathbb{C}^N$ is a minimizer of (8). Then, with probability exceeding $1 - k\varepsilon$, where $k = k_1 + \dots + k_r$, we have that

$$\|c - \hat{c}\|_2 \leq C \left(\eta \sqrt{D} (1 + E \sqrt{k}) + \sigma_{\mathbf{k}, \mathbf{M}}(c)_1 \right), \quad (11)$$

for some constant C , where $\sigma_{\mathbf{k}, \mathbf{M}}(c)_1$ is as in (7), $D = 1 + \frac{\sqrt{\log_2(6\varepsilon^{-1})}}{\log_2(4En\sqrt{s})}$ and $E = \max_{j=1, \dots, r} \{(N_j - N_{j-1})/m_j\}$. If $m_j = N_j - N_{j-1}$, $j = 1, \dots, r$, then this holds with probability 1.

A similar theorem can be stated and proved in the infinite-dimensional setting [2]. For brevity, we shall not do this.

The key feature of Theorem 1 is that the bounds (9) and (10) involve only local quantities: namely, local sparsities k_j , local coherences $\mu(j, l)$, local measurements m_j and relative sparsities K_j . Note that the estimate (11) directly generalizes a standard compressed sensing estimate for sampling with incoherent bases [1, 11] to the case of multiple levels. Having said this, it is not immediately obvious how to understand these bounds in terms of how many measurements m_j are actually required in the j^{th} level. However, it can be shown that these bounds are in fact sharp for a large class of matrices U [2]. Thus, little improvement is possible. Moreover, in the important case Fourier/wavelets, one can analyze the local coherences $\mu(j, l)$ and relative sparsities K_j to get such explicit estimates. We consider this next.

4.3 The case of Fourier sampling with wavelets

Let us consider the example of Section 3.3, where the matrix U arises from the Fourier/Haar wavelet pair, the sampling levels are correspond to the aforementioned frequency bands W_j and the sparsity levels are the Haar wavelet scales.

Theorem 2.³ Let U and Ω be as in Section 3.3 (recall that we index over $j, l = 0, \dots, r-1$) and suppose that $x \in \mathbb{C}^n$. Let $\varepsilon \in (0, e^{-1}]$ and suppose that

$$m_j \gtrsim \left(k_j + \sum_{\substack{l=0 \\ l \neq j}}^{r-1} 2^{-\frac{|j-l|}{2}} k_l \right) \log(\varepsilon^{-1}) \log^2(n), \quad j = 0, \dots, r-1. \quad (12)$$

Then, with probability exceeding $1 - k\varepsilon$, where $k = k_0 + \dots + k_{r-1}$, any minimizer \hat{x} of (1) satisfies

$$\|x - \hat{x}\|_2 \leq C \left(\eta \sqrt{D} (1 + E \sqrt{k}) + \sigma_{\mathbf{k}, \mathbf{M}}(\Phi^* x)_1 \right),$$

where $\sigma_{\mathbf{k}, \mathbf{M}}(\Phi^* x)_1$ is as in (7), $D = 1 + \frac{\sqrt{\log_2(6\varepsilon^{-1})}}{\log_2(4En\sqrt{k})}$ and $E = \max_{j=0, \dots, r-1} \{(N_j - N_{j-1})/m_j\}$. If $m_j = |W_j|$, $j = 0, \dots, r-1$, then this holds with probability 1.

The key part of this theorem is (12). Recall that if U were exactly block diagonal, then $m_j \gtrsim k_j$ would suffice (up to log factors). The estimate (12) asserts that we require only slightly more samples, and this is due to interferences from the other sparsity levels. However, as $|j-l|$ increases, the effect of these levels decreases exponentially. Thus, the number of measurements m_j required in the j^{th} frequency band is determined predominantly by the sparsities in the scales $l \approx j$. Note that $k_l \approx \mathcal{O}(k_j)$ when $l \approx j$ for typical signals and images, so the estimate (12) is typically on the order of k_j in practice.

The estimate (12) both agrees with the conclusion of the flip test in Figure 1 and explains the results seen. Flipping the wavelet coefficients changes the local sparsities k_1, \dots, k_r . Therefore to recover the flipped image to the same accuracy as the unflipped image, (12) asserts that one must change the local numbers of measurements m_j . But in Figure 1 the same sampling pattern was used in both cases, thereby leading to the worse reconstruction in the flipped case. Note that (12) also highlights why the optimal sampling pattern must depend on the image, and specifically, the local sparsities. In particular, there can be no optimal sampling strategy for all images.

Note that Theorem 2 is a simplified version, presented here for the purposes of elucidation, of a more general result found in [2] which applies to all compactly-supported orthonormal wavelets in the infinite-dimensional setting.

5 Structured sampling and structured recovery

Structured sparsity within the context of compressed sensing has been considered in numerous previous works. See [5, 8, 13, 20, 22, 23, 29, 30, 41, 48, 51, 46] and

³ For a proof, we refer to [3].




Fig. 6: Recovery from 12.5% measurements at 256×256 . Comparison between random sampling with structured recovery and structured sampling with ℓ^1 -minimization recovery.

references therein. For the problem of reconstructing wavelet coefficients, most efforts have focused on their inherent tree structure (see Remark 3). Three well-known algorithmic approaches for doing this are model-based compressed sensing [5], TurboAMP [46], and Bayesian compressed sensing [29, 30]. All methods use Gaussian or Bernoulli random measurements, and seek to leverage the wavelet tree structure – the former deterministically, the latter two in a probabilistic manner – by appropriately-designed recovery algorithms (based on modifications of existing iterative algorithms for compressed sensing). In other words, structure is incorporated solely in the recovery algorithm, and not in the measurements themselves.




Fig. 7: The flip test in levels for the image considered in Figure 1.

In Figure 6 we compare these algorithms with the previously-described method of multilevel Fourier sampling (similar results are also witnessed with the Hadamard matrix). Note that the latter, unlike other three methods, exploits structure by taking appropriate measurements, and uses an unmodified compressed sensing algorithm (ℓ^1 minimization). As is evident, this approach is able to better exploit the sparsity structure, leading to a significantly improved reconstruction. This experiment is representative of a large set tested. In all cases, we find that exploiting structure by sampling with asymptotically incoherent Fourier/Hadamard bases outperforms such approaches that seek to leverage structure in the recovery algorithm.

6 The Restricted Isometry Property in levels

The flip test demonstrates that the subsampled Fourier/wavelets matrix $P_\Omega U$ does not satisfy a meaningful RIP. However, due to the sparsity structure of the signal, we are still able to reconstruct, as was confirmed in Theorem 1. The RIP is therefore too crude a tool to explain the recoverability properties of structured sensing matrices. Having said that, the RIP is a useful for deriving uniform, as opposed to nonuniform, recovery results, and for analyzing other compressed sensing algorithms besides convex optimization, such as greedy methods [24]. This raises the question of whether there are alternatives which are satisfied by such matrices. One possibility which we now discuss is the RIP in levels. Throughout this section we shall work in the finite-dimensional setting. For proofs of the theorems, see [6].

Definition 6. Given an r -level sparsity pattern (\mathbf{k}, \mathbf{M}) , where $M_r = n$, we say that the matrix $A \in \mathbb{C}^{m \times n}$ satisfies the *RIP in levels* (RIP_L) with RIP_L constant $\delta_{\mathbf{k}} \geq 0$ if for all x in $\Sigma_{\mathbf{k}, \mathbf{M}}$ we have

$$(1 - \delta_{\mathbf{k}})\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta_{\mathbf{k}})\|x\|_2^2.$$

The motivation for this definition is the following. Suppose that we repeat the flip test from Figure 1 except that instead of completely flipping the coefficients we only flip them within levels corresponding to the wavelet scales. We will refer to this as the *flip test in levels*. Note the difference between Figure 1 and Figure 7, where latter presents the flip test in levels: clearly, flipping within scales does not alter the reconstruction quality. In light of this experiment, we propose the above RIP in levels definition so as to respect the level structure.

We now consider the recovery properties of matrices satisfying the RIP in levels. For this, we define the *ratio constant* $\lambda_{\mathbf{k}, \mathbf{M}}$ of a sparsity pattern (\mathbf{k}, \mathbf{M}) to be $\lambda_{\mathbf{k}, \mathbf{M}} := \max_{j,l} k_j/k_l$. We assume throughout that $k_j \geq 1, \forall j$, so that $\eta_{\mathbf{k}, \mathbf{M}} < \infty$, and also that $M_r = n$. We now have the following:

Theorem 3. *Let (\mathbf{k}, \mathbf{M}) be a sparsity pattern with r levels and ratio constant $\lambda = \lambda_{\mathbf{k}, \mathbf{M}}$. Suppose that the matrix A has RIP_L constant $\delta_{2\mathbf{k}}$ satisfying*

$$\delta_{2\mathbf{k}} < \frac{1}{\sqrt{r(\sqrt{\lambda} + 1/4)^2 + 1}}. \quad (13)$$

Let $x \in \mathbb{C}^n$, $y \in \mathbb{C}^m$ be such that $\|Ux - y\|_2 \leq \eta$, and let \hat{x} be a minimizer of

$$\min_{z \in \mathbb{C}^n} \|z\|_1 \quad \text{s.t.} \quad \|y - Az\|_2 \leq \eta.$$

Then

$$\|x - \hat{x}\|_1 \leq C\sigma_{\mathbf{k}, \mathbf{M}}(x)_1 + D\sqrt{k}\eta, \quad (14)$$

where $k = k_1 + \dots + k_r$ and the constants C and D depend only on $\delta_{2\mathbf{k}}$.

This theorem is a generalization of a known result in standard (i.e. one-level) compressed sensing. Note that (13) reduces to the well-known estimate $\delta_{2k} \leq 4/\sqrt{41}$ [24] when $r = 1$. On the other hand, in the multiple level case the reader may be concerned that the bound ceases to be useful, since the right-hand side of (13) deteriorates with both the number of levels r and the sparsity ratio λ . As we show in the following two theorems, the dependence on r and λ in (13) is sharp:

Theorem 4. *Fix $a \in \mathbb{N}$. There exists a matrix A with two levels and a sparsity pattern (\mathbf{k}, \mathbf{M}) such that the RIP_L constant $\delta_{a\mathbf{k}}$ and ratio constant $\lambda = \lambda_{\mathbf{k}, \mathbf{M}}$ satisfy*

$$\delta_{a\mathbf{k}} \leq 1/|f(\lambda)|, \quad (15)$$

where $f(\lambda) = o(\sqrt{\lambda})$, but there is an $x \in \Sigma_{\mathbf{k}, \mathbf{M}}$ such that x is not the minimizer of

$$\min_{z \in \mathbb{C}^n} \|z\|_1 \quad \text{s.t.} \quad Az = Ax.$$

Roughly speaking, Theorem 4 says that if we fix the number of levels and try to replace (13) with a condition of the form

$$\delta_{2\mathbf{k}} < \frac{1}{C\sqrt{r}}\lambda^{-\frac{\alpha}{2}}$$

for some constant C and some $\alpha < 1$ then the conclusion of Theorem 3 ceases to hold. In particular, the requirement on $\delta_{2\mathbf{k}}$ cannot be independent of λ . The parameter a in the statement of Theorem 4 also means that we cannot simply fix the issue by changing $\delta_{2\mathbf{k}}$ to $\delta_{3\mathbf{k}}$, or any further multiple of \mathbf{k} .

Similarly, we also have a theorem that shows that the dependence on the number of levels r cannot be ignored.

Theorem 5. *Fix $a \in \mathbb{N}$. There exists a matrix A and an r -level sparsity pattern (\mathbf{k}, \mathbf{M}) with ratio constant $\lambda_{\mathbf{k}, \mathbf{M}} = 1$ such that the RIP_L constant $\delta_{a\mathbf{k}}$ satisfies*

$$\delta_{a\mathbf{k}} \leq 1/|f(r)|,$$

where $f(r) = o(\sqrt{r})$, but there is an $x \in \Sigma_{\mathbf{k}, \mathbf{M}}$ such that x is not the minimizer of

$$\min_{z \in \mathbb{C}^n} \|z\|_1 \quad \text{s.t.} \quad Az = Ax.$$

These two theorems suggests that, at the level of generality of the RIP_L , one must accept a bound that deteriorates with the number of levels and ratio constant. This begs the question: what is the effect of such deterioration? To understand this, consider the case studied earlier, where the r levels correspond to wavelet scales. For typical images, the ratio constant λ grows only very mildly with n , where $n = 2^r$ is the dimension. Conversely, the number of levels is equal to $\log_2(n)$. This suggests that estimates for the Fourier/wavelet matrix that ensure an RIP in levels (thus guaranteeing uniform recovery) will involve at least several additional factors of $\log(n)$ beyond what is sufficient for nonuniform recovery (see Theorem 2). Proving such estimates for the Fourier/wavelets matrix is work in progress.

Acknowledgements The authors thank Andy Ellison from Boston University Medical School for kindly providing the MRI fruit image, and General Electric Healthcare for kindly providing the brain MRI image. BA acknowledges support from the NSF DMS grant 1318894. ACH acknowledges support from a Royal Society University Research Fellowship. ACH and BR acknowledge the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L003457/1.

References

1. B. Adcock and A. C. Hansen. Generalized sampling and infinite-dimensional compressed sensing. *Technical report NA2011/02, DAMTP, University of Cambridge*, 2011.
2. B. Adcock, A. C. Hansen, C. Poon, and B. Roman. Breaking the coherence barrier: A new theory for compressed sensing. *arXiv:1302.0561*, 2014.
3. B. Adcock, A. C. Hansen, and B. Roman. A note on compressed sensing of structured sparse wavelet coefficients from subsampled Fourier measurements. *arXiv*, 2014.
4. B. Adcock, A. C. Hansen, B. Roman, and G. Teschke. Generalized sampling: stable reconstructions, inverse problems and compressed sensing over the continuum. *Advances in Imaging and Electron Physics (to appear)*, Elsevier, 2013.
5. R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hedge. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56(4):1982–2001, 2010.

6. A. Bastounis and A. C. Hansen. On the Restricted Isometry Property in Levels. *Preprint*, 2014.
7. P. Binev, W. Dahmen, R. A. DeVore, P. Lamby, D. Savu, and R. Sharpley. Compressed sensing and electron microscopy. In T. Vogt, W. Dahmen, and P. Binev, editors, *Modeling Nanoscale Imaging in Electron Microscopy*, Nanostructure Science and Technology, pages 73–126. Springer US, 2012.
8. A. Bourrier, M. E. Davies, T. Peleg, P. Pérez, and R. Gribonval. Fundamental performance limits for ideal decoders in high-dimensional linear inverse problems. *arXiv:1311.6239*, 2013.
9. E. Candès and D. L. Donoho. Recovering edges in ill-posed inverse problems: optimality of curvelet frames. *Ann. Statist.*, 30(3):784–842, 2002.
10. E. J. Candès and D. L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. *Comm. Pure Appl. Math.*, 57(2):219–266, 2004.
11. E. J. Candès and Y. Plan. A probabilistic and RIPless theory of compressed sensing. *IEEE Trans. Inform. Theory*, 57(11):7235–7254, 2011.
12. E. J. Candès and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969–985, 2007.
13. W. R. Carson, M. Chen, M. R. D. Rodrigues, R. Calderbank, and L. Carin. Communications-inspired projection design with application to compressive sensing. *SIAM J. Imaging Sci.*, 5(4):1185–1212, 2012.
14. R. Chartrand, Y. Sidky, and X. Pan. Nonconvex compressive sensing for X-ray CT: an algorithm comparison. In *Asilomar Conference on Signals, Systems, and Computers*, 2013.
15. M. S. Crouse, R. D. Nowak, and R. G. Baraniuk. Wavelet-based statistical signal processing using hidden Markov models. *IEEE Trans. Signal Process.*, 46:886–902, 1998.
16. S. Dahlke, G. Kutyniok, G. Steidl, and G. Teschke. Shearlet coorbit spaces and associated Banach frames. *Appl. Comput. Harmon. Anal.*, 27(2):195–214, 2009.
17. Stephan Dahlke, Gitta Kutyniok, Peter Maass, Chen Sagiv, Hans-Georg Stark, and Gerd Teschke. The uncertainty principle associated with the continuous shearlet transform. *Int. J. Wavelets Multiresolut. Inf. Process.*, 6(2):157–181, 2008.
18. R. A. DeVore. Nonlinear approximation. *Acta Numer.*, 7:51–150, 1998.
19. M. N. Do and M. Vetterli. The contourlet transform: An efficient directional multiresolution image representation. *IEEE Trans. Image Proc.*, 14(12):2091–2106, 2005.
20. D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.
21. M. F. Duarte, M. A. Davenport, D. Takhar, J. Laska, K. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.*, 25(2):83–91, 2008.
22. M. F. Duarte and Y. C. Eldar. Structured compressed sensing: from theory to applications. *IEEE Trans. Signal Process.*, 59(9):4053–4085, 2011.
23. Y. C. Eldar. Robust recovery of signals from a structured union of subspaces. *IEEE Trans. Inform. Theory*, 55(11):5302–5316, 2009.
24. S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhauser, 2013.
25. L. Gan, T. T. Do, and T. D. Tran. Fast compressive imaging using scrambled hadamard ensemble. *Proc. EUSIPCO*, 2008.
26. M. Guerquin-Kern, M. Häberlin, K. P. Pruessmann, and M. Unser. A fast wavelet-based reconstruction method for Magnetic Resonance Imaging. *IEEE Trans. Med. Imaging*, 30(9):1649–1660, 2011.
27. M. Guerquin-Kern, L. Lejeune, K. P. Pruessmann, and M. Unser. Realistic analytical phantoms for parallel Magnetic Resonance Imaging. *IEEE Trans. Med. Imaging*, 31(3):626–636, 2012.
28. J. Haldar, D. Hernando, and Z. Liang. Compressed-sensing MRI with random encoding. *IEEE Trans. Med. Imaging*, 30(4):893–903, 2011.
29. L. He and L. Carin. Exploiting structure in wavelet-based Bayesian compressive sensing. *IEEE Trans. Signal Process.*, 57(9):3488–3497, 2009.
30. L. He, H. Chen, and L. Carin. Tree-structured compressive sensing with variational Bayesian analysis. *IEEE Signal Process. Letters*, 17(3):233–236, 2010.
31. M. A. Herman and T. Strohmer. High-resolution radar via compressed sensing. *IEEE Trans. Signal Process.*, 57(6):2275–2284, 2009.

32. G. Huang, H. Jiang, K. Matthews, and P. Wilford. Lensless imaging by compressive sensing. *arXiv:1305.7181*, 2013.
33. F. Krahmer and R. Ward. Stable and robust recovery from variable density frequency samples. *IEEE Trans. Image Proc. (to appear)*, 2014.
34. Gitta Kutyniok, Jakob Lemvig, and Wang-Q Lim. Compactly supported shearlets. In Marian Neamtu and Larry Schumaker, editors, *Approximation Theory XIII: San Antonio 2010*, volume 13 of *Springer Proceedings in Mathematics*, pages 163–186. Springer New York, 2012.
35. Rowan Leary, Zineb Saghi, Paul A. Midgley, and Daniel J. Holland. Compressed sensing electron tomography. *Ultramicroscopy*, 131(0):70–91, 2013.
36. D. Liang, G. Xu, H. Wang, K. F. King, D. Xu, and L. Ying. Toeplitz random encoding MR imaging using compressed sensing. In *Proc. IEEE. Int. Symp. Biomed. Imag.*, pages 270–273, June 2009.
37. T. Lin and F. J. Herrman. Compressed wavefield extrapolation. *Geophysics*, 72(5):SM77–SM93, 2007.
38. M. Lustig, D. L. Donoho, and J. M. Pauly. Sparse MRI: the application of compressed sensing for rapid MRI imaging. *Magn. Reson. Imaging*, 58(6):1182–1195, 2007.
39. M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed Sensing MRI. *IEEE Signal Process. Mag.*, 25(2):72–82, March 2008.
40. S. G. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way*. Academic Press, 3 edition, 2009.
41. M. Mishali, Y. C. Eldar, and A. J. Elron. Xampling: Signal acquisition and processing in union of subspaces. *IEEE Trans. Signal Process.*, 59(10):4719–4734, 2011.
42. D. D.-Y. Po and M. N. Do. Directional multiscale modeling of images using the contourlet transform. *IEEE Trans. Image Proc.*, 15(6):1610–1620, June 2006.
43. G. Puy, J. P. Marques, R. Gruetter, J. Thiran, D. Van De Ville, P. Vanderghenst, and Y. Wiaux. Spread spectrum Magnetic Resonance Imaging. *IEEE Trans. Med. Imaging*, 31(3):586–598, 2012.
44. J. Romberg. Imaging via compressive sampling. *IEEE Signal Process. Mag.*, 25(2):14–20, 2008.
45. F. Seibert, Y. M. Xou, and L. Ying. Compressed sensing MRI with random B1 field. *Proc. Int. Soc. Magn. Reson. Med*, 16:3151, 2008.
46. S. Som and P. Schniter. Compressive imaging using approximate message passing and a markov-tree prior. *IEEE Trans. Signal Process.*, 60(7):3439–3448, 2012.
47. V. Studer, J. Bobin, M. Chahid, H. Moussavi, E. Candès, and M. Dahan. Compressive fluorescence microscopy for biological and hyperspectral imaging. *Submitted*, 2011.
48. Y. Tsaig and D. L. Donoho. Extensions of compressed sensing. *Signal Process.*, 86(3):549–571, 2006.
49. E. van den Berg and M. P. Friedlander. SPGL1: A solver for large-scale sparse reconstruction, June 2007. <http://www.cs.ubc.ca/labs/scl/spgl1>.
50. 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.
51. L. Wang, D. Carlson, M. R. D. Rodrigues, D. Wilcox, R. Calderbank, and L. Carin. Designed measurements for vector count data. In *Advances in Neural Information Processing Systems*, pages 1142–1150, 2013.
52. Z. Wang and G. R. Arce. Variable density compressed image sampling. *IEEE Trans. Image Proc.*, 19(1):264–270, 2010.
53. Y. Wiaux, L. Jacques, G. Puy, A. M. M. Scaife, and P. Vanderghenst. Compressed sensing imaging techniques for radio interferometry. *Mon. Not. R. Astron. Soc.*, 395(3):1733–1742, 2009.
54. E. C. Wong. Efficient randomly encoded data acquisition for compressed sensing. In *Proc. Int. Soc. Magn. Reson. Med.*, page 4893, 2010.