

Utilizing the structure of a redundant dictionary comprised of wavelets and curvelets with compressed sensing

Nicholas Dwork^{a,*} and Peder E. Z. Larson^b

^aUniversity of Colorado Anschutz, Departments of Biomedical Informatics and Radiology, Aurora, Colorado, United States

^bUniversity of California, Department of Radiology and Biomedical Imaging, San Francisco, California, United States

Abstract The discrete curvelet transform decomposes an image into a set of fundamental components that are distinguished by direction and size and a low-frequency representation. The curvelet representation of a natural image is approximately sparse; thus, it is useful for compressed sensing. However, with natural images, the low-frequency portion is seldom sparse. This manuscript presents a method to modify the redundant sparsifying transformation comprised of the wavelet and curvelet transforms to take advantage of this fact for compressed sensing image reconstruction. Instead of relying on sparsity for this low-frequency estimate, the Nyquist–Shannon sampling theorem specifies a rectangular region centered on the 0 frequency to be collected, which is used to generate a blurry estimate. A basis pursuit denoising problem is solved to determine the details with a modified sparsifying transformation. Improvements in quality are shown on magnetic resonance and optical images. © 2022 SPIE and IS&T [DOI: [10.1117/1.JEI.31.6.063043](https://doi.org/10.1117/1.JEI.31.6.063043)]

Keywords: compressed sensing; curvelet; wavelet; imaging; magnetic resonance imaging.

Paper 220148G received Feb. 1, 2022; accepted for publication Nov. 16, 2022; published online Dec. 6, 2022.

1 Introduction

Fourier sensing applications include magnetic resonance imaging (MRI), optical coherence tomography, computed tomography, and radio astronomy. In each of these applications, the data collected are values of the Fourier transform of the image of interest. To reconstruct the image, one must perform an approximation to the inverse Fourier transform. It is necessarily an approximation because the samples are a finite set of isolated points in the Fourier domain. If the samples are collected on a uniform grid, then the inverse discrete Fourier transform serves as a fast and high-quality method for image reconstruction. However, if the samples are not on a uniform grid or if the grid is not fully sampled (meaning that the samples are not sufficiently dense to satisfy the Nyquist–Shannon sampling theorem), then other inversion algorithms have shown improved results.

Compressed sensing is one such reconstruction method (compressed sensing is also known as compressive sampling); it permits accurate reconstructions of images with fewer samples than the number required to satisfy the Nyquist–Shannon theorem.¹ Compressed sensing can achieve this feat because it incorporates additional knowledge into the reconstruction. It relies on the *a priori* knowledge that the image can be reconstructed well with a linear transformation applied to a sparse vector. Common choices include the inverse wavelet transform,^{2,3} the inverse curvelet transform,⁴ and learned dictionaries.⁵ It is often beneficial to have fast algorithms that implement the sparsifying transformation and fast algorithms exist for the wavelet^{6,7} and curvelet transforms.⁸ Therefore, we focus on these transforms for this manuscript. Image reconstruction is formulated as a basis pursuit denoising (BPD) convex optimization problem.⁹ Compressed

*Address all correspondence to Nicholas Dwork, nicholas.dwork@cuanschutz.edu

sensing has benefited many applications including medical imaging,¹⁰ holography,¹¹ photography,¹² and communications.^{13,14}

Though the wavelet transforms of natural images are often sparse,¹⁵ the portion corresponding to the lowest-frequency bin is seldom sparse. In Ref. 16, Dwork et al. used this property with Fourier sampling to specify the size and shape of a region centered on the 0 frequency that is fully sampled and to alter the optimization problem to solve for a modified vector. The resulting problem is then converted into the standard form of the BPD problem. Thus, image reconstruction becomes a two-step process: (1) estimate a blurry image with the fully-sampled low-frequency portion of the sampled region and (2) enhance the image with high frequency details that result from solving a BPD problem. The system matrix of the optimization problem remains the same, thus all theoretical guarantees of compressed sensing that pertain to the original BPD problem apply to the modified problem. Because the sparsity of the optimization variable in the new problem is higher than that of the original problem, the error of the result is reduced.¹⁷

Redundant dictionaries (also called overcomplete bases) offer a representation that permit a more sparse representation than a basis.⁹ In this work, we consider the overcomplete basis comprised of the inverse wavelet and inverse curvelet transformations for two-dimensional image reconstruction problems. Notably, like the wavelet transform, although the curvelet transform of a natural image is sparse there is a low-frequency portion of the transform is not sparse. In this work, we extend the technique of¹⁶ and simultaneously take advantage of this low-frequency structure in both transforms. (An early version of this work was presented at the 2021 annual meeting of the Society for Industrial and Applied Mathematicians.) This leads to an incremental improvement in quality at the cost of additional computations required for the curvelet transform and its inverse. Again, the theoretical guarantees of compressed sensing that pertain to reconstruction with a redundant vector set remain. The increased sparsity attainable with the redundant dictionary further reduces the theoretical bound on the error. We will demonstrate a reduced error empirically with several images.

2 Methods

In this section, we review the BPD optimization problem and the theoretical results of compressed sensing. We will then discuss compressed sensing using the redundant dictionary comprised of wavelet and curvelet vectors. And we will show a method to utilize the structure of the sparsifying transform to better sample the Fourier domain and increase the sparsity of the optimization variable. Results of reconstructing images with these algorithms are presented in Sec. 4.

2.1 Background

Let x be the desired image and b be the measured data where $Ax + e = b$. Here, A is called the system matrix, and e is an additive Gaussian noise vector of independent and identically distributed values. Let Ψ be the sparsifying linear transformation. Let y be a sparse vector such that the desired image is approximated well by Ψ^*y . Then one can estimate y by solving the following optimization problem

$$\underset{y}{\text{Minimize}} \ \|y\|_0 \quad \text{subject to} \ (1/2)\|Ay - b\|_2 \leq \sigma, \quad (1)$$

where σ^2 is a bound on the noise power. With this notation, Ψ^* is included in A (see below). This problem, called the sparse signal recovery problem, is an np-hard combinatorial optimization problem. The BPD problem is a related problem where the L_0 penalty is replaced with the L_1 norm

$$\underset{y}{\text{Minimize}} \ \|y\|_1 \quad \text{subject to} \ (1/2)\|Ay - b\|_2 \leq \sigma. \quad (2)$$

This is a convex optimization problem and can be solved with known efficient algorithms.

One might hope that solving Eq. (2) gets us close to an answer for Eq. (1). Indeed, the theory of compressed sensing dictates that when A satisfies specific properties, (e.g., the mutual

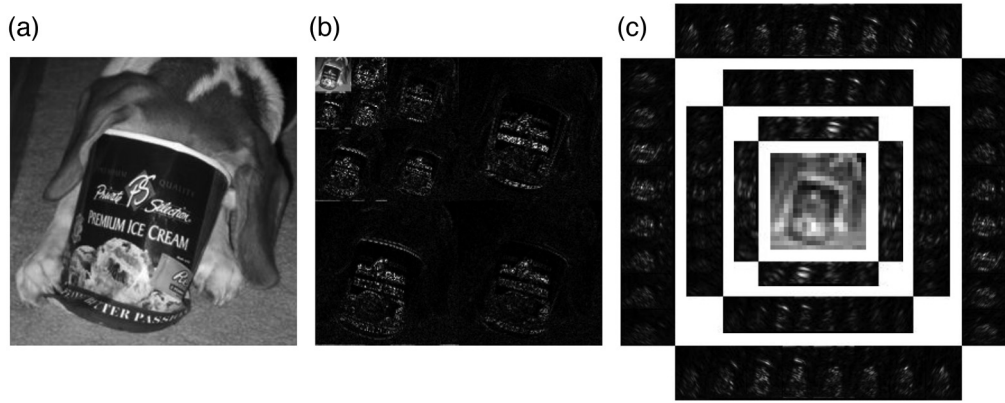


Fig. 1 (a) Original image, (b) magnitude of the DDWT-4 coefficients, and (c) magnitude of the curvelet transform coefficients.

coherence conditions,¹⁸ the restricted isometry property,¹⁹ or the restricted isometry property in levels¹⁷) then the solution to the BPD problem is optimal for the sparse signal recovery problem with very high probability. Moreover, the more sparse y is, the fewer the number of measurements that are required to attain the solution.^{17,19} This latter fact is the leverage that we will exploit to generate high quality images with fewer samples in this manuscript.

With standard compressed sensing and Fourier sensing, b is the vector of measured Fourier values and $A = MF\Psi^*$ where Ψ is the sparsifying transformation (which may have more rows than columns), $*$ represents the adjoint of the transformation, F is the unitary discrete Fourier transform, and M is the sampling mask that identifies those samples that were collected. An effective choice for Ψ is the discrete Daubechies wavelet transform (DDWT). Figure 1(a) shows an image of a dog eating from an ice cream container, and Fig. 1(b) shows the magnitude of the DDWT of order 4 (DDWT-4) coefficients when applied recursively to the lowest-frequency bin four times. The vast majority of the resulting coefficients appear black, which demonstrates the sparsifying behavior of the DDWT-4. The upper left corner of the transform is not black; this is a low-pass filter and downsampling of the original image and is seldom sparse with natural images.

Let W represent the orthogonal DDWT-4. To satisfy the Nyquist–Shannon sampling theorem and accurately estimate the lowest frequency bin of the DDWT-4, one must collect a fully sampled region (FSR) with size equal to that of the lowest-frequency bin.¹⁶ This is the same FSR as specified by the two-level sampling scheme in.¹⁷ From this FSR, a blurry image can be reconstructed.¹⁶ When using the DDWT-4 as the sparsifying transformation, optimizing for the remaining details increases the sparsity of the resulting optimization variable. The missing details of the image can be estimated by solving the following problem

$$\underset{y}{\text{Minimize}} \|y\|_1 \quad \text{subject to} \quad \|MF(\Psi^*y + x_L) - b\|_2 \leq \sigma, \quad (3)$$

where $\Psi = W$ and $x_L = F^*K_B M_L b$. Here, K_B is a Kaiser–Bessel window used to reduce ringing in the blurry estimate and M_L is a mask that specifies the fully sampled low-frequency region. Equation (3) is equivalent to

$$\underset{y}{\text{Minimize}} \|y\|_1 \quad \text{subject to} \quad (1/2)\|MF\Psi^*y - \beta\|_2 \leq \sigma, \quad (4)$$

where $\beta = b - MFx_L$. By letting $A = MF\Psi^*$, it becomes apparent that this problem has the form of the BPD Eq. (2) where b has been replaced by β . Once the solution y^* is determined, the completely reconstructed image is $x^* = x_L + \Psi^*y^*$. We call Eq. (4) the structured BPD (S-BPD) problem.

2.2 Curvelets

Curvelets offer an opportunity to further increase the sparsity of the optimization variable. The curvelet transformation permits a representation of an image using a multiscale pyramid with several directions available at each position (as opposed to wavelets, which decompose an image into horizontal and vertical components). Note that the discrete curvelet transform is a tall transform (it is a tight frame²⁰) with more coefficients than there are pixels in the original image. Using curvelets as the sparsifying transform with compressed sensing image reconstruction has been shown to reduce the blocking artifacts common to wavelet representations.²¹ Figure 1(c) shows the magnitude of the discrete curvelet coefficients of Fig. 1(a). (Note that the display of the curvelet coefficients includes white-space to ease interpretation; it was generated using the CurveLab software.²²)

Similar to the low-frequency bin of the wavelet transform, the center portion of the curvelet transform is the result of windowing, low-pass filtering, and downsampling the original image.⁸ Thus, like the wavelet transform, this region will not be sparse for most natural images.

2.3 Structured Sparsity with Curvelets and Wavelets

Consider the redundant dictionary comprised of curvelets and wavelets. The number of significant nonzero coefficients would never increase with the addition of curvelets into the dictionary for the image can always be represented with the same wavelet coefficients and with all curvelet coefficients equal to 0. Intuitively, one would expect that increasing the redundancy of the set of vectors would permit an even more sparse representation of an image. Suppose, for example, one wanted to represent $v = w_i + c_j$ where w_i is an element of the wavelet basis, and c_j is an element of the curvelet basis.²³ If one were to express v as a linear combination of vectors from both bases then the number of nonzero linear coefficients required would be 2. However, if one were to try to represent the same image v using only the wavelet basis, then it would require a larger number of significant linear coefficients.

Recall that, in accordance with the theorems of compressed sensing, increased sparsity permits an accurate reconstruction with fewer data samples. Indeed, after the benefit was shown empirically,^{24,25} the theory was developed to show that accurate reconstruction was possible with a redundant (or overcomplete) incoherent set of vectors.^{9,26} We will now present a method analogous to that of¹⁶ that simultaneously takes advantage of the low-frequency structure present in both the curvelet and wavelet transforms of natural images.

Because sparsity cannot be assumed for low-frequency region of the curvelet transform, we rely on the Nyquist–Shannon sampling theorem. This dictates a FSR with spacing equal to the inverse of the full size of the image and with size equal to the center portion of the curvelet transform. The values of the FSR are used to create a blurry image by applying a low-pass window (W_0 as specified in Ref. 8) and then performing an inverse discrete Fourier Transform. Figure 2(a) shows this blurry estimate for the image of Fig. 1(a). Figure 2(b) shows the result of subtracting this blurry estimate from the original image, and Fig. 2(c) shows the curvelet transform of Fig. 2(b). When comparing Fig. 1(c) with Fig. 2(c), the increased sparsity of the low-frequency region is apparent.

This is usually the case for natural images: after subtracting the blurry estimate from the sparsity of the curvelet transform is increased.

As discussed, both wavelets and curvelets have a region that is a low-pass filter and downsampling of the original image. The larger of these two regions dictates an FSR centered on 0 frequency to be collected that satisfies the Nyquist–Shannon sampling theorem. From this FSR, a blurry frequency estimate is created. If the center of the curvelet is larger than the lowest-frequency bin of the wavelet transform, then the blurry estimate is generated according to Sec. 2.3. Otherwise, the blurry estimate is generated according to Ref. 16. Denote this blurry estimate as x_L . Let $\Psi = (W, C)$, the block-matrix comprised of the wavelet and curvelet transforms. As discussed, the columns of Ψ^* may offer a more sparse representation of the image. This redundancy makes a more sparse representation of the optimization variable available, leading to a reconstruction with a reduced error bound. The image is reconstructed by solving Eq. (4) for y^* and setting $x^* = x_L + \Psi^*y^*$.




Fig. 2 (a) Blurry estimate of Fig. 1(a), (b) result of subtracting the blurry estimate from the original image, and (c) magnitude of the curvelet transform coefficients of panel b. Note that the low-frequency center region in panel c is much more sparse than the corresponding region of Fig. 1(c).

3 Experiment Design

The six images studied in this manuscript, along with their sizes, are shown in Fig. 3. All images were scaled by their maximum so that pixel values lied within $[0, 1]$. The data for Fig. 3(a) and 3(b) were collected with MRI machines. The data for Fig. 3(a) were collected with a 3 Tesla clinical MRI using a Cartesian trajectory, a field of view of $25.6 \times 25.6 \text{ cm}^2$ and a 1-mm slice thickness. The data of Fig. 3(a) were gathered with institutional review board (IRB) approval, health insurance portability and accountability act (HIPAA) compliance, and patient informed assent/consent. The data for Fig. 3(b) were made available to the public at Ref. 27. The data for Figs. 3(c)–3(f) were made publicly available from various sources and are often studied in image processing research papers. All data were initially fully sampled and retrospectively downsampled for processing with the algorithms presented in this manuscript.

The MRI data of the knee [Fig. 3(b)] are from a fully sampled three-dimensional dataset. The data were transformed into a the hybrid-space of spatial frequency in the anterior–posterior and superior–inferior dimensions and location in the left–right dimension. Afterwards, a two-dimensional slice located near the center of the knee in the left–right dimension was isolated for further processing in this manuscript.

Other works have attempted to address the difference in sparsity of the lowest-frequency bin of the sparsifying transform by minimizing an objective function with a weighted norm, as follows

$$\text{Minimize } \|y\|_{\omega,1} \quad \text{subject to } (1/2)\|Ay - \beta\|_2 \leq \sigma, \quad (5)$$

where $\|y\|_{\omega,1} = \omega_1|y_1| + \omega_2|y_2| + \dots + \omega_N|y_N|$. In Ref. 28, Candes et al. iteratively update the weight vector ω by setting it inversely proportional to the intensity of the corresponding




Fig. 3 The six images analyzed in this manuscript along with their sizes in pixels. The data of panels (a) and (b) were collected with MRI machines; the data of panels (c)–(f) were collected with optical cameras. The data of panel (b) were made available to the public at Ref. 27.

coefficient. In Ref. 29, Varela-Mattatall et al. set the weight ω for each bin of the wavelet transform independently by analyzing the zero-filled reconstruction and they set the elements of ω for the lowest-frequency bin to 0. We will analyze a noniterative version of these techniques where $\omega_i = 0$ if the pixel in the transformed domain corresponds to the lowest-frequency bin and $\omega_i = 1$ otherwise. That is, ω is a mask that specifies whether or not the pixel is an element of a high-frequency bin of the sparsifying transform. We will denote solving Eq. (5) with this definition of ω as BPD with mask.

For all results presented, the fast iterative shrinkage threshold algorithm (FISTA) with line search^{30,31} run for 100 iterations was used to solve the equivalent Lagrangian form of the problem. For example, to solve Eq. (2), the following equivalent optimization problem was solved

$$\text{Minimize } (1/2)\|Ay - \beta\|_2^2 + \|y\|_{\lambda,1}, \quad (6)$$

where λ is the regularization vector. The iterative reweighting scheme of Ref. 28 was used to automatically determine λ with the number of reweighting iterations equal to 5. The initial value of λ was set to the corresponding values of the wavelet coefficients determined with the zero-filled reconstruction (the reconstruction assuming all values not sampled are equal to 0). The error metrics reported are relative error, defined as $e = \|\text{true} - \text{estimate}\|_2 / \|\text{true}\|_2$, mean squared error (MSE), and the structural similarity metric (SSIM).³²

Three different sparsifying transformations were tested: DDWT-4, the wrapping version of the discrete curvelet transform, and a redundant dictionary comprised of both of these transforms; we denote these transformations as wavelet, curvelet, and wavCurv, respectively. For each transformation, two different sampling patterns were tested: a variable density sampling pattern and a variable density sampling pattern with an FSR. (Due to spatial limitations, not all results are presented.) The sampling patterns used for the results in this manuscript were generated using a separable Laplacian distribution. Examples are shown in Fig. 4. When an FSR was included, the number of variable density samples was reduced to retain the same total number of samples.




Fig. 4 Example variable density sampling patterns generated from a separable Laplacian distribution with a standard deviation of 0.3 the length of the sides. The top/bottom row shows sampling patterns without/with the fully sampled center region corresponding to the size of the low-frequency region of the curvelet transform, respectively. From left to right, the sampling patterns include 10,000, 12,000, 15,000, and 20,000 points, respectively; this corresponds to 15%, 18%, 23%, and 31% of the total number of samples.




Fig. 5 (Left) Sampling masks without and with the fully sampled center region. (Right) Reconstructions of the knee with 8192 samples (8% of the total number of pixels). The relative error for each reconstruction is written in the corresponding image.

4 Results

Figure 5 shows reconstructions of Fig. 3(b) from 8192 samples (8% of the number required for full sampling). The sampling mask was generated from a separable Laplacian distribution with a standard deviation of 0.3 the length of the sides of the image in each dimension. The reconstructions from sampling patterns that include the FSR are far superior to those without it. When the FSR is included, the relative error for the reconstruction from the structured BPD problem is slightly smaller than those reconstructed with BPD or BPD with mask.

Figure 6 shows reconstructions of Fig. 3(b) from 8192 samples with sampling masks generated from separable Laplacian distributions with standard deviations of 0.2, 0.3, and 0.4 along with the magnitude of the difference images. The standard deviation has a significant effect on the quality of the reconstruction and the metric value.

Figure 7 shows the results of solving the Lagrangian form of Eq. (4) for three different sparsifying transformations: wavelets, curvelets, and the redundant dictionary made of both wavelets and curvelets. We see that the wavelet sparsifying transformation achieves a result with less relative error than the curvelet sparsifying transformation. Notably, though, the masts of the boat appear sharper when using curvelets than when using wavelets. However, there are striations in the clouds in the curvelet reconstruction, which are not present in the original image or in the wavelet reconstruction. The lowest relative error is attained using both wavelets and curvelets, and the reconstruction is able to reduce the block artifacts on the boat mast associated with wavelets without adding striations into the clouds. The reconstruction from wavelets alone presents blocking artifacts, and the reconstruction from curvelets alone presents striping artifacts with isolated high level values; both artifacts are well known.³³ The reconstruction of the redundant dictionary presents a combination of both artifacts and does so in a way that recapitulates the details of the reference image better.

Figure 8 shows plots of MSE and SSIM for reconstructions of Fig. 3(a) generated with structured BPD using sparsifying transformations of wavelets, curvelets, and the redundant dictionary versus number of samples. The redundant dictionary generates reconstructions with lower (better) relative error and higher (better) structural similarity for all numbers of samples.

Table 1 shows the MSE and structural similarity for reconstructions generated with BPD, BPD with a regularization mask, and structured BPD using sampling masks without and with the fully sampled center region and 10% of the samples required to satisfy the Nyquist-Shannon sampling theorem. For all the images of Fig. 3, structured BPD attains the best quality metrics.




Fig. 6 Reconstructions of the knee with 8192 sampling patterns generated from a separable Laplacian distribution with a standard deviation of (left) 0.2, (center) 0.3, and (right) 0.4. The relative error for each reconstruction is written in the corresponding image.

Figure 9 presents differences in the values of SSIM and the sum of mean absolute error (MAE) with MSE for reconstructions from solving Eq. (6) when using the redundant dictionary of wavelets and curvelets plotted against a dictionary comprised exclusively of wavelets and a dictionary comprised exclusively of curvelets. SSIM is a perceptually motivated metric while MAE + MSE attempts to present a more balanced metric than either MAE or MSE alone.³⁵ The results are plotted for 158 different images from the ImageNet database where each image was randomly selected from a separate class.³⁴ The Fourier values of each square image were sampled with an FSR added to a variable density Laplacian distribution using a parameter value of 0.3 the length of the image's side. As can be seen, the SSIM value when using the redundant dictionary is almost always higher than when using wavelets alone or curvelets alone, indicating that the quality is improved when using the redundant dictionary.

For the majority of images, the redundant dictionary improves the quality of the result with respect to either of the metrics evaluated. When comparing the redundant dictionary to wavelets, the improvement is usually more pronounced when a small fraction of samples was used. When comparing the redundant dictionary to curvelets, the improvement is usually more pronounced when a larger fraction of samples was used. Notably, for some images, the redundant dictionary reduces the value of the quality metric. This is more often the case when comparing the redundant dictionary to curvelets and in spite of the fact that curvelet regularization can cause striping artifacts and isolated high values (sometimes described as salt and pepper noise).³³ Thus, even




Fig. 7 Reconstructions of Fig. 3(e) using the structured compressed sensing algorithm using three different sparsifying transformations: wavelets, curvelets, and wavelets and curvelets. An auto-calibration region was included as shown in the second row of Fig. 4 and the Laplacian variable density sampling pattern was generated with a standard deviation of 75 pixels.




Fig. 8 Plots of error metrics versus sampling percentage for brain reconstructions.

though the metrics are reduced, the images may appear to be better quality to the human eye. This is the case for the boat reconstruction of Fig. 7

5 Discussion and Conclusion

In this work, we show that the low-frequency structure present in the wavelet and curvelet transforms can be simultaneously exploited to improved the quality when using a redundant dictionary comprised of both transforms. This exploitation comes in two forms: (1) specifying the size

Table 1 Structural similarity and MSE metric values for reconstructions from 10% of samples using a sampling mask generated from a separable Laplacian distribution with a standard deviation of 0.3 excluding and including an autocalibration region using a sparsifying transformation comprised of wavelets and curvelets. The image labels are according to Fig. 3.

Image	Without FSR				With FSR					
	SSIM		MSE · 10 ³		SSIM			MSE · 10 ³		
	BPD	BPD with mask	BPD	BPD with mask	BPD	BPD with mask	S-BPD	BPD	BPD with mask	S-BPD
a	0.687	0.613	6.20	7.20	0.968	0.967	0.969	0.656	0.670	0.638
b	0.876	0.871	7.88	8.18	0.981	0.981	0.981	1.24	1.25	1.22
c	0.790	0.788	19.3	19.4	0.973	0.972	0.973	2.95	3.05	2.92
d	0.00	0.00	236	236	0.939	0.939	0.940	6.03	6.02	6.00
e	0.823	0.821	12.2	12.2	0.947	0.946	0.948	4.08	4.12	4.03
f	0.700	0.700	26.4	26.5	0.969	0.969	0.969	3.63	3.68	3.62




Fig. 9 Improvements in the SSIM and MAE plus MSE for images reconstructed with structured BPD with the redundant dictionary of wavelets and curvelets plotted against exclusively wavelets and plotted against exclusively curvelets for 158 different images (each from a separate class) taken from the ImageNet dataset.³⁴ The color indicates the fraction of Fourier domain pixels that were sampled.

and shape of a fully sampled center region to satisfy the Nyquist–Shannon sampling theorem and (2) a structured BPD problem that determines the details. We show that this improves the quality over standard compressed sensing reconstructions and that it improves the quality over a masked version of compressed sensing where the values of the lowest-frequency bins of the wavelet and curvelet transforms are not penalized. Taking advantage of the fully sampled center region improved the quality dramatically for the images studied in this manuscript. The structured BPD problem improved the results further still but by a much smaller amount.

The MRI data presented were from three-dimensional data collections that made a hybrid space (two dimensions of frequency and one of location) possible. For some MRI experiments, the data is collected in such a way that a more general three-dimensional sparsifying transformation is required (e.g., using the cones³⁶ or yarn ball³⁷ trajectories). For these applications,

the methods of this manuscript can be extended using the three-dimensional discrete curvelet transformation.³⁸

Notably, the low-frequency filter of the the fast discrete curvelet transform of CurveLab has a much higher order than that of the DDWT-4 wavelet transform. This suggests that a lower order filter could be used, which would permit a nonnegligible amount of aliasing, but where the quality of the reconstruction would be retained when using a fully sampled center region. This would reduce the computations required to implement the curvelet transform and its adjoint, which would reduce the time required to attain a reconstruction.

FISTA was used to solve the optimization problems for this manuscript. However, after exploiting the low-frequency structure, the optimization variable becomes much more sparse. It may be sparse enough that an orthogonal matching pursuit (OMP) algorithm could be used to solve the problem well. Candidate algorithms include compressive sampling matching pursuit³⁹ and stagewise OMP.⁴⁰ These are greedy algorithms, so the result may be attained with less time than when using FISTA.

Finally, additional gains in reconstruction quality may be had by either replacing the curvelet sparsifier with wave-atoms^{41,42} or hexagonal wavelets,⁴³ or by augmenting the redundant dictionary with these sets. We expect these gains to be small, and so their usage would depend on the computational power available and the importance of fine details.

We leave all of these prospects as future work.

6 Compliance with Ethical Standards

All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards. MR data of humans were gathered with IRB approval and HIPAA compliance. Informed consent was obtained from all individual participants included in the study.

Acknowledgments

The authors would like to thank Daniel O'Connor for helpful conversations on frames and optimization. ND has received postdoctoral training funding from the American Heart Association (Grant No. 20POST35200152). ND has received funding from the Quantitative Biosciences Institute at UCSF. ND and PL have received funding from the the National Institutes of Health (Grant No. NIH R01 HL136965). No conflicts of interest, financial or otherwise, are declared by the authors.

References

1. B. Adcock and A. C. Hansen, *Compressive Imaging: Structure, Sampling, Learning*, Cambridge University Press (2021).
2. M. Lustig et al., "Compressed sensing MRI," *IEEE Signal Process. Mag.* **25**(2), 72–82 (2008).
3. C. A. Baron et al., "Rapid compressed sensing reconstruction of 3D non-cartesian MRI," *Magn. Reson. Med.* **79**(5), 2685–2692 (2018).
4. J. Ma, "Improved iterative curvelet thresholding for compressed sensing and measurement," *IEEE Trans. Instrum. Meas.* **60**(1), 126–136 (2010).
5. H. Lee et al., "Efficient sparse coding algorithms," in *Adv. Neural Inf. Process. Syst.*, pp. 801–808 (2007).
6. S. Mallat, *A Wavelet Tour of Signal Processing*, Elsevier (1999).
7. G. Beylkin, R. Coifman, and V. Rokhlin, *Fast Wavelet Transforms and Numerical Algorithms*, pp. 741–783. Princeton University Press (2009).
8. E. Candes et al., "Fast discrete curvelet transforms," *Multiscale Model. Simul.* **5**(3), 861–899 (2006).

9. E. J. Candès et al., “Compressed sensing with coherent and redundant dictionaries,” *Appl. Comput. Harmon. Anal.* **31**(1), 59–73 (2011).
10. M. Lustig, D. Donoho, and J. M. Pauly, “Sparse MRI: the application of compressed sensing for rapid MR imaging,” *Magn. Reson. Med.* **58**(6), 1182–1195 (2007).
11. D. J. Brady et al., “Compressive holography,” *Opt Express* **17**(15), 13040–13049 (2009).
12. Y. Oike and A. El Gamal, “CMOS image sensor with per-column $\sigma\delta$ ADC and programmable compressed sensing,” *IEEE J. Solid-State Circuits* **48**(1), 318–328 (2012).
13. H. Huang et al., “Applications of compressed sensing in communications networks,” arXiv 1305.3002 (2013).
14. R. N. Mahalati, R. Y. Gu, and J. M. Kahn, “Resolution limits for imaging through multi-mode fiber,” *Opt. Express* **21**(2), 1656–1668 (2013).
15. A. Majumdar and R. K. Ward, “On the choice of compressed sensing priors and sparsifying transforms for MR image reconstruction: an experimental study,” *Signal Process. Image Commun.* **27**(9), 1035–1048 (2012).
16. N. Dwork et al., “Utilizing the wavelet transform’s structure in compressed sensing,” *Signal, Image Video Process.* **15**(7), 1407–1414 (2021).
17. B. Adcock et al., “Breaking the coherence barrier: a new theory for compressed sensing,” in *Forum of Mathematics, Sigma*, Vol. 5, Cambridge University Press (2017).
18. D. L. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization,” *Proc. Natl. Acad. Sci. U. S. A.* **100**(5), 2197–2202 (2003).
19. E. J. Candès and M. B. Wakin, “An introduction to compressive sampling,” *Signal Process. Mag.* **25**(2), 21–30 (2008).
20. B. Adcock and D. Huybrechs, “Frames and numerical approximation,” *SIAM Rev.* **61**(3), 443–473 (2019).
21. A.-C. Liew and H. Yan, “Blocking artifacts suppression in block-coded images using overcomplete wavelet representation,” *Trans. Circuits Syst. Video Technol.* **14**(4), 450–461 (2004).
22. E. Candès, “Curvelab,” <http://www.curvelet.org/software.html> (accessed 2021-07-20).
23. J.-L. Starck, E. J. Candès, and D. L. Donoho, “The curvelet transform for image denoising,” *Trans. Image Process.* **11**(6), 670–684 (2002).
24. G. Peyre, “Best basis compressed sensing,” *IEEE Trans. Signal Process.* **58**(5), 2613–2622 (2010).
25. M. Doneva et al., “Compressed sensing reconstruction for magnetic resonance parameter mapping,” *Magn. Reson. Med.* **64**(4), 1114–1120 (2010).
26. H. Rauhut, K. Schnass, and P. Vandergheynst, “Compressed sensing and redundant dictionaries,” *Trans. Inf. Theory* **54**(5), 2210–2219 (2008).
27. F. Ong et al., “Mridata.org: an open archive for sharing MRI raw data,” in *Proc. Int. Soc. Mag. Reson. Med.*, 26 (2018).
28. E. J. Candès, M. B. Wakin, and S. P. Boyd, “Enhancing sparsity by reweighted ℓ_1 minimization,” *J. Fourier Anal. Appl.* **14**(5), 877–905 (2008).
29. G. Varela-Mattatall, C. A. Baron, and R. S. Menon, “Automatic determination of the regularization weighting for wavelet-based compressed sensing MRI reconstructions,” *Magn. Reson. Med.* **86**(3), 1403–1419 (2021).
30. A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *J. Imaging Sci.* **2**(1), 183–202 (2009).
31. K. Scheinberg, D. Goldfarb, and X. Bai, “Fast first-order methods for composite convex optimization with backtracking,” *Found. Comput. Math.* **14**(3), 389–417 (2014).
32. Z. Wang, E. P. Simoncelli, and A. C. Bovik, “Multiscale structural similarity for image quality assessment,” in *The Thirty-Seventh Asilomar Conf. Signals, Syst. & Comput.*, IEEE, Vol. 2, pp. 1398–1402 (2003).
33. M. Wiecek et al., “X-ray computed tomography using curvelet sparse regularization,” *Med. Phys.* **42**(4), 1555–1565 (2015).
34. O. Russakovsky et al., “Imagenet large scale visual recognition challenge,” *Int. J. Comput. Vis.* **115**(3), 211–252 (2015).
35. G. Wang et al., “Stochastic optimization of 3D non-cartesian sampling trajectory (SNOPY),” arXiv:2209.11030 (2022).

36. P. T. Gurney, B. A. Hargreaves, and D. G. Nishimura, "Design and analysis of a practical 3D cones trajectory," *Magn. Reson. Med.* **55**(3), 575–582 (2006).
37. R. W. Stobbe and C. Beaulieu, "Three-dimensional yarnball k-space acquisition for accelerated MRI," *Magn. Reson. Med.* **85**(4), 1840–1854 (2021).
38. L. Ying, L. Demanet, and E. Candes, "3D discrete curvelet transform," *Proc. SPIE* **5914**, 591413 (2005).
39. D. Needell and J. A. Tropp, "CoSaMP: iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harmon. Anal.* **26**(3), 301–321 (2009).
40. D. L. Donoho et al., "Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit," *IEEE Trans. Inf. Theory* **58**(2), 1094–1121 (2012).
41. L. Demanet and L. Ying, "Curvelets and wave atoms for mirror-extended images," *Proc. SPIE* **6701**, 67010J (2007).
42. L. Demanet and L. Ying, "Wave atoms and sparsity of oscillatory patterns," *Appl. Comput. Harmon. Anal.* **23**(3), 368–387 (2007).
43. W. Zhu and I. Daubechies, "Constructing curvelet-like bases and low-redundancy frames," arXiv 1910.06418 (2019).

Nicholas Dwork received his BS degree in electrical engineering from California Polytechnic State University, San Luis Obispo, California, United States, in 2001 and received his MS in biomedical engineering from the University of California, Los Angeles in 2004. He received a PhD in electrical engineering from Stanford University in 2019; his advisor was John Pauly. Currently, he is working as an assistant professor in the Department of Biomedical Informatics with a joint appointment in the Department of Radiology at the University of Colorado School of Medicine. He previously worked as a postdoctoral scholar in the Surbeck Advanced Imaging Laboratory, the University of California, San Francisco, California, United States, advised by Peder Larson. His research interests include reducing the scan time of MRI and quantification in medical imaging.

Peder E. Z. Larson, PhD, is working as an associate professor in residence and a principal investigator in the Department of Radiology and Biomedical Imaging at the University of California, San Francisco, California, United States. His research program is primarily centered around developing new MRI scanning and reconstruction technology for improved clinical outcomes. These developments in his group are aimed at several applications: metabolic imaging methods using hyperpolarized carbon-13, semisolid tissue MRI, and PET/MRI. He is an active member of International Society for Magnetic Resonance in Medicine, the Institute for Electrical and Electronics Engineering, the UC Berkeley and UCSF Graduate Group in Bioengineering, and the California Institute for Quantitative Biosciences.