

# Benchmarking Compressed Sensing, Super-Resolution, and Filter Diagonalization

Thomas Markovich, Samuel M. Blau, Jacob N. Sanders, and Alán Aspuru-Guzik\*

Signal processing techniques have been developed that use different strategies to bypass the Nyquist sampling theorem in order to recover more information than a traditional discrete Fourier transform. Here we examine three such methods: filter diagonalization, compressed sensing, and super-resolution. We apply them to a broad range of signal forms commonly found in science and engineering in order to discover when and

how each method can be used most profitably. We find that filter diagonalization provides the best results for Lorentzian signals, while compressed sensing and super-resolution perform better for arbitrary signals. © 2016 Wiley Periodicals, Inc.

DOI: 10.1002/qua.25144

## Introduction

The reconstruction of frequency-resolved spectra from limited and undersampled measurements in the time domain is a significant problem throughout the physical sciences. The standard approach to solving such a problem is the discrete Fourier transform, which decomposes a time series in terms of its component frequencies (or, more generally, decomposes a series into its conjugate domain). The discrete Fourier transform offers two major advantages: no *a priori* knowledge about the signal is required, and the computation can be implemented very efficiently via the fast Fourier transform. Unfortunately, a major disadvantage is that the discrete Fourier transform imposes a natural bound on the maximum frequency resolution possible given the nature of the time series, known as the Shannon–Nyquist condition.<sup>[1]</sup>

A natural question to ask is whether the Shannon–Nyquist condition can be bypassed by exploiting any additional knowledge we may have about the signal. Recent advances in signal processing have provided many such techniques for leveraging additional *a priori* knowledge about the signal to improve reconstruction. Our goal in this article is to compare three such methods, filter diagonalization,<sup>[2–7]</sup> compressed sensing,<sup>[3,8–14]</sup> and super-resolution,<sup>[1,15–19]</sup> against a series of test signals in order to understand their relative strengths and weaknesses. Our comparison will be based on a subset of the signals contained in the Sparco toolbox,<sup>[20]</sup> a Gaussian, a sum of random Lorentzians, and the Jacob's Ladder signal.<sup>[2]</sup> The Sparco toolbox provides a standard set of signal processing benchmarks while the other signals are commonly encountered throughout the physical sciences.

Filter diagonalization, one of the earliest techniques for bypassing the Shannon–Nyquist condition, assumes that the time series is generated by an underlying dynamical system with a frequency spectrum modeled by a sum of Lorentzians. It attempts to express the frequency spectrum as a sum of Lorentzian peaks by finding the optimal frequencies, linewidths, and intensities that fit the time series. Filter diagonalization has been applied to a broad

range of signals that vary from NMR spectra<sup>[2,6]</sup> to scattering data<sup>[5]</sup> and image analysis.<sup>[4]</sup>

More recently,  $\mathcal{L}_1$  minimization techniques, such as compressed sensing and super-resolution, have also been proposed as an alternative technique for sampling below the rate imposed by the Nyquist–Shannon condition. Rather than assuming a particular kind of underlying dynamical system, these techniques simply assume that the signal is sparse in some *a priori* known basis. The two methods differ in both sampling strategy and the particular optimization problem to be solved. Compressed sensing is designed to recover sparse frequency spectra (or other signals) by randomly undersampling data over the entire time domain, and then minimizing the  $\mathcal{L}_1$  norm of an underdetermined system of linear equations. Compressed sensing has been successfully applied to data acquisition in many different areas,<sup>[21]</sup> including the improvement of the resolution of medical magnetic-resonance imaging<sup>[22]</sup> and the experimental study of atomic and quantum systems.<sup>[11,14,23]</sup>

Super-resolution is a related technique that shares the spirit of compressed sensing, but with a different sampling technique.<sup>[3,8–10,12,22,24–28]</sup> Super-resolution was developed to recover sparse frequency spectra (or other signals) from regularly sampled data over a short segment of the time domain. It provides a provably convergent algorithm for the reconstruction of signals from these limited time-domain measurements by using a total-variation minimization procedure. Like compressed sensing, super-

T. Markovich, S. M. Blau, J. N. Sanders, Alán Aspuru-Guzik  
Department of Chemistry and Chemical Biology, Harvard University,  
Cambridge, Massachusetts

E-mail: [aspuru@chemistry.harvard.edu](mailto:aspuru@chemistry.harvard.edu)

Contract grant sponsor: Defense Advanced Research Projects Agency;  
contract grant number: N66001-10-1-4063.

Contract grant sponsor: Defense Threat Reduction Agency;  
contract grant number: HDTRA1-10-1-0046.

Contract grant sponsor: National Science Foundation (NSF) through the Graduate Research Fellowship Program (GRFP), Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program, Department of Energy (DoE) through the Computational Sciences Graduate Fellowship (CSGF), Corning Foundation..

© 2016 Wiley Periodicals, Inc.

resolution has been applied to a broad range of scientific problems, including image<sup>[18]</sup> and video compression,<sup>[29]</sup> image denoising,<sup>[30]</sup> atomistic modeling of open quantum systems,<sup>[31]</sup> astronomy,<sup>[17]</sup> microscopy,<sup>[16]</sup> and medical imaging.<sup>[15]</sup>

The goal of this article is to elucidate the strengths and weaknesses of the aforementioned signal processing techniques to provide a clear and coherent aid in choosing a method. To achieve this, we will first introduce the theory that underlies each method and outline our procedure for benchmarking the methods. Then, we will introduce the test signals and compare the performance of each method on each signal. Finally, we will present some general conclusions.

## Theory

### Discrete Fourier transform

The Fourier transform is a cornerstone method in signal processing, as it provides a technique for decomposing an arbitrary function of time into its component frequencies:

$$\hat{f}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\omega t} f(t) dt. \quad (1)$$

When treating a problem numerically, we often only have access to the values of the signal  $f(t)$  on an equally-spaced  $N$ -point grid. Accordingly, we discretize the continuous Fourier transform to obtain the discrete Fourier transform:

$$\hat{f}(\omega_i) = \frac{1}{\sqrt{N}} \sum_j f(t_j) e^{i\omega_i t_j}. \quad (2)$$

which can be reformulated as a matrix multiplication according to:

$$\hat{f}_i = \sum_j \mathcal{F}_{ij} f_j, \quad (3)$$

where  $f_j \equiv f(t_j)$ ,  $\hat{f}_i \equiv \hat{f}(\omega_i)$ , and  $\mathcal{F}_{ij}$  is the Fourier operator. Here we have assumed a uniform frequency grid with a spacing of  $f_s/N$ , a time sampling rate of  $\Delta t$ , the maximum frequency that can be sampled is  $1/\Delta t$ ,  $N$  is the number of time points, and  $T$  is the time length of the signal. The Nyquist–Shannon sampling theorem states that if a function is band limited with maximum frequency  $\Omega$ , it is completely characterized with a uniform series of time points spaced by  $1/2 \Omega$ . It is often more convenient to use the converse statement in signal reconstruction, which claims that with a sampling rate of  $\Delta t$  the maximum frequency that can be recovered is  $1/2 \Delta t$ . This is a direct consequence of discretizing the Fourier transform. A major disadvantage of the discrete Fourier transform is that a long and uniformly-sampled time series is required to obtain high resolution in the frequency domain.<sup>[1,32]</sup>

### $\mathcal{L}_1$ Minimization

$\mathcal{L}_1$  minimization methods, including compressed sensing and super-resolution, have emerged as a powerful technique for bypassing the constraint of the Shannon–Nyquist theorem in

the special case where the signal is known to be sparse in a particular basis.<sup>[8,19]</sup>

To illustrate this, suppose we have an unknown function  $f(t)$  that we wish to recover with as few samples as possible. Suppose further that we can find a set of basis functions  $\{g_j(t)\}$  such that  $f(t)$  is sparse when expanded in this basis. That is,

$$f(t) = \sum_j \lambda_j g_j(t), \quad (4)$$

where most of the  $\lambda_j$  expansion coefficients are equal to zero (or near zero). Our goal is to find the set of coefficients  $\{\lambda_j\}$ , since this would in turn identify the function  $f(t)$ . All we know *a priori* is that most of the  $\lambda_j$  are zero; we do not know which of them are zero, and we do not know their values in general. By sampling  $f(t)$  at a set of points  $\{t_i\}$ , we can obtain a set of linear equations,

$$f_i = \sum_j \lambda_j g_{ij}, \quad (5)$$

where  $f_i \equiv f(t_i)$  and  $g_{ij} \equiv g_j(t_i)$ , and our goal is to solve these equations for the set of coefficients  $\{\lambda_j\}$ . Since we are trying to obtain accurate resolution by taking as few time samples as possible, this system of equations will be underdetermined and we must impose additional constraints to pick out the desired solution. One such constraint is the minimization of the  $\mathcal{L}_1$  norm.

In  $\mathcal{L}_1$  optimization methods, including compressed sensing and super-resolution, the desired solution to the underdetermined system of Eq. (5) is chosen by solving the following  $\mathcal{L}_1$  minimization problem:

$$\operatorname{argmin}_{\lambda_j} \|\lambda_j\|_1 \text{ subject to } \left\| f_i - \sum_j \lambda_j g_{ij} \right\|_2 < \eta, \quad (6)$$

where  $\eta$  is a small thresholding parameter. In this minimization problem, the  $\mathcal{L}_1$  norm serves as a proxy for the sparsity-enforcing  $\mathcal{L}_0$  norm by selecting the sparsest set of coefficients  $\{\lambda_j\}$  such that the system of Eq. (5) are satisfied to within  $\eta$ . It is important to note that each  $g_j(t)$  should be normalized to unity so that no single basis function is privileged.

Compressed sensing and super-resolution differ in the sampling strategy, which, in turn, is often determined by computational and experimental constraints. Compressed sensing addresses the case where the value of the function  $f(t)$  is sampled at random points  $\{t_i\}$  over the entire domain. This random sampling of points  $f_i$  ensures that each point provides the maximum possible amount of information for the reconstruction of the signal.<sup>[8]</sup> A key result from compressed sensing is that the number of time samples  $f_i$  which must be measured for accurate recovery scales roughly with the sparsity of the basis expansion (i.e., the number of nonzero  $\lambda_j$ ), rather than the total size of the basis expansion (i.e., the total number of  $\lambda_j$ ).<sup>[8,32]</sup> In particular, Donoho showed if the basis obeyed a condition called the “restricted isometry property” (RIP), this basis expansion problem can be solved exactly provided that:

$$M \geq Ck \log(N/k), \quad (7)$$

where  $M$  is the number of measurements (frequently in the time domain),  $C$  is an instance specific constant,  $N$  is number of desired conjugate variables (frequently the frequencies), and  $k$  is the degree of sparsity. Many matrices obey RIP, including all sub-Gaussian random matrices and the Fourier matrix.<sup>[8]</sup>

A related method to compressed sensing is super-resolution. Unlike compressed sensing, which applies to randomly-sampled data, super-resolution applies to data that is regularly sampled in the time domain. Making no assumptions about the basis, aside from sparsity, it has been proven that super-resolution enables the recovery of signals with frequencies at one quarter of the Shannon–Nyquist condition reliably.<sup>[19]</sup>

A major advantage of both compressed sensing and super-resolution is that we can recover  $f(t)$  in any basis in which the signal is sparse. The methods work with bases as varied as wavelets,<sup>[33,34]</sup> treelets,<sup>[35]</sup> geometric harmonics,<sup>[36]</sup> and polynomials.<sup>[13,37]</sup> All that is required is that we know the sparse basis ahead of time. Although this may seem like a strong restriction, for many scientific problems physical intuition often leads to a sparse basis. One does not need to pick the optimal basis; any reasonably sparse basis will work. Moreover, both compressed sensing and super-resolution are robust to choosing an overcomplete basis, which allows for a lot more freedom in finding a sparse basis.<sup>[14]</sup>

For example, in computational chemistry, we are often interested in resolving spectra which are known to be sparse directly in the frequency domain (i.e., the spectrum is mostly zero except for a few sharp frequency peaks). In this case, we might choose a basis of complex exponentials  $g_j(t) = \frac{1}{2\pi} e^{i\omega_j t}$ . After time sampling, the matrix  $g_{ij} = \frac{1}{2\pi} e^{i\omega_j t_i}$  simply becomes an undersampled set of rows of the discrete Fourier transform matrix. Once the sparse coefficients  $\lambda_j$  have been found by solving Eq. (6), the final spectrum may be plotted as

$$\hat{f}(\omega) = \sum_j \lambda_j \hat{g}_j(\omega) = \sum_j \lambda_j \delta(\omega - \omega_j), \quad (8)$$

Other similar bases commonly used when applying compressed sensing or super-resolution to Fourier analysis are sine functions  $g_j(t) = \frac{1}{2\pi} \sin(\omega_j t)$  and cosine functions  $g_j(t) = \frac{1}{2\pi} \cos(\omega_j t)$ .

To take another common example in the physical sciences, we often find damped oscillatory signals which may be expressed as a sum of damped cosines:

$$g_{jk}(t) = e^{-\gamma_k t} \cos(\omega_j t). \quad (9)$$

Compressed sensing and super-resolution are easily adapted to this overcomplete basis and, once the sparse coefficients  $\lambda_{jk}$  have been found via Eq. (6), the final spectrum may be plotted as

$$\begin{aligned} \hat{f}(\omega) &= \sum_{j,k} \lambda_{jk} \hat{g}_{jk}(\omega) \\ &= \sum_{j,k} \frac{\lambda_{jk}}{\sqrt{2\pi}} \left( \frac{\gamma_k}{\gamma_k^2 + (\omega - \omega_j)^2} + \frac{\gamma_k}{\gamma_k^2 + (\omega + \omega_j)^2} \right). \end{aligned} \quad (10)$$

In short, compressed sensing and super-resolution both enable the recovery of an undersampled signal by using a customized, sparse basis that is appropriate to the problem at hand. The central difference between the two methods is the sampling technique, which defines the prior information available to the reconstruction. Namely, random sampling over the entire time domain is appropriate for compressed sensing, while regular sampling over the time domain is the defining characteristic of super-resolution. While the methods are closely algorithmically, the differences in sampling provide very different theoretical error bounds. In this sense, it is important to note that both techniques have rigorously defined scaling characteristics and computational bounds.

### Filter diagonalization

Filter diagonalization is another approach to circumvent the Shannon–Nyquist condition.<sup>[2,4–6]</sup> Inspired by quantum mechanics, the method assumes that the signal  $f(t)$  to be recovered is generated by the time evolution of a unitary propagator,<sup>[5]</sup>

$$f(t) = (\Phi_0, e^{-i\hat{U}t} \Phi_0). \quad (11)$$

If we sample  $f(t)$  on an equally spaced grid  $t_n = n\tau$ , we can discretize this equation as

$$f(t_n) = (\Phi_0, e^{-in\tau} \Phi_0). \quad (12)$$

where  $\hat{U} = e^{-i\tau}$  is the unitary propagator. By expanding the propagator in terms of its eigenvalues and (possibly complex) eigenvectors,

$$e^{-i\tau} = \sum_j e^{-i\omega_j \tau} |u_j\rangle \langle u_j|, \quad (13)$$

and substituting this expansion into Eq. (12), we obtain

$$f(t_n) = \sum_j |(u_j, \Phi_0)|^2 e^{-i\omega_j n\tau}, \quad (14)$$

which is the equation for a Lorentzian signal with (possibly damped) frequencies  $\omega_j$  and amplitudes  $\lambda_j \equiv |(u_j, \Phi_0)|^2$ . Hence, resolving a signal  $f(t)$  into a sum of Lorentzian peaks is reduced to the standard linear algebra problem of finding the eigenvalues and eigenvectors of the propagator  $\hat{U} = e^{-i\tau}$ .<sup>[5]</sup>

The key insight of filter diagonalization is that the propagator to be diagonalized,  $e^{-i\tau}$ , may be expressed entirely in terms of time samples of the signal  $f_n \equiv f(t_n)$ . A common approach is to write the propagator in a so-called Krylov basis,

$$\Psi_k = \sum_{n=0}^N \left( \frac{\hat{U}}{z_k} \right)^n \Phi_0, \quad (15)$$

where  $z_k = e^{-i\nu_k \tau}$  is a complex value chosen along the unit circle.<sup>[2]</sup> By selecting the  $\nu_k$  close to the frequencies we wish to resolve, it is possible to filter  $f(t)$  and recover only those frequency components near the  $\nu_k$ ; this is where the name filter

diagonalization comes from. It is important to include more basis vectors  $|\Psi_k\rangle$  than there are frequencies we wish to resolve.<sup>[2]</sup>

Expressing the propagator in the Krylov basis yields

$$U_{kk'} = (\Psi_k, \hat{U}\Psi_{k'}) = \sum_{n=0}^N \sum_{n'=0}^N f_{n+n'+1} z_k^{-n} z_{k'}^{-n'}, \quad (16)$$

which is expressed completely in terms of time samples of the signal. Because the Krylov basis is not orthonormal, we also need the overlap matrix

$$S_{kk'} = (\Psi_k, \Psi_{k'}) = \sum_{n=0}^N \sum_{n'=0}^N f_{n+n'+1} z_k^{-n} z_{k'}^{-n'}, \quad (17)$$

after which the eigenvalues and eigenvectors of  $\hat{U}$  may be found by solving the generalized eigenvalue problem,

$$UB_j = u_j SB_j. \quad (18)$$

For computational efficiency, the double sums in Eqs. (16) and (17) are rewritten as single sums, as shown in Ref. 4. With the eigenvalues  $u_j$  and eigenvectors

$$u_j = \sum_k B_{kj} \Psi_k \quad (19)$$

in hand, the frequencies  $\omega_j$  and amplitudes  $\lambda_j$  in the signal  $f(t)$  may be reconstructed according to the formulas

$$u_j = e^{-i\omega_j \tau}, \text{ and} \quad (20)$$

$$\lambda_j = |(u_j, \Phi_0)|^2 = \left| \sum_k B_{kj} (\Psi_k, \Phi_0) \right|^2. \quad (21)$$

Because the Krylov basis is often close to becoming linearly dependent, we include a numerical conditioning step to remove possible spurious frequencies. In particular, we select a value of  $p$  and resolve the generalized eigenvalue equation with  $\hat{U}^{p+1}$  and  $\hat{U}^p$  (used in place of  $\hat{U}$  and  $\hat{S}$ ). We remove the eigenvalues that are not shared in the two spectra, and then select a filtering grid with frequencies  $\nu_j$  located only at the nonspurious eigenvalues. We rerun filter diagonalization one more time on this adaptive frequency grid,<sup>[7]</sup> and these are the results we report below.

While this technique was initially derived with quantum mechanics in mind, it is not limited to such applications. Indeed, with generalizations such as 2D filter diagonalization<sup>[4]</sup> and multi-resolution filter diagonalization,<sup>[38]</sup> the method has been expanded to be applicable to a broad range of signals.

## Methods

As our goal in this article is to compare the performance of compressed sensing, super-resolution, and filter diagonalization in recovering sparse signals, we began by obtaining a series of sparse signals from the Sparco toolkit, which is a well-known set of sparse signals used for benchmarking various signal process-

ing techniques.<sup>[20]</sup> We also generated a few other signals of interest to highlight particular properties of each technique. Unless otherwise stated, each signal began as a continuous function of time  $f(t)$  and, to generate a discrete time series, we sampled  $f(t)$  at 4096 time points ranging uniformly from  $t = 0$  to  $t = 1$  second. This gave a grid separation of  $1/4096$  s, with a maximum recoverable frequency of 2048 Hz.

For each sparse signal processing method, we varied how many of the 4096 time points we sampled (in increments of 64) and investigated the dependence of the recovery error on the extent of undersampling. As a measure of the recovery error, we employed the relative 2-norm error over all 4096 time points (regardless of the extent of undersampling):

$$\text{Recovery Error} = \frac{\sum_{i=1}^{4096} |f_{\text{recovered}}(t_i) - f_{\text{original}}(t_i)|^2}{\sum_{i=1}^{4096} |f_{\text{original}}(t_i)|^2} \quad (22)$$

We consistently obtained similar results with the 1-norm error and the  $\infty$ -norm error, but the 2-norm error has the advantage that, by Parseval's theorem, it is the same whether it is measured in the time domain or the frequency domain. Therefore, we adopted the 2-norm as our primary benchmark. In some instances, filter diagonalization has been marketed as a parameter estimation technique, but this problem is identical to reconstruction of the signal as a whole. That is, if a method can extract the characteristic parameters of the signal, then these parameters can be used to reconstruct the signal.

For compressed sensing and super-resolution, we attempted to recover each signal in an appropriate sparse basis. The basis used depends on the signal and is discussed in the individual sections below. Because super-resolution requires a grid of equally spaced sample points, we began our analysis by examining the full signal and computing the errors. We then repeated our analysis for the signal by successively undersampling in powers of two, taking care to ensure that our sample points were always equally spaced. In contrast, our analysis with compressed sensing involved randomly selecting the same number of points that were included in the super-resolution analysis at each step.

For filter diagonalization, we used the same regular sampling strategy as for super-resolution, and we monitored the recovery error as a function of the sampling. Filter diagonalization requires specifying a grid of frequencies on which we expect the components of the signal to lie, so we specified a frequency range of 0 kHz to 20 kHz. To find the appropriate grid density and number of frequencies, we tuned these two parameters for optimal reconstruction with the full time signal and assumed these parameters would be valid for the entire numerical experiment. To ensure the robustness of our results, we also performed the analysis from 0 to 5, 10, 50, 100, 200, and 500 kHz, with a similar density of frequencies.

From a numerics standpoint, compressed sensing and super-resolution require a fast, memory-efficient  $\mathcal{L}_1$  solver. For all results in this article, we implemented the two-step iterative

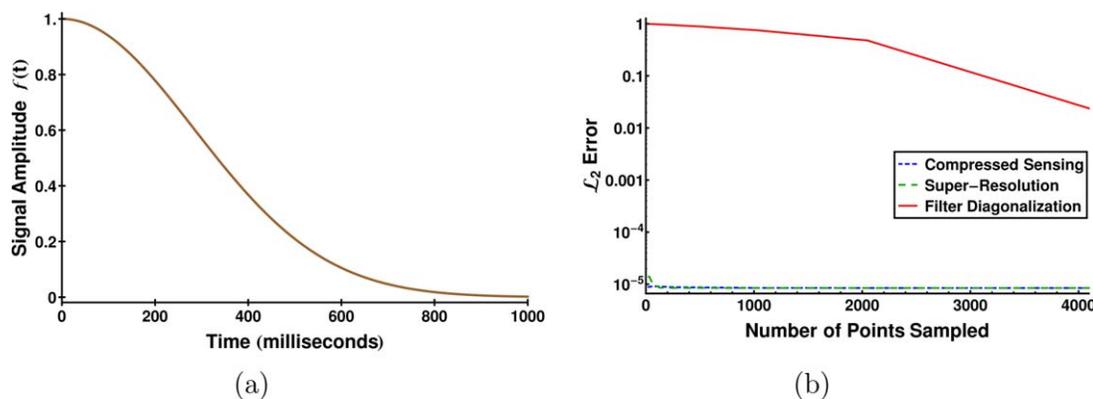


Figure 1. a) Time series consisting of a Gaussian given by Eq. (23) centered at  $t = 0$  with standard deviation  $\sigma = 0.4$ . b) Comparison of the relative 2-norm error in the reproduction of a Gaussian signal as a function of undersampling between compressed sensing, super-resolution, and filter diagonalization.

shrinkage/thresholding (TwIST) algorithm in Python.<sup>[39,40]</sup> Our TwIST solver is capable of solving arbitrary optimization problems given a measurement matrix, signal vector, and objective function, and gives numerically identical answers to the Matlab version for a wide range of test signals, including all those in this article.

We implemented filter diagonalization in Python, performing all required matrix diagonalizations using the `zgeev` function from LAPACK. We benchmarked our implementation against Harminv, a freely-available C++ implementation based on the methods described in Ref. [7], and found that they give the same answers to within numerical precision for a wide range of signals, including all of those presented in this article.

### Gaussian signal

We begin with one of the most ubiquitous signals throughout signal processing, a simple Gaussian centered at time  $t = 0$  with  $\sigma = 0.4$  (Fig. 1a),

$$f(t) = e^{-\frac{t^2}{0.4^2}}. \quad (23)$$

To recover this signal with compressed sensing and super-resolution, we employ a basis of displaced Gaussians

$$g_{jk}(t) = e^{-\frac{(t-t_j)^2}{\sigma_k^2}}, \quad (24)$$

with 100 centers  $t_j$  ranging uniformly from 0 to 1, and 100 standard deviations  $\sigma_k$  also ranging uniformly from 0 to 1, for a total of 10,000 different basis functions. It is clear that the function we hope to recover,  $f(t)$ , is sparse in this basis.

Figure 1b compares the performance of compressed sensing, super-resolution, and filter diagonalization in recovering the Gaussian signal. Compressed sensing and super-resolution both converge quickly to the correct signal, and as more time-domain information is sampled, the signal becomes more obviously composed of a single Gaussian. Moreover, compressed sensing converges more quickly than super-resolution, indicating that randomly sampling over the entire time domain provides more complete information about the overall shape of the signal than sampling uniformly with a coarse grid. Both compressed sensing

and super-resolution recover a single strongly converged, correct, peak with amplitude 1 and a few spurious peaks with amplitudes smaller than  $10^{-6}$ . This represents a small numerical instability in our implementation of TwIST but these spurious features are easy to identify and disregard.

By contrast, filter diagonalization fails to converge completely because it attempts to recover the Gaussian as a sum of Lorentzian peaks, rather than taking advantage of the natural sparsity of the signal in a Gaussian basis. This example highlights the basis set agnosticism of the  $\mathcal{L}_1$  minimization techniques, which is one of their principal advantages.

### Sparco problem 1

For our second signal, we consider a sinusoid that is “disrupted” by two Heaviside step functions (Fig. 2a),

$$f(t) = 4\sin(4\pi t) - \Theta(t - 0.3) - \Theta(0.72 - t), \quad (25)$$

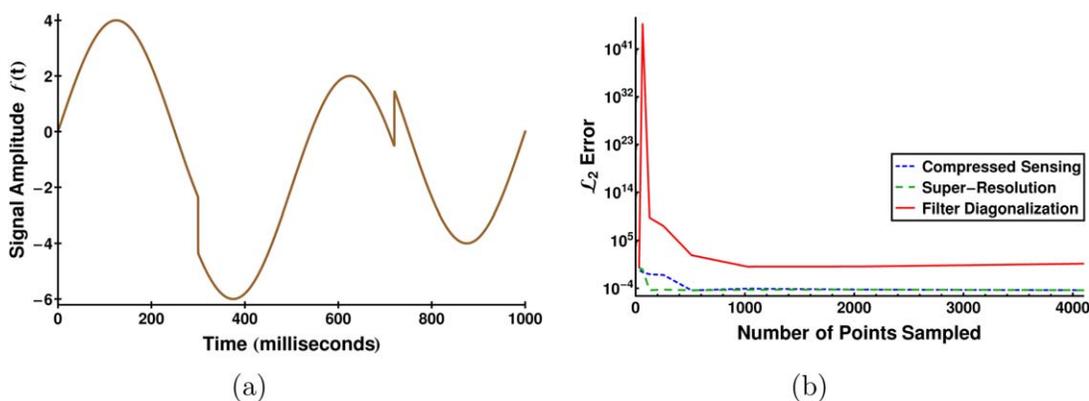
one of the earliest signals used to benchmark wavelet and compressed sensing techniques.<sup>[8,41,42]</sup> This signal is problem 1 in the Sparco toolbox of sparse signals.<sup>[20]</sup>

To recover this signal via compressed sensing and super-resolution, we employ a composite basis of sine functions and Heaviside step functions

$$g_j(t) = \sin(\omega_j t)$$

$$h_j(t) = \Theta(t - t_j),$$

with the spectral spacing  $\omega_j$  of the sine functions ranging uniformly from 0 to  $4096\pi$  kHz in units of  $\frac{\pi}{10}$  kHz, and the unit steps  $t_j$  of the Heaviside step functions ranging uniformly from 0 to 1 s in units of 0.01 s. While either basis  $g_j(t)$  or  $h_j(t)$  by itself would provide a complete basis for recovery of the signal  $f(t)$  (to within numerical precision), the function  $f(t)$  would not be sparse in either basis on its own. On the other hand, there is no problem in  $\mathcal{L}_1$  minimization techniques with using the combined basis, which affords the additional advantage that  $f(t)$  is sparse in this combined basis. However, when building a composite basis with different functional forms, it is important



**Figure 2.** a) The combination of a sinusoid and a Heaviside signal in the time domain as given by Eq. (25). b) Comparison of the 2-norm error in the reproduction of problem 1 from the Sparco toolbox as a function of undersampling between compressed sensing, super-resolution, and filter diagonalization. We attribute the large spike in error by filter diagonalization to the recovery of spurious exponentially divergent solutions.

to ensure that each basis function is normalized to the same value, for which we chose unity.

Figure 2b compares the performance of compressed sensing, super-resolution, and filter diagonalization in recovering the signal  $f(t)$ . For super-resolution and filter diagonalization, both of which involve regular sampling over a coarse time domain grid, most of the initial error simply comes from the fact that the methods cannot reproduce aspects of the signal that have not been sufficiently sampled.

Super-resolution and compressed sensing are able to identify that the signal has some underlying sine structure, but initially fails to recognize the exact position of the step functions. As more samples are included in the analysis, both techniques are able to quickly converge to the exact location of the step function. This convergence leads to a very sharp phase transition characteristic of an  $\mathcal{L}_1$  analysis, and represents the minimum amount of information required to exactly reproduce the full signal. This phase transition is a well known aspect of  $\mathcal{L}_1$  minimization techniques, and provides a useful and valid check on convergence and accuracy.

This stands in contrast to filter diagonalization, which attempts to match the Heaviside step functions by creating a signal that contains exponentially growing components, eventually resulting in an explosion of error. To better understand this behavior, we varied the magnitude of the Heaviside step functions, but found that the creation of an exponentially growing signal persisted even when the Heaviside step function was 0.1% of the amplitude of the oscillating sine wave.

Not surprisingly, compressed sensing fares better than both super-resolution and filter diagonalization. This is easily explained by the fact that compressed sensing randomly samples the entire domain, so it can quickly “recognize” all features of the signal and recover them accurately. After roughly one sixty-fourth of the signal has been sampled, the error changes only marginally, and this effect is robust across different runs of random sampling.

### Sparco problem 5

For our third signal, we consider the sum of three cosines with the addition of 40 spikes at random time points  $\{t_i\}$  (Fig. 3a),

$$f(t) = 2\cos(2\pi t) + 3\cos(9\pi t) - \cos(20\pi t) + \sum_i \alpha_i \delta(t - t_i), \quad (26)$$

where  $\delta(t - t_i)$  is regarded here as the Kronecker delta function (equal to 1 at the time point  $t = t_i$ , 0 otherwise) and  $\alpha_i$  is a uniform random number between 0 and 1. This signal is problem 5 in the Sparco toolbox of sparse signals,<sup>[20]</sup> and including it in our comparison is particularly useful for benchmarking the ability of compressed sensing, super-resolution, and filter diagonalization to deal with random noise.

To recover this signal via compressed sensing and super-resolution, we employ a cosine basis:

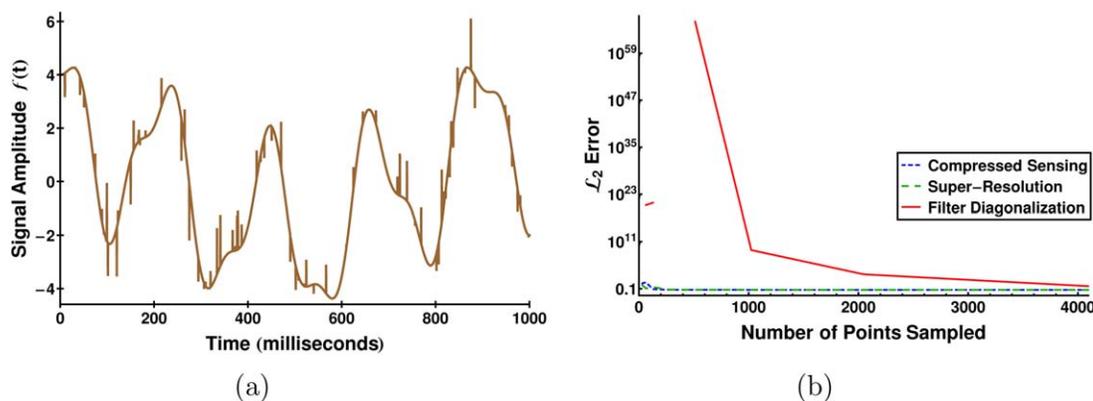
$$g_j(t) = \cos(\omega_j t)$$

with the spectral spacing  $\omega_j$  of the cosine functions ranging uniformly from 0 to  $4096\pi$  kHz in units of  $\frac{\pi}{10}$  kHz. Note that we do not include Kronecker delta functions  $\delta(t - t_i)$  in our basis, since our goal is to see whether our signal processing methods can recover the underlying cosine functions despite the random noise.

As shown in Figure 3b, both compressed sensing and super-resolution successfully recover the underlying cosine functions in spite of the noise peaks (the noise peaks simply get absorbed into the denoising parameter  $\eta$ ). As expected, compressed sensing recovers the signal with less sampling than super-resolution, since randomly sampled points over the entire time domain effectively contribute more information than regularly sampled points over that same time domain.

In contrast, filter diagonalization does not include a robust denoising procedure, and the method struggles with the Kronecker delta peaks because they represent sharp deviations from the underlying cosine signal. In particular, in attempting to match the Kronecker delta peaks, filter diagonalization creates a signal that contains exponentially growing components rather than exponentially damped Lorentzians.

In summary, compressed sensing and super-resolution both pick out the underlying cosine signals by denoising the Kronecker delta peaks, whereas filter diagonalization does not.



**Figure 3.** a) The combination of three cosines and random spikes in the time domain as described in Eq. (26). b) Comparison of the 2-norm recovery error for the problem 5 from the Sparco toolbox as a function of undersampling for compressed sensing, super-resolution, and filter diagonalization. We attribute the large error peaks found by filter diagonalization to the recovery of spurious exponentially divergent solutions.

### Jacob's Ladder

Next, we consider a time series devised by some of the original developers of filter diagonalization for benchmarking sparse signal processing techniques.<sup>[2]</sup> This signal is known as Jacob's Ladder, and it consists of a very large number of Lorentzian peaks (Fig. 4a):

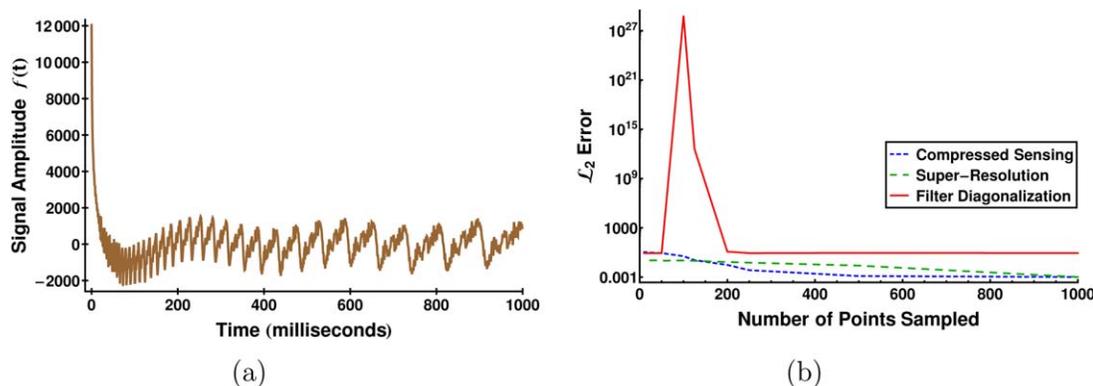
$$f_n = \sum_{m=0}^{49} e^{-1.8m\pi \times 10^{-4} \cdot n} [\cos(1.8m\pi \times 10^{-4} \cdot 2500n) + \cos(1.8m\pi \times 10^{-4} \cdot 2487.5n) + \cos(1.8m\pi \times 10^{-4} \cdot 2475.0n)] \quad (27)$$

In this experiment, we created a signal with 1000 data points that ranged, uniformly, from 0 to 1 s using the formula in Eq. (27). We first analyzed this signal with a numerical implementation of filter diagonalization using a frequency range of 0 kHz to  $\pi$  kHz and assumed that we had a maximum of 4500 frequencies in our signal.

We started by performing the filter diagonalization analysis with the above frequency grid using the entire signal, and used the recovered frequencies and expansion coefficients to

construct the recovered signal. This allowed us to compute the 2-norm error between the recovered signal and the exact signal over the entire range. We then repeated the analysis, while successively undersampling, first taking every second point, then every fourth, fifth, eighth, tenth, twentieth, twenty-fifth, fortieth, fiftieth, and finally, hundredth. We then performed the same analysis with super-resolution and compressed sensing. Because compressed sensing involved random sampling, we took care to randomly sample the same number of points that were included in the super-resolution analysis at each step.

In order to perform the  $\mathcal{L}_1$  analysis we examined the functional form and the signal itself and concluded that it should be sparse in the basis of damped oscillators. Thus we constructed such a basis with a spectral spacing of 0.01 Hz and a maximum of  $4\pi$  Hz. We also scanned exponential decay parameters ranging from 0 to  $\pi$  Hz in steps of 0.01 Hz. For the super resolution analysis, we performed the same time addition procedure that was performed with filter diagonalization. We began the compressed sensing analysis by picking 50 random time points, and at each subsequent step an additional



**Figure 4.** a) The Jacob's Ladder signal, a combination of many Lorentzian peaks, in the time domain as given by Eq. (27). b) Comparison of the 2-norm error for the reproduction of the Jacob's Ladder signal as a function of undersampling for compressed sensing, super-resolution, and filter diagonalization. We have no explanation clear for why our implementation of filter diagonalization fails to reproduce this signal, and we note that further work may be necessary.

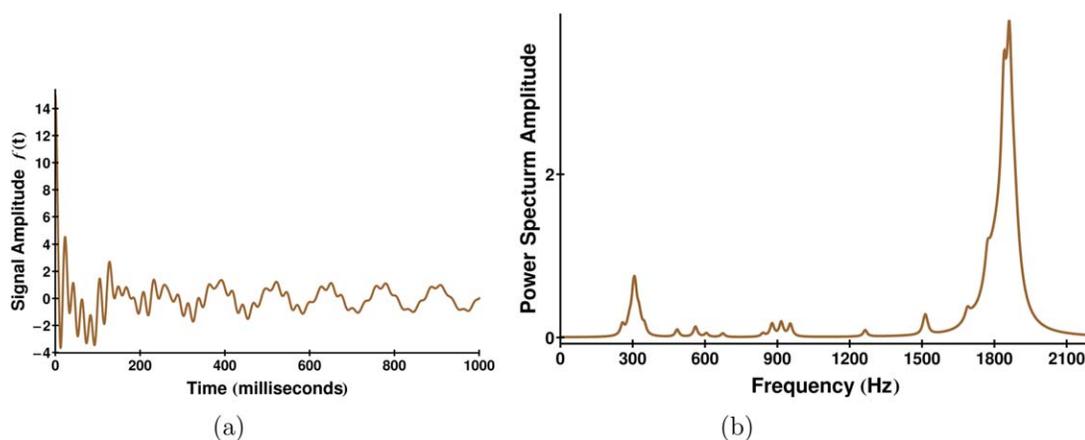


Figure 5. a) Combination of twenty random damped cosines as described in Eq. (28). b) The power spectrum for this particular random sum of damped cosines.

50 random points were taken from those remaining until we were sampling the full signal.

From the errors given in Figure 4b, both compressed sensing and super resolution converge to a better answer more rapidly than filter diagonalization. We attribute these errors to the recovery of exponentially divergent solutions but we have no way of accounting for the difference between our results and those obtained elsewhere. Our working theory is a sensitivity of filter diagonalization to the frequency grid and parameter choice, and we have not found the correct combination of parameters that allows us to completely recover the desired signal. This suggests that our implementation, as well as our standard of comparison suffered from significant numerical instability. We are unsure whether this is explained better by fundamental instabilities in the method at hand, or simply instabilities in the current implementations.

By contrast, super-resolution and compressed sensing do not suffer from these same problems. Many current methods are extremely stable. These techniques only include the basis functions that are explicitly chosen. Unfortunately, we are lim-

ited to the recovery of Lorentzian parameters (frequency and line-width), that are on the grid, which requires a sufficiently dense set of parameters for accurate recovery. As a result,  $\mathcal{L}^1$ -optimization methods become more and more memory intensive as the number of basis functions increases.

#### Sum of random Lorentzians

For the final comparison, we created a sum of twenty random damped cosines:

$$f(t) = \sum_{n=1}^{20} e^{-\gamma_n t} \cos \omega_n t \quad (28)$$

with  $\gamma_n$  drawn from a uniform random distribution ranging from 0 to 20 Hz and  $\omega_n$  ranging from 0 to  $50\pi$  Hz. This type of autocorrelation signal is ubiquitous not only in chemistry applications, but in signal processing at large. Extracting parameters from the power spectrum of an autocorrelation function is one of the most powerful tools in the whole of

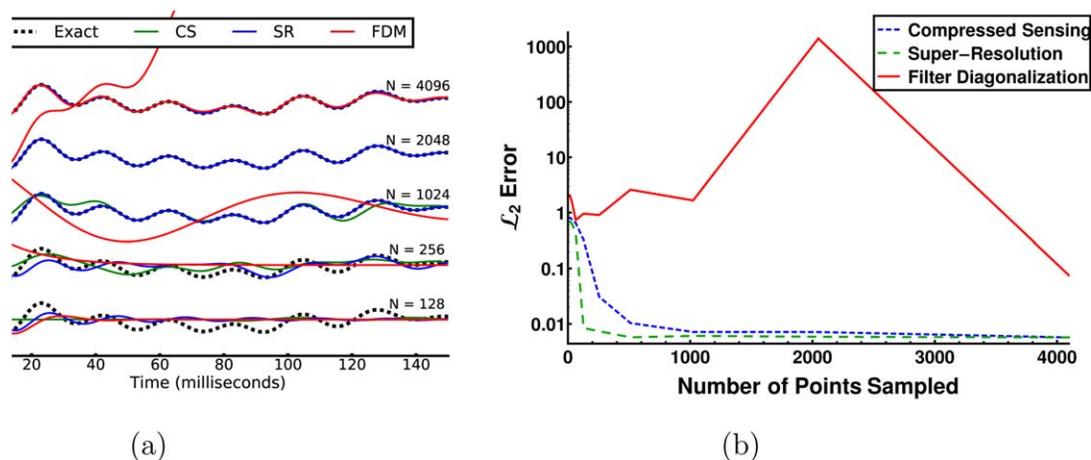


Figure 6. a) Reconstructed signals for compressed sensing, super-resolution, and filter diagonalization as a function of the number of points included in the reconstruction algorithm. The full signal has 4096 points. b) Comparison among compressed sensing, super-resolution, and filter diagonalization of the 2-norm error in the reproduction of the combination of random damped cosines as a function of undersampling. We attribute the large error peaks found by filter diagonalization to the recovery of spurious exponentially divergent solutions.

chemistry. While no explicit chemical application is present in this manuscript, we have included the power spectrum of the sum of random Lorentzians to emphasize the importance of this problem. Indeed, the power spectrum plotted in Figure 5b is similar to chemically relevant spectra, such as those obtained from NMR experiments.

To recover this signal via compressed sensing and super-resolution, we employ a basis of damped cosine functions:

$$g_{jk}(t) = e^{-\gamma_k t} \cos(\omega_j t).$$

Here, the spectral spacing  $\omega_j$  ranges uniformly from 0 to  $40\pi$  Hz in steps of  $\pi/24$  Hz and the damping parameters  $\gamma_k$  ranges uniformly from 0 to 20 Hz in steps of 1/2 Hz. For filter diagonalization, we selected a frequency range from 0 to  $100\pi$  Hz and chose a basis of 1200 frequencies. This was chosen because it gave near perfect reconstruction of the full signal, while bases smaller than this were prone to numerical instability.

In Figure 6a we present the signal that was recovered by each method when restricted only to the number of points indicated. With only 128 included points, Super-resolution and filter diagonalization are both capable of recovering the general shape of the signal, but miss many of the key features, including the dominant frequencies. Compressed sensing isn't able to recover a sensible signal and thus, returns a vector of zeros. As more points are added Super-resolution recovers a more faithful representation of the signal. In fact, the visual difference between the signal recovered with 1024 and 4096 points is negligible. While compressed sensing systematically improves the recovered signal as more points are added, it requires 2048 points to recover a signal that is visually indistinguishable from the original. In contrast, filter diagonalization does not display the characteristics of systematic improvement. In fact, with 2048 included points filter diagonalization returns a recovered signal with exponentially growing components. With the exception of the recovery with the full signal, even when filter diagonalization returns bounded signals, it fails to recover a signal that is a faithful representation of the original. While filter diagonalization ends with a better representation of the signal, it is unclear how to develop diagnostics for systematic improvement.

As shown in Figure 6b, both compressed sensing and super-resolution successfully recover the underlying damped cosine structure but is restricted to the functions on the grid. This is the most significant source of error. As expected, compressed sensing recovers the signal with less sampling than super-resolution, since randomly sampled points over the entire time domain effectively contribute more information than regularly sampled points over a short time.

In contrast, filter diagonalization can recover off-grid frequencies extremely efficiently. Because of this, the final errors should be smaller than the errors from both compressed sensing and super resolution. Unfortunately, we encountered significant stability issues during many of our decompositions that resulted in exponentially growing solutions. These solutions gave  $\mathcal{L}_2$  errors on the order of  $10^4$  at times (note that

we have chosen to only plot a few orders of magnitude in Fig. 6b) While filter diagonalization is capable of giving a much better answer, the technique is significantly more sensitive to slight deviations in the operational parameters chosen.

## Conclusions

In conclusion, we have performed a broad comparison of three different signal processing techniques that attempt to "beat" the Shannon–Nyquist limit. With prior information about a reasonable basis for your signal,  $\mathcal{L}_1$  minimization techniques provide a robust and faithful reproduction of the signal. We emphasize that the difference between super-resolution or compressed sensing is simply a choice of sampling procedure and normally is determined by the data acquisition technique.

Additionally, we found that if the signal at hand was Lorentzian, filter diagonalization was capable of significantly outperforming both compressed sensing and super resolution because of its ability to sample off the grid. Even still, the technique was sensitive to a broad range of parameters, which were capable of making it divergent if chosen incorrectly. Given enough tuning and the appropriate signal form, however, filter diagonalization is the superior method for these types of signals.

## Acknowledgment

Authors acknowledge Professor Vladimir Mandelshtam for useful conversations.

**Keywords:** compressed sensing · super resolution · filter diagonalization

How to cite this article: T. Markovich, S. M. Blau, J. N. Sanders, A. Aspuru-Guzik. *Int. J. Quantum Chem.* **2016**, *116*, 1097–1106. DOI: 10.1002/qua.25144

- [1] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed.; Academic Press, Amsterdam, **2008**.
- [2] H. Hu, Q. N. Van, V. A. Mandelshtam, A. J. Shaka, *J Magn. Reson.* **1998**, *134*, 76.
- [3] C. Qiu, W. Lu, N. Vaswani, in ICASSP 2009, Taipei, Apr **2009**, p. 393.
- [4] R. Chen, H. Guo, *J. Chem. Phys.* **1999**, *111*, 464.
- [5] M. R. Wall, D. Neuhauser, *J. Chem. Phys.* **1995**, *102*, 8011.
- [6] V. A. Mandelshtam, *Prog. Nucl. Magn. Reson. Spectrosc.* **2001**, *38*, 159.
- [7] V. A. Mandelshtam, H. S. Taylor, *J. Chem. Phys.* **1997**, *107*, 6756.
- [8] D. L. Donoho, *IEEE Trans. Inform. Theory* **2006**, *52*, 1289.
- [9] J. Ma, *IEEE Trans. Instrum. Meas.* **2010**, *59*, 1600.
- [10] A. Oka, L. Lampe, in ICUWB 2009, Vancouver, Sep **2009**, p. 279.
- [11] J. N. Sanders, S. K. Saikin, S. Mostame, X. Andrade, J. R. Widom, A. H. Marcus, A. Aspuru-Guzik, *J. Phys. Chem. Lett.* **2012**, *3*, 2697.
- [12] M. A. Herman, T. Strohmmer, *IEEE Trans. Signal Process.* **2009**, *57*, 2275.
- [13] V. Kekatos, G. B. Giannakis, *IEEE Trans. Signal Process.* **2011**, *59*, 5907.
- [14] X. Andrade, J. N. Sanders, A. Aspuru-Guzik, *Proc. Natl. Acad. Sci. USA* **2012**, *109*, 13928.
- [15] D. Kouame, M. Ploquin, in ISBI '09, Boston, Jun **2009**, p. 249.
- [16] C. W. McCutchen, *J. Opt. Soc. Am.* **1967**, *57*, 1190.
- [17] K. G. Puschmann, F. Kneer, *Astron. Astrophys.* **2005**, *436*, 373.
- [18] W. T. Freeman, T. R. Jones, E. C. Pasztor, *IEEE Comput. Graph. Appl.* **2002**, *22*, 56.

- [19] E. J. Candès, C. Fernandez-Granda, *Comm. Pure Appl. Math.* **2014**, *67*, 906.
- [20] E. Van Den Berg, M. P. Friedlander, G. Hennenfent, F. Herrmann, R. Saab, O. Yilmaz, Sparco: A Testing Framework for Sparse Reconstruction; Dept Comput Sci, Univ British Columbia: Vancouver, **2007**. Tech Rep TR-2007-20.
- [21] R. G. Baraniuk, E. Candès, M. Elad, Y. Ma, *Proc. IEEE* **2010**, *98*, 906.
- [22] M. Lustig, D. Donoho, J. M. Pauly, *Magn. Reson. Med.* **2007**, *58*, 1182.
- [23] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, A. G. White, *Phys. Rev. Lett.* **2011**, *106*, 100401.
- [24] M. Nagahara, D. E. Quevedo, J. Ostergaard, In CDC 2012, Grand Wailea, Maui, Hawaii, **2012**, p. 1362.
- [25] T. Tuma, S. Rooney, P. Hurley, In ICECCS 2009, IEEE, Paris, Jun **2009**, p. 210.
- [26] M. Mishali, Y. C. Eldar, *IEEE J. Sel. Top. Signal Process.* **2010**, *4*, 375.
- [27] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, R. G. Baraniuk, *IEEE Signal Process. Mag.* **2008**, *25*, 83.
- [28] W. K. Coulter, C. J. Hillar, G. Isley, F. T. Sommer, In ICASSP 2009, Taipei, Mar **2010**, p. 5494.
- [29] A. J. Patti, M. I. Sezan, A. Murat Tekalp, *IEEE Trans. Image Process.* **1997**, *6*, 1064.
- [30] M. Elad, A. Feuer, *IEEE Trans. Image Process* **1997**, *6*, 1646.
- [31] T. Markovich, S. M. Blau, J. Parkhill, C. Kreisbeck, J. N. Sanders, X. Andrade, A. Aspuru-Guzik, arXiv, Jul **2013**.
- [32] E. J. Candès, J. Romberg, T. Tao, *IEEE Trans. Inform. Theory* **2006**, *52*, 489.
- [33] G. Plonka, J. Ma, *Int J Wavelets Multiresolution* **2011**, *9*, 79.
- [34] M. F. Duarte, M. B. Wakin, R. G. Baraniuk, In ICASSP 2009, Taipei, Mar **2008**, p. 5137.
- [35] A. B. Lee, B. Nadler, L. Wasserman, *Ann. Appl. Stat.* **2008**, *2*, 435.
- [36] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, S. W. Zucker, *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 7432.
- [37] X. Li, R. Zhao, S. Hu, In WiCOM 2010, Chengdu, Sep **2010**, p. 1.
- [38] K. Aizikov, P. B. O'Conner, *J. Am. Soc. Mass Spectrom.* **2006**, *17*, 836.
- [39] J. M. Bioucas-Dias, M. A. Figueiredo, *IEEE Trans. Image Process.* **2007**, *16*, 2992.
- [40] J. M. Bioucas-Dias, M. A. T. Figueiredo, In ICIP 2007, **2007**, p. I-105.
- [41] J. B. Buckheit, D. L. Donoho, *Wavelets and Statistics*; Springer: New York, **1995**.
- [42] S. S. Chen, D. L. Donoho, M. A. Saunders, *SIAM Rev* **2001**, *43*, 129.

---

Received: 9 February 2016

Revised: 11 March 2016

Accepted: 22 March 2016

Published online 15 April 2016