SPARSE RECOVERY OF COMPLEX PHASE-ENCODED VELOCITY IMAGES USING ITERATIVE THRESHOLDING

Tim Roberts[*]*, Nick Kingsbury*

Signal Processing & Communication Lab
Dept. of Engineering
University of Cambridge, U.K.
Email: {tr331,ngk10}@cam.ac.uk

Daniel J Holland[†]

Magnetic Resonance Research Centre
Dept. of Chem. Engineering & Biotechnology
University of Cambridge, U.K.
Email: djh79@cam.ac.uk

ABSTRACT

In this paper we propose a new algorithm for reconstructing phase-encoded velocity images of catalytic reactors from undersampled NMR acquisitions. Previous work on this application has employed total variation and nonlinear conjugate gradients which, although promising, yields unsatisfactory, unphysical visual results. Our approach leverages prior knowledge about the piecewise-smoothness of the phase map and physical constraints imposed by the system under study. We show how iteratively regularizing the real and imaginary parts of the acquired complex image separately in a shift-invariant wavelet domain works to produce a piecewise-smooth velocity map, in general. Using appropriately defined metrics we demonstrate higher fidelity to the ground truth and physical system constraints than previous methods for this specific application.

Index Terms— magnetic resonance, velocity imaging, compressed sensing, iterative algorithms, sparsity.

1. INTRODUCTION

Magnetic resonance imaging (MRI) has been used effectively to non-invasively image the distribution of velocity in liquid and gas flows. While most MRI applications seek to accurately recover the voxel magnitudes, the quantity of interest here is the velocity of a fluid, which is proportional to the phase of the complex image. In this application, we investigate the accurate recovery of phase information from subsampled MRI measurements, used in the characterization and prediction of flow behavior in packed bed catalytic reactors. Subsampling enables one to reduce acquisition times and thus study detailed features, such as fine vortices. Recovery of phase is also useful in cardiovascular MRI [1], rheology, and the study of porous media and granular flows [2].

Holland *et al* [3] have demonstrated the accurate recovery of liquid-gas flow images in this application via total variation (TV) minimization using a nonlinear conjugate gradients (NCG) solver. This method, referred to as TV-NCG, was adapted from the CS MRI methods originally proposed by Lustig *et al* [4]. However, the TV regularization sometimes leads to unphysical blocking/staircasing artifacts [5], limiting our ability to quantify flow structures. Another technique developed by Zhao *et al* [6] regularizes the phase images (using preconditioned conjugate gradients) and magnitude images (using iterative soft thresholding) separately at each iteration. This method is well-adapted to producing a smooth phase-map while preserving edges. However it requires switching between different regularizers for the phase cost function in early iterations in order to converge to a desirable local minimum, with no theoretical guarantees.

We approach the problem of recovering phase maps with the same piecewise-smooth assumption employed in [3, 6], and develop an algorithm that incorporates the ideas in each. Like [3], we consider the recovery of phase by regularizing the real and imaginary parts separately, which is an approach justified in section 2.2. Unlike [3] however, we employ a smooth dual-tree wavelet-based regularization in order to avoid the blocking artifacts introduced by TV minimization. Also we adopt a reweighted least squares approach to encourage sparsity, together with a continuation strategy as in [6], in order to converge smoothly towards good sparse solutions, despite the non-convexity of the problem space.

With these requirements in mind, iterative soft thresholding in the wavelet domain is a natural approach. Wavelets can sparsely represent piecewise-smooth signals very well, and the rapid convergence of iterative sparse recovery via soft thresholding is well studied [7, 8]. Generally, subband-adaptive thresholding algorithms have produced promising results for image deconvolution [9, 10] and, more recently, for CS MRI [11]. Therefore, we propose utilizing the modified subband-adaptive iterative shrinkage/thresholding (MSIST) algorithm [12], which finds a sparse solution using a reweighted l_2 approximation of the l_0 norm in the wavelet domain. However, because our input signal is complex, rather than being purely real, we need to adapt the basic MSIST algorithm to handle complex, phase-encoded data, which is explained in section 2.2.

For the proposed algorithm the dual-tree complex wavelet transform (DT-CWT) is chosen as an overcomplete sparsifying basis, because of the following desirable properties which make it well-suited for this wavelet regularization application:

1. Shift-invariance,
2. limited redundancy ($4:1$ for 2-D images), and
3. an efficient filter-bank implementation.

Shift invariance ensures that the effects of wavelet domain thresholding remain spatially localized and consistent. The dual-tree approach is a more computationally efficient way to achieve this than other techniques (e.g. random shifting [9, 11]) for mitigating unwanted artifacts such as ringing.

Following the development of this algorithm, we show how the physical constraints for the system under observation are incorporated in the reconstruction in section 2.3.

[*]Supported by Cambridge Overseas Trust.
[†]Supported by EPSRC (Grants no. EP/F047991/1 and EP/K008218/1).

Fig. 1. k-space mask reduced vertically 4:1 to save space. The k-space origin is centered and the sampled phase-encode lines are shown in white.

2. METHODS

2.1. Compressed Sensing for Velocity Imaging

In this application, the velocity of a fluid, passing through a reactor tube filled with spherical catalyst beads, is acquired by applying two successive gradient pulses, using magnetic field gradients of $+g$ and $-g$, respectively, and separated by an observation time Δ. A thorough treatment of MRI phase-encoded velocity imaging can be found in [2].

After applying the appropriate imaging gradients, the acquired complex image f is modeled as

$$f = m \exp(-ix) \tag{1}$$

where $i = \sqrt{-1}$, the magnitude m is proportional to the signal strength, and the phase x encodes the velocity V according to

$$x = \gamma g \delta \Delta V \tag{2}$$

Here, γ is the gyromagnetic ratio, Δ is the time interval between the two gradient pulses, and δ is the time duration of each gradient pulse.

We model the observation of the k-space representation of f via the undersampled Fourier transform F_u:

$$y = F_u f + b \tag{3}$$

The term b represents the effects of noise or scanner imprecision. Because the acquisition time is proportional to the sampling fraction, it is desirable to increase the degree of undersampling in F_u in order to improve the spatio-temporal resolution of the imaging system.

Using the Monte-Carlo incoherent sampling design in [4], k-space patterns are simulated according to the polynomial decaying density in eq. (4), and the sampling pattern with the smallest peak interference is chosen.

$$p(k_r = 1) = (1 - r)^\rho \tag{4}$$

Here ρ is the decay factor, typically chosen to be $5/2$, r is the normalized distance from the k-space origin, and k_r takes on a binary value, or mask value, according to this distribution (see fig. 1). Fully-sampling near the k-space origin with random sampling according to eq. (4) gives an optimal trade-off between the fidelity to the observed signal and the necessary incoherence for artifact-free CS recovery from fully-random sampling [4, 13].

2.2. MSIST Algorithm for Sparse Recovery of Phase Images

In [12], the MSIST algorithm has been shown to achieve superior deconvolution results relative to other subband-adaptive algorithms [9, 14] for approximately piece-wise smooth natural images. In this application, we also assume that the velocity image is piecewise-smooth, with a sparse representation in the wavelet domain. By regularizing the real and imaginary parts of f, we can promote a piecewise-smooth distribution of phase values.

We justify regularizing the real and imaginary parts in order to obtain a piecewise-smooth phase x in eq. (1) as follows. If $f_{\mathrm{Re}} = m\cos(x)$ and $f_{\mathrm{Im}} = m\sin(x)$ are smooth, then f will be smooth, and so $x = \angle f$ will also be smooth as long as $m = |f|$ does not become small. Hence, if f_{Re} and f_{Im} are piecewise-smooth, then x will be piecewise-smooth as long as m is not small.

To avoid confusion because the input signal f and DT-CWT coefficients are both complex-valued, we summarize how the DT-CWT coefficients are stored and used. The two uses of complex values need to be kept separate since they are using *complex* to represent different things. The DT-CWT produces $2N$ j-indexed, complex-valued coefficients $\tilde{w}_j$ from a real image of size N pixels, via 4 parallel filter-banks [15]. Arranging the real and imaginary parts of $\tilde{w}$ into a real column vector w such that $w_{2j-1} = \mathrm{Re}(\tilde{w}_j)$ and $w_{2j} = \mathrm{Im}(\tilde{w}_j)\ \forall j$, we take W and W' to be the forward and inverse DT-CWT operators and these can now be purely real matrices of size $4N \times N$ and $N \times 4N$ respectively. This effectively converts the DT-CWT into a *purely real* overcomplete transform, which can then be applied separately to the real and imaginary parts of f to generate the complex $4N$-vector, $w_{\mathrm{Re}} + iw_{\mathrm{Im}} = Wf = Wf_{\mathrm{Re}} + iWf_{\mathrm{Im}}$; and similarly $f = W'w_{\mathrm{Re}} + iW'w_{\mathrm{Im}}$.

In order to find a sparse solution f^* which minimizes the discrepancy between the observed signal y and $F_u f$, we start with minimization of the following penalty function $J(f)$, representing $-2\nu^2$ log of the pdf, where ν^2 is the measurement noise variance:

$$J(f) = \|y - F_u f\|^2 - 2\nu^2 \log(p(f)) \tag{5}$$

Because f is complex, we must be careful how we define the sparsity-inducing term, $-2\nu^2 \log(p(f))$ in eq. (5). To achieve this, we assume that the real and imaginary parts of f are both sparsely represented in the wavelet-domain, each having a Gaussian distribution of coefficient magnitudes with zero mean and spatially adaptable variance, as in [10].

Following the same stabilized Gaussian scale mixture model and majorization minimization (MM) approach detailed in [10], eq. (5) can be expressed in terms of new variables w, S and z, where $w = w_{\mathrm{Re}} + iw_{\mathrm{Im}}$, S is the diagonal matrix of inverse variances in the model, and z represents the previous estimate of w for MM. Hence:

$$\begin{aligned} J(w, S, z)) =& \|y - F_u W' w\|^2 \\ &+ \nu^2 \left[w^H S w + \ln|S| + \epsilon^2 \mathrm{Tr}(S) \right] \\ &+ (w - z)^H (\Lambda_\alpha - W F_u' F_u W')(w - z) \end{aligned} \tag{6}$$

The MM expression on the third line of this equation is designed to cancel out the term $w^H W F_u' F_u W' w$ from expanding the first line. This leaves only linear terms in w plus a single quadratic term, $w^H(\Lambda_\alpha + \nu^2 S)w$, which avoids cross-coupling between the elements of w because Λ_α and S are diagonal and real. To achieve fast and stable convergence, Λ_α must be chosen in a subband-dependent way to be the minimum value which makes $(\Lambda_\alpha - W F_u' F_u W')$ positive definite.

Now $J(w, S, z)$ is minimized iteratively in the conventional way by taking derivatives with respect to w, z and S in turn and setting them to zero. Since w is complex, we take derivatives with respect to w_{Re} and w_{Im}, and set each to zero. This leads to the following

expressions for updating w_{Re} and w_{Im} at iteration $n+1$:

$$w_{\mathrm{Re},n+1} = (\Lambda_\alpha + \nu^2 S)^{-1}[W\,\mathrm{Re}\{F_u'(y - F_u f_n)\} + \Lambda_\alpha z_{\mathrm{Re},n}]$$
$$w_{\mathrm{Im},n+1} = (\Lambda_\alpha + \nu^2 S)^{-1}[W\,\mathrm{Im}\{F_u'(y - F_u f_n)\} + \Lambda_\alpha z_{\mathrm{Im},n}] \quad (7)$$

Because $(\Lambda_\alpha - WF_u'F_uW')$ is positive definite, $z_{n+1} = w_{n+1}$, as this minimizes the MM component of J, making it zero.

Then the diagonal elements of S are updated from w_{n+1} by minimizing the second line of eq. (6), subject to the tying constraint that $S_{2j} = S_{2j-1}$ so that the four components of each *doubly-complex* wavelet coefficient pair are drawn in a phase-unconstrained way from an isotropic 4-D Gaussian distribution. This gives, for all j, indexing elements $2j-1$ and $2j$ from $w_{\mathrm{Re,n+1}}$ and $w_{\mathrm{Im,n+1}}$:

$$S_{2j} = S_{2j-1} = \frac{1}{\frac{1}{4}(w_{2j-1,\mathrm{Re}}^2 + w_{2j-1,\mathrm{Im}}^2 + w_{2j,\mathrm{Re}}^2 + w_{2j,\mathrm{Im}}^2) + \epsilon^2} \quad (8)$$

Finally, to obtain the reconstructed complex image, we calculate

$$f_{n+1} = W' w_{\mathrm{Re},n+1} + iW' w_{\mathrm{Im},n+1} \quad (9)$$

These steps from eqs. (7) to (9) are then iterated to convergence.

Note that ϵ^2 in eqs. (6) and (8) is a stabilizing parameter which makes J locally convex, and is typically shrunk from a large to a small value as the minimization of J proceeds so as to encourage convergence to the globally minimal solution (or a good approximation to it). For fastest convergence, it was reported in [10, 16] that ν^2 should be reduced with ϵ^2 so that the ratio ν/ϵ remains approximately constant. Reducing ν and ϵ changes the shape of the sparse penalty to mimic an l_2 norm in earlier iterations, and an l_0 norm in later iterations. This slowly reduces the *soft* sparsity map of wavelet coefficients and avoids convergence to poor local minima. Following the suggestions in [10,16], a rapid reduction in early iterations of the initial choice of ν_0 and ϵ_0 and a slower reduction in later iterations to final values of ν_N and ϵ_N, ensures that the algorithm converges well.

Here we have chosen an exponentially decaying continuation rule for ν and ϵ which works well for this application:

$$\nu_n = \nu_0 \exp\left(-\frac{n}{K}\right) + \nu_N \quad (10)$$

and

$$\epsilon_n = \epsilon_0 \exp\left(-\frac{n}{K}\right) + \epsilon_N \quad (11)$$

where N is the total number of iterations (typically 20 to 50) and K is a decay constant.

2.3. Enforcing physical system constraints and initial solution

For the application at hand, we consider improving the reconstruction by taking into account the physical constraints of a system of fluid flowing through a tube packed with solid spheres, and, more generally, any laminar flow system where solid and fluid interfaces exist. From fluid dynamics, we expect the velocity of the fluid at the fluid/solid interface to be zero [17], barring limitations imposed by the resolution or imprecision of the acquisition. We refer to this as the *boundary layer constraint*.

We start with a logical mask M, which indicates the presence (1) or absence (0) of fluid. This mask is obtained by acquiring a fully-sampled k-space image in the desired position along the reactor

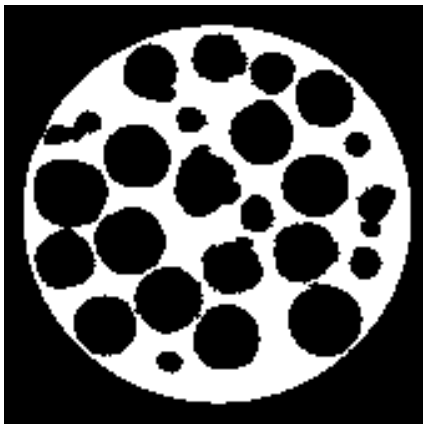

Fig. 2. Mask used to enforce physical system constraints.

cylinder, and determining the regions where there is negligible signal magnitude (absence of fluid) by using a simple thresholding rule. This threshold is set by choosing a value halfway between the peaks of a bimodal distribution of signal magnitudes, which correspond to the solid and fluid regions. In practice, this mask M is available for most experiments. It is derived from a separate, fully-sampled k-space acquisition, where the velocity distribution of the fluid is of no interest, and only the magnitude image is considered.

In order to provide ground truth for our results, we simulate the effect of k-space undersampling by applying F_u to a fully sampled k-space acquisition f to obtain y. We then obtain a zero-filled initial solution $f_0 = F_u'y$.

Before each iteration n of our reconstruction, we enforce zero phase on all of the non-fluid elements of f_n via $\bar{M}(\angle f_n) = 0$, but allow the magnitude to be unconstrained so it achieves smoothness with the surrounding fluid elements and consistency with the measured data, y. Here we have used $\bar{M}(f_n)$ to indicate selection of those elements of f_n corresponding to the non-fluid regions (black in fig. 2). In this way, we incorporate a spatial domain velocity prior into the wavelet regularization algorithm.

3. RESULTS

In our experiments, we have used the same raw data and k-space sampling technique as in [3]. We start with a fully-sampled MR acquisition of water flowing through a packed tube (27mm diameter) of borosilicate glass beads (5mm diameter), the phase image of which is shown in fig. 3a.

We quantify the performance of our algorithm using the same l_2 error metric defined in [3] for the ground truth x and the reconstructed phase x^*:

$$\mathrm{L2E} = \frac{\|M(x - x^*)\|_2}{\|M(x)\|_2} \quad (12)$$

Using a 30% sampling fraction for y, and setting $\nu_0 = \epsilon_0 = 5 \times 10^{-3}$, $\nu_N = \epsilon_N = 5 \times 10^{-4}$, $K = 1.5$, and $N = 50$, we have obtained an L2E of 7.4%. This compares favorably with the previous L2E of 11% reported in [3] for TV-NCG. Both our method and TV-NCG benefit from incorporating the mask prior shown in fig. 2. Without using this mask to enforce the zero phase constraint explained in sec. 2.3, our L2E worsens to 11.3%, with a similar effect for TV-NCG indicated in [3].

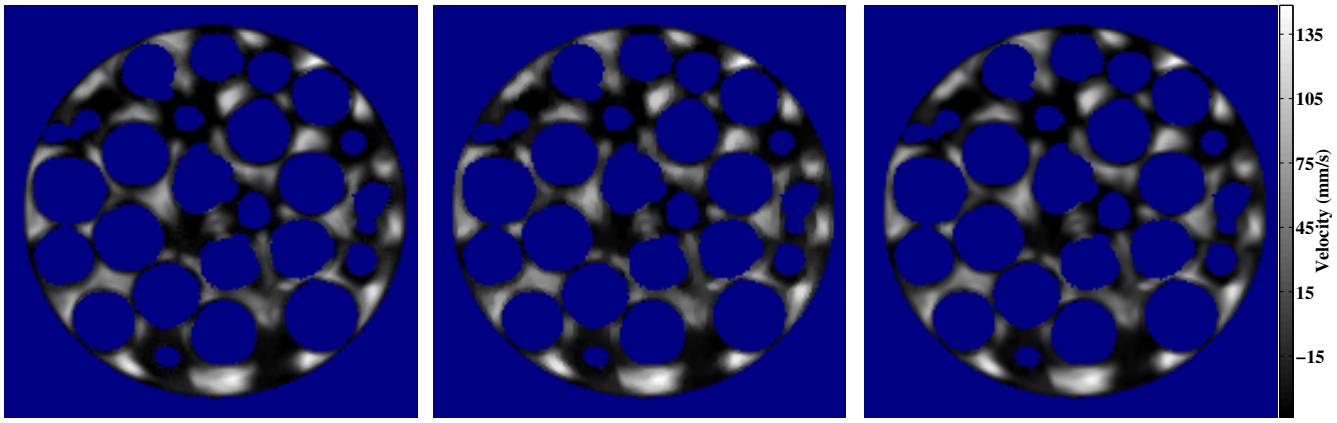

(a) Ground truth. (b) TV-NCG. (c) MSIST.

Fig. 3. Ground truth velocity image and reconstruction using TV-NCG and MSIST. Non-fluid regions are shown in blue.

The proposed method improves upon the visual quality as well, as shown in figs. 3 and 4. In fig. 3, typical blocking artifacts which arise from TV minimization are not present in the reconstructions using our proposed wavelet regularization algorithm. This difference more clearly illustrated by the line plots in figs. 4d and 4f corresponding to central horizontal rows in figs. 4c and 4e, respectively. Here the MSIST reconstruction exhibits smoother transitions than TV-NCG between low-flow and high-flow regions, and yet these transitions are still sharp, as in the ground truth data.

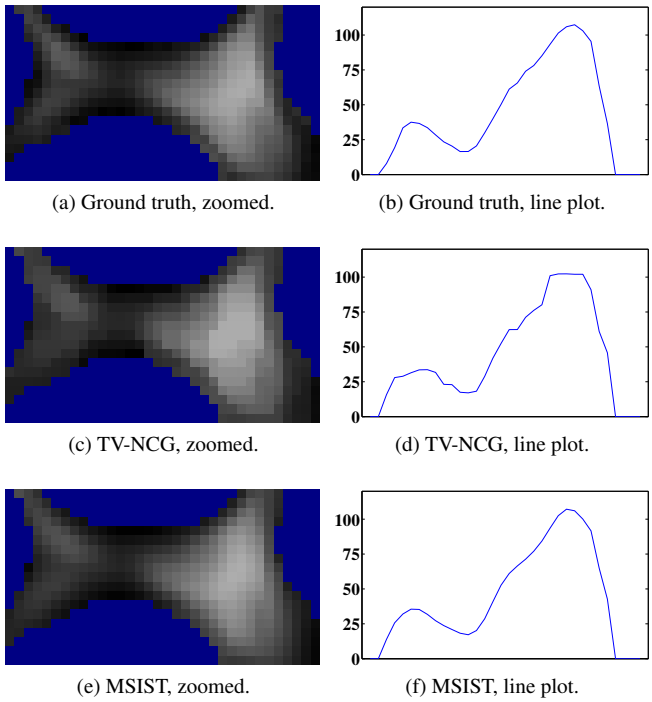

(a) Ground truth, zoomed. (b) Ground truth, line plot.

(c) TV-NCG, zoomed. (d) TV-NCG, line plot.

(e) MSIST, zoomed. (f) MSIST, line plot.

Fig. 4. Zoomed ground truth and reconstruction images with plots from the central image line, where velocity (mm/s) is shown on the vertical axis. Blocking/staircasing artifacts in (c) and (d) are mitigated in (e) and (f).

4. CONCLUSION

We have shown that the MSIST algorithm can be modified to accommodate complex images with velocity phase encoding by regularizing the wavelet-domain real and imaginary parts separately. We have designed this algorithm to reconstruct a piecewise-smooth phase map obtained from a subsampled MR acquisition, and we have demonstrated how the reconstruction of these images can be accomplished in general. By employing an overcomplete wavelet transform (DT-CWT) in a subband-dependent regularization algorithm with appropriate continuation rules, we have minimized potential artifacts and allowed for fast convergence towards a more accurate sparse solution.

For the specific application of catalytic bead reactor flow imaging, the presence of physical velocity map constraints was handled by enforcing spatial domain phase constraints in this wavelet regularization approach. Consequently, we have improved upon the result in [3] obtained using TV regularization, both quantitatively and qualitatively.

Future improvements using this algorithm on phase-encoded velocity applications could potentially be realized by exploiting inter-scale dependencies and stronger priors on the wavelet coefficients [18, 19]. Improving the convergence speed by finding better parameters for Λ_α [12] and better continuation rules for ϵ and ν may also be worth pursuing.

5. REFERENCES

[1] P.J. Kilner, P.D. Gatehouse, and D.N. Firmin, "Flow measurement by magnetic resonance: a unique asset worth optimising," *Journal of Cardiovascular Magnetic Resonance*, vol. 9, no. 4, pp. 723–728, 2007.

[2] E. Fukushima, "Nuclear magnetic resonance as a tool to study flow," *Annual review of fluid mechanics*, vol. 31, no. 1, pp. 95–123, 1999.

[3] DJ Holland, DM Malioutov, A. Blake, AJ Sederman, and LF Gladden, "Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing," *Journal of Magnetic Resonance*, vol. 203, no. 2, pp. 236–246, 2010.

[4] M. Lustig, D. Donoho, and J.M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[5] J. Romberg, "Variational methods for compressive sampling," in *Electronic Imaging 2007*. International Society for Optics and Photonics, 2007, pp. 64980J–64980J.

[6] F. Zhao, D.C. Noll, J.F. Nielsen, and J.A. Fessler, "Separate magnitude and phase regularization via compressed sensing," *Medical Imaging, IEEE Transactions on*, vol. 31, no. 9, pp. 1713–1723, 2012.

[7] P.L. Combettes, V.R. Wajs, et al., "Signal recovery by proximal forward-backward splitting," *Multiscale Modeling and Simulation*, vol. 4, no. 4, pp. 1168–1200, 2006.

[8] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on pure and applied mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.

[9] C. Vonesch and M. Unser, "A fast thresholded landweber algorithm for wavelet-regularized multidimensional deconvolution," *Image Processing, IEEE Transactions on*, vol. 17, no. 4, pp. 539–549, 2008.

[10] Y. Zhang and N. Kingsbury, "A bayesian wavelet-based multidimensional deconvolution with sub-band emphasis," in *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*. IEEE, 2008, pp. 3024–3027.

[11] M. Guerquin-Kern, M. Haberlin, KP Pruessmann, and M. Unser, "A fast wavelet-based reconstruction method for magnetic resonance imaging," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 9, pp. 1649–1660, 2011.

[12] Y. Zhang and N. Kingsbury, "Improved bounds for modified subband-adaptive iterative shrinkage/thresholding algorithms," *Image Processing, IEEE Transactions on*, to appear 2013.

[13] B. Adcock and A.C. Hansen, "Generalized sampling and infinite-dimensional compressed sensing," Tech. Rep. NA2011/02, University of Cambridge, DAMTP, 2011.

[14] I. Bayram and I.W. Selesnick, "A subband adaptive iterative shrinkage/thresholding algorithm," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1131–1143, 2010.

[15] I.W. Selesnick, R.G. Baraniuk, and N.G. Kingsbury, "The dual-tree complex wavelet transform," *Signal Processing Magazine, IEEE*, vol. 22, no. 6, pp. 123–151, 2005.

[16] Y. Zhang and N. Kingsbury, "Restoration of images and 3D data to higher resolution by deconvolution with sparsity regularization," in *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 2010, pp. 1685–1688.

[17] G.K. Batchelor, *An Introduction to Fluid Dynamics*, Cambridge Mathematical Library. Cambridge University Press, 2000.

[18] L. He and L. Carin, "Exploiting structure in wavelet-based bayesian compressive sensing," *Signal Processing, IEEE Transactions on*, vol. 57, no. 9, pp. 3488–3497, 2009.

[19] L. Sendur and I.W. Selesnick, "Bivariate shrinkage with local variance estimation," *Signal Processing Letters, IEEE*, vol. 9, no. 12, pp. 438–441, 2002.