

PRACTICAL NONCONVEX COMPRESSIVE SENSING RECONSTRUCTION OF HIGHLY-ACCELERATED 3D PARALLEL MR ANGIOGRAMS

Joshua Trzasko, Clifton Haider, and Armando Manduca

Center for Advanced Imaging Research
Mayo Clinic, Rochester, MN, USA
{trzasko.joshua,haider.clifton,manduca}@mayo.edu

ABSTRACT

In this work, a nonconvex Compressive Sensing model targeted at true 3D reconstructions of highly undersampled MR angiograms acquired with parallel imaging is proposed. When combined with the Max-CAPR acquisition sequence, it is demonstrated that high quality, non-view-shared, 3D images of the contrast-filled neurovasculature can be acquired (at acceleration factors exceeding the number of coils) in just over 2 seconds and reconstructed in as few as 14 minutes on a high-performance workstation.

Index Terms— Compressed Sensing, Compressive Sensing, MRI, Angiography, Parallel Imaging

1. INTRODUCTION

Contrast-enhanced magnetic resonance angiography (CE-MRA) is a powerful technique for non-invasively studying vessel morphology. State-of-the-art time-resolved image acquisition strategies such as CAPR [1] and HYPR [2] now offer true 3D imaging of vascular morphology and flow dynamics. In pursuit of higher frame rates, most contemporary time-resolved imaging methods employ “view-sharing” where only a subset of k -space is updated in each frame. Reducing the subset size permits an increased frame rate but the “temporal footprint” [1] of the sequence is consequently enlarged. Ideally, an inverse relationship would exist between these variables.

Increasing frame rate without widening the “temporal footprint” necessitates auxiliary sources of acceleration. One approach is Compressive Sensing (CS) [3, 4], a novel sampling theory asserting that near-exact signal recovery from only a limited sample set may be possible if the signal is sparse (or compressible) in some *a priori* known transform domain. Relaxed sampling conditions and the ability to recover highly-undersampled signals practically translates to faster imaging [5]. Moreover, as view-sharing is no longer necessary, the “temporal footprint” of a sequence would simply be the width of a single time frame. But these gains come at a cost and, unlike conventional MR imaging, a computationally-expensive nonlinear optimization routine

is now required to reconstruct the image. Single-channel true 3D CS MRI reconstructions are reported to take several hours [6, 7, 8]. Lustig et al. [9, 5] have suggested separating the Cartesian 3D problem into a series of 2D problems. While this approach can improve computational performance, powerful 3D connectivity information is inherently lost.

In [1], parallel imaging is used as the auxiliary source of acceleration. Recently, several authors [10, 11, 12, 13, 14, 15] have promoted combining parallel imaging with CS. While this union can lead to even greater accelerations, integrating multisensor data handling into the CS framework only augments the computational burden. To the authors’ knowledge only 2D parallel CS MRI applications have been reported to date, with reconstruction times of several minutes [10, 11, 12, 13, 14, 15].

2. THEORY

Suppose we are interested in imaging a discrete object, \mathbf{f}_0 , defined on Ω ; additionally, let $\mathbf{N} = |\Omega|$. In standard Cartesian MR imaging, the measured signal is $\mathbf{F} = \Phi \mathbf{f}_0 + \mathbf{n}$, where $\Phi : \mathbb{C}^{\mathbf{N}} \rightarrow \mathbb{C}^{\mathbf{K}}$ is a \mathbf{K} -point Discrete Fourier Transform (DFT) operator and \mathbf{n} is complex additive white Gaussian noise (AWGN). When Φ is full rank, direct Fourier inversion is performed. In CE-MRA, pre and post-contrast images are both measured using Φ and complex subtraction is performed prior to reconstruction to remove immaterial background information.

2.1. MR Imaging with Multiple Sensors

During parallel MRI acquisition, an array of images of \mathbf{f}_0 is collected. Each of the \mathbf{C} coil sensors exhibits different sensitization and thus imposes a pointwise spatial modulation, Γ_c , onto the object. All of the modulated images in the array are then sampled using Φ . Denote the array sampling operator as $\Lambda = [\Phi \Gamma_1 \ \Phi \Gamma_2 \ \dots \ \Phi \Gamma_C]^T$. Ignoring intersensor correlation (following [16]), the new signal model is $\mathbf{F} = \Lambda \mathbf{f}_0 + \mathbf{n}$, where $\mathbf{F} = [\mathbf{F}_1 \ \mathbf{F}_2 \ \dots \ \mathbf{F}_C]^T$ and \mathbf{n} is, again, complex AWGN.

For the given signal model, the maximum-likelihood (ML) estimate of \mathbf{f}_0 is the solution of the ordinary least-squares (OLS) problem

$$\mathbf{u}_{\text{LS}} = \arg \min_{\mathbf{u}} \|\mathbf{\Lambda}\mathbf{u} - \mathbf{F}\|_2^2; \quad (1)$$

this formulation was independently proposed as generalized SENSE [16] and SPACE-RIP [17]. If $\mathbf{\Lambda}$ is full-rank (this is generally true if $\mathbf{N} \leq \mathbf{C} \cdot \mathbf{K}$), (1) is well-posed albeit potentially numerically ill-conditioned. The conditioning of (1) can typically be improved through numerical regularization [18] or increasing \mathbf{K} . Loss of acceleration due to an increased \mathbf{K} can be partially recovered via homodyne detection [19, 20, 21].

2.2. MR Imaging via Compressive Sensing

From a Bayesian perspective, numerical regularization coerces an estimator to accord with an *a priori* defined model. Following [22], the *maximum a posteriori* (MAP) estimate of \mathbf{f}_0 after parallel acquisition is given by

$$\mathbf{u}_{\text{reg}} = \arg \min_{\mathbf{u}} \left\{ \mathbf{P}(\mathbf{\Psi}\mathbf{u}) + \frac{\lambda}{2} \|\mathbf{\Lambda}\mathbf{u} - \mathbf{F}\|_2^2, \right\} \quad (2)$$

where $\mathbf{\Psi}$ is an (optional) analysis operator (e.g. wavelet, finite difference) and $\mathbf{P}(\cdot)$ is a model of the expected log-distribution of $\mathbf{\Psi}\mathbf{u}$. λ controls the relative weightings of the prior and objective functionals. In most imaging applications, $\mathbf{P}(\mathbf{\Psi}\mathbf{u})$ is a measure of spatial variation (i.e. “roughness”).

Suppose $\exists \mathbf{\Psi}$ such that $\mathbf{\Psi}\mathbf{f}_0$ is sparse, i.e. $|\text{supp}\{\mathbf{\Psi}\mathbf{f}_0\}| = \mathbf{S} \ll \mathbf{N}$. If \mathbf{f}_0 is measured (via $\mathbf{\Lambda}$) in a domain that is incoherent with $\mathbf{\Psi}$, it can often be exactly reconstructed or well-approximated from as few as $\mathbf{K} = 2\mathbf{S}$ data samples by solving the optimal ℓ_0 -minimization problem, namely, $\mathbf{P}(\mathbf{\Psi}\mathbf{u}) = \|\mathbf{\Psi}\mathbf{u}\|_0 = \sum_{\mathbf{x} \in \Omega} \mathbf{1}(|\mathbf{\Psi}\mathbf{u}(\mathbf{x})| \neq 0)$. However, this combinatorial search is NP-complete [23] and thus intractable for most applications.

Candès et al. [3] and Donoho [4] showed that, at the expense of slight oversampling ($\mathbf{K} \approx 3\mathbf{S} - 5\mathbf{S}$) relative to the ℓ_0 problem, exact or well-approximated recovery is also possible with $\mathbf{P}(\mathbf{\Psi}\mathbf{u}) = \|\mathbf{\Psi}\mathbf{u}\|_1$, a convex and thus tractable functional. Reconstruction of sparse signals via ℓ_1 , now known as Compressive Sensing, is the standard formulation utilized by the MRI community.

Several recent works [24, 25, 26] have explored nonconvex priors that more closely resemble the ℓ_0 measure in attempt to eliminate the oversampling factor required by ℓ_1 -minimization. The cost of this migration is that algorithmic global convergence is no longer assured. Nonetheless, Trzasko and Manduca [26] have demonstrated that concave metric (formal) priors consistently allow high-quality MRI reconstructions from fewer samples than required by conventional CS methods. Theoretical arguments proving that sampling conditions are relaxed when the ℓ_1 penalty is replaced with

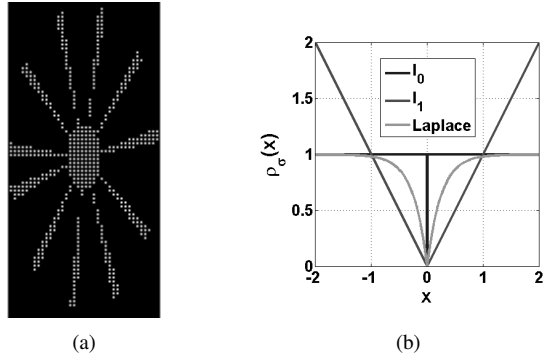


Fig. 1. (a) Example Max-CAPR undersampling pattern ($\mathbf{N}_y = 160, \mathbf{N}_z = 80$). (b) A comparison of sparsity-promoting prior functionals.

one of a large class of concave metric priors have since been developed [27].

3. METHODS

3.1. Data Acquisition and Sampling

The most effective sampling trajectories for CS MRI are those exhibiting incoherent point spread functions (PSF) [5]. In CE-MRA, elliptic-centric (EC) view ordering is commonly used to address signal transiency. Max-CAPR [28] is a rapid ($< 5\text{s}$) 3DFT gradient recalled echo (GRE) sequence possessing both of these properties. Following uniform 2D downsampling for parallel imaging, additional acceleration is achieved by further undersampling the phase-encoding plane in a radial-like fashion (see Fig. 1a). Under Max-CAPR [1], view-sharing is not used and thus the “temporal footprint” and acquisition time are equivalent.

3.2. Sparsity Priors for CE-MRA

CE-MRA images are dominated by morphological information and thus possess sparse spatial gradient representations [5, 26]. Consequently, a $6\times$ redundant (albeit non-invertible) transformation is created using both forward and backward finite differences along each dimension to fully capture 3D vascular connectivity.

Additionally, the near-binary distribution of intensities (and thus gradients) in CE-MRA images can be exploited by employing the nonconvex Laplace prior [26]. For appropriate values of the scale parameter, σ , this functional, defined as

$$\rho_{\sigma}(\cdot) = 1 - \exp(-\sigma^{-1}|\cdot|), \quad (3)$$

closely mimics the ℓ_0 measure over the populated intensity windows (see Fig. 1b).

Let η denote the set of 6 first-order cardinal neighbors. The proposed model for 3D parallel CE-MRA image recovery

is now given by

$$\mathbf{u}_* = \arg \min_{\mathbf{u}} \mathbf{J}(\mathbf{u}), \quad (4)$$

where the energy functional

$$\mathbf{J}(\mathbf{u}) = \sum_{\mathbf{n} \in \eta} \sum_{\mathbf{x} \in \Omega} \rho_{\sigma}(\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})) + \frac{\lambda}{2} \|\Lambda\mathbf{u} - \mathbf{F}\|_2^2 \quad (5)$$

and $\mathbf{D}_{\mathbf{n}}$ is the finite difference operator (Neumann boundary conditions) for neighbor \mathbf{n} . In effort to reduce computation, continuation on σ [26] is not performed.

3.3. Matrix-Free Optimization

Let \mathbf{A}^* denote the adjoint of \mathbf{A} . The (weak) complex gradient of $\mathbf{J}(\mathbf{u})$ with respect to $\bar{\mathbf{u}}$ [29] is defined as

$$\mathbf{L}(\mathbf{u}) = \sum_{\mathbf{n} \in \eta} \mathbf{D}_{\mathbf{n}}^* \mathbf{W}(\mathbf{D}_{\mathbf{n}}\mathbf{u}) \mathbf{D}_{\mathbf{n}}\mathbf{u} + \lambda \Lambda^* (\Lambda\mathbf{u} - \mathbf{F}), \quad (6)$$

where the Hadamard product (element-wise) operator

$$\begin{aligned} \mathbf{W}(\mathbf{D}_{\mathbf{n}}\mathbf{u})_{\mathbf{x},\mathbf{x}} &= \frac{1}{\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})} \cdot \frac{d\rho_{\sigma}(\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x}))}{d\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})} \\ &= \frac{\exp\{-\sigma^{-1}|\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})|\}}{2\sigma|\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})|}. \end{aligned} \quad (7)$$

As (5) is nonconvex, stationary points are not guaranteed to be associated with globally-optimal solutions. Nonetheless, locally-optimal solutions of (4) can be readily found as elements of the kernel of (6) [26].

In practice, inexact quasi-Newton methods are the workhorse of large-scale nonlinear optimization [30]. First recall the standard complex Newton iteration,

$$\mathbf{u}^{t+1} = \mathbf{u}^t - \mathbf{H}^{-1}(\mathbf{u}^t)\mathbf{L}(\mathbf{u}^t), \quad (8)$$

for solving $\mathbf{L}(\mathbf{u}) = \mathbf{0}$, where $\mathbf{H}(\mathbf{u}^t)$ is the complex Hessian [31] of (5). The term ‘‘inexact’’ is used when $\mathbf{H}^{-1}(\mathbf{u}^t)\mathbf{L}(\mathbf{u}^t)$ is not determined exactly but instead estimated, e.g. via truncated Conjugate Gradient (CG) iteration.

Formally, assuming the definition of $\mathbf{L}(\mathbf{u})$ in (6),

$$\mathbf{H}(\mathbf{u}) = \sum_{\mathbf{n} \in \eta} \mathbf{D}_{\mathbf{n}}^* \left[\mathbf{W}(\mathbf{D}_{\mathbf{n}}\mathbf{u}) + \mathbf{W}'(\mathbf{D}_{\mathbf{n}}\mathbf{u}) \right] \mathbf{D}_{\mathbf{n}} + \lambda \Lambda^* \Lambda, \quad (9)$$

where

$$\mathbf{W}'(\mathbf{D}_{\mathbf{n}}\mathbf{u})_{\mathbf{x},\mathbf{x}} = \mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x}) \cdot \frac{d\mathbf{W}(\mathbf{D}_{\mathbf{n}}\mathbf{u})_{\mathbf{x},\mathbf{x}}}{d\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})}; \quad (10)$$

however, as (7) is discontinuous at the origin, the second-order term in (10), and consequently the Hessian, is not well-defined. This instability can be eliminated by considering a linear approximation of (9), namely

$$\tilde{\mathbf{H}}(\mathbf{u}) = \sum_{\mathbf{n} \in \eta} \mathbf{D}_{\mathbf{n}}^* \mathbf{W}(\mathbf{D}_{\mathbf{n}}\mathbf{u}) \mathbf{D}_{\mathbf{n}} + \lambda \Lambda^* \Lambda; \quad (11)$$

this is a generalization of the ‘‘Lagged Diffusivity’’ [32, 26] method. The penalty for stabilizing the Newton iteration in this manner is that convergence is now only expected to be linear (rather than quadratic). Hessian inversion can be further stabilized by relaxing the denominator term of (7) to $|\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})| \approx \sqrt{|\mathbf{D}_{\mathbf{n}}\mathbf{u}(\mathbf{x})|^2 + \epsilon}$. Following [24], continuation on ϵ towards zero is performed within our code implementation.

Naturally, all discrete Fourier Transform (DFT) calls should be carried out using the Fast Fourier Transform (FFT) algorithm. Additionally, note that the finite difference operator (Neumann boundary conditions) can be defined as $\mathbf{D}_{\mathbf{n}} = \mathbf{I} - \mathbf{S}_{\mathbf{n}} - \mathbf{C}_{\mathbf{n}}$, where $\mathbf{S}_{\mathbf{n}}$ is a non-wrapped shift operator (towards neighbor \mathbf{n}) with zero-filling, $\mathbf{C}_{\mathbf{n}}$ is an operator that copies (non-shifted) boundary elements, and \mathbf{I} is the identity operator. By definition, both \mathbf{I} and $\mathbf{C}_{\mathbf{n}}$ are self-adjoint, and the adjoint of $\mathbf{S}_{\mathbf{n}}$ is simply a shift in the negative direction of \mathbf{n} , denoted $\mathbf{S}_{-\mathbf{n}}$. Consequently, $\mathbf{D}_{\mathbf{n}}^* = \mathbf{I} - \mathbf{S}_{-\mathbf{n}} - \mathbf{C}_{\mathbf{n}}$. Unlike most definitions of $\mathbf{D}_{\mathbf{n}}$, the proposed subroutine definition wholly avoids the use of sparse matrices.

4. EXPERIMENTAL RESULTS AND DISCUSSION

All images were acquired on a 3T GE Signa machine (V14.0) using a 14-channel receive coil (sagittal plane). Max-CAPR was employed with the following parameters: TR=4.4ms; TE=2.0ms; FA=30°; BW=±62.5kHz; SENSE acceleration of $R_Y=2, R_Z=2$; FOV = 22cm×22cm×16cm; sampling matrix=256×160×80. The resultant spatial resolution was 0.8594mm×1.375mm×2mm. Only 492 views were measured every 2.1648s for a net acceleration of 26.0163×. Intravenous contrast administration was Multihance[®] injected at 3mL/s followed by 20mL of saline at 3mL/s. The sequence was typically run continuously for 90s, with the first frame used as a the background subtraction mask.

All reconstructions were performed in Matlab[®] on a dual 3.0Ghz Intel Quad-Core Xeon processor computing server with 24MB L2 cache and 32GB 800Mhz DDR2 memory. No view sharing was employed. Each 3D volume (frame) was reconstructed via the proposed technique with fixed 5 outer iterations and 10 inner (CG) iterations; additionally, $\mathbf{u}^0 = \mathbf{0}$, $\lambda = 1 \times 10^{-3}$, and $\sigma=0.25$ as the dynamic range of the data was roughly unity. ϵ was initialized to 1×10^{-1} and decreased by a factor of 10 at each outer iteration. On average, 3D volume reconstruction took less than 14 minutes, substantially lower than contemporary results reported for similar problems. For comparison, each volume was also reconstructed using Tikhonov-SENSE [16, 18] with $\lambda = 1 \times 10^2$ and 15 CG iterations.

Fig. 3 shows sagittal maximum intensity projection (MIP) images for 4 sequential time frames reconstructed using Tikhonov-SENSE (Fig. 3a-d) and the proposed method (Fig. 3e-h). Each time pair was windowed and leveled identically. Note that results obtained with the proposed method contain significantly less noise than Tikhonov-SENSE results. Ad-




Fig. 2. Sagittal MIP images depicting contrast filling the neurovasculature with 2.1648s temporal resolution (no view sharing). All images were acquired at $\sim 26\times$ acceleration using a 14-channel array. Fig. 3a-d show reconstructions via standard SENSE with Tikhonov regularization and Fig. 3e-h using the proposed method.

ditionally, many vessel structures are visible in Fig. 3e-h that are simply not apparent in Fig. 3a-d. By combining the rapid Max-CAPR acquisition strategy with the proposed CS-type reconstruction model, we are able to provide high quality, high temporal fidelity, true 3D visualization of vascular flow dynamics with acceleration factors exceeding the number of coils and computational performance that is finally approaching clinical practicality.

5. ACKNOWLEDGEMENTS

The authors would like to thank Stephen Riederer and Eric Borisch (Mayo Clinic) for their valuable input and making data available for experimentation.

6. REFERENCES

- [1] C. Haider et al. *MRM*, 60(3):749–760, 2008.
- [2] C. Mistretta et al. *MRM*, 55(1):30–40, 2006.
- [3] E. Candès et al. *IEEE TIT*, 52(2):489–509, 2006.
- [4] D. Donoho. *IEEE TIT*, 52(4):1289–1306, 2006.
- [5] M. Lustig et al. *MRM*, 58(6):1182–1195, 2007.
- [6] A. Bilgin et al. In *Proc. ISMRM*, page 337, 2008.
- [7] M. Doneva et al. In *Proc. ISMRM*, page 336, 2008.
- [8] Y. Kim et al. In *Proc. ISMRM*, page 2003, 2008.
- [9] M. Lustig et al. In *Proc. ISMRM*, page 695, 2006.
- [10] K. Block et al. *MRM*, 57(6):1086–1098, 2007.
- [11] B. Liu et al. In *Proc. ISMRM*, page 3154, 2008.
- [12] D. Liang et al. In *Proc. EMBS*, pages 1667–1670, 2008.
- [13] K. King. In *Proc. ISMRM*, page 1488, 2008.
- [14] B. Wu et al. In *Proc. ISMRM*, page 1480, 2008.
- [15] C. Zhao et al. In *Proc. ISMRM*, page 1478, 1478.
- [16] K. Pruessmann et al. *MRM*, 46(4):638–651, 2001.
- [17] W. Kyriakos et al. *MRM*, 44(2):301–308, 2000.
- [18] K. King and L. Angelos. In *Proc. ISMRM*, page 1771, 2000.
- [19] D. Noll et al. *IEEE TMI*, 10(2):154–163, 1991.
- [20] H. Hu et al. *MRM*, 55(1):16–22, 2006.
- [21] M. Bydder and M. Robson. *MRM*, 53(6):1393–1401, 2005.
- [22] A. Raj et al. *MRM*, 57:8–21, 2007.
- [23] B. Natarajan. *SIAM J. Sci. Comp.*, 24(2):227–234, 1995.
- [24] R. Chartrand. *IEEE SPL*, 14(10):707–710, 2007.
- [25] E. Candès et al. *J. Fourier Anal.*, to appear.
- [26] J. Trzasko and A. Manduca. *IEEE TMI*, 28(1):106–121, 2009.
- [27] J. Trzasko and A. Manduca. *submitted to IEEE TSP*.
- [28] C. Haider et al. In *Proc. 20th MRA Workshop*, page 86, 2008.
- [29] D. Brandwood. *IEE Proc.: Comm., Radar, and Sig. Proc.*, 130(1):11–16, 1983.
- [30] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 2006.
- [31] A. van den Bos. *IEE Proc.: Vis., Image, and Sig. Proc.*, 141(6):380–383, 1994.
- [32] C. Vogel and M. Oman. *IEEE TIP*, 7(6):813–824, 1998.