Published in final edited form as:
Magn Reson Med. 2011 October ; 66(4): 1019–1032. doi:10.1002/mrm.22892.

Sparse-CAPR: Highly-Accelerated 4D CE-MRA with Parallel Imaging and Nonconvex Compressive Sensing

Joshua D. Trzasko[1], **Clifton R. Haider**[1], **Eric A. Borisch**[2], **Norbert G. Campeau**[3], **James F. Glockner**[3], **Stephen J. Riederer**[3], and **Armando Manduca**[1]

[1]Department of Physiology and Biomedical Engineering, Mayo Clinic, Rochester, MN, USA

[2]Information Technology, Mayo Clinic, Rochester, MN, USA

[3]Department of Radiology, Mayo Clinic, Rochester, MN, USA

Abstract

CAPR is a SENSE-type parallel 3DFT acquisition paradigm for 4D contrast-enhanced magnetic resonance angiography (CE-MRA) that has been demonstrated capable of providing high spatial and temporal resolution, diagnostic-quality images at very high acceleration rates. However, CAPR images are typically reconstructed online using Tikhonov regularization and partial Fourier methods, which are prone to exhibit noise amplification and undersampling artifacts when operating at very high acceleration rates. In this work, a sparsity-driven offline reconstruction framework for CAPR is developed and demonstrated to consistently provide improvements over the currently-employed reconstruction strategy against these ill-effects. Moreover, the proposed reconstruction strategy requires no changes to the existing CAPR acquisition protocol, and an efficient numerical optimization and hardware system are described that allow for a 256×160×80 volume CE-MRA volume to be reconstructed from an 8-channel data set in less than two minutes.

Keywords

CAPR; Angiography; Compressed Sensing; Compressive Sensing; Image Reconstruction

INTRODUCTION

CE-MRA [1, 2] is a minimally-invasive imaging technique widely used to study vessel lumen morphology and hemodynamics for such clinical tasks as identification of atherosclerotic stenotic and thrombotic lesions, assessment of vascular malformations, and vascular surgery planning. However, to be carried out with diagnostic certainty, these and analogous tasks demand image sequences with simultaneously high spatial and temporal resolution. Given the direct relationship between scan duration and the number of acquired image samples (and thus spatial resolution and field-of-view), generating a CE-MRA image sequence that possesses both of these attributes using conventional acquisition and reconstruction strategies is typically infeasible. Consequently, many techniques have since been proposed for accelerating CE-MRA exams, including parallel imaging [3, 4, 5], novel undersampling and view-ordering methodologies [6, 7, 8], and constrained reconstruction models [9, 10].

Please direct all correspondances to: Armando Manduca, Ph.D., Mayo Clinic (Opus 2-101), 200 First Street SW, Rochester, MN 55905, USA, Phone: (507) 284-8163, Fax: (507) 284-9778, manduca.armando@mayo.edu.

CAPR (Cartesian Acquisition with Projection–Reconstruction-like sampling) [11] is a SENSE-type parallel [12, 13] acquisition paradigm for time-resolved 3D CE-MRA that has permitted high spatial and temporal resolution, diagnostic quality images of varied anatomical regions to be obtained even at high acceleration rates. Previously, both view-shared [11, 14] and non-view-shared [15] CAPR image sequences have been reconstructed using a combination of Tikhonov regression [16] and partial Fourier methods [17, 18]. While this strategy facilitates online clinical assessment and requires little to no user interaction, images reconstructed using these techniques can exhibit significant noise amplification and/or undersampling artifacts when operating at extreme rates of acceleration. As such, it may be beneficial to consider auxiliary use of a more advanced (albeit slower and possibly offline) reconstruction routine that is less subject to these ill-effects to avoid potential complications in diagnosis.

For the highly-undersampled MRI problem, reconstruction strategies based on the *a priori* assumption that images in the class of interest are sparse or compressible (either intrinsically or following mathematical transformation) have been particularly successful [9]. Compressive Sensing (CS) theory [19, 20] asserts that the number of samples needed to generate an accurate approximation of an image is largely determined by its information complexity. For example, images such as angiograms can often be formed from quite limited k-space measurement sets when sparsity-promoting reconstruction techniques such as ℓ_1-minimization [9], homotopic ℓ_0 [21] or ℓ_p-minimization ($0 < p < 1$) [22], or orthogonal matching pursuit (OMP) [23, 24], are employed.

In this work, an offline, sparsity-driven reconstruction framework for CAPR time-series acquisitions is proposed and demonstrated to provide both superior vessel conspicuity and resilience against noise amplification relative to the current clinical reconstruction protocol. Following definition of the recovery model, a robust and efficient quasi-Newton iterative scheme for executing the nonconvex Compressive Sensing (NCCS) reconstruction process is described and a computational framework for handling large data sets in times approaching clinical usefulness is discussed. Several examples comparing the results of both the current and proposed reconstruction strategy across are proffered.

THEORY

A Forward Model for Parallel MRI

Suppose we are interested in forming a discrete image estimate of some anatomy of interest using three-dimensional Cartesian (3DFT) parallel magnetic resonance imaging. Denoting f as the discrete image of interest, a commonly-assumed forward model for this acquisition process is

$$\begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_C \end{bmatrix} = \begin{bmatrix} \Phi\mathscr{F}\Gamma_1 \\ \Phi\mathscr{F}\Gamma_2 \\ \vdots \\ \Phi\mathscr{F}\Gamma_C \end{bmatrix} f + n, \tag{1}$$

where g_c is the signal observed by the c^{th} coil sensor, Γ_c is the element-wise (i.e. diagonal) spatial sensitivity function for the c^{th} coil sensor, $\mathscr{F}$ is the discrete 3D Fourier transform (3DFT) operator, and Φ is a binary operator that identifies the subset of k-space measured during the imaging experiment. The vector n represents system noise and is hereafter assumed to be a complex additive white Gaussian process (AWGN) [13].

Time-Varying Signals and CAPR

In CE-MRA, the signal of interest is inherently transient and thus routinely probed at multiple different time points to characterize patient hemodynamics in addition to vascular morphology. Given the direct relationship between the number of k-space indices measured during an MRI exam and the duration of the exam, spatiotemporal undersampling is often employed to accelerate dynamic MRI exams such as CE-MRA. Reconstruction techniques such as view-sharing [28], HYPR-type processing [8, 29], and constraint or regularization methods [30, 31, 9, 32, 33, 23, 34, 35, 36] that rely on *a priori* assumptions about spatial and/or temporal image structure may be employed to avoid generation of images with significant artifacts.

Assuming that both the signal of interest and the sampling process may be temporally variant in dynamic MRI, (1) can be generalized to

$$\begin{bmatrix} g_1(t) \\ g_2(t) \\ \vdots \\ g_C(t) \end{bmatrix} = \begin{bmatrix} \Phi(t)\mathscr{F}\Gamma_1 \\ \Phi(t)\mathscr{F}\Gamma_2 \\ \vdots \\ \Phi(t)\mathscr{F}\Gamma_C \end{bmatrix} f(t)+n(t), \tag{2}$$

where $t \geq 0$ is an integer corresponding to the frame number. During a CAPR acquisition, k-space is partitioned into two distinct regions: 1) a circularly-symmetric low-pass region that is sampled during each temporal update; 2) a high-pass region of which a different subset is sampled during each temporal update. Thus, the effective sampling operator for CAPR can be described as

$$\Phi(t)=\begin{bmatrix} \Phi_{LO} \\ \Phi_{HI}(t) \end{bmatrix}. \tag{3}$$

The dynamic high-pass sampling operator $\Phi_{HI}(t)$ is also strictly W-periodic and so, $\forall t \geq 0$, $\Phi_{HI}(t) = \Phi_{HI}(t + W)$, and $\forall \tau \in (0, W)$, Trace $(\Phi^{*}_{HI}(t+\tau)\Phi_{HI}(t))=0$. A pictorial example of a CAPR acquisition sequence is given in Figure 1. For a more detailed description of the CAPR acquisition protocol, the reader is referred to [11].

Data acquired using the CAPR strategy is usually reconstructed using a combination of least squares regression (i.e. standard SENSE [12]) with Tikhonov regularization [16] and partial Fourier methods [17, 18]. View-sharing [28] is also routinely performed prior to reconstruction to reduce the presence of spatial undersampling artifacts in the generated images, although zero-filling has been employed for studies of the lower peripheral vasculature [15] in effort to reduce the "temporal footprint" of the exam. Reference or background image subtraction is also typically performed prior to reconstruction of CAPR data.

Constrained Image Reconstruction

Although the current CAPR reconstruction protocol is straightforward to execute, permits online image generation [37], and has proven successful for several highly-accelerated clinical applications, the effectiveness of this strategy can be compromised when operating at very high rates of acceleration. Specifically, the lack of strong regularization in the recovery model can lead to noise amplification during the reconstruction process and, for non-view-shared reconstruction, permit significant undersampling artifacts to manifest. In such scenarios, it may prove beneficial to employ an auxiliary reconstruction protocol that is less prone to these ill-effects.

First, consider the reconstruction of a single time frame volume without any spatial or temporal regularization or preprocessing. The standard SENSE-type parallel MRI recovery model (and the one underlying the current CAPR reconstruction protocol) is

$$\tilde{u}(t)=\arg\min_{u}\sum_{c=1}^{C}\|\Phi(t)\mathscr{F}\Gamma_c u - g_c(t)\|_2^2, \tag{4}$$

where u is the optimization variable. In practice, however, (4) is often rank-deficient due to undersampling or inaccuracies in coil sensitivity profile estimates, and some form of regularization is needed to stabilize the recovery problem. Typically, generalizations of (4) of the form

$$\tilde{u}(t)=\arg\min_{u}\left\{\alpha P(\Psi u)+\sum_{c=1}^{C}\|\Phi(t)\mathscr{F}\Gamma_c u - g_c(t)\|_2^2\right\}, \tag{5}$$

are considered, where

$$P(\Psi u)=\sum_{s\in\Omega}\rho\left([\Psi u](s)\right) \tag{6}$$

is a separable functional (defined element-wise by the penalty functional $\rho(\cdot)$) that promotes *a priori* assumptions about the structure of the target signal, α is a mixing parameter, and Ω is the domain of the (transformed) image. Tikhonov regression, for example, employs $P(\cdot)=\|\cdot\|_2^2$ (or equivalently, $\rho(\cdot)=|\cdot|^2$) to promote low (transform) energy solutions [16]. Analogously, Markov Random Field (MRF) based priors [30, 32] may be used to promote piecewise homogeneous solutions.

Following the work of Lustig et al. [9], the use of prior functionals that promote image sparsity (either intrinsically or in some mathematical transform domain) has become increasingly popular for the undersampled MRI reconstruction problem. Recall that Compressive Sensing (CS) theory [19, 20] asserts that the number of samples sufficient to form an accurate approximation of an image is dominantly determined by its underlying complexity. If an efficient means of representing the image of interest (i.e. via transformation by Ψ) is *a priori* known, a sparsity-promoting recovery model such as ℓ_1-minimization ($P(\cdot)=\|\cdot\|_1$, or $\rho(\cdot)=|\cdot|$) can allow for that image (or an accurate approximation of it) to be generated from many fewer samples than required by traditional Nyquist-limited linear strategies. Block et al. [31], and later several other groups (e.g., [38, 39, 40]), noted that, in the context of parallel MRI, CS provides a mechanism for potentially accelerating by factors that significantly exceed the number of coils.

Recall that a signal is said to be "sparse" if only a few of its elements are non-zero, and that its (relative) sparsity can be assessed by counting how many non-zero elements it has. This suggests that the ℓ_0-norm1 ($P(\cdot)=\|\cdot\|_0$, or $\rho(\cdot)=\mathbf{1}(|\cdot|\neq 0)$, where $\mathbf{1}(\cdot)$ is an indicator functional) should be used as the regularizing functional or sparsity measure. In [19] and [20], it was asserted that sparse signals could (theoretically) be accurately estimated from many fewer samples than traditionally required (e.g., to satisfy the Nyquist criterion) by ℓ_0-minimization. However, since directly solving the ℓ_0-minimization problem is NP-hard [41] and thus impractical, alternative sparsity measures are needed. Most applications of CS theory for undersampled MRI reconstruction employ $P(\cdot)=\|\cdot\|_1$ as the sparsity measure,

[1]Due to lack of homogeneity, the ℓ_0-norm is technically not a true norm.

since the convexity of the ℓ_1-norm guarantees that the global optimum of the optimization problem will be identifiable. The cost of this ability is an increase in the sampling rate required to guarantee a prescribed level of performance [19, 20]. Several authors (e.g. [25, 26, 21, 27]) have also suggested the use of nonconvex sparsity-promoting penalty functionals that more closely resemble the ℓ_0-norm than does the ℓ_1-norm. Although they lack global convergence guarantees, in practice reconstruction models employing (nonconvex) concave metric prior functionals [27] such as $\rho(\cdot) = \log(\sigma^{-1}|\cdot| + 1)$, for any $\sigma \in [0, \infty)$, often yield higher quality reconstructions than ℓ_1-based methods without significantly increasing computational expense [25, 26, 22, 21, 42]). It is noted that Raj et al.'s [30] truncated linear penalty functional is closely related to these so-called nonconvex CS (NCCS) methods; however, reconstruction problems with acceleration factors exceeding the number of coils were not considered in that work.

Specifying the Reconstruction Model for CE-MRA

Noting that the information of interest in background-suppressed angiograms (acquired using any technique) is dominantly morphological [43], Lustig et al. [9] asserted that transforming these images via finite spatial difference operators can lead to even more highly compressible image representations. Çukur et al. [44] later employed this model in steady-state free precession (SSFP) MRA, rapidly acquiring only a small set of phase-encoded k-space views after each magnetization preparation step to ensure contrast retention and reconstructing a relatively artifact-free image from this limited view set using ℓ_1-minimization. In a standard time-resolved CE-MRA exam, pre-contrast images are typically available. The pre-contrast image, or background, is subsequently subtracted from the contrast-enhanced image to reveal images (ideally) depicting only the gadolinium-enhanced vessel lumens, and possibly perfused tissue regions [45]. Regardless of the whether either source image is itself sparse or compressible, a finite spatial difference transformed subtraction image certainly will be.

Let r be a pre-injection background image (assumed fully sampled) and D_n a finite spatial difference operator for offset direction n. Given the proposed sparsifying transformation for CE-MRA images, the constrained parallel MRI reconstruction model in (5) can be adapted as

$$\tilde{u}(t)=\arg\min_u J(u)=\arg\min_u \left\{\alpha\sum_{n\in\eta}P(D_n(u-r))+\sum_{c=1}^{C}\|\Phi(t)\mathscr{F}\Gamma_c u-g_c(t)\|_2^2\right\}, \qquad (7)$$

where η is the set of all offset directions over which spatial differences are to be computed. For 3D volume reconstructions, it is recommended that this set span all three cardinal directions to maximally exploit vascular connectivity. It is noted that use of a reference image within the penalty functional of a constrained MRI reconstruction model as done above has been previously considered by several authors, both in the context of a generalized Tikhonov regression model for parallel MRI reconstruction [46, 47] and within CS reconstruction models [48, 49, 50] as a means of sparsifying the image of interest.

Although reconstruction models like (7) are intuitively attractive, the general unavailability of a fully sampled reference image makes employing them in practice challenging. Akin to what is done during HYPR reconstruction [8], Chen et al. [48], and Samsonov et al. [50] and Wu et al. [49], independently suggested that the reference image be generated by compositing a subset of the individual undersampled frames of the time sequence. By construction, the composite reference image will contain both background and intravascular signal. It will also exhibit characteristics similar to a single long acquisition image. Correspondingly, there are two potential drawbacks to using this strategy. First, prolonged

sampling of a time-varying process imposes a nonuniform k-space modulation that results in spatial blurring of the reference image [51, 52, 53]. Second, due to the unavoidable disparity in intravascular signal intensity between any individual time frame and the composite image, a subtraction image formed using the latter will inherently contain residual signal in areas of either previous or eventual contrast enhancement. A subtraction image formed using a pure background reference should not contain this residual signal, and is thus expected to possess a higher degree of sparsity or compressibility than its composite-based analog.

To circumvent these complications, consider the following. Define the subtraction image $\upsilon = u - r$. Since r is presumed constant, observe that (7) can be recast as

$$\tilde{\upsilon}(t)=\arg\min_{\upsilon} J(\upsilon)=\arg\min_{\upsilon}\left\{\alpha\sum_{n\in\eta}P(D_n\upsilon)+\sum_{c=1}^{C}\|\Phi(t)\mathcal{F}\Gamma_c\upsilon-h_c(t)\|_2^2\right\}, \quad (8)$$

where $h_c(t) = g_c(t) - \Phi(t)\mathcal{F}\Gamma_c r$ [49, 50]. If r is completely known, the non-subtracted reconstruction image can of course be recovered simply by computing $\tilde{u}(t) = \tilde{\upsilon}(t) + r$ and the reformulation in (8) proffers no real benefit over (7). Note, however, that in many CE-MRA applications only the subtraction image is actually of interest. When $\tilde{\upsilon}(t)$ is all that is requested, observe that only a set of sampled versions of the background reference image is required by the reconstruction model. In a time-resolved CE-MRA exam like CAPR that employs a W-periodic sampling operator, these data sets can be obtained simply by initiating the acquisition process W cycles prior to contrast injection such that r is sampled under every utilized $\Phi(t)$. Since even the complete sample set needed by $h_c(t)$ indexes only a small fraction of k-space, it can be accumulated quickly and in significantly less time than would be needed to form a complete image of the contrast-free background. Thus, the effective benefits of possessing a complete, unbiased, and accurate background reference image can be practically realized with minimal protocol overhead and simple complex k-space subtraction of the undersampled background signal from the raw coil data.

Embedding View-Sharing

View-sharing [28] can improve the numerical conditioning of an undersampled dynamic series reconstruction albeit at the expense of some loss of spatiotemporal resolution [52, 53]. As both view-shared [11] and non-view-shared reconstructions [15] of CAPR sequences have been previously considered, the use of view-sharing in conjunction with the proposed reconstruction strategy is discussed. For a detailed discussion about the implications of view-sharing on the temporal resolution of a CE-MRA exam and its appropriateness for different anatomical studies, the reader is referred to [11, 15].

As with background subtraction, view-sharing can be performed directly on the raw k-space data on a coil-by-coil basis and so requires only trivial modification of the reconstruction model. Let $\Theta \in [1, W]$ be the temporal window over which view-sharing is performed (typically, only $\Theta = 1$ [15] or $\Theta = W$ [11] are considered). For $t \geq \Theta - 1$, performing view-sharing on a non-background-subtracted CAPR data series generates a signal described by

$$j_c(t)=\Phi_{VS}(t,\Theta)\left[\Phi_{LO}^{*}g_c(t)+\sum_{\theta=0}^{\Theta-1}\Phi_{HI}^{*}(t-\theta)g_c(t-\theta)\right], \quad (9)$$

where the view-shared sampling operator

$$\Phi_{VS}(t,\Theta)=\begin{bmatrix}\Phi_{LO}\\ \Phi_{HI}(t)\\ \vdots\\ \Phi_{HI}(t-\Theta+1)\end{bmatrix} \tag{10}$$

and the notation A^* denotes the adjoint of the operator A. Similarly, performing view-sharing on a background-subtracted CAPR data series generates a signal described by $k_c(t) = j_c(t) - \Phi_{VS}(t, \Theta)\mathcal{F}\Gamma_c r$.

For $\Theta = W$, observe that $\Phi_{VS}(t, \Theta)$ will always select the same set of k-space indices. Although the first complete view-shared pre-injection data set could be used as the background signal (i.e., letting $\Phi_{VS}(t, \Theta)\mathcal{F}\Gamma_c r = k_c(\Theta - 1)$) [11], Borisch et al. [54] have reported that this strategy can lead to Θ-periodic temporal intensity fluctuations in the reconstructed image sequence. Conjecturing that this phenomenon results from subtracting volumes with differing magnetization histories, the usage of a history-matched subtraction strategy was suggested as a remedy for this problem. This approach, which requires $\Theta + 1$ pre-contrast-arrival time frames be available to form the background signal set, is also adopted in this work. Generalizing the reconstruction problem in (8) to incorporate view-sharing thus yields

$$\tilde{v}(t)=\arg\min_{v} J(v)=\arg\min_{v}\left\{\alpha\sum_{n\in\eta}P(D_n v)+\sum_{c=1}^{C}\|\Phi_{VS}(t,\Theta)\mathcal{F}\Gamma_c v-k_c(t)\|_2^2\right\}, \tag{11}$$

which is applicable for $t \geq \Theta + 1$.

Numerical Optimization

Adopting the model in (11) for the reconstruction of CAPR CE-MRA data, what remains to be discussed is how to numerically compute solutions to such problems. For the prior functional class of interest, (11) generally does not possess a closed-form solution and thus iterative methods are needed to identify minima of $J(\upsilon)$. In this work, an inexact quasi-Newton iteration similar to that described in [32, 21, 27] is adopted for this task; however, other strategies such as nonlinear conjugate gradient iteration [31, 9] or discrete optimization with graph cuts [30] are also possible.

Complex quasi-Newton iterations are typically of the form

$$v_{i+1}=v_i+\Delta_i=v_i - B^{-1}(v_i)L(v_i), \tag{12}$$

where $i \geq 0$ is the iteration number, the cost functional gradient $L(\upsilon_i) = \nabla_{\bar{\upsilon}} J(\upsilon)|_{\upsilon=\upsilon_i}$ [56, 57], and the Hessian approximator $B(\upsilon_i) \approx H(\upsilon_i) = \nabla_{\bar{\upsilon}}\nabla_{\upsilon} J(\upsilon)|_{\upsilon=\upsilon_i}$ [58]. . The term “inexact” arises when Δ_i is estimated (e.g. via truncated Conjugate Gradient (CG) iteration) rather than exactly determined. As $P(\cdot)$ iscontinuous but nonsmooth (this is true for both convex and nonconvex prior functionals), a weak gradient can be employed to yield

$$L(v_i)=\alpha\sum_{n\in\eta}D_n^*\Lambda(D_n v_i)D_n v_i+\sum_{c=1}^{C}\Gamma_c^*\mathcal{F}^*\Phi_{VS}^*(t,\Theta)\left(\Phi_{VS}(t,\Theta)\mathcal{F}\Gamma_c v_i-k_c(t)\right), \tag{13}$$

where the diagonal operator

$$\Lambda(D_n \upsilon_i)_{(s,s)} = \frac{1}{2|[D_n \upsilon_i](s)|} \cdot \frac{\partial \rho\left([D_n \upsilon_i](s)\right)}{\partial |[D_n \upsilon_i](s)|}. \tag{14}$$

Similarly, the utilized linear approximation of the true Hessian operator is given by

$$B(\upsilon_i) = \frac{\alpha}{2} \sum_{n \in \eta} D_n^* \Lambda(D_n \upsilon_i) D_n + \sum_{c=1}^{C} \Gamma_c^* \mathcal{F}^* \Phi_{VS}^*(t, \Theta) \Phi_{VS}(t, \Theta) \mathcal{F} \Gamma_c, \tag{15}$$

which is a generalization of the model used by Vogel and Oman [59] for total variation (TV) denoising. .

As noted in [25] and [21], spurious local minima exist in nonconvex optimization problems and specific effort must be put forth to avoid those that are far from the global optimum. In this work, two numerical strategies are used to address this challenge: 1) warm-starting and 2) ε-continuation [60, 61]. Warm-starting consists of providing an iterative estimation process with an initial guess, υ_0, that is already close to the targeted solution. Although this strategy is generally inapplicable in static imaging applications, during causal reconstruction of a time-series one reconstructed time frame volume can simply be used as the initial estimate for the subsequent time frame. Assuming that the first frame of the (background subtracted) time-series was acquired prior to contrast arrival, a zero image can be used as the initial estimate for the first time frame. The second strategy, ε-continuation, uses functional smoothing to manipulate the targeted solution space such that only significant minima remain. Suppose that every $|\cdot|$ inherent to $\rho(\cdot)$ is replaced with $|\cdot|_\varepsilon = \sqrt{|\cdot|^2 + \varepsilon}$, for some smoothing parameter $\varepsilon > 0$. When ε is large, many spurious local minima will be "smoothed away" [61]; however, recall that it is the nonsmootheness of $\rho(\cdot)$ that ultimately permits the generation of sparse solutions. As such, ε is made a decreasing function of the Newton iteration number (i.e. $\lim_{i \to \infty} \varepsilon_i = 0$). In this work, the diminishing sequence $\varepsilon_{i+1} = \beta \varepsilon_i$ ($0 < \beta < 1$) is used. Since warm-starting theoretically initializes each reconstruction problem close to its solution, we typically prescribe each time frame reconstruction to consist of only a few Newton steps in conjunction with a moderately-aggressive $\beta = 0.1$. Under ε-continuation, (14) is correspondingly replaced by

$$\Lambda(D_n \upsilon_i)_{(s,s)} = \frac{1}{2|[D_n \upsilon_i](s)|_{\varepsilon_i}} \cdot \frac{\partial \rho\left(|[D_n \upsilon_i](s)|_{\varepsilon_i}\right)}{\partial |[D_n \upsilon_i](s)|_{\varepsilon_i}}. \tag{16}$$

Figure 2 displays an example progression of $|\cdot|_{\varepsilon_i}$.

METHODS AND MATERIALS

Implementation of the Reconstruction Procedure

As discussed in the Theory section, the acquired raw k-space data is first (optionally) view-shared, after which (optionally) history-matched k-space reference subtraction is performed. The resultant signal set is then passed to the reconstruction engine along with both the set of coil sensitivity functions (acquisition discussed in the next section) and the warm start initializer. After reconstruction, the generated image volume is both saved off and passed back into the iterative quasi-Newton solver to be used as the warm-start initializer for reconstruction of the subsequent time frame volume.

Considering specification of the targeted cost functional in (11), a normalized variant of the Laplace functional [21],

$$\rho(\cdot)=\frac{1-\mathrm{e}^{-\frac{|\cdot|}{\sigma}}}{1-\mathrm{e}^{-\frac{1}{\sigma}}}, \tag{17}$$

was employed as the sparsity measure for all experiments in this work. (17) was selected because its derivative has a simple form and was thus easy to work with. We note that other nonconvex functional forms could be employed and would be expected to perform similarly [61, 62, 27]. Alternatively, a standard (ℓ_1-based) CS reconstruction could be performed by defining $\rho(\cdot) = |\cdot|$, but this option was not considered here. Setting the scale parameter $\sigma \in [0, \infty)$ to be one-quarter of the anticipated dynamic range of the reconstructed image was observed to be effective in practice. The same σ was used for all reconstructions of similar anatomy. It is again noted that since this functional is nonconvex, all presented reconstruction results are only formally guaranteed to be local optima of the targeted cost functional in (11). The specific form of the quasi-Newton iteration strategy used for all reconstructions is described below:

For all experiments, $\beta = 0.1$, $\mathrm{maxIter}_{outer} = 5$, $\mathrm{maxIter}_{inner} = 1$, and $\varepsilon_0=\mathrm{pow}\left(10, \lfloor \log_{10}\left(\frac{\sigma^2}{10}\right)\rfloor\right)$. The regularization parameter α implicit to $B(\cdot)$ and $L(\cdot)$ in Algorithm 1 was manually assigned during each reconstruction. As with σ, the same α was used for all reconstructions of similar anatomy. Between 15–25 linear Conjugate Gradient (CG) iterations were used to solve for the Newton update step, Δ_i. It is noted that, by fixing the number of iterations per reconstruction cycle, the time-to-completion for an entire time-series reconstruction can be estimated based on the reconstruction time for the first frame. All DFT operations needed to form $B(\cdot)$ and $L(\cdot)$ should, of course, be carried out using the FFT algorithm. Analogously, a generic operator formalism (described in the Appendix) was adopted for all finite spatial difference transformations. While this action does not change the complexity of the transformation, it facilitates the production of an optimized code implementation. For all experiments, both forward and backward finite differences were calculated along the three cardinal axes

Algorithm 1

Fixed-Stage Quasi-Newton Iteration with ε-Continuation and Warm-Start Initialization

define: $\varepsilon_0 > 0$, $\beta \in (0, 1)$, $\mathrm{maxIter}_{outer}$, $\mathrm{maxIter}_{inner}$;

let: $i = 0, \upsilon_0 = \tilde{\upsilon}(t)$, $\mathrm{count}_{outer} = 0$;

repeat

  $\mathrm{count}_{inner} = 0$;

  repeat

    solve: $B(\upsilon_i)\Delta_i = -L(\upsilon_i)$;

    $\upsilon_{i+1} = \upsilon_i + \Delta_i$;

    $\mathrm{count}_{inner} = \mathrm{count}_{inner} + 1$;

    $i = i + 1$;

  until $\mathrm{count}_{inner} \geq \mathrm{maxIter}_{inner}$;

  $\varepsilon_{i+1} = \beta\varepsilon_i$;

  $\mathrm{count}_{outer} = \mathrm{count}_{outer} + 1$;

until $\mathrm{count}_{outer} \geq \mathrm{maxIter}_{outer}$;

$\tilde{\upsilon}(t+1) = \upsilon_i$;

(x, y, z), such that $\eta = \{(1, 0, 0), (-1, 0, 0), (0, 1, 0), (0, -1, 0), (0, 0, 1), (0, 0, -1)\}$. Although use of higher-order and off-axis neighbors is also possible under the proposed numerical framework, these options were not considered in this work.

All reconstructions were executed offline on a dedicated 8-node cluster system, where each node holds two 3.4 GHz Xeon processors and 16GB memory. The entire reconstruction system is currently implemented in C++ under the templated class framework developed by Borisch et al. in [37], and makes use of the message passing interface (MPI) for inter-node and intra-node communication and the OpenMP library for intra-node parallelism as well as FFTW. On this system, reconstruction of a single 256×160×88 neurovasculature volume from an 8-coil data set using the proposed computation paradigm requires only about 100 seconds of computation, or about one hour for the entire time series of 36 volumes.

In Vivo Experiments

All imaging experiments were conducted according to a protocol approved by the Institutional Review Board (IRB). Five time-resolved CE-MRA studies of the neurovasculature or peripheral vasculature performed at 3.0T (Signa; GE Healthcare, Milwaukee, WI) using the CAPR acquisition protocol are included here. Included neurovascular studies were performed using a commercially-available receive-only head coil while peripheral vascular studies were performed using in-house built modular receive-only arrays [63].

For all experiments, an elliptical-centric (EC) ordered [6] 3D spoiled gradient-echo sequence was employed with the following parameters: bandwidth = ±62.5kHz; flip angle = 30°; readout direction = superior-inferior (SI). Experiment specific parameters such as scanner software version, repetition time (*TR*), echo time (*TE*), number of coils (*C*), number of views inherent to Φ_{LO} (M_{LP}), number of views inherent to $\Phi_{HI}(t)$ (M_{HP}), reconstructed volume dimensions ($N_x \times N_y \times N_z$), number of reconstructed time frames (N_t), sampling periodicity (*W*), number of view-sharing frames (Θ), spatial resolution ($\delta_x \times \delta_y \times \delta_z$), temporal update rate (δ_t), temporal footprint (*TF*), acceleration factor (*AF*), and undersampling factor (*USF*) can be found in Table 1. The SENSE acceleration factor (*R*) underlying the CAPR sampling strategy is also included for completeness.

The acceleration factor is herein defined as

$$AF = \frac{N_y N_z}{M_{LP} + M_{HP}}, \tag{18}$$

which is simply the inverse of the fraction of total views investigated during a single CAPR update. This quantity, however, is not indicative of how underdetermined reconstruction of a single time frame volume will be given that several data sets and view-sharing may be utilized. Also note that (18) includes undersampling the corners of k-space, whereas prior measures of acceleration for CAPR acquisitions typically do not. Alternatively, the undersampling factor,

$$USF = 100\% \times \left[1 - \min\left\{\frac{C(M_{LP} + \Theta M_{HP})}{N_y N_z}, 1\right\}\right], \tag{19}$$

accounts for the total number of measurements (from any time) during reconstruction of a single image volume. By construction, any recovery problem involving image estimation from more samples than there are reconstructed volume elements achieves a score of USF = 0%. For an example acqusition with $N_y = N_z = 256$, $M_{LP} = 250$, $M_{HP} = 3000$, and $W = 4$, a non-view-shared reconstruction (i.e. $\Theta = 1$) would have $USF = 60.33\%$,, whereas a view-shared reconstruction with $\Theta = 4$ would have $USF = 0.00\%$. In particular, conventional parallel acquisitions with acceleration factors smaller than the number of coils are not undersampled by this definition, since the resulting reconstruction problem is not mathematically underdetermined.

Prior to initialization of each time-resolved sequence, a high-resolution calibration scan (sequence = GRE; bandwidth = ±31.25kHz; flip angle = 10°; full spatial resolution for neurovascular exams, 2 × 2 reduction in phase-encoded plane for peripheral vascular exams) was performed and coil sensitivity profiles were estimated using root-sum-of-squares demodulation. Time-series acquisition was initiated prior to contrast injection to allow the earlier described set of pure background images to be obtained for reference subtraction. Intravenous administration of 20 mL of a gadolinium-based contrast agent (MultiHance; Bracco Diagnostics, Princeton, NJ) at 3mL/s, followed by 20mL of saline at 3mL/s, was performed using a commercially-available power injector. Following completion of an exam, time series data was (optionally) view-shared and reference subtracted. The processed data set was then reconstructed using both the standard clinical protocol described in [11] (hereafter referred to as “SENSE+Homodyne”) and the proposed iterative reconstruction protocol (hereafter referred to as “Nonconvex Compressive Sensing”). Following [14], images were 3× sinc-interpolated to make small features more apparent (albeit while producing subtle zipper artifacts) when a region-of-interest was enlarged [64].

RESULTS

Example 1: Figure 3 shows a subset of results obtained for a (view-shared) neurovasculature study of a pediatric patient. The NCCS reconstruction time was slightly less than 2 minutes per 3D volume. Observe that both the NCCS cross-sectional and MIP images exhibit better vessel-to-background contrast than do the corresponding SENSE +Homodyne results. Also, note in (f) and (l), and their corresponding enlargements (m) and (n), a mass-like region of increased-contrast appears near the posterior right frontal lobe of the brain. The enhancing abnormality corresponds to post-ischemic hyperperfusion of a large section of the cortex several days following stroke. Although these regions of contrast uptake could likely be localized from the SENSE+Homodyne results, they can be identified with higher confidence using images generated under the proposed technique.

Example 2: Figure 4 shows a subset of results obtained for a (view-shared) patient study of the feet. For improved visualization, only one foot is shown in the renderings. The NCCS reconstruction time was slightly less than 10 minutes per 3D volume. In both the early and middle filling stages, major vessels such as the distal anterior tibial and dorsalis pedis arteries are clearly visible in both reconstruction results; however, observe that many smaller distal vessel branches are obscured by noise in the SENSE-Homodyne results,, whereas they are apparent in the NCCS reconstruction results. This phenomenon is highlighted in (g) and (h), the respective enlargements of (b) and (e). Also note that the multiple stenotic lesions throughout the distal anterior and tibial artery, as well in several secondary branching vessels, are not distorted by the regularized reconstruction procedure.

Example 3: Figure 5 shows a subset of results obtained for a (view-shared) neurovasculature volunteer study. The NCCS reconstruction time was slightly less than 2 minutes per 3D volume. Observe in this example that a flow-based artifact (dominantly

along the anterior-posterior (AP) direction) has manifested from the transverse sinus in the SENSE+Homodyne image, but that this artifact is (anecdotally) suppressed in the NCCS result. Additionally, notice the superior contrast and homogeneity of the contrast-perfused parotid, submandibular, and sublingual salivary glands as well as the paranasal sinuses in the NCCS result over the SENSE+Homodyne result. The enlargements in (c) and (d) further highlight both of these image properties.

Example 4: Figure 6 shows a subset of results obtained for a (non-view-shared) study of the hands. The NCCS reconstruction time was slightly less than 3 minutes per 3D volume. In both the early and middle filling stage volumes, note that there is strong ringing orthogonal to the radial and ulnar arteries in the SENSE+Homodyne results, but these artifacts are not apparent in the NCCS results. In the enlargements of the middle filling stage volume in (g) and (h), notice that vessels (in particular, the proper and common palmar digital arteries) appear sharper in the NCCS result than in the SENSE+homodyne result, and that background haze has been removed, making the deep palmar arch and palmar metacarpal arteries more apparent. As before, background signal homogeneity is also superior in the NCCS reconstruction.

Example 5: Figure 7 shows a subset of results obtained for a (non-view-shared) volunteer study of the vasculature of the calves. The NCCS reconstruction time was slightly less than 6 minutes per 3D volume. Observe in (i) and (j), which are enlargements of the trifurcation region of the middle filling stage MIP images, the improved sharpness of the major vessels as well as the resolution of secondary branching vessels originating off the posterior tibial and peroneal arteries. Evidence of vessel sharpening is also apparent in the cross-section enlargements in (k) and (l). In the axial plane of the SENSE+Homodyne result, radial-like artifacts matching the point-spread function of the CAPR sampling pattern permeate the image; these artifacts are not apparent in the NCCS result. Also, again note that superior background homogeneity of the NCCS result and, in this example, suggestion of muscular perfusion.

DISCUSSION

In this work, we have presented a sparsity-driven reconstruction framework for time-resolved 3D (4D) CAPR CE-MRA image sequences that consistently generates angiograms with less noise amplification and fewer undersampling artifacts than does the currently-employed clinical reconstruction protocol (SENSE+Homodyne). Moreover, the presented strategy requires no changes to the existing CAPR acquisition protocol, and can be executed in times that are approaching clinical usefulness. Generalizing the reconstruction strategy for similar CE-MRA acquisition sequences (e.g. [65]) would also be straightforward.

Even though the proposed reconstruction framework has several demonstrated advantagesover the current CAPR reconstruction protocol, it is not without limitation. Further development is undoubtedly required to reduce the computation time of the reconstruction routine such that it is truly clinically practical. As also recently demonstrated by Murphy et al. [66], optimized use of advanced computing hardware such as graphics processing units (GPU) may be the key to making iterative sparsity-driven reconstruction algorithms practical for routine clinical use. We are continuing to investigate mapping the proposed reconstruction algorithm to alternative hardware platforms as well as algorithmic modifications that would allow for increased parallelizability.

Another limitation of the proposed reconstruction strategy is that both the recovery model and the numerical optimization routine each possess several parameters that must be reasonably assigned to achieve high quality results. For example, assigning too small a value

to the regularization parameter, α, can allow noise amplification during the reconstruction, whereas assigning too large a value to α can result in over-sparsification and thus oversmoothing of the image, and a corresponding loss of features. However, it is noted that once an effective setting is found for a particular image scenario, it can be reused (e.g. all neurovascular studies are run under the same settings).

Although beyond the scope of this manuscript, it is expected that the improved vessel-to-background conspicuity (i.e. contrast) and homogeneity of vessels, enhancing structures, and background tissue resulting from the NCCS reconstruction strategy for CAPR CE-MRA image series will lead to improved diagnosis, with better sensitivity for detection of abnormalities, and fewer false positive/negative interpretations. Nonetheless, a formal radiological comparison of these two reconstruction strategies is still needed and will be the subject of a separate, future work.

APPENDIX

FINITE SPATIAL DIFFERENCE COMPUTATION

The adopted finite difference spatial model assumes that any finite spatial difference centered inside Ω but for which the neighbor of interest is outside Ω will be zero. More specifically, $\forall s \in \Omega$ and any n,

$$[D_n u](s) = \begin{cases} u(s) - u(s+n), & \text{if } (s+n) \in \Omega \\ 0, & \text{else} \end{cases} .$$

Note that this operator can be abstracted as $D_n = I - S_n - C_n$, where I is the identity operator, S_n is a non-wrapped shift operator (towards neighbor n) with zero-filling, and C_n is an operator that copies non-shifted boundary elements. Let $n = \{\Delta_x, \Delta_y, \Delta_z\}$ and $s = \{x, y, z\}$, where $x \in [0, N_x)$, $y \in [0, N_y)$, and $z \in [0, N_z)$. Assuming the notation $u(s + n) = u(x + \Delta_x, y + \Delta_y, z + \Delta_z)$, the component operators of D_n can then be defined as

$$[S_n u](s) = \begin{cases} u(s+n), & \text{if } (x \in \Upsilon_x) \wedge (y \in \Upsilon_y) \wedge (z \in \Upsilon_z) \\ 0, & \text{else} \end{cases} \quad (20)$$

and

$$[C_n u](s) = \begin{cases} 0, & \text{if } (x \in \Upsilon_x) \wedge (y \in \Upsilon_y) \wedge (z \in \Upsilon_z) \\ u(s), & \text{else} \end{cases} , \quad (21)$$

where

$$\begin{aligned} \Upsilon_x &= [\max(0, -\Delta_x), \min(N_x, N_x - \Delta_x)) \\ \Upsilon_y &= [\max(0, -\Delta_y), \min(N_y, N_y - \Delta_y)) \\ \Upsilon_z &= [\max(0, -\Delta_z), \min(N_z, N_z - \Delta_z)) \end{aligned} .$$

Note both I and C_n are inherently self-adjoint, and application of the adjoint of S_n simply generates a shift in the opposite direction of n. Consequently, $D_n^* = I - S_{-n} - C_n$. This trivial construction allows for consistent boundary handling during iteration, and correspondingly facilitates code optimization.

References

1. Prince M, Yucel E, Kaufman J, Harrison D, Geller S. Dynamic gadolinium enhanced three-dimensional abdominal MR arteriography. J. Magnetic Resonance Imaging. 1993; 3(6):877–881.
2. Korosec F, Frayne R, Grist T, Mistretta C. Time-resolved contrast-enhanced 3D MR angiography. Magnetic Resonance in Medicine. 1996; 36(3):345–351. [PubMed: 8875403]
3. Weiger M, Pruessmann K, Kassner A, Roditi G, Lawton T, Reid A, Boesiger P. Contrast-enhanced 3D MRA using SENSE. J. Magnetic Resonance Imaging. 2000; 12(5):671–677.
4. Golay X, Brown S, Itoh R, Melhem E. Time-resolved contrast-enhanced carotid MR angiography using sensitivity encoding (SENSE). American Journal of Neuroradiology. 2001; 22(8):1615–1619. [PubMed: 11559518]
5. Hu H, Madhuranthakam A, Kruger D, Huston J III, Riederer S. Improved venous suppression and spatial resolution with SENSE in elliptical centric 3D contrast-enhanced MR angiography. Magnetic Resonance in Medicine. 2004; 52(4):761–765. [PubMed: 15389954]
6. Wilman A, Riederer S. Performance of an elliptical centric view order for signal enhancement and motion artifact suppression in breath-hold three-dimensional gradient echo imaging. Magnetic Resonance in Medicine. 1997; 38(5):793–802. [PubMed: 9358454]
7. Peters D, Korosec F, Grist T, Block W, Holden J, Vigen K, Mistretta C. Undersampled projection reconstruction applied to MR angiography. Magnetic Resonance in Medicine. 2000; 43(1):91101.
8. Mistretta C, Wieben O, Velikina J, Block W, Perry J, Wu Y, Johnson K, Wu Y. Highly constrained backprojection for time-resolved MRI. Magnetic Resonance in Medicine. 2006; 55(1):30–40. [PubMed: 16342275]
9. Lustig M, Donoho D, Pauly J. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine. 2007; 58(6):1182–1195. [PubMed: 17969013]
10. Trzasko, J.; Haider, C.; Manduca, A. Practical nonconvex compressive sensing reconstruction of highly-accelerated 3D parallel MR angiograms; Proc. IEEE International Symposium on Biomedical Imaging (ISBI); 2009. p. 274-277.
11. Haider C, Hu H, Campeau N, Huston J, Riederer S. 3D high temporal and spatial resolution contrast-enhanced MR angiography of the whole brain. Magnetic Resonance in Medicine. 2008; 60(3):749–760. [PubMed: 18727101]
12. Pruessmann K, Weiger M, Scheidegger M, Boesiger P. SENSE: sensitivity encoding for fast MRI. Magnetic Resonance in Medicine. 1999; 42:952–962. [PubMed: 10542355]
13. Pruessmann K, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. Magnetic Resonance inMedicine. 2001; 46(4):638–651.
14. Haider C, Glockner J, Stanson A, Riederer S. High-temporal and high-spatial-resolution three dimensional contrast-enhanced MR angiography. Radiology. 2009; 253(3):831–843. [PubMed: 19789238]
15. Haider C, Borisch E, Glockner J, Mostardi P, Rossmann P, Young P, Riederer S. Max CAPR: High-resolution 3D contrast-enhanced MR angiography with acquisition times under 5 seconds. Magnetic Resonance in Medicine. 2010; 64(6):1171–1181. [PubMed: 20715291]
16. King K, Angelos L. SENSE image quality improvement using matrix regularization. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2001:1771.
17. King K, Angelos L. SENSE with partial Fourier homodyne reconstruction. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2000:153.
18. Hu H, Madhuranthakam A, Kruger D, Glockner J, Riederer S. Combination of 2D sensitivity encoding and 2D partial Fourier techniques for improved acceleration in 3D contrast-enhanced MR angiography. Magnetic Resonance in Medicine. 2006; 55(1):16–22. [PubMed: 16342155]
19. Candès E, Romberg J, Tao T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Information Theory. 2006; 52(2):489–509.
20. Donoho D. Compressed sensing. IEEE Trans. Information Theory. 2006; 52(4):1289–1306.
21. Trzasko J, Manduca A. Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0-minimization. IEEE Trans. on Medical Imaging. 2009; 28(1):106–121.
22. Chartrand, R. Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data; Proc. IEEE International Symposium on Biomedical Imaging (ISBI); 2009.

23. Gamper U, Boesiger P, Kozerke S. Compressed sensing in dynamic MRI. Magnetic Resonance in Medicine. 2008; 59(2):365–373. [PubMed: 18228595]
24. Otazo R, Sodickson D. Distributed compressed sensing for accelerated MRI. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:378.
25. Chartrand R. Exact reconstruction of sparse signal via nonconvex minimization. IEEE Signal Processing Letters. 2007; 14(10):707–710.
26. Candès E, Wakin M, Boyd S. Enhancing sparsity by reweighted ℓ_1 minimization. J. Fourier Analysis and Applications. 2008; 14(5–6):877–905.
27. Trzasko J, Manduca A. Relaxed conditions for sparse signal recovery with general concave priors. IEEE Trans. Signal Processing. 2009; 57(11):4347–4354.
28. Riederer S, Tasciyan T, Farzaneh F, Lee J, Wright R, Herfkens R. MR fluoroscopy: technical feasibility. Magnetic Resonance in Medicine. 1988; 8(1):1–15. [PubMed: 3173063]
29. Mistretta C. Undersampled radial MR acquisition and highly constrained back projection (HYPR) reconstruction: Potential medical imaging applications in the post-Nyquist era. J. Magnetic Resonance Imaging. 2009; 29(3):501–516.
30. Raj A, Singh G, Zabih R, Kressler B, Wang Y, Schuff N, Weiner M. Bayesian parallel imaging with edge-preserving priors. Magnetic Resonance in Medicine. 2007; 57:8–21. [PubMed: 17195165]
31. Block K, Uecker M, Frahm J. Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. Magnetic Resonance in Medicine. 2007; 57:1086–1098. [PubMed: 17534903]
32. Ying L, Liu B, Steckner G, Wu M, Wu M, Li S. A statistical approach to SENSE regularization with arbitrary k-space trajectories. Magnetic Resonance in Medicine. 2008; 60(2):414–421. [PubMed: 18666100]
33. Tsao J, Boesiger P, Pruessmann K. k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. Magnetic Resonance in Medicine. 2003; 50(5):1031–1042. [PubMed: 14587014]
34. Jung H, Ye J, Kim E. Improved k-t BLAST and k-t SENSE using FOCUSS. Physics in Medicine and Biology. 2007; 52:3201–3226. [PubMed: 17505098]
35. Jung H, Sung K, Nayak K, Kim E, Ye J. k-t FOCUSS: a general compressed sensing framework for high resolution dynamic MRI. Magnetic Resonance in Medicine. 2009; 61(1):103–116. [PubMed: 19097216]
36. Otazo R, Kim D, Axel L, Sodickson D. Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. Magnetic Resonance in Medicine. 2010; 64(3):767–776. [PubMed: 20535813]
37. Borisch E, Grimm R, Rossman P, Haider C, Riederer S. Real-time high-throughput scalable MRI reconstruction via cluster computing. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:1492.
38. King K. Combining compressed sensing and parallel imaging. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:1488.
39. Kim Y, Narayanan S, Nayak K. Accelerated three-dimensional upper airway MRI using compressed sensing. Magnetic Resonance in Medicine. 2009; 61(6):1434–1440. [PubMed: 19353675]
40. Liang D, Liu B, Wang J, Ying L. Accelerating SENSE using compressed sensing. Magnetic Resonance in Medicine. 2009; 62(6):1574–1584. [PubMed: 19785017]
41. Natarajan B. Sparse approximate solutions to linear systems. SIAM J. Scientific Computing. 1995; 24(2):227–234.
42. Fischer A, Breuer F, Blaimer M, Seiberlich N, Jakob P. Introduction of a nonconvex compressed sensing algorithm for MR imaging. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:1487.
43. Schneider, G.; Prince, M.; Meaney, J.; Ho, V. Magnetic resonance angiography: techniques, indications, and practical applications. Springer-Verlag; 2005.

44. Çukar T, Lustig M, Nishimura D. Improving non-contrast-enhanced steady-state free precession angiography with compressed sensing. Magnetic Resonance in Medicine. 2009; 61(5):1122–1131. [PubMed: 19230013]
45. Wang Y, Johnston D, Breen J, Huston J III, Julsrud P, Kiely M, King B, Riederer S, Ehman R. Dynamic MR digital subtraction angiography using contrast enhancement, fast data acquisition, and complex subtraction. Magnetic Resonance in Medicine. 2006; 36(4):551–556. [PubMed: 8892206]
46. Lin F, Kwong K, Belliveau J, Wald L. Parallel imaging reconstruction using automatic regularization. Magnetic Resonance in Medicine. 2004; 51(3):559–567. [PubMed: 15004798]
47. Ying L, Xu D, Liang Z. On Tikhonov regularization for image reconstruction in parallel MRI. Proc. IEEE Engineering in Medicine and Biology Society (EMBS). 2004:1056–1059.
48. Chen G, Tang J, Leng S. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. Medical Physics. 2008; 35(3):660–663. [PubMed: 18383687]
49. Wu H, Block W, Samsonov A. HYPR-constrained compressed sensing reconstruction for accelerated time resolved imaging. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:339.
50. Samsonov A, Jung Y, Alexander A, Block W, Field A. MRI compressed sensing via sparsifying images. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2008:341.
51. Mistretta C, Grist T, Korosec F, Frayne R, Peters D, Mazaheri Y, Carroll T. 3D time-resolved contrast-enhanced MR DSA: Advantages and tradeoffs. Magnetic Resonance in Medicine. 1998; 40(4):571–581. [PubMed: 9771574]
52. Fain S, Riederer S, Bernstein M, Huston J III. Theoretical limits of spatial resolution in elliptical-centric contrast-enhanced 3D-MRA. Magnetic Resonance in Medicine. 1999; 42(6):1106–1116. [PubMed: 10571932]
53. Riederer H, Hu S, Kruger D, Haider C, Campeau N, Huston J III. Intrinsic signal amplification in the application of 2D SENSE parallel imaging to 3D contrast-enhanced elliptical centric MRA and MRV. Magnetic Resonance in Medicine. 2007; 58(5):855–864. [PubMed: 17969124]
54. Borisch E, Haider C, Grimm R, Riederer S. Subtraction in view-shared 3D contrast-enhanced MRA. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2010:2932.
55. Nocedal, J.; Wright, S. Numerical Optimization. Springer; 2006.
56. Brandwood D. A complex gradient operator and its application in adaptive array theory. IEE Proceedings F: Communications, Radar, and Signal Processing. 1983; 130(1):11–16.
57. Hayes, M. Statistical digital signal processing and modeling. Wiley; 1996.
58. van den Bos A. Complex gradient and Hessian . IEE Proceedings: Vision, Image, and Signal Processing. 1994; 141(6):380–383.
59. Vogel C, Oman M. Fast, robust total variation-based reconstruction of noisy, blurred images. IEEE Trans. Image Processing. 1998; 7(6):813–824.
60. Chan T, Zhou H, Chan R. Continuation method for total variation denoising. UCLA CAM Reports. 1995 95-28.
61. Chartrand, R.; Yin, W. Iteratively reweighted algorithms for compressive sensing; Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP); 2008.
62. Candès E. The restricted isometry property and its implications for compressed sensing. Comptes Rendus de l'Acadèmie des Sciences, Sèrie I. 2008; 346:589–592.
63. Haider T, Hulshizer C, Johnson C, Mostardi P, Rossman P, Riederer S. Modular coil array for highly-accelerated 2D parallel acquisition. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2010:1492.
64. Du Y, Parker D, Davis W, Cao G. Reduction of partial-volume artifacts with zero-filled interpolation in three-dimensional MR angiography. J. Magnetic Resonance Imaging. 1994; 4(5): 733–741.
65. Rettmann D, Saranathan M, Glockner J. Differential subsampling with cartesian ordering (DISCO): A novel k-space ordering scheme for dynamic MRI. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2010:3044.

66. Murphy M, Keutzer K, Vasanawala S, Lustig M. Clinically feasible reconstruction time for L1-SPIRiT parallel imaging and compressed sensing MRI. Proc. International Society for Magnetic Resonance in Medicine (ISMRM). 2010:4854.

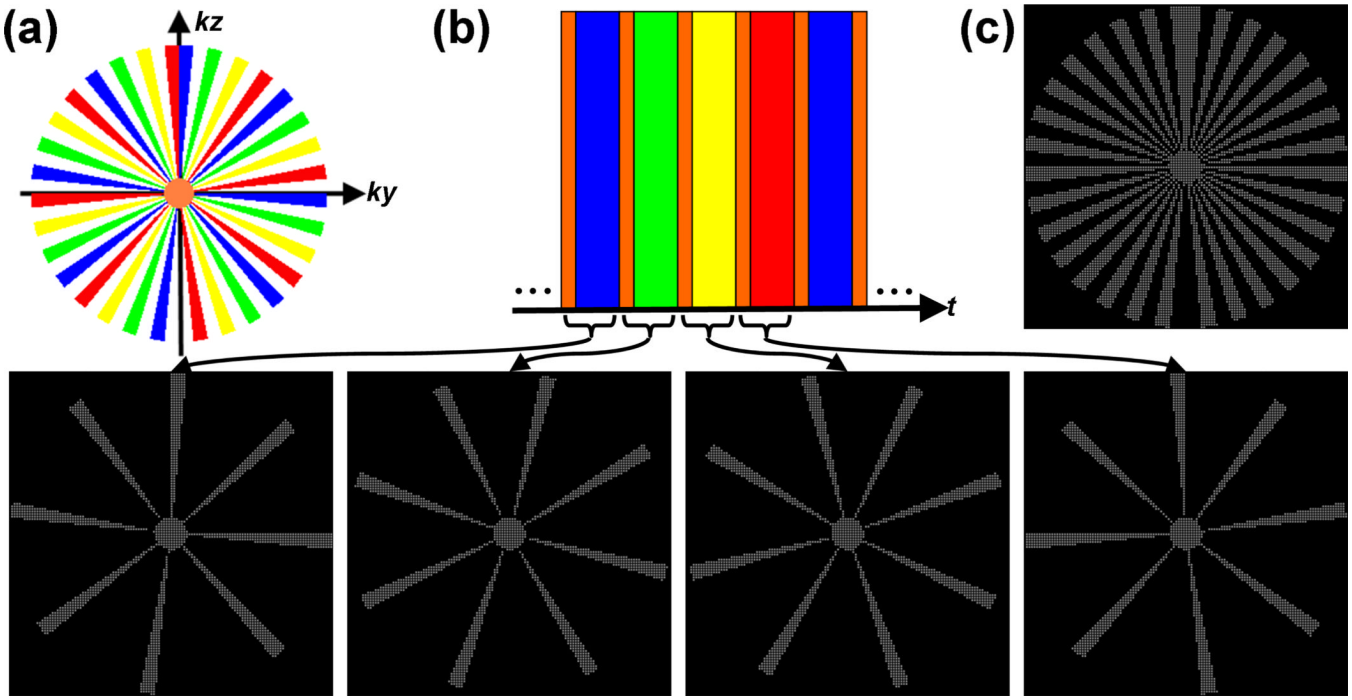

Figure 1.

An example $W = 4$ CAPR acquisition sequence ($N_y = N_z = 256$). During this acquisition, the phase-encoded plane of k-space (*ky–kz*) is partitioned into a distinct low-pass region, shown in orange, and a high-pass region (a). The high-pass annulus is itself further partitioned azimuthally into W subsets of “vanes”, shown here in blue, green, yellow, and red, which are placed asymmetrically about the origin. Although Cartesian, the CAPR sampling operator, by construction, tends to exhibit properties similar to non-Cartesian radial trajectories (within the phase-encoded plane). The sampling order or schedule is shown in (b). During each temporal update, uniformly-spaced samples (here, spaced 2 × 2 apart) from within the low-pass region and a single high-pass vane set are acquired. This same set of k-space indices is reinvestigated after *W* temporal updates. The set of all k-space indices investigated at any point during the entire exam is shown in (c), and is simply the union of the sample sets shown in (b). Note that, for this example, only about 20% of k-space is investigated during the exam, a property that facilitates practical formation of the reference signal needed for background subtraction.

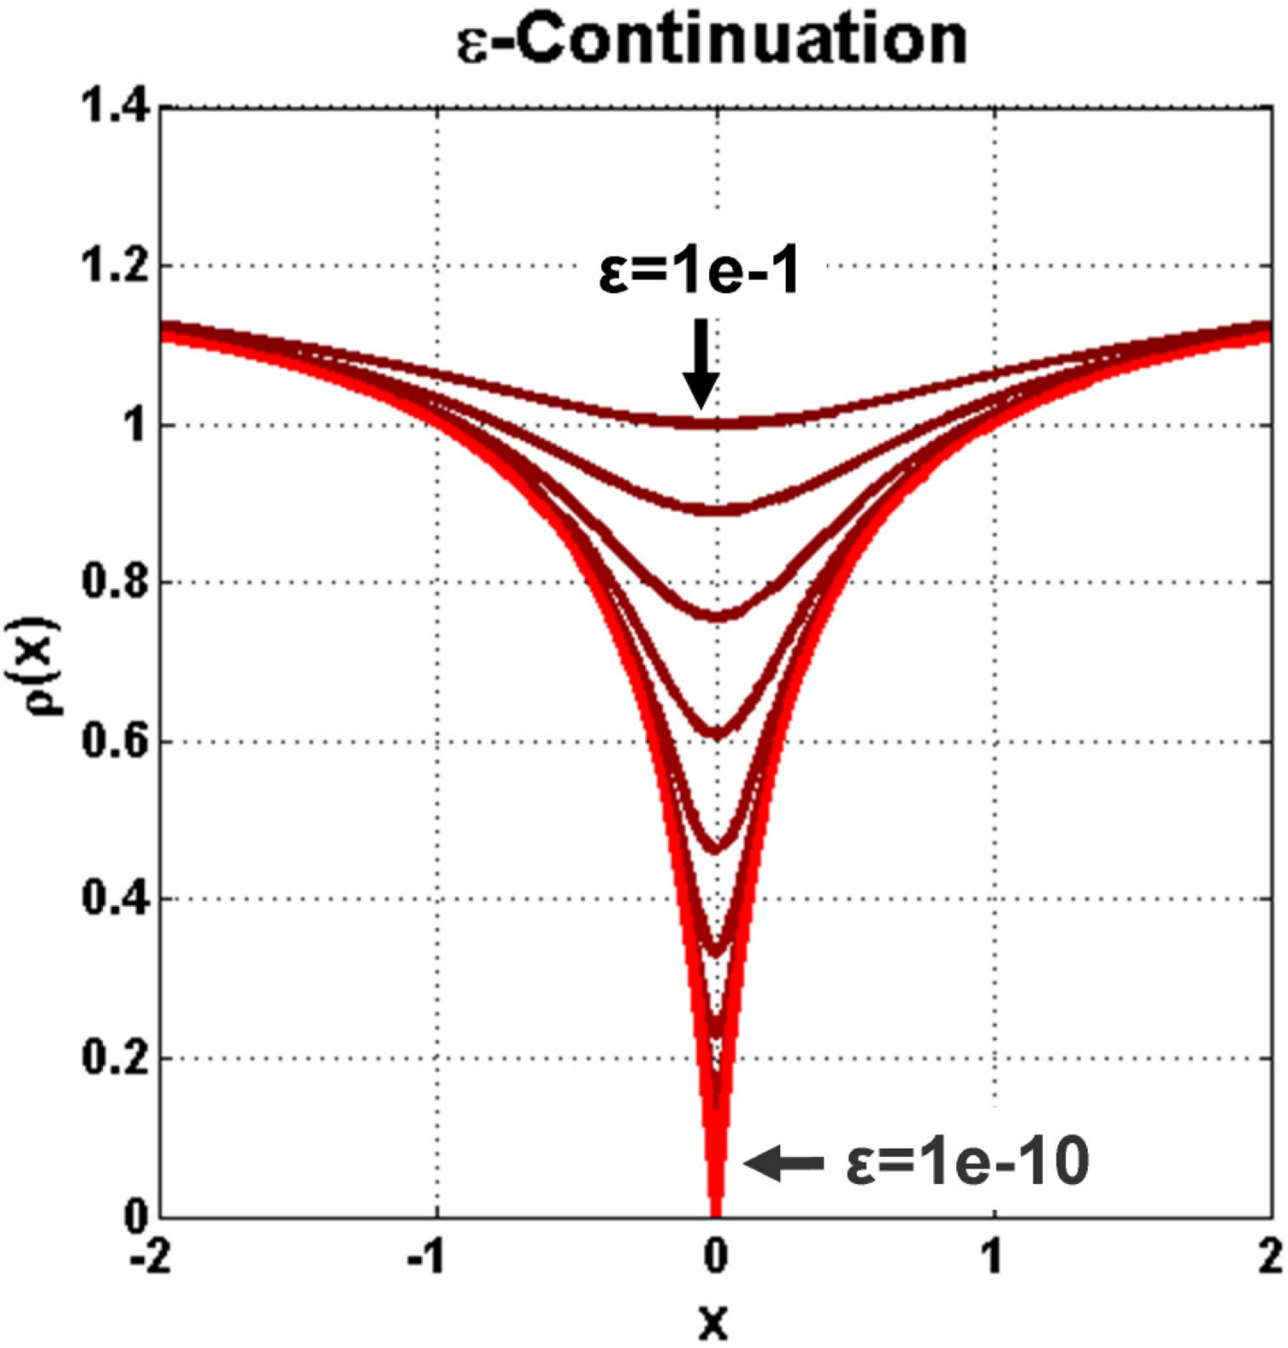

Figure 2.
Deformation of the Laplace error functional as it undergoes ε-continuation (ε → 0).

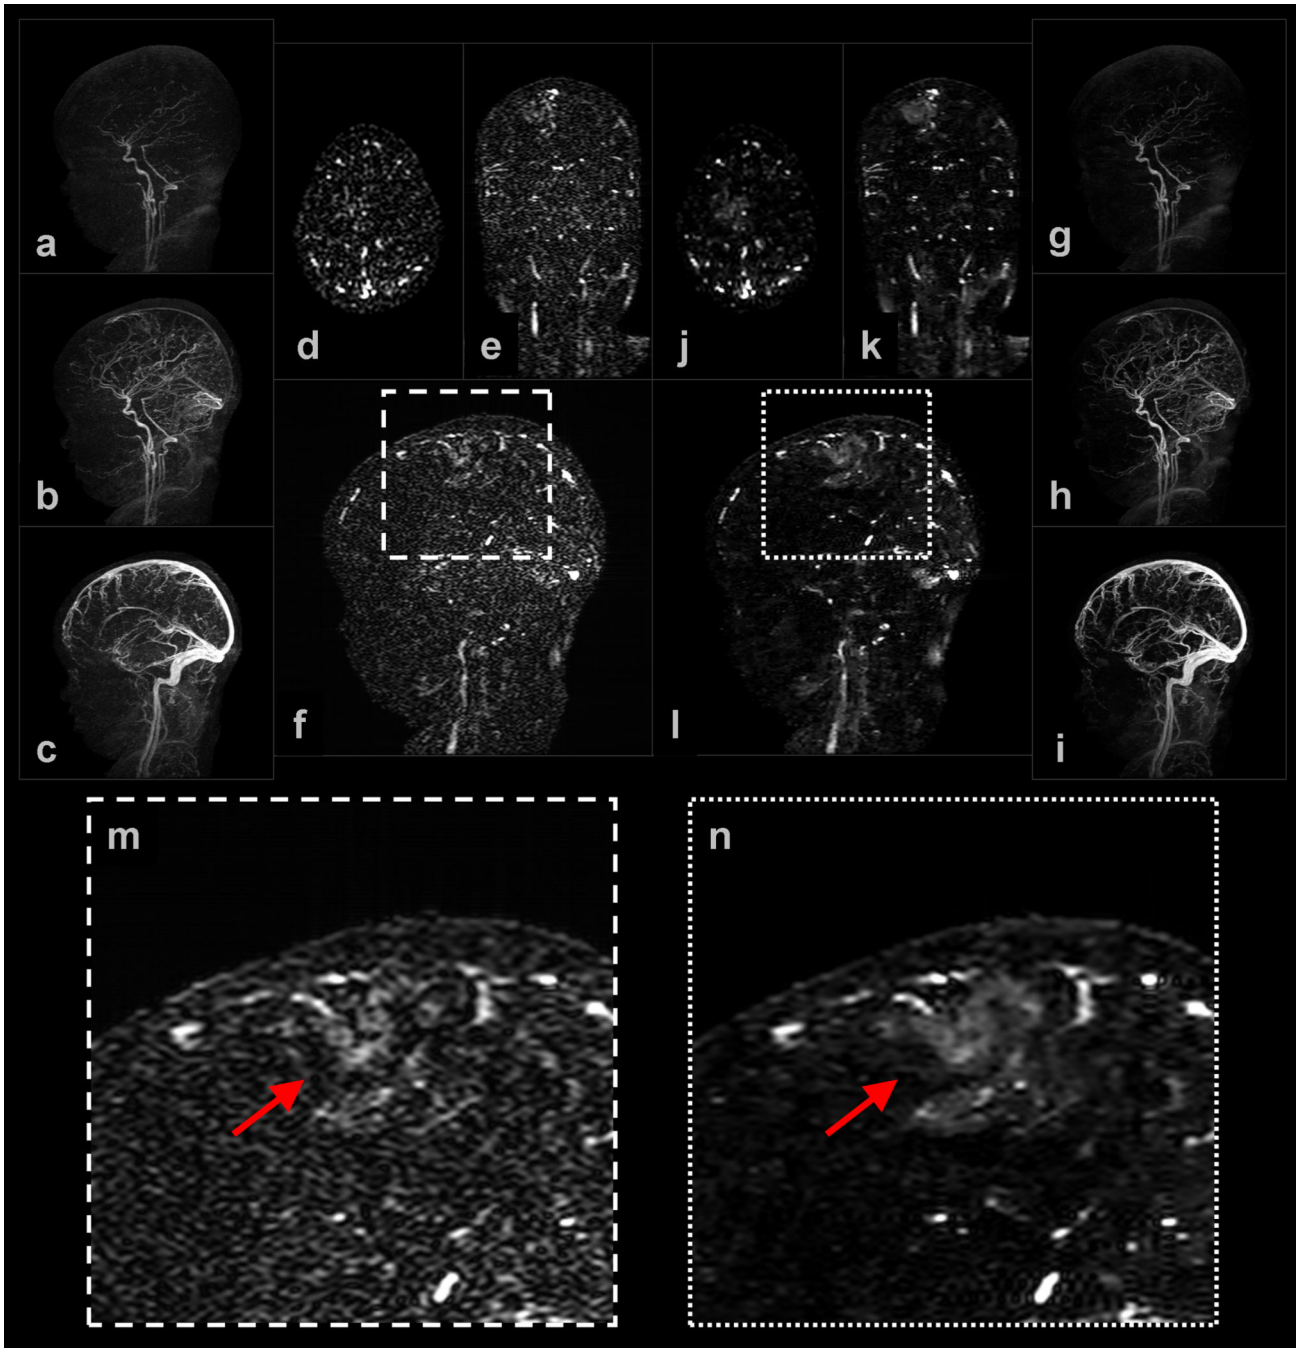

Figure 3.
Selected images from Example 1: (a–c) are sagittal MIPs of early, middle, and late filling stage volumes from a view-shared SENSE+Homodyne reconstruction of the neurovasculature of a pediatric patient, and (d–f) are axial, coronal, and sagittal slices from a late-filling stage SENSE+Homodyne volume; (g–i) and (j–l) are the corresponding images from the NCCS reconstructions. (m) and (n) are the respective enlargements of a region of interest (ROI) in (f) and (l).

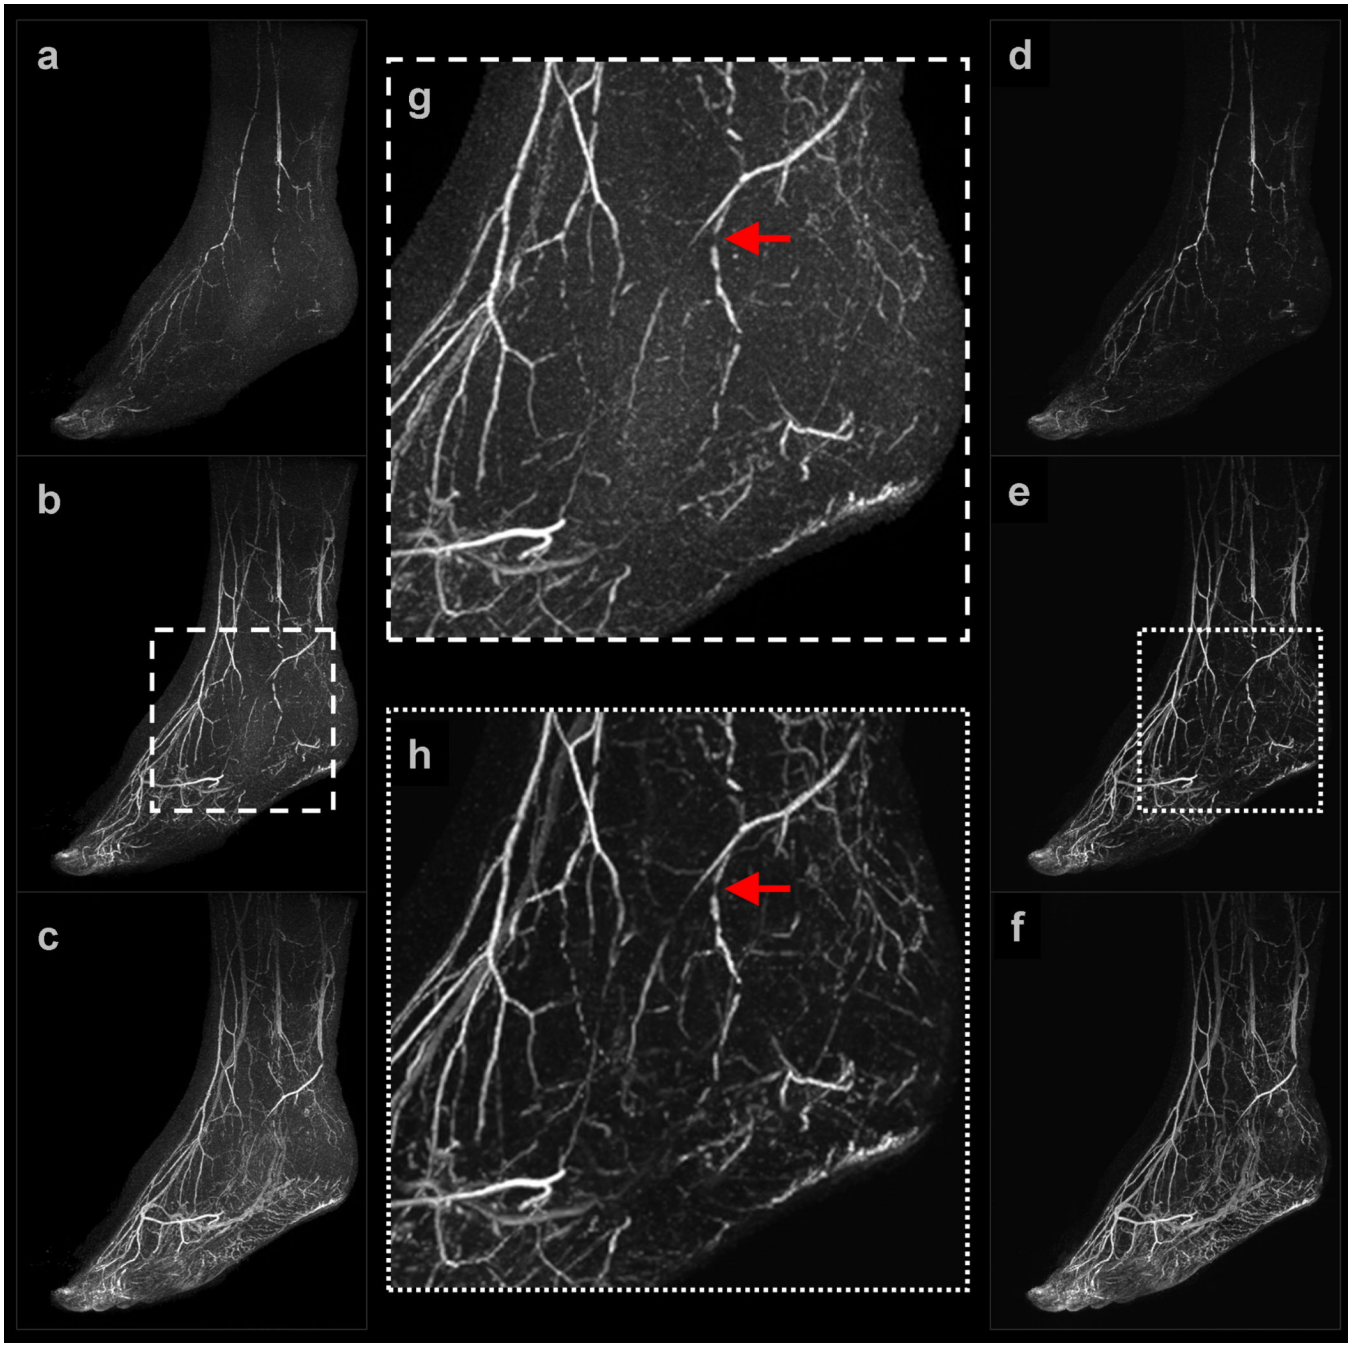

Figure 4.
Selected images from Example 2: (a–c) are sagittal maximum intensity projections (MIP) of early, middle, and late filling stage volumes from a view-shared SENSE+Homodyne reconstruction of the feet; (d–f) are from the corresponding NCCS reconstruction. (g) and (h) are the respective enlargements of an ROI in (b) and (e).

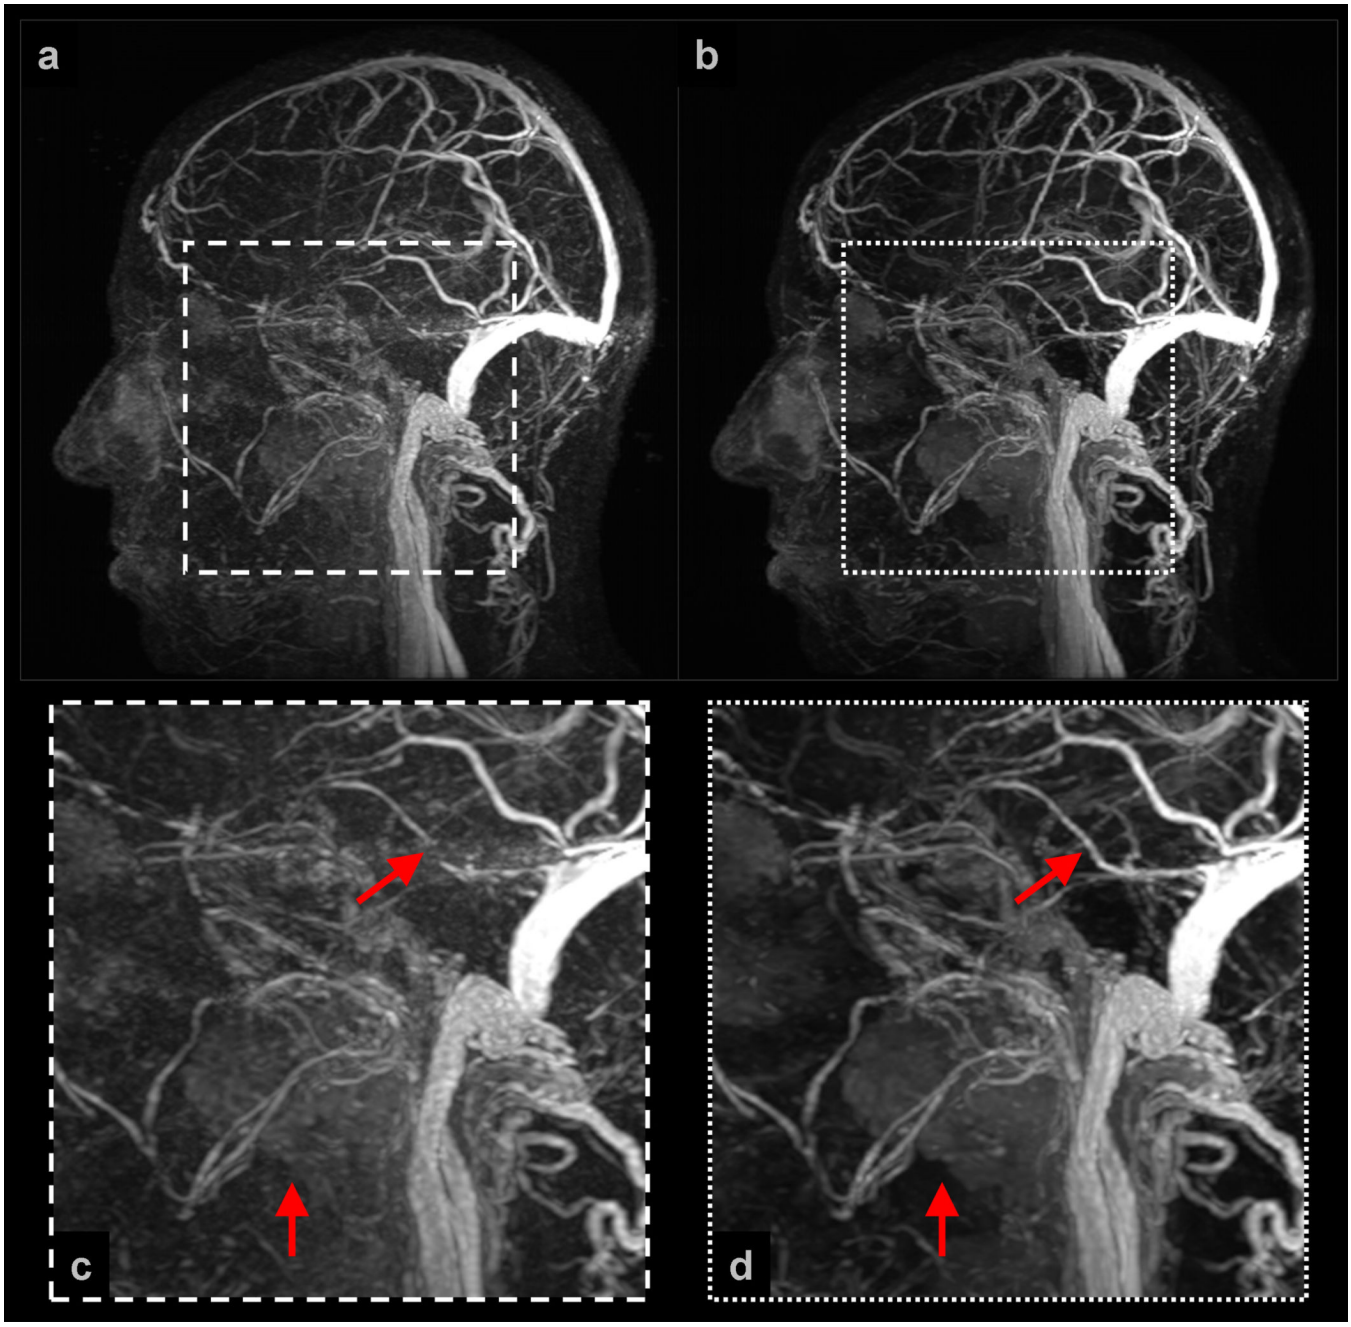

Figure 5.
Selected images from Example 3: (a) is a sagittal MIP of a late filling stage volume from a view-shared SENSE+Homodyne reconstruction of the neurovasculature of a volunteer; (b) is from the corresponding NCCS reconstruction. (c) and (d) are the respective enlargements of an ROI in (a) and (b).

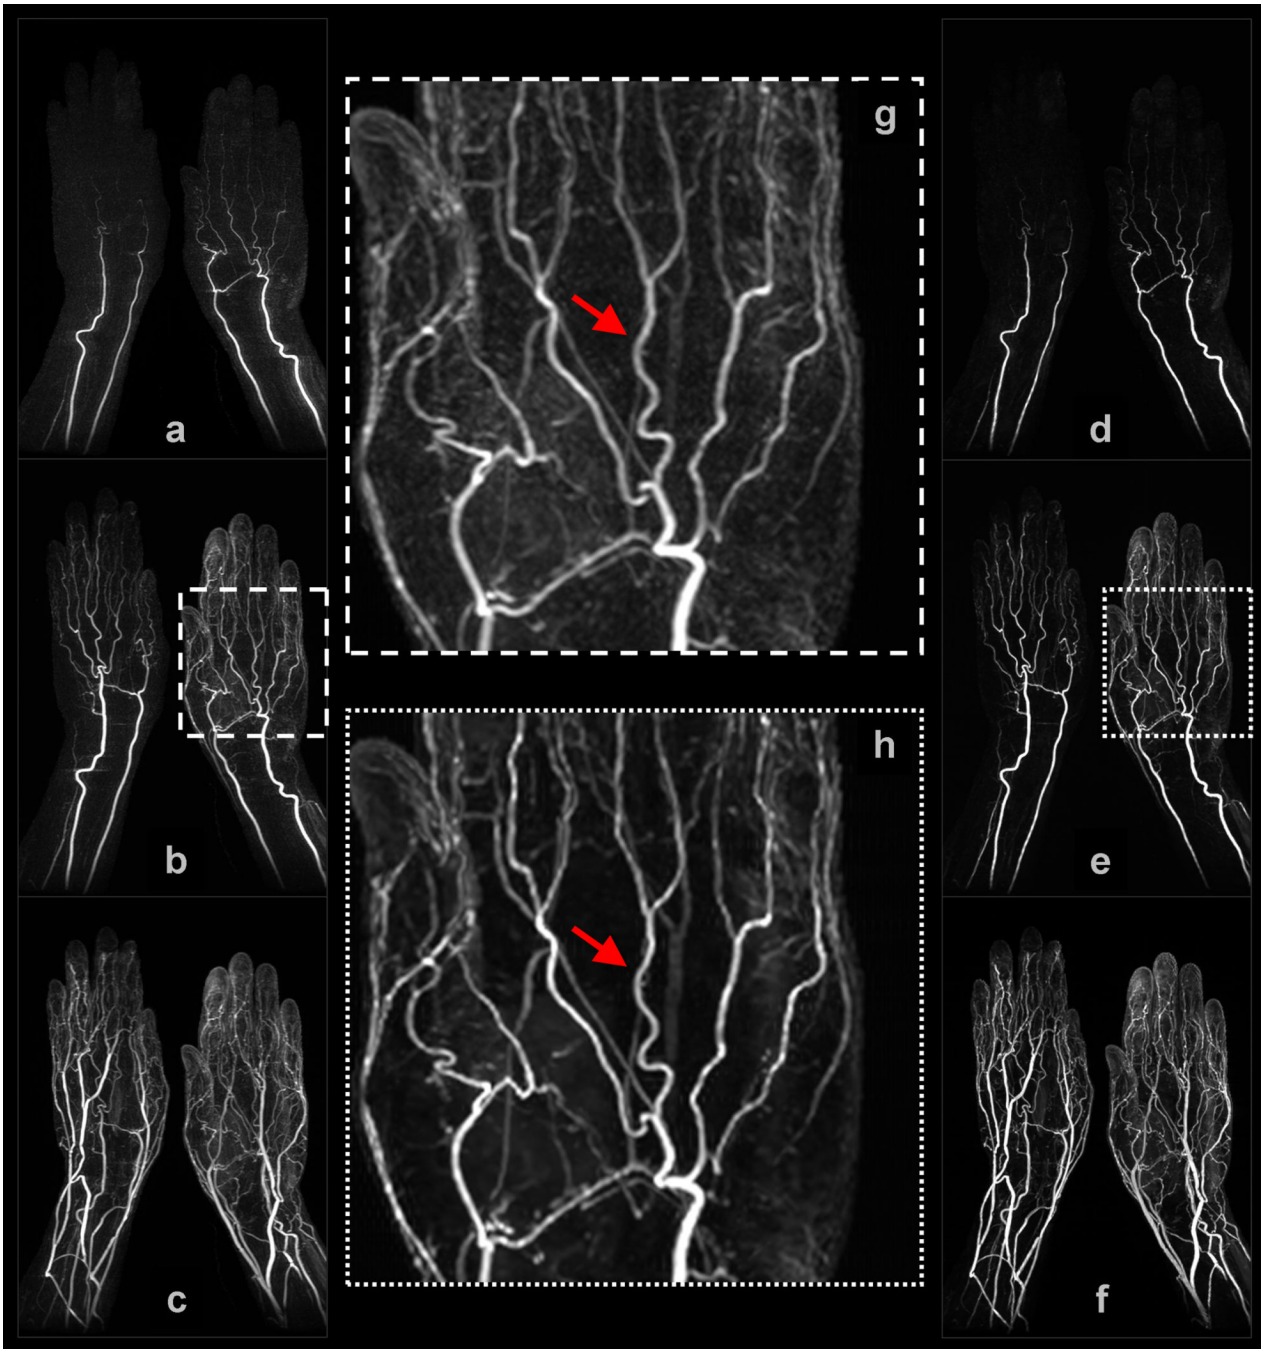

Figure 6.
Selected images from Example 4: (a–c) are coronal MIPs of early, middle, and late filling stage volumes from a non-view-shared SENSE+Homodyne reconstruction of the hands; (d–f) are from the corresponding NCCS reconstruction. (g) and (h) are the respective enlargements of a region-of-interest (ROI) in (b) and (e).

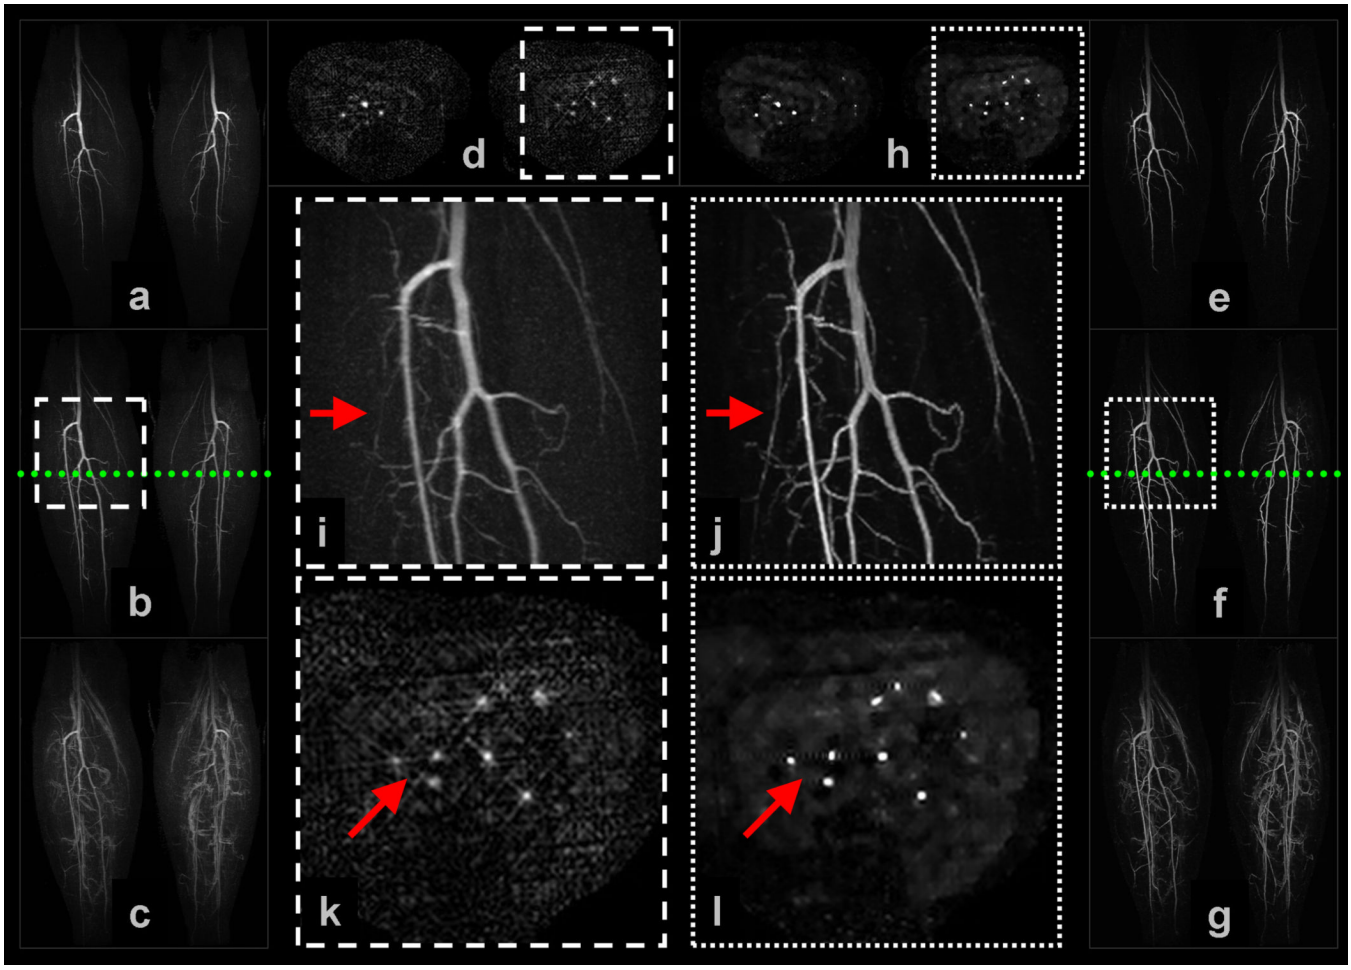

Figure 7.
Selected images from Example 5: (a–c) are coronal MIPs of early, middle, and late filling stage volumes from a non-view-shared SENSE+Homodyne reconstruction of the calves of a volunteer, and (d) is an axial cross-section of a middle filling stage volume; (e–g) and (h) are from the corresponding NCCS reconstruction. The location of the axial cross-section is marked by the dotted line in (b) and (f). (i) and (j) are the respective enlargements of an ROI in (b) and (f); (k) and (l) are the respective enlargements of an ROI in (d) and (h).

Table 1

Acquisition and reconstruction parameters for the five reconstruction examples.

Example	**1**	**2**	**3**	**4**	**5**
Anatomy	Head	Feet	Head	Hands	Legs
Software v.	20.0	14.0	14.0	14.0	15.0
TR (ms)	4.85	5.64	4.43	5.56	5.73
TE (ms)	2.25	2.62	2.02	2.58	2.66
C	8	8	14	12	8
R	4	8	8	8	8
M_{LP}	111	400	111	400	400
M_{HP}	341	795	230	145	432
N_x	256	400	256	400	400
N_y	160	320	160	280	320
N_z	72	220	100	60	132
N_t	38	27	45	26	35
W	3	4	3	4	4
Θ	3	4	3	1	1
δ_x (mm)	0.86	0.80	0.86	0.75	1.00
δ_y (mm)	1.38	0.80	1.38	0.75	1.00
δ_z (mm)	2.00	1.00	2.00	1.80	1.00
δ_t (s)	2.2	6.7	1.5	3.0	4.9
TF (s)	6.6	26.8	4.5	3.0	4.9
AF	25.49	59.81	46.92	30.83	49.58
USF (%)	21.25	59.32	29.92	61.07	83.86