A GRAPPA algorithm for arbitrary 2D/3D non-Cartesian sampling trajectories with rapid calibration

Tianrui Luo[1] | Douglas C. Noll[1] | Jeffrey A. Fessler[2] | Jon-Fredrik Nielsen[1]

[1]Department of Biomedical Engineering, University of Michigan, Ann Arbor, Michigan

[2]Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, Michigan

Correspondence
Tianrui Luo, Department of Biomedical Engineering, University of Michigan, 2360 Bonisteel Blvd, Ann Arbor, MI 48109.
Email: tianrluo@umich.edu

Purpose: GRAPPA is a popular reconstruction method for Cartesian parallel imaging, but is not easily extended to non-Cartesian sampling. We introduce a general and practical GRAPPA algorithm for arbitrary non-Cartesian imaging.

Methods: We formulate a general GRAPPA reconstruction by associating a unique kernel with each unsampled k-space location with a distinct constellation, that is, local sampling pattern. We calibrate these generalized kernels using the Fourier transform phase shift property applied to fully gridded or separately acquired Cartesian Autocalibration signal (ACS) data. To handle the resulting large number of different kernels, we introduce a fast calibration algorithm based on nonuniform FFT (NUFFT) and adoption of circulant ACS boundary conditions. We applied our method to retrospectively under-sampled rotated stack-of-stars/spirals *in vivo* datasets, and to a prospectively under-sampled rotated stack-of-spirals functional MRI acquisition with a finger-tapping task.

Results: We reconstructed all datasets without performing any trajectory-specific manual adaptation of the method. For the retrospectively under-sampled experiments, our method achieved image quality (i.e., error and g-factor maps) comparable to conjugate gradient SENSE (cg-SENSE) and SPIRiT. Functional activation maps obtained from our method were in good agreement with those obtained using cg-SENSE, but required a shorter total reconstruction time (for the whole time-series): 3 minutes (proposed) vs 15 minutes (cg-SENSE).

Conclusions: This paper introduces a general 3D non-Cartesian GRAPPA that is fast enough for practical use on today's computers. It is a direct generalization of original GRAPPA to non-Cartesian scenarios. The method should be particularly useful in dynamic imaging where a large number of frames are reconstructed from a single set of ACS data.

KEYWORDS

dynamic imaging, GRAPPA, g-factor, non-iterative reconstruction, non-cartesian imaging, NUFFT

1 | INTRODUCTION

Multi-shot 3D volumetric imaging is a potential alternative to single-shot simultaneous multislice (SMS)[1] imaging in, for example, fMRI.[2] Compared to 2D (slice-selective) acquisitions, 3D volumetric imaging is free from slice profile artifacts, has reduced spin-history effects (due to, e.g., in-flow or motion),[3,4] and can provide improved image signal-to-noise

ratio (SNR).[2,5] However, a drawback of multi-shot acquisitions is increased susceptibility to physiological signal fluctuations between shots. To mitigate physiological noise artifacts, it is beneficial to accelerate the acquisition using parallel imaging (PI) such that the total acquisition time per volume is reduced (to, e.g., one-half the heartbeat interval[6,7]). In addition, non-Cartesian readout trajectories are desirable due to their high sampling efficiency.

Iterative parallel imaging (PI) methods such as cg-SENSE[8,9] and SPIRiT[10] are often chosen due to their ability to handle arbitrary readout trajectories. However, in applications such as fMRI where a time-series of several hundred images must be reconstructed, the total reconstruction time using iterative methods can become prohibitive (since each image is reconstructed independently). To speed up reconstruction, several groups have adopted dedicated computational hardware, for example, general purpose GPUs,[11] but working with such hardware usually requires specialized programming expertise making such methods less widely transferable. Therefore, there is currently a need for an easy-to-implement and robust reconstruction technique for arbitrary non-Cartesian trajectories that scales well with the total number of image frames (i.e., short overall reconstruction time for a time-series) even on common hardware (e.g., CPUs).

The foremost non-iterative alternative to the above-mentioned techniques is GRAPPA,[12-14] which (like SPIRiT) does not require explicit knowledge of receive coil sensitivity maps. Unfortunately, using GRAPPA with non-Cartesian trajectories has so far been somewhat awkward. Incomplete remedies have been explored for certain kinds of readouts. For example,[15,16] splits stack-of-stars or stack-of-spirals readouts into segments, and then performs conventional Cartesian GRAPPA within each segment. However, this procrustean approach cannot be used with arbitrary 3D non-Cartesian trajectories, and requires a large amount of ACS data for kernel estimation. In addition, a trade-off must be made between trajectory segmentation and reconstruction quality. Alternatively, one can grid a Cartesian dataset from a non-Cartesian one and apply ordinary GRAPPA thereafter (based on GROG[17,18]). However, reliable and accurate gridding may be problematic for under-sampled trajectories due to its 1D interpolation nature that limits its effectiveness for interpolation beyond 1 k-space sample distance. Another approach that has been used in some applications is through-time GRAPPA,[19] which alternates sampling patterns along time and extends the reconstruction along the temporal dimension. However, its calibration may become ill-conditioned when dynamic imaging contrast is predominantly stationary, for example, as in fMRI. Finally, PARS[20] and kSPA[21] are non-iterative k-space reconstruction methods that are closely related to GRAPPA, however, they both rely on explicit knowledge of receive coil sensitivity maps. PARS needs to calibrate a large number of kernels which can be time-consuming in practice, while kSPA requires extra tuning of the order of its polynomial approximations. Neither of the approaches has become widely adopted.

We propose a conceptually simple method for generalizing GRAPPA to arbitrary 3D non-Cartesian PI acquisitions, and provide an efficient algorithm for its calibration.[22] For each unsampled k-space location (with a distinct local sampling constellation), our method assigns a unique GRAPPA kernel, whose calibration is efficiently implemented by utilizing the phase-shift property of the Fourier Transform (FT) and the NUFFT.[23,24] Like Cartesian GRAPPA, our method does not require explicit coil sensitivity information, and reconstruction per image volume (once weights have been calibrated) is rapid. Apart from choice of GRAPPA kernel size and Tikhonov regularization coefficient (which is also typically used in conventional GRAPPA), our proposed method is fully automatic and does not require manual parameter selection based on, for example, segments or other trajectory-dependent aspects. Thus, our method (once coded) requires minimal user expertise and should be broadly applicable to arbitrary non-Cartesian PI applications. We demonstrate our method in 3D rotated stack-of-stars[25] and rotated stack-of-spirals[26] structural (T$_1$-weighted) imaging, and 3D rotated stack-of-spirals fMRI, using the same implementation parameters for all 3 cases.

2 | METHODS

2.1 | Background: General principle of GRAPPA

It is helpful to consider 2 aspects of the GRAPPA kernel: the local sampling *constellation*, and the associated *weights*. The constellation captures the relative positions between the unsampled ("center") location that we wish to reconstruct and its sampled local neighbors, and the weights are the coefficients for later reconstruction. Depending on the sampling trajectory, there are usually many different constellations in 1 dataset, and thus multiple sets of weights will be required for the reconstruction.

GRAPPA works by first identifying all distinct constellations. Then, for each constellation, to solve for its weights, GRAPPA collects all combinations of data from the ACS region whose relative positions match that of the constellation. Figure 1B shows a simple 2-neighbor (upper-left and lower-right) illustrative example for reconstructing an unsampled (center) location. For each combination, the neighbors across N_c coils that surround its center form 1 row of matrix $\tilde{A} = [\tilde{A}_1, \ldots, \tilde{A}_{N_c}]$, where the tilde indicates that these signals reside in k-space. The submatrix $\tilde{A}_c \in \mathbb{C}^{N_k \times N_n}$ collects the neighbors from the c^{th} coil. Here N_k is the number of combinations (within the ACS region) one can collect that match the constellation being calibrated, and N_n is the number of neighbors inside the constellation.
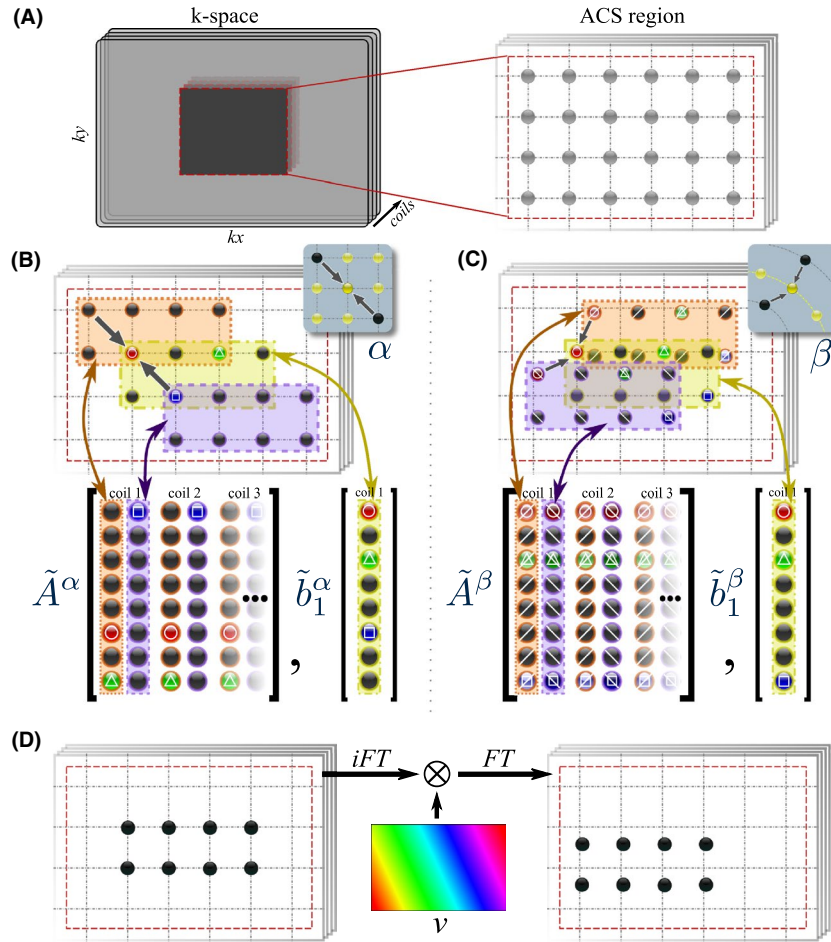
FIGURE 1 Comparison of (B) conventional Cartesian GRAPPA, and (C,D) the proposed non-Cartesian GRAPPA calibration. In both methods, the ACS region consists of densely (Nyquist) sampled (or gridded) Cartesian data, illustrated in (A). (B) Conventional (Cartesian) GRAPPA reconstruction using the constellation "α", that is composed of 2 sampled neighbors and an unsampled center. Several "combinations" are identified from the ACS region that match the desired constellation (α); in (B), there are 8 such combinations, whose center points form the yellow rectangle in (B). Each combination gives rise to 1 row in $\tilde{A}$, and 1 element in $\tilde{b}_c$ (cf. Equation (1)). As a result, the orange and purple rectangles form columns of $\tilde{A}$ (after vectorizing), and $\tilde{b}_c$ is formed from the yellow rectangle. (C,D) Proposed non-Cartesian GRAPPA reconstruction. For a non-Cartesian (local) constellation "β", we synthesize off-grid ACS data using the phase-shift property of the Fourier transform (d). The GRAPPA weights for this constellation are then obtained as in conventional GRAPPA (cf. Equation (1))

In addition, $\tilde{b}_c$ denotes the vector of center values from the c^{th} coil. The weights are obtained by solving the least square (LS) problem

$$\underset{w_c}{\text{argmin}} \, \|\tilde{A}w_c - \tilde{b}_c\|_2^2 + \lambda \|w_c\|_2^2 \Rightarrow w_c^\star = (\tilde{A}^H \tilde{A} + \lambda I)^{-1} \tilde{A}^H \tilde{b}_c, \quad (1)$$

where $w_c \in \mathbb{C}^{N_c \cdot N_n}$ denotes the vector of weights for reconstructing an unsampled value for the c^{th} coil, and λ is the Tikhonov regularization coefficient.

When forming $\tilde{A}$ and $\tilde{b}_c$ in this way, the columns of $\tilde{A}$ contain the (vectorized) signals from partially overlapping rectangles, for example, as shown in Figure 1B. In the limit of a very large ACS region, these rectangles resemble shifted replicas of each other, apart from differences at the boundary.

2.2 | Proposed Non-Cartesian GRAPPA

The procedure described above is well suited for Cartesian acquisitions, where all locations, whether sampled or not, are on a Cartesian grid. Hence, for any peripheral constellation, one can always find matching combinations within the ACS region, and arrange them into the LS form Equation (1). For non-Cartesian sampling, given a fixed size, constellations can be efficiently identified using kd-tree[27]: Specifically, this algorithm forms a tree-structure given all k-space points of interest; It returns all neighboring points within a given distance (i.e., prescribed GRAPPA kernel-size) when queried with a certain (unsampled) location. Following constellation identification, the calibration can no longer be directly conducted in the conventional way. However, while there

may be no direct match from the "on-grid" (Cartesian) ACS region for an off-grid non-Cartesian GRAPPA constellation, we propose to obtain such matches through the phase-shift property of the discrete Fourier transform (DFT). We first inverse FT the on-grid ACS data, and then modulate the resulting image by a linear phase corresponding to the wanted off-grid shift. This phase-modulated image is then Fourier transformed back to k-space. The whole procedure is equivalent to a periodic sinc interpolation. A simple example of our approach is illustrated in Figure 1C, for a typical constellation β from a spiral readout, where 2 off-grid (black) neighbors are assumed sampled and selected to reconstruct the unsampled center. The key idea is to "synthesize" the corresponding non-Cartesian (off-grid) combinations from an on-grid ACS dataset. These interpolated ACS data are again arranged as in Equation (1), yielding the associated GRAPPA weights.

Following the calibration phase, we use the GRAPPA weights (1 set for each unique constellation) to reconstruct unsampled k-space coil data and restore a full non-Cartesian k-space. We then reconstruct coil images using the NUFFT adjoint method with Voronoi density compensation.[23,28] Final images are obtained using sum-of-squares[12] or linear[29] coil image combinations.

2.3 | Efficient implementation of the proposed method

Our non-Cartesian approach must calibrate a different set of GRAPPA weights for every distinct constellation. In the most general case, the number of distinct constellations equals the number of unsampled points. The workload for many applications may be less, since commonly used trajectories, such as rotated stack-of-stars, stack-of-spirals, and stack-of-cones,[30] often possess certain k-space regularities (e.g., regular sampling along k_z) that reduce the number of distinct constellations that must be calibrated. Nevertheless, depending on acquisition parameters, such as image matrix size and undersampling factor, calibration can still be computationally demanding or even impractical. To address this, we propose the following approximate algorithm for efficient implementation of our method. Specifically, we accelerate calibration by avoiding any explicit interpolation to ACS datasets.

In the formation of $\tilde{A}$ and $\tilde{b}_c$ in Equation (1), the signals at the boundary of the ACS region are normally excluded, as forming matching combinations for them would require signals residing outside the region. In our fast algorithm, boundaries are included, by adopting circulant boundary conditions. In other words, we allow the ACS region to wrap around as needed to complete a given constellation. The underlying assumption of this approach is that when the ACS region is sufficiently large, the impact of these wrap-around combinations is marginal. With the circulant boundary assumption, the rectangles described above are now fully overlapping in

a circularly-shifted manner, and the columns formed from them, that assemble the matrix $\tilde{A}_c$ and the center vector $\tilde{b}_c$, are now circularly shifted replicates. This allows efficient calibration as follows.

Figure 2 illustrates our algorithm for efficient weight calibration. We denote unitary $F \in \mathbb{C}^{N_k \times N_k}$ as the (2D or 3D) DFT of the ACS region size. The solution to (1) can now be written as:

$$\begin{aligned} w_c^\star &= (\tilde{A}^H \tilde{A} + \lambda I)^{-1} \tilde{A}^H \tilde{b}_c \\ &= (A^H F^H F A + \lambda I)^{-1} A^H F^H F b_c = (A^H A + \lambda I)^{-1} A^H b_c, \end{aligned} \quad (2)$$

where $A = [A_1, \ldots, A_{N_c}]$ and b_c are in the image domain (tilde symbol removed). In particular, b_c is a low-resolution coil image obtained by inverse FT of all of the ACS data. Then, we represent $A^H A$ and $A^H b_c$ with block matrices:

$$A^H A = \begin{pmatrix} A_1^H A_1 & \cdots & A_1^H A_{N_c} \\ \vdots & \ddots & \vdots \\ A_{N_c}^H A_1 & \cdots & A_{N_c}^H A_{N_c} \end{pmatrix}, \quad A^H b_c = \begin{pmatrix} A_1^H b_c \\ \vdots \\ A_{N_c}^H b_c \end{pmatrix}. \quad (3)$$

Due to the circulant attribute, the columns in A_c are essentially b_c modulated by different linear phases, that is, $\text{diag}(v)b_c$, where $v \in \mathbb{C}^{N_k}$ is a phase vector. Analytically, these phase vectors are formed following the Fourier relation. That is, for a k-space neighbor at position p relative to its constellation center as origin, the qth element in its corresponding v, which modulates image domain location q, is $[v]_q = \exp(1i \cdot 2\pi p^T q)$. Let $m, n = 1, \ldots, N_n$. In MATLAB notation:

$$\begin{aligned} A_i(:,m)^H A_j(:,n) &= b_i^H \text{diag}(v_m)^H \text{diag}(v_n) b_j \\ &= (v_m^H \text{diag}(v_n))(\text{diag}(b_i)^H b_j) \quad (4) \\ A_i(:,m)^H b_j &= (v_m^H)(\text{diag}(b_j)^H b_j), \end{aligned}$$

where $i, j = 1, \ldots, N_c$ are coil indices. In other words. the elements of $A_i^H A_j$ and $A_i^H b_j$ are inner products between various linear phases and fixed (low-resolution) "coil product" images, that is, these elements are (spatial) frequency components of $\text{diag}(b_i)^H b_j$. Analytically, that is, if for any image domain location q, the element, $[v_m]_q$, can be expressed as $\exp(1i \cdot 2\pi p^T q)$ with the same p, then $A_i(:,m)^H b_j = [F \text{diag}(b_j)^H b_j]_p$, the frequency p component. An important property of this formalism is that b_c, and hence all coil product images, are shared across all constellations; therefore we pre-compute and cache the FT of $\text{diag}(b_i)^H b_j$. For non-Cartesian constellation calibration involving off-grid frequency components of $\text{diag}(b_i)^H b_j$, we prepare and cache the dense frequency spectrum of $\text{diag}(b_i)^H b_j$ by zero-padding and Fourier transforming (which only needs to be done once), and then linearly interpolate to the desired off-grid frequency (as commonly done in NUFFT[23]). This algorithm reduces the complexity of computing $A^H A$ and $A^H b_c$ from $\Theta(N_k N_n^2 N_c^2 + N_k N_n N_c)$ to $\Theta(N_n^2 N_c^2 + N_n N_c)$, where N_k can reach several thousand (e.g., for a 3D ACS region of size $20 \times 20 \times 20$).
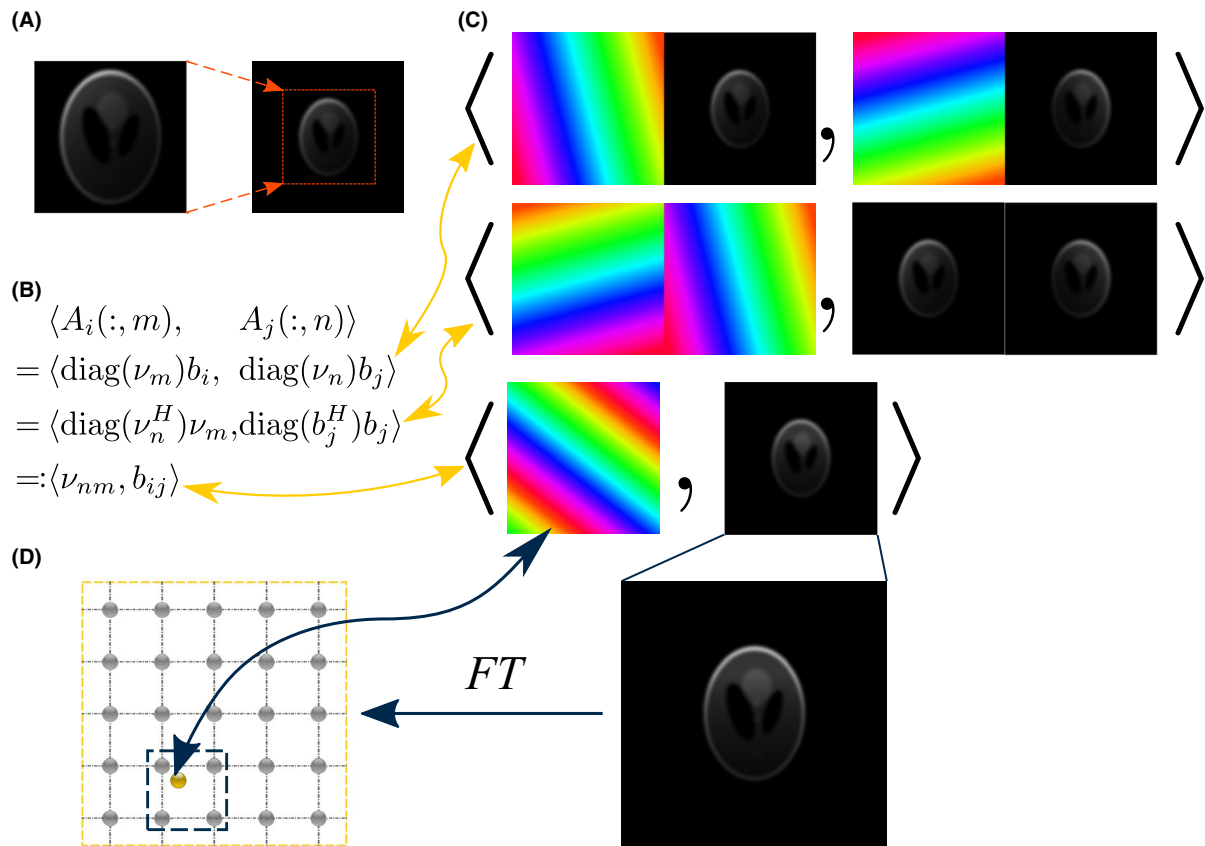
FIGURE 2 Illustration of our method for fast computation of $A^H A$ (Equation (2)). (A) A (low-resolution) coil image is reconstructed from the ACS data, and this image is then zero-padded to increase the spatial frequency (Fourier) sampling density. (B,C) The (m, n)th element of $A_i^H A_j$ is a pixel-wise multiplication of coil images b_i and b_j, each modulated by a linear phase v_m and v_n, respectively. The result is a single "coil-product" image b_{ij} modulated by the combined phase image v_{nm}. (D) The frequency component corresponding to the combined linear phase is efficiently calculated by interpolating the Fourier transform of the zero-padded coil-product images, in NUFFT-like fashion

We implemented the proposed calibration in C as it requires efficient indexing for NUFFT interpolation. Compared to the MATLAB inner-product approach (Equation (4)), our NUFFT interpolation-based approach can speed up the calibration about 20-fold (results not shown). Once the calibration is done, the non-iterative reconstruction stage for each image frame requires only a few seconds, as opposed to minutes using iterative methods.

2.4 | Algorithm and implementation details

The proposed non-Cartesian GRAPPA reconstruction procedure comprises the following list:

2.4.1 | Calibration

1. Prepare a Cartesian ACS dataset: This can be from either direct Cartesian acquisition at the center of the k-space, or from gridding a Nyquist sampled non-Cartesian k-space center. An ACS dataset of size $20 \times 20 \times 20 \times N_c$ has been found sufficiently large in this work. For consistent

Tikhonov regularization behavior, normalize the ACS dataset by the mean square root energy across coils.

2. Identify all the constellations for calibration: From the complete readout trajectories, build a kd-tree including all sampled k-space locations. For each un-sampled k-space point, query the kd-tree with the prescribed distance (i.e., GRAPPA kernel size). This will identify all the neighbors of each center, and constellations can then be formed with the relative positions between neighbors and centers. For the un-sampled points located near the center of (the highly oversampled) k-space, the signals from close neighbors are nearly linearly dependent, and can cause ill-conditioning in calibration when Tikhonov coefficient is not chosen carefully. It is hence helpful to sift the crowded neighbors, that is, by grouping the neighbors by Nyquist Δk rounded relative positions, and arbitrarily picking 1 from each group to keep as a neighbor. By sifting the neighbors such that the relative distances of the remaining neighbors become sufficiently large, we avoid any linear-dependence issue in calibration. The unique constellations identified in this step are to be calibrated. The number of
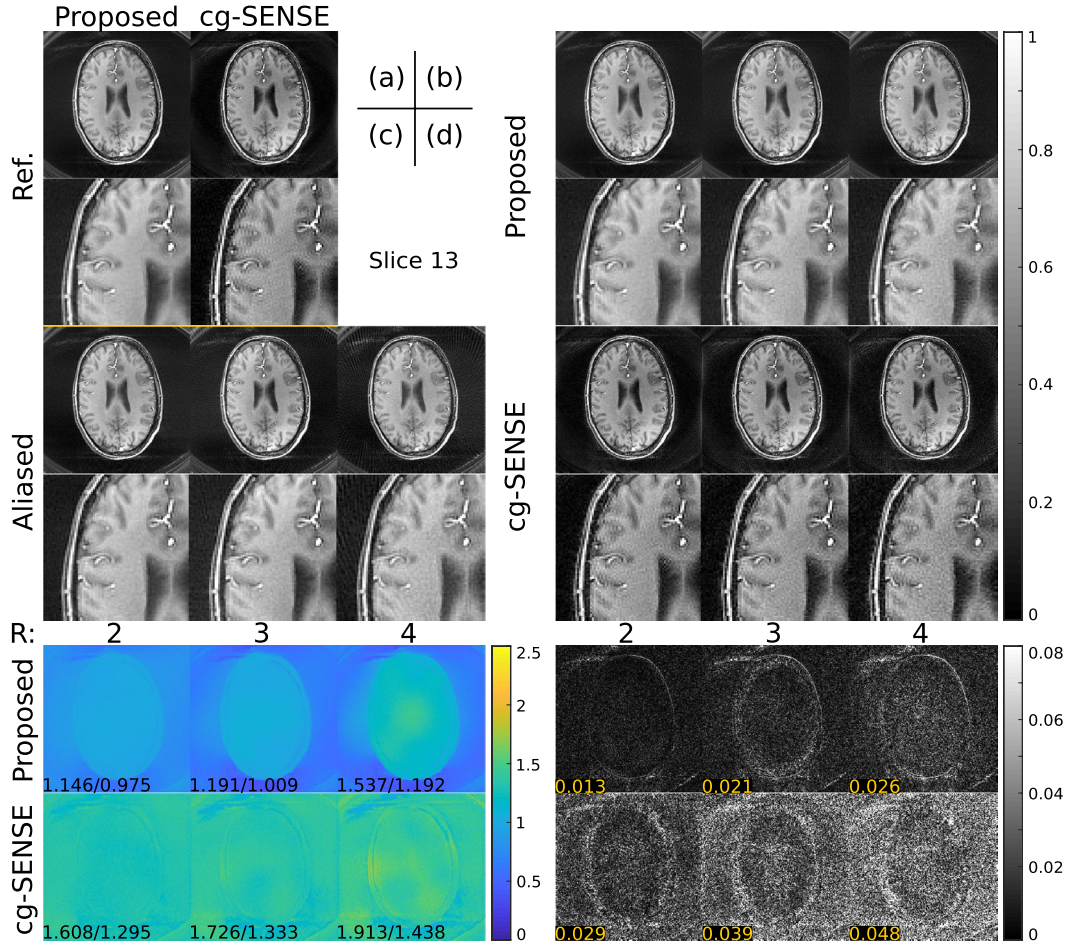
FIGURE 3 3D reconstruction comparison between the proposed non-Cartesian GRAPPA and cg-SENSE, from retrospectively under-sampled rotated stack-of-stars dataset. One out of a total 20 slices are shown. Coil images reconstructed with the proposed method are linearly combined into final images. (a) (fully sampled) reference images and under-sampled aliased images (acceleration factor $R = 2, 3, 4$) with zoom-in details. The reference for the proposed method is reconstructed with NUFFT adjoint.[23,28] (B) reconstructed under-sampled images of our proposed method and cg-SENSE. (C) g-factor maps of the 2 methods. The reported numbers are "max/mean" g-factors within the whole 3D head. (D) error map of the 2 methods. The reported numbers are "root mean squared error (RMSE)" within the whole 3D head. Our non-Cartesian GRAPPA algorithm rivals cg-SENSE in reconstruction quality. As expected, reconstruction error (gray-scale windowed 12.5x) is largest near object edges, since there are fewer samples available for reconstruction at higher spatial frequencies. The g-factor maps for the 2 methods are within a similar range, as observed in previous work[29]

neighboring points in each constellations varies from a few points (e.g., 3), when it is in the peripheral k-space, to $5 \times 5 \times 5 = 125$ capped by sifting with kernel size of 5, when it is in the center k-space.

3. Maintain a table of un-sampled points and the constellations they correspond to. With the recorded relative positions, GRAPPA coefficients are calibrated by the algorithm described in Section 2.3. With a normalized ACS, Tikhonov coefficient $\lambda = 5 \times 10^{-7}$ is used in all experiments in the following sections, which was found to produce consistent reconstruction quality.

2.4.2 | Reconstruction

1. Index the un-sampled points from the under-sampled dataset by the table maintained in the calibration step. For each un-sampled point, extract its sampled neighbors

in the constellation it corresponds to, and reconstruct the missing signal using the calibrated coefficients.

2. Use NUFFT adjoint method to resolve the reconstructed coil images, and combine them using either sum-of-squares or linear combinations to yield the final images.

Matlab and C code are opensourced on GitHub[a].

2.5 | Experiments

To evaluate our proposed algorithm from Section 2.3 in terms of reconstruction time and image quality, we first acquired 3D fully sampled stack-of-stars and stack-of-spirals spoiled gradient echo (GRE) datasets in healthy volunteers on a GE 3T scanner with an 8-channel receive-only head coil. Both the fully sampled stack-of-stars and stack-of-spirals readouts contained 20 k_z platters. For the stack-of-stars dataset, each platter

contained 315 spokes. Its image matrix size and field-of-view (FOV) were $200 \times 200 \times 20$ and $24 \times 24 \times 10\,\mathrm{cm}^3$, respectively. For the stack-of-spirals dataset, each platter contained 12 interleaves. The image matrix size and FOV were $240 \times 240 \times 20$ and $24 \times 24 \times 10\,\mathrm{cm}^3$, respectively. 8 out of 20 slices are aliased in z direction, due to the excitation profile being slightly larger than FOV_z. Other acquisition parameters were: TR = 15 ms, minimum TE, and flip-angle 8°.

We retrospectively under-sampled these 2 non-Cartesian acquisitions. Spokes and interleaves in-plane were regularly skipped to simulate different acceleration factors. For 3D reconstruction experiments, in the through-plane direction, we rotated the under-sampling pattern to improve reconstruction quality.[25,26] This rotation produces a through-plane under-sampling factor that is separately the same as the in-plane and the overall under-sampling factor. While this rotation may complicate existing non-iterative methods (e.g.,[26]), it does not impact our proposed non-Cartesian method (from the algorithmic viewpoint; it does of course impact the constellations that will be identified for a given kernel size). For 2D reconstruction experiments, prior to the

in-plane regular under-sampling, an inverse Fourier transform in k_z is applied to convert the 3D k-space into 2D.

To compare with iterative methods (cg-SENSE), and evaluate the feasibility of reducing total reconstruction time for a non-Cartesian time-series acquisition, we acquired a prospectively 3D under-sampled (acceleration factor $R = 3$) multi-interleaf rotated stack-of-spiral-in fMRI dataset of the motor cortex, with finger tapping stimulus. The readout has 20 k_z platters, each containing 9 interleaves (under-sampled to 3 interleaves). Other sequence parameters are: TR = 36.2 ms, TE = 26.3 ms, flip-angle 8°, image matrix size $88 \times 88 \times 20$ with 131 temporal frames, FOV $22 \times 22 \times 6\,\mathrm{cm}^3$. In this case, the excitation bandwidth was set to match FOV^z. In[13], it was suggested that GRAPPA kernels can be pre-calibrated using a separate dataset with possibly different contrasts. Accordingly, in the subject undergoing fMRI we also acquired a fully sampled stack-of-spiral-out dataset with the same FOV as the fMRI acquisition but with a different contrast: TR = 30 ms, minimum TE, flip-angle 8°, image matrix size $220 \times 220 \times 20$. Again, the experiments were conducted on a GE 3T scanner using an 8-channel receive-only head coil.
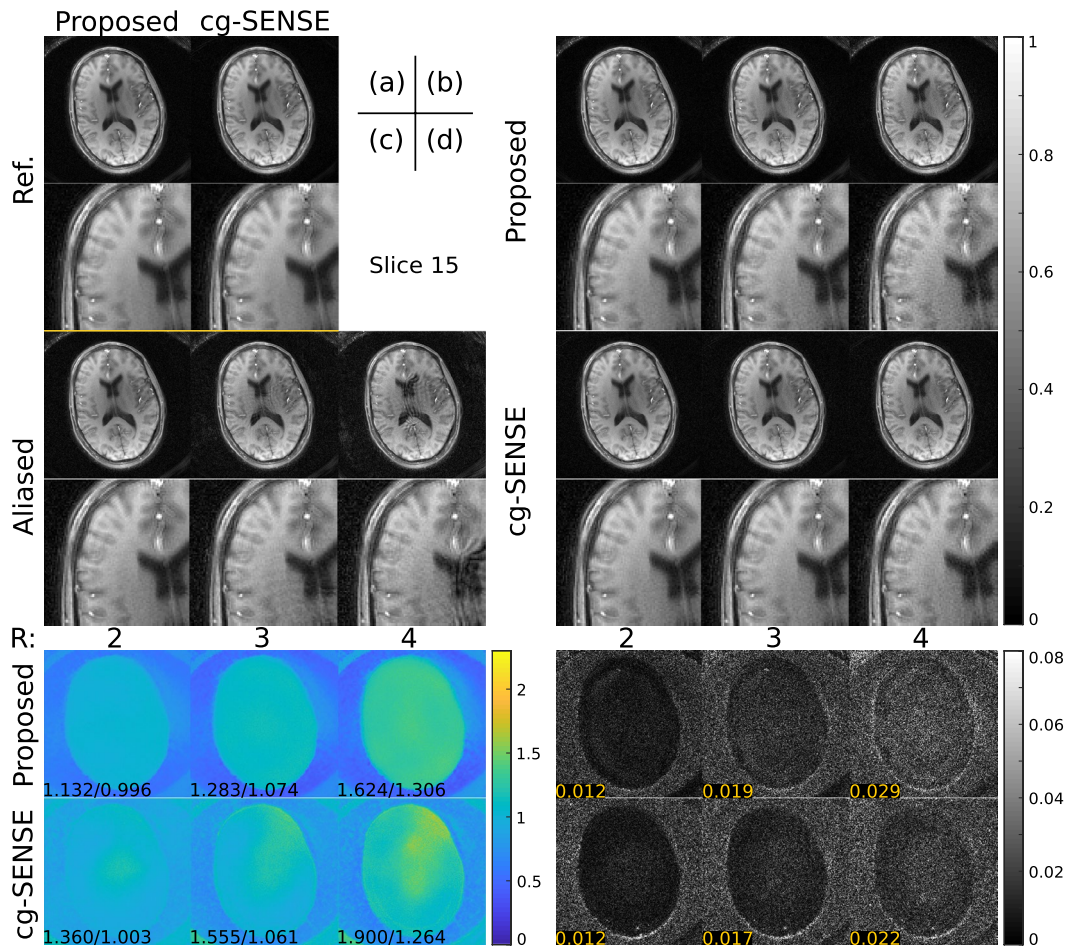

FIGURE 4 A similar comparison as in Figure 3, but for a retrospectively under-sampled rotated stack-of-spirals dataset. One slice out of 20 is shown. Without modifying our algorithm and implementation from 3, our non-Cartesian GRAPPA algorithm again achieves comparable reconstruction quality to cg-SENSE (error images' gray-scale windowed 12.5x)
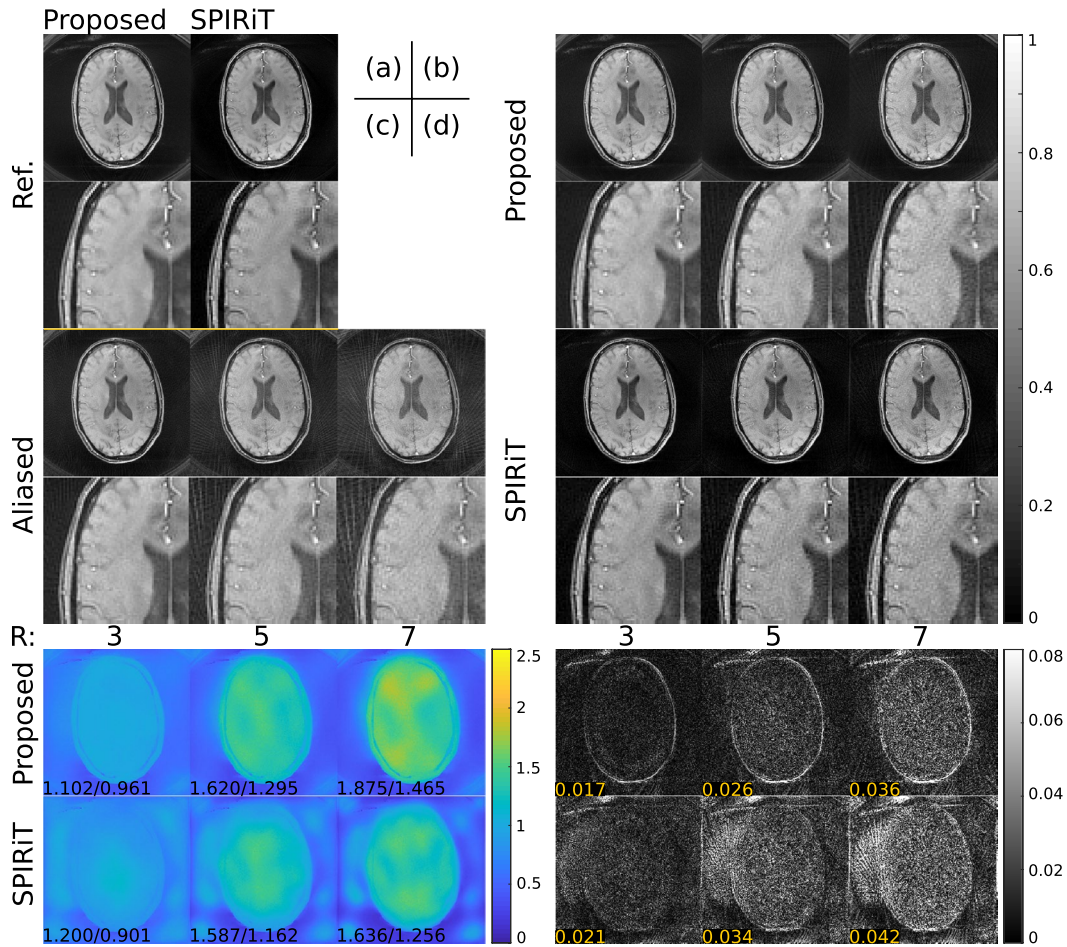
FIGURE 5 A similar comparison as in Figure 3, but in 2D, between the proposed non-Cartesian GRAPPA and SPIRiT, with large retrospective undersampling factors ($R = 3, 5, 7$). Without modifying our algorithm and implementation, our method achiews comparable reconstruction quality to SPIRiT. The "max/mean" g-factors, and RMSE are computed within the 2D head region

In all experiments, our proposed non-Cartesian GRAPPA only used the k-space center data for autocalibration. With these data, we compute low-resolution coil-product images of size $20 \times 20 \times n_z$, (n_z slices); and zero-pad them to $128 \times 128 \times n_z$ (~5× larger than ACS size, rounded up to power of 2) to attain the dense frequency spectrum needed for our efficient algorithm. A smaller padding size suffers from reduced accuracy, and a larger padding size trades off efficiency for only marginal improvement in quality. The reconstructed coil images of our method are linearly combined into final images. The sensitivity maps used in coil-image linear combination, cg-SENSE, and g-factor simulations are estimated using the methods developed in[31]. cg-SENSE reconstruction was implemented using MIRT[b], and was preconditioned with density compensation function[28] for acceleration. A quadratic roughness penalization was used to avoid overfitting the readouts.[9] SPIRiT reconstruction was done using its reference implementation[c]. We used the g-factor to measure noise amplification. Since a direct g-factor calculation for non-Cartesian imaging is intractable, we pseudo-replicated noisy reconstructions[29] in simulation (1000 times) using coil noise covariance information measured in vivo.

3 | RESULTS

Figures 3 and 4 compare our non-Cartesian GRAPPA with cg-SENSE on the retrospectively under-sampled 3D rotated stack-of-stars and rotated stack-of-spiral spoiled GRE datasets, respectively. For each comparison, 1 slice that is rich in structural detail is displayed (out of total 20 slices). The GRAPPA kernel size is $5 \times 5 \times 5$ in the units of the Nyquist sampling distance. Figures 5 and 6 compare our non-Cartesian GRAPPA with SPIRiT on the retrospectively under-sampled 2D star and spiral spoiled GRE datasets, respectively. The kernel size used in both non-Cartesian GRAPPA and SPIRiT is 7×7. As shown in both sets of comparisons, our non-Cartesian GRAPPA can rival cg-SENSE and SPIRiT in reconstruction quality in terms of error map and g-factor map. There is no modification to our algorithm or implementation across the comparisons, demonstrating the generality of our method.
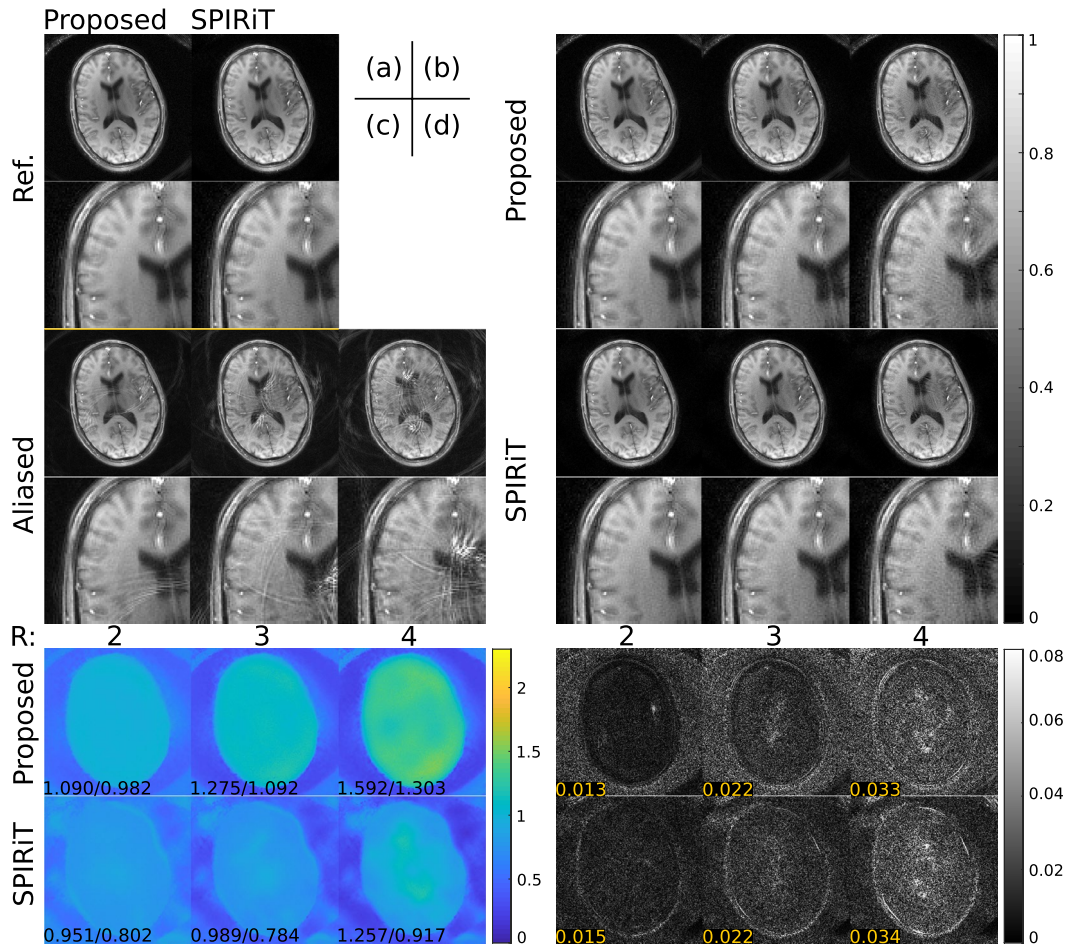
FIGURE 6 A similar comparison as in Figure 5, but for retrospectively under-sampled spiral dataset, with retrospective under-sampling acceleration factor ($R = 2, 3, 4$). Our proposed method still attains comparable quantities in terms of "max/mean" g-factor, and RMSE within 2D head image support. The reconstruction errors still contains high-frequency components as expected

From these 2 comparisons, besides noise, we observe edge-like (high-frequency) reconstruction errors with our method, which is also the case for the cg-SENSE and SPIRiT reconstructed images. This error behavior is expected, since for high-frequency (peripheral) k-space regions the non-Cartesian readout generally becomes sparser, such that fewer (sampled) neighbors are available for reconstruction. In addition, Supporting Information Figure S1 shows that in our non-Cartesian GRAPPA calibration, the influence of adopting a circulant boundary condition on reconstructed image quality is small.

In the pseudo-replica simulated g-factor maps, we observe that there are regions in the non-Cartesian GRAPPA case that fall below 1, as observed by others in Cartesian imaging.[29] We do notice that certain regions for the cg-SENSE case also fall below one. While Pruessmann et al[8] have rigorously proven the g-factor to always be greater than 1 for unfolding (non-iterative) SENSE, the derivation does not apply to non-Cartesian iterative SENSE (e.g., cg-SENSE). Supporting Information Figure S2 addresses this "g-factor smaller than one" behavior of cg-SENSE by illustrating a simulated 2D SENSE example where the g-factor falls below 1 analytically, obtained by calculating noise covariance directly. Basically, when under-sampling improves sampling density uniformity (e.g., over-sampled center k-space from fully-sampled trajectory), it is possible for g-factor to fall below 1.

Figure 7 compares the functional response obtained from the prospectively under-sampled image time-series reconstructed using our proposed algorithm, against that obtained from cg-SENSE. Three slices out of 20 are displayed. For cg-SENSE, different amounts of roughness penalization are evaluated. In this finger tapping experiment, as expected, we observe activations in the motor cortex, and, for cg-SENSE, the activation maps become increasingly blurred as roughness penalization increases. While the ground truth activation map is unknown, in terms of Dice similarity coefficient (dsc) (computed from all active voxels), our result most resembles the outcome of cg-SENSE at quadratic roughness penalty level $\lambda = 5$. However, our method is considerably faster: for this dataset it required 1 minute for calibration and 2 minutes for time-series
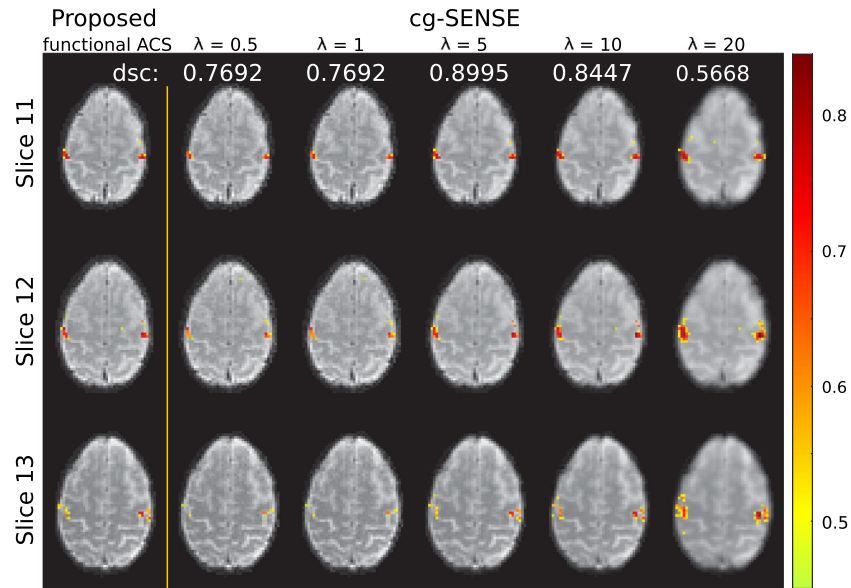
FIGURE 7 Functional imaging results from a *prospectively* under-sampled ($R = 3$) rotated stack-of-spirals 3D acquisition, (dsc: Dice coefficient). The volunteer performed a block (stimulus on/off) finger-tapping task, that is known to reliably activate motor cortex. Shown are activation maps in 3 consecutive slices (out of 20) covering the active region of motor cortex. Images from left to right are reconstructed with: Proposed non-Cartesian GRAPPA with kernels calibrated using the functional imaging ACS dataset (a fully sampled k-space center was obtained by combining the first 3 under-sampled frames of the fMRI time series); cg-SENSE with different levels of l2-roughness regularization parameter λ. While the ground truth activation map is unknown, in terms of Dice coefficient, our result most resembles the outcome of cg-SENSE at quadratic roughness penalty level $\lambda = 5$; and the activity maps from these 2 sets matches well with our expectation (motor cortex), while other activity maps are either incomplete or blurred. Reconstruction times were 1 minute for calibration plus 2 minutes for reconstruction for the proposed method, and 15 minutes for cg-SENSE

reconstruction, significantly shorter than the total cg-SENSE reconstruction time (15 minutes) (timed using MATLAB R2016a on a RHEL7.4 server with 2 Intel E5-2630V4 processors. The NUFFT operations of cg-SENSE are also implemented in C subroutines.). Supporting Information Figure S3 presents the activation map reconstructed using GRAPPA kernels calibrated with the separate distinct-contrast higher resolution dataset. As expected from,[13] we obtain an essentially identical reconstruction compared to the reconstruction using the original GRAPPA kernels (obtained by combining 3 subsequent fMRI data frames to form a full-sampled k-space), with a dsc equal to 1. The image from the separate dataset is presented in Supporting Information Figure S3.

4 | DISCUSSION

Our proposed non-Cartesian GRAPPA reconstruction is general and applicable to arbitrary readout trajectories, without modifying or adapting the code. Here, we demonstrated its use with rotated stack-of-stars and rotated stack-of-spirals, using 3D kernels. As stated above, a convenient feature of these trajectories in the context of the proposed method is that the under-sampling pattern is regular along k_z, which reduces the number of different (unique) constellations. However, in

the extreme case, the number of different kernels can equal the number of unsampled k-space points, which can require large computer memory. For instance, in 3D high-resolution imaging with 32 receive coils, there are about 400 000 3D kernels to be calibrated, each of which could contain about 50 neighbors on average, resulting in approximately 76 GB memory/storage demand for the coefficients (for single float precision). Coil compression[32,33] can reduce this demand. Another mitigating approach may be to "interleave" calibration with reconstruction of each distinct constellation in a time-series; in other words, one can calibrate and cache the GRAPPA weights for a single constellation (or group of constellations) at a time, reconstruct all the data corresponding to this constellation throughout the whole time-series, then discard these weights and move on to the next constellation.

We have shown that using circulant boundary conditions (concerning the ACS region) permits fast GRAPPA weight calibration, without degrading image quality. From our retrospectively under-sampled reconstruction experiments (in Supporting Information), we observe that for certain calibration set-ups (e.g., ACS region size, GRAPPA kernel size, and Tikhonov regularizer coefficient), allowing the ACS to wrap around can sometimes even improve reconstruction quality slightly (in terms of image error and g-factor). This behavior can be dependent on specific subjects and sampling trajectories.

However, we certainly do not expect this behavior to be a general features of our method; whether the image error is slightly improved or worsened may depend on various factors, such as the sampling trajectory used, or the subject or receive coil configuration. Overall, the circulant boundary condition appears to have a marginal influence on reconstructed image quality.

We demonstrated our method by reconstructing retrospectively unsampled k-space locations. While these locations are a natural choice, our method is not limited to reconstruct only the "unsampled" locations; for example, one possibility would be to reconstruct points on a Cartesian grid. The optimal choice of k-space locations to reconstruct with respect to, for example, image quality and overall computation time, is an open problem.

k-Space reconstruction methods, such as GRAPPA and SPIRiT, are known to be robust against aliasing when FOV is limited (i.e., smaller than the object being imaged). Our proposed non-Cartesian GRAPPA is exactly equivalent to ordinary GRAPPA when the sampling is Cartesian; and hence would be expected to also be robust to limited FOV. In fMRI, limited-FOV is usually not an issue, as brains are typically small enough to be fully covered in scanning. However, this issue could be significant in cardiac and other body imaging, where the region of interest is not that isolated from other parts of body. We will study the performance of our method in cardiac imaging in the future.

5 | CONCLUSION

This paper introduced a flexible and rapid non-iterative reconstruction method that is a true non-Cartesian generalization of the canonical GRAPPA method. The method works with arbitrary sampling trajectories, and may be particularly beneficial in applications such as dynamic imaging (e.g., fMRI) where a large number of images must be reconstructed.

NOTES

[a] Code available at https://github.com/tianrluo/NonCrtGRAPPA.

[b] Michigan Image Reconstruction Toolbox, available at http://web.eecs.umich.edu/fessler/code.

[c] SPIRiT reference implementation, available at http://people.eecs.berkeley.edu/mlustig/Software.html.

REFERENCES

1. Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn Reson Med*. 2012;67:1210–1224.
2. Stirnberg R, Huijbers W, Brenner D, Poser BA, Breteler M, Stöcker T. Rapid whole-brain resting-state fMRI at 3 T: efficiency-optimized three-dimensional EPI versus repetition time-matched simultaneous-multi-slice EPI. *NeuroImage* 2017;163:81–92.
3. Feinberg DA, Yacoub E. The rapid development of high speed, resolution and precision in fMRI. *Neuroimage* 2012;62:720–725.
4. Schneider JT, Kalayciyan R, Haas M, et al. Inner-volume imaging in vivo using three-dimensional parallel spatially selective excitation. *Magn Reson Med*. 2013;69:1367–1378.
5. Luo T, Nielsen JF, Noll DC. Improved aliasing suppression in steady-state, parallel imaging using inner volume excitation. In: the 24th Annual Meeting of ISMRM, Singapore, 2016:3221.
6. Peltier SJ, Noll DC. Physiological noise in multi-shot functional MRI. In: the 25th Annual Meeting of ISMRM, Honolulu, HI, USA, 2002:118.
7. Peltier SJ. Characterization and compensation of systematic noise in functional magnetic resonance imaging. PhD thesis, University of Michigan, Ann Arbor; 2003.
8. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*. 1999;42:952–962.
9. Pruessmann KP, Weiger M, Börnert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med*. 2001;46:638–651.
10. Lustig M, Pauly JM. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med*. 2010;64:457–471.
11. Cerjanic A, Holtrop JL, Ngo GC, et al. PowerGrid: a open source library for accelerated iterative magnetic resonance image reconstruction. In: the 24th Annual Meeting of ISMRM, Singapore, 2016:0525.
12. Griswold MA, Jakob PM, Heidemann RM, et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*. 2002;47:1202–1210.
13. Talagala SL, Sarlls JE, Liu S, Inati SJ. Improvement of temporal signal-to-noise ratio of GRAPPA accelerated echo planar imaging using a FLASH based calibration scan. *Magn Reson Med*. 2016;75:2362–2371.
14. Polimeni JR, Bhat H, Witzel T, et al. Reducing sensitivity losses due to respiration and motion in accelerated echo planar imaging by reordering the autocalibration data acquisition. *Magn Reson Med*. 2016;75:665–679.
15. Griswold MA, Heidemann RM, Jakob PM. Direct parallel imaging reconstruction of radially sampled data using GRAPPA with relative shifts. In: the 11th Annual Meeting of ISMRM, Toronto, Ontario, Canada, 2003;11:2349.
16. Heidemann RM, Griswold MA, Seiberlich N, et al. Direct parallel image reconstructions for spiral trajectories using GRAPPA. *Magn Reson Med*. 2006;56:317–326.
17. Seiberlich N, Breuer FA, Blaimer M, Barkauskas K, Jakob PM, Griswold MA. Non-Cartesian data reconstruction using GRAPPA operator gridding (GROG). *Magn Reson Med*. 2007;58:1257–1265.
18. Seiberlich N, Breuer FA, Ehses P, et al. Using the GRAPPA operator and the generalized sampling theorem to reconstruct undersampled non-Cartesian data. *Magn Reson Med*. 2009;61:705–715.
19. Seiberlich N, Lee G, Ehses P, Duerk JL, Gilkeson R, Griswold M. Improved temporal resolution in cardiac imaging using through-time spiral GRAPPA. *Magn Reson Med*. 2011;66:1682–1688.
20. Yeh EN, McKenzie CA, Ohliger MA, Sodickson DK. Parallel magnetic resonance imaging with adaptive radius in k-space (PARS): constrained image reconstruction using k-space locality in radiofrequency coil encoded data. *Magn Reson Med*. 2005;53:1383–1392.

21. Liu C, Bammer R, Moseley ME. Parallel imaging reconstruction for arbitrary trajectories using k-space sparse matrices (kSPA). *Magn Reson Med*. 2007;58:1171–1181.

22. Luo T, Noll DC, Fessler JA, Nielsen JF. A fast and general non-Cartesian GRAPPA reconstruction method. In Proceedings of the 26th Annual Meeting of ISMRM, Paris, France, 2018:2821.

23. Fessler JA, Sutton BP. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Trans Signal Processing* 2003;51:560–574.

24. Tian Y, Erb KC, Adluru G, et al. Technical note: evaluation of pre-reconstruction interpolation methods for iterative reconstruction of radial k-space data. *Med Phys*. 2017;44:4025–4034.

25. Zhou Z, Han F, Yan L, Wang DJJ, Hu P. Golden-ratio rotated stack-of-stars acquisition for improved volumetric MRI. *Magn Reson Med*. 2017;78:2290–2298.

26. Deng W, Zahneisen B, Stenger VA. Rotated stack-of-spirals partial acquisition for rapid volumetric parallel MRI. *Magn Reson Med*. 2016;76:127–135.

27. Bentley JL. Multidimensional binary search trees used for associative searching. *Commun ACM* 1975;18:509–517.

28. Rasche V, Proksa R, Sinkus R, Bornert P, Eggers H. Resampling of data between arbitrary grids using convolution interpolation. *IEEE Trans Med Imaging* 1999;18:385–392.

29. Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, McKenzie CA. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn Reson Med*. 2008;60:895–907.

30. Gurney PT, Hargreaves BA, Nishimura DG. Design and analysis of a practical 3D cones trajectory. *Magn Reson Med*. 2006;55:575–582.

31. Allison MJ, Ramani S, Fessler JA. Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods. *IEEE Trans Med Imaging* 2013;32:556–564.

32. Zhang T, Pauly JM, Vasanawala SS, Lustig M. Coil compression for accelerated imaging with Cartesian sampling. *Magn Reson Med*. 2013;69:571–582.

33. Chu A, Noll DC. Coil compression in simultaneous multislice functional MRI with concentric ring slice-GRAPPA and SENSE. *Magn Reson Med*. 2016;76:1196–1209.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

FIGURE S1 Reconstruction quality comparison of our non-Cartesian GRAPPA method using different ACS boundary conditions (circulant and trimmed), for 2D star (left panels) and spiral (right panels) datasets. The top row plots the absolute error maps and the digits are their averages within the object support. The center row plots the g-factor maps and the digits are their max/average g-factors within the support. The bottom row plots the absolute error viewed from k-space. We observe that circulant boundaries can produce similar reconstruction error as trimmed boundaries. Moreover, for certain Tikhonov regularization setups, circulant boundary outperforms trimmed boundaries

FIGURE S2 An illustration of a reconstruction with g-factor smaller than one (in some regions of the image), using realistic (in vivo) sensitivity maps. Here, "full" sampling consists of both blue and red locations, while "under-sampling" only contains the blue locations. The central k-space region is oversampled, as is typically the case in non-Cartesian acquisitions. The center k-space oversampling ratio 4.1 in this example produces off-grid sampling. In this example, the g-factor is just below 1.0 near the right and left parts of the image (white regions in the binary black/white image on the lower left)

FIGURE S3 Reconstruction quality demonstration of GRAPPA kernel calibrated with ACS of a different contrast. Dice coefficients (dsc) are labeled for convenience of assessment. This figure is the same as Figure 7, except with 2 extra columns (from left to right): Proposed non-Cartesian GRAPPA with kernels calibrated using the structural imaging ACS dataset. The high-resolution structural image of the same subject, acquired along with the fMRI scanning

A GRAPPA algorithm for arbitrary 2D/3D non-Cartesian sampling trajectories with rapid calibration

Tianrui Luo | Douglas C. Noll | Jeffrey A. Fessler | Jon-Fredrik Nielsen

1 | CALIBRATION BOUNDARY CONDITION: CIRCULANT VS TRIMMED

Fig. S1 examines the influence of adopting circulant ACS boundary conditions on 2D reconstruction quality. In general, with proper Tikhonov regularization, the circulant boundary assumption produces similar reconstruction error as conventional ("trimmed") boundaries. The g-factor behavior is more complicated: depending on the choice of regularization parameter λ, circulant boundaries can do either slightly better or slightly worse than trimmed boundaries. In this comparison, reconstruction error with $\lambda = 5 \times 10^{-7}$ is low and difference is small between circulant and trimmed boundary conditions. This λ represents a compromise between g-factor and image/k-space error. We used this value throughout our experiments in the main text.
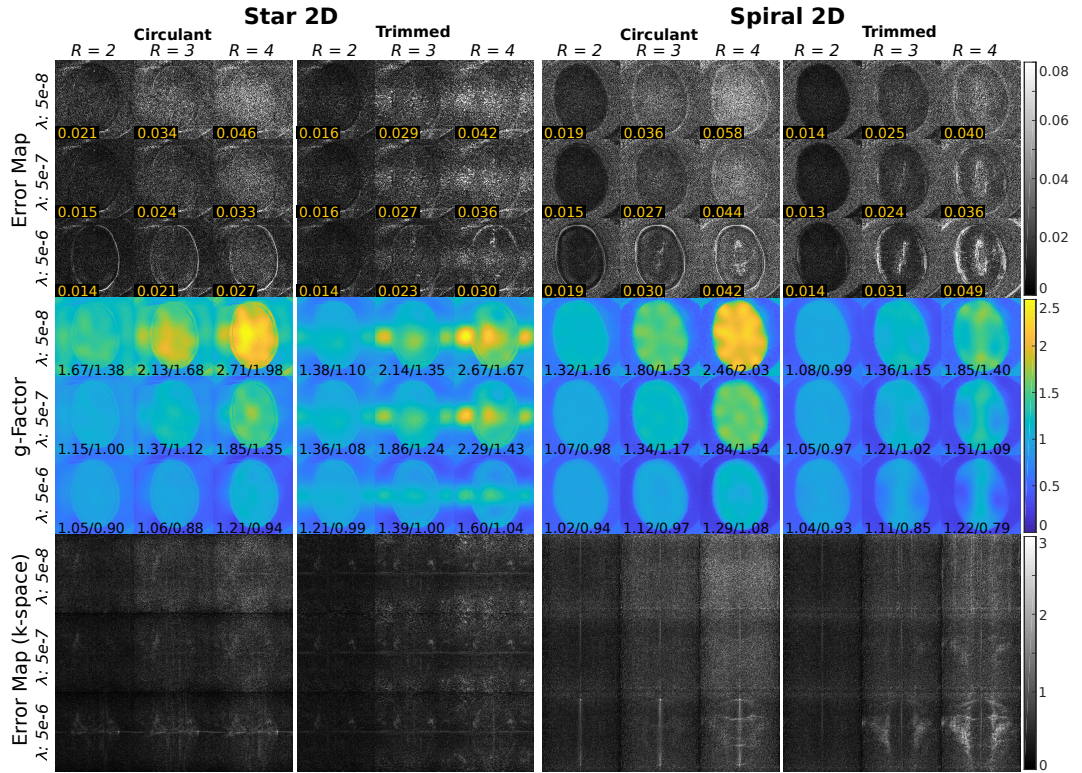

FIGURE S1 Reconstruction quality comparison of our non-Cartesian GRAPPA method using different ACS boundary conditions (circulant and trimmed), for 2D star (left panels) and spiral (right panels) datasets. The top row plots the absolute error maps and the digits are their averages within the object support. The center row plots the g-factor maps and the digits are their max/average g-factors within the support. The bottom row plots the absolute error viewed from k-space. We observe that circulant boundaries can produce similar reconstruction error as trimmed boundaries. Moreover, for certain Tikhonov regularization setups, circulant boundary outperforms trimmed boundaries.

2 | G-FACTOR FOR SENSE WITH OVER-SAMPLED CENTER OF K-SPACE

Oversampling, which occurs often near the k-space center for non-Cartesian acquisitions, is a key reason why the g-factor can fall below 1 for SENSE. To see this, consider a simple 2-pixel example with two noise-uncorrelated coil sensitivities of [1, 0.5] and [0.3, 1], respectively. For the "fully sampled" acquisition, we sample three times at k=0, and once at k=1; for the "under-sampled" acquisition, we sample once at both k=0 and k=1. The acceleration factor, R, in this example is thus 2. Following the definition of g-factor, a simple calculation would yield $g = \sqrt{3/4} < 1$.

For a more realistic example, in Fig. S2 we simulate a sampling pattern where the phase-encoding direction is over-sampled by a factor of 4.1 near the k-space center. The sensitivity maps are estimated from an *in vivo* data set; a low-resolution (64 × 64) matrix size is chosen to enable direct calculation. In this sampling pattern, "full"-sampling consists of both blue and red locations, while under-sampling only contains the blue locations. Again, we observe that g falls below 1 in some parts of the analytically calculated g-factor map. In summary, we conclude that oversampling, which is common for general non-Cartesian sampling, can cause the g-factor to fall below 1.
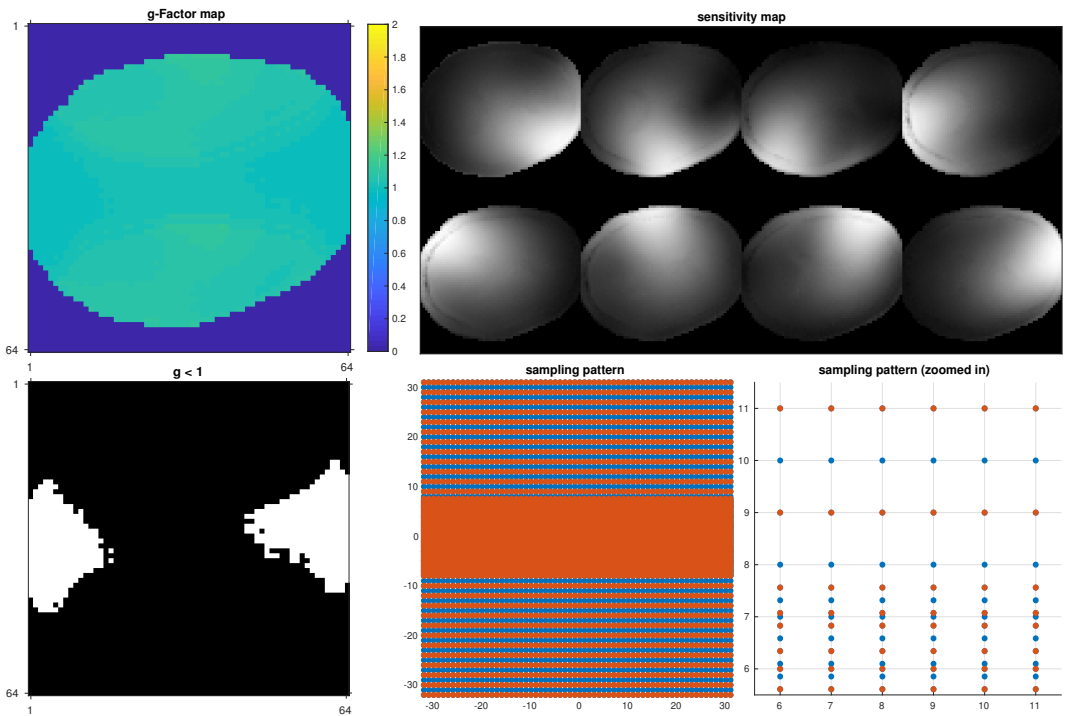

FIGURE S2 An illustration of a reconstruction with g-factor smaller than one (in some regions of the image), using realistic (*in vivo*) sensitivity maps. Here, "full" sampling consists of both blue and red locations, while "under-sampling" only contains the blue locations. The central k-space region is oversampled, as is typically the case in non-Cartesian acquisitions. The center k-space oversampling ratio 4.1 in this example produces off-grid sampling. In this example, the g-factor is just below 1.0 near the right and left parts of the image (white regions in the binary black/white image on the lower left).

3 | GRAPPA CALIBRATION USING ACS WITH DIFFERENT CONTRAST

Talagala *et al* observed that GRAPPA kernels calibrated with one acquisition can faithfully reconstruct other datasets of possibly different contrasts and resolution configurations (Ref. [13] of the maintext). This protocol can be useful for dynamic imaging, e.g., fMRI, where a structural dataset is commonly acquired alongside a number of functional activity datasets. To evaluate the use of an ACS dataset with contrast different from the to-be-reconstructed (undersampled) images, we tested reconstructing a finger-tapping dataset with kernels calibrated using a separate structural dataset. The fMRI results are shown in Fig. 7 of the main text. Here, we again show the results from Fig 7, but in addition we also show fMRI activation maps obtained with the proposed method based on ACS data from the structural dataset. The acquisition parameters are included in the **Experiments** section of the main text, and are repeated here for convenience: "…on the same subject from the fMRI experiment at the identical FOV, we also acquired a fully-sampled stack-of-spiral-out dataset with a different contrast: This acquisition still have 20 k_z platters, each containing 9 interleaves. Its other parameters are: TR=30ms, TE=2.17ms, flip-angle 8°, image matrix size $220 \times 220 \times 20$. Again, the experiments were conducted on a GE 3T scanner using an 8-channel receive-only head coil…"
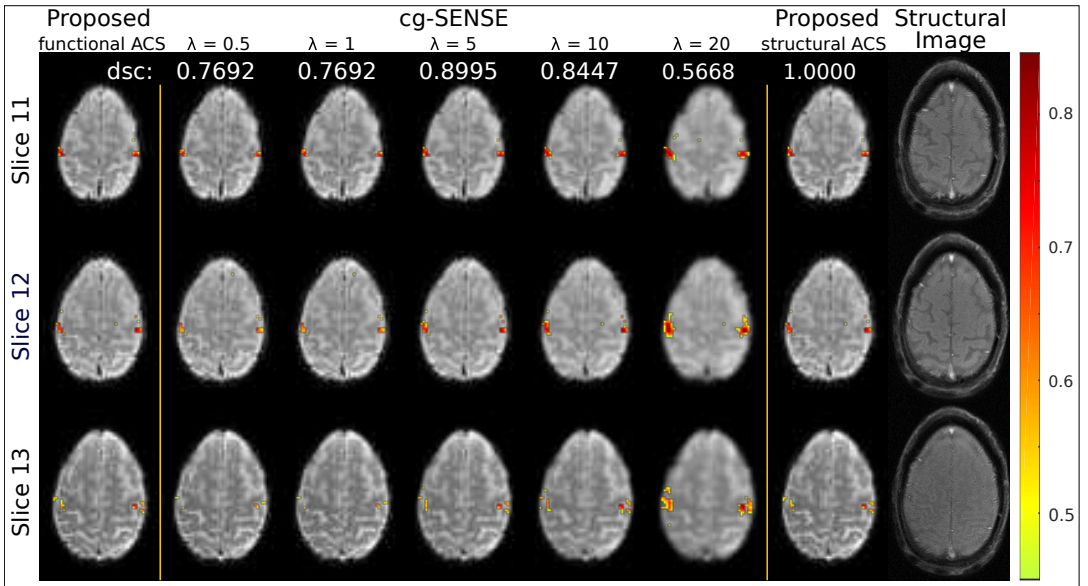

FIGURE S3 Reconstruction quality demonstration of GRAPPA kernel calibrated with ACS of a different contrast. Dice coefficients (dsc) are labeled for convenience of assessment. This figure is the same as Fig. 7, except with two extra columns (from left two right): Proposed non-Cartesian GRAPPA with kernels calibrated using the structural imaging ACS dataset; The high-resolution *structural* image of the same subject, acquired along with the fMRI scanning.