Magnetic Resonance in Medicine

Robust autocalibrated structured low-rank EPI ghost correction

Rodrigo A. Lobos[1,2] (iD) | **W. Scott Hoge**[4,5] | **Ahsan Javed**[1,2] (iD) | **Congyu Liao**[5,6] (iD) |
Kawin Setsompop[5,6] | **Krishna S. Nayak**[1,2,3] (iD) | **Justin P. Haldar**[1,2,3] (iD)

[1]Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA

[2]Signal and Image Processing Institute, University of Southern California, Los Angeles, CA, USA

[3]Department of Biomedical Engineering, University of Southern California, Los Angeles, CA, USA

[4]Department of Radiology, Brigham and Women's Hospital, Boston, MA, USA

[5]Department of Radiology, Harvard Medical School, Boston, MA, USA

[6]Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA, USA

Correspondence
Rodrigo A. Lobos, Signal and Image Processing Institute, University of Southern California, University Park Campus, 3710 McClintock Avenue, Ronald Tutor Hall (RTH) #317, Los Angeles, CA 90036, USA.
Email: rlobos@usc.edu

Funding information
NSF, Grant/Award Number: CCF-1350563; NIH, Grant/Award Number: R01-MH116173, R21-EB022951, R01-NS089212, R01-NS074980 and R01-HL130494

Purpose: We propose and evaluate a new structured low-rank method for echo-planar imaging (EPI) ghost correction called Robust Autocalibrated LORAKS (RAC-LORAKS). The method can be used to suppress EPI ghosts arising from the differences between different readout gradient polarities and/or the differences between different shots. It does not require conventional EPI navigator signals, and is robust to imperfect autocalibration data.

Methods: Autocalibrated LORAKS is a previous structured low-rank method for EPI ghost correction that uses GRAPPA-type autocalibration data to enable high-quality ghost correction. This method works well when the autocalibration data are pristine, but performance degrades substantially when the autocalibration information is imperfect. RAC-LORAKS generalizes Autocalibrated LORAKS in two ways. First, it does not completely trust the information from autocalibration data, and instead considers the autocalibration and EPI data simultaneously when estimating low-rank matrix structure. Second, it uses complementary information from the autocalibration data to improve EPI reconstruction in a multi-contrast joint reconstruction framework. RAC-LORAKS is evaluated using simulations and in vivo data, including comparisons to state-of-the-art methods.

Results: RAC-LORAKS is demonstrated to have good ghost elimination performance compared to state-of-the-art methods in several complicated EPI acquisition scenarios (including gradient-echo brain imaging, diffusion-encoded brain imaging, and cardiac imaging).

Conclusions: RAC-LORAKS provides effective suppression of EPI ghosts and is robust to imperfect autocalibration data.

KEYWORDS
constrained reconstruction, echo-planar imaging, multi-contrast reconstruction, Nyquist ghost correction, structured low-rank matrix recovery

1 | INTRODUCTION

Echo-planar imaging (EPI) is a widely used high-speed MRI acquisition strategy,[1] but is subject to several undesirable artifacts.[2] Nyquist ghosts are one of the most common EPI artifacts, and occur because of systematic differences between the interleaved lines of k-space that are acquired with different readout gradient polarities, and/or because of systematic differences between interleaved lines of k-space data that are acquired with different shots in a multi-shot acquisition. Despite substantial efforts over several decades to solve this problem,[2-20] the widely deployed modern ghost correction schemes are still prone to incomplete ghost suppression, as illustrated in Supporting Information Figure S1.

Recently, structured low-rank matrix methods for ghost correction[16-20] have received increasing attention for their ability to provide excellent ghost-suppression performance without the need for additional "navigator" information (ie, reference scans collected alongside each EPI readout that allow estimation of the systematic inconsistencies between different gradient polarities or different shots). These methods can suppress ghosts better than navigator-based methods, and eliminate the need to acquire navigators for each EPI readout. Among different structured low-rank matrix approaches, a ghost correction method based on Autocalibrated LORAKS (AC-LORAKS)[21] was previously demonstrated to yield high-quality results across a range of different scenarios.[18] To eliminate a fundamental ambiguity in structured low-rank matrix recovery from uniformly undersampled EPI data,[18] AC-LORAKS makes use of parallel imaging subspace information estimated from autocalibration (ACS) data acquired in a pre-scan. This ACS-based approach is similar to standard autocalibrated parallel imaging methods like GRAPPA,[22] SPIRiT,[23] and PRUNO.[24]

While the AC-LORAKS approach to ghost correction generally works well when the ACS data are pristine and well-matched to the EPI data to be reconstructed, there are many situations where experimental conditions (eg, subject motion, eddy currents, etc.) can lead to artifacts within the ACS data or mismatches between the ACS and EPI data. The performance of the AC-LORAKS ghost correction procedure degrades in the presence of these ACS artifacts and mismatches. Note that this kind of issue is not unique to AC-LORAKS or to ghost correction, and imperfect ACS/calibration data is a long-standing and commonly reported problem for all calibration-based image reconstruction methods.[11,25-29] For AC-LORAKS ghost correction, imperfect ACS data can be especially troublesome in contexts where the prescan would be done once before acquiring a long sequence of multiple EPI images (eg, in BOLD fMRI or diffusion MRI applications), and then used to reconstruct each image in the sequence.

In this paper, we propose an extension of AC-LORAKS for EPI ghost correction that is more robust to imperfections in the ACS data. The new method, called Robust Autocalibrated LORAKS (RAC-LORAKS), has two major differences from the previous AC-LORAKS approach. First, RAC-LORAKS does not completely trust the subspace information learned from the ACS data, but rather uses a novel structured low-rank matrix formulation that learns subspace information jointly from both the (imperfect) ACS data and the EPI data being reconstructed. To the best of our knowledge, no previous methods have used this kind of approach to address the longstanding issue of imperfect ACS data. And second, RAC-LORAKS uses the ACS data to provide additional complementary information for the reconstruction of the EPI data within a multi-contrast joint reconstruction framework.[30] Preliminary accounts of the first strategy were previously reported in recent conferences,[31,32] although we have not previously reported the combination with the second strategy.

2 | THEORY

Due to space constraints, our descriptions in this paper will assume that the reader is already familiar with the basic physics of EPI. Readers interested in a more detailed explanation are referred to classic references.[1,2] For simplicity, our description of EPI ghost correction will generally assume that we are correcting ghosts associated with the differences between data acquired with different readout gradient polarities in a single-shot EPI experiment. However, since the ghost model for bipolar gradients is nearly identical to the ghost model for multi-shot acquisition, the same approach is easily adapted *mutatis mutandis* to multi-shot acquisition with an arbitrary number of shots.[18]

2.1 | Background: Structured low-rank EPI ghost correction

Structured low-rank matrix methods for EPI ghost correction[16-20] can be viewed as an extension of structured low-rank matrix methods for conventional MR image reconstruction,[24,33-39] and are based on the same underlying theoretical principles. In particular, it has been shown that when MRI images have limited support, smooth phase variations, multi-channel correlations, or transform-domain sparsity, then the MRI k-space data will be linearly predictable,[40] which means that convolutional Hankel- or Toeplitz-structured matrices formed from the k-space data will possess low-rank characteristics. This observation means that MRI reconstruction can be reformulated as structured low-rank matrix recovery. Importantly, these structured low-rank matrix recovery

methods can even be successful in calibrationless scenarios where ACS data or other prior information about the spatial support, phase, or multi-channel sensitivity profiles is not available.[35-37]

Structured low-rank EPI ghost correction methods combine these principles with the earlier observation that EPI data acquired from different gradient polarities or different shots can be treated as coming from different effective "channels" in a parallel imaging experiment, where the systematic differences between different polarities or shots lead to different phase or magnitude modulations of the underlying EPI image.[9,11,12,15] Since structured low-rank methods for conventional image reconstruction automatically account for the unknown sensitivity maps that modulate the underlying image in a parallel imaging experiment, it is reasonable to apply these same types of methods to handle the unknown polarity- or shot-dependent modulations that manifest in EPI ghost correction.

For the sake of brevity, we will focus the remainder of our review on the AC-LORAKS method for EPI ghost correction, [18] since our proposed RAC-LORAKS method is a generalization of AC-LORAKS. The AC-LORAKS method for EPI ghost correction is based on solving the following regularized optimization problem subject to exact data consistency constraints:

There are theoretical benefits to choosing a nonconvex low-rank regularization penalty,[18] and the previous AC-LORAKS approach for ghost correction[18] used the nonconvex function proposed in the original LORAKS paper[36] defined by

$$J_r(\mathbf{X}) = \min_{\mathbf{Y}} \|\mathbf{X} - \mathbf{Y}\|_F^2 \text{ s.t. } \text{rank}(\mathbf{Y}) \leq r, \qquad (2)$$

where r is a user-selected rank parameter, $\mathbf{X}$ is a matrix representing the point at which we are evaluating the function $J_r(\mathbf{X})$, and $\mathbf{Y}$ is an optimization variable of the same size as $\mathbf{X}$. This penalty encourages matrices that have accurate rank-r approximations.

The first two terms appearing on the right-hand side of Equation (1), respectively, impose limited support and parallel imaging constraints on the reconstructions of the positive and negative readout gradient polarities. The constraints that are used in these terms are implicit in the low-rank characteristics of the structured LORAKS matrices, as captured by the nullspace matrix $\mathbf{N}$. The nullspace matrix is learned in advance from the ACS data, and as a result, there is an implicit assumption that the support and parallel imaging constraints that were valid for the ACS data are also valid for the EPI data to be reconstructed. Note that if the third term were removed from Equation (1), then these first two terms would reduce to performing separate PRUNO[24] or conven-

$$\left\{ \widehat{\mathbf{k}}^+, \widehat{\mathbf{k}}^- \right\} = \underset{\{\mathbf{k}^+, \mathbf{k}^-\}}{\text{argmin}} \left\| \mathscr{P}_C(\mathbf{k}^+)\mathbf{N} \right\|_F^2 + \left\| \mathscr{P}_C(\mathbf{k}^-)\mathbf{N} \right\|_F^2 + \lambda J_r([\mathscr{P}_S(\mathbf{k}^+) \ \mathscr{P}_S(\mathbf{k}^-)]). \qquad (1)$$

In this expression, $\mathbf{k}^+$ and $\mathbf{k}^-$, respectively, represent the ideal fully sampled multi-channel Cartesian k-space data for the positive and negative readout gradient polarities; $\mathscr{P}_C(\cdot)$ is the LORAKS operator that maps the k-space data into a structured matrix that should possess low-rank if the multi-channel image possess limited support and/or interchannel parallel imaging correlations; $\mathscr{P}_S(\cdot)$ is the LORAKS operator that maps the k-space data into a structured matrix that should possess low-rank if the multi-channel image possess limited support, smooth phase, and/or interchannel parallel imaging correlations; the matrix $\mathbf{N}$ comprises an orthonormal (ie, $\mathbf{N}^H\mathbf{N} = \mathbf{I}$) basis for the nullspace of the matrix $\begin{bmatrix} \mathscr{P}_C(\mathbf{k}_{acs+}) \\ \mathscr{P}_C(\mathbf{k}_{acs-}) \end{bmatrix}$, where $\mathbf{k}_{acs}^+$ and $\mathbf{k}_{acs}^-$, respectively, represent the ACS data for the positive and negative readout gradient polarities; λ is a regularization parameter; $J_r(\cdot)$ is a regularization penalty that promotes low-rank characteristics; and $\|\cdot\|_F$ denotes the Frobenius norm. Due to space constraints, this paper will not provide a detailed recipe for implementing the LORAKS operators $\mathscr{P}_C(\cdot)$ and $\mathscr{P}_S(\cdot)$, and simply note that our implementations for this paper are identical to those that are described in detail in earlier LORAKS papers.[36,37]

tional AC-LORAKS[21] reconstructions of the data from each polarity. Acquiring ACS/calibration data is relatively fast and easy, and is already a standard part of most modern parallel imaging protocols, so is not very burdensome on the acquisition. Using ACS data can also be important in this context, since it has been mathematically proven that structured low-rank matrix methods for ghost correction suffer from fundamental ambiguities unless some form of side information is available.[18] While other options exist for removing ambiguity (eg, using SENSE-like[41] image-domain constraints[17,18]), it was previously observed that the AC-LORAKS approach (ie, using GRAPPA-like[22] Fourier-domain constraints) offered better performance.[18]

The third term of Equation (1) couples the reconstruction of the two polarities together, allowing the reconstruction of one polarity to benefit from information from the other polarity, while also introducing phase constraints to allow the reconstruction to benefit from k-space conjugate symmetry characteristics. In particular, the third term implicitly and automatically imposes the following constraints whenever they are compatible with the measured data: limited image support, smooth phase, interchannel parallel imaging correlations, and interpolarity correlations. Notably,

these constraints are all imposed implicitly through the nullspace of a structured matrix, and if a given constraint is not compatible with the measured data, then that constraint will automatically not be imposed by the reconstruction procedure.[40]

The ACS data for AC-LORAKS ghost correction have typically been acquired using the same process used by dual-polarity GRAPPA (DPG).[11,42-44] In particular, assuming a parallel imaging acceleration factor of R, DPG employs a $2R$-shot EPI prescan. The data from different shots and different gradient polarities are then rearranged and interleaved to form one fully sampled ACS dataset comprised only of data acquired with a positive readout gradient polarity ($\mathbf{k}_{acs}^{+}$) and another fully sampled ACS dataset comprised only of data acquired with a negative readout gradient polarity ($\mathbf{k}_{acs}^{-}$). Since this ACS acquisition strategy is based on a multi-shot approach, it therefore may be prone to ghosting artifacts due to shot-to-shot variations. In addition, since the ACS data are often acquired only once at the beginning of a long multi-image EPI scan (eg, in BOLD fMRI or diffusion MRI experiments), the ACS data acquired at the beginning of the experiment may gradually become mismatched with data acquired at later time points due to scanner drift, subject motion, etc. As noted previously, the ghost correction performance of AC-LORAKS can be substantially degraded when there are mismatches between the ACS data and the EPI data to be reconstructed. Although a preprocessing procedure has been previously developed to correct for shot-to-shot variations in the ACS data for DPG,[11] this approach is not sufficient for the present context. In particular, this approach undesirably modifies the magnitude and phase characteristics of the ACS data in ways that are not well-suited for AC-LORAKS, and only addresses ACS artifacts without accounting for mismatches that may exist between the ACS data and the EPI data.

2.2 | RAC-LORAKS

Our proposed RAC-LORAKS method is based on solving the following optimization problem

$$
\begin{aligned}
\left\{ \widehat{\mathbf{k}}^{+}, \ \widehat{\mathbf{k}}^{-}, \ \widehat{\mathbf{N}} \right\} = & \operatorname*{argmin}_{\{\mathbf{k}^{+}, \mathbf{k}^{-}, \mathbf{N}\}} \ \left\| \mathscr{P}_C(\mathbf{k}^{+})\mathbf{N} \right\|_F^2 + \left\| \mathscr{P}_C(\mathbf{k}^{-})\mathbf{N} \right\|_F^2 \\
& + \eta \left\| \mathscr{P}_C(\mathbf{k}_{acs}^{+})\mathbf{N} \right\|_F^2 + \eta \left\| \mathscr{P}_C(\mathbf{k}_{acs}^{-})\mathbf{N} \right\|_F^2 \\
& + \lambda J_r \left(\left[\mathscr{P}_S(\mathbf{k}^{+}) \ \mathscr{P}_S(\mathbf{k}^{-}) \ \mathscr{P}_S(\mathbf{k}_{acs}^{+}) \ \mathscr{P}_S(\mathbf{k}_{acs}^{-}) \right] \right)
\end{aligned}
$$

(3)

subject to exact data consistency constraints on $\mathbf{k}^{+}$ and $\mathbf{k}^{-}$ and subject to orthonormality constraints on $\mathbf{N}$ such that $\mathbf{N}^H\mathbf{N} = \mathbf{I}$. This optimization problem involves four user-selected

parameters: the regularization parameters η and λ, the rank parameter r, and the number of columns p of the matrix $\mathbf{N}$ (which determines the dimension of the approximate nullspace).

Equation (3) has two main differences from Equation (1). First, instead of choosing a predetermined value of the approximate nullspace matrix $\mathbf{N}$ that depends only on the ACS data, $\mathbf{N}$ is now an optimization variable that depends on both the ACS data and the EPI data to be reconstructed. This allows the reconstruction to be more robust against possible imperfections in the ACS data. The extent to which the ACS data is trusted is controlled by the user-selected parameter η. In the limit as $\eta \rightarrow \infty$, the approximate nullspace matrix $\mathbf{N}$ will converge to the fixed matrix from Equation (1).

The second difference is that the final term of Equation (3) now includes structured matrices formed from the ACS data, in addition to the previous structured matrices formed from the EPI data to be reconstructed. By concatenating the ACS data in this way, we are essentially treating the ACS data in the same way that we would treat additional channels in a parallel imaging experiment. Although the ACS data may not have the same contrast as the EPI data to be reconstructed, it has previously been shown that treating multi-contrast information like additional channels in a parallel imaging experiment often leads to improved reconstruction performance.[30] While this improvement has been justified empirically, some level of theoretical justification for this approach can be obtained by modeling different image contrasts as different modulations of some latent image.[40]

Algorithmically, Equation (3) can be minimized using existing algorithms for LORAKS optimization.[18,36,37,45] In particular, it is not hard to show that the solution to Equation (3) can be equivalently obtained by solving:

$$
\begin{aligned}
\left\{ \widehat{\mathbf{k}}^{+}, \ \widehat{\mathbf{k}}^{-} \right\} = & \operatorname*{argmin}_{\{\mathbf{k}^{+}, \mathbf{k}^{-}\}} J_{(C-p)} \left(\begin{bmatrix} \mathscr{P}_C(\mathbf{k}^{+}) \\ \mathscr{P}_C(\mathbf{k}^{-}) \\ \sqrt{\eta}\mathscr{P}_C(\mathbf{k}_{acs}^{+}) \\ \sqrt{\eta}\mathscr{P}_C(\mathbf{k}_{acs}^{-}) \end{bmatrix} \right) \\
& + \lambda J_r \left(\left[\mathscr{P}_S(\mathbf{k}^{+}) \ \mathscr{P}_S(\mathbf{k}^{-}) \ \mathscr{P}_S(\mathbf{k}_{acs}^{+}) \ \mathscr{P}_S(\mathbf{k}_{acs}^{-}) \right] \right),
\end{aligned}
$$

(4)

where $J_{(C-p)}(\cdot)$ is the same as $J_r(\cdot)$ but replacing the rank parameter r with the rank parameter $(C-p)$, where C is the number of columns of the LORAKS matrix formed by $\mathscr{P}_C(\cdot)$. Equation (4) is convenient because it takes the same form as previous LORAKS optimization problems involving multiple $J_r(\cdot)$ terms.[36] For this paper, we use a multiplicative half-quadratic majorize-minimize algorithm to minimize this objective function,[45] which takes advantage of FFT-based matrix multiplications to improve computational complexity.[46]

The RAC-LORAKS solution is obtained through the optimization of a nonconvex cost function. As such, the algorithm has the potential to converge to an undesirable local minimum.

For the results shown in this paper, we initialize RAC-LORAKS using a naive initialization with minimal processing cost as explained in the next section. Other choices could potentially result in even higher performance, but are not considered here.

3 | METHODS

3.1 | Datasets used for evaluation

As described below, we evaluated the characteristics of RAC-LORAKS using data from several different contexts. All in vivo data were acquired under IRB-approved written informed consent.

3.1.1 | Gradient-echo EPI brain data

In one set of experiments, we acquired in vivo human brain data using a gradient-echo EPI sequence with parameters that are somewhat similar to a BOLD fMRI experiment. Data were acquired on a Siemens 3T Prisma Fit scanner using a standard 32-channel receiver array. The data were acquired using FOV = 220 mm × 220 mm; matrix size = 128 × 128; slice thickness = 3 mm; and TR = 2 seconds In one subject, data were acquired without acceleration ($R = 1$) with TE = 47 ms. From this same subject, data were also acquired for parallel imaging acceleration factors of $R = 2, 3, 4$ with TE = 35 ms. In a second subject, data were acquired for parallel imaging acceleration factors of $R = 5, 6$ with TE = 35 ms. In all cases, fully sampled ACS data were acquired using the same interleaved $2R$-shot EPI prescan as used for DPG.[11]

The previous datasets were acquired with a conventional axial slice orientation. However, because Nyquist ghost problems tend to be more extreme with oblique acquisitions due to concomitant fields that can produce substantial nonlinear 2D phase differences between positive and negative readout polarities,[10,11,47-49] we also acquired an additional dataset with a double-oblique slice orientation from a third subject to test performance in a more challenging scenario. The slice orientation in this case is nonstandard and likely difficult to interpret for many readers, so we have depicted its position in Supporting Information Figure S2. For this case, the data were acquired with TR = 2.08 seconds and TE = 35 ms for parallel imaging acceleration factors of $R = 1, 2, 3, 4, 5, 6$.

3.1.2 | Diffusion-encoded EPI brain data

In another set of experiments, we acquired in vivo human brain data using a diffusion-encoded spin-echo EPI sequence. Diffusion EPI data might be considered more challenging than the previous gradient-echo EPI data, due to the fact that diffusion MRI data usually suffers from random image-to-image phase variations, and can also have lower SNR than gradient-echo EPI. In addition, the rapid switching of strong diffusion gradients can introduce substantial additional eddy current effects that can cause systematic differences between the ACS data and the diffusion EPI data if they are acquired with different diffusion gradient settings.[50]

A first diffusion dataset was acquired on a Siemens 3T Prisma Fit scanner using a standard 32-channel receiver array. For the sake of computational complexity, this data was subsequently reduced to 16 channels using standard coil-compression techniques. The data were acquired using FOV = 220 mm × 220 mm; matrix size = 220 × 220; slice thickness = 5 mm; TR = 2.8 seconds; TE = 63 ms; b-values of 0 sec/mm^2 and 1000 sec/mm^2; 6 diffusion encoding directions; parallel imaging acceleration factor $R = 3$; and 6/8ths partial Fourier sampling. ACS data were acquired using the same interleaved $2R$-shot EPI prescan as used for DPG,[11] except that the data were acquired with lower resolution along the phase encoding dimension (ie, we only acquired 45 phase-encoding lines for the ACS data). Due to the random phase variations associated with diffusion encoding gradients, the ACS data were acquired without diffusion weighting, which means that the ACS data have very different contrast characteristics from the EPI data. To show results across a broader range of acceleration factors, a second set of acquisitions was performed with $R = 2, 3, 4, 5$. Other parameters were identical to the previous case, except for matrix size = 110 × 110; slice thickness = 2 mm; TR = 11.4 seconds; TE = 73 ms; and fully sampled ACS data.

3.1.3 | Cardiac EPI data

In a third set of experiments, we acquired in vivo human cardiac data during diastole using a spin-echo EPI sequence with parameters that are typical for a myocardial arterial spin labeling experiment.[51] Data were acquired on a GE 3T Signa HDx scanner with an 8-channel cardiac coil. The acquisition used FOV = 280 mm × 140 mm; matrix size = 128 × 64; slice thickness = 10 mm; TR = 55 ms; TE = 32.9 ms; velocity cutoff = 5 cm/s; no parallel imaging acceleration ($R = 1$); and 5/8ths partial Fourier sampling. ACS data were acquired using the same interleaved $2R$-shot EPI prescan as used for DPG,[11] but with 5/8ths partial Fourier sampling. Data were acquired with a double-oblique slice orientation to achieve a mid-short axis view.

3.2 | Simulations

In addition to in vivo data, the different methods were also evaluated using simulations where a gold standard was present.
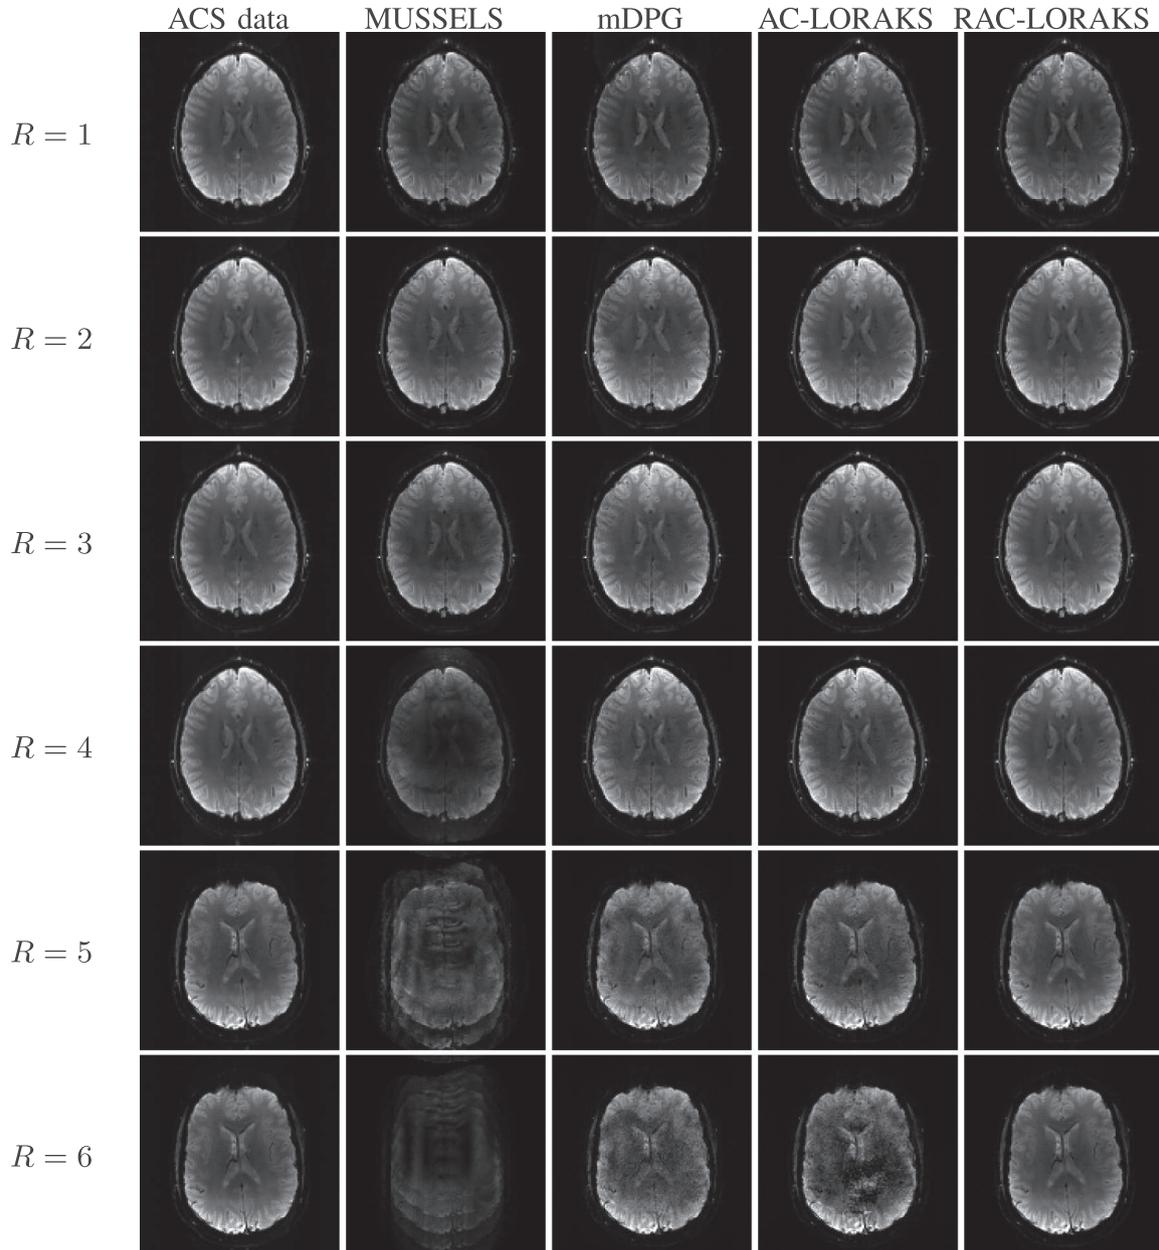
FIGURE 1 ACS data and reconstruction results for in vivo gradient-echo EPI brain data with an axial slice orientation for different parallel imaging acceleration factors. Note that the first four acceleration factors ($R = 1$-4) were acquired from one subject during a single scan session while the last two acceleration factors ($R = 5$, 6) were acquired from a different subject on a different day, which explains the visual discontinuity between these cases

To form a gold standard with realistic EPI characteristics, we took two in vivo gradient-echo EPI brain datasets (as described in Section 3.1.1) with axial slice orientation and $R = 1$ from the same scan session, and reconstructed them both using SENSE. Each gradient polarity was reconstructed separately, providing a realistic representation of typical interpolarity image differences. This procedure provides two sets of fully sampled multi-channel dual-polarity gold standard images. One of these sets was used for ACS data, while the other was undersampled (including parallel imaging acceleration, along with interleaving the data from positive and negative gradient polarities) to simulate EPI data. These datasets were acquired roughly

5 minutes apart, allowing time for mismatches to evolve. Since ghost correction is frequently more difficult for EPI datasets with 2D nonlinear phase differences between the two polarities, we applied an additional 2D nonlinear phase pattern to make the problem more challenging. This additional phase difference was designed to be roughly 3×-larger than we observed in the real data from Figure 2.

In a first set of simulations, to mimic the situation where a localized image feature is different between the ACS data and EPI data (eg, as may happen in a dynamic experiment), we added a Gaussian-shaped additive image hyperintensity to the EPI data that we did not add to the ACS data. The
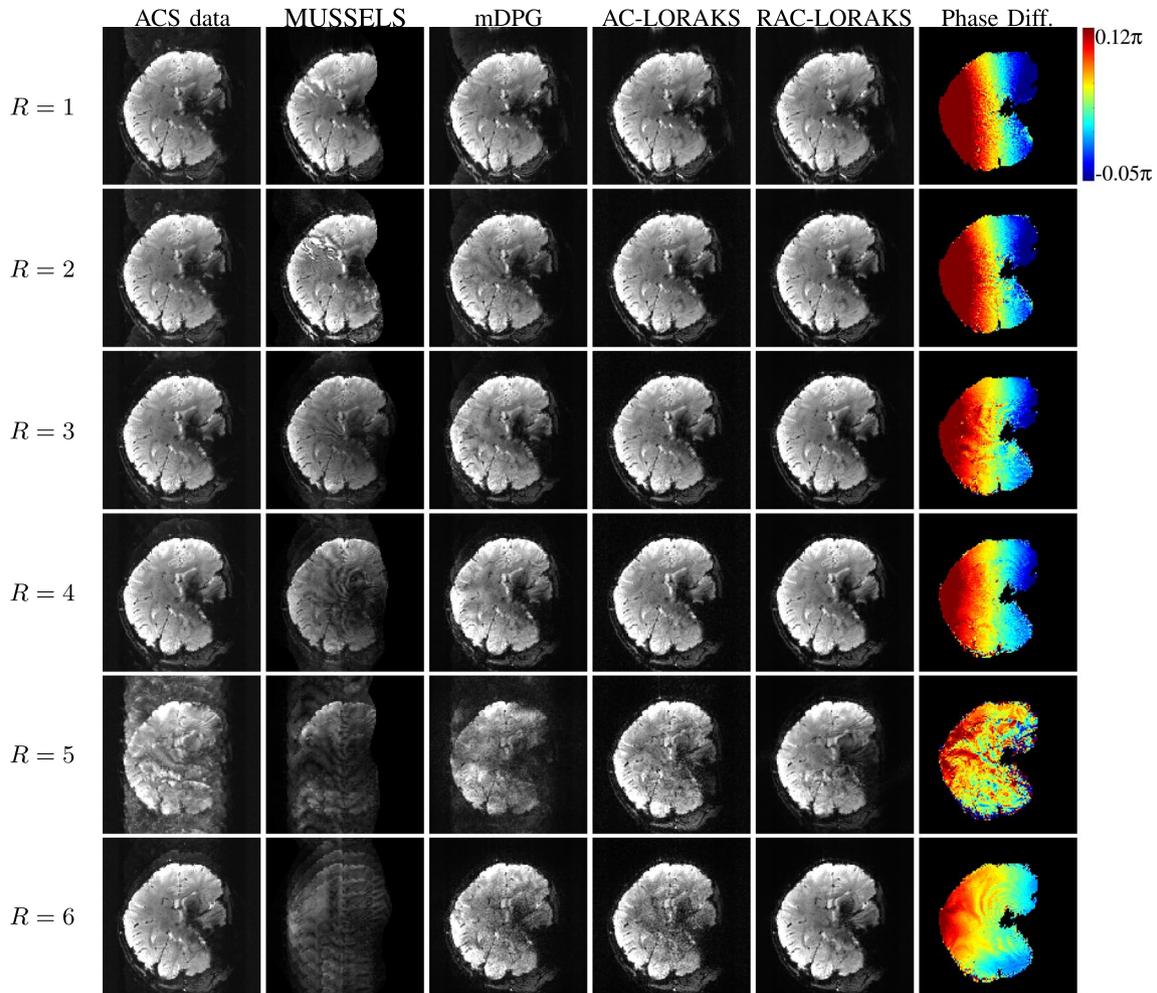
FIGURE 2 ACS data and reconstruction results for in vivo gradient-echo EPI brain data with a double-oblique slice orientation for different parallel imaging acceleration factors. For reference, we also show the interpolarity phase difference as estimated from a coil-combined RAC-LORAKS result. The degree of phase nonlinearity is an indicator of how difficult ghost correction is expected to be. As can be seen, complicated 2D nonlinear phase differences are present in many of these cases

hyperintensity was designed to follow both the coil sensitivity maps (obtained by applying ESPIRiT[34]) and the phase characteristics of the original data. We also performed simulations with these two datasets interchanged, that is, with the hyperintensity in the ACS data but not in the EPI data.

In another set of simulations, to mimic the situation where the ACS data and EPI data have very different contrasts, we inverted the magnitude image for the ACS data to create an image with different contrast,[52] while still following the coil sensitivity maps and the phase characteristics of the original data.

In addition to performing multi-channel simulations, we also performed a simulation in a very challenging single-channel setting. For this, single-channel data were obtained by a linear combination of the multi-channel data.[53] Single-channel ghost correction is a difficult setting where only a few previous methods have had any success.[16,18] This case is hard because even with unaccelerated data ($R = 1$), each polarity has an effective acceleration factor of $R = 2$ when the data for each

readout gradient polarity are separated, and it can be difficult to reconstruct $R = 2$ data without multi-channel information.

The fully sampled ACS and EPI datasets used for all three simulations are illustrated in Supporting Information Figure S3.

3.3 | Data processing

RAC-LORAKS was applied to perform reconstruction and ghost correction on these datasets. For comparison against existing methods, the datasets were also reconstructed using the previous AC-LORAKS ghost-correction method,[18] DPG,[11] and MUSSELS.[17]

For some of the datasets we consider, the ACS data may be incomplete due to low-resolution ACS acquisition (ie, the first brain EPI diffusion data) or partial Fourier ACS acquisition (ie, the cardiac EPI data). In such cases, we modify

RAC-LORAKS to consider the fully sampled ACS data vectors $\mathbf{k}_{acs}^+$ and $\mathbf{k}_{acs}^-$ as additional variables to be optimized in Equation (4), subject to ACS data consistency constraints.

For RAC-LORAKS and AC-LORAKS, the regularization parameters λ and η were selected manually based on subjective visual inspection reconstruction quality and ghost-reduction performance for in vivo data, and to minimize quantitative error measures for simulated data. The rank-related parameters p and r were selected based on the singular value characteristics of LORAKS matrices formed from the ACS data. The rank parameters were set based on the points at which the singular value curves begin to flatten out, which is a common rank estimation technique for noisy matrices. This decision was made manually (based on visual inspection) for the results shown later in the paper, although fully automatic approaches would also be viable.

DPG is a ghost correction method that treats different gradient polarities like different coils in a parallel imaging experiment, and uses a dual GRAPPA kernel estimated from ACS data for image reconstruction.[11] In order to use DPG for the initialization of RAC-LORAKS, we have adapted DPG to output two sets of images (with calibration based on the raw uncorrected multi-channel ACS data), one for the original $\mathbf{k}^+$ data and one for the original $\mathbf{k}^-$ data. This is different than the original DPG implementation, which applies ACS preprocessing to try and correct for errors in the ACS data, and then directly fuses information from the two polarities together into a single virtual "hybrid" output.[11] This hybrid output can have different magnitude and phase characteristics than the original $\mathbf{k}^+$ and $\mathbf{k}^-$ data, so is not useful as an initialization for RAC-LORAKS. We refer to our adapted version as modified DPG (mDPG) from now on. In some cases, we also compare against the original version of DPG (including the original ACS preprocessing procedure to correct for shot-to-shot variations in the ACS data[11]), although note that such comparisons are necessarily qualitative, since the magnitude and phase characteristics of the hybrid output DPG images do not match the images generated using other methods.

MUSSELS is a structured low-rank matrix recovery method that uses SENSE-type parallel imaging constraints together with nuclear norm regularization to impose low-rank constraints.[17] While MUSSELS was originally developed for multi-shot EPI ghost correction, it can apply equally well to the ghost correction problem associated with different gradient polarities. Sensitivity maps for MUSSELS were estimated by applying ESPIRiT[34] to the same ACS data used for the other methods. The regularization parameter for MUSSELS was selected manually based on subjective visual inspection of reconstruction quality and ghost-reduction performance in the case of in vivo data, or to minimize quantitative error measures for simulated data.

Note that DPG and MUSSELS were both developed for the multi-channel setting. We can adapt DPG to the

single-channel setting in straightforward ways,[18] and we apply this adaptation to the single-channel simulated data. We did not adapt MUSSELS to the single-channel case. Note that the SENSE-based constraints used by MUSSELS would reduce to a simple spatial-domain support constraint in the single-channel case, which is not strong enough to yield good performance results.

For all methods, results were visualized by using a standard square-root sum-of-squares technique to combine the images from different coils and different gradient polarities into a single image. Results from in vivo experiments were evaluated qualitatively, since a gold standard reference was not available in these cases.

Simulation results were evaluated quantitatively using the normalized root mean-squared error (NRMSE):

$$\text{NRMSE} \triangleq \frac{\sqrt{\left\| \hat{\mathbf{k}}^+ - \mathbf{k}_{\text{gold}}^+ \right\|_2^2 + \left\| \hat{\mathbf{k}}^- - \mathbf{k}_{\text{gold}}^- \right\|_2^2}}{\sqrt{\left\| \mathbf{k}_{\text{gold}}^+ \right\|_2^2 + \left\| \mathbf{k}_{\text{gold}}^- \right\|_2^2}}, \quad (5)$$

where $\mathbf{k}_{\text{gold}}^+$ and $\mathbf{k}_{\text{gold}}^-$ are, respectively, the gold standard values for the positive and negative gradient polarities. We also plotted Fourier error spectrum plots (ESPs) to gain further insight into how the errors were distributed across different spatial resolutions scales.[54] An ESP is designed to reveal the spectral characteristics of the error, and for example, can discriminate between methods that make more errors in the low-resolution features of an image vs methods that make more errors in high-resolution features.

4 | RESULTS

Figure 1 shows ACS data and reconstruction results from the in vivo gradient-echo EPI brain data with an axial slice orientation. The ACS data in this case do not have strong artifacts, although close inspection does reveal that ACS ghost artifacts are present. This can be further appreciated in Supporting Information Figure S4 where the same images are shown with amplified image intensity to highlight ghost characteristics in the image background. As can be seen, all ghost correction methods work well at smaller acceleration factors, although performance begins to degrade at larger acceleration factors. We observe that, compared to other methods, the visual quality of the MUSSELS reconstruction seems to degrade most rapidly as a function of acceleration factor, which is consistent with previous observations.[18] The mDPG method had qualitatively better performance than MUSSELS in this case. However, a close inspection of the images reveals that the mDPG results are not entirely ghost-free even for the unaccelerated ($R = 1$) case. This may be expected due to the artifacts and mismatches that are present in the ACS data.

Although mDPG does not attempt to correct the ACS artifacts, it should be noted that the original DPG method does try to correct them through pre-processing. Results showing the qualitative performance of the original DPG method are shown in Supporting Information Figure S5, where we observe that the ghost artifacts still exist, though as expected, are less prominent than were observed for mDPG. In spite of the ACS artifacts, the AC-LORAKS reconstruction still has good performance at low acceleration factors and does a good job of suppressing ghosts in the background regions of the image at all acceleration factors, although exhibits substantial degradation in image quality at the highest acceleration factors (with artifacts similar to those observed for highly accelerated parallel imaging reconstructions). However, the RAC-LORAKS reconstruction appears to have much higher quality than the other methods, even at very high acceleration factors like $R = 6$. (Note that when $R = 6$, the effective acceleration factor is $R = 12$ when each readout gradient polarity is considered separately. This leads to a highly ill-posed inverse problem.)

Figure 2 shows results from the in vivo gradient-echo EPI brain data with a double-oblique slice orientation. This case is more challenging than the previous one due to the complicated nonlinear 2D spatial phase differences we observed between data acquired with positive and negative polarities

(as visualized in the last column of Figure 2), the proximity to air-tissue interfaces that result in substantial magnetic field inhomogeneity effects, as well as more substantial ghosting artifacts present in the ACS data. Note that the ACS data corresponding to the $R = 5$ case is particularly corrupted, which can be attributed to the unpredictable shot-to-shot variations that frequently occur in these kinds of multi-shot acquisitions. Despite the more extreme scenario, the different ghost reconstruction methods have similar characteristics to those observed in the previous case, with RAC-LORAKS appearing to demonstrate the cleanest overall results.

Figure 3 shows reconstruction results from the first set of multi-channel simulations (with similar contrast between ACS and EPI data, but with a hyperintensity added to the EPI data). Quantitative NRMSE values are reported in Table 1 with corresponding ESPs shown in Figure 4. Qualitatively, the results from Figure 3 have similar characteristics to the results observed with in vivo data. Notably, RAC-LORAKS is able to consistently reconstruct a high-quality image that bears close resemblance to the gold standard image, while methods like mDPG and AC-LORAKS have artifacts due to the small mismatches between the ACS and EPI data. The visual assessment of reconstruction quality matches well with the quantitative NRMSE assessment shown in Table 1. AC-LORAKS and RAC-LORAKS have a similar performance
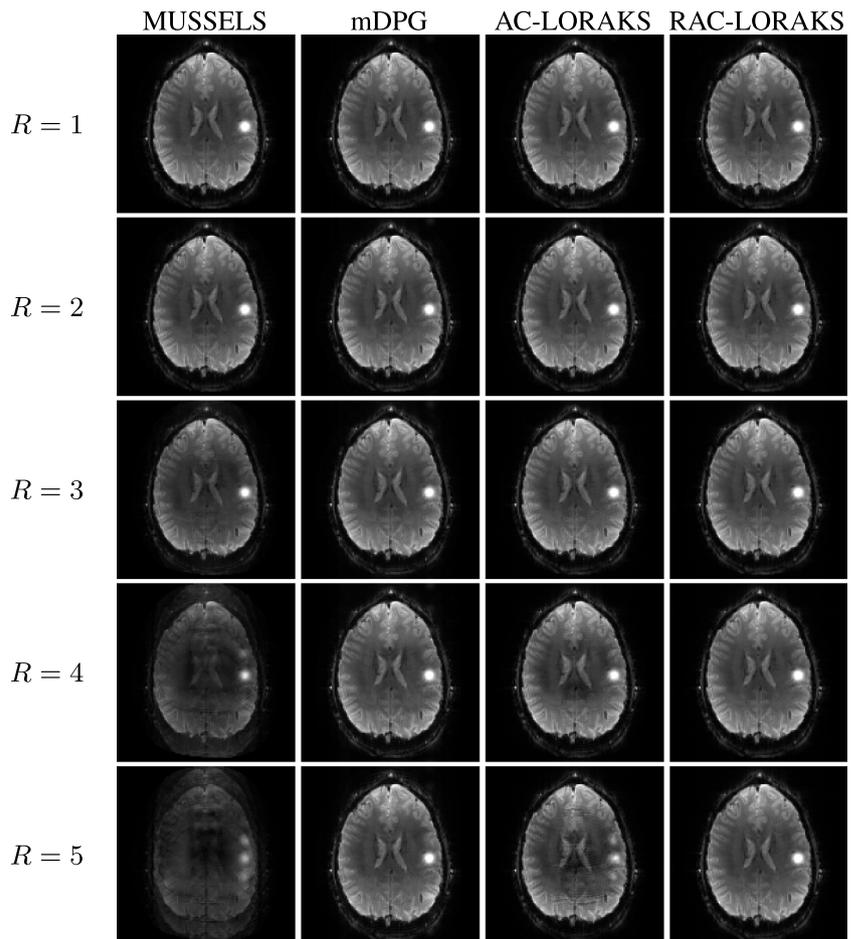

FIGURE 3 Reconstruction results for the first set of multi-channel simulations (with similar contrast between ACS and EPI data, but with a hyperintensity added to the EPI data) with different parallel imaging acceleration factors
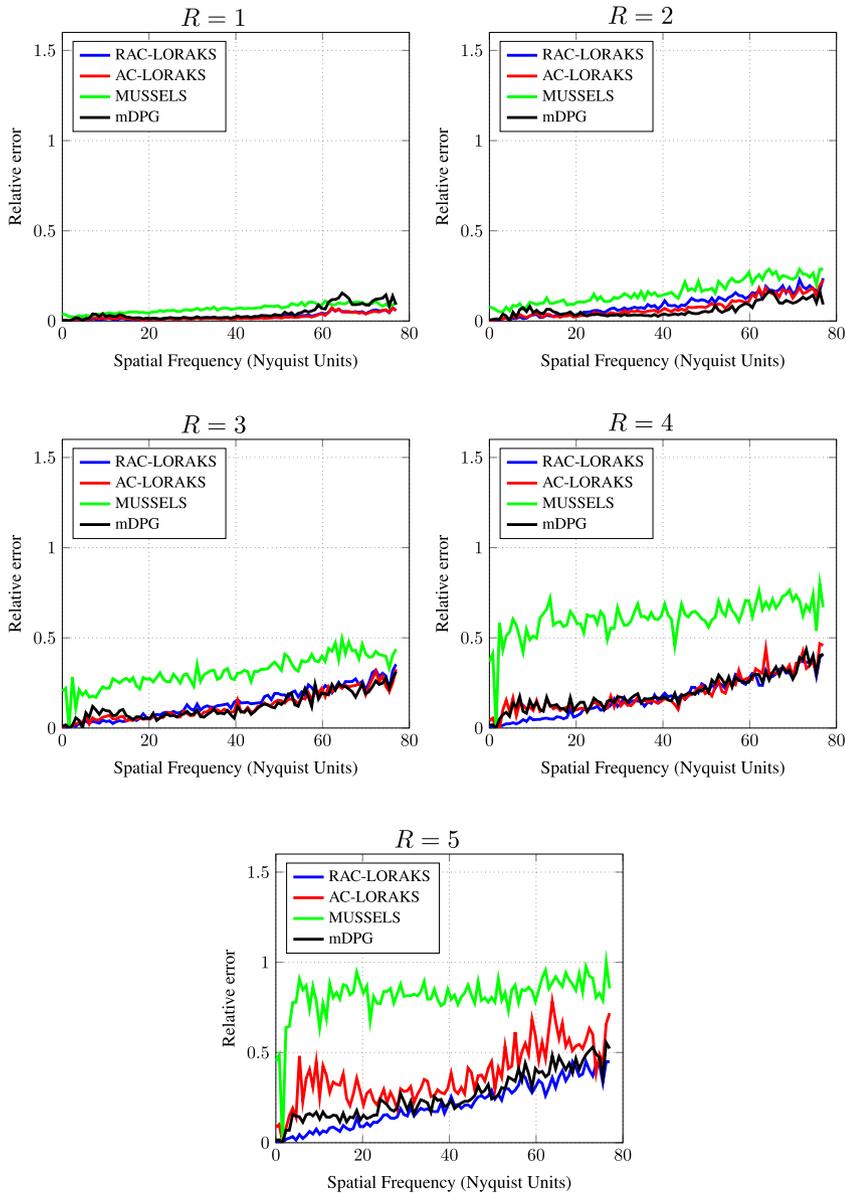
FIGURE 4 ESPs for the multi-channel simulation results shown in Figure 3. The vertical axis of each ESP uses a consistent range to enable comparisons between different acceleration factors

TABLE 1 NRMSEs for the multi-channel simulation results shown in Figure 3

	MUSSELS	mDPG	AC-LORAKS	RAC-LORAKS
$R = 1$	0.059	0.024	**0.016**	0.020
$R = 2$	0.104	0.045	**0.035**	0.042
$R = 3$	0.271	0.083	0.056	**0.055**
$R = 4$	0.572	0.127	0.132	**0.064**
$R = 5$	0.741	0.161	0.269	**0.085**

Note: For each acceleration factor, the smallest values are highlighted in bold.

at $R = 1$ and 2, with RAC-LORAKS having the best performance at high acceleration factors.

Reconstructions were also performed using the original DPG formulation as shown in Supporting Information Figure S6. In this case, DPG has similar ghost artifacts to mDPG, which is expected because there are no artifacts in the ACS data, while there is a problematic mismatch between the ACS data and the EPI data that neither DPG nor mDPG address. Notably, for both DPG and mDPG, we observe aliasing artifacts that seem to be associated with the hyperintensity that was present in the EPI data but was not in the ACS data.

The ESP plots in Figure 4 enable a more nuanced analysis. These results suggest that RAC-LORAKS has good (ie, among the best, even if it is not always the best) performance at all spatial frequencies, meaning that it is good at reconstructing image features across the whole range of resolution scales.

Supporting Information Figure S7 shows a similar simulation result to that shown in Figure 3, with the main difference being that the previous EPI images (with the hyperintensity) were used as ACS data and the previous ACS images (without the hyperintensity) were used to generate EPI data. Consistent with the previous case, we observe good

performance for RAC-LORAKS, and do not observe the features of the hyperintensity being erroneously transferred into the RAC-LORAKS reconstruction results.

Supporting Information Figure S8 and Table S1 show simulation results for the case where the ACS data has an even more substantial contrast difference (ie, inverted contrast) with respect to the EPI data. For this case, we observe a degradation in performance for all methods compared to the previous cases, although RAC-LORAKS still showed the best overall qualitative and quantitative performance. This result suggests that RAC-LORAKS may have better performance when the contrast is similar between the ACS and EPI data, although can still provide benefits when the contrast difference is substantial.

Figure 5 shows reconstruction results from the single-channel simulation, with quantitative NRMSE values reported in Supporting Information Table S2. While previous work[18] reported that mDPG and AC-LORAKS can be reasonably successful for single-channel data with $R = 1$ when the ACS data is pristine, our new results demonstrate that this performance can be sensitive to the quality of the ACS data. In particular, we observe strong ghost artifacts for both of these methods, even though we do observe that the AC-LORAKS reconstruction has successfully suppressed ghost artifacts in the image background (outside of the support of the true image). In contrast, RAC-LORAKS is substantially more successful for $R = 1$. Notably, RAC-LORAKS also performed well for the even more challenging $R = 2$ case, unlike the other methods. For reference, note that even with high-quality ACS data, the previous AC-LORAKS method did not yield good results with similar single-channel $R = 2$ data.[18]

Figure 6 shows reconstruction results from the first set of in vivo diffusion EPI brain data, including a 10× intensity amplification to highlight the ghost characteristics. As can be seen, the ACS data has ghost artifacts in all cases, and both MUSSELS and mDPG reconstructions also exhibit

unsuppressed ghosting artifacts. On the other hand, both AC-LORAKS and RAC-LORAKS are relatively ghost-free in this example and have only minor differences from one another (it might be argued that the RAC-LORAKS result has a slightly less-noisy appearance than the AC-LORAKS result, but if so, this difference is very subtle). While this result does not demonstrate an obvious advantage for RAC-LORAKS over AC-LORAKS, it should be observed that this diffusion result is at least consistent with the previous gradient-echo EPI data results, in which we also did not observe a substantial difference between RAC-LORAKS and AC-LORAKS when $R = 3$. In addition, this case involves a very substantial contrast difference between the ACS data and the EPI data. This difference does not appear to have adversely affected the performance characteristics of these methods in substantial ways.

Figure 7 shows reconstruction results from the second set of in vivo diffusion EPI brain acquisitions (with different acceleration factors), with zoom-ins shown in Supporting Information Figure S9 for improved visibility. Consistent with the results shown for the gradient-echo EPI data in Figure 1, we observe that all methods perform well for low acceleration factors. As the acceleration factor increases, the performance of each method degrades, with RAC-LORAKS showing a lower qualitative degradation in comparison to the other methods at the very high acceleration factors $R = 4, 5$. Note that at high acceleration factors (eg, $R = 4, 5$) the reconstruction quality for RAC-LORAKS is not quite as good as for the gradient-echo EPI dataset shown in Figure 1. We believe that this should be expected, since as mentioned before, diffusion EPI data can be considered more challenging than the gradient-echo EPI data due to SNR issues, eddy current effects, motion-induced phase effects, and contrast mismatches between the ACS and EPI datasets.

Finally, Figure 8 shows results from the in vivo cardiac EPI data. While this data was not accelerated ($R = 1$), this case is challenging because of the double-oblique slice
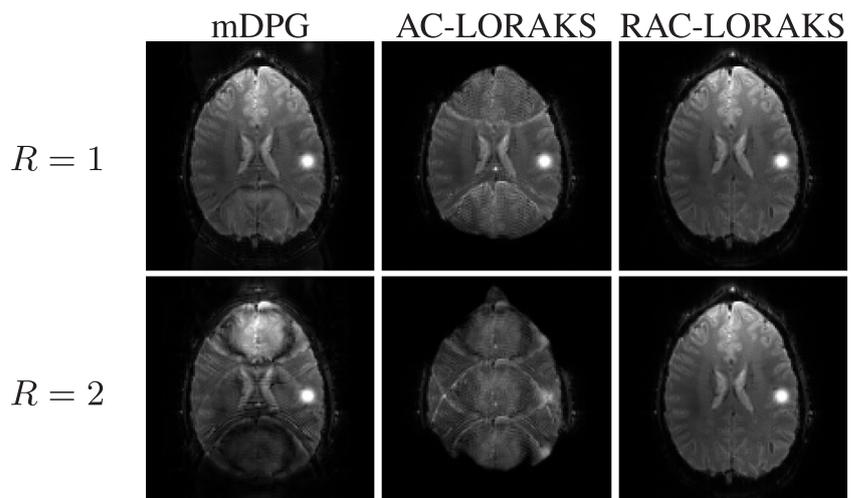

FIGURE 5 Reconstruction results for single-channel simulated data with different acceleration factors
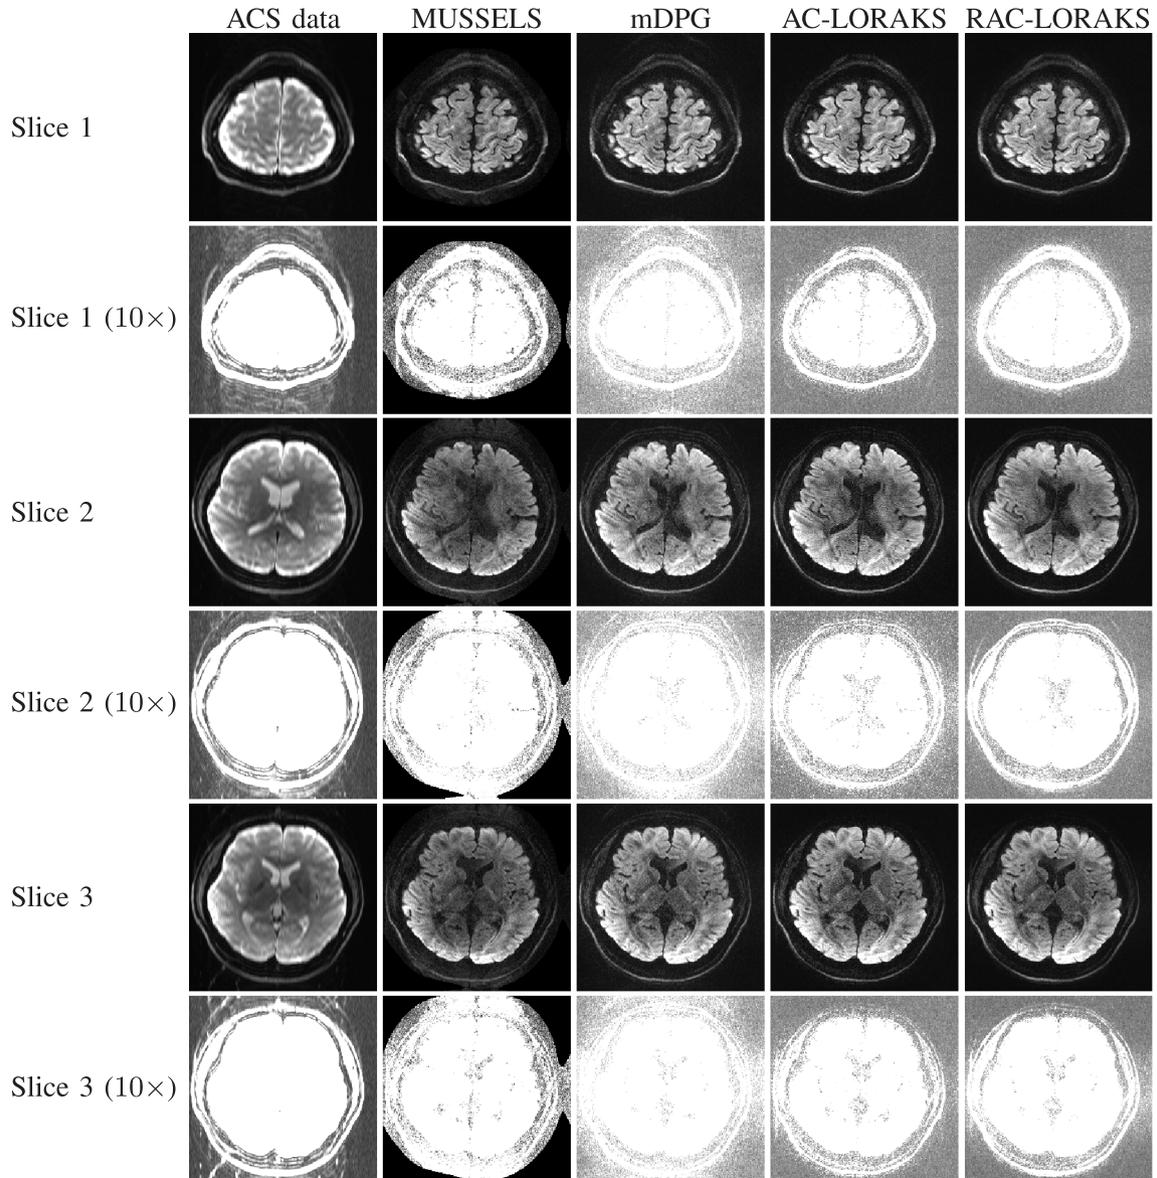
FIGURE 6 ACS data and reconstruction results for three representative slices from in vivo diffusion brain data ($R = 3$). A 10× intensity amplification is also shown for each slice to better highlight the ghosting characteristics

orientation as well as the substantial artifacts present in the ACS data resulting from cardiac motion-induced shot-to-shot variations. In addition, this case can also be challenging for SENSE-based methods (like MUSSELS), due to the use of a small FOV with aliasing. When aliasing is present within the FOV, it violates the standard SENSE modeling assumption of one sensitivity map value per spatial location, which generally leads to artifacts if not properly accounted for. The results demonstrate that both MUSSELS and mDPG have substantial residual ghosting artifacts, which might not be surprising given the high degree of corruption that is present in the ACS data. On the other hand, both AC-LORAKS and RAC-LORAKS are more successful at suppressing the ghosts. Without a gold standard reference, it is hard to establish definitively whether AC-LORAKS or RAC-LORAKS is better in this example, although we believe that the RAC-LORAKS

result demonstrates slightly less ghosting than AC-LORAKS, particularly on the left side of the image where the ACS data and mDPG both have particularly strong ghost artifacts.

5 | DISCUSSION

The results in the previous section demonstrated that, in the presence of imperfect ACS data, RAC-LORAKS frequently offers similar or better performance to the previous AC-LORAKS ghost correction method that it generalizes, while both of these methods perform substantially better than methods like MUSSELS or DPG. We also observed that RAC-LORAKS appears to have the biggest advantage over AC-LORAKS in scenarios where the parallel imaging acceleration factor was high. For these cases, we observed that
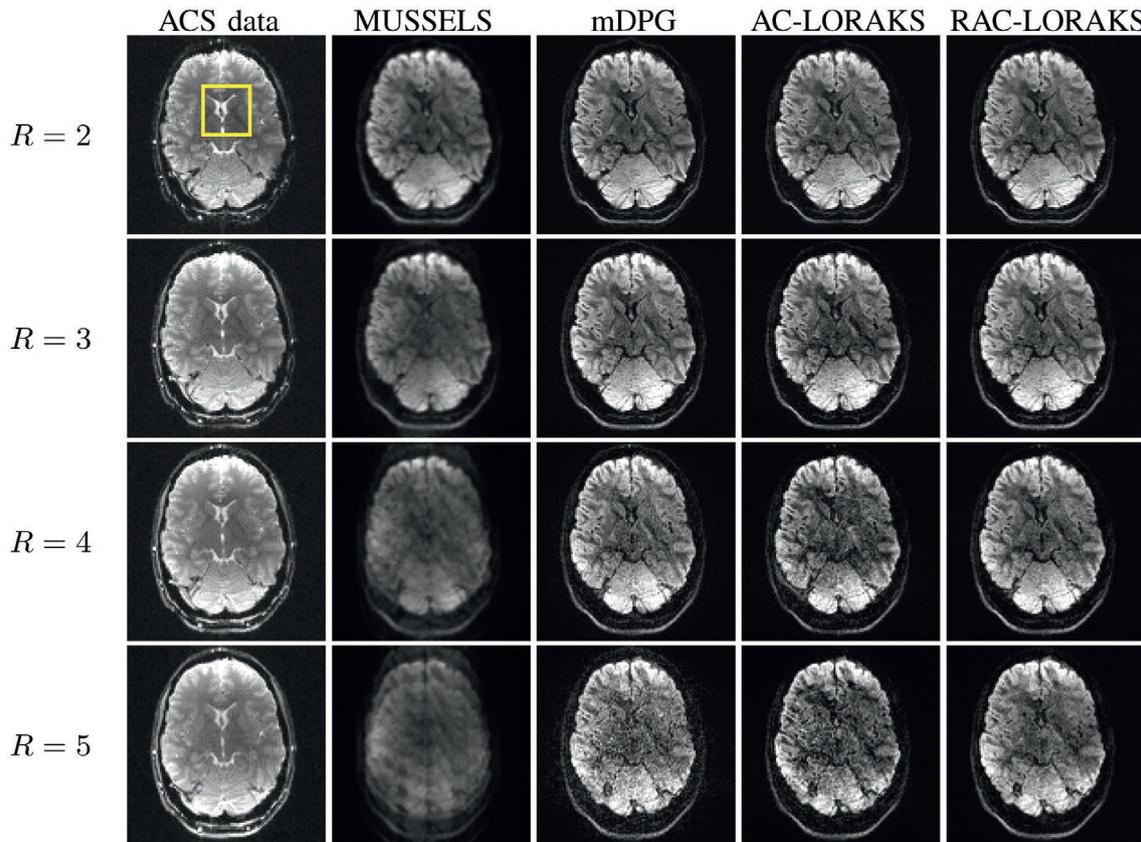
FIGURE 7 ACS data and reconstruction results for in vivo diffusion EPI brain data for different parallel imaging acceleration factors. For improved visualization, zoomed-in versions of these results (corresponding to the spatial region marked with a yellow rectangle in the first column and first row) are shown in Supporting Information Figure S9. It should be noted that the subject appears to have slightly moved between scans, so that there is not perfect correspondence between anatomical image features across different acceleration factors

RAC-LORAKS was able to mitigate ghost artifacts both inside and outside the support of the original image, while AC-LORAKS was only able to mitigate ghost artifacts outside the support but not inside. This advantage for RAC-LORAKS is likely the result of its improved robustness to ACS errors combined with the multi-contrast linear predictability constraints which help to make the reconstruction problem less ill-posed. However, it should be noted that RAC-LORAKS has one more regularization parameter than AC-LORAKS (ie, η, which controls the level of trust placed in the information from the ACS data). In our experience, manual tuning of this parameter is not hard (ie, we always started from the small value $\eta = 10^{-3}$, and frequently did not have to modify this value to achieve satisfying results). The method would be easier to use if the selection of η were automated.

Both RAC-LORAKS and AC-LORAKS also depend on the choice of rank parameters, and as described previously, the results shown in this work made a heuristic choice based on the empirical rank characteristics of the ACS data. Even though the low-rank characteristics of the structured matrices might vary between the ACS data and the acquired EPI data due to systematic phenomena (eg, thermal noise, subject motion, respiration, artifacts in the ACS data, etc.), we have not

observed major problems associated with inappropriate rank selection in our empirical results. This might be expected, based on the observation that LORAKS reconstruction results are frequently not very sensitive to small variations in the rank parameter.[36,37] Nevertheless, the development of improved automatic RAC-LORAKS parameter selection methods would be an interesting topic for future work.

Although RAC-LORAKS offers good performance, it should be noted that our current implementation of RAC-LORAKS can be more computationally expensive than existing methods. For example, for the results shown with $R = 1$ in Figure 1, RAC-LORAKS used $\approx$45 min of reconstruction time, while MUSSELS, mDPG, and AC-LORAKS, respectively, used $\approx$15, $\approx$2, and $\approx$100 minutes. All methods were implemented in MATLAB on a standard desktop computer with an Intel Xeon E5-1603 2.8 GHz quad core CPU processor and 32GB of RAM. While this relatively long computation time may be a concern, it should be noted that we are reporting the results of a simple proof-of-principle implementation, and we did not spend much time to optimize the computational efficiency of this approach. We believe that major improvements may be possible by leveraging better computational hardware, smarter algorithms, and more efficient implementations. Given
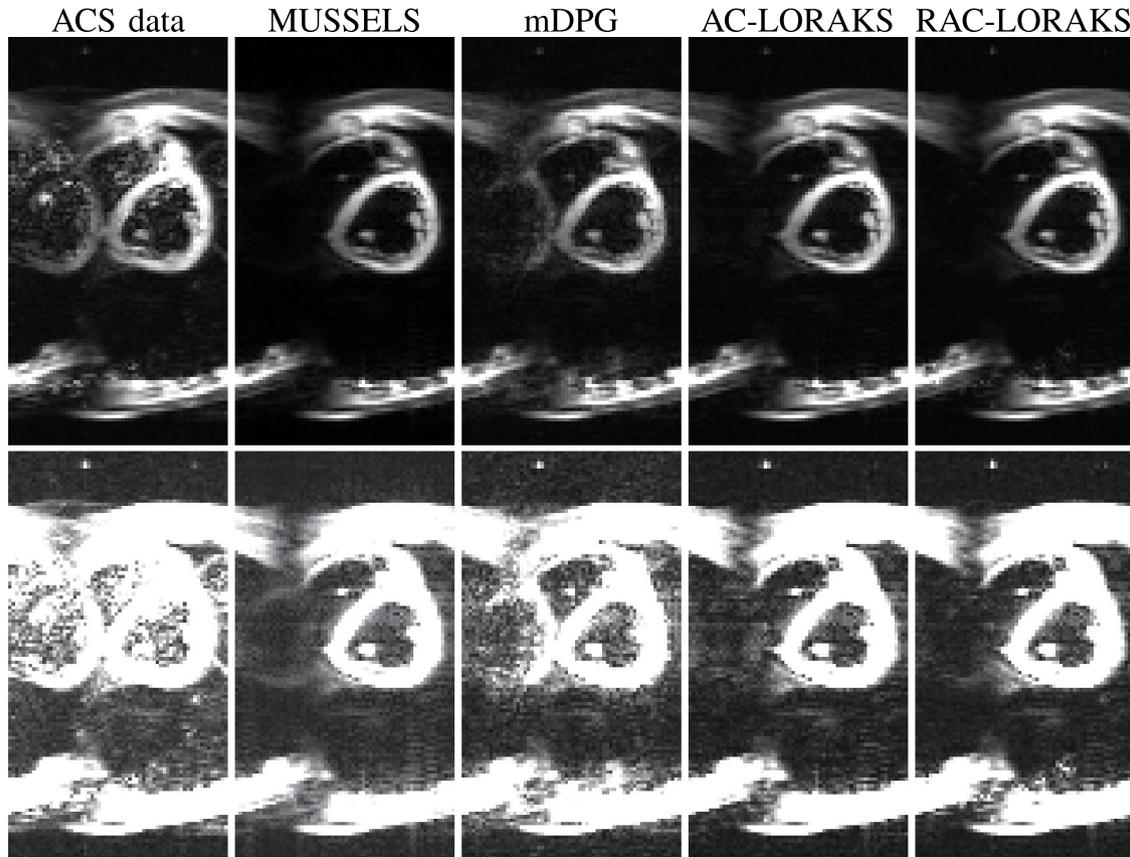
ACS data   MUSSELS   mDPG   AC-LORAKS RAC-LORAKS

FIGURE 8 ACS data and reconstruction results for unaccelerated in vivo cardiac EPI data. The two rows show the same results, but the second row has 5× intensity amplification to better highlight the ghosting characteristics

the reconstruction performance offered by RAC-LORAKS, we believe that improving its computational performance is a promising topic for future research. However, RAC-LORAKS is notably faster than AC-LORAKS, and it appears that this speed difference results from the fact that RAC-LORAKS has consistently faster convergence than AC-LORAKS in this setting. The reason for this faster convergence is unclear at this stage, although we believe that a detailed analysis of convergence characteristics is beyond the scope of the present paper.

While this paper focused on EPI ghost correction for standard single-slice excitation, we believe that the extension of these ideas to simultaneous multi-slice EPI acquisitions (similar to Refs. [9,20,55,56]) is a very promising research direction.

Finally, although the techniques we developed in this work were described and evaluated in the context of EPI ghost correction, we believe that the overall approach is likely to be useful across a wide range of parallel imaging applications, particularly those for which the measured ACS data are not adequate to resolve all of the reconstruction artifacts. Specifically, we believe that the key principles employed by RAC-LORAKS (ie, using structured low-rank matrix methods to avoid placing complete trust in the accuracy of ACS data, and leveraging ACS data to provide additional information in a multi-contrast framework)

are both novel ideas that are applicable to arbitrary image reconstructions involving ACS data, and are not exclusive to ghost correction settings. In addition, we are encouraged by the high-quality reconstruction results that RAC-LORAKS produces even in very highly accelerated scans. These results suggest to us that there may be value in exploring the usefulness of RAC-LORAKS to other parallel imaging experiments in future work.

6 | CONCLUSIONS

This work proposed and evaluated RAC-LORAKS, a new structured low-rank matrix method for EPI ghost correction that integrates multiple constraints (including parallel imaging constraints, support constraints, phase constraints, and inter-image linear predictability constraints) to not only mitigate artifacts resulting from imperfect ACS data and Nyquist ghosts, but also accounting for partial Fourier acquisition and reducing parallel imaging artifacts and noise in an integrated fashion. RAC-LORAKS uses ACS data and k-space domain linear predictive modeling to stabilize the solution of the ill-posed inverse problem, and was observed to offer advantages relative to state-of-the-art ghost correction methods like AC-LORAKS, DPG, and MUSSELS.

ORCID

Rodrigo A. Lobos https://orcid.org/0000-0003-0425-8976
Ahsan Javed https://orcid.org/0000-0003-1311-1247
Congyu Liao https://orcid.org/0000-0003-2270-276X
Krishna S. Nayak https://orcid.org/0000-0001-5735-3550
Justin P. Haldar https://orcid.org/0000-0002-1838-0211

REFERENCES

1. Stelhing MK, Turner R, Mansfield P. Echo-planar imaging: magnetic resonance imaging in a fraction of a second. *Science*. 1991;254:43-50.
2. Bernstein MA, King KF, Zhou XJ. *Handbook of MRI Pulse Sequences*. Burlington, VT: Elsevier Academic Press; 2004.
3. Bruder H, Fischer H, Reinfelder HE, Schmitt F. Image reconstruction for echo planar imaging with nonequidistant k-space sampling. *Magn Reson Med*. 1992;23:311-323.
4. Hu X, Le TH. Artifact reduction in EPI with phase-encoded reference scan. *Magn Reson Med*. 1996;36:166-171.
5. Buonocore MH, Gao L. Ghost artifact reduction for echo planar imaging using image phase correction. *Magn Reson Med*. 1997;38:89-100.
6. Chen NK, Wyrwicz AM. Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction. *Magn Reson Med*. 2004;51:1247-1253.
7. Skare S, Clayton D, Newbould R, Moseley M, Bammer R. A fast and robust minimum entropy based non-interactive Nyquist ghost correction algorithm. *Proc Int Soc Magn Reson Med*. 2006;2349.
8. Xiang QS, Ye FQ. Correction for geometric distortion and N/2 ghosting in EPI by phase labeling for additional coordinate encoding (PLACE). *Magn Reson Med*. 2007;57:731-741.
9. Hoge WS, Tan H, Kraft RA. Robust EPI Nyquist ghost elimination via spatial and temporal encoding. *Magn Reson Med*. 2010;64:1781-1791.
10. Xu D, King KF, Zur Y, Hinks RS. Robust 2D phase correction for echo planar imaging under a tight field-of-view. *Magn Reson Med*. 2010;64:1800-1813.
11. Hoge WS, Polimeni JR. Dual-polarity GRAPPA for simultaneous reconstruction and ghost correction of echo planar imaging data. *Magn Reson Med*. 2016;76:32-44.
12. Chang HC, Chen NK. Joint correction of Nyquist artifact and minuscule motion-induced aliasing artifact in interleaved diffusion weighted EPI data using a composite two-dimensional phase correction procedure. *Magn Reson Imaging*. 2016;34:974-979.
13. Yarach U, In MH, Chatnuntawech I, et al. Model-based iterative reconstruction for single-shot EPI at 7T. *Magn Reson Med*. 2017;78:2250-2264.

14. Ianni JD, Welch EB, Grissom WA. Ghost reduction in echo-planar imaging by joint reconstruction of images and line-to-line delays and phase errors. *Magn Reson Med*. 2018;79:3114-3121.
15. Xie VB, Lyu M, Liu Y, Feng Y, Wu EX. Robust EPI Nyquist ghost removal by incorporating phase error correction with sensitivity encoding (PEC-SENSE). *Magn Reson Med*. 2018;79:943-951.
16. Lee J, Jin KH, Ye JC. Reference-free EPI Nyquist ghost correction using annihilating filter-based low rank Hankel matrix for k-space interpolation. *Magn Reson Med*. 2016;76:1775-1789.
17. Mani M, Jacob M, Kelley D, Magnotta V. Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS). *Magn Reson Med*. 2017;78:494-507.
18. Lobos RA, Kim TH, Hoge WS, Haldar JP. Navigator-free EPI ghost correction with structured low-rank matrix models: new theory and methods. *IEEE Trans Med Imaging*. 2018;37:2390-2402.
19. Lyu M, Barth M, Xie VB, et al. Robust SENSE reconstruction of simultaneous multislice EPI with low-rank enhanced coil sensitivity calibration and slice-dependent 2D Nyquist ghost correction. *Magn Reson Med*. 2018;80:1376-1390.
20. Liu Y, Lyu M, Barth M, et al. PEC-GRAPPA reconstruction of simultaneous multislice EPI with slice-dependent 2D Nyquist ghost correction. *Magn Reson Med*. 2019;81:1924-1934.
21. Haldar JP. Autocalibrated LORAKS for fast constrained MRI reconstruction. *Proc IEEE Int Symp Biomed Imag*. 2015;910-913.
22. Griswold MA, Jakob PM, Heidemann RM, et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*. 2002;47:1202-1210.
23. Lustig M, Pauly JM. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med*. 2010;65:457-471.
24. Zhang J, Liu C, Moseley ME. Parallel reconstruction using null operations. *Magn Reson Med*. 2011;66:1241-1253.
25. Ying L, Sheng J. Joint image reconstruction and sensitivity estimation in SENSE (JSENSE). *Magn Reson Med*. 2007;57:1196-1202.
26. Sheltraw D, Inglis B, Deshpande V, Trumpis M. Simultaneous reduction of two common autocalibration errors GRAPPA EPI time series data. Preprint 2012;arXiv:1208.0972.
27. Baron CA, Beaulieu C. Motion robust GRAPPA for echo-planar imaging. *Magn Reson Med*. 2016;75:1166-1174.
28. 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.
29. Holme HCM, Rosenzweig S, Ong F, Wilke RN, Lustig M, Uecker M. ENLIVE: an efficient nonlinear method for calibrationless and robust parallel imaging. *Sci Rep*. 2019;9:1-13.
30. Bilgic B, Kim TH, Liao C, et al. Improving parallel imaging by jointly reconstructing multi-contrast data. *Magn Reson Med*. 2018;80:619-632.
31. Lobos RA, Javed A, Nayak KS, Hoge WS, Haldar JP. Robust autocalibrated LORAKS for EPI ghost correction. *Proc IEEE Int Symp Biomed Imaging*. 2018;663-666.
32. Lobos RA, Javed A, Nayak KS, Hoge WS, Haldar JP. Robust autocalibrated LORAKS for improved EPI ghost correction with structured low-rank matrix models. *Proc Int Soc Magn Reson Med*. 2018;3533.
33. Liang ZP, Haacke EM, Thomas CW. High-resolution inversion of finite Fourier transform data through a localised polynomial approximation. *Inverse Probl*. 1989;5:831-847.

34. Uecker M, Lai P, Murphy MJ, et al. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn Reson Med*. 2014;71:990-1001.

35. Shin PJ, Larson PEZ, Ohliger MA, et al. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magn Reson Med*. 2014;72:959-970.

36. Haldar JP. Low-rank modeling of local-space neighborhoods (LORAKS) for constrained MRI. *IEEE Trans Med Imaging*. 2014;33:668-681.

37. Haldar JP, Zhuo J. P-LORAKS: low-rank modeling of local k-space neighborhoods with parallel imaging data. *Magn Reson Med*. 2015;75:1499-1514.

38. Ongie G, Jacob M. Off-the-grid recovery of piecewise constant images from few Fourier samples. *SIAM J Imaging Sci*. 2016;9:1004-1041.

39. Jin KH, Lee D, Ye JC. A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix. *IEEE Trans Comput Imaging*. 2016;2:480-495.

40. Haldar JP, Setsompop K. Linear predictability in magnetic resonance imaging reconstruction: leveraging shift-invariant Fourier structure for faster and better imaging. *IEEE Signal Process Mag*. 2020;37:69-82.

41. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*. 1999;42:952-962.

42. 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.

43. Hoge WS, Setsompop K, Polimeni JR. Dual-polarity slice-GRAPPA for concurrent ghost correction and slice separation in simultaneous multi-slice EPI. *Magn Reson Med*. 2018;80:1364-1375.

44. Liao C, Manhard MK, Bilgic B, et al. Phase-matched virtual coil reconstruction for highly accelerated diffusion echo-planar imaging. *NeuroImage*. 2019;194:291-302.

45. Kim TH, Haldar JP. *LORAKS Software Version 2.0: Faster Implementation and Enhanced Capabilities*. Technical Report USC-SIPI-443. Los Angeles, CA: University of Southern California; 2018.

46. Ongie G, Jacob M. A fast algorithm for convolutional structured low-rank matrix recovery. *IEEE Trans Comput Imaging*. 2017;3:535-550.

47. Reeder SB, Atalar E, Faranesh AZ, McVeigh ER. Referenceless interleaved echo-planar imaging. *Magn Reson Med*. 1999;41:87-94.

48. Grieve SM, Blamire AM, Styles P. Elimination of Nyquist ghosting caused by read-out to phase-encode gradient cross-terms in EPI. *Magn Reson Med*. 2002;47:337-343.

49. Chen NK, Avram AV, Song AW. Two-dimensional phase cycled reconstruction for inherent correction of echo-planar imaging Nyquist artifacts. *Magn Reson Med*. 2011;66:1057-1066.

50. Le Bihan D, Poupon C, Amadon A, Lethimonnier F. Artifacts and pitfalls in diffusion MRI. *J Magn Reson Imaging*. 2006;24:478-488.

51. Javed A, Nayak KS. Single-shot EPI for ASL-CMR. *Magn Reson Med*. 2020;84:738-750.

52. Bhushan C, Haldar JP, Choi S, Joshi AA, Shattuck DW, Leahy RM. Co-registration and distortion correction of diffusion and anatomical images based on inverse contrast normalization. *Neuroimage*. 2015;115:269-280.

53. Buehrer M, Boesiger P, Kozerke S. Virtual body coil calibration for phased-array imaging. *Proc Int Soc Magn Reson Med*. 2009;760.

54. Kim TH, Haldar JP. The Fourier radial error spectrum plot: a more nuanced quantitative evaluation of image reconstruction quality. *Proc IEEE Int Symp Biomed Imaging*. 2018;61-64.

55. Mani M, Jacob M, McKinnon G, et al. SMS-MUSSELS: a navigator-free reconstruction for slice-accelerated multi-shot diffusion imaging. *Proc Int Soc Magn Reson Med*. 2019;233.

56. Bilgic B, Liao C, Manhard MK, et al. Robust high-quality multi-shot EPI with low-rank prior and machine learning. *Proc Int Soc Magn Reson Med*. 2019;1250.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the Supporting Information section.

FIGURE S1 Illustration of EPI ghost correction. The top row of this figure shows EPI images obtained from different methods, while the bottom row shows the same images with 10× intensity amplification to highlight the ghost characteristics. If EPI data are naively reconstructed without accounting for the systematic differences between data acquired with positive and negative readout gradient polarities ("Uncorrected"), then strong Nyquist ghosts appear in the image as indicated with arrows. Modern EPI techniques frequently try to eliminate these artifacts using navigator information to estimate the systematic differences between the data collected with different readout polarities. In the navigator-based example, we show ("Navigator") that the navigator information was collected using a three-line EPI acquisition with the phase encoding gradients turned off, and the difference between positive and negative gradient polarities was modeled using constant and 1D linear phase terms. Although this approach substantially reduces Nyquist ghosts, it is common for some amount of residual ghosting to still be present in the images, particularly in cases where simple 1D phase modeling is inadequate to capture the differences between the two gradient polarities. We also show an example of our proposed approach ("RAC-LORAKS"), which can account for more complicated variations between the different gradient polarities, and which is substantially more successful at suppressing Nyquist ghosts in this example

FIGURE S2 Illustration of the orientation of the double-oblique gradient-echo EPI dataset. The double-oblique slices are shown in red, overlaid on a structural T1-weighted image of the same subject. The double-oblique slice used for the results in Figure 2 is shown with a yellow rectangle

FIGURE S3 Illustration of the EPI and ACS datasets used in simulation. The first and second top rows show coil-combined multi-channel data for the case when the EPI and ACS data have similar and inverted contrast, respectively, while the bottom row shows representative single-channel images. We also show the interpolarity phase difference for the coil-combined EPI data, as well as the difference in the interpolarity phase difference between the coil-combined EPI and ACS data

FIGURE S4 The same results shown in Figure 1, but with a 5× intensity amplification to highlight the ghost characteristics

FIGURE S5 DPG results corresponding to the same data shown in Figure 1 and Supporting Information Figure S4. The same mDPG results shown in Figure 1 and Supporting Information Figure S4 are also reproduced in this figure for reference. Note that the processing steps of DPG cause the image intensities to be mismatched from the intensities of mDPG and the other reconstruction methods, which precludes a quantitative comparison

FIGURE S6 mDPG and DPG results corresponding to the same multi-channel simulated data from Figure 3

FIGURE S7 Reconstruction results for multi-channel simulated data with different parallel imaging acceleration factors. These simulations are identical to those reported in Figure 3, except that the images used to generate EPI data and the images used to generate ACS data were interchanged

FIGURE S8 Reconstruction results for the second set of multi-channel simulations (with inverted contrast between ACS and EPI data) with different parallel imaging acceleration factors

FIGURE S9 The same results shown in Figure 7, but zoomed-in to a region of interest for improved visualization

TABLE S1 NRMSEs for the multi-channel inverted contrast simulation results shown in Supporting Information Figure S8. For each acceleration factor, the smallest values are highlighted in bold

TABLE S2 NRMSEs for the single-channel simulation results shown in Figure 5. For each acceleration factor, the smallest values are highlighted in bold