

A BAYESIAN APPROACH TO GRAPPA PARALLEL FMRI IMAGE RECONSTRUCTION INCREASES SNR AND POWER OF TASK DETECTION

BY CHASE J. SAKITIS^a AND DANIEL B. ROWE^b

Mathematical and Statistical Sciences, Marquette University, ^achase.sakitis@marquette.edu, ^bdaniel.rowe@marquette.edu

In fMRI, capturing brain activation during a task is dependent on how quickly k -space arrays are obtained. Acquiring full k -space arrays, which are reconstructed into images using the inverse Fourier transform (IFT), that make up volume images can take a considerable amount of scan time. Undersampling k -space reduces the acquisition time but results in aliased, or “folded,” images. GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) is a parallel imaging technique that yields full images from subsampled arrays of k -space. GRAPPA uses localized interpolation weights, which are estimated prescan and fixed over time, to fill in the missing spatial frequencies of the subsampled k -space. Here we propose a Bayesian approach to GRAPPA (BGRAPPA) where prior distributions for the unacquired spatial frequencies, localized interpolation weights, and k -space measurement uncertainty are assessed from the a priori calibration k -space arrays. The prior information is utilized to estimate the missing spatial frequency values from the posterior distribution and reconstruct into full field-of-view images. Our BGRAPPA technique successfully reconstructed both a simulated and experimental time series resulting in reduced noise leading to an increased signal-to-noise ratio (SNR) and stronger power of task detection.

1. Introduction.

1.1. Background. Functional magnetic resonance imaging (fMRI) is a type of medical imaging developed in the early 1990s as a technique to noninvasively observe human brain activity without exogenous contrast agents (Bandettini et al. (1993)). This procedure examines the brain in action by detecting changes in the brain using the blood-oxygen-level dependent (BOLD) contrast (Ogawa et al. (1990)). The increase in the BOLD contrast in the area of a neuron is a correlate for neuronal firing. Measurements from the machine are arrays of complex-valued spatial frequencies called k -space (Kumar, Welti and Ernst (1975)). These complex-valued k -space arrays are then reconstructed into images using a 2D discrete inverse Fourier transform (IFT) producing complex-valued brain images. The magnitude and phase of the complex-valued reconstructed images can be utilized for activation analysis (Rowe and Logan (2004), Rowe (2005)), but generally, only the magnitude is used (Bandettini et al. (1993)).

In fMRI, measuring full arrays of data for all the slices that form each volume image typically takes about one to two seconds, limiting the temporal resolution of the obtained images and potentially diminishing brain activity detection. Shortening the time it takes to acquire the data required for volume images would improve capturing brain activity. A great deal of work has been dedicated to reducing the acquisition time of the MRI process by accelerating the number of images obtained per unit of time. Hyde et al. (1986), Pruessmann et al. (1999), and Griswold et al. (2002) explored parallel imaging techniques, while Li (2008) subsampled three dimensional k -space data and filtered to expand into the full volume k -space. The purpose of each of these techniques was to reduce the acquisition time in MRI.




FIG. 1. Full k -space array (top left), sequence of the subsampled k -space array with $n_A = 2$ (top right), the acquired subsampled k -space array (bottom right), and aliased brain image after IFT of the subsampled k -space array (bottom left).

1.2. *Previous approach.* Historically, a single channel coil receiver has been utilized in fMRI to measure full-sampled k -space data arrays. Reducing time between successive volume images is the primary goal of parallel imaging, which can also reduce total scan time. More recently, the technological development focus has been to reduce acquisition time by measuring less data without losing the ability to form a full image. This can be achieved by skipping the acquisition of lines in the k -space array, that is, subsampling. To accomplish this, multiple receiver coils are utilized in parallel to obtain spatial frequency arrays, which are reconstructed into coil-specific brain images.

Skipping lines in k -space introduces what is called an acceleration factor. The acceleration factor indicates which lines of k -space data are measured. For example, with an acceleration factor of $n_A = 2$, every other line horizontally in k -space is measured. Figure 1 shows the sequential pattern for a fully sampled k -space array (top left) compared to a subsampled k -space array with an acceleration factor of $n_A = 2$ (top right). This acceleration factor will cause the reconstructed coil images to appear as if they are folded over, because the Fourier transform cannot uniquely map the downsampled signals. We can see an example of this in the bottom left of Figure 1 where the IFT of the subsampled k -space causes the brain image to be aliased.

To obtain a full field-of-view (FOV) image, the unacquired spatial frequencies need to be estimated to have full coil k -space arrays. The full k -space arrays (acquired plus estimation) for each coil are averaged to yield a single, full spatial frequency array. Then the averaged, full k -space array is inverse Fourier transformed into a full brain image. A common method that estimates the unacquired coil spatial frequencies is GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) and was introduced by Griswold et al. (2002). GRAPPA operates in the spatial frequency domain before the IFT, utilizing localized weights to interpolate the missing values in each coil k -space array. GRAPPA has its deficiencies, such as low image quality, a low SNR, and diminished task detection power with higher acceleration factors. Bayesian methodologies have been utilized in k -space to improve spatial resolution

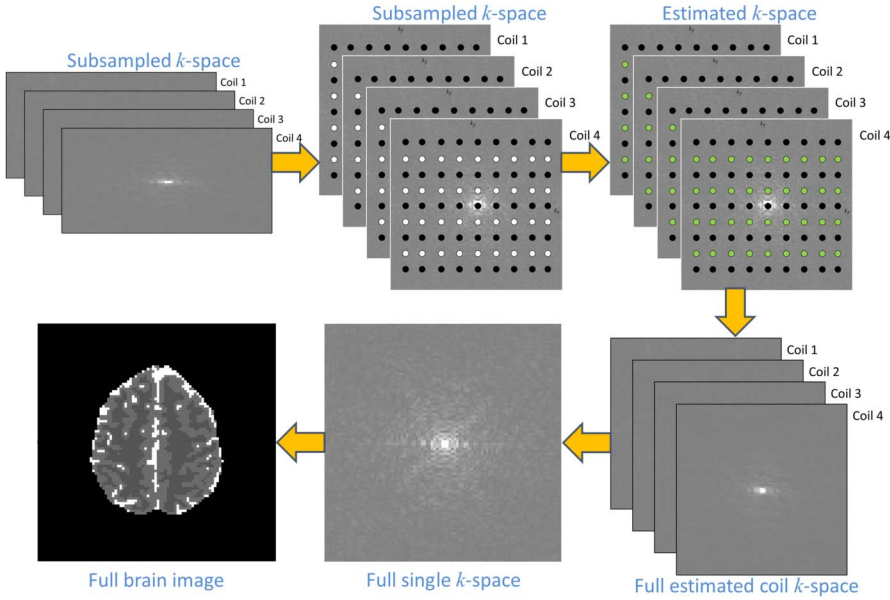


FIG. 2. Subsampled k -space coil arrays (top left) that are spread out to show a full k -space array where the black dots are the acquired spatial frequencies and the white dots are the unacquired spatial frequencies (top middle). The missing spatial frequencies are then estimated (green dots in the top right) yielding full coil k -space arrays (bottom right). The full coil k -space arrays are averaged together to produce a full spatial frequency array (bottom middle), which is then transformed into a full brain image (bottom left) using the IFT.

and image quality (Kornak et al. (2010)), but here we aim to use it for reconstructing subsampled k -space data to produce full brain images. Similar to BSENSE (Sakitis, Brown and Rowe (2025)), here we propose a Bayesian approach to GRAPPA that will incorporate prior information, yielding increased SNR and image quality with improved task detection power.

1.3. Overview. The second section of this paper will explain the model of GRAPPA image reconstruction and formulate the complex-valued problem as a real-valued isomorphic representation. This will lead into our proposed Bayesian approach presented in Section 3. Section 4 will show results from comparing traditional GRAPPA and our new BGRAPPA approach to simulated nontask and task fMRI data. Section 5 presents a similar comparison with experimental task fMRI data. We will conclude in Section 6 with an overview of the important results of the paper and a discussion of future work.

2. GRAPPA technique.

2.1. Reconstruction process. As mentioned in Section 1.1, to measure less k -space data and still produce a full brain image, $n_C > 1$ receiver coils must be utilized. The process for GRAPPA is exhibited in Figure 2 with an illustrative example of using $n_C = 4$ coils. The machine acquires subsampled spatial frequency arrays for each of the four coils shown in the top left of Figure 2. The top middle of Figure 2 displays the subsampled k -space arrays as full arrays with the black dots indicating the acquired spatial frequencies and the white dots indicating the unacquired spatial frequencies (the points skipped from the subsampling process). The unacquired spatial frequencies are estimated using GRAPPA image reconstruction, displayed as the green dots in the top right of Figure 2. The GRAPPA technique utilizes localized kernel weights to interpolate these unacquired points, which is further explained in Section 2.2. This yields full coil k -space arrays, as shown in the bottom right of Figure 2. To

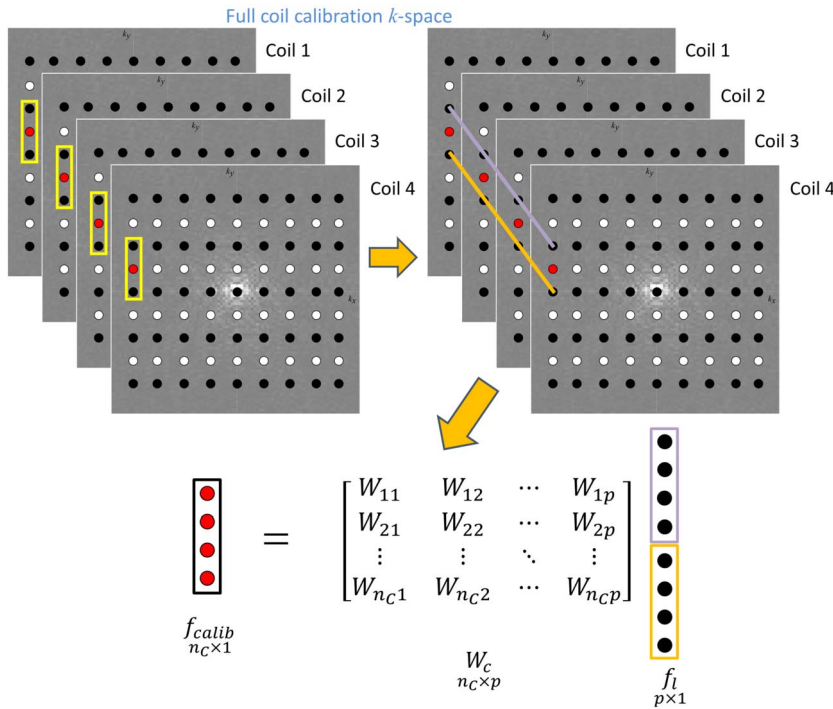


FIG. 3. The k -space coil arrays in the top left are fully sampled where the darker dots are treated as the acquired spatial frequencies and the calibration points (f_{calib}) are utilized to calculate the weights for those coil spatial frequencies. The boxes (in the k -space array in the top right) shows a 2×1 kernel indicating which points are utilized to estimate the weights. From this we get an acquired point above and below each coil calibration point. The points above the calibration points are then stacked by coil (1 through 4), which is then placed above the stacked dots below the calibration points (displayed as f_l), setting up the system of linear operations. The weights, W_c , for those “unacquired” coil spatial frequencies are then estimated using least squares. Once the weights have been estimated, the calibration points move to the next unacquired points to estimate the set of weights for the next unacquired spatial frequencies.

get a single full spatial frequency array (bottom middle), the full coil spatial frequency arrays are averaged together. The full spatial frequency is then IFT reconstructed into a single, full FOV brain image (bottom left of Figure 2).

2.2. Model. In GRAPPA the complex-valued localized interpolation weights are estimated using prescan coil calibration spatial frequency arrays. These coil calibration k -space arrays are fully sampled coil spatial frequency arrays that are collected prior to the actual fMRI experiment. Kernels of varying sizes can be used to estimate the weights, creating a system of linear equations. Figure 3 illustrates how a 2×1 kernel is utilized to estimate the weights from the full coil calibration spatial frequencies with a four-channel coil array. In Figure 3 all the complex-valued data points are acquired but are treated differently, depending on the location of the data point. The black data points, f_l , are utilized as the “acquired” complex-valued spatial frequency values, the red points, f_{calib} , are the complex-valued calibration spatial frequency points, and the white points are ignored for the calculation of those weights associated with the current f_{calib} points. The white dots represent the spatial frequencies that would be unacquired during the fMRI experiment but are used as calibrations points to estimate the complex-valued weights for those spatial frequencies.

The calibration points f_{calib} and the “acquired” spatial frequencies f_l along with the unacquired complex-valued weights, W_c , create a system of linear equations, as displayed in

Figure 3 (bottom). From the linear equations, we can estimate the weights w_c using equation (2.1),

$$(2.1) \quad W_c^{(v)} = f_{\text{calib}}^{(v)} f_l^{(v)H} (f_l^{(v)} f_l^{(v)H})^{-1}, \quad v = 1, \dots, K,$$

where $W_c \in \mathbb{C}^{n_c \times p}$ represents complex-valued weights, $f_{\text{calib}} \in \mathbb{C}^{n_c \times 1}$ represents complex-valued calibration spatial frequencies, $f_l \in \mathbb{C}^{p \times 1}$ represents “acquired” complex-valued spatial frequencies, $p = n_c k_{\text{rows}} k_{\text{cols}}$, k_{rows} represents number of rows in the kernel, k_{cols} represents number of columns in the kernel, H represents Hermitian or conjugate transpose, and K represents total number of unacquired spatial frequencies in the subsampled k -space array. The process is repeated for each spatial frequency point that would be unacquired during the actual fMRI experiment (the white dots in Figure 3), yielding different weights for each unacquired spatial frequency.

Once the weights for each of the unacquired coil spatial frequencies are estimated from the calibration k -space arrays, those weights are then utilized to interpolate the unacquired spatial frequencies in the actual fMRI experiment. The GRAPPA model with the estimated weights becomes

$$(2.2) \quad f_{ec}^{(v)} = W_c^{(v)} f_{kc}^{(v)} + \eta_c^{(v)}, \quad v = 1, \dots, K,$$

where $f_{ec} \in \mathbb{C}^{n_c \times 1}$ is the complex-valued interpolated k -space values, $f_{kc} \in \mathbb{C}^{p \times 1}$ is the complex-valued acquired k -space values, and $\eta_c \in \mathbb{C}^{n_c \times 1}$ is the additive complex-valued noise with $\eta_c \sim N(0, \tau^2(1 + i))$. The interpolated coil k -space values, f_{ec} , are inserted in the respective locations of each coil yielding full coil k -space arrays (top right of Figure 2).

With GRAPPA image reconstruction, however, the resulting reconstructed brain images can have diminished SNR, which is a consequence of either a decreased signal intensity, increased temporal noise variance, or a combination of the two. With an increase in the temporal noise variance, this can lead to reduced power in task detection as well. These deficiencies motivate our Bayesian approach, which will allow for a more automated method for image reconstruction without having to potentially store and use large matrices. Unlike GRAPPA, our Bayesian approach will utilize all available prior information from the calibration spatial frequency arrays and provide full distributions for the unacquired spatial frequencies, the weights, and the residual k -space variance.

3. Bayesian approach to GRAPPA. For our proposed Bayesian approach, we use the same linear model as GRAPPA as expressed equation (2.2), except the acquired spatial frequencies will be the f_{ec} variable instead of the f_{kc} variable. This creates a model where the design matrix and the coefficients can both be treated as unknown parameters, allowing us to take a Bayesian approach to the linear regression. Then the weights, W_c , and the unacquired spatial frequencies, f_{kc} , along with the residual k -space variance, τ^2 , are treated as unknowns with prior distributions placed on them. We also use an isomorphic real-valued representation of the linear GRAPPA model in equation (2.2) and is given by

$$(3.1) \quad \begin{bmatrix} f_{eR} \\ f_{eI} \end{bmatrix} = \begin{bmatrix} W_R & -W_I \\ W_I & W_R \end{bmatrix} \begin{bmatrix} f_{kR} \\ f_{kI} \end{bmatrix} + \begin{bmatrix} \eta_R \\ \eta_I \end{bmatrix},$$

where $f_{eR} \in \mathbb{R}^{n_c \times 1}$ and $f_{eI} \in \mathbb{R}^{n_c \times 1}$ are the real and imaginary components, respectively, of f_{ec} , $W_R \in \mathbb{R}^{n_c \times p}$ and $W_I \in \mathbb{R}^{n_c \times p}$ are the real and imaginary components of W_c , $f_{kR} \in \mathbb{R}^{p \times 1}$ and $f_{kI} \in \mathbb{R}^{p \times 1}$ are the real and imaginary components of f_{kc} , and $\eta_R \in \mathbb{R}^{n_c \times 1}$ and $\eta_I \in \mathbb{R}^{n_c \times 1}$ are the real and imaginary components of η_c with $(\eta_R, \eta_I)' \sim N(0, \tau^2 I_{2n_c})$. This equation is a latent variable model with complex values and can be more compactly written as $f_e = W f_k + \eta$ where $f_e \in \mathbb{R}^{2n_c \times 1}$, $W \in \mathbb{R}^{2n_c \times 2p}$, $f_k \in \mathbb{R}^{2p \times 1}$, and $\eta \in \mathbb{R}^{2n_c \times 1}$ are the real-valued isomorphic representations of f_{ec} , W_c , f_{kc} , and η_c , respectively.

In this method, two different representations of the weights will be used. The first representation is the proper skew-symmetric design matrix W , as shown in equation (3.1). The second representation is $D = [W_R, W_I]$, which is used in the prior distribution and for parameter estimation of the weights. This is to ensure W_R and W_I are uniquely estimated for W and do not need to be duplicated.

3.1. Data likelihood, prior, and posterior distributions. Like GRAPPA, we assume that the residual spatial frequency error is normally distributed in the real and imaginary components, since the real and imaginary components of fMRI data are commonly assumed to be normally distributed (Lindquist (2008)). The data likelihood for the acquired spatial frequencies for the n_c coils is

$$(3.2) \quad P(f_e|W, f_k, \tau^2) \propto (\tau^2)^{-\frac{2n_c}{2}} \exp\left[-\frac{1}{2\tau^2}(f_e - Wf_k)'(f_e - Wf_k)\right].$$

We can quantify available prior information about the unacquired spatial frequencies f_k , the weights W , and the residual k -space variance τ^2 with assessed hyperparameters of prior distributions. The unacquired spatial frequencies f_k are specified to have a normal prior distribution, expressed in equation (3.3). The weights D are also specified to have a normal prior distribution (equation (3.4)), and the k -space noise variance τ^2 is specified to have an inverse gamma prior distribution (equation (3.5)),

$$(3.3) \quad P(f_k|n_k, f_{k0}, \tau^2) \propto (\tau^2)^{-\frac{2p}{2}} \exp\left[-\frac{n_k}{2\tau^2}(f_k - f_{k0})'(f_k - f_{k0})\right],$$

$$(3.4) \quad P(D|n_w, D_0, \sigma^2) \propto (\tau^2)^{-\frac{2n_{CP}}{2}} \exp\left[-\frac{n_w}{2\tau^2} \text{tr}(D - D_0)(D - D_0)'\right],$$

$$(3.5) \quad P(\tau^2|\alpha_k, \delta) \propto (\tau^2)^{-(\alpha_k+1)} \exp\left[-\frac{\delta}{\tau^2}\right],$$

where tr is the trace of the $(D - D_0)(D - D_0)'$ matrix and the hyperparameters n_k , f_{k0} , n_w , D_0 , α_k , and δ are assessed from the prescan calibration spatial frequencies, as outlined in Section 3.2. The joint posterior distribution of the unacquired spatial frequencies f_k , the weights W , and the residual k -space variance τ^2 is

$$(3.6) \quad P(f_k, D, \tau^2|f_e) \propto P(f_e|W, f_k, \tau^2)P(f_k|n_k, f_{k0}, \tau^2)P(D|n_w, D_0, \tau^2)P(\tau^2|\alpha_k, \delta),$$

with the distributions specified from equations (3.2), (3.3), (3.4), and (3.5).

A technique that can be utilized for parameter estimation is using Markov chain Monte Carlo (MCMC) Gibbs sampling. The Gibbs sampler uses the posterior conditionals to generate the entire distribution for each parameter at each time point yielding more information that can be used for statistical analysis. However, the computation time is longer compared to using an iterative maximum a posteriori (MAP) method. For this paper we only use the MAP estimate since we would only be interested in the mean of the distributions for each parameter, which is theoretically equal to the mode for the unacquired spatial frequencies f_{k0} and the weights W .

3.2. Hyperparameter determination. The hyperparameters can be appropriately assessed in an automated way using the full prescan coil calibration spatial frequencies. For the BGRAPPA hyperparameter assessment, the same full calibration spatial frequencies and $f_{\text{calib}} = Wf_l$ model are used like in GRAPPA reconstruction, but each spatial frequency point is treated differently than GRAPPA. As shown in Figure 4, the calibration spatial frequencies f_{calib} for BGRAPPA are in the location of the data points where the acquired spatial frequencies are in the actual fMRI experiment. For GRAPPA these data points are assigned to the f_l




FIG. 4. Full calibration k -space arrays that indicate which data points are used as f_{calib} points and the f_l points for GRAPPA (left) and BGRAPPA (right).

variable in the $f_{\text{calib}} = W_c f_l$ model shown at the bottom of Figure 4. Using equation (2.1), this will result in the prior for the weights in BGRAPPA, D_0 , to be different than the estimated weights utilized in GRAPPA image reconstruction. The f_l points used for estimating the prior mean for the weights are averaged to obtain the prior mean of the unacquired spatial frequencies, f_{k0} .

The hyperparameters n_k and n_w , which are the prior scalars of the prior means, are assessed to be the number of calibration time points n_{cal} . The average residual k -space variance over the coil spatial frequency arrays is calculated to obtain a prior mean for the residual k -space variance τ_0^2 . The hyperparameters α_k (shape parameter of the inverse gamma) and δ (scale parameter of the inverse gamma) are assessed to be $\alpha_k = n_{\text{cal}} - 1$ and $\delta = (n_{\text{cal}} - 1)\tau_0^2$. This prior information is incorporated to estimate the unacquired spatial frequencies in the subsampled k -space arrays.

3.3. Posterior estimation. Using the posterior distribution in equation (3.6), the full posterior conditional distributions for the the unacquired spatial frequencies f_k , the weights D , and the residual k -space variance τ^2 are given by

$$(3.7) \quad f_k | W, \tau^2, f_e \sim N\{\hat{f}_k, \tau^2(W'W + n_k I_{2p})^{-1}\},$$

$$(3.8) \quad D | f_k, \tau^2, f_e \sim N\{\hat{D}, \tau^2(F_k' F_k + n_w I_{2p})^{-1}\},$$

$$(3.9) \quad \tau^2 | f_k, W, f_e \sim IG\{\alpha_{k*}, \delta_*\},$$

where $\alpha_{k*} = n_C p + n_C + p + \alpha_k$ and $\delta_* = [(f_e - W f_k)'(f_e - W f_k) + n_k(f_k - f_{k0})'(f_k - f_{k0}) + n_w \text{tr}((D - D_0)(D - D_0)') + 2\delta]/2$. MAP estimates for f_k , D , and τ^2 are estimated via the iterated conditional modes (ICM) optimization algorithm (Lindley and Smith (1972), O'Hagan (1994)). Beginning with the prior means for each parameter as initial estimates, the ICM algorithm iterates over the parameters, calculating its posterior conditional mode until convergence at the joint posterior mode. The posterior conditional modes are

$$(3.10) \quad \hat{f}_k = (W'W + n_k I_{2p})^{-1}(W' f_e + n_k f_{k0}),$$

$$(3.11) \quad \hat{D} = (F_e F_k' + n_w D_0)(F_k F_k' + n_w I_{2p})^{-1},$$

$$(3.12) \quad \hat{\tau}^2 = \frac{\Theta}{2(n_C + 2p + 2n_C p + 1)},$$

where $\Theta = (f_e - W f_k)'(f_e - W f_k) + n_k(f_k - f_{k0})'(f_k - f_{k0}) + \alpha_k \delta + n_w \text{tr}[(D - D_0)(D - D_0)']$, $F_e = [f_{eR}, f_{eI}]$, and $F_k \in \mathbb{R}^{2p \times 2}$ is a skew symmetric matrix representation of the

unaliaised voxel values f_k , as expressed by

$$(3.13) \quad F_k = \begin{bmatrix} f_{kR} & f_{kI} \\ -f_{kI} & f_{kR} \end{bmatrix}.$$

4. Simulation study.

4.1. Nontask spatial frequency data. A noiseless nontask image was used to create two separate series of $n_{\text{TR}} = 510$ simulated full coil images for one slice to mimic real-world MRI experimental data. The last n_{cal} time points of the first time series served as the calibration information utilized for hyperparameter assessment. The second time series was used for simulating a subsampled nontask experiment. The complex-valued nontask image was multiplied by a complex-valued designed sensitivity map with $n_C = 8$ coils. Figure 5 illustrates the real and imaginary parts of the full simulated brain image (first and second columns) being voxelwise multiplied by the real and imaginary components of the sensitivities for each of the $n_C = 8$ coils (third and fourth columns). This results in the real and imaginary components of the complex-valued full coil-weighted images (fifth and sixth columns).

In real-world MRI experiments, the first few images in an fMRI time series have increased signal as the magnetization reaches a stable state (Steinhoff et al. (2001)). To mimic this, the first three of both nontask time series of $n_{\text{TR}} = 510$ time points of the simulated nontask time series were scaled with the signal slightly decreasing from the first to the third time point before reaching a stable signal in the fourth time point. The scaling was determined by dividing the first three images of the experimental data by the twenty-first image, separately. After dividing the three images, the signal increase for each tissue type (white matter, grey matter, and CSF) was averaged together for each of the three divided images, calculating the average signal increase for each matter type. For example, the average signal increase in the first image for the white matter was 40%, 55% for the grey matter, and 75% for the CSF giving multiplication factors of 1.40, 1.55, and 1.75 for the matter types, respectively. This process was repeated for the second and third image in the series with the multiplication factors decreasing from the first to the third image.

The series of images for both the n_{cal} calibration images and the full simulated images were then Fourier transformed into noiseless full coil k -space arrays. The time series of coil k -space arrays were simulated by adding separate $N(0, 0.0036n_y n_x)$ noise, where n_y and n_x are the number of rows and columns, respectively, in the full k -space array, to the real and imaginary parts of full coil k -space arrays, corresponding to the noise in the real-world fMRI experimental data. To mimic the fMRI experiment, the first 20 time points of the second time series were omitted, leaving 490 time points of spatial frequency arrays for the single slice. However, the first 10 time points of an fMRI experiment can be used to estimate a T_1 map which efficiently segments the different tissue types. The next 10 time points can be utilized to estimate a static magnetic field map to adjust for geometric distortions (Karaman, Bruce and Rowe (2015)). The remaining 490 time points in the time series were subsampled by censoring lines in k -space according to an acceleration factor of $n_A = 3$. An example of the real and imaginary components of subsampled k -space arrays for $n_C = 8$ coils and an acceleration factor of $n_A = 3$ at one time point is exhibited in Figure 6.

4.2. Nontask reconstruction results. To analyze the reconstruction performance of BGRAPPA vs. GRAPPA, we first reconstructed subsampled k -space arrays at one time point, yielding a single unalised image for both methods. For calibration analysis the last $n_{\text{cal}} = 30$ time points from the first nontask time series were utilized for hyperparameter assessment. The first time point of the 490 subsampled, simulated nontask time series with an acceleration factor of $n_A = 3$, shown in Figure 6. The results of reconstructing the first time point in




FIG. 5. *Real and imaginary components of the true complex-valued simulated image (first and second columns) voxelwise multiplied by the real and imaginary parts of the complex-valued coil sensitivities for each of the $n_C = 8$ coils (third and fourth columns) yielding the real and imaginary components of the complex-valued coil-weighted images (fifth and sixth columns).*

the subsampled time series using both BGRAPPA and GRAPPA are shown in Figures 7, 8, 9, and 10.

The prior means from the calibration information for the unacquired spatial frequency arrays f_{k0} and the localized weights D_0 were used as initial values for f_k and D . These initial values were used to generate a τ^2 value from the posterior conditional mode from equation (3.12), initializing the ICM optimization algorithm. The simulated subsampled coil k -space arrays were reconstructed into a single, full brain image using the BGRAPPA MAP estimate from the ICM algorithm, and traditional GRAPPA estimate. For the ICM algorithm, only three iterations were needed for estimating the parameters with a computation time of ap-




FIG. 6. *Simulated acquired noisy subsampled coil spatial frequency arrays for the first time point in the nontask time series with an acceleration factor of $n_A = 3$.*

proximately half a second per time point. Figure 7 shows the reference magnitude and phase images (first column) that the reconstructed images will be compared to. These reference images were determined by simply averaging the full, noisy coil k -space arrays in the time series yielding a single spatial frequency array and then applying the IFT resulting in a full brain image. This provides us what image reconstruction would look like without applying an acceleration factor. This process is used for all of the reference images provided in the figures displaying simulated results for the remainder of the paper. Figure 7 also shows the BGRAPPA MAP unaliased image (second column) and the GRAPPA unaliased image (third

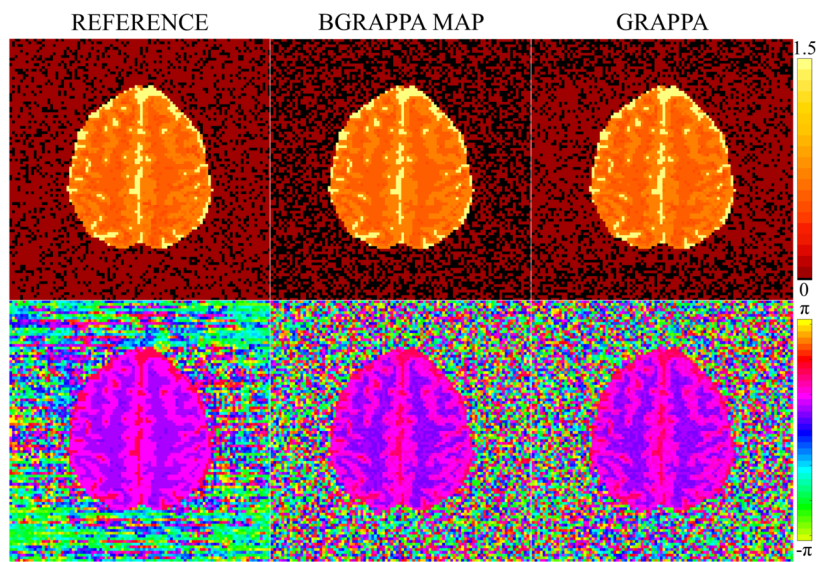


FIG. 7. *True nontask unaliased images (first column), reference nontask reconstructed images (second column), and BGRAPPA MAP unaliased nontask images (third column) using ICM and GRAPPA nontask images (fourth column) with magnitude images in the first row and phase images in the second row. Due to the circular nature of phase angles, the color bar for the phase images have wraparound.*




FIG. 8. Reconstructed magnitude images for different number of calibration images using BGRAPPA MAP estimate (top row of the right three columns) and GRAPPA (second row of the right three columns) with the true simulated magnitude image (top left), and the reference magnitude image (bottom left).

column) for the first time point in the simulated nontask series. We can see that the joint MAP estimate from BGRAPPA and the GRAPPA estimate both produce magnitude and phase images that closely resemble the true nonaliased, reference images in Figure 7. Visually the BGRAPPA image is slightly more accurate and less noisy than the GRAPPA image, which is further analyzed in Section 1 of the Supplementary Material (Sakitis and Rowe (2025)).

To quantify the differences between the true and reconstructed magnitude and phase images, we use the mean squared error, $MSE = \frac{1}{K} \sum_{j=1}^K (v_j - \bar{v}_j)^2$, where K is the number of voxels (either inside or outside the brain) in the full reconstructed image, v_j is the reconstructed magnitude or phase value of the j th voxel, and \bar{v}_j is the true magnitude or phase value of the j th voxel. This measure will indicate the accuracy of a single reconstructed image, compared to the true simulated image, with lower MSE indicating a more accurate reconstructed image. The MSE for BGRAPPA for inside and outside the brain was lower for the magnitude and phase images compared to GRAPPA. The MSE for the magnitude reconstructed image of GRAPPA was 114% and 51% higher for inside and outside the brain, respectively, compared to BGRAPPA. For the MSE of the phase reconstructed images, GRAPPA was 12% and 3% higher for inside and outside the brain compared to BGRAPPA.

Next, we evaluated how the number of calibration time points, n_{cal} , affected the reconstructed images. For the prescan calibration analysis, we fixed the acceleration factor to be $n_A = 3$ for the subsampled k -space coil arrays of the simulated nontask time series with $n_{IMG} = 490$ time points. Then we set the number of calibration time points to be $n_{cal} = 5, 10, 15, 20, 25, 30$ for separate hyperparameter assessments. After assessing the hyperparameters using each number of calibration time points, the simulated nontask time series with the subsampled coil spatial frequency arrays were reconstructed using BGRAPPA MAP and GRAPPA.

The results, displayed in Figure 8, indicate that increasing the number of calibration time points does not noticeably affect the noise level inside or outside the brain for either BGRAPPA or GRAPPA. This means we can have short calibration scans and do not need to take up valuable scanner time. For each of the number of calibration time points, GRAPPA, visually, is slightly noisier than BGRAPPA (shown in Section 1 of the Supplementary Material (Sakitis and Rowe (2025))). To further analyze the differences between the BGRAPPA and GRAPPA reconstructed magnitude images, the MSE and entropy for BGRAPPA and

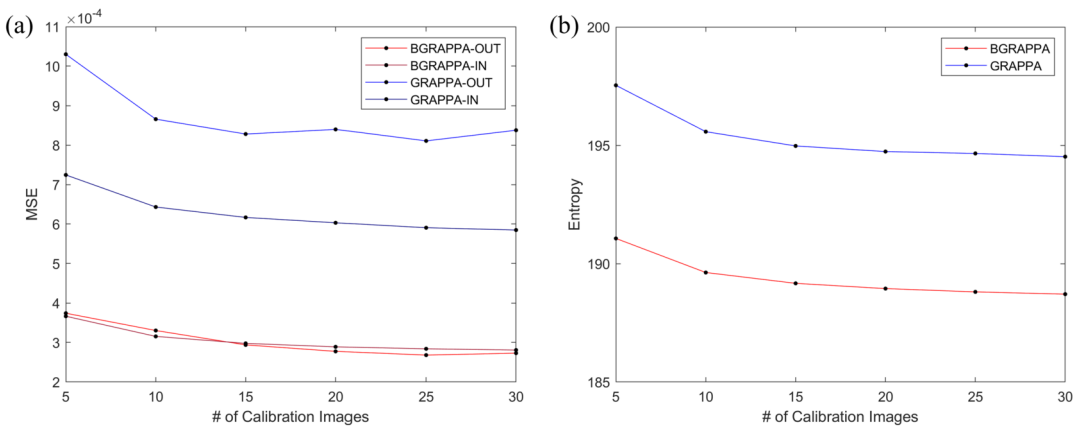


FIG. 9. (a) MSE for inside and outside the brain for BGRAPPA and GRAPPA comparing both method’s reconstructed images to the true simulated magnitude image for each number of calibration images. (b) Entropy plot for BGRAPPA and GRAPPA for each number of calibration images.

GRAPPA for each number of calibration time points were calculated to quantify this result. Entropy analyzes uncertainty and smoothness across a single image with lower entropy meaning less uncertainty throughout the image. The equation for entropy is given by $E = -\sum_{j=1}^N [\frac{v_j}{v_{\max}} \ln(\frac{v_j}{v_{\max}})]$, where \ln is the natural log, N is the number of voxels in the full reconstructed image, v_j is the reconstructed magnitude value of the j th voxel, and v_{\max} is the voxel intensity if all the image intensities were in one pixel (Atkinson et al. (1997)), given by $v_{\max} = \sqrt{\sum_{j=1}^N v_j^2}$.

Shown in Figure 9a, the MSE for inside and outside the brain for the BGRAPPA MAP reconstructed magnitude images was markedly smaller than the GRAPPA reconstructed magnitude images for both inside and outside the brain. BGRAPPA also had noticeably smaller entropy values, compared to the GRAPPA reconstructed magnitude images, displayed in Figure 9b. Lower MSE for BGRAPPA indicates a more precise reconstructed image while smaller entropy means less uncertainty with image reconstruction. For both BGRAPPA and




FIG. 10. Reconstructed magnitude images for different acceleration factors using BGRAPPA MAP estimate (top row of the right three columns) and GRAPPA (second row of the right three columns) with the true simulated magnitude image (top left), and the reference magnitude image (bottom left).

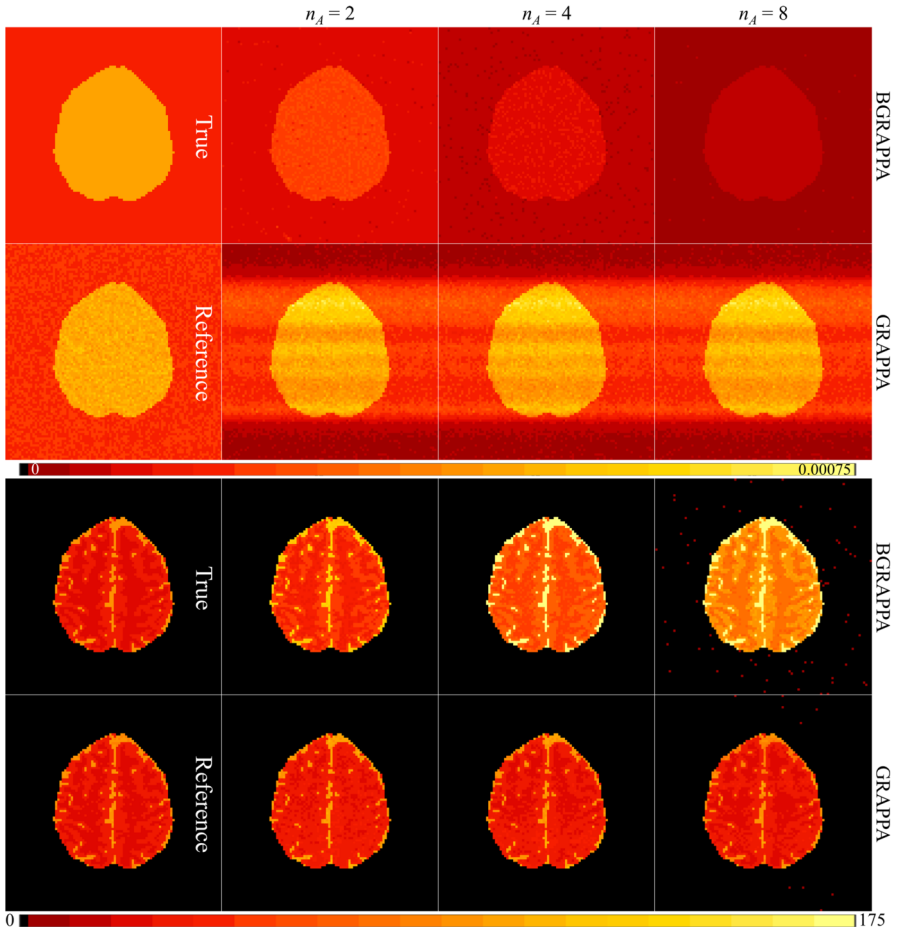


FIG. 11. Temporal variance and SNR images for different acceleration factors using BGRAPPA MAP estimate (first row and third row, respectively, of the right three columns) and GRAPPA (second row and fourth row, respectively, of the right three columns). The first column shows the true variance and SNR (first and third rows, respectively) and the reference variance and SNR (second and fourth rows, respectively).

GRAPPA, increasing the number of calibration images does not meaningfully affect the temporal variance, resulting in similar SNR for each n_{cal} , which again means we can have a short calibration scan. In all cases the temporal variance for BGRAPPA is substantially lower than for GRAPPA, demonstrating that BGRAPPA mitigates noise in the reconstructed image. The phase of the reconstructed images for the different calibration time points can be found in Section 1 of the Supplementary Material (Sakitis and Rowe (2025)).

Along with analysis of the number of calibration time points, we evaluated how different acceleration factors, n_A , affected the reconstructed magnitude and phase images. Here we fixed the number of calibration time points to be $n_{\text{cal}} = 30$ for hyperparameter assessment and set the acceleration factors of the nontask time series to be $n_A = 2, 3, 4, 6, 8, 12$. We only show results for $n_A = 2, 4, 8$ just to see how increasing the acceleration factor affects the reconstruction results. These subsampled coil k -space arrays with separate acceleration factors were reconstructed into full images using the BGRAPPA MAP estimate and GRAPPA, again comparing the results for both methods.

The results, exhibited in Figure 10, showed that inside the brain of the reconstructed magnitude images from BGRAPPA and GRAPPA are negligibly affected by increasing the acceleration factor with BGRAPPA visually slightly more accurate. The noise level outside the brain for BGRAPPA does decrease as the acceleration factor increases while it only slightly

decreases for GRAPPA. Again, the GRAPPA magnitude reconstructed images have slightly more noise than the BGRAPPA magnitude reconstructed images (shown in Section 1 of the Supplementary Material (Sakitis and Rowe (2025))). The phase of the reconstructed images for the different acceleration factors can be found in Section 1 of the Supplementary Material (Sakitis and Rowe (2025)).

In Figure 11 we examine the temporal variance of the reconstructed time series and the SNR images for BGRAPPA, GRAPPA, the true images, and the reference images when no acceleration factor is applied. The reference reconstruction is simply averaging the full coil k -space arrays and applying the IFT to get full brain images for the full time series. The temporal variance for BGRAPPA decreases and for GRAPPA increases as the acceleration factor increases (first and second rows of Figure 11 of the right three columns). The temporal variance, overall, from Figure 11 for BGRAPPA is substantially lower than GRAPPA, the theoretically true variance (first row, first column), and the reference (second row, first column), showing that BGRAPPA reduces the noise through of the reconstructed time series. This also leads to higher SNR for BGRAPPA compared to GRAPPA (third and fourth row of Figure 11 of the right three columns), the true SNR (first column, third row), and the reference SNR (first column, fourth row). The average of the BGRAPPA reconstructed time series was also taken, and the result magnitude image looks similar to the true simulated magnitude image.

4.3. Task activation model. In task-based fMRI, the nontask reconstructed images create a baseline value for each voxel. This yields an intercept only simple linear regression model $y = \beta_0 + \varepsilon$ where y is the magnitude of the reconstructed voxel value. By adding in task activation to select images in the series of images, we have a simple linear regression model $y = \beta_0 + x\beta_1 + \varepsilon$ for the unaliased voxel values. In this regression, β_0 is the baseline voxel value from the nontask reconstructed images determining the $\text{SNR} = \beta_0/\sigma$, as demonstrated in the previous subsection. The β_1 value is the estimated task related signal increase from β_0 determining the contrast-to-noise ratio $\text{CNR} = \beta_1/\sigma$. The vector $x \in \{0, 1\}^{n_{\text{IMG}}}$, where n_{IMG} is the number of reconstructed images in the series, is a vector such that the zeros correspond to the images in the series without task activation and the ones correspond to the images with task activation. We can write this regression as $y = XB + \varepsilon$, where $X = [1, x] \in \mathbb{R}^{n_{\text{IMG}} \times 2}$ and $B = [\beta_0, \beta_1]'$.

The task is not usually visible on the reconstructed images since the CNR is typically much lower than the SNR. Instead, a right-tailed t -test is carried out with $\beta_1 \leq 0$ as the null hypothesis and $\beta_1 > 0$ as the alternative. The reason for the one-sided hypothesis test is because we only anticipate an increased signal from the task activation. To simulate added task, a $\beta_1 = 0.045$ magnitude-only signal increase is added to select voxels of the true noiseless nontask image. This added task activation is located in the left motor cortex to resemble the region of interest (ROI) of brain activity from the fMRI unilateral right-hand finger tapping experiment used in the Section 5 (Karaman, Bruce and Rowe (2014)).

Similar to magnitude-only task activation, we can also use the phase images for task detection. A simulated phase task of $\pi/120$ was also added to the simulated true simulated task image. A simple linear regression model, $\phi = \theta_0 + \theta_1 x + \epsilon$, can be used for the phase activation as well. In this regression, ϕ is the phase of the reconstructed voxel, θ_0 is the baseline phase voxel value from the nontask reconstructed images, and θ_1 is the estimated increase from θ_0 . We then use a one-tailed t -test, $t = \hat{\theta}_1/SE(\hat{\theta}_1)$, to determine which voxels contain statistically significant θ_1 values indicating which voxels experience phase task activation (Rowe, Meller and Hoffmann (2007)).

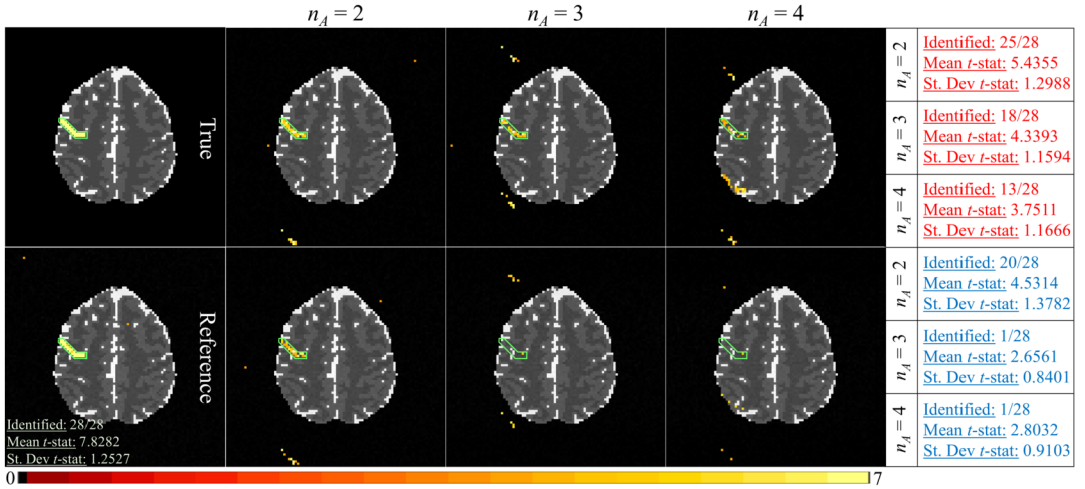


FIG. 12. Statistically significant voxels in the ROI using FDR for BGRAPPA reconstructed images (first row of the right three columns), significant voxels in the ROI using FDR for GRAPPA (second row of the right three columns), and analysis of the t -statistics in the boxes on the right. The true (first row) and the reference (second row) magnitude-only task activation is shown in the first column with the analysis of t -statistics of the reference reconstruction shown in the image.

4.4. fMRI spatial frequency data. A noiseless task image was used along with the noiseless nontask image to create a series of $n_{TR} = 510$ simulated full coil images for one slice mimicking real-world fMRI data. The simulated task activation was designed to mimic tapping of the subject's right fingers leading to activity in the left motor cortex which becomes our ROI for analyzing task detection in this experiment, as mentioned above. Knowing this, artificial signal increase was added to the voxels in the ROI (as mentioned in Section 4.3) for task images.

The true images were multiplied by the same $n_C = 8$ coil sensitivity maps used for the nontask simulated time series (as illustrated in Figure 5), and then the series of images were Fourier transformed in full coil k -space arrays. This series was also generated by adding separate $N(0, 0.0036n_y n_x)$ noise to the real and imaginary parts of the full coil k -space arrays and were then inverse Fourier transformed back into full coil images, yielding a CNR of 0.75. To simulate our real-world fMRI experimental process, of the first 20 time points in the series, 20 were nontask. The scaling for the first few images in the fMRI simulated data is similar to that outlined in Section 4.1 for each of the tissue types. The initial 20 nontask time points are followed by 16 epochs alternating between 15 nontask and 15 task time points. An epoch is a stimulation period with time points of the subject at rest (nontask) and the subject performing an action or task. The series culminated with 10 nontask time points producing the simulated fMRI series of $n_{TR} = 510$ images. To mimic the fMRI experiment in Section 5, the first 20 time points were omitted, leaving 490 time points in the series. This series is transformed into the spatial frequency domain and then subsampled according to the acceleration factor to simulate subsampling of a real fMRI experiment. The last n_{cal} full coil FOV time points in the second nontask time series from Section 4.1 were utilized as full FOV coil calibration k -space arrays to assess the hyperparameters. For this simulation, we evaluate both BGRAPPA and GRAPPA using $n_{cal} = 5, 10, 15, 20, 25, 30$ calibration time points. The results for the different calibration time points were similar to that of Section 4.2 where the different number of calibration time points had negligible effects on the results. Different acceleration factors of $n_A = 2, 3, 4$ were also tested in this simulated fMRI experiment and are shown in the next subsection.




FIG. 13. Statistically significant voxels in the ROI using FDR for BGRAPPA phase reconstructed images (first row of the right three columns), significant voxels in the ROI using FDR for GRAPPA (second row of the right three columns), and analysis of the t -statistics in the boxes on the right. The true (first row) and the reference (second row) phase-only task activation is shown in the first column with the analysis of t -statistics of the reference reconstruction shown below in the image.

4.5. *FMRI reconstruction results.* The hypothesis test described in Section 4.3 was utilized to determine voxels with a statistically significant magnitude-only signal increase. The statistically significant voxels for different acceleration factors were analyzed for the BGRAPPA MAP reconstructed time series and the GRAPPA reconstructed time series using the 5% false discovery rate (FDR) threshold procedure (Benjamini and Hochberg, 1995; Genovese, Lazar and Nichols, 2002; Logan and Rowe, 2004). The ROI here consists of 28 voxels located in the left motor cortex.

Figure 12 shows the statistically significant magnitude-only voxels from the BGRAPPA MAP reconstructed time series (first row of the right three columns) and the GRAPPA reconstructed time series (second row of the right three columns) for the different acceleration factors compared true and reference activations in the first column. Figure 12 also summarizes the t -statistics in the ROI for each acceleration factor. BGRAPPA identified more statistically significant voxels in the ROI for each acceleration factor. For acceleration factors of 3 and 4, the task activation is virtually undetected using the GRAPPA method. The mean value for the t -statistics was also substantially higher for BGRAPPA, compared to GRAPPA, demonstrating that BGRAPPA has a stronger task detection power. Increasing the acceleration factor decreases number of voxels identified and the mean of the t -statistics for both BGRAPPA and GRAPPA, but much more activation is captured from BGRAPPA than GRAPPA.

Using the 5% FDR threshold, Figure 13 shows phase activation for BGRAPPA and GRAPPA reconstructed time series using acceleration factors of 2, 3, and 4. Like the BGRAPPA reconstructed magnitude images, we can see that it captures the simulated task activation in the ROI for each acceleration factor. For GRAPPA the phase task activation is captured using an acceleration of 2, but noticeably diminishes when applying acceleration factors of 3 and 4. With higher mean t -statistic values for BGRAPPA, this indicates that BGRAPPA has more power in phase task detection. Phase activation is a topic of study as previously described in Section 4.2.




FIG. 14. BGRAPPA MAP unaliased nontask magnitude images for each acceleration factor (first row of the right three columns) using the ICM algorithm, and GRAPPA unaliased nontask magnitude images for each acceleration factor (second row of the right three columns) with the magnitude reference image (left column).

5. Experimental data.

5.1. Data description. A 3.0 T General Electric Signa LX magnetic resonance imager was used to conduct an fMRI experiment on a single subject. A right-handed finger-tapping task was performed in a block design with an initial 20 s rest followed by 16 epochs with 15 s off (rest state) and 15 s on (task performed). The experiment was concluded with 10 s of rest giving us a series of $n_{TR} = 510$ repetitions with each repetition being 1 s, a flip angle of 90° and an acquisition bandwidth of 125 kHz. The data set consists of nine 2.5 mm thick axial slices with $n_C = 8$ receiver coils that have a 96×96 dimension for a 24 cm full FOV, with a posterior to anterior phase encoding direction. For this paper the time series for all nine slices was used to analyze the effects of applying acceleration factors of $n_A = 2, 3, 4$ for both BGRAPPA and GRAPPA, but only the time series of the second slice is shown. Note that the simulation study in Section 4 directly mimics this experimental data.

Typically, the magnetic fields in an fMRI experiment will induce a drift in the phase over time, which we correct before reconstruction to give us a stable phase through time. Once the phase was corrected, the last $n_{cal} = 30$ full k -space arrays of a nontask series of $n_{TR} = 510$ time points performed on the subject were used for hyperparameter assessment. The fMRI experimental series described above was used for fMRI analysis. The first 20 images of each series were discarded due to varying echo times and magnetization stability. Like the simulation study, the subsampled coil k -space arrays came from artificially skipping lines in the full coil k -space arrays of the fMRI experimental time series, mimicking the effect of actually subsampling the coil k -space arrays. Before subsampling the time series, a reference image (left image in Figure 14) was produced by averaging the $n_C = 8$ full coil spatial frequency arrays at the first time point. This provides a magnitude and phase image with which to compare to GRAPPA and our proposed BGRAPPA.

5.2. Experimental results. Similar to the process for the simulated data described in Section 4, the unacquired spatial frequencies at each time point in the entire time series of subsampled coil k -space arrays were estimated using BGRAPPA and GRAPPA separately. The estimated full coil k -space arrays were then averaged together and transformed into image

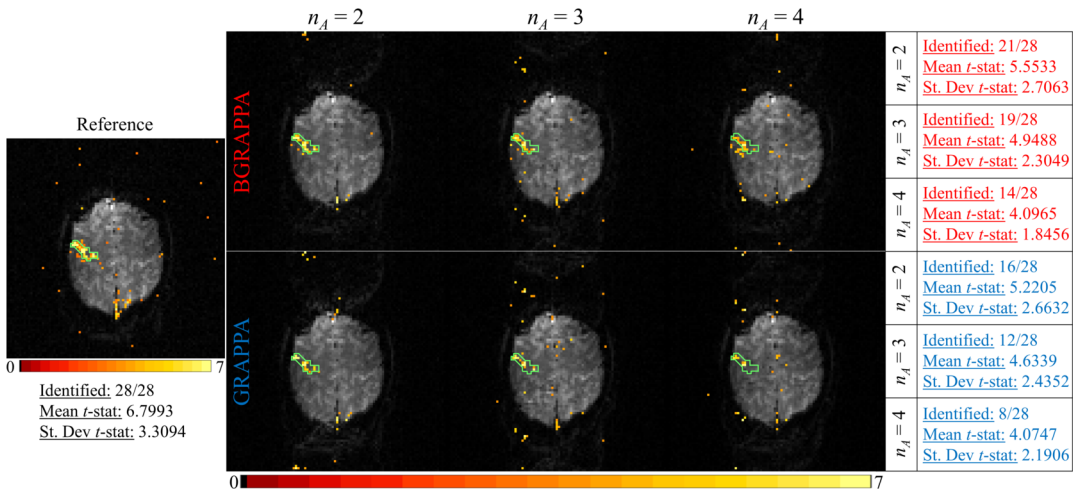


FIG. 15. Statistically significant voxels in the ROI using FDR for BGRAPPA reconstructed images (first row of the right three columns) for three different acceleration factors, significant voxels in the ROI using FDR for GRAPPA (second row of the right three columns) for three different acceleration factors, and analysis of the t -statistics to the right of the images. The reference magnitude-only task activation is shown on the left column with the analysis of t -statistics below.

space resulting in a single composite brain image. Figure 14 displays the BGRAPPA MAP reconstructed images (top row) and the GRAPPA reconstructed images (bottom row) of the first time point of the 490 images using acceleration factors 2, 3, and 4. Just as the simulated results in Figure 10 demonstrated, the BGRAPPA reconstruction method in Figure 14 produced visually similar magnitude reconstructed images, compared to GRAPPA reconstruction, but with slightly less noise.

MSE was again utilized to quantify the differences between the reference image and reconstructed images. The MSE for inside the brain for GRAPPA was approximately 12%, 10%, and 3% higher for each acceleration factor, respectively, compared to BGRAPPA. GRAPPA having a larger MSE inside the brain for each acceleration factor, respectively, reflects decreased noise from BGRAPPA vs. GRAPPA. The entropy for BGRAPPA (214.1026, 207.5331, and 204.1746, respectively) was also lower than the entropy for GRAPPA (216.0362, 212.3556, and 210.3667, respectively), indicating that the BGRAPPA reconstructed images are more smooth. The phase of the experimental reconstructed images for the different acceleration factors can be found in Section 2 of the Supplementary Material (Sakitis and Rowe (2025)).

For the detection of magnitude task activation, the hypothesis test outlined in Section 4.3 was carried out. Figure 15 shows the statistically significant voxels under BGRAPPA (top row) and GRAPPA (bottom row) reconstruction. The images for the statistically significant voxels in Figure 15 for both methods use the 5% FDR threshold. Voxels outside the brain are usually masked out, meaning the statistically significant voxels shown outside the brain in Figure 15 would typically be discarded. Figure 15 also summarizes the t -statistics with BGRAPPA (red) and GRAPPA (blue). BGRAPPA correctly detected more voxels than GRAPPA as task activation in the ROI for all three acceleration factors. Our BGRAPPA approach also had a much higher mean t -statistic and a lower standard deviation for all the acceleration factors.

6. Discussion. Parallel imaging techniques such as GRAPPA (Griswold et al. (2002)) have utilized subsampling of k -space to reduced the acquisition time for MR imaging. This allows practitioners to reconstruct higher resolution images, decrease the time between each

image, increase the number of images or slices in an fMRI experiment, or a combination of both in the same time as fully sampled k -space, depending on the acceleration factor. The acceleration factor in an fMRI experiment is determined by how important time is in completing a scan. The number of coils used in an experiment is dependent on the coil configurations that facility possesses.

Applying an acceleration factor in an fMRI experiment can significantly reduce the acquisition time of spatial frequency arrays and volume images, but taking the IFT of the subsampled k -space yields aliased images. GRAPPA parallel image reconstruction estimates the unacquired spatial frequencies that are skipped during the acquisition of the subsampled k -space arrays yielding full FOV coil spatial frequency arrays. However, GRAPPA has its drawbacks, which include low image quality, low SNR, and weakened task detection power at higher acceleration factors. Hence, we introduce a Bayesian approach to estimate the unacquired spatial frequencies. Using more available information from the calibration spatial frequencies to assess the hyperparameters, our proposed approach successfully reconstructed a series of simulated nontask images without any aliasing artifacts. The BGRAPPA reconstructed images were shown to more accurately reconstruct the truth compared to GRAPPA. The number of calibration time points had minimal effect on the GRAPPA reconstructed images and its performance against BGRAPPA reconstructed images. The results also indicated that the different acceleration factors had little effect on the reconstruction of the images but did have lesser task detection power for both methods. Our BGRAPPA approach had better performance when detecting the signal increase in the voxels that experienced task activation, which is demonstrated from both the simulated and experimental data.

For this paper, only the MAP estimate using the ICM algorithm was used to reconstruct the time series both the simulated and experimental data. Since we have posterior conditionals for each of the parameters, this allows us to use other estimation techniques such as the MCMC Gibbs sampling method. We chose not to present this method due to the Gibbs sampler being more computationally expensive when running a long series of images so it is not be as practical to use compared to evaluating the MAP estimate. This does not mean there is no value in running a Gibbs sampler, as it has the additional benefit of quantifying uncertainty. For instance, it can be utilized on a shorter series of images, provide us more statistical information about any voxel, hypothesis testing between two reconstructed images, or identifying which voxels are outside the brain for masking. We could also hybridize the ICM and Gibbs sampler where we start with a few iterations of the ICM algorithm followed by a short, no-burn Gibbs sampler. Our Bayesian approach allows for more options of how to run an fMRI experiment based on the objective of the scan compared to GRAPPA.

In this paper a magnitude-only and phase-only activation model was utilized to detect task activation. Due to the high noise in the experimental data set, there were no phase active voxels as there were in the idealized simulation. Since the reconstructed images are complex-valued, our proposed model is expected to be applicable for complex activation models for task detection (Rowe and Logan (2004), Rowe (2005, 2009)) as well as magnetic field mapping. Further, our proposed procedure can also be repeated for vertical aliasing, as opposed to the horizontal aliasing used here.

Acknowledgments. The authors thank the Wehr Foundation, as this research was funded by the Computational Sciences Summer Research Fellowship (CSSRF) at Marquette University in the Department of Mathematical and Statistical Sciences. This research was also funded by the Regular Research Grant (RRG) provided by the Committee on Research at Marquette University.

SUPPLEMENTARY MATERIAL

Supplement to “Bayesian approach to GRAPPA parallel image reconstruction increases SNR and power of task detection” (DOI: [10.1214/24-AOAS1962SUPP](https://doi.org/10.1214/24-AOAS1962SUPP); .zip). The supplement to this paper provides additional results for magnitude and phase reconstructed images, subsampling the calibration time points for separate hyperparameter assessment, and details about the experimental data.

REFERENCES

- ATKINSON, D., HILL, D. L. G., STOYLE, P. N. R., SUMMERS, P. E. and KEEVIL, S. F. (1997). Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion. *IEEE Trans. Med. Imag.* **16** 903–910.
- BANDETTINI, P., JESMANOWICZ, A., WONG, E. and HYDE, J. S. (1993). Processing strategies for time-course data sets in functional MRI of the human brain. *Magn. Reson. Med.* **30** 161–173.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](https://doi.org/10.1111/j.2517-6161.1995.tb00331.x)
- GENOVESE, C. R., LAZAR, N. A. and NICHOLS, T. E. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* **15** 870–878.
- GRISWOLD, M. A., JAMOB, P. M., HEIDEMANN, R. M., NITTKA, M., JELLUS, V., WANG, J., KIEFER, B. and HAASE, A. (2002). Generalized autocalibrating partially parallel acquisition (GRAPPA). *Magn. Reson. Med.* **47** 1202–1210.
- HYDE, J. S., JESMANOWICZ, A., FRONCISZ, W., KNEELAND, J. B., GRIST, T. M. and CAMPAGNA, N. F. (1986). Parallel image acquisition from noninteracting local coils. *J. Magn. Reson.* (1969) **70** 512–517.
- KARAMAN, M. M., BRUCE, I. P. and ROWE, D. B. (2014). A statistical fMRI model for differential T_2^* contrast incorporating T_1 and T_2^* of gray matter. *Magn. Reson. Imaging* **32** 9–27.
- KARAMAN, M. M., BRUCE, I. P. and ROWE, D. B. (2015). Incorporating relaxivities to more accurately reconstruct MR images. *Magn. Reson. Imaging* **33** 374–384.
- KORNAK, J., YOUNG, K., SCHUFF, N., DU, A., MAUDSLEY, A. A. and WEINER, M. W. (2010). K-Bayes reconstruction for perfusion MRI. I: Concepts and application. *J. Digit. Imag.* **23** 277–286.
- KUMAR, A., WELTI, D. and ERNST, R. R. (1975). NMR Fourier zeugmatography. *J. Magn. Reson.* (1969) **18** 69–83.
- LINDLEY, D. V. and SMITH, A. F. M. (1972). Bayes estimates for the linear model. *J. Roy. Statist. Soc. Ser. B* **34** 1–41. [MR0415861](https://doi.org/10.1111/j.2517-6161.1972.tb00792.x)
- LINDQUIST, M. A. (2008). The statistical analysis of fMRI data. *Statist. Sci.* **23** 439–464. [MR2530545](https://doi.org/10.1214/09-STS282) <https://doi.org/10.1214/09-STS282>
- LINDQUIST, M. A., ZHANG, C., GLOVER, G. and SHEPP, L. (2008). Rapid three-dimensional functional magnetic resonance imaging of the initial negative BOLD response. *J. Magn. Reson.* **191** 100–111.
- LOGAN, B. R. and ROWE, D. B. (2004). An evaluation of thresholding techniques in fMRI analysis. *NeuroImage* **22** 95–108. <https://doi.org/10.1016/j.neuroimage.2003.12.047>
- O’HAGAN, A. (1994). *Kendall’s Advanced Theory of Statistics. Vol. 2B. Kendall’s Library of Statistics.* Wiley, New York. [MR1285356](https://doi.org/10.1080/01621459408838535)
- OGAWA, S., LEE, T. M., NAYAK, A. S. and GLYNN, P. (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn. Reson. Med.* **14** 68–78.
- PRUESSMANN, K. P., WEIGER, M., SCHEIDEGGER, M. B. and BOESIGER, P. (1999). SENSE: Sensitivity encoding for fast MRI. *Magn. Reson. Med.* **42** 952–962.
- ROWE, D. B. (2005). Modeling both the magnitude and phase of complex-valued fMRI data. *NeuroImage* **25** 1310–1324. <https://doi.org/10.1016/j.neuroimage.2005.01.034>
- ROWE, D. B. (2009). Magnitude and phase signal detection in complex-valued fMRI data. *Magn. Reson. Med.* **62** 1356–1357.
- ROWE, D. B. and LOGAN, B. R. (2004). A complex way to compute fMRI activation. *NeuroImage* **23** 1078–1092. <https://doi.org/10.1016/j.neuroimage.2004.06.042>
- ROWE, D. B., MELLER, C. P. and HOFFMANN, R. G. (2007). Characterizing phase-only fMRI data with an angular regression model. *J. Neurosci. Methods* **161** 331–341. <https://doi.org/10.1016/j.jneumeth.2006.10.024>
- SAKITIS, C. J., BROWN, D. A. and ROWE, D. B. (2025). A Bayesian complex-valued latent variable model applied to functional magnetic resonance imaging. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **74** 100–125. [MR4863100](https://doi.org/10.1093/jrssc/qlae046) <https://doi.org/10.1093/jrssc/qlae046>
- SAKITIS, C. J. and ROWE, D. B. (2025). Supplement to “A Bayesian approach to GRAPPA parallel fMRI image reconstruction increases SNR and power of task detection.” <https://doi.org/10.1214/24-AOAS1962SUPP>

STEINHOFF, S., ZAITSEV, M., ZILLES, K. and SHAH, N. J. (2001). Fast T_1 mapping with volume coverage. *Magn. Reson. Med.* **46** 131–140.