Bayesian Magnitude and Phase Estimation of *k*-Space Data Enhances Reconstructed Image Quality

John C. Bodenschatz, Department of Mathematical and Statistical Sciences, Marquette University

Daniel B. Rowe, Department of Mathematical and Statistical Sciences, Marquette University

Abstract

Many real-world applications in digital signal processing demand the use of mathematical and statistical methods that support complex-valued data. We examine the case of signal degradation over time, in particular, as it applies to functional magnetic resonance imaging (fMRI). This attenuated signal can be amplified through a variety of methods. We outline a formal Bayesian approach to enhance the noisy signal, increasing the signal-to-noise ratio. This model was evaluated through simulated data and applied to experimental fMRI data.

1. Introduction

Many real-world applications, in particular digital signal processes, require support for complex-valued data. It is also a common problem in digital signal applications to have a loss or degradation of signal quality over time as a result of various factors. One real-world application that experiences both of these challenges is functional magnetic resonance imaging (fMRI). In fMRI, complex-valued spatial frequency coefficients are measured over time. Initially these frequency coefficients have a very strong signal. However, due to physical properties of the bodily tissues and the operation of the machine, the signal quality quickly degrades to a steady state at a much lower strength than the initial measurements. This will be explored further in Section 4.

We choose to employ a Bayesian model to address this problem. The Bayesian approach provides three main advantages: computational expense, interpretable procedures, and results. We will show that *a priori* data that is of higher quality than the remaining measured data can be used to assess hyperparameters of known distributions that can be used to form *a posteriori* estimates of the data. This posterior estimated data will have increased signal magnitude and decreased noise, resulting in a significantly greater signal-to-noise ratio (SNR). Resulting data will have improved fidelity, a more accurate representation of the true object being imaged, which enhances the ability to draw conclusions and make determinations about the data.

2. Bayesian Approach

2.1 Likelihoods

As has been demonstrated, the real and imaginary parts of the measured complex-valued signal, x_R and x_I, are independent and identically normally distributed (den Dekker and Sijbers, 2005; Gudbjartsson and Patz, 1995; Rowe, 2023). The means are given by $\mu_R = \rho\cos\theta$ and $\mu_I = \rho\sin\theta$ with a common variance σ^2, where ρ is the true signal magnitude and θ is the true signal phase. The joint distribution of the real and imaginary components is given by

$$f(x_R, x_I|\rho,\theta,\sigma^2) = \frac{1}{(2\pi\sigma^2)^{1/2}}\exp\left[-\frac{(x_R-\rho\cos(\theta))^2}{2\sigma^2}\right]\frac{1}{(2\pi\sigma^2)^{1/2}}\exp\left[-\frac{(x_I-\rho\sin(\theta))^2}{2\sigma^2}\right]. \quad (1)$$

A transformation of variables from Cartesian to polar $(x_R, x_I) \to (r, \phi)$ using $x_R = r\cos(\phi)$ and $x_I = r\sin(\phi)$, on Equation 1 is performed. The Jacobian $J = r$ is calculated, and we arrive at the joint distribution

$$\mathrm{f}(r,\phi|\rho,\theta,\sigma^2) = \frac{r}{2\pi\sigma^2}\exp\left[-\frac{1}{2\sigma^2}[r^2+\rho^2-2r\rho\cos(\phi-\theta)]\right]. \tag{2}$$

We can integrate out the measured phase ϕ from Equation 2 to arrive at the Ricean marginal distribution for the measured magnitude r (Rice, 1944)

$$\mathrm{f}(r|\rho,\sigma^2) = \frac{r}{\sigma^2}\exp\left[-\frac{r^2+\rho^2}{2\sigma^2}\right] I_0\left(\frac{r\rho}{\sigma^2}\right). \tag{3}$$

Here I_0 is the zeroth order modified Bessel function of the first kind. The mean of this Ricean distribution is $\sigma\sqrt{\pi/2}L_{1/2}(-\rho^2/2\sigma^2)$ and the variance is given by $2\sigma^2+\rho^2-\frac{\pi\sigma^2}{2}L_{1/2}^2(-\rho^2/2\sigma^2)$, where $L_{1/2}$ is a Laguerre polynomial. If the true signal ρ is small, $\rho \approx 0$, Equation 3 becomes the Rayleigh distribution with mean $\sigma\sqrt{\pi/2}$ and variance $(4-\pi)\sigma^2/2$ (Rayleigh, 1880). If the true signal ρ is high, this becomes the normal distribution with mean ρ and variance σ^2. Integrating out the measured magnitude r from the joint distribution in Equation 2 gives an unnamed non-normal distribution marginal distribution for the measured phase ϕ,

$$f(\phi|\rho,\theta,\sigma^2) = \frac{1}{2\pi}\exp\left[-\frac{\rho^2}{2\sigma^2}\right]\left[1+\frac{\rho}{\sigma}\sqrt{2\pi}\cos(\phi-\theta)\exp\left[\frac{\rho^2\cos^2(\phi-\theta)}{2\sigma^2}\right]\Phi\left(\frac{\rho\cos(\phi-\theta)}{\sigma}\right)\right], \tag{4}$$

where $\Phi(x)$ is the cumulative distribution function of the standard normal distribution (Lathi, 1983; Rowe and Bodenschatz, 2025). When the true signal ρ is near zero, the measured phase will be uniformly distributed on $[-\pi,\pi]$ with mean 0 and variance $\pi^2/3$. When the true signal ρ becomes large, the distribution of the measured phase becomes normal with mean θ and variance σ^2/ρ^2. It is between these zero and high SNR scenarios that we are concerned with.

The joint distribution for the measured magnitude and phase as in Equation 2 can be expressed as the product of the marginal distribution for the magnitude and the conditional distribution for the phase,

$$f(r,\phi|\rho,\theta,\sigma^2) = f(r|\rho,\sigma^2)f(\phi|r,\rho,\theta,\sigma^2). \tag{5}$$

Dividing the joint distribution $f(r,\phi)$ in Equation 2 by the marginal distribution for ϕ in Equation 4, we see that the conditional distribution for ϕ is von Mises (Von Mises, 1918)

$$f(\phi|r,\rho,\theta,\sigma^2) = \frac{\exp\left[\frac{r\rho\cos(\phi-\theta)}{\sigma^2}\right]}{2\pi I_0\left(\frac{r\rho}{\sigma^2}\right)}. \tag{6}$$

This von Mises distribution has mean θ and variance $1-I_1(r\rho/\sigma^2)/I_0(r\rho/\sigma^2)$. In summary, the likelihoods for the measured data are $r \sim Rice(\rho,\sigma)$ and $\phi|r \sim VM(\theta, r\rho/\sigma^2\})$.

2.2 Priors

Using the same family of distributions, we select the priors to be conjugate to the likelihoods for a simpler estimation of the parameters and because of well-known data distributions. The prior distribution for the true magnitude ρ is Ricean,

$$f(\rho|\rho_0,\sigma^2,\gamma) = \frac{\rho}{\sigma^2/\gamma}\exp\left[-\frac{\rho^2+\rho_0^2}{2\sigma^2/\gamma}\right] I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right). \tag{7}$$

with hyperparameters ρ_0 and γ. Similarly to the marginal distribution for the measured magnitude r as in Equation 3, the distribution in Equation 7 becomes the Rayleigh distribution when $\rho_0 \sim 0$ and the normal distribution when ρ_0 is large. The prior distribution for the true phase θ is von Mises,

$$f(\theta|\rho,\sigma^2,\theta_0,\gamma) = \frac{\exp\left[\frac{\rho\rho_0\cos(\theta-\theta_0)}{\sigma^2/\gamma}\right]}{2\pi I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right)}, \tag{8}$$

with location hyperparameter θ_0. The von Mises distribution in Equation 8 becomes the uniform distribution on $[-\pi, \pi]$ when $\rho_0 \sim 0$ and the normal distribution when ρ_0 is large. Both the Ricean distribution in Equation 7 and the von Mises distribution in Equation 8 become vague with a small γ. Lastly, we asses a prior for the variance σ^2 with an inverse gamma distribution

$$f(\sigma^2|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)(\sigma^2)^{\alpha+1}} \exp\left[-\frac{\beta}{\sigma^2}\right], \tag{9}$$

with hyperparameters of shape α and scale β. The inverse gamma distribution becomes vague when α and β are small. In summary, the prior distributions are $\rho|\sigma \sim Rice(\rho_0, \sigma/\sqrt{\gamma})$, $\theta|\rho, \sigma \sim VM(\theta_0, \rho\rho_0/(\sigma^2/\gamma))$, and $\sigma \sim IG(\alpha, \beta)$.

2.3 Posteriors

The likelihoods from Equations 3 and 6 can be combined with the priors in Equations 7-9 to get the posterior distribution

$$f(\rho, \theta, \sigma^2|r, \phi) \propto f(r|\rho, \sigma^2) f(\phi|r, \rho, \theta, \sigma^2)$$

$$f(\rho|\sigma^2, \gamma) f(\theta|\rho, \sigma^2, \theta_0, \gamma) f(\sigma^2|\alpha, \beta), \tag{10}$$

Or

$$f(\rho, \theta, \sigma^2|r, \phi) \propto \frac{r}{\sigma^2} \exp\left[-\frac{r^2 + \rho^2}{2\sigma^2}\right] I_0\left(\frac{r\rho}{\sigma^2}\right) \frac{\exp\left[\frac{r\rho\cos(\phi - \theta)}{\sigma^2}\right]}{2\pi I_0\left(\frac{r\rho}{\sigma^2}\right)} \frac{\rho}{\sigma^2} \exp\left[-\frac{\rho^2 + \rho_0^2}{2\sigma^2/\gamma}\right]$$

$$\times I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right) \frac{\exp\left[\frac{\rho\rho_0\cos(\theta - \theta_0)}{\sigma^2/\gamma}\right]}{2\pi I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right)} \frac{\beta^\alpha}{\Gamma(\alpha)(\sigma^2)^{\alpha+1}} \exp\left[-\frac{\beta}{\sigma^2}\right]. \tag{11}$$

Some simplification leads to the joint posterior distribution

$$f(\rho, \theta, \sigma^2|r, \phi) \propto \frac{\beta^\alpha}{\Gamma(\alpha)} \exp\left[-\frac{1}{2\sigma^2}(\gamma\rho_0^2 + r^2 + 2\beta)\right] \frac{\rho}{(\sigma^2)^{\alpha+3}} \exp\left[-\frac{1}{2\sigma^2}\Big((\gamma + 1)\rho^2 - 2\rho(\rho_0\cos(\theta - \theta_0) + r\cos(\phi - \theta))\Big)\right]. \tag{12}$$

The posterior conditional distributions can be found by considering only the terms that contain the parameter of interest. The posterior conditional distribution for the magnitude ρ given the phase θ and variance σ^2 can be shown to be Modified-Half-Normal (MHN) distributed (Sun et al., 2023)

$$f(\rho|\theta, \sigma^2, r, \phi) = \frac{2B^{A/2}\rho^{A-1}}{\Psi\left(\frac{A}{2}, \frac{C}{\sqrt{B}}\right)} \exp[-B\rho^2 + C\rho], \tag{13}$$

where Ψ is the Fox-Wright Psi function (Fox, 1928; Wright, 1935), $A = 2$, $B = (\gamma + 1)/(2\sigma^2)$, and $C = [\rho_0\gamma\cos(\theta - \theta_0) + r\cos(\phi - \theta)]/\sigma^2$. The MHN is a relatively new distribution with few known applications at the time of this writing. Further, it can be shown that the posterior conditional distribution for the phase θ given the magnitude ρ and variance σ^2 is von Mises

$$f(\theta|\rho, \sigma^2, r, \phi) = \frac{\exp[\kappa\cos(\theta - \lambda)]}{2\pi I_0(\kappa)}, \tag{14}$$

where $\kappa = c\rho/\sigma^2$, $a = \rho_0\gamma\cos(\theta_0) + r\cos(\phi)$, $b = \rho_0\gamma\sin(\theta_0) + r\sin(\phi)$, $c = \text{sign}(a)\sqrt{a^2 + b^2}$, and $\lambda = \arctan(b/a)$. Lastly, the posterior distribution of the variance σ^2 given the magnitude ρ and phase θ is an inverse gamma distribution

$$f(\sigma^2|\rho, \theta, r, \phi) = \frac{(\beta^*)^{\alpha^*}}{\Gamma(\alpha^*)(\sigma^2)^{\alpha^*+1}} \exp\left[-\frac{\beta^*}{\sigma^2}\right], \tag{15}$$

where $\beta^* = [(\gamma+1)\rho^2 - 2\rho(\rho_0\gamma\cos(\theta-\theta_0) + r\cos(\phi-\theta)) + (\gamma\rho_0^2 + r^2 + 2\beta)]/2$ and $\alpha^* = \alpha + 2$. In summary, the posterior conditional distributions are $\rho | \theta, \sigma^2 \sim MHN(A, B, C)$, $\theta | \rho, \sigma^2 \sim VM(\lambda, \kappa)$, and $\sigma^2 \mid \rho, \theta \sim IG(\alpha^*, \beta^*)$. To our knowledge, the MHN distribution has recently been developed but has not been applied.

2.4 Estimation

We will partition the data such that the first n_0 complex-valued measurements are used as calibration measurements. These will be used to objectively assess the hyperparameters of the prior distributions which will then be combined with the steady-state signal to form an enhanced posterior signal. The average magnitude of the complex-valued calibration data is computed to obtain the prior magnitude mean $\rho_0 = \left(\overline{x_R^2} + \overline{x_I^2}\right)^{1/2}$. The hyperparameter $\gamma = n_0$ is assessed to be the number of calibration measurements. The four-quadrant arctangent of the averaged complex-valued calibration data is computed to obtain the prior phase mean $\theta_0 = \arctan_4(\overline{x_R}/\overline{x_I})$. Lastly, the prior distribution quantity $\sigma_0^2 = (s_R^2 + s_I^2)/2$ is computed as the average variance of the real and imaginary parts of the calibration measurements. The hyperparameters of shape and scale are assessed to be $\alpha = n_0 - 1$ and $\beta = (n_0 - 1)\sigma^2$ respectively.

When only one value is needed, such as the mode, we can use the iterated conditional modes (ICM) algorithm to obtain a maximum *a posteriori* (MAP) estimate (Lindley and Smith, 1972; O'Hagan, 2003). For this method, we will use the ICM algorithm to get MAP estimates of the parameters. The modes of the conditional distributions in Equations 13-15 are given by

$$\hat{\rho} = \frac{C + \sqrt{C^2 + 8B(A-1)}}{4B}, \tag{16}$$

where $A = 2$, $B = (\gamma+1)/(2\sigma^2)$, and $C = [\rho_0\gamma\cos(\theta-\theta_0) + r\cos(\phi-\theta)]/\sigma^2$,

$$\hat{\theta} = \arctan\left(\frac{b}{a}\right), \tag{17}$$

where $a = \rho_0\gamma\cos(\theta_0) + r\cos(\phi)$, $b = \rho_0\gamma\sin(\theta_0) + r\sin(\phi)$, and

$$\widehat{\sigma^2} = \frac{\beta^*}{\alpha^*+1}, \tag{18}$$

where $\beta^* = [(\gamma+1)\rho^2 - 2\rho(\rho_0\gamma\cos(\theta-\theta_0) + r\cos(\phi-\theta)) + (\gamma\rho_0^2 + r^2 + 2\beta)]/2$ and $\alpha^* = \alpha + 2$.

The ICM algorithm starts with an initial value for the parameters $\hat{\rho}$ and $\widehat{\sigma^2}$. Based on these values, an estimation of the parameter $\hat{\rho}$ is made using Equation 16. Each iteration of the algorithm will continue to obtain the MAP estimate for each of these parameters based on the most recent value of the other, using Equations 16-18. This is performed for each of the parameters until convergence is reached. In this work, this could be performed in as few as $L = 3$ ICM iterations but is generally allowed up to $L = 10$ at a negligible computational cost. If there exists a unique global maxima, as is the case with each of our unimodal posterior conditional distributions, then the iterative parameter estimates will converge to this maximum. Note that since Equation 17 does not depend on ρ or σ, the mode of the phase $\hat{\theta}$ can be immediately estimated without the ICM algorithm.

Since we have the posterior conditional distributions in Equations 13-15, we could also perform a Gibbs sampling to obtain an entire distribution and compute any desired quantity such as marginal posterior means (MPM) (Geman and Geman, 1984; Gelfand and Smith, 1990). An initial estimate for the parameters ρ, θ, and σ^2 is made. Following this, random samples are collected from each of the posterior conditional distributions given by Equations 13-15 using the initial estimates. This step is repeated for L iterations, each using the latest sample for each parameter. The first B samples are often burned, leaving a final set of $L - B$ samples for each parameter. The samples are then averaged for each parameter, resulting in MPM estimates. Random samples from the MHN distribution for the magnitude ρ are generated using an acceptance-rejection method described by

Algorithm 1 from Sun et al. (2023). Random samples from the von Mises distribution for the phase θ are also generated from an acceptance-rejection method that employs the wrapped Cauchy distribution, as outlined in Section 4 from Best and Fisher (1979). Lastly, random samples from the inverse gamma distribution for the variance σ^2 are generated directly using MATLAB's built in gamrnd() function.

3. Illustrative Example

To demonstrate the effectiveness of the described method in enhancing complex-valued signals, a simple simulated experiment will be performed. Consider a single complex-valued signal that is measured over 100 time points where the first three measurements are of higher signal magnitude than the remaining 97 (which are at a steady state with 1/3 of the intensity of the first measurement). The magnitude and phase of this original data set are shown in black in Figure 1. Prior estimates for the magnitude and phase were calculated using the method described in Section 2.4, with the first three measurements serving as calibration data. The remaining 97 measurements serve individually as the likelihood. The ICM algorithm and Gibbs sampling technique were then both implemented to obtain the MAP and MPM estimates respectively. The ICM algorithm was run using $L = 15$ iterations and took about 0.02 seconds total. Generally, convergence was reached after only three iterations. There were 5,000 Gibbs samples including a 500-sample burn-in phase, resulting in 4500 samples for MPM estimation at each time point. This took a total of 2 seconds to simulate using MATLAB's parallel computing toolbox (it took 8 seconds to run in serial). The MAP and MPM estimates are shown in blue and red in Figure 1. Both posterior estimation methods similarly improve the original data with a noticeable suppression of noise.

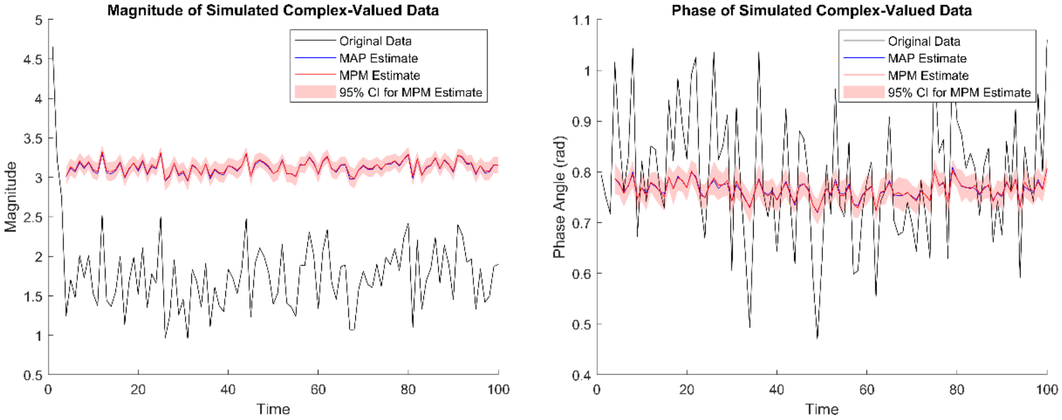

Figure 1: Magnitude and phase of the simulated complex-valued data. The original data has a much lower SNR than the improved MAP and MPM estimates.

4. FMRI Application

FMRI is a commonly used non-invasive imaging technique that allows physicians and scientists to observe the functionality of organs such as the human brain. This is done by exciting hydrogen nuclei in the various water molecules that make up the different tissues of the organ with a radio frequency (RF) pulse, then measuring a resulting net change in magnetization as determined by an induced current in a loop of wire that surrounds the patient. The intensity of the net magnetization in different voxels of the region of interest (ROI) is associated with complex-valued spatial frequency coefficients that fill *k*-space; a high order approximation of the Fourier transform of the image of the organ. A representation of *k*-space is depicted in Figure 2, showing that each element in *k*-space is composed of a complex-valued spatial frequency coefficient. The *k*-space is then inverse discrete Fourier transformed (IDFT) to reconstruct an image of the slice that is being observed. In fMRI experiments, this process is repeated many times; each image taking about 1 second to collect. A series of “task” and “non-task” images are collected where the patient is instructed to perform some task, such as right-hand finger tapping. This results in an expected increase of signal in the “task” images in certain regions of the brain, in this case- the left motor cortex.

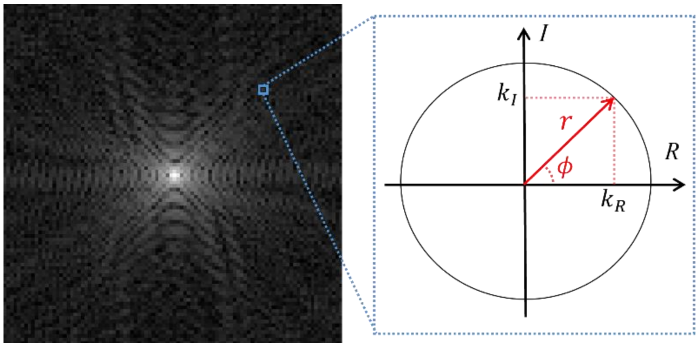

Figure 2: Representation of the two main coordinate systems to work with k*-space measurements. The real/imaginary form* (k_R, k_I) *lies in a Cartesian space, while the magnitude/phase form* (r, ϕ) *lies in radial space.*

It is well known in fMRI studies that the first three or so images in a time series have much higher signal than the remainder of the time series. Assuming the experiment starts in thermal equilibrium, the net magnetization is entirely in the longitudinal direction, so $M_Z(0) = M_0$ and $M_{XY}(0) = 0$. After the first pulse RF_1 at some flip angle α, the longitudinal and transverse magnetization will become $M_Z(RF_1) = M_0 \cos(\alpha)$ and $M_{XY}(RF_1) = M_0 \sin(\alpha)$. Due to the fact that $TR \approx T_1$ for fMRI experiments, the longitudinal component of the magnetization, M_Z, will not yet be fully recovered by the time of the second pulse RF_2. After a sufficient number of RF pulses (generally two or three), the system will reach a steady state where each RF pulse beyond this point results in the same magnetization value for the longitudinal direction, called M_{SS}. An example of longitudinal magnetization after the first few RF pulses is shown graphically in Figure 3. Since these first three reconstructed images do not match the observed signal from the remainder of the time series, they are often discarded, and only the steady-state images are analyzed.

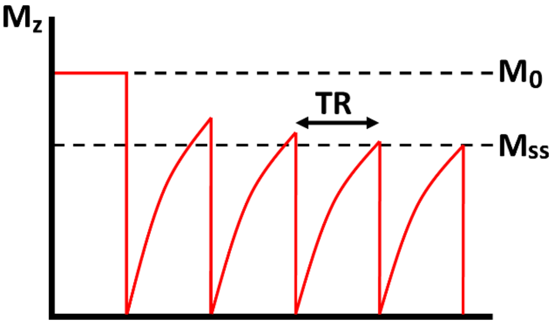

Figure 3: Magnetic saturation through the first few images of an fMRI time series. The intensity of M_z *directly correlates to image signal contrast. The steady state* M_{SS} *is usually achieved after three images.*

When Bayesian analysis of fMRI data is done on image-space data, voxels may be spatially correlated, demanding the use of a spatial model (Penny et al., 2005; Wang et al., 2024; Yu et al., 2023). It is well understood, however, that the Fourier transform of correlated voxels results in uncorrelated spatial frequency coefficients (Kornak et al., 2024). Thus, we chose to employ our Bayesian method on the uncorrelated *k*-space data since the spatial frequency coefficients can be treated independently of each other. Since the analog-to-digital converters (ADCs) collect *k*-space measurements independently, it is understood that the real and imaginary parts of a given *k*-space measurement (k_R, k_I) are independent and identically distributed (*iid*) normally for each spatial frequency location. The transformation $k_R = r_k \cos(\phi_k)$ and $k_I = r_k \sin(\phi_k)$ allows us to work with the *k*-space measurements in polar space, as visualized in Figure 2. This experimental setup precisely matches the theory established in Section 2.

Previous work on this Bayesian image analysis in Fourier space (BIFS) has been done on polar coordinates (Kornak et al., 2024). The BIFS method operates in a Fourier space that is constructed from the Fourier transform of some noisy real-valued image, resulting in a Hermitian symmetric Fourier space. The method then applies non-conjugate priors which increase computational burden when using Markov Chain Monte Carlo (MCMC) methods like the Metropolis-Hastings algorithm. This is then inverse Fourier transformed back into some altered (often de-noised) image. In fMRI the complex-valued spatial frequency coefficients in *k*-space are measured directly and then inverse Fourier transformed into image space. The methods outlined in Section 2 as applied to this problem are similar to BIFS, but done directly on a measured Fourier space (*k*-space) and use mathematically correct conjugate prior distributions which leads to simpler estimations. All distributions and parameters in this section are analyzed in *k*-space unless stated otherwise.

4.1 Simulated and Experimental Data

4.1.1 Data Description

The following fMRI time series data was simulated using SHAKER v1.1 (Bodenschatz and Rowe, 2025). The simulated time series is of slice 83 from a size 128 × 128 phantom in the Axial plane. The MRI parameters were set to be the following: Acceleration Factor = 1, Field Strength = 3 T, TE = 50 ms, TR = 1000 ms, Flip Angle = 90 deg, EESP = 0.832 ms, and Number of Coils = 1. The data was simulated with the Gradient Echo signal equation using a Cartesian *k*-space trajectory. The experimental design involved an initial 16 rest images followed by 19 epochs, each consisting of 16 task images followed by 16 rest images, for a total of 624 images. In voxels designed to be active, the SNR was set to 5 and the CNR was set to 0.25. There were 3 degrees of phase added to the task-related activation. Images were reconstructed using the inverse Fourier transform. These options were chosen to maintain consistency with the experimental data. Figure 4 shows the magnitude and phase of the first four *k*-space arrays and corresponding reconstructed images in the time series. The remainder of the time series appears similarly to TR 4.

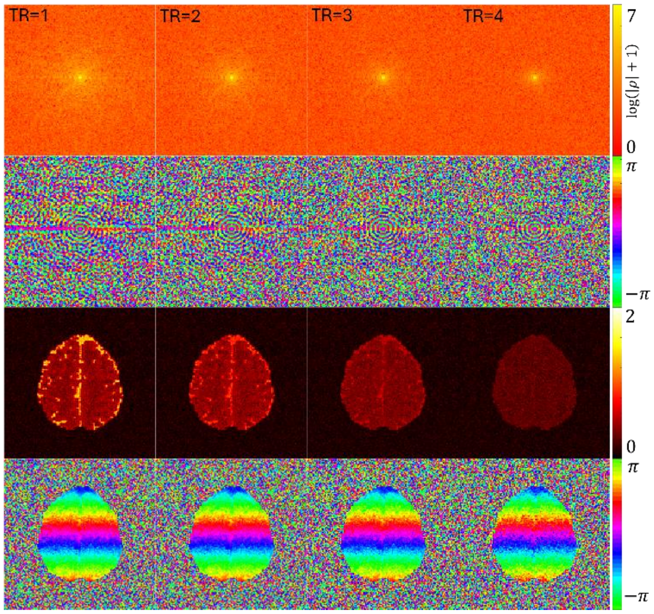

Figure 4: The magnitude and phase of both k-space and reconstructed images for the first four TRs in the experimental time series. The color bar is consistent across TRs so the signal degradation is clear.

The experimental data used in this section is from a block-design unilateral right-hand finger tapping experiment collected on a 3.0-Tesla General Electric Signa LX MRI scanner. Brain image data was collected using a body coil that resulted in a lower SNR than what is typically expected. There were seven excited axial slices $n_s = 7$ of 2.5 mm thickness with an array size of 128 × 128 with a field-of-view FOV = 24 cm, echo time TE = 50 ms, effective echo spacing EESP = 0.832 ms, and repetition time TR = 1 s. For this paper, we look at the time series for only the sixth slice of the seven collected. The experiment included 16 initial seconds of rest followed by 19 epochs each consisting of 16 seconds of task followed by 16 seconds of rest, resulting in 624 total images $n_t = 624$. It is often observed that 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 (Sakitis and Rowe, 2025). First, the angular phase temporal mean of each k-space element's time-series is calculated and angularly subtracted. A local second order polynomial was spatially fit to the resultant difference of the k-space element phase time-series. Then the spatially fitted phase is angularly subtracted from the original phase, the residual phase was unwrapped and linearly detrended, then the angular phase temporal mean was angularly added back producing a linearly stable phase over time for each spatial frequency. Figure 5 shows the magnitude and phase of the first four k-space arrays and corresponding reconstructed images in the time series. The remainder of the time series appears similarly to TR 4.

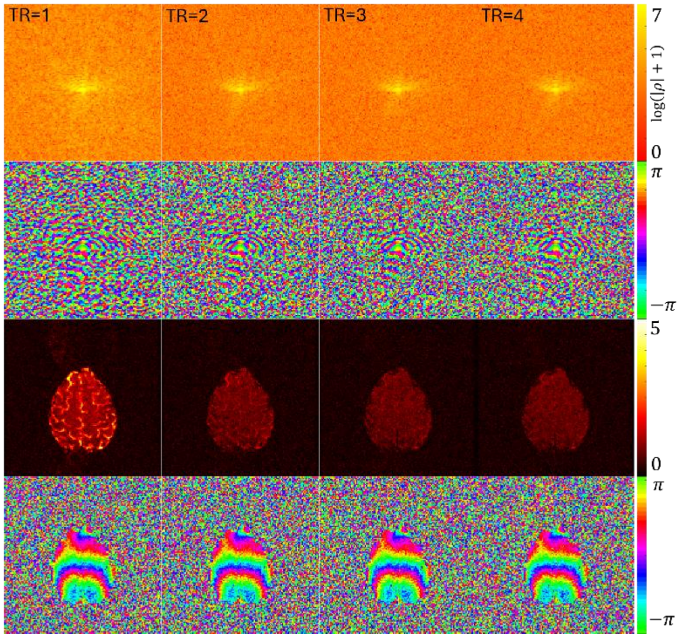

Figure 5: The magnitude and phase of both k-space and reconstructed images for the first four TRs in the experimental time series. The color bar is consistent across TRs so the signal degradation is clear.

4.1.2 Estimation

The first three k-space arrays have decreasing intensity and are not consistent with the remaining arrays. In common practice, these measurements would be omitted prior to computing activation. Instead, we will use the first $n_0 = 3$ k-space arrays as calibration data, or "prior" information. These will be used to assess the hyperparameters of the prior distributions that will be combined with the steady-state k-space observations to form enhanced posterior reconstructed images. The remaining 621 k-space arrays will be considered as the "observed" data. The estimation method will be similar to that described in Section 2.4 and exemplified in Section 3, but applied to each k-space element through time individually ($128 \times 128 \times 621 = 10{,}174{,}464$ complex-valued data points in total). At each k-space spatial frequency, the magnitude of the averaged prior

k-space measurements is computed for the prior mean $\rho_0 = \left(\overline{k_R^2} + \overline{k_I^2}\right)^{1/2}$. The hyperparameter $\gamma = n_0$ is assessed to be the number of calibration arrays. We set $\gamma = n_0 = 3$ due to the previous theoretical description on magnetic saturation. The four-quadrant arctangent of the averaged prior k-space measurements is computed for the prior mean $\theta_0 = \arctan_4\left(\overline{k_R}/\overline{k_I}\right)$. Lastly, the prior distribution quantity $\sigma_0^2 = (s_R^2 + s_I^2)/2$ is computed as the average of the variance of the real and imaginary parts of the prior k-space measurements. The prior hyperparameters are assessed to be $\alpha = n_0 - 1 = 2$ and $\beta = (n_0 - 1)\sigma^2 = 2\sigma^2$.

The magnitude and phase of each k-space spatial frequency location is individually estimated. We implement both the ICM algorithm and Gibbs sampling techniques to obtain posterior estimates. Similarly to Section 3, the ICM algorithm was run with $L = 15$ iterations at each time point for each k-space element, but in practice it took as few as 3 iterations to reach convergence. For the simulated data there were 200 Gibbs samples generated including a 50-sample burn-in phase, resulting in 150 samples for MPM estimation at each time point for each k-space element. For the experimental data, there were 5,000 Gibbs samples generated including a 500-sample burn-in phase, resulting in 4500 samples for MPM estimation at each time point for each k-space element. This was done using MATLAB's parallel computing toolbox. The experimental ICM data takes only a few seconds to run on a single core, while the Gibbs samples took several days in parallel across 20 cores. The precomputed prior k-space arrays and corresponding reconstructed images as well as the likelihood and posterior estimates for the final TR in steady state are presented in Figures 6 and 7. The second column, "Observed" is representative of the latter 621 images in the time series. Both Bayesian methods enhance image contrast and decrease noise when compared to the observed data. Increased image contrast can lead to improved image segmentation, feature extraction, and classification (Zhou et al., 2019).

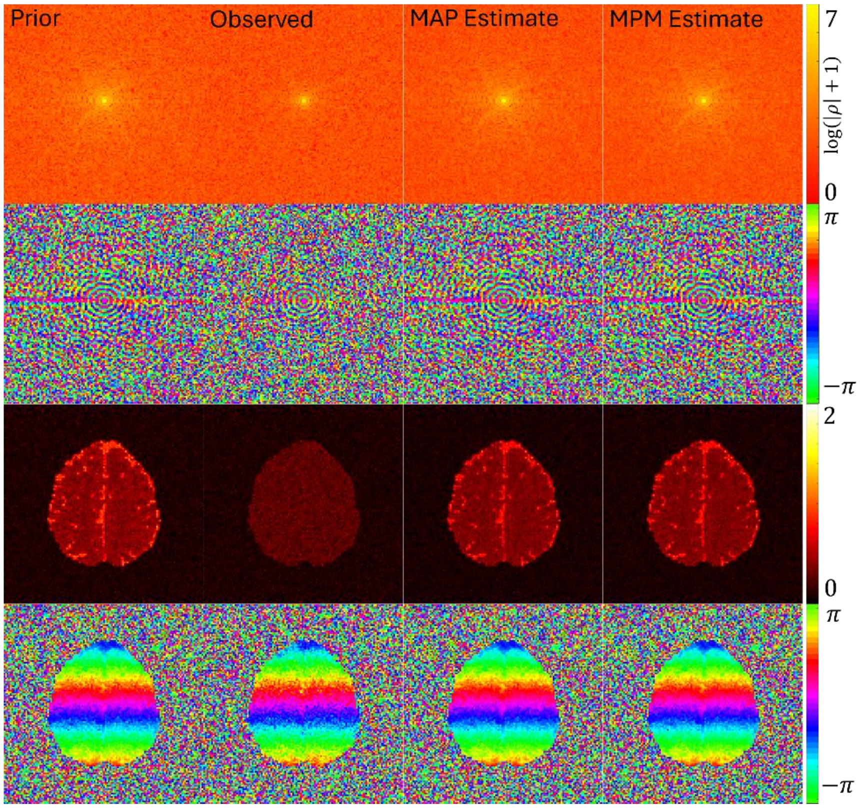

Figure 6: The magnitude and phase of both k-space and reconstructed images for the prior distribution, likelihood (simulated data at $TR = 624$), and posterior estimates of the simulated data. Both Bayesian methods enhance image contrast when compared to the observed data.

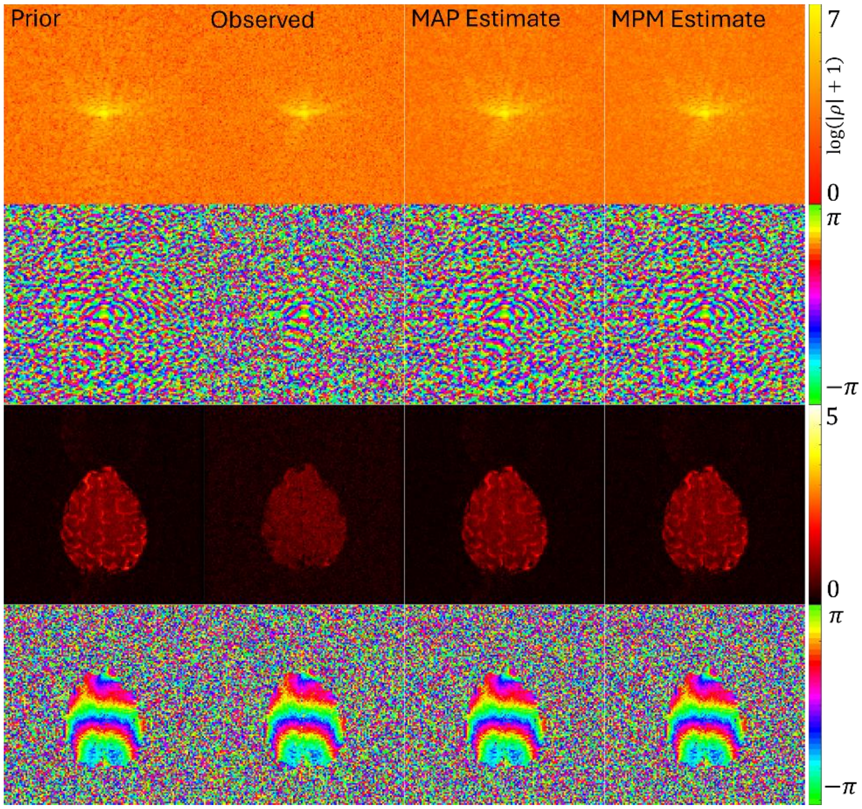

Figure 7: The magnitude and phase of both k-space and reconstructed images for the prior distribution, likelihood (measured data at $TR = 624$), and posterior estimates of the experimental data. Both Bayesian methods enhance image contrast when compared to the observed data.

4.1.3 Results

Task-based fMRI for an individual voxel's magnitude time series r_t can be expressed as the linear equation

$$r_t = \beta_0 + \beta_1 x_t + \varepsilon_r. \tag{19}$$

The additive noise ε_r is Ricean distributed with variance σ_r^2. Here, $\beta_0 \in R^+$ is the baseline signal which determines the signal-to-noise ratio $SNR = \beta_0/\sigma_r$, and $\beta_1 \in R$ is the task-related signal increase which determines the contrast-to-noise ratio $CNR = \beta_1/\sigma_r$. The design vector $x_t \in \{0,1\}^{n_t}$ has length equal to the number of reconstructed images in the time series, n_t. In x, indices corresponding to a non-task image have an element of 0, while indices corresponding to a task-active image have an element of 1.

The SNR for the original data, the MAP estimates, and the MPM estimates is shown in Figures 8 and 9 (top row). Note that the upper scaling for the color bar is increased from 5 to 30 for the Bayesian posterior estimates. Both Bayesian methods similarly improved the SNR when compared to the original data. Slight timing differences when measuring spatial frequency coefficients in k-space result in a Nyquist ghost above and below the brain, this is apparent in the SNR of the Bayesian methods in the experimental data. The CNR is an indication of where the increase in task activation-related signal is measured. The CNR calculated from the original data and the two Bayesian methods is shown in Figures 8 and 9 (bottom row) with a fixed color bar. Both posterior estimations accurately capture the task activation in the left motor cortex associated with the right hand finger tapping, as indicated by the yellow region. When comparing the posterior estimated data to the original simulated and experimental data, it is clear that we have significantly decreased the effect of noise on the data, while still maintaining fidelity to detect task-related activation. Having improved contrast in images, like those produced by the MAP and MPM estimation methods, will allow for better image registration when correcting for motion as well as assist image segmentation when labeling regions of interest. We recommend the use of the ICM algorithm to generate MAP estimates for this application. The temporal cost is minimal with vastly improved signal quality when compared to the original data. This approach preserves activation while decreasing noise and increasing SNR.

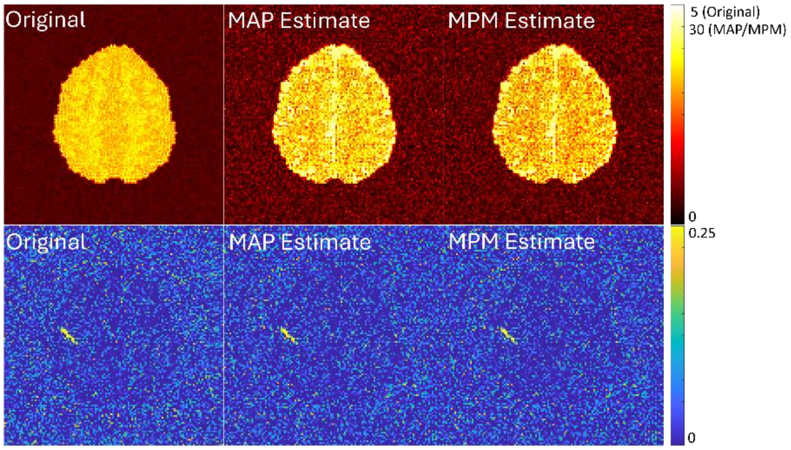

Figure 8: The SNR for the original data, MAP and MPM estimates (top row). Note the differing color bar scales- the MAP and MPM estimates have significantly increased SNR. The CNR for the original data, MAP and MPM estimates (bottom row). Both posterior estimates have similar CNR to the original data, indicating no loss in ability to detect task related activation.

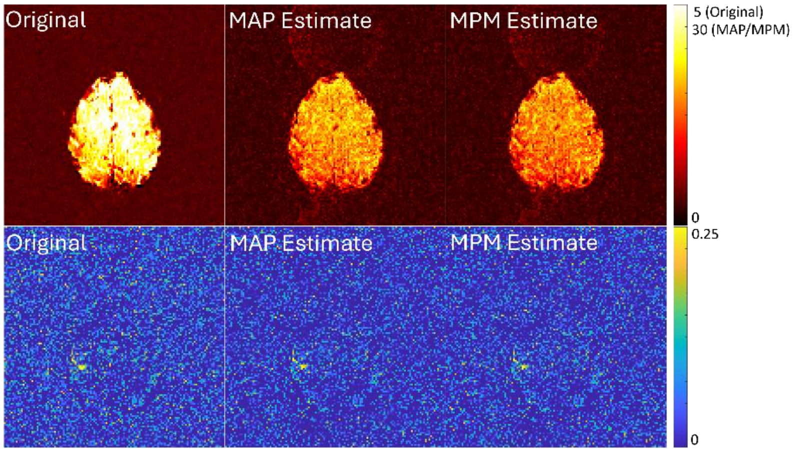

Figure 9: The SNR for the original data, MAP and MPM estimates (top row). Note the differing color bar scales- the MAP and MPM estimates have significantly increased SNR. The CNR for the original data, MAP and MPM estimates (bottom row). Both posterior estimates have similar CNR to the original data, indicating no loss in ability to detect task related activation.

5. Discussion

Systems that receive complex-valued signals can often suffer from signal loss and degradation over time. Here, we implement a Bayesian approach to enhance the noisy signal using only a small number of prior strong measurements. This process is demonstrated on simple simulated data in Section 3. The results of the simulated and experimental fMRI studies in Section 4 demonstrate that the model can increase the SNR of the original data, leading to higher contrast reconstructed images.

We used full posterior distributions for the estimation of parameters ρ, θ, and σ^2. This includes, to our knowledge, one of the first applications of the recently described MHN distribution (Sun et al., 2023). This allowed us to apply both the ICM algorithm and a Gibbs sampling technique to our model. The signal amplification for the experimental data was done using the ICM algorithm with $L = 15$ iterations and Gibbs sampling with 5,000 samples; both methods were applied at each *k*-space element at each time point. Due to the significant difference in computational expense with minimal effect on results, we recommend the use of the ICM algorithm to obtain MAP estimates. The Gibbs sampler has the added benefit of quantifying uncertainty since it generates distributions of estimates as compared to a point estimate from the ICM algorithm, so if computational time is not a concern this may be worth performing.

The ability to detect task in the MAP and MPM reconstructed images is not hindered by the use of rest images to form the prior. The MAP and MPM reconstructed images have an increased SNR and improved contrast from the original data, allowing for better image registration and segmentation. This will help physicians correct for motion artifacts and to more easily segment regions of interest within the brain.

References

Best, D.J., Fisher, N.I., 1979. Efficient simulation of the von Mises distribution. Journal of the Royal Statistical Society Series C: Applied Statistics 28, 152-157. doi: https://doi.org/10.2307/2346732.

Bodenschatz, J.C., Rowe, D.B., 2025. Simulation and harmonic analysis of *k*-space readout (SHAKER). URL: https://arxiv.org/abs/2502.17620,arXiv:2502.17620.

den Dekker, A.J., Sijbers, J., 2005. Implications of the Rician distribution for fMRI generalized likelihood ratio tests. Magnetic Resonance Imaging 23, 953-959. doi: https://doi.org/10.1016/j.mri.2005.07.008.

Fox, C., 1928. The asymptotic expansion of generalized hypergeometric functions. Proceedings of the London Mathematical Society s2-27, 389-400. doi: https://doi.org/10.1112/plms/s2-27.1.389.

Gelfand, A.E., Smith, A.F.M., 1990. Sampling-based approaches to calculating marginal densities. Journal of the American Statistical Association 85, 398-409. URL: http://www.jstor.org/stable/2289776.

Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-6, 721-741. doi: https://doi.org/10.1109/TPAMI.1984.4767596.

Gudbjartsson, H., Patz, S., 1995. The Rician distribution of noisy MRI data. Magnetic Resonance in Medicine 34, 910914. doi: https://doi.org/10.1002/mrm.1910340618.

Kornak, J., Young, K., Friedman, E., 2024. Bayesian image analysis in Fourier space. URL: https://arxiv.org/abs/2305.19481, arXiv:2305.19481.

Lathi, B., 1983. Modern Digital and Analog Communication. Oxford University Press.

Lindley, D.V., Smith, A.F.M., 1972. Bayes estimates for the linear model. Journal of the Royal Statistical Society. Series B (Methodological) 34, 1-41. URL: http://www.jstor.org/stable/2985048.

O'Hagan, A., 2003. Kendall's advanced theory of statistic 2B. Kendall's advanced theory of statistics, Wiley-Blackwell, Hoboken, NJ.

Penny, W.D., Trujillo-Barreto, N.J., Friston, K.J., 2005. Bayesian fmri time series analysis with spatial priors. NeuroImage 24, 350362. doi: https://doi.org/10.1016/j.neuroimage.2004.08.034.

Rayleigh, L., 1880. Xii. on the resultant of a large number of vibrations of the same pitch and of arbitrary phase. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 10, 7378. doi: https://doi.org/10.1080/14786448008626893.

Rice, S.O., 1944. Mathematical analysis of random noise. Bell System Technical Journal 23, 282332. doi: https://doi.org/10.1002/j.1538-7305.1944.tb00874.x.

Rowe, D.B., 2023. Statistics of intrinsic fMRI data. doi: https://doi.org/10.5281/zenodo.10002334.

Rowe D.B., Bodenschatz J.C., 2025. Bayesian k-Space Estimation Decreases Image Noise and Increased Activation Detection. Proc. Joint Stat. Meet., Section on Statistics in Imaging, Nashville, TN. doi: https://doi.org/10.5281/zenodo.16921786

Sakitis, C.J., Rowe, D.B., 2025. Bayesian merged utilization of grappa and sense (BMUGS) for in-plane accelerated reconstruction increases fMRI detection power. Magnetic Resonance Imaging 115, 110252. doi: https://doi.org/10.1016/j.mri.2024.110252.

Sun, J., Kong, M., Pal, S., 2023. The Modified-Half-Normal distribution: Properties and an efficient sampling scheme. Communications in Statistics - Theory and Methods 52, 15911613. doi: https://doi.org/10.1080/03610926.2021.1934700.

Von Mises, R., 1918. Uber die 'ganzzahligkeit' der atomgewicht und verwandte fragen. Phys. Z. 19, 490-500. URL: https://cir.nii.ac.jp/crid/1572543024364690176.

Wang, Z., Rowe, D.B., Li, X., Brown, D.A., 2024. A fully Bayesian approach for comprehensive mapping of magnitude and phase brain activation in complex-valued fMRI data. Magnetic Resonance Imaging 109, 271-285.doi:https://doi.org/10.1016/j.mri.2024.03.029.

Wright, E.M., 1935. The asymptotic expansion of the generalized hypergeometric function. Journal of the London Mathematical Society s1-10, 286-293. doi: https://doi.org/10.1112/jlms/s1-10.40.286.

Yu, C., Prado, R., Ombao, H., Rowe, D., 2023. Bayesian spatiotemporal modeling on complex-valued fMRI signals via kernel convolutions. Biometrics 79, 616628. doi: https://doi.org/10.1111/biom.13631.

Zhou, Y., Shi, C., Lai, B., Jimenez, G., 2019. Contrast enhancement of medical images using a new version of the world cup optimization algorithm. Quantitative Imaging in Medicine and Surgery 9. doi: https://doi.org/10.21037/qims.2019.08.19.