

IMPROVED SIMULATION, BAYESIAN ESTIMATION,
PHASE ACTIVATION, AND NON-CARTESIAN RECONSTRUCTION IN FMRI

by

John C. Bodenschatz, M.S.

A Dissertation submitted to the faculty of the Graduate School,
Marquette University,
in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy

Milwaukee, Wisconsin

April 24, 2026

ABSTRACT

IMPROVED SIMULATION, BAYESIAN ESTIMATION, PHASE ACTIVATION, AND NON-CARTESIAN RECONSTRUCTION IN fMRI

John C. Bodenschatz, M.S.

Marquette University,

April 24, 2026

In fMRI it is desirable to obtain images with high contrast and low noise at a high spatio-temporal resolution. Doing so allows for a more reliable assessment of capturing brain dynamics in higher detail. Many techniques have been employed and tested to collect more images per unit of time and/or increase spatial resolution, while maintaining a high quality image. This has been done using in-plane and through-plane acceleration with promising results (SENSE, GRAPPA, CAIPIRINHA, CAIPIVAT). Following a description of work done towards a unified fMRI simulation software package, this dissertation will present work towards three methods to enhance fMRI time series image quality using simulated data as the benchmark. The simulation tool presented is titled *Simulation and Harmonic Analysis of k -Space Readout* (SHAKER) and is designed to simulate fMRI time series for MR scientists. At present, most researchers involved in fMRI simulate their own data in-house and do not necessarily do so in a way that is representative of the machine. SHAKER aims to provide software that allows for fast simulations to be done without the need for advanced understanding of how an MRI machine works. The images and time series simulated by SHAKER will be used to demonstrate the efficacy of the three new techniques presented that will improve fMRI images. The first method makes use of the first three k -space arrays in fMRI time series that are often discarded due to having a higher signal than the remaining images. These brighter images will be used to assess hyperparameters for prior distributions that will be combined with distributionally accurate likelihood images from the steady-state time series to form posterior images that have higher signal and contrast, with lower noise. The second method focuses on phase-only activation in complex-valued fMRI time series. The phase half of the data is often discarded, and only the magnitude is studied. We will show that the phase part of the data contains biological information, in particular task-related signal change, that has exciting physiological implications. The third method introduced in this dissertation will operate in a radial k -space where instead of sampling on a Cartesian grid, points are collected on spokes that each pass through the center of k -space. This method of sampling k -space has many proven benefits, but is not without its challenges- in particular, reconstructing the k -space arrays back into images. A few, fully-sampled radial k -space arrays can be used to assess hyperparameters for prior distributions that can be combined with subsampled likelihood images from the fMRI time series to form posterior images. It is expected that these images will have higher signal and contrast, lower noise, and measured at significantly higher temporal resolutions than conventional subsampling techniques offer. The software tool SHAKER and the three techniques described are novel and significant contributions to fMRI analysis.

ACKNOWLEDGEMENTS

John Bodenschatz, M.S.

I would like to extend my sincerest gratitude to my advisor, Dr. Daniel B. Rowe, for being an incredible mentor throughout the Ph.D. process. Among the many wisdoms he has imparted on me, he has taught me how to be a more effective learner, researcher, and communicator. Dr. Rowe's experience, expertise, and guidance have provided a path for me to move forward in my career that I will be forever grateful for. I would like to thank Dr. Mehdi Maadooliat, Dr. Cheng-Han Yu, and Dr. Tugan Muftuler for their continued support and willingness to be a part of my doctoral committee and assist me in meeting my scholarly goals.

I am grateful for the Department Chair Dr. Anne Clough, Graduate Chair Dr. Sarah Hamilton and administrative assistant Oanh Pham for the help they have provided in navigating the Ph.D. process. I am thankful to all of the professors and faculty in the Computational Mathematical and Statistical Sciences Department for the help they have provided. I am appreciative of the wonderful professors at the University of Cincinnati, in particular Dr. Michael Sokoloff, who have helped me to build a foundation of knowledge that has allowed me to thrive at Marquette.

I am grateful to have learned from former lab members Dr. Chase Sakitis and Dr. Ke Xu, who have gone on to find great success. I am also happy to have received the help of current lab members Ozy McEvoy and Dean Lang. I would like to thank the many friends that I have made who have provided endless support during my studies. In particular: Maya, Joey, Emily, Josh, Mahi, and Aaron.

TABLE OF CONTENTS

LIST OF TABLES	iv
LIST OF FIGURES	v
CHAPTER 1: SIMULATION AND HARMONIC ANALYSIS OF <i>K</i> -SPACE READ-OUT (SHAKER)	
1.1 Introduction	1
1.1.1 Nuclear Magnetic Resonance	2
1.1.2 <i>k</i> -Space and the Fourier Transform	5
1.2 Design	7
1.2.1 Digital Phantom	8
1.2.2 MRI Parameters	10
1.2.3 fMRI Experimental Design and Noise	13
1.2.4 Time Series Data Analysis	18
1.3 Example fMRI Experiment	19
1.4 Discussion and Future Directions	22
CHAPTER 2: A BAYESIAN MODEL TO ENHANCE COMPLEX-VALUED MEASUREMENTS APPLIED TO FUNCTIONAL MAGNETIC RESONANCE IMAGING	
2.1 Introduction	24
2.2 Bayesian Approach	24
2.2.1 Likelihoods	24
2.2.2 Priors	26
2.2.3 Posteriors	27
2.2.4 Estimation	28
2.3 Illustrative Example	30
2.4 fMRI Application	32
2.4.1 Simulated and Experimental Data	35
2.5 Discussion and Future Directions	43
CHAPTER 3: A MATHEMATICALLY ACCURATE ANGULAR REGRESSION MODEL OPTIMIZES PHASE ACTIVATION AND YIELDS ADDITIONAL PHYSIOLOGICAL INFORMATION IN FMRI	
3.1 Introduction	46
3.2 Methods	49
3.2.1 Distributions	49
3.2.2 Estimation	50
3.3 Results	52
3.3.1 Simulated Results	52
3.3.2 Experimental Results	55
3.4 Connection to Chapter 2	61
3.5 Discussion and Future Directions	63

CHAPTER 4: BAYESIAN ESTIMATION AND RECONSTRUCTION OF NON-CARTESIAN K-SPACE FOR FMRI	66
4.1 Introduction	66
4.2 Radial Encoding of k -Space	67
4.3 Subsampling of k -Space	68
4.4 Bayesian Methodology	71
4.5 Simulated Results	72
4.6 Discussion and Future Directions	74
BIBLIOGRAPHY	75
APPENDIX A: SHAKER VALUES AND DATA STRUCTURES	85
A.1 SHAKER Default Values	85
A.2 SHAKER Structure and Objects	85
APPENDIX B: FMRI TIME SERIES PHASE CORRECTION AND FURTHER USES	88
B.1 Phase Correction in fMRI Time Series	88
B.2 T_1 Map Estimation	89
APPENDIX C: A FURTHER LOOK INTO PHASE-ONLY ACTIVATION	90
C.1 Power analysis of alternative distributions	90
C.2 Grid search approach to parameter estimation	92
C.3 Experimental results for all seven slices	95

LIST OF TABLES

LIST OF FIGURES

1.1	The process of obtaining an image from the machine. a) MRI machine with main axes indicated; the z -direction is referred to as longitudinal, the xy -plane is referred to as the transverse plane. b) Acquired k -space array of spatial frequencies. c) Inverse discrete Fourier transform of k -space; a reconstructed image.	1
1.2	Depicted in (a) is the net magnetization and precession of the magnetic moment about the central axis. Representations of the T_1 (b) and T_2^* (c) relaxivities. T_1 is relaxation back into the longitudinal direction, T_2 is relaxation in the transverse direction.	4
1.3	Three examples of k -space trajectory. a) Standard Cartesian encoding. b) Radial encoding. c) Spiral encoding.	6
1.4	A screenshot of the working SHAKER GUI. In the top left pane, users can view the raw phantom data that will be input to the signal equation. In the bottom left pane are the adjustable MRI parameters and the fMRI experimental setup. The right pane displays two views: the top presents unaltered data from the simulated time series, and the bottom reflects an example of a statistical data set created from the time series data (in this case: a histogram of one voxel's magnitude time series).	8
1.5	Maps taken from an axial slice of the digital phantom. a) Net Magnetization, M_0 (dimensionless). b) Longitudinal relaxation, T_1 (seconds). c) Transverse relaxation, T_2^* (seconds). d) Field inhomogeneity, ΔB (Tesla).	9
1.6	<i>Data Viewer</i> : the top left pane of SHAKER. Here the phantom size can be selected, each of the four maps can be viewed, and slice orientation/location can be selected.	10
1.7	<i>Data Viewer</i> : the top left pane of SHAKER. Here the phantom size can be selected, each of the four maps can be viewed, and slice orientation/location can be selected.	10
1.8	<i>MRI Parameters</i> : the first tab of the bottom left pane of SHAKER. Here relevant MRI parameters can be set. A single k -space can be simulated and observed to check that settings are correct before simulating an entire time series.	11
1.9	Coil sensitivities (top, right, bottom, left) and sensitivity weighted images (corners) for the case of $n_c = 4$. The center image is the true, unweighted slice image. The red outlines indicate the location of the slice as seen by each of the coil sensitivities.	13
1.10	<i>Task Design</i> : the second tab of the bottom left pane of SHAKER. This is where users may adjust the experimental design of the time series, as well as specify <i>SNR</i> and <i>CNR</i>	14
1.11	<i>Time Series</i> : the top right pane of SHAKER. Here the simulated time series can be observed and checked for expected results.	18
1.12	<i>Statistical Analysis</i> : the bottom right pane of SHAKER. This is where users may analyze statistical properties from the data, such as: t statistic for activation, <i>SNR</i> maps, voxel histograms, etc.	20

1.13	Screenshot of the input to SHAKER for the example simulation described in Section 1.3. Not pictured: the fMRI time series and SNR/CNR options set in the second tab of the bottom left panel.	21
2.1	Magnitude and phase of the simulated complex-valued data. The original data has a much lower SNR than the improved MAP and MPM estimates. . .	32
2.2	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.	33
2.3	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.	34
2.4	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.	35
2.5	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.	36
2.6	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.	38
2.7	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.	39
2.8	The magnitude of reconstructed images for the likelihood (experimental data at $TR = 624$), prior distribution, and MAP estimates for all seven slices of the experimental data. Followed by the SNR and CNR for all seven slices. Color bars are consistent across slices. Experimental images used throughout this paper come from slice 6 data.	40
2.9	The SNR for the original simulated 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. . .	41
2.10	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.	41
2.11	T statistics for activation related changes in magnitude (top) and phase (bottom) signals for the slice 6 original data, MAP and MPM estimates. The highlighted region, the left motor cortex, corresponds to the expected area of activation when a right-hand finger tapping experiment is conducted.	43
2.12	T statistics for activation related changes in magnitude and phase signals for the MAP estimates of each axial slice. Activation in the left motor cortex is most noticeable in slices 6 and 7.	44

3.1	Possible signal changes for a voxel as a results of task-related activation. In red is the baseline signal (r, ϕ) , in blue is the task-active signal (r', ϕ') , and in purple is the task-related signal change $(\delta r, \delta \phi)$	48
3.2	Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.7 with resulting SNR (c). Average angular measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the within brain banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.	53
3.3	Simulated data estimations from the von Mises approximation for $\hat{\theta}_0$ (a), $\hat{\theta}_1$ (b), and $\hat{\sigma}^2$ (c) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$. Note the high scale of variance in space compared to the Normal distribution in Figure 3.2 and later for the Lathi distribution in Figure 3.4.	54
3.4	Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta = \theta_0$ with anatomical image (c). Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$	55
3.5	Z-statistics from each of the three distributions, thresholded by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The results from the Normal approximation (a), von Mises approximation (b), and mathematically correct Lathi distribution (c). The Lathi distribution was most successful in determining task-active voxels.	56
3.6	Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.7 with resulting SNR (c). Average angular measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the within brain banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.	57
3.7	Experimental data estimations from the von Mises approximation for $\hat{\theta}_0$ (a), $\hat{\theta}_1$ (b), and $\hat{\sigma}^2$ (c) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$. Note the high scale of variance in space compared to the Normal distribution in Figure 3.6 and later for the Lathi distribution in Figure 3.8.	58
3.8	Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta = \theta_0$ with anatomical image (c). Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$	58
3.9	Z-statistics from each of the three distributions, thresholded by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The results from the Normal approximation (a), von Mises approximation (b), and mathematically correct Lathi distribution (c). The Lathi distribution was most successful in determining task-active voxels.	59
3.10	Experimental data time series (center) of the magnitude (blue) and phase (red) of voxels located in the left motor cortex (top), right motor cortex (bottom), and posterior midline (bottom). The gray shaded regions in the time series plots indicate TRs of subject task activation. Left and right are the significant z-statistics for magnitude- and phase-only activation, respectively.	60
3.11	Experimental data log Fourier coefficients (center) of the magnitude (blue) and phase (red) of voxels located in the left motor cortex (top), right motor cortex (middle), and posterior midline (bottom). The arrow in the top right plot indicates a large Fourier coefficient at a frequency of $19/621 \text{ s}^{-1}$ or about 0.03 Hz. Left and right are the significant F-statistics for the magnitude- and phase-only Fourier coefficients, respectively.	61

3.12	Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.8 with resulting SNR (c) for the MAP estimates of the experimental data. Average measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.	63
3.13	Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta_t = \theta_0$ with anatomical image (c) for the MAP estimates of the experimental data. Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta_t = \theta_0 + \theta_1 x_t$	64
3.14	Z-statistics from the likelihood ratio test statistic described in Equation 3.11 for the MAP estimates of the experimental data, filtered by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The critical z-statistic was found to be 4.18, as indicated by the black lines on the color bar. The left motor cortex is enhanced to the left. This is the region of expected activation for right hand finger tapping experiments.	65
4.1	An example of the gradient echo pulse sequence that would result in a Cartesian trajectory in k -space.	67
4.2	Cartesian sampling (a) is the most common method of traversing k -space. Radial encoding (b) of k -space has benefits, in particular, every spoke passes through the center of k -space.	68
4.3	A fully scanned Cartesian k -space (a) and an example of a subsampled Cartesian k -space with an acceleration factor of $n_a = 2$. The resulting reconstructed image (c) suffers from an aliasing artifact.	69
4.4	A fully scanned (151 spokes) radial k -space (a) and a subsampled (48 spokes) radial k -space resulting in an acceleration factor n_a just over 3. The original reconstructed image (c) compared to the reconstructed image from the subsampled k -space (d). The most noticeable effect of using fewer spokes in reconstruction is a blurring within the brain.	70
4.5	The projection slice theorem in two dimensions states that the Fourier transform of a projection of a 2D object into 1D is equivalent to spoke of the 2D Fourier transform of the object that is orthogonal to the projection.	71
4.6	The proposed method for Bayesian estimation of radial k -space. A fully sampled k -space for calibration (a) and a subsampled k -space during experimentation (b). Estimation of missing spokes (c) and enhancement of collected spokes (d).	72
4.7	Preliminary simulated data from SHAKER. Measured k -space arrays with 128 samples per spoke, and (a) 202 spokes, (b) 101 spokes, (c) 51 spokes, and (d) 26 spokes.	73
4.8	Preliminary simulated data from SHAKER. The M_0 map of the slice used from the digital phantom (a) next to reconstructed magnitude images using (b) 202 spokes, (c) 101 spokes, (d) 51 spokes, and (e) 26 spokes.	73
B1	A depiction of phase drift in k -space measurements over time. For a given spatial frequency coefficient in steady state images the magnitude will remain approximately the same but the phase will drift over time due to changing magnetic fields.	88
B2	A single voxel’s time series through the stages of phase-drift correction.	89

B3	Left is the digital phantom T_1 map. To the right are the estimations of the T_1 maps using the first image of an fMRI time series in addition to steady-state images at TR s 6-16 for both a simulated and experimental time series. The units are in seconds and estimation is consistent with known T_1 values.	90
C1	Theoretical distributions based on maximum likelihood estimates from an active voxel in Figure 3.9. The null hypothesis (black) assumes the Lathi distribution in Equation 3.6. The alternative hypothesis using the Lathi (red), Normal (blue), and von Mises (green) distributions. The power is the highlighted area to the right of the critical value for each alternative hypothesis distribution.	91
C2	Power curves for Lathi (red), Normal (blue), and von Mises (green) alternative hypotheses under a null hypothesis that assumes the Lathi distribution in Equation 3.6. The vertical line at about 0.18 indicates where the maximum likelihood estimate $\hat{\theta}_1$ was located for this particular active voxel. The inset image is a zoomed in region of interest to highlight the difference in statistical power of the three distributions.	92
C3	Grid search results for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta_t = \theta_0$ with anatomical image (c). Grid search results for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta_t = \theta_0 + \theta_1 x_t$	94
C4	Z -statistics from the likelihood ratio test statistic described in Equation 3.11, filtered by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The critical z -statistic was found to be 4.27, as indicated by the black lines on the color bar. The left motor cortex is enhanced to the left. This is the region of expected activation for right hand finger tapping experiments.	94
C5	Summary images for all seven experimental slices (rows). The columns indicate magnitude estimation $\hat{\rho}$, baseline phase $\hat{\theta}_0$, task-related phase change $\hat{\theta}_1$, and variance $\hat{\sigma}^2$ under the alternative hypothesis, and FDR=0.05 filtered z -statistics. Slice 6 was explored further in Section 3.3.2.	97
C6	Summary images for MAP estimates of all seven experimental slices (rows). The columns indicate magnitude estimation $\hat{\rho}$, baseline phase $\hat{\theta}_0$, task-related phase change $\hat{\theta}_1$, and variance $\hat{\sigma}^2$ under the alternative hypothesis, and $\alpha = 0.05$ thresholded z -statistics. Slice 6 was explored further in Section 3.4.	98
C7	The left column is the average phase $\bar{\phi}$ in each voxel for each of the seven experimental slices. The two columns to the left are the linear regression estimates $\hat{\gamma}_0$ and $\hat{\gamma}_1$ that are formed under the Normal assumption. Each slice suffers from phase wrap-around banding artifacts that leads to poor estimation.	99
C8	Resulting statistics for each of the seven experimental slices. The leftmost column is the likelihood ratio statistic, λ described in Section 3.2.2. This is followed by the the statistic λ which is known to have a χ^2 distribution with one degree of freedom. Lastly is the unfiltered z -statistic given by Equation 3.11.	100

CHAPTER 1: SIMULATION AND HARMONIC ANALYSIS OF K -SPACE READOUT (SHAKER)

1.1 Introduction

Functional magnetic resonance imaging (fMRI) is a non-invasive imaging technique that allows physicians and scientists to observe functionality of organs, in particular- the human brain. This is done by exciting protons in the various molecules that make up the different tissues of the organ, then determining a net change in magnetization as determined by an induced current in a loop of wire surrounding the patient. This net magnetization in different voxels of the region of interest (ROI) is associated with complex-valued spatial frequencies that fill an array of k -space; a high order approximation of the Fourier transform of the voxel image of the organ. The k -space array is then inverse discrete Fourier transformed (IDFT) to reconstruct an image. Figure 1.1a shows a simple MRI machine and the major axes, Figure 1.1b depicts an example of the magnitude of a measured complex-valued k -space array, and Figure 1.1c shows the magnitude of the complex-valued image reconstructed from complex-valued k -space, the magnitude of which is in Figure 1.1b.

To perform experiments in the machine is both financially and temporally costly; demanding machine time and obtaining Institutional Review Board (IRB) approvals can slow down the process of investigating new statistical techniques to extract information from




Figure 1.1: The process of obtaining an image from the machine. a) MRI machine with main axes indicated; the z -direction is referred to as longitudinal, the xy -plane is referred to as the transverse plane. b) Acquired k -space array of spatial frequencies. c) Inverse discrete Fourier transform of k -space; a reconstructed image.

fMRI data. Consequently, researchers will test their methods on simulated data as a cost-effective way of assessing potential. Currently, simulated fMRI data are largely developed in-house for each researcher using a variety of methods. There have been studies to develop a more standardized method to simulate fMRI time series data using various languages such as Python and R. However, many such methods disregard the complex-valued nature along with the true statistical and physical properties of the data output by the machine, in addition to returning magnitude-only images from simulations (Welvaert et al., 2011). Some of these methods may also demand some form of in-line coding or require external files to support simulation (Comby et al., 2024). It will be beneficial to provide a complete software tool to researchers that allows the simulation of complex-valued fMRI time series data with the ability to tune various parameters relating to the scan to match future experimental data, that will allow for proper testing of developing models. We present the current work on such an fMRI simulation software tool titled *Simulation and Harmonic Analysis of k-Space Readout* (SHAKER). SHAKER, a GUI simulator built on the MATLAB programming language, is built on the physics-based principles of the MRI machine and is designed so that both new and well-versed researchers in the field can simulate data with ease. The sections that follow will give a brief overview of the physics being applied in the simulator, followed by an in-depth description of each of the parts of the simulator. This will be examined in an example simulation study at the end.

1.1.1 Nuclear Magnetic Resonance

A primary aim of SHAKER is to provide a realistic simulation of fMRI measurements. It is important to understand the physical principles and phenomena that determine the measured signal which is later reconstructed into an image. The MR machine creates a very strong magnetic field B_0 along the z direction of the scanner as indicated in Figure 1.1a (1.5, 3, and 7 Tesla are common). This field aligns the spins of hydrogen nuclei within the body to become parallel or anti-parallel with the direction of the scanner. A slightly higher number of spins align parallel to the field, resulting in a net magnetization M_0 which is measured in imaging. These hydrogen nuclei precess (resonate) at the Larmor frequency

which is proportional to the external magnetic field they are exposed to,

$$f_0 = \gamma B_0, \quad (1.1)$$

where γ is the gyromagnetic ratio, a constant unique to each nucleus (Larmor, 1897). In the case of hydrogen, we have $\gamma = 42.58 \text{ MHz/T}$. To excite these nuclei, a radio frequency (RF) burst of energy is sent into the system at this resonant frequency. The nuclei enter a higher energy state where their spins tip with respect to the direction of the main magnetic field B_0 at some flip angle α determined by the length of the RF pulse. An $\alpha = 90^\circ$ or less flip angle is common for fMRI. In the time that follows the RF pulse, these nuclei emit energy through two relaxation processes- T_1 and T_2 . The longitudinal or spin-lattice relaxation time, T_1 , is the recovery time for the parallel component of M_0 , M_Z , back to equilibrium. The transverse or spin-spin relaxation time, T_2 , is the decay of M_{XY} , the transverse component of M_0 to zero. In fMRI experiments a gradient echo pulse sequence is used, where T_2^* is measured. The relationship between T_2 and T_2^* is defined by

$$1/T_2^* = 1/T_2 + 1/T_2', \quad (1.2)$$

where $1/T_2' = \gamma \Delta B$ is the dephasing of the hydrogen nuclei as a result from hydrogen nuclei precessing at slightly different frequencies due to inhomogeneities in the magnetic field, ΔB . The two effects, T_1 and T_2^* are visualized in Figure 1.2b-c. Figure 1.2a shows the net magnetization change, a vector sum of the T_1 and T_2^* relaxivities. These relaxivities result in a changing magnetic field within the tissue that is measured through current via Faraday's law of induction in one or more coils of wire that are placed locally to the region of interest. This measured signal is then later transformed into complex-valued images via the inverse discrete Fourier transform.

In fMRI, the blood-oxygen-level-dependent (BOLD) signal is interrogated to determine regions of activation (Ogawa et al., 1990; Bandettini et al., 1992; Kwong et al., 1992). The BOLD signal is a measure of localized brain blood level and oxygenation changes which are correlates for neuronal activity. These changes occur as a result of certain stimuli or

tasks, e.g., right-hand finger tapping, that activate known regions of the brain. The BOLD signal presents itself as a T_2^* effect since the change in magnetic properties of oxygenated and deoxygenated hemoglobin in blood causes a perturbation in the local magnetic field, ΔB . Hence, fMRI time series are T_2^* -weighted images, highlighting regions of the brain with significant T_2^* effects.

Every image from an MRI machine comes from a predetermined “pulse sequence” of RF bursts and changing of magnetic gradients within the machine. In fMRI, images are most often collected via single shot echo planar imaging (EPI); “single shot” meaning only one RF excitation is applied per k -space array. The most commonly used pulse sequence used in EPI is gradient echo (GRE) (Kumar et al., 1975; Bernstein et al., 2004). In general, a given signal equation yields complex signal s received at a given point (k_x, k_y) in k -space. The GRE signal equation is given by

$$s(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{M_0 \sin(\alpha)}{(1 - \cos(\alpha)e^{-TR/T_1})} \left(1 - e^{-TR/T_1}\right) e^{-t/T_2^*} e^{i\gamma\Delta B t} e^{-i2\pi(k_x x + k_y y)} dx dy, \quad (1.3)$$

where $M_0(x, y)$, $T_1(x, y)$, $T_2^*(x, y)$, and $\Delta B(x, y)$ are functions of locations (x, y) within the physical object (or phantom) and $t(k_x, k_y)$ is the time at which the point (k_x, k_y) in k -space is scanned. The simplification of replacing $t(k_x, k_y)$ with TE , echo time, is often used and is equivalent to assuming that all data are acquired simultaneously at the TE . SHAKER does not impose this assumption. The repetition time, TR , is the time between successive RF pulses of the same slice, or equivalently, the time between successive measured k -space




Figure 1.2: Depicted in (a) is the net magnetization and precession of the magnetic moment about the central axis. Representations of the T_1 (b) and T_2^* (c) relaxivities. T_1 is relaxation back into the longitudinal direction, T_2 is relaxation in the transverse direction.

arrays of the same slice. The flip angle α is commonly set to 90° , which simplifies the first term in Equation 1.3 to just M_0 . While GRE is most commonly used because of the high BOLD contrast it provides, some higher strength scanners (7 T+) may opt to use the spin echo (SE) pulse sequence to detect BOLD signal (Chen and Glover, 2015; Nencka and Rowe, 2005). The SE signal equation is given by

$$s(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} M_0 \left(1 - e^{-\text{TR}/T_1}\right) e^{-t/T_2} e^{i\gamma\Delta B t} e^{-i2\pi(k_x x + k_y y)} dx dy, \quad (1.4)$$

noting the use of T_2 instead of T_2^* . It has been shown that SE pulse sequences correct for the large scale dephasing caused by larger veins, which may not be as closely related to activation as capillaries (Kida et al., 2000). Closely related to the SE pulse sequence, but not generally used for fMRI experiments, is the inversion recovery (IR) pulse sequence. The signal equation for IR is given by

$$s(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} M_0 \left(1 - 2e^{-\text{TI}/T_1} + e^{-\text{TR}/T_1}\right) e^{i\gamma\Delta B t} e^{-i2\pi(k_x x + k_y y)} dx dy, \quad (1.5)$$

where TI is the inversion time. The IR pulse sequence is more commonly used for T_1 -weighted images, as compared to the T_2^* -weighted image that is standard in fMRI, but is still included in SHAKER. Both the SE and IR signal equations assume a 90° flip angle from the initial RF pulse. Other pulse sequences such as diffusion weighted imaging (DWI-fMRI), and saturation recovery (SR) may be included in the future plans for SHAKER development.

1.1.2 k -Space and the Fourier Transform

The signal equations from Section 1.1.1 are measured in the spatial-frequency domain called k -space (magnitude images of k -space are presented in Figure 1.3). Each of Equations 1.3-1.5 could be condensed to

$$s(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \rho_0(x, y) e^{-i2\pi(k_x x + k_y y)} dx dy. \quad (1.6)$$

In this form, we can see that the signal equation is the Fourier transform of ρ_0 , the net magnetization after having been weighted by the relevant relaxivities. In practice, however,

k -space is only measured at a finite set of discrete points. So, we can discretize Equation 1.6 into

$$s(k_x, k_y) = \frac{1}{N_x N_y} \sum_{m=0}^{N_x-1} \sum_{n=0}^{N_y-1} \rho_0(x, y) e^{-i2\pi \left(\frac{k_x}{N_x} x_m + \frac{k_y}{N_y} y_n \right)}, \quad (1.7)$$

where N_x and N_y are the number of points in image-space in the x and y directions, respectively. In fMRI, it is common that $N_x = N_y = 64, 96, 128$. Thus, we arrive at k -space equating to the *discrete* Fourier transform of image-space.

The objects and phantoms being imaged are composed of real-valued voxels only. So, ideally, the Fourier transform of these objects would result in a k -space that maintains Hermitian symmetry. However due to the terms involving T_2^* and ΔB , the expected Hermitian symmetry of k -space is broken. In fact, when looking at Equations 1.3-1.5, the inclusion of these terms necessarily implies that k -space is only a very close *approximation* to the Fourier transform of image-space, since the terms are time-dependent. This leads to possible distortions and artifacts when reconstructing k -space into images using the inverse discrete Fourier transform

$$\rho_0(x, y) = \frac{1}{N_{k_x} N_{k_y}} \sum_{m=0}^{N_{k_x}-1} \sum_{n=0}^{N_{k_y}-1} s(k_x, k_y) e^{-i2\pi \left(\frac{x}{N_{k_x}} k_{x m} + \frac{y}{N_{k_y}} k_{y n} \right)}, \quad (1.8)$$

however it is still the most common method of image reconstruction in MRI. The incorporation of prior knowledge regarding the relaxivities as well as magnetic field inhomogeneity has been implemented to enhance image reconstruction (Karaman et al., 2015).




Figure 1.3: Three examples of k -space trajectory. a) Standard Cartesian encoding. b) Radial encoding. c) Spiral encoding.

As a result of physical limitations, k -space must be scanned, or traversed, in one continuous path. The most conventional method is to scan horizontal rows, often referred to as the frequency-encoding direction, in alternating directions working up (or down) k -space in the phase-encoding direction as shown in Figure 1.3a. This involves a set of “turnaround points” at the end of each row that are often discarded or not measured, resulting in dead scan time that decreases the rate of useful data acquisition. Modern scanners perform ramp sampling and then interpolate to a Cartesian grid, thus eliminating the dead time in normal EPI acquisitions. This has proven to be a convenient way to scan k -space as it results in a Cartesian encoding of the spatial frequencies which allows for the simple inverse discrete Fourier transform to reconstruct k -space back into an image. Other k -space trajectories, including non-Cartesian methods such as PROPELLER, radial as in Figure 1.3b (equivalent to PROPELLER with blade width 1), and spiral as in Figure 1.3c have been implemented for various reasons such as reducing scan time and increasing robustness to artifacts due to motion (Pipe, 1999; Block et al., 2014).

1.2 Design

SHAKER is an all-inclusive fMRI simulation software package built with the user in mind. SHAKER is built using the MATLAB programming language and presents as a GUI (Figure 1.4), with no scripting or external data required (The MathWorks Inc., 2022). In the top left pane of Figure 1.4, users can view the pre-loaded digital phantom. The bottom left pane involved customization of MRI parameters and fMRI experimental design options. The top right pane presents a view of the simulated time series data. The bottom right pane is where statistical maps and measurements from the time series data can be observed. This is also where any models in development may be tested on the simulated data. Preset values for SHAKER are described in Appendix A.1. All code and data used to operate SHAKER are publicly available on GitHub to encourage a better understanding and allow customizations to be made. The contents of this section explain in detail the functionality of SHAKER in each of these panes.

1.2.1 Digital Phantom

SHAKER comes pre-loaded with a full volume digital phantom that was simulated with realistic M_0 , T_1 , and T_2^* values based on a 3 T machine (Karaman, 2016). The tissues included in the phantom are gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF). The ΔB map was considered as a gradient along each of the dimensions of the scanner, combined with some biological detail from the T_2^* map. This included digital phantom is stored as a MATLAB structure array: `Phantom: M0, T1, T2, deltaB`. Users may also load in their own M_0 , T_1 , T_2^* , ΔB , and activation threshold maps.

An axial slice of the phantom is shown in Figure 1.5. In some cases, a higher or lower sampling density requirement is needed in k -space. To support this, the digital phantom can be rescaled to $64 \times 64 \times 64$ or $128 \times 128 \times 128$ by changing the *Phantom Size* option. Should one want to implement their own digital phantom, it will be necessary to create a




Figure 1.4: A screenshot of the working SHAKER GUI. In the top left pane, users can view the raw phantom data that will be input to the signal equation. In the bottom left pane are the adjustable MRI parameters and the fMRI experimental setup. The right pane displays two views: the top presents unaltered data from the simulated time series, and the bottom reflects an example of a statistical data set created from the time series data (in this case: a histogram of one voxel's magnitude time series).

structure with the same naming convention, having four maps whose dimensions all agree. This custom phantom can then be imported as a `.mat` file from the toolbar located at the top of the GUI. Additionally, there is an activation map included with SHAKER that carries the same dimension as the phantom. The activation map is a binary array, with ones only at the intended location(s) of simulated activation. It is designed such that it roughly resembles the left primary motor cortex region of the brain- the area that is expected to be active during right-hand finger tapping. Custom activation maps can also be imported as `.mat` files with type `double` and name `ActMap` and should share dimensions with the phantom being used for simulation.

At present, SHAKER is equipped to handle single-slice excitations in any of the three major planes: axial, sagittal, or coronal. Support for echo-volume imaging could be supported in future versions. Slice selection and orientation are both chosen and viewed in the top left pane of the GUI, titled *Data Viewer* as in Figure 1.6. The size of the phantom is also adjusted from this pane. Two other options that have no effect on the simulation: viewing each of the maps and visualizing where the activation is expected, are available from this panel as well. Making use of the recently developed “volshow” function in MATLAB, users can get a 3-D view of the four maps that make up the phantom, sliced at the indicated location and orientation. A 3-D view of the M_0 map from the digital phantom at the same slice as in Figure 1.6 is depicted in Figure 1.7.




Figure 1.5: Maps taken from an axial slice of the digital phantom. a) Net Magnetization, M_0 (dimensionless). b) Longitudinal relaxation, T_1 (seconds). c) Transverse relaxation, T_2^* (seconds). d) Field inhomogeneity, ΔB (Tesla).




Figure 1.6: *Data Viewer*: the top left pane of SHAKER. Here the phantom size can be selected, each of the four maps can be viewed, and slice orientation/location can be selected.




Figure 1.7: *Data Viewer*: the top left pane of SHAKER. Here the phantom size can be selected, each of the four maps can be viewed, and slice orientation/location can be selected.

1.2.2 MRI Parameters

Found in the bottom left pane of the GUI are tabs for *MRI Parameters* and *Task Design* as shown in Figures 1.8 and 1.10. *MRI Parameters* is the part of SHAKER where users will make selections similar to that of an MRI technician. *Task Design* will be discussed in Section 1.2.3. All settings in this pane should be confirmed before initializing any simulations. From the top toolbar of SHAKER, the MRI structure being used for simulation can be saved into a `.mat` file which contains a MATLAB structure array named `MRI`. This file

may be imported to future instances of SHAKER for ease of reproducibility of results.

The first two options for MRI parameters are *Signal Equations* and *Trajectory*. These refer to a choice of signal equation as described in Section 1.1.1 and a k -space trajectory from Section 1.1.2. The k -space trajectory functions are designed to receive the **MRI** object inherent to SHAKER that contains all pertinent information regarding the scanner properties. The trajectory functions then return three arrays: one each for the k_x and k_y locations at which k -space is sampled as well as an array noting the time at which the points are sampled $t(k_x, k_y)$. These arrays are stored in an object within SHAKER for later reference. Following this, the user can select their choice of signal equation function, which receives input data about the phantom, k -space sampling, and the **MRI** object, then return a simulated array of k -space measurements. Both the signal equation and k -space trajectory are two files that users can create their own version of, using the templates provided by SHAKER,

MRI Parameters | Task Design

Signal Equation: GradientEcho_...
 Trajectory: Cartesian_ksp...
 Acceleration Factor: 1
 Field Strength (T): 3
 TE (ms): 50
 TR (ms): 1000
 Flip Angle (deg): 90

	M0	T1	T2*
WM	0.7100	832	60
GM	0.8300	1331	60
CSF	1	4000	2200

EESP (ms): 0.72
 Number of Coils: 1
 Reconstruction: CartesianIFFT.m
☒ Include B0 Inhomogeneity

Generate Single kSpace

Domain: kSpace
 Voxel Data: Magnitude
 Coil Number: 1
☐ Show Activated Slice

Export Single k-Space

Figure 1.8: *MRI Parameters*: the first tab of the bottom left pane of SHAKER. Here relevant MRI parameters can be set. A single k -space can be simulated and observed to check that settings are correct before simulating an entire time series.

to sample k -space in their own preferred way. Further details regarding expected inputs and outputs of SHAKER functions can be found in Appendix A.2.

The next option is *Acceleration Factor*. This can mean different things depending on the context of the k -space sampling method. For example, in Cartesian trajectories of k -space an acceleration factor of n_a is often implemented as a measurement of every n_a lines in the frequency-encode direction. In the single-spoke radial trajectory of k -space, this is commonly the measurement of every n_a spokes. There is no restriction on how this might be implemented in one's own k -space trajectory file. Since SHAKER currently supports single-slice imaging, this acceleration factor should be interpreted as an *in-plane acceleration* (IPA) as done in SENSE, as compared to GRAPPA acceleration (Pruessmann et al., 1999; Griswold et al., 2002). Following this is the choice to change the simulated magnetic field strength. The values found in the table located at the top right of Figure 1.8 can be altered to produce an effect on other k -space features at the will of the user. The four options that follow, *TE*, *TR*, *Flip Angle*, and *EESP* are direct inputs to the signal equation as described in Section 1.1.1.

Parallel imaging in fMRI has received a lot of attention recently due to its ability to accelerate the rate at which images are acquired in fMRI experiments (Pruessmann et al., 1999; Griswold et al., 2002). In practice, each coil measures a sensitivity-weighted image of the brain, or phantom, at no additional temporal cost. SHAKER supports the use of a single, uniform coil, or multiple coils aligned with the bore of the machine. Users may specify any number of coils to simulate their data by changing the value for the *Number of Coils* option. The n_c coils have sensitivity matrices that match the dimensions of the phantom. In SHAKER, each of the n_c coil sensitivities is constructed by placing a point at a each coil location, all of which are equidistant from the center of image-space and angularly equidistant from each other. The sensitivity of each coil array then decreases proportional to the inverse of the distance from this point. An example of the simulated coils, coil sensitivity weighted images, and averaged image is shown in Figure 1.9.

The last option is the choice of *Reconstruction Algorithm*. While SHAKER is a k -space simulation tool, the reconstruction of images is supported for the more common k -space trajectories. Similar to the k -space trajectory and signal equation, the choice of reconstruc-




Figure 1.9: Coil sensitivities (top, right, bottom, left) and sensitivity weighted images (corners) for the case of $n_c = 4$. The center image is the true, unweighted slice image. The red outlines indicate the location of the slice as seen by each of the coil sensitivities.

tion algorithm can be user-created based on an included template. These templates can be found in the subdirectories for each of the respective steps in the simulation process. More detail on this can be found in Appendix A.2.

1.2.3 FMRI Experimental Design and Noise

Task-based fMRI generally starts with an initial set of rest images that allow the tissue to reach a steady state in the magnetic field (Elster et al., 2001). Following this, many epochs of rest / task images are taken. For example, an experiment may include: 16 initial rest images followed by 19 epochs of 16 task images followed by 16 rest images for a total of 624 images. It is often the case that some or all of the initial rest images are discarded for fMRI analysis due to the fact that they yield a higher signal than the remaining steady state images. In fact many modern scanners simply do not generate these images anymore due to their limited use. Instead, they may apply a few RF pulses without data collection to bring the spins to steady state and only then begin image acquisition. This can be circumvented by increasing




Figure 1.10: *Task Design*: the second tab of the bottom left pane of SHAKER. This is where users may adjust the experimental design of the time series, as well as specify SNR and CNR .

the flip angle for the first few images so that the amount of transverse magnetization excited in each image is approximately the same (Haase et al., 1986). For this example, discarding the first 16 from analysis would give $n_{IMG} = 608$ images in the fMRI time series. It has been shown, however, that the first few images can be used to aid in analysis of the measured fMRI data, e.g., T_1 map estimation. In SHAKER, users may choose a set number of initial rest images, the number of epochs, and number of rest/task images per epoch. This is then stored as a design vector that can be used for later analysis of the simulated time series. This is all done from the second tab of the bottom left pane of SHAKER, as in Figure 1.10.

Complex-valued measurements k_c in fMRI are composed of both a real and imaginary part, $k_c = k_R + ik_I$ where $i = \sqrt{-1}$. The measured magnitude r_k and phase ϕ_k of the spatial frequency coefficient come from the transformation $k_R = r_k \cos(\phi_k)$ and $k_I = r_k \sin(\phi_k)$. To better model the process of the machine, SHAKER adds noise to k -space directly rather than to reconstructed images as is often done. Since the analog-to-digital converters (ADCs) sam-

ple k -space measurements independently, it is understood that the real and imaginary parts of k -space measurements are independently Normally distributed for each spatial frequency. Thus, the joint distribution is given as

$$f(k_R, k_I) = \frac{1}{(2\pi\sigma_k^2)^{1/2}} \exp \left[-\frac{(k_R - \mu_{k,R})^2}{2\sigma_k^2} \right] \frac{1}{(2\pi\sigma_k^2)^{1/2}} \exp \left[-\frac{(k_I - \mu_{k,I})^2}{2\sigma_k^2} \right], \quad (1.9)$$

where $\mu_{k,R}$ and $\mu_{k,I}$ are the *true* real and imaginary components of the spatial frequency (Henkelman, 1985; Lindquist, 2008). The inverse Fourier transform of the real and imaginary components of the noise from k -space into image space will also be normally distributed with a scaled variance. This relationship is given by

$$\sigma^2 = \frac{\sigma_k^2}{n_x n_y}, \quad (1.10)$$

where σ^2 is the variance of the normally distributed real/imaginary noise in image space, σ_k^2 is the variance of the normally distributed real/imaginary noise in k -space, and n_x , n_y are the dimensions of reconstructed image space (Rowe, 2016). This fact reveals that the joint distribution of a voxels real and imaginary parts in image-space can be written similarly to Equation 1.9 as

$$f(y_R, y_I) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(y_R - \mu_R)^2}{2\sigma^2} \right] \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(y_I - \mu_I)^2}{2\sigma^2} \right]. \quad (1.11)$$

The true real and imaginary components, μ_r and μ_I , can be expressed in terms of the true magnitude and phase, ρ and θ , by the transformation $\mu_R = \rho \cos(\theta)$ and $\mu_I = \rho \sin(\theta)$. Since the magnitude of voxels is preferred over the real/imaginary values when looking at an image, we can transform the measured random variables (y_R, y_I) to (r, ϕ) where $y_R = r \cos(\phi)$ and $y_I = r \sin(\phi)$ (Rowe, 2023). Calculating the Jacobian to be $J = r$, this gives the joint distribution

$$f(r, \phi) = \frac{r}{2\pi\sigma^2} \exp \left(-\frac{1}{2\sigma^2} [r^2 + \rho^2 - 2r\rho \cos(\phi - \theta)] \right). \quad (1.12)$$

By integrating out ϕ from Equation 1.12, we get a Ricean marginal distribution for the voxel's magnitude r (Rice, 1944; Gudbjartsson and Patz, 1995; Rowe, 2005a; Adrian et al.,

2013),

$$f(r) = \frac{r}{\sigma^2} \exp \left[-\frac{r^2 + \rho^2}{2\sigma^2} \right] I_0 \left(\frac{r\rho}{\sigma^2} \right). \quad (1.13)$$

Here I_0 is the zeroth order modified Bessel function of the first kind. The mean of the Ricean distribution is $\sigma\sqrt{\pi/2}L_{1/2}(-\rho^2/2\sigma^2)$ where $L_{1/2}$ is a Laguerre polynomial. The variance of the Ricean distribution, denoted as σ_r^2 , has the following relationship with the variance of the real and imaginary components of voxels in image space, σ^2 ,

$$\sigma_r^2 = 2\sigma^2 + \rho^2 - \frac{\pi\sigma^2}{2}L_{1/2}^2(-\rho^2/2\sigma^2). \quad (1.14)$$

The subscript r is used to indicate the observed magnitude. In regions of empty space where the true signal ρ is small, $\rho \approx 0$, this is reduced to the Rayleigh distribution with mean $\sigma\sqrt{\pi/2}$ and variance $\frac{4-\pi}{2}\sigma^2$ (Rayleigh, 1880). In regions of space with high true signal ρ , this becomes the Normal distribution with mean ρ and variance σ^2 . Integrating out the magnitude r from the joint distribution in Equation 1.12 gives the unnamed non-Normal distribution marginal distribution for the phase ϕ ,

$$f(\phi) = \frac{1}{2\pi} \exp \left[-\frac{\rho^2}{2\sigma^2} \right] \left[1 + \frac{\rho}{\sigma} \sqrt{2\pi} \cos(\phi - \theta) \exp \left[\frac{\rho^2 \cos^2(\phi - \theta)}{2\sigma^2} \right] \Phi \left(\frac{\rho \cos(\phi - \theta)}{\sigma} \right) \right], \quad (1.15)$$

where $\Phi(x)$ is the cumulative distribution function of the standard Normal distribution (Lathi, 1983). When the signal ρ is near zero, the phase will be uniformly distributed on $[-\pi, \pi]$ with mean 0 and variance $\frac{\pi^2}{3}$. When the signal ρ becomes large, the distribution of the phase becomes normal with mean θ and variance $\frac{\sigma^2}{\rho^2}$.

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

$$\rho_t = \beta_0 + \beta_1 x_t. \quad (1.16)$$

Here, $\beta_0 \in \mathbf{R}^+$ is the baseline signal which determines the signal-to-noise ratio $SNR = \beta_0/\sigma_r$, and $\beta_1 \in \mathbf{R}$ is the task-related signal change 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; this is equivalent to a block design hemodynamic response function (HRF). At present, SHAKER supports block design HRFs only. Future development considerations include the option to convolve user-defined functions with the block design to create custom HRFs.

When multiple coils are introduced to support parallel imaging, the statistical properties of the data are altered. Considering C receiver coils, the composite magnitude signal is given by $M_C = \sqrt{\sum_{j=1}^C [M_{jR}^2 + M_{jI}^2]}$ where $M_{jR} = \rho_{jC} \cos(\theta_{jC})$ and $M_{jI} = \rho_{jC} \sin(\theta_{jC})$ denote the real and imaginary signals reconstructed from the j^{th} receiver coil. It has been shown that the probability density of this composite magnitude is the non-central Chi distribution, which can be written as:

$$f(M_C) = \frac{\rho_C}{\sigma^2} \left(\frac{M_C}{\rho_C} \right)^C \exp \left(-\frac{(\rho_C^2 + M_C^2)}{2\sigma^2} \right) I_{C-1} \left(\frac{M_C \rho_C}{\sigma^2} \right) \quad (1.17)$$

where ρ_C is the true signal magnitude and I_k is the k^{th} order modified Bessel function (Koay and Basser, 2006). The mean of this non-central Chi distribution is given by

$$\mu_{M_C} = \sigma \beta_C {}_1F_1(-1/2, C, -\rho_C/(2\sigma^2))$$

where $\beta_C = \sqrt{\pi/2} \frac{(2C-1)}{2^{C-1}(C-1)!}$ and ${}_1F_1$ is the confluent hypergeometric function (Koay et al., 2009). The variance is given by

$$\sigma_{M_C} = \xi(\rho_C|\sigma, C)\sigma^2,$$

where ξ is a scaling factor given by

$$\xi(\rho_C|\sigma, C) = 2C + \frac{\rho_C^2}{\sigma^2} - [\beta_C {}_1F_1(-1/2, C, -\rho_C/(2\sigma^2))]^2.$$

When designing an fMRI experiment in SHAKER, users may specify an *SNR* and *CNR* which will be used to calculate the baseline signal β_0 and magnitude signal increase β_1 , respectively¹. In addition to specifying *SNR* and *CNR*, users may choose to specify some

¹While not previously mentioned, now is a good time to acknowledge that SHAKER should be used with caution if studying quantitative MRI, as the final values have been through at least one layer of scaling and

amount of task-related-phase-change (TRPC) in degrees. It has been shown that there is biological information contained in the phase of an image, and as such it may be desirable to consider it in statistical models (Rowe, 2005a; Rowe et al., 2007). In SHAKER, magnitude activation is determined through the *CNR*, and an additional phase angle in the activated areas can be specified. Once all MRI parameters and fMRI options are confirmed, users may select the *Generate Time Series* button found in the *Task Design* tab to initialize the simulation.

1.2.4 Time Series Data Analysis

The top right panel of SHAKER as in Figure 1.11, labeled as *Time Series*, allows for visualization and examination of the simulated time series data. Users may look through individual images in the time series, in either *k*-space or image-space (if reconstructed), viewing the real, imaginary, magnitude, or phase parts of an image. There is also the option to monitor the time series of individual voxels which can be helpful to determine regions of activation and activation structure (magnitude/phase).

Below this, in the bottom right panel of SHAKER as in Figure 1.12, labeled *Statistical Analysis*, is where statistical maps and models of the simulated time series may be assessed. it may be difficult to extract quantitative meaning.




Figure 1.11: *Time Series*: the top right pane of SHAKER. Here the simulated time series can be observed and checked for expected results.

True to the MRI machine, SHAKER supports analysis of complex-valued data. Users may closely examine the real and imaginary part of an image, or magnitude and phase part of the image. This can be done in both k -space and reconstructed image-space. SHAKER comes preloaded with two simple statistical measures: a voxel-wise t test for activation detection and an SNR calculation to evaluate image fidelity (Ardekani and Kanno, 1998). The t test compares the task-active images in a simulated time series to the mean rest image in order to determine some change in magnitude or phase over some ROI. This is what is expected with the BOLD signal due to the increased T_2^* effect. The SNR calculation estimates the SNR of each voxel throughout the time series. This helps determine the quality of k -space trajectory and reconstruction method by highlighting any regions of leakage (signal leaking between previously aliased voxels) or other artifacts. Statistical maps can be superimposed onto an anatomical image of the excited slice for better viewing of the activated regions or other ROIs. Additionally, there is the option to look at the histogram of any voxel’s magnitude/phase/real/imaginary component throughout the time series with theoretical probability density functions (PDF) overlain. This may be used to confirm expected distributions of voxel’s time series as described in Section 1.2.3. As described in Appendix A.2, SHAKER supports the use of custom statistical methods and models to analyze the simulated data. At present it is recommended that for advanced models the data be exported and examined in a more controlled environment.

1.3 Example fMRI Experiment

This section will carry out an example simulated fMRI experiment, detailing all steps taken in SHAKER. Suppose that a right-hand finger tapping experiment was performed on a subject. The 3T machine is set to scan an axial slice in the center of the brain. The data matrix size is 96×96 , with $TE=60.4$ ms, $EESP=0.832$ ms, $TR=1000$ ms, and no acceleration factor ($n_a = 1$). The experimental timing starts with an initial 16s of rest followed by 19 epochs of 16s of task alternating with 16s of rest for a total of $n_{IMG} = 624$ images.

Firstly, the phantom size and slice orientation and number would be set. This is done in the top left pane of SHAKER by setting the options **Phantom Size: 96**, **View: Axial**, and **Slice: 48**. Following this, the MRI parameters should be set. Without further knowl-




Figure 1.12: *Statistical Analysis*: the bottom right pane of SHAKER. This is where users may analyze statistical properties from the data, such as: t statistic for activation, SNR maps, voxel histograms, etc.

edge of the experiment, it may be safe to assume that a GRE pulse sequence is used and measured along the standard Cartesian trajectory. So the options **Signal Equation:** `GradientEcho_SigEq.m` and **Trajectory:** `Cartesian_kspace.m` should be selected. Following this, the acceleration factor, field strength, TE , TR , and $EESP$ can be input directly from the experimental setup data. It can be assumed that the flip angle is $\alpha = 90^\circ$ (this is not always the case experimentally, but, unless other information is known, is a reasonable specification). For simplicity it can also be assumed that the machine is using a single, uniform coil. Since the k -space trajectory is the standard Cartesian path, images can be reconstructed by setting **Reconstruction:** `CartesianIFFT.m`. To better represent the machine, the box for B_0 inhomogeneity may be checked to include the ΔB effect into the simulation.

In the *Task Design* tab next to *MRI Parameters*, the options for the fMRI experimental design can be set. The HRF can be set to block and the four values that follow: initial rest images, epochs, and task/rest images per epoch, can be filled in directly from the experimental setup. The *Plot Design* button can be pressed to visualize and ensure the experimental




Figure 1.13: Screenshot of the input to SHAKER for the example simulation described in Section 1.3. Not pictured: the fMRI time series and SNR/CNR options set in the second tab of the bottom left panel.

timing is setup correctly. To be consistent with empirical data it is recommended that the SNR is set somewhere in the range of $[1, 10]$ and the CNR is set somewhere between $[0.1, 2]$. For this example, the two are set to be $SNR: 5$ and $CNR: 0.5$. There will be no phase activation added in this example, so **Phase: 0**. Once all inputs are confirmed to be correct, the time series is simulated by pressing the *Generate Time Series Data* button.

Once the simulation is complete, a summary of simulation will be displayed and the *Time Series* panel will be populated with data. This summary can be copied and pasted into a paper for publishing. From here, the images can be observed and sorted through to check for any errors in simulation. A time series of a voxel in an active region may be generated to observe any noticeable patterns. The data may then be analyzed using custom or built-in tools. This can be done for task detection, noise analysis, etc. The results for this example are shown in Figure 1.13. The summary of simulation for this example reads:

The following fMRI time series data was simulated on 11-Nov-2024 at 17:07:23.

The simulated time series is of slice 48 from a size 96 phantom in the Axial

plane. The MRI parameters were set to be the following: Acceleration Factor = 1, Field Strength = 3T, TE = 60.4ms, TR = 1000ms, Flip Angle = 90deg, EESP = 0.832ms, and Number of Coils = 1. The data was simulated with the GradientEcho_SigEq.m signal equation using the Cartesian_kspace.m 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. The SNR was set to 5 and the CNR was set to 0.5. There were 0 degrees of phase added to the activation. Images were reconstructed using the CartesianIFFT.m algorithm.

1.4 Discussion and Future Directions

SHAKER is a one-stop shop for fMRI simulation. The GUI-based approach to the simulator allows for specially simulated data to be placed quickly into the hands of researchers as compared to long wait times for experimental data. The physics-based approach entrusts that the simulated data is representative of what a proper MRI machine might theoretically output. The most recent version of SHAKER can be found on GitHub at the following URL: <https://github.com/bodensjc/SHAKER>. Due to the public availability of SHAKER, it can remain in a constant state of development as users contribute ideas and needs for the simulation tool. The default values inserted in SHAKER can be found in Appendix A.1 and a description of the structure of SHAKER objects can be found in Appendix A.2.

Future consideration for this work involve the addition of more MRI features and parameters as well as additional fMRI experimental design components. Control of Field-of-View (FOV) and bandwidth could be helpful for generating zoomed images. More control of the noise generation through temporal variation as well as spatial. The inclusion of intra- and inter-imaging motion for fMRI experiments would help test registration and motion correction algorithms. While acceleration sampling is provided, standard reconstruction methods such as SENSE and GRAPPA can be implemented to facilitate accelerated parallel imaging (Pruessmann et al., 1999; Griswold et al., 2002). Proprietary non-uniform inverse Fourier transforms may also be implemented to facilitate the reconstruction of non-Cartesian based k -space trajectories. Simultaneous multi-slice (SMS) techniques such as

CAIPIRINHA (Breuer et al., 2005; Xu and Rowe, 2025; Karaman, 2016) could be supported in the future as well as other full- or partial-volume imaging methods. Additional support for the processing of archival data can be added, including techniques such as: Nyquist ghost correction, zero-filling, apodization, motion correction, static B_0 correction, etc.

CHAPTER 2: A BAYESIAN MODEL TO ENHANCE COMPLEX-VALUED MEASUREMENTS APPLIED TO FUNCTIONAL MAGNETIC RESONANCE IMAGING

2.1 Introduction

Many real-world applications, in particular digital signal processing, require the 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, arrays of 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 2.4.

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.

2.2 Bayesian Approach

2.2.1 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

$$f(x_R, x_I | \rho, \theta, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(x_R - \rho \cos(\theta))^2}{2\sigma^2} \right] \times \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(x_I - \rho \sin(\theta))^2}{2\sigma^2} \right]. \quad (2.1)$$

A transformation of variables from Cartesian to polar coordinates $(x_R, x_I) \rightarrow (r, \phi)$ using $x_R = r \cos(\phi)$ and $x_I = r \sin(\phi)$, on Equation 3.3 is performed. The Jacobian $J = r$ is calculated, and we arrive at the joint distribution

$$f(r, \phi | \rho, \theta, \sigma^2) = \frac{r}{2\pi\sigma^2} \exp \left[-\frac{1}{2\sigma^2} [r^2 + \rho^2 - 2r\rho \cos(\phi - \theta)] \right]. \quad (2.2)$$

We can integrate out the measured phase ϕ from Equation 2.2 to arrive at the Ricean marginal distribution of the measured magnitude r (Rice, 1944)

$$f(r | \rho, \sigma^2) = \frac{r}{\sigma^2} \exp \left[-\frac{r^2 + \rho^2}{2\sigma^2} \right] I_0 \left(\frac{r\rho}{\sigma^2} \right). \quad (2.3)$$

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 2.3 becomes the Rayleigh distribution with mean $\sigma \sqrt{\frac{\pi}{2}}$ and variance $\frac{4-\pi}{2} \sigma^2$ (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 2.2 gives an unnamed non-Normal distribution marginal distribution of the measured phase ϕ ,

$$f(\phi | \rho, \theta, \sigma^2) = \frac{1}{2\pi} \exp \left[-\frac{\rho^2}{2\sigma^2} \right] \times \left[1 + \frac{\rho}{\sigma} \sqrt{2\pi} \cos(\phi - \theta) \exp \left[\frac{\rho^2 \cos^2(\phi - \theta)}{2\sigma^2} \right] \Phi \left(\frac{\rho \cos(\phi - \theta)}{\sigma} \right) \right], \quad (2.4)$$

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 $\frac{\pi^2}{3}$. When the true signal ρ becomes large, the distribution of the measured phase becomes

normal with mean θ and variance $\frac{\sigma^2}{\rho^2}$. 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 2.2 can be expressed as the product of the marginal distribution of the magnitude and the conditional distribution of the phase given the magnitude r ,

$$f(r, \phi | \rho, \theta, \sigma^2) = f(r | \rho, \sigma^2) f(\phi | r, \rho, \theta, \sigma^2). \quad (2.5)$$

Partitioning the joint distribution $f(r, \phi)$ in Equation 2.2 by the marginal distribution for ϕ in Equation 2.4, we see that the conditional distribution of the phase ϕ given magnitude r is von Mises (Von Mises, 1918)

$$f(\phi | r, \rho, \theta, \sigma^2) = \frac{\exp \left[\frac{r \rho \cos(\phi - \theta)}{\sigma^2} \right]}{2\pi I_0 \left(\frac{r \rho}{\sigma^2} \right)}. \quad (2.6)$$

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

2.2.2 Priors

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

$$f(\rho | \rho_0, \sigma^2, \gamma) = \frac{\rho}{\sigma^2/\gamma} \exp \left[-\frac{\rho^2 + \rho_0^2}{2\sigma^2/\gamma} \right] I_0 \left(\frac{\rho \rho_0}{\sigma^2/\gamma} \right), \quad (2.7)$$

with hyperparameters ρ_0 and γ . Similarly to the marginal distribution for the measured magnitude r as in Equation 2.3, the distribution in Equation 2.7 becomes the Rayleigh distribution when $\rho_0 \sim 0$ and the Normal distribution when ρ_0 is large. The prior distribution of the true phase θ given the magnitude ρ is von Mises,

$$f(\theta | \rho, \sigma^2, \theta_0, \gamma) = \frac{\exp \left[\frac{\rho \rho_0 \cos(\theta - \theta_0)}{\sigma^2/\gamma} \right]}{2\pi I_0 \left(\frac{\rho \rho_0}{\sigma^2/\gamma} \right)}, \quad (2.8)$$

with location hyperparameter θ_0 . The von Mises distribution in Equation 2.8 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 2.7 and the von Mises distribution in Equation 2.8 become vague with a small γ . Lastly, we assess a prior of the variance σ^2 with an inverse gamma distribution

$$f(\sigma^2|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)(\sigma^2)^{\alpha+1}} \exp\left[-\frac{\beta}{\sigma^2}\right], \quad (2.9)$$

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 $\rho|\sigma \sim \text{Rice}(\rho_0, \sigma/\sqrt{\gamma})$, $\theta|\rho, \sigma \sim \text{VM}(\theta_0, \frac{\rho\rho_0}{\sigma^2/\gamma})$, and $\sigma \sim \text{IG}(\alpha, \beta)$.

2.2.3 Posteriors

The likelihoods from Equations 2.3 and 2.6 can be combined with the priors in Equations 2.7-2.9 to obtain the posterior distribution

$$\begin{aligned} f(\rho, \theta, \sigma^2|r, \phi) &\propto f(r|\rho, \sigma^2)f(\phi|r, \rho, \theta, \sigma^2) \\ &\times f(\rho|\sigma^2, \gamma)f(\theta|\rho, \sigma^2, \theta_0, \gamma)f(\sigma^2|\alpha, \beta), \end{aligned} \quad (2.10)$$

$$\begin{aligned} f(\rho, \theta, \sigma^2|r, \phi) &\propto \frac{r}{\sigma^2} \exp\left[-\frac{r^2 + \rho^2}{2\sigma^2}\right] I_0\left(\frac{r\rho}{\sigma^2}\right) \frac{\exp\left[\frac{r\rho \cos(\phi - \theta)}{\sigma^2}\right]}{2\pi I_0\left(\frac{r\rho}{\sigma^2}\right)} \\ &\times \frac{\rho}{\sigma^2} \exp\left[-\frac{\rho^2 + \rho_0^2}{2\sigma^2/\gamma}\right] I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right) \frac{\exp\left[\frac{\rho\rho_0 \cos(\theta - \theta_0)}{\sigma^2/\gamma}\right]}{2\pi I_0\left(\frac{\rho\rho_0}{\sigma^2/\gamma}\right)} \\ &\times \frac{\beta^\alpha}{\Gamma(\alpha)(\sigma^2)^{\alpha+1}} \exp\left[-\frac{\beta}{\sigma^2}\right]. \end{aligned} \quad (2.11)$$

Some simplification leads to the joint posterior distribution

$$\begin{aligned} f(\rho, \theta, \sigma^2|r, \phi) &\propto \frac{\beta^\alpha}{\Gamma(\alpha)} \exp\left[-\frac{1}{2\sigma^2}(\gamma\rho_0^2 + r^2 + 2\beta)\right] \frac{\rho}{(\sigma^2)^{\alpha+3}} \\ &\times \exp\left[-\frac{1}{2\sigma^2}((\gamma + 1)\rho^2 - 2\rho(\rho_0 \cos(\theta - \theta_0) + r \cos(\phi - \theta)))\right]. \end{aligned} \quad (2.12)$$

The posterior conditional distributions can be found by considering only the terms that contain the parameter of interest. The posterior conditional distribution of 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)

$$f(\rho|\theta, \sigma^2, r, \phi) = \frac{2B^{A/2}\rho^{A-1}}{\Psi\left(\frac{A}{2}, \frac{C}{\sqrt{B}}\right)} \exp[-B\rho^2 + C\rho], \quad (2.13)$$

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 and, as its name would suggest, is a modification of the half-Normal distribution with support for three parameters instead of one. The MHN has no known applications at the time of this writing. The MHN is an exponential family of distributions with support on $(0, \infty)$ (Casella and Berger, 2024). Further, it can be shown that the posterior conditional distribution of the phase θ given the magnitude ρ and variance σ^2 is von Mises

$$f(\theta|\rho, \sigma^2, r, \phi) = \frac{\exp[\kappa \cos(\theta - \lambda)]}{2\pi I_0(\kappa)}, \quad (2.14)$$

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

$$f(\sigma^2|\rho, \theta, r, \phi) = \frac{(\beta^*)^{\alpha^*}}{\Gamma(\alpha^*)(\sigma^2)^{\alpha^*+1}} \exp\left[-\frac{\beta^*}{\sigma^2}\right], \quad (2.15)$$

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 $\sigma^2|\rho, \theta \sim IG(\alpha^*, \beta^*)$. To our knowledge, the MHN distribution has recently been developed but not applied.

2.2.4 Estimation

We will partition the data such that the first n_0 high contrast 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 $\rho_0 = \sqrt{\bar{x}_R^2 + \bar{x}_I^2}$. 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 $\theta_0 = \arctan_4(\bar{x}_R/\bar{x}_I)$. 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, 1994). The modes of the conditional distributions in Equations 2.13-2.15 are given by

$$\hat{\rho} = \frac{C + \sqrt{C^2 + 8B(A - 1)}}{4B}, \quad (2.16)$$

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

$$\hat{\theta} = \arctan\left(\frac{b}{a}\right), \quad (2.17)$$

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

$$\hat{\sigma}^2 = \frac{\beta^*}{\alpha^* + 1}, \quad (2.18)$$

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 $\hat{\sigma}^2$. Based on these values, an estimation of the parameter $\hat{\rho}$ is made using Equation 2.16. 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 2.16-2.18. This is performed for each of the parameters until convergence is reached, as shown in Algorithm 1. 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 2.17 does not depend on ρ or σ , the mode of the phase $\hat{\theta}$ can be immediately estimated without the ICM algorithm.

Algorithm 1 ICM Algorithm for L iterations

$\hat{\theta} \leftarrow f(a, b)$	▷ Equation 2.17
$\hat{\rho}_{(1)} \leftarrow 1$	▷ Initialize guess
$\hat{\sigma}^2_{(1)} \leftarrow 1$	▷ Initialize guess
for $l = 2 : L$ do	
$\hat{\rho}_{(l)} \leftarrow f(\hat{\theta}, \hat{\sigma}^2_{(l-1)})$	▷ Equation 2.16
$\hat{\sigma}^2_{(l)} \leftarrow f(\hat{\rho}_{(l)}, \hat{\theta})$	▷ Equation 2.18
end for	
$\hat{\rho} \leftarrow \hat{\rho}_{(L)}$	▷ Final parameter estimate
$\hat{\sigma}^2 \leftarrow \hat{\sigma}^2_{(L)}$	▷ Final parameter estimate

Since we have the posterior conditional distributions in Equations 2.13-2.15, we could also perform a Gibbs sampling to obtain an entire distribution and compute any desired quantity such as marginal posterior means (MPM) or variances (Geman and Geman, 1984; Gelfand and Smith, 1990). An initial estimate for the parameters ρ , θ , and σ^2 is made. Following this, random observations are generated from each of the posterior conditional distributions given by Equations 2.13-2.15 using the initial estimates. This step is repeated for L iterations, each using the latest observation for each parameter. The first B observations are often burned, leaving a final set of $L - B$ observations for each parameter. The observations are then averaged for each parameter, resulting in MPM estimates. This process is outlined in Algorithm 2 below. Random observations 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 observations 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 observations from the inverse gamma distribution for the variance σ^2 are generated directly using MATLAB's built in `gamrnd()` function.

2.3 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

Algorithm 2 Gibbs Sampler for L iterations with $B < L$ burned observations

$\rho_{(1)} \leftarrow 1$	▷ Initialize guess
$\theta_{(1)} \leftarrow 1$	▷ Initialize guess
$\sigma_{(1)}^2 \leftarrow 1$	▷ Initialize guess
for $l = 2 : L$ do	
$\rho_{(l)} \leftarrow f(\theta_{(l-1)}, \sigma_{(l-1)}^2)$	▷ Equation 2.13
$\theta_{(l)} \leftarrow f(\rho_{(l-1)}, \sigma_{(l-1)}^2)$	▷ Equation 2.14
$\sigma_{(l)}^2 \leftarrow f(\rho_{(l-1)}, \theta_{(l-1)})$	▷ Equation 2.15
end for	
$\hat{\rho} \leftarrow \bar{\rho}_{(B:L)}$	▷ Final parameter estimate
$\hat{\theta} \leftarrow \bar{\theta}_{(B:L)}$	▷ Final parameter estimate
$\hat{\sigma}^2 \leftarrow \bar{\sigma}_{(B:L)}^2$	▷ Final parameter estimate

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 2.1. Prior estimates for the magnitude and phase were calculated using the method described in Section 2.2.4, 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, as described in Algorithms 1 and 2, 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 observations generated including a 500 observation burn-in phase, resulting in 4500 observations 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 2.1. Both posterior estimation methods similarly improve the original data with a noticeable suppression of noise.

The original data had an SNR of 5.3 while the MAP and MPM estimates had improved SNRs of 45.1 and 46.7, respectively. This is an improvement by a factor of about 9 for both methods.




Figure 2.1: Magnitude and phase of the simulated complex-valued data. The original data has a much lower SNR than the improved MAP and MPM estimates.

2.4 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 primarily hydrogen nuclei in water molecules in tissues 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 2.2, showing that each element in k -space is composed of a complex-valued spatial frequency coefficient. The k -space array 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.

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. As discussed previously in Section 1.2.3 many modern machines perform some preparation pulses to get the spins to steady state before image acquisition begins. Assuming the experiment starts in thermal equilibrium, the




Figure 2.2: 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.

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 illustration of longitudinal magnetization after the first few RF pulses is shown in Figure 2.3. Since the 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.

When Bayesian analysis of fMRI data is done on image-space data, voxels may be spatially correlated, often due to image processing (Nencka and Rowe, 2007; Karaman et al., 2014). Working with this spatially correlated data demands 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., 2026). Thus, we chose to employ our Bayesian method on the uncorrelated k -space data prior to IFT 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




Figure 2.3: 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.

k -space measurement (k_R, k_I) are independently Normally distributed 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 coordinates, as visualized in Figure 2.2. This experimental setup precisely matches the theory established in Section 2.2.

Previous work on Bayesian image analysis in Fourier space (BIFS) has been done on polar coordinates (Kornak et al., 2026; Bakas et al., 2025). The BIFS method operates in a Fourier space that is reconstructed from the Fourier transform of a noisy real-valued image, resulting in a Hermitian symmetric Fourier space. An MCMC Bayesian method to enhance the complex-valued spatial frequency coefficients is conducted. The result is then inverse Fourier transformed back into some de-noised complex-valued image. In fMRI the complex-valued spatial frequency coefficients in k -space are measured directly and then inverse Fourier transformed into image space. In our problem and formulation we do not assume Hermitian symmetry. The method outlined in Section 2.2 as applied to this problem can be seen as an extension to BIFS that is applied directly to a measured Fourier space (k -space) and makes use of mathematically correct conjugate prior distributions which leads to simpler estimation. All distributions and parameters in this section are analyzed in k -space unless stated otherwise.

2.4.1 Simulated and Experimental Data

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 discrete inverse Fourier transform. These options were chosen to maintain consistency with the experimental data shown later. Figure 2.4 shows the magnitude and phase of the first four k -space arrays and corresponding reconstructed




Figure 2.4: 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.

images in the time series. The remainder of the time series appears similarly to TR 4.

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 low SNR. 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 $\text{FOV} = 24$ cm, echo time $\text{TE} = 50$ ms, effective echo spacing $\text{EESP} = 0.832$ ms, and repetition time $\text{TR} = 1$ 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 changing global 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, 2025b). First, the angular phase temporal mean of each k -space element's time-series is




Figure 2.5: 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.

calculated and angularly subtracted. A second order local polynomial is then 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 is unwrapped and linearly detrended, then the angular phase temporal mean is angularly added back, producing a linearly stable phase over time for each spatial frequency. This is described further in Appendix B.1. Figure 2.5 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.

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 2.2.4 and exemplified in Section 2.3, 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 $\rho_0 = \sqrt{\bar{k}_R^2 + \bar{k}_I^2}$. 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 $\theta_0 = \arctan_4(\bar{k}_R/\bar{k}_I)$. 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 as described by Algorithms 1 and 2 to obtain posterior estimates. Similarly to Section 2.3, 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 observations generated including a 50-observation burn-in phase, resulting in 150 observations for MPM estimation at each time point for each k -space element. For the experimental data, there were 5,000 Gibbs observations generated including a 500-observation burn-in phase, resulting in 4500 observations 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 observations 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 2.6 and 2.7. 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




Figure 2.6: 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 2.7: 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.

extraction, and classification (Zhou et al., 2019). Likelihood, prior, and MAP estimates for all seven experimental slices can be found in Figure 2.8.

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

$$\rho_t = \beta_0 + \beta_1 x_t \quad (2.19)$$

where $\beta_0 \in \mathbf{R}^+$ is the baseline signal which determines the signal-to-noise ratio $SNR = \beta_0/\sigma_r$, and $\beta_1 \in \mathbf{R}$ 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




Figure 2.8: The magnitude of reconstructed images for the likelihood (experimental data at $TR = 624$), prior distribution, and MAP estimates for all seven slices of the experimental data. Followed by the SNR and CNR for all seven slices. Color bars are consistent across slices. Experimental images used throughout this paper come from slice 6 data.

Figures 2.9 and 2.10 (top row). Note that the upper scaling for the color bar is increased from 5 to 30 for the Bayesian posterior estimates highlighting the improved SNR. Both Bayesian methods similarly improved the SNR when compared to the original data. Slight timing




Figure 2.9: The SNR for the original simulated 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.

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. This ghost could be masked out since it exists in what is really empty space. The CNR is an indication of where the increase in task activation-related signal




Figure 2.10: 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.

is measured. The CNR calculated from the original data and the two Bayesian methods is shown in Figures 2.9 and 2.10 (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.

The model supports a complex-valued task activation related blood oxygenation level dependent response (Ogawa et al., 1990; Bandettini et al., 1992). When a voxel is active as a result of the BOLD response, the magnitude can change, the phase can change, or a combination of both can occur. A simple t test is conducted on the experimental images at rest compared to the images collected while the task was active. The results, presented in Figure 2.11, highlight detected activation in the left motor cortex (the region where task activation is expected for the right-hand finger tapping experiment). It is clear that voxels in this region exhibit both an increase in magnitude as well as an associated increase in phase that correlate to task performance. Further, both estimation methods preserve the ability to detect task in the posterior enhanced data.

It is worth noting that while both estimation methods preserve task-related activation, the average values for the z -statistics in the task-active voxels are smaller than the values from the original data. It is believed that this is caused from the prior images being collected while the patient is at rest, thus “drowning out” some of the task-related signal change when combined with the likelihood images. Ideally in an experimental setting there could be two runs of the machine- one where the patient is at rest while the machine reaches thermodynamic equilibrium, and another where the patient is performing the task. This would enable the creation of a “rest” prior and an additional “task” prior. The rest prior

could be used to enhance rest images, and the task prior could be used to enhance images where the task is active. It is believed that this would in fact *enhance* the ability to detect task in steady-state images. Unfortunately, we did not have access to such experimental data at the time of writing.

Figure 2.12 shows the separate task-activation related t-statistics for magnitude and phase that results from the MAP estimates of all seven experimental slices. Slice 6 is the slice used previously for demonstration. Slices 5-7 have the most noticeable active region, in particular slice 6. Gibbs observations were only generated for slice 6 due to the significant computational time required.




Figure 2.11: T statistics for activation related changes in magnitude (top) and phase (bottom) signals for the slice 6 original data, MAP and MPM estimates. The highlighted region, the left motor cortex, corresponds to the expected area of activation when a right-hand finger tapping experiment is conducted.

2.5 Discussion and Future Directions

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 2.3. The results of the simulated and experimental fMRI studies in Section 2.4 demonstrate that the model can increase the SNR of the original data, leading to higher contrast reconstructed images.




Figure 2.12: T statistics for activation related changes in magnitude and phase signals for the MAP estimates of each axial slice. Activation in the left motor cortex is most noticeable in slices 6 and 7.

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 observations; 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. An additional use case of the first several images in an fMRI time series to estimate the T_1 values for voxels is presented in Appendix B.2. If

the first few TRs include a clever use of variation of echo time, it is also possible to make estimations of the static ΔB field as well as an estimated T_2^* map (Hahn and Rowe, 2012).

The ability to detect task in the MAP and MPM reconstructed images is not significantly 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.

Future considerations for this work involve the comparison of hypotheses $H_0 : \mathbf{C}\beta = 0, \mathbf{D}\gamma = 0$ vs. $H_1 : \mathbf{C}\beta \neq 0, \mathbf{D}\gamma \neq 0$ where $\mathbf{C}$ and $\mathbf{D}$ are linear constraint matrices for the magnitude coefficients β and the phase coefficients γ (Rowe, 2005a). It is also worth considering a weighting scheme for the first few images when forming the prior distributions. For example, a given spatial frequency prior value might be given by $\bar{k}_R = \alpha_1 k_{R_1} + \alpha_2 k_{R_2} + \alpha_3 k_{R_3}$ where $\sum \alpha_i = 1$. This would shift even more weighting of the prior distribution into the first, best image. Additionally, as has been previously mentioned, it would be worth exploring the idea of a “rest” versus “task” prior image. The rest prior would be formed from the first few images in a time series when the patient is at rest, and the task prior would be formed from the first few images from a separate time series where the patient is performing the task. These could then be used to separately enhance images during experimentation when the patient is either at rest or performing the task.

CHAPTER 3: A MATHEMATICALLY ACCURATE ANGULAR REGRESSION MODEL OPTIMIZES PHASE ACTIVATION AND YIELDS ADDITIONAL PHYSIOLOGICAL INFORMATION IN FMRI

3.1 Introduction

In functional magnetic resonance imaging (fMRI), images are most often collected via single shot echo planar imaging (EPI)—“single shot” meaning only one radio frequency (RF) excitation is applied per measured k -space array. The most commonly used pulse sequence in EPI for fMRI is gradient recalled echo (GRE) (Kumar et al., 1975; Bernstein et al., 2004). In general, a given signal equation yields a complex-valued signal s received at a given point (k_x, k_y) in k -space. The steady state GRE signal equation is given by

$$s(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{M_0 \sin(\alpha)}{(1 - \cos(\alpha)e^{-TR/T_1})} \times \left(1 - e^{-TR/T_1}\right) e^{-t/T_2^*} e^{i\gamma\Delta B t} e^{-i2\pi(k_x x + k_y y)} dx dy, \quad (3.1)$$

where proton spin density $M_0(x, y)$, longitudinal relaxation time $T_1(x, y)$, observed transverse relaxation time $T_2^*(x, y)$, and $\Delta B(x, y)$ are functions of voxels (x, y) within the physical object being imaged and $t(k_x, k_y)$ is the time at which the point (k_x, k_y) in k -space is scanned (Hargreaves, 2012; Wang et al., 2020). The simplification of replacing $t(k_x, k_y)$ with echo time TE is often used and is equivalent to assuming that all data are acquired simultaneously at the TE . The repetition time TR is the time between successive RF pulses of the same slice, or equivalently the time between successive measured k -space arrays of the same slice. The flip angle α is commonly set to 90° , which simplifies the quotient term in Equation 3.1 to M_0 .

A point of interest in the GRE signal equation is the term $e^{i\gamma\Delta B t}$ where ΔB is a catch-all term for any spatio-temporal inhomogeneity in the magnetic field local to a voxel. These deviations from the intended uniform magnetic field can be caused by scanner imperfections, gradient coil imperfections and eddy currents, tissue composition and interfaces with air, and more (Bernstein et al., 2004; Brown et al., 2014; Holland et al., 2010). Since the exponential term that includes ΔB is the only non-Fourier term that includes an imaginary part, the

local biological information contained within ΔB can be found within the local phase of reconstructed images. It is often the case that the true complex-valued nature of fMRI data is ignored, and only magnitude images are studied.

In fMRI, the blood-oxygen-level-dependent (BOLD) signal is interrogated to determine regions of activation (Ogawa et al., 1990; Bandettini et al., 1992, 1993). The BOLD signal is a measure of localized brain blood volume and oxygenation changes which are correlates for neuronal activity. These changes occur as a result of certain stimuli or tasks, e.g., right-hand finger tapping, that activate known regions of the brain. The BOLD signal presents itself as a T_2^* effect since the change in magnetic properties of oxygenated and deoxygenated hemoglobin in blood causes a perturbation in the local magnetic field, ΔB , which causes a transverse dephasing of spins. Hence, fMRI time series are T_2^* -weighted images, highlighting regions of the brain with significant T_2^* effects. It is known that

$$1/T_2^* = 1/T_2 + 1/T_2' \quad (3.2)$$

where T_2 is the “true” transverse relaxation effect due to intrinsic molecular interactions, and T_2' is transverse relaxation due to magnetic field inhomogeneities ΔB . In fact, $1/T_2' = \gamma\Delta B$, which is located in the non-Fourier complex exponential term in Equation 3.1 (Chavhan et al., 2009). The BOLD signal has a measurable effect in both the magnitude (from the term e^{-t/T_2^*}) and phase (from the term $e^{i\gamma\Delta B t}$) of the signal within a voxel. Therefore one would expect to observe any detectable task-related signal change within a voxel in both the magnitude and phase of the reconstructed voxel time series. Traditionally the BOLD signal is analyzed in magnitude-only images, discarding any information contained within the phase. Figure 3.1 depicts the complex plane in which a voxel measurement exists. In red is the baseline signal (r, ϕ) , in blue is the task-active signal influenced by BOLD effects (r', ϕ') , and in purple is the task-related signal change $(\delta r, \delta \phi)$. It is possible that there is change in one or both of the magnitude and phase (Rowe, 2005a). If the phase data is ignored, only changes in magnitude can be detected, which leads to the loss of physiological information found in the phase.

Historically in fMRI, a general linear model with Normally distributed errors was used




Figure 3.1: Possible signal changes for a voxel as a results of task-related activation. In red is the baseline signal (r, ϕ) , in blue is the task-active signal (r', ϕ') , and in purple is the task-related signal change $(\delta r, \delta \phi)$.

for the marginal distribution of the magnitude component of the time series (Bandettini et al., 1993). However, the Normal assumption is not valid at low signal-to-noise ratios (SNRs). Recently, efforts have been made to use the proper Ricean marginal distribution for the magnitude component in fMRI (Rowe, 2005b; Adrian et al., 2013, 2025). It should be noted that there is also a line of research utilizing the full complex-valued time series (Lai and Glover, 1997; Nan and Nowak, 1999; Rowe and Logan, 2004; Rowe, 2005a, 2009). In this line of research, task related magnitude, phase, or magnitude and phase activation is computed within the full complex-valued time series. It is believed that by using the full $2n$ complex-valued measurements there is increased detection power over using n magnitude measurements even at high SNR (Rowe, 2005b). Further, when there are task related changes in both the magnitude and phase, a model that detects both yields increased detection of task activation (Rowe, 2005a). Ever increasing relaxation of the assumed model parameters was made (Adrian et al., 2018, 2025, 2024) as well as Bayesian inference (Yu et al., 2018, 2023; Wang et al., 2024, 2025). There has been some effort to properly estimate and perform inferences on the phase component of the time series (Rowe et al., 2007). In past work, researchers centered each voxel time series, unwrapped it, and simply fit a linear regression model with Normally distributed errors. However this can have challenges when

there are one or more phase transitions on the $[-\pi, \pi]$ boundary and Type I and II errors when the incorrect distribution is used. In this work, we will use Lathi's mathematically correct phase distribution to compute task-related phase activation (Lathi, 1983). This will be compared to using a Normal distribution and a more accurate angular regression model to make estimations (Rowe et al., 2007).

3.2 Methods

3.2.1 Distributions

It has been well established that MRI voxel values are complex-valued consisting of real and imaginary parts. The real y_R and imaginary y_I parts of a given voxel value at a specified time contain independent additive Normally distributed noise, $y_R \sim N(\rho \cos \theta, \sigma^2)$ and $y_I \sim N(\rho \sin \theta, \sigma^2)$ (den Dekker and Sijbers, 2005; Gudbjartsson and Patz, 1995; Rowe, 2023). The parameters ρ and θ represent the true magnitude and phase signal within the voxel. The joint distribution of the real and imaginary parts can be expressed as a bivariate Normal distribution with phase-coupled means

$$f(y_R, y_I | \rho, \theta, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(y_R - \rho \cos(\theta))^2}{2\sigma^2} \right] \times \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{(y_I - \rho \sin(\theta))^2}{2\sigma^2} \right]. \quad (3.3)$$

A transformation of variables from the Cartesian random variables (y_R, y_I) to the polar random variables (r, ϕ) can be performed, where $y_R = r \cos \phi$, $y_I = r \sin \phi$ and the Jacobian is found to be $J = r$. This results in the following joint distribution for measured magnitude r and phase ϕ

$$f(r, \phi | \rho, \theta, \sigma^2) = \frac{r}{2\pi\sigma^2} \exp \left[-\frac{1}{2\sigma^2} [r^2 + \rho^2 - 2r\rho \cos(\phi - \theta)] \right]. \quad (3.4)$$

From this bivariate distribution, the Ricean marginal distribution for magnitude with location ρ and scale σ can be extracted by integrated out the phase ϕ , yielding

$$f(r | \rho, \sigma^2) = \frac{r}{\sigma^2} \exp \left[-\frac{r^2 + \rho^2}{2\sigma^2} \right] I_0 \left(\frac{r\rho}{\sigma^2} \right) \quad (3.5)$$

where I_0 is the zeroth order modified Bessel function of the first kind (Rice, 1944; Gudbjartsson and Patz, 1995; Rowe and Logan, 2004). The mathematically correct marginal distribution for phase, which will further be referred to as the “Lathi” distribution, is found to be

$$f(\phi|\rho, \theta, \sigma^2) = \frac{1}{2\pi} \exp\left[-\frac{\rho^2}{2\sigma^2}\right] \times \left[1 + \frac{\rho}{\sigma} \sqrt{2\pi} \cos(\phi - \theta) \exp\left[\frac{\rho^2 \cos^2(\phi - \theta)}{2\sigma^2}\right] \Phi\left(\frac{\rho \cos(\phi - \theta)}{\sigma}\right)\right], \quad (3.6)$$

where $\Phi(x)$ is the cumulative distribution function of the standard Normal distribution (Lathi, 1983; Rowe and Logan, 2004; Rowe and Bodenschatz, 2025). This Lathi distribution is often approximated to be the Normal distribution. This approximation is only supported in voxels where the true signal ρ is high relative to the noise σ . The von Mises distribution, a close approximation to the wrapped Normal distribution, has also been used to characterize the phase (Rowe et al., 2007). In regions of empty space where $\rho \approx 0$, the Lathi distribution becomes uniform on the interval $[-\pi, \pi]$. A power analysis of the distribution in Equation 3.6 compared to the Normal and von Mises approximations as well as a visual comparison of the three distributions are explored in C.1. At a fixed Type I error rate of $\alpha = 0.05$, the Lathi distribution provides the largest statistical power over the Normal and von Mises distributions, indicating an increased ability to correctly identify task-active voxels.

3.2.2 Estimation

In fMRI, there is a time series of n complex-valued measurements collected at each voxel location. An association measure between the time series and the expected response from the fMRI experiment is computed, usually in the form of a linear regression that is applied to only the magnitude of the complex-valued time series. Without task, the parameters in Equation 3.5 can be estimated via maximum likelihood estimation. The likelihood function is given by

$$L(\rho, \sigma^2) = \prod_{t=1}^n f(r_t|\rho, \sigma^2), \quad (3.7)$$

which upon insertion of the Rice distribution becomes

$$L(\rho, \sigma^2) = \prod_{t=1}^n \frac{r_t}{\sigma^2} \exp \left[-\frac{r_t^2 + \rho^2}{2\sigma^2} \right] I_0 \left(\frac{r_t \rho}{\sigma^2} \right), \quad (3.8)$$

where $f(r_t|\rho, \sigma^2)$ is the Ricean distribution described in Equation 3.5.

Similarly, without task, the parameters in Equation 3.6 can be estimated via maximum likelihood estimation. The likelihood function is given by

$$L(\rho, \theta, \sigma^2) = \prod_{t=1}^n f(\phi_t|\rho, \theta, \sigma^2), \quad (3.9)$$

with $\theta = \theta_0$ where θ_0 is the baseline phase. In the case where task is present, $\theta = \theta_0 + \theta_1 x_t$ where θ_1 is the additive task-related signal and $x_t \in \{0, 1\}$ is an indicator of expected response at time t .

This can be structured as a hypothesis test of $H_0 : \theta_1 = 0$ versus $H_1 : \theta_1 \neq 0$. The likelihood for the null and alternative hypotheses can be maximized and the parameters estimated via maximum likelihood estimation. This can be done with a multitude of optimization methods. We choose to equivalently minimize the negative log likelihood function using MATLAB's "fmincon" function which employs a sequential quadratic programming method (The MathWorks Inc., 2022). While perhaps more verbose, this alternative method of maximization was deemed necessary due to the insufficient float precision that came with the large numbers resulting from MATLAB's "mle" function. Since it is well known that the task-related magnitude signal change is extremely small compared to the baseline, an aggregated magnitude MLE $\hat{\rho}$ is estimated using MATLAB's "mle" function and used when computing the phase MLEs from Equation 3.9. This also provides an estimate for $\hat{\sigma}^2$ which is updated during the phase estimation. Utilizing the distribution in Equation 3.6 for each of the n observations at time t , the phase likelihood becomes

$$L(\theta, \sigma^2) = \prod_{t=1}^n \left\{ \frac{1}{2\pi} \exp \left(-\frac{\hat{\rho}^2}{2\sigma^2} \right) \left[1 + \frac{\hat{\rho}}{\sigma} \sqrt{2\pi} \cos(\phi_t - \theta) \exp \left(\frac{\hat{\rho}^2 \cos^2(\phi_t - \theta)}{2\sigma^2} \right) \Phi \left(\frac{\hat{\rho} \cos(\phi_t - \theta)}{\sigma} \right) \right] \right\} \quad (3.10)$$

where $\theta = \theta_0$ under the null hypothesis and $\theta = \theta_0 + \theta_1 x_t$ under the alternative hypothesis. A grid search for the parameters the maximum likelihood estimates was also performed and results are consistent with the optimization procedures used.

A voxel-wise likelihood ratio statistic $\lambda = L_{H_0}/L_{H_1}$ can be computed as the ratio of the null likelihood over the alternative likelihood, and the $\Lambda = -2\ln(\lambda) \sim \chi^2(1)$ test statistic formed (Wilks, 1938). Since there is a single degree of freedom, a z -statistic can be formed

$$z = \text{sign}(\hat{\theta}_1) \sqrt{-2\ln(\lambda)} \quad (3.11)$$

and utilized to detect “task-active” voxels (Severini, 2000).

3.3 Results

3.3.1 Simulated Results

This method was tested on both simulated and experimental fMRI data. The simulated data was made using the same MR parameters as the experimental data for consistency of methods and comparability of results. The simulated fMRI time series data was generated using SHAKER v1.1 (Bodenschatz and Rowe, 2025). The simulated time series is of slice 91 from a size 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°, EESP = 0.832 ms, and Number of Coils = 1. The data was simulated with the Gradient Echo signal equation (Equation 3.1) 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. The initial 3 images are omitted from analysis. The SNR was set to be about 5 in the brain and the CNR was set to 0.25 in the region of task-activation. There were 6 degrees of phase added to the region of activation.

Figure 3.2 are voxel-wise estimates that use the Normal distribution ordinary least squares (OLS) and not the Lathi distribution in Equation 3.6. Figures 3.2a and 3.2b depict the MLE estimates for magnitude and variance from Equation 3.7. Figure 3.2c shows the estimated SNR defined by $\widehat{SNR} = \hat{\rho}/\hat{\sigma}$, which was intentionally simulated to be about 5 within the brain. Figure 3.2d depicts the mean angular phase $\bar{\phi}$ of the time series. Figures

3.2e and 3.2f are the simple linear regression coefficient estimates for baseline phase θ_0 and task-related phase θ_1 . The baseline phase coefficient estimate appears reasonable within the brain when compared to the mean angular phase, but poorly estimates the regions in empty space. Similarly, the task related phase coefficient estimate is very unreasonable in empty space, and within the brain there is a very clear banding artifact that comes from the phase wraparound near the $[-\pi, \pi]$ boundary. This artifact can make it very difficult to determine with confidence both the existence of task-related phase change and the true value of the angular difference from baseline phase.

Another phase model described by Rowe (2007) uses the von Mises distribution to characterize the phase. This distribution has support on $[-\pi, \pi]$ and is often compared to the wrapped Normal distribution. The results of applying this model to the simulated data under the alternative hypothesis are shown in Figure 3.3. The baseline phase estimate $\hat{\theta}_0$ (a) is spatially smoother at the wraparound regions, and the estimate for task-related phase




Figure 3.2: Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.7 with resulting SNR (c). Average angular measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the within brain banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.




Figure 3.3: Simulated data estimations from the von Mises approximation for $\hat{\theta}_0$ (a), $\hat{\theta}_1$ (b), and $\hat{\sigma}^2$ (c) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$. Note the high scale of variance in space compared to the Normal distribution in Figure 3.2 and later for the Lathi distribution in Figure 3.4.

change $\hat{\theta}_1$ is free of the banding artifacts that were observed with Normal estimation. The variance of estimation is increased, largely in regions of empty space.

Figure 3.4 shows the maximum likelihood estimates of $\hat{\theta}_0$, $\hat{\sigma}^2$, and $\hat{\theta}_1$ under the null and alternative hypothesis in the top and bottom rows respectively. Contrary to Figure 3.2, these estimations are based on the more accurate Lathi distribution in Equation 3.6 and formed via the maximum likelihood presented in Equation 3.10. In the top row, Figures 3.4a and 3.4b depict the estimates $\hat{\theta}_0$ and $\hat{\sigma}^2$ under the null hypothesis that $\theta = \theta_0$. In the bottom row, Figures 3.4d, 3.4e, and 3.4f depict the estimates $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ under the alternative hypothesis that $\theta = \theta_0 + \theta_1 x_t$. The estimate $\hat{\theta}_0$ is approximately the same for both hypotheses while the estimate for $\hat{\sigma}^2$ appears to be slightly decreased in empty space under the alternative hypothesis. The estimate for $\hat{\theta}_1$ looks nearly uniform at about 0 inside the brain, and noisy in empty space.

While the image of $\hat{\theta}_1$ in Figure 3.4e appears fairly spatially smooth just like with the simulated data in Figure 3.3, closer examination reveals details about the nature of θ_1 . Calculating the likelihood ratio statistic as described in Section 3.2.2 and corresponding z -statistic as described by Equation 3.11, we can highlight exactly where the task-related phase coefficient θ_1 is significant. The z -statistics produced from this method, presented in Figure 3.5, are thresholded using the Benjamini-Hochberg procedure to control the false




Figure 3.4: Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta = \theta_0$ with anatomical image (c). Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$.

discovery rate at the $\alpha = 0.05$ level (Benjamini and Hochberg, 2018; Logan et al., 2008; Logan and Rowe, 2004). The resultant critical z -statistics are found to be ± 3.28 , ± 4.40 , and ± 4.45 for the Normal, von Mises, and Lathi distributions, respectively. Here we can clearly identify a region of task-related phase change. The average value for $\hat{\theta}_1$ in the voxels that were simulated to have an additional 6 degrees of task-related phase change was found to be 5.34 degrees. In each of Figures 3.2f, 3.3b, and 3.4f, the measure for $\hat{\theta}_1$ is a direct measure of the change in magnetic field local to the space that make up each of the voxels.

3.3.2 Experimental Results

Experimental data is from a single subject block design right-hand finger tapping experiment on a 3.0-Tesla General Electric Signa LX MRI scanner. The imaging parameters were $n_z=7$ slices of 2.5 mm thick and 128×128 array size with a field-of-view FOV=24.0 cm, echo time TE=60.4 ms, effective echo spacing EESP=0.832 ms, and time-of-repetition TR=1 s. The experiment timing followed an initial 16 s of rest followed by 19 epochs of 16 s of




Figure 3.5: Z -statistics from each of the three distributions, thresholded by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The results from the Normal approximation (a), von Mises approximation (b), and mathematically correct Lathi distribution (c). The Lathi distribution was most successful in determining task-active voxels.

task alternating with 16 s of rest resulting in a total of $n_t=624$ total image volumes. The initial 3 images are omitted from analysis. Each slice image at each time point, was Nyquist ghost corrected (Nencka et al., 2009). Images were phase drift corrected by subtracting each voxel’s angular phase temporal mean. A local second order polynomial was spatially fit to the resultant difference of each phase image in the time-series. The spatially fitted phase is angularly subtracted from the original time series (Sakitis and Rowe, 2025b). A simple linear regression model was fit to each voxel’s unwrapped phase time series and angularly subtracted off. Then the angular mean was added to each voxel’s time series, resulting in a stable global phase over time. This is described further in Appendix B.1.

Figure 3.6 are voxel-wise estimates that use the Normal distribution OLS and not the Lathi distribution in Equation 3.6; this figure is the experimental equivalent to Figure 3.2. Figures 3.6a and 3.6b depict the MLE estimates for magnitude and variance from Equation 3.7. Figure 3.6c shows the SNR which was found to be about 5 within the brain. These brain images were collected using a body receive coil with decreased SNR, resulting in particularly noisy data. Figure 3.6d depicts the mean angular phase $\bar{\phi}$ of the experimental time series. Figures 3.6e and 3.6f are the simple linear regression coefficient estimates for baseline phase θ_0 and task-related phase θ_1 . Similar to the simulated data, the baseline phase coefficient estimate appears reasonable within the brain when compared to the mean angular phase,




Figure 3.6: Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.7 with resulting SNR (c). Average angular measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the within brain banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.

but poorly estimates the regions in empty space. Also the task related phase coefficient estimate is very unreasonable in empty space, and within the brain there is a very clear banding artifact that comes from the phase wraparound near the $[-\pi, \pi]$ boundary. This artifact makes it impossible to determine with confidence the existence of task-related phase change.

The results of applying the previously described von Mises model to the experimental data under the alternative hypothesis are shown in Figure 3.7 (Rowe et al., 2007). Similarly to the simulated data, we see improved estimation of both the baseline phase and task-related phase change. The results are more consistent with what we would expect based on the mean angular phase $\bar{\phi}$.

Figure 3.8 shows the maximum likelihood estimates of $\hat{\theta}_0$, $\hat{\sigma}^2$, and $\hat{\theta}_1$ under the null and alternative hypothesis in the top and bottom rows respectively; this is the experimental equivalent of Figure 3.4. In the top row, Figures 3.8a and 3.8b depict the estimates $\hat{\theta}_0$ and




Figure 3.7: Experimental data estimations from the von Mises approximation for $\hat{\theta}_0$ (a), $\hat{\theta}_1$ (b), and $\hat{\sigma}^2$ (c) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$. Note the high scale of variance in space compared to the Normal distribution in Figure 3.6 and later for the Lathi distribution in Figure 3.8.




Figure 3.8: Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta = \theta_0$ with anatomical image (c). Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta = \theta_0 + \theta_1 x_t$.

$\hat{\sigma}^2$ under the null hypothesis that $\theta = \theta_0$. In the bottom row, Figures 3.8d, 3.8e, and 3.8f depict the estimates $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ under the alternative hypothesis that $\theta = \theta_0 + \theta_1 x_t$. The




Figure 3.9: Z -statistics from each of the three distributions, thresholded by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The results from the Normal approximation (a), von Mises approximation (b), and mathematically correct Lathi distribution (c). The Lathi distribution was most successful in determining task-active voxels.

estimate $\hat{\theta}_0$ is approximately the same for both hypotheses while the estimate for $\hat{\sigma}^2$ appears to be slightly decreased in empty space under the alternative hypothesis. The estimate for $\hat{\theta}_1$ looks mostly uniform at about 0 inside the brain and noisy in empty space, as one might expect.

The z -statistics produced from the method described in Section 3.2.2, presented in Figure 3.9, are thresholded using the Benjamini-Hochberg procedure to control the false discovery rate at the $\alpha = 0.05$ level. The resultant critical z -statistics are found to be ± 3.18 , ± 4.1 , and ± 3.89 for the Normal, von Mises, and Lathi distributions, respectively. The left motor cortex is expected to have a task-related signal change as a result of right-hand finger tapping (the performed task for the experimental data). There is a distinct region of voxels that have task-related phase increase exactly at this location. The average value for $\hat{\theta}_1$ in the voxels that were determined to have task-related phase change was found to be 6.33 degrees. In each of Figures 3.6f, 3.7b, and 3.8f, the measure for $\hat{\theta}_1$ is a direct measure of the change in magnetic field local to the space that make up each of the voxels.

An experimental signal time series selected from voxels in the left and right motor cortices as well as from the posterior midline are presented in Figure 3.10 (center). A stable magnitude and phase are seen through time for both voxels (left). Neither the observed task-related phase change nor any task-related magnitude change are visually evident from




Figure 3.10: Experimental data time series (center) of the magnitude (blue) and phase (red) of voxels located in the left motor cortex (top), right motor cortex (bottom), and posterior midline (bottom). The gray shaded regions in the time series plots indicate TRs of subject task activation. Left and right are the significant z -statistics for magnitude- and phase-only activation, respectively.

the time series alone without further investigation. The gray shaded regions in the plots indicate TRs of subject task activation. The task-associated signal change is very small for both the magnitude and phase, making it difficult to visually infer from the time series alone. The magnitude (left) and phase (right) FDR thresholded z -statistics are shown as well as the locations of the selected voxels. The magnitude activation statistics are found from a magnitude-only model and thresholded to control the FDR at 0.05 (Rowe, 2005b).

The logarithm of the amplitudes of the Fourier coefficients of the Fourier transformed magnitude and phase time series from the same voxels in Figure 3.10 are presented in Figure 3.11. The zero-frequency has been omitted to further illustrate the differences between higher frequencies. The arrow in the top center plot for the voxel in the left motor cortex indicates a large Fourier coefficient at a frequency of about 0.03 Hz or $19/621 \text{ s}^{-1}$. The experimental timing had 19 epochs of the task on-off sequence that repeated every 32 seconds. No large Fourier coefficients are observed in the expected range for respiration (0.2-0.4 Hz). The TR of 1 s for the experimental data is consistent with fMRI experiments. As a result of the Nyquist criterion, any phenomenon occurring at a frequency greater than 0.5 Hz is unable to be isolated (e.g. cardiac rhythms in the 1-1.5 Hz range). The magnitude (left) and phase




Figure 3.11: Experimental data log Fourier coefficients (center) of the magnitude (blue) and phase (red) of voxels located in the left motor cortex (top), right motor cortex (middle), and posterior midline (bottom). The arrow in the top right plot indicates a large Fourier coefficient at a frequency of $19/621 \text{ s}^{-1}$ or about 0.03 Hz. Left and right are the significant F -statistics for the magnitude- and phase-only Fourier coefficients, respectively.

(right) FDR thresholded F -statistics for the Fourier coefficient at the previously mentioned 0.03 Hz are shown as well as the locations of the selected voxels. The F -statistics are computed by comparing the zero-frequency only model with a Fourier regression model. At each frequency, the F statistic is computed by

$$F_{2,n_t-2-1} = \frac{(RSS_1 - RSS_2)/2}{RSS_2/(n_t - 2 - 1)},$$

where RSS_1 is the residual sum of squares for the zero-frequency only model and RSS_2 is the residual sum of squares for the Fourier model that includes coefficients for specified, measurable frequencies. The F -statistics are thresholded at the $\alpha = 0.05$ level of significance. As indicated by the left and right figures, the time series for voxels in the left motor cortex were found to have statistically significant Fourier coefficients at a frequency of about 0.03 Hz.

3.4 Connection to Chapter 2

It would follow that the work presented in the previous sections on phase activation can also be applied to the Bayesian enhanced data as described in Chapter 2. The same maximum

likelihood estimation method described in Section 3.2 is applied to the iterated conditional modes maximum *a posteriori* estimates from Chapter 2. This is performed for the experimental data and results are shown below. Figure 3.12 shows the voxel-wise estimates that do not use the Lathi distribution in Equation 3.6. Figures 3.12a and 3.12b depict the MLE estimates for magnitude and variance from Equation 3.8. The apparent improvement in estimation for ρ when compared to Figure 3.6a is a direct results of the Bayesian enhancement of k -space. The variance estimate, particularly in space, is significantly reduced from the original experimental data. Figure 3.12c shows the signal-to-noise ratio (SNR) which is improved by a factor of about 5-6 from the original data. Figure 3.12d depicts the mean angular phase $\bar{\phi}$ of the MAP time series. Figures 3.12e and 3.12f are the simple linear regression coefficient estimates for baseline phase θ_0 and task-related phase θ_1 . Both the baseline phase and task-related phase change appear improved when compared to the original experimental data. The banding artifact, while still clearly present, has been significantly reduced.

Figure 3.13 shows the maximum likelihood estimates of $\hat{\theta}_0$, $\hat{\sigma}^2$, and $\hat{\theta}_1$ under the null and alternative hypothesis in the top and bottom rows respectively for the MAP estimates. In the top row, Figures 3.13a and 3.13b depict the estimates $\hat{\theta}_0$ and $\hat{\sigma}^2$ under the null hypothesis that $\theta_t = \theta_0$. The only noticeable difference from the original data here is the decrease in estimation for the variance, which is a direct result of the denoising of the data. In the bottom row, Figures 3.13d, 3.13e, and 3.13f depict the estimates $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ under the alternative hypothesis that $\theta_t = \theta_0 + \theta_1 x_t$. The estimate $\hat{\theta}_0$ and $\hat{\sigma}^2$ are both approximately the same for both hypotheses. The estimate for $\hat{\theta}_1$ looks largely comparable to the original experimental data aside from a new tendency in estimation that leans towards 0 in empty space.

The z -statistics produced from the method described in Section 3.2.2 when applied to the MAP estimated data are presented in Figure 3.14. The z -statistics are again filtered using the Benjamini-Hochberg procedure to control the false discovery rate at the $\alpha = 0.05$ level which resulted in a critical z -statistic of 4.18. The zoomed in region indicates the left motor cortex. There is a distinct region of voxels that have task-related phase increase, although a couple fewer voxels have been identified when compared to the original experimental data.




Figure 3.12: Maximum likelihood estimation for $\hat{\rho}$ (a) and $\hat{\sigma}^2$ (b) from Equation 3.8 with resulting SNR (c) for the MAP estimates of the experimental data. Average measured voxel phase $\bar{\phi}$ (d). Linear regression estimate for $\hat{\theta}_0$ (e) and $\hat{\theta}_1$ (f). Notice the banding effect in (f) that is a result of phase wrap-around. This can make diagnosing regions that contain task-related activity difficult.

The average value for $\hat{\theta}_1$ in the voxels that were determined to have task-related phase change was found to be 1.2 degrees, a significant decrease from the original data.

3.5 Discussion and Future Directions

It is rarely the case in fMRI studies that the whole complex-valued data is studied, leaving out the phase information to look only at the magnitude data. Further, when phase data *is* analyzed, the time series distribution of measurements is often simplified to the Normal distribution, which is only valid in voxels with high SNR. The Normal distribution also lacks support on the unit circle in favor of the real line, which creates artifacts when the measured phase is near the $[-\pi, \pi]$ boundary. In this work the mathematically correct non-Normal distribution from Lathi, which is valid for all SNRs, was used to estimate task-related changes in the phase signal. This maximum likelihood estimation model was tested on complex-valued simulated data and promising results were demonstrated on




Figure 3.13: Maximum likelihood estimations for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta_t = \theta_0$ with anatomical image (c) for the MAP estimates of the experimental data. Maximum likelihood estimations for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta_t = \theta_0 + \theta_1 x_t$.

complex-valued experimental data. In both the simulated and experimental data it was demonstrated that the non-Normal Lathi distribution works well at all SNRs and has detected task-related phase changes in the left motor cortex from a unilateral right-handed finger tapping experiment. This mathematically correct model has stronger detection ability than the popular Normal assumption model. It also performs slightly better than a von Mises approximation.

Work has been done to show that the sensitivity of the phase data can contain physiologically meaningful and spatially distinct information (Menon, 2002; Bandettini et al., 2005; Syväoja et al., 2025), though the exact statistics may be different due to multiple coils and a more elaborate preprocessing pipeline. This work continues to show exciting results for the ability to detect additional biological information contained in the phase: task-related activation. This motivates further investigations into phase activation resulting from local magnetic field changes due to vascularity or even direct neuronal current (Shmuel




Figure 3.14: Z -statistics from the likelihood ratio test statistic described in Equation 3.11 for the MAP estimates of the experimental data, filtered by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The critical z -statistic was found to be 4.18, as indicated by the black lines on the color bar. The left motor cortex is enhanced to the left. This is the region of expected activation for right hand finger tapping experiments.

and Leopold, 2008; Harris et al., 2011; Mullinger et al., 2014).

CHAPTER 4: BAYESIAN ESTIMATION AND RECONSTRUCTION OF NON-CARTESIAN K-SPACE FOR FMRI

4.1 Introduction

To yield an image of scientific and clinical value, the machine must measure many points in k -space. This can and has been done using a wide variety of pulse-sequences. Each pulse sequence consists of a series of RF bursts and controlled changes in magnetic gradients within the machine. The specifics of each pulse sequence determine both the manner in which k -space is traversed and the contrast enhanced for the image. As previously mentioned, fMRI images are most often collected via single shot echo planar imaging (EPI) and images are collected in two-dimensional slices (really each slice has some thickness on the order of about 2 mm). The most common in-plane method to sample k -space is to collect points on a uniform Cartesian grid. This allows for simple image reconstruction by way of the conventional two-dimensional inverse discrete Fourier transform (IDFT). An example GRE pulse sequence that would result in a Cartesian scan of k -space is presented in Figure 4.1. The top line “RF” indicates the timing of radio frequency bursts that flip the direction of the net magnetization M_0 by some angle α . The next three lines correspond to magnetic gradients in the slice select (z), phase-encode, and frequency-encode directions. The slice selection gradient G_{SS} isolates a certain slice within the brain. The phase-encode gradient G_{PE} bumps the sampling location up (or down) horizontal rows in k -space. The frequency-encode gradient G_{FE} pushes sampling along a row in k -space. The peak signal occurs exactly at the echo time TE, and the next images is collected exactly one repetition time (TR) apart from the last.

There exist many alternative pulse-sequences that involve nonuniform sweeps of k -space, which often require more thought when it comes to reconstructing images (Pipe, 1999; Block et al., 2014). One such technique, sometimes referred to as projection imaging (PI), fills in k -space using radial spokes rather than a series of horizontal and vertical segments as in conventional Cartesian encoding. This method of sampling k -space presents many benefits and some challenges that will be explored in this chapter.




Figure 4.1: An example of the gradient echo pulse sequence that would result in a Cartesian trajectory in k -space.

4.2 Radial Encoding of k -Space

PI as a method relates directly the projection-slice theorem. The projection-slice theorem (of 2-dimensions) states that the Fourier transform of the projection of a 2-D image into 1-D space is the equivalent to taking the 2-D Fourier transform of the image and selecting the spoke that goes through the origin and is orthogonal to the projection, as shown in Figure 4.5. Although current PI reconstruction techniques generally do not use the projection-slice theorem, it helps to provide a motivating force why one might want to fill k -space in this manner. Some PI models and other imaging techniques, such as Computed Tomography (CT), use the projection-slice theorem in a modified format called Filtered Back Projection (Ersoz et al., 2013). A feature of PI that is very desirable and a large motivation for why this project is of interest is that each spoke has “equal” value. This is due to the oversampling of the central region of k -space, where most of the information for the large details of the reconstructed image is contained, since every spoke goes through the center of k -space. As a result, this oversampling allows the reconstruction to be more resistant towards ghosting artifacts that are often observed resulting from motion of the patient during the measurements in the machine. This comes with the requirement that, for an image reconstructed via Cartesian acquisition with the same FOV and resolution, PI would

require a factor of $\pi/2$ additional sampling points (about 57% more points) (Bernstein et al., 2004). Since a primary objective is to enhance temporal resolution while maintaining image fidelity, it is desirable to implement some form of subsampling when measuring k -space in the radial scheme. Additionally, because the low spatial frequency region of k -space is so densely sampled compared to the regions of high spatial frequencies, a density compensation function must be factored in when reconstructing so that the resultant image is not blurry or biased by the low frequencies.




Figure 4.2: Cartesian sampling (a) is the most common method of traversing k -space. Radial encoding (b) of k -space has benefits, in particular, every spoke passes through the center of k -space.

4.3 Subsampling of k -Space

The most commonly performed methods to increase temporal resolution in fMRI involve either parallel imaging using either multiple coils or simultaneous multi-slice (SMS) imaging, subsampling k -space, or a combination of the two. Sensitivity encoding (SENSE) is a parallel imaging technique that employs multiple receive coils with unique sensitivity profiles (Pruessmann et al., 1999). These coils each generate an image that is sensitive to various regions of the brain determined by coil shape, size, and location. It is common to have upwards of 32 coils surrounding the brain when performing parallel imaging in this way. Controlled aliasing in parallel imaging (CAIPIRINHA) and multi-slice CAIPIRINHA using view angle tilting (CAIPIVAT) are popular methods for SMS imaging (Breuer et al., 2005; Kim et al., 2016). Loosely, these methods involve imaging multiple excited slices at the same time resulting in one image with several slices aliased together. The clever use of some phase modulations in k -space enable the anti-aliasing of the simultaneously collected slices.

Work towards a Bayesian model has been applied to CAIPIVAT which further enhanced the efficiency and image quality of the SMS imaging technique (Xu and Rowe, 2025). In both of these parallel imaging methods (or in the simpler case of one slice measured with one uniform coil) it is also possible to “skip” lines in k -space to speed up the imaging process. Generally, every n^{th} line is measured, as determined by the acceleration factor n_a . The resulting inverse Fourier reconstructed image has aliasing artifacts where it appears as if the brain has been cut into three strips which have been overlaid onto each other. Generalized autocalibrating partially parallel acquisitions (GRAPPA) is a technique that applies the skipping of lines in k -space to images collected with multiple coils. A kernel is formed in calibration images that uses the spatial frequency coefficients in the neighborhood of the missing k -space values that is used during experimentation to estimate the missing values, resulting in a full resolution image free of aliasing artifacts. There have been Bayesian approaches to enhance images using SENSE, GRAPPA, and a merged utilization of both techniques (Sakitis et al., 2024; Sakitis and Rowe, 2025a,b).

Figure 4.3 depicts a Cartesian k -space array that is fully sampled (as indicated by the red lines) (a), next to the same array sampled with an acceleration factor of $n_a = 2$ (b). The reconstructed magnitude image (c) has a noticeable aliasing artifact where the top of the brain has wrapped onto the bottom of the brain. Without techniques such as SENSE or GRAPPA, it can be extremely difficult to resolve regions of task-related activation with much confidence.




Figure 4.3: A fully scanned Cartesian k -space (a) and an example of a subsampled Cartesian k -space with an acceleration factor of $n_a = 2$. The resulting reconstructed image (c) suffers from an aliasing artifact.

The most time-impactful method to increase temporal resolution in radial k -space is to simply measure fewer spokes. This is functionally similar to skipping lines in a Cartesian

k -space. Aliasing is generally not a problem when it comes to subsampling a radial k -space. The largest negative effect of sampling fewer spokes is blurring within the brain, since it necessarily increases the sparsity of high-frequency coefficients that are measured. This weakens the ability to define sharp details and high contrast in the object being imaged. A version of GRAPPA has been developed and applied to subsampled radial k -space (Griswold et al., 2003; Seiberlich et al., 2011). Figure 4.4 shows a fully scanned (151 spokes) radial k -space (a) and a subsampled (48 spokes) radial k -space resulting in an acceleration factor n_a just over 3. The reconstructed image will have a resolution of 96×96 , hence the $\lceil 96 \cdot \pi/2 \rceil = 151$ spokes for the fully sampled k -space array. The original reconstructed image (c) compared to the reconstructed image from the subsampled k -space (d). The most noticeable effect of using fewer spokes in reconstruction is a blurring within the brain. The warping in the space surrounding the brain is a common artifact of radial reconstructions.




Figure 4.4: A fully scanned (151 spokes) radial k -space (a) and a subsampled (48 spokes) radial k -space resulting in an acceleration factor n_a just over 3. The original reconstructed image (c) compared to the reconstructed image from the subsampled k -space (d). The most noticeable effect of using fewer spokes in reconstruction is a blurring within the brain.

The *Highly Accelerated Projection Imaging* (HAPI) algorithm was introduced as a way to produce high quality images using as little as one densely sampled spoke to reconstruct images (Ersoz et al., 2013). This technique reflects the experimental ability to use only a few densely sampled spokes through k -space to reconstruct an image with high spatial resolution and remarkably high temporal resolution. A novel idea that sets HAPI apart is that it considers the amount each spoke passes through each voxel for each receiver coil. This allows the researchers to work with coil sensitivity-weighted projections of the image, which is shown to support high fidelity image reconstruction using fewer spokes. The HAPI

algorithm presents the linear equation

$$p = Af + \varepsilon, \quad (4.1)$$

where the vector p (of length $n_c n_p n_s$) represents the projections of the image, f is a vectorization of the object being imaged, the rows of matrix A are determined by receiver coil sensitivities while the columns correspond to voxels in the object f . If n_c is the number of coils, n_p is the number of projections (spokes), n_s is the number of samples per projection, and $m \times m$ is the dimension of the image, then A will have dimensions $[n_c n_p n_s, m^2]$ and f will have length m^2 . Lastly, $\varepsilon \sim N(0 + 0i, \sigma^2(I + iI))$ is the random noise that is measured by the machine. Complex-valued least squares regression is used to solve Equation 4.1 for f , which is necessary due to the noise and general invertibility of A .




Figure 4.5: The projection slice theorem in two dimensions states that the Fourier transform of a projection of a 2D object into 1D is equivalent to spoke of the 2D Fourier transform of the object that is orthogonal to the projection.

4.4 Bayesian Methodology

The methods that will be employed for this portion of the project are ongoing. Section 3 will serve as an introduction to the ideas used in this work. The final work may present as something similar to BGRAPPA, but in a radial k -space and further informed by the coil

sensitivities through the HAPI algorithm (Sakitis and Rowe, 2025a; Ersoz et al., 2013). Continuing work for this section goes specifically into the method of subsampling. In Cartesian encoding, it is most common that the *same* lines of k -space are omitted in every image for subsampling. Cartesian k -space encoding also has the restriction that the center line(s) of k -space *must* be sampled. Radial encoding does not have this limitation due to the fact that *every* spoke passes through the center of k -space. This allows for more freedom in designing subsampling methods that may allow for further enhancement of images that is not possible in conventional Cartesian encoding of k -space. A rough outline of the proposed method is presented in Figure 4.6. The imaging time series starts with a few fully sampled radial k -space arrays (a) that are used for calibration. During experimentation only a small subset of spokes will be collected (b), resulting in a significant temporal resolution increase. The Bayesian model will be employed to perform one or both of the following tasks: estimate the missing spokes based on calibration data (c), and enhance the collected spokes using the calibration data (d). This section is expected to be completed in Spring 2026.




Figure 4.6: The proposed method for Bayesian estimation of radial k -space. A fully sampled k -space for calibration (a) and a subsampled k -space during experimentation (b). Estimation of missing spokes (c) and enhancement of collected spokes (d).

4.5 Simulated Results

Simulated results for this project are forthcoming. SHAKER has been used to simulate an fMRI experiment that is closely representative of real data that we would expect to see. At present, work is being done to further understanding of the reconstruction methods used with radially sampled k -space data. This section is expected to be completed at the end of Spring 2026.

Preliminary Simulated data is presented in Figure 4.7. The MR parameters were set to be the same as in Section 3.3.1, but with a radial trajectory of k -space. The TR was

significantly increased as well to simulate the first image in an fMRI times series. The spokes are stacked vertically, left to right, so that the center of k -space observed as a narrow bright band going through the middle of the arrays shown. For a 128×128 reconstructed image, 202 spokes are needed for a complete reconstruction. Figures 4.7a, 4.7b, 4.7c, and 4.7d show a simulated radial k -space for the exact same image but with 202, 101, 51, and 26 spokes respectively. This corresponds to acceleration factors of $n_a = 1$, $n_a = 2$, $n_a = 4$, and $n_a = 8$. Reconstructed magnitude images, as well as M_0 map of the original slice are presented in Figure 4.8 using the Michigan Image Reconstruction Toolbox (MIRT) (Fessler, sent; Fessler and Sutton, 2003). Signal degradation and blurring are evident as the amount of spokes used in reconstruction decreases. Figure 4.8e is reconstructed from the k -space array from Figure 4.7d which measured only 26 spokes. The blurring and lack of ability to make out sharp details is apparent, especially in contrast to Figure 4.8a. It is our goal to use a similar amount of data as in Figure 4.7d but reconstruct an image with quality closer to that of Figure 4.8a.




Figure 4.7: Preliminary simulated data from SHAKER. Measured k -space arrays with 128 samples per spoke, and (a) 202 spokes, (b) 101 spokes, (c) 51 spokes, and (d) 26 spokes.




Figure 4.8: Preliminary simulated data from SHAKER. The M_0 map of the slice used from the digital phantom (a) next to reconstructed magnitude images using (b) 202 spokes, (c) 101 spokes, (d) 51 spokes, and (e) 26 spokes.

4.6 Discussion and Future Directions

Present work in the realm of radial sampling of k -space has shown many advantages over the conventional Cartesian encoding. Future work for this project will further explore this idea. Similar methods to those presented in Chapter 3 will be employed to develop a Bayesian model to enhance subsampled radial k -space arrays which will reconstruct high fidelity images with a significant improvement to temporal resolution. Considerations for this project will include selecting a reconstruction algorithm that best suits our model (at present, HAPI). It has been shown that not using the same spokes for every image can increase image quality (Feng et al., 2014). Consequently, we will study the different methods of subsampling radial k -space and determining a method that is optimal. Using a few, fully-sampled radial k -space arrays will allow us to assess hyperparameters that will lead to brighter signal, higher contrast, and less noise in images that are also robust to motion artifacts in the machine. This work is expected to be completed by Summer 2026.

BIBLIOGRAPHY

- Adrian, D. W., Maitra, R., and Rowe, D. B. (2013). Ricean over Gaussian Modelling in Magnitude fMRI Analysis-Added Complexity with Negligible Practical Benefits. *Stat*, 2(1):303–316.
- Adrian, D. W., Maitra, R., and Rowe, D. B. (2018). Complex-Valued Time Series Modeling for Improved Activation Detection in fMRI Studies. *The Annals of Applied Statistics*, 12(3):1451–1478.
- Adrian, D. W., Maitra, R., and Rowe, D. B. (2024). Magnitude and Phase Activation Detection in Functional MRI Time Series. *Proceedings of the Joint Statistical Meeting*.
- Adrian, D. W., Maitra, R., and Rowe, D. B. (2025). Rice-Distributed Autoregressive Time Series Modeling of Magnitude Functional MRI Data. *The Annals of Applied Statistics*, 19.
- Ardekani, B. A. and Kanno, I. (1998). Statistical Methods for Detecting Activated Regions in Functional MRI of the Brain. *Magnetic Resonance Imaging*, 16(10):1217–1225.
- Bakas, K., Kornak, J., and Ombao, H. C. (2025). An MCMC Approach to Bayesian Image Analysis in Fourier Space. *Data Science in Science*, 4(1):2452603.
- Bandettini, P. A., Jesmanowicz, A., Wong, E. C., and Hyde, J. S. (1993). Processing Strategies for Time-Course Data Sets in Functional MRI of the Human Brain. *Magnetic Resonance in Medicine*, 30(2):161–173.
- Bandettini, P. A., Petridou, N., and Bodurka, J. (2005). Direct Detection of Neuronal Activity with MRI: Fantasy, Possibility, or Reality? *Applied Magnetic Resonance*, 29(1):65–88.
- Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S., and Hyde, J. S. (1992). Time Course EPI of Human Brain Function During Task Activation. *Magnetic Resonance in Medicine*, 25(2):390–397.

- Benjamini, Y. and Hochberg, Y. (2018). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.
- Bernstein, M., King, K., and Zhou, X. (2004). *Handbook of MRI Pulse Sequences*. Elsevier Inc. Publisher Copyright: © 2004 Elsevier Inc. All rights reserved.
- Best, D. J. and Fisher, N. I. (1979). Efficient Simulation of the von Mises Distribution. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 28(2):152–157.
- Block, K., Chandarana, H., Milla, S., Bruno, M., Mulholland, T., Fatterpekar, G., Hagiwara, M., Grimm, R., Geppert, C., Kiefer, B., and Sodickson, D. (2014). Towards Routine Clinical Use of Radial Stack-of-Stars 3D Gradient-Echo Sequences for Reducing Motion Sensitivity. *Journal of the Korean Society of Magnetic Resonance in Medicine*, 18:87.
- Bodenschatz, J. C. and Rowe, D. B. (2025). Simulation and Harmonic Analysis of k-Space Readout (SHAKER). *arXiv*.
- Breuer, F. A., Blaimer, M., Heidemann, R. M., Mueller, M. F., Griswold, M. A., and Jakob, P. M. (2005). Controlled Aliasing in Parallel Imaging Results in Higher Acceleration (CAIPIRINHA) for Multi-Slice Imaging. *Magnetic Resonance in Medicine*, 53(3):684–691.
- Brown, R. W., C. Norman Cheng, Y., Haacke, E. M., Thompson, M. R., and Venkatesan, R. (2014). *Magnetic Resonance Imaging*. Wiley-Blackwell, Chichester, England, 2 edition.
- Casella, G. C. and Berger, R. (2024). *Statistical Inference*. Chapman and Hall/CRC.
- Chavhan, G. B., Babyn, P. S., Thomas, B., Shroff, M. M., and Haacke, E. M. (2009). Principles, Techniques, and Applications of T_2^* -based MR Imaging and its Special Applications. *RadioGraphics*, 29(5):1433–1449.
- Chen, J. E. and Glover, G. H. (2015). Functional Magnetic Resonance Imaging Methods. *Neuropsychology review*, 25(3):289–313.

- Comby, P.-A., Vignaud, A., and Ciuciu, P. (2024). SNAKE-fMRI: A Modular fMRI Data Simulator from the Space-Time Domain to k-Space and Back. *arXiv*.
- den Dekker, A. J. and Sijbers, J. (2005). Implications of the Rician Distribution for fMRI Generalized Likelihood Ratio Tests. *Magnetic Resonance Imaging*, 23(9):953–959.
- Elster, A. D., Burdette, J. H., and Field, A. (2001). *Questions & answers in magnetic resonance imaging*. Mosby, St. Louis, second edition. edition.
- Ersoz, A., Arpınar, V. E., and Muftuler, L. T. (2013). Highly Accelerated Projection Imaging with Coil Sensitivity Encoding for Rapid MRI. *Med. Phys.*, 40(2):022305.
- Feng, L., Grimm, R., Block, K. T., Chandarana, H., Kim, S., Xu, J., Axel, L., Sodickson, D. K., and Otazo, R. (2014). Golden-Angle Radial Sparse Parallel MRI: Combination of Compressed Sensing, Parallel Imaging, and Dolden-Angle Radial Sampling for Fast and Flexible Dynamic Volumetric MRI. *Magnetic Resonance in Medicine*, 72(3):707–717.
- Fessler, J. A. (2007–present). Michigan Image Reconstruction Toolbox (MIRT). <https://web.eecs.umich.edu/~fessler/code/>. Online; accessed YEAR-MONTH-DAY.
- Fessler, J. A. and Sutton, B. P. (2003). Nonuniform Fast Fourier Transforms Using min-max Interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574.
- Fox, C. (1928). The asymptotic expansion of generalized hypergeometric functions. *Proceedings of the London Mathematical Society*, s2-27(1):389–400.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-Based Approaches to Calculating Marginal Densities. *Journal of the American Statistical Association*, 85(410):398–409.
- Geman, S. and Geman, D. (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741.
- Griswold, M. A., Heidemann, R. M., and Jakob, P. M. (2003). Direct Parallel Imaging Reconstruction of Radially Sampled Data Using GRAPPA with Relative Shifts.

- Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., and Haase, A. (2002). Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–1210.
- Gudbjartsson, H. and Patz, S. (1995). The Rician Distribution of Noisy MRI Data. *Magnetic Resonance in Medicine*, 34(6):910–914.
- Haase, A., Frahm, J., Matthaei, D., Hanicke, W., and Merboldt, K.-D. (1986). FLASH Imaging. Rapid NMR Imaging Using Low Flip-Angle Pulses. *Journal of Magnetic Resonance (1969)*, 67(2):258–266.
- Hahn, A. D. and Rowe, D. B. (2012). Physiologic Noise Regression, Motion Regression, and TOAST Dynamic Field Correction in Complex-Valued fMRI Time Series. *NeuroImage*, 59(3):2231–2240.
- Hargreaves, B. A. (2012). Rapid Gradient-Echo Imaging. *Journal of Magnetic Resonance Imaging*, 36(6):1300–1313.
- Harris, J. J., Reynell, C., and Attwell, D. (2011). The Physiology of Developmental Changes in BOLD Functional Imaging Signals. *Developmental Cognitive Neuroscience*, 1(3):199–216.
- Henkelman, R. M. (1985). Measurement of Signal Intensities in the Presence of Noise in MR Images. *Medical Physics*, 12(2):232–233.
- Holland, D., Kuperman, J. M., and Dale, A. M. (2010). Efficient Correction of Inhomogeneous Static Magnetic Field-Induced Distortion in Echo Planar Imaging. *NeuroImage*, 50(1):175–183.
- Karaman, M. M. (2016). *Improving fMRI Analysis and MR Reconstruction With the Incorporation of MR Relaxivities and Correlation Effect Examination*. PhD thesis, Marquette University.
- Karaman, M. M., Bruce, I. P., and Rowe, D. B. (2014). A statistical fMRI model for differential T2* contrast incorporating T1 and T2* of gray matter. *Magnetic Resonance Imaging*, 32(1):9–27.

- Karaman, M. M., Bruce, I. P., and Rowe, D. B. (2015). Incorporating relaxivities to more accurately reconstruct MR images. *Magnetic Resonance Imaging*, 33(4):374–384.
- Kida, I., Kennan, R. P., Rothman, D. L., Behar, K. L., and Hyder, F. (2000). High-Resolution CMRO2 Mapping in Rat Cortex: A Multiparametric Approach to Calibration of BOLD Image Contrast at 7 Tesla. *Journal of Cerebral Blood Flow & Metabolism*, 20(5):847–860. PMID: 10826536.
- Kim, M.-O., Hong, T., and Kim, D.-H. (2016). MultiSlice CAIPIRINHA Using View Angle Tilting Technique (CAPIVAT). *Tomography*, 2(1):43–48.
- Koay, C. G. and Basser, P. J. (2006). Analytically Exact Correction Scheme for Signal Extraction From Noisy Magnitude MR Signals. *Journal of magnetic resonance*, 179 2:317–22.
- Koay, C. G., Özarslan, E., and Basser, P. J. (2009). A Signal Transformational Framework for Breaking the Noise Floor and its Applications in MRI. *Journal of Magnetic Resonance*, 197(2):108–119.
- Kornak, J., Young, K., Friedman, E., and Bakas, K. (2026). Bayesian Image Analysis in Fourier Space. *Journal of the American Statistical Association*, 0(0):1–14.
- Kumar, A., Welte, D., and Ernst, R. R. (1975). NMR Fourier Zeugmatography. *Journal of Magnetic Resonance*, 18(1):69–83.
- Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D. N., Hoppel, B. E., Cohen, M. S., and Turner, R. (1992). Dynamic Magnetic Resonance Imaging of Human Brain Activity During Primary Sensory Stimulation. *Proceedings of the National Academy of Sciences*, 89(12):5675–5679.
- Lai, S. and Glover, G. H. (1997). Detection of BOLD fMRI Signals Using Complex Data. *Proceedings of the International Society for Magnetic Resonance in Medicine*, 1997(S1):1671.
- Larmor, J. (1897). A Dynamical Theory of the Electric and Luminiferous Medium.— Part III. Eelations With Material media. *Philosophical Transactions of the Royal Society of*

- London. Series A, Containing Papers of a Mathematical or Physical Character*, 190:205–300.
- Lathi, B. P. (1983). *Modern Digital and Analog Communication*. Oxford University Press.
- Lindley, D. V. and Smith, A. F. M. (1972). Bayes Estimates for the Linear Model. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(1):1–41.
- Lindquist, M. A. (2008). The Statistical Analysis of fMRI Data. *Statistical Science*, 23(4):439–464.
- Logan, B. R., Geliazkova, M. P., and Rowe, D. B. (2008). An Evaluation of Spatial Thresholding Techniques in fMRI Analysis. *Hum. Brain Mapp.*, 29(12):1379–1389.
- Logan, B. R. and Rowe, D. B. (2004). An Evaluation of Thresholding Techniques in fMRI Analysis. *Neuroimage*, 22(1):95–108.
- Menon, R. S. (2002). Postacquisition suppression of large-vessel bold signals in high-resolution fMRI. *Magnetic Resonance in Medicine*, 47(1):1–9.
- Mullinger, K., Mayhew, S., Bagshaw, A., Bowtell, R., and Francis, S. (2014). Evidence that the Negative BOLD Response is Neuronal in Origin: A Simultaneous EEG–BOLD–CBF Study in Humans. *NeuroImage*, 94:263–274.
- Nan, F. Y. and Nowak, R. D. (1999). Generalized Likelihood Ratio Detection for fMRI Using Complex Data. *IEEE Transactions on Medical Imaging*, 18(4):320–329.
- Nencka, A. S., Hahn, A. D., and Rowe, D. B. (2009). A Mathematical Model for Understanding the STatistical effects of k-space (AMMUST-k) Preprocessing on Observed Voxel Measurements in fcMRI and fMRI. *Journal of Neuroscience Methods*, 181(2):268–282.
- Nencka, A. S. and Rowe, D. B. (2005). Complex Constant Phase Activation Model Removes Venous BOLD Contribution in fMRI.
- Nencka, A. S. and Rowe, D. B. (2007). Reducing the Unwanted Draining Vein BOLD Contribution in fMRI with Statistical Post-Processing Methods. *NeuroImage*, 37(1):177–188.

- Ogawa, S., Lee, T. M., Kay, A. R., and Tank, D. W. (1990). Brain Magnetic Resonance Imaging with Contrast Dependent on Blood Oxygenation. *Proceedings of the National Academy of Sciences U. S. A.*, 87(24):9868–9872.
- O’Hagan, A. (1994). *Kendall’s Advanced Theory of Statistics: Bayesian inference. vol. 2B*. Kendall’s library of statistics. Edward Arnold.
- Penny, W. D., Trujillo-Barreto, N. J., and Friston, K. J. (2005). Bayesian fMRI time series analysis with spatial priors. *NeuroImage*, 24(2):350–362.
- Pipe, J. G. (1999). Motion Correction with PROPELLER MRI: Application to Head Motion and Free-Breathing Cardiac Imaging. *Magnetic Resonance in Medicine*, 42(5):963–969.
- Pruessmann, K. P., Weiger, M., Scheidegger, M. B., and Boesiger, P. (1999). SENSE: Sensitivity Encoding for Fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962.
- Rayleigh, J. (1880). 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(60):73–78.
- Rice, S. O. (1944). Mathematical Analysis of Random Noise. *Bell System Technical Journal*, 23(3):282–332.
- Rowe, D. B. (2005a). Modeling Both Magnitude and Phase of Complex fMRI Data. *NeuroImage*, 25:1310–24.
- Rowe, D. B. (2005b). Parameter Estimation in the Magnitude-Only and Complex-Valued fMRI Data Models. *Neuroimage*, 25(4):1124–1132.
- Rowe, D. B. (2009). Magnitude and Phase Signal Detection in Complex-Valued fMRI Data. *Magnetic Resonance in Medicine*, 62(5):1356–7; 1358–60.
- Rowe, D. B. (2016). *Image Reconstruction in Functional MRI*. CRC Press, first edition. edition.
- Rowe, D. B. (2023). Statistics of Intrinsic FMRI Data. *Proceedings of the Joint Statistical Meeting*.

- Rowe, D. B. and Bodenschatz, J. C. (2025). Distributionally Accurate fMRI Phase Activation. *Proceedings of the Joint Statistical Meeting*.
- Rowe, D. B. and Logan, B. R. (2004). A Complex Way to Compute fMRI Activation. *Neuroimage*, 23(3):1078–1092.
- Rowe, D. B., Meller, C. P., and Hoffmann, R. G. (2007). Characterizing Phase-Only fMRI Data with an Angular Regression Model. *Journal of Neuroscience Methods*, 161(2):331–341.
- Sakitis, C. J., Brown, D. A., and Rowe, D. B. (2024). A Bayesian Complex-Valued Latent Variable Model Applied to Functional Magnetic Resonance Imaging. *Journal of the Royal Statistical Society Series C: Applied Statistics*, page qlae046.
- Sakitis, C. J. and Rowe, D. B. (2025a). A Bayesian Approach to GRAPPA Parallel fMRI Image Reconstruction Increases SNR and Power of Task Detection. *The Annals of Applied Statistics*, 19(2).
- Sakitis, C. J. and Rowe, D. B. (2025b). Bayesian Merged Utilization of GRAPPA and SENSE (BMUGS) for In-Plane Accelerated Reconstruction Increases fMRI Detection Power. *Magnetic Resonance Imaging*, 115:110252.
- Seiberlich, N., Ehse, P., Duerk, J., Gilkeson, R., and Griswold, M. (2011). Improved Radial GRAPPA Calibration for Real-Time Free-Breathing Cardiac Imaging. *Magn. Reson. Med.*, 65(2):492–505.
- Severini, T. A. (2000). *Likelihood Methods in Statistics*. Oxford Statistical Science Series. Oxford University Press, London, England.
- Shmuel, A. and Leopold, D. A. (2008). Neuronal Correlates of Spontaneous Fluctuations in fMRI Signals in Monkey Visual Cortex: Implications for Functional Connectivity at Rest. *Human Brain Mapping*, 29(7):751–761.
- Sun, J., Kong, M., and Pal, S. (2023). The Modified-Half-Normal distribution: Properties and an Efficient Sampling Scheme. *Communications in Statistics - Theory and Methods*, 52(5):1591–1613.

- Syv  oja, S. S., Raitamaa, L., Helakari, H., Kantola, J., J  rvel  , M., Kananen, J., Isokoski, V., Korhonen, V., V  yrynen, T., Tuovinen, T., Hennig, J., and Kiviniemi, V. (2025). Ultrafast Complex-Valued 4D fMRI Reveals Sleep-Induced Brain Respiratory Pulsation Changes in Both Magnitude and Phase Signals. *NeuroImage*, 317:121290.
- The MathWorks Inc. (2022). MATLAB version: 24.1.0 (R2024a).
- Von Mises, R. (1918). Uber die 'Ganzzahligkeit' der Atomgewicht und Verwandte Fragen. *Physikalische Zeitschrift*, 19:490–500.
- Wang, X., Hernando, D., and Reeder, S. B. (2020). Phase-based T_2 Mapping with Gradient Echo Imaging. *Magnetic Resonance in Medicine*, 84(2):609–619.
- Wang, Z., Rowe, D. B., Li, X., and 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.
- Wang, Z., Rowe, D. B., Li, X., and Brown, D. A. (2025). Efficient Fully Bayesian Approach to Brain Activity Mapping with Complex-Valued fMRI Data. *Journal of Applied Statistics*, 52(6):1299–1314.
- Welvaert, M., Durnez, J., Moerkerke, B., Berdoolaege, G., and Rosseel, Y. (2011). neuRosim: An R Package for Generating fMRI Data. *Journal of Statistical Software*, 44(10):1–18.
- Wilks, S. S. (1938). The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62.
- Wright, E. M. (1935). The Asymptotic Expansion of the Generalized Hypergeometric Function. *Journal of the London Mathematical Society*, s1-10(4):286–293.
- Xu, K. and Rowe, D. B. (2025). A Bayesian CAIPIVAT Approach with Through-Plane Acceleration to Enhance Efficiency of Simultaneously Encoded Slice Acquisition in FMRI. *Magnetic Resonance Imaging*, 124:110540.
- Yu, C.-H., Prado, R., Ombao, H. C., and Rowe, D. B. (2018). A Bayesian Variable Selection

- Approach Yields Improved Detection of Brain Activation From Complex-Valued fMRI. *Journal of the American Statistical Association*, 113(524):1395–1410.
- Yu, C.-H., Prado, R., Ombao, H. C., and Rowe, D. B. (2023). Bayesian Spatiotemporal Modeling on Complex-Valued fMRI Signals via Kernel Convolutions. *Biometrics*, 79(2):616–628.
- Zhou, Y., Shi, C., Lai, B., and 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(9).

APPENDICES

APPENDIX A: SHAKER VALUES AND DATA STRUCTURES

A.1 SHAKER Default Values

This section quickly defines the default values entered for SHAKER upon startup. In the *Data Viewer* pane, we have:

```
Phantom Size: 96
Map: M0
View: Axial
Slice: 48
```

The *MRI Parameters* are set to:

```
Signal Equation: GradientEcho_SigEq.m*
Trajectory: Cartesian_kspace.m*
Acceleration Factor: 1
Field Strength: 3
TE: 50
TR: 1000
Flip Angle: 90
EESP: 0.72
Number of Coils: 1
Reconstruction: CartesianIFFT.m*
Include B0 Inhomogeneity: 1
```

*: or first file alphabetically in respective folders. Lastly, the default *Task Design* options are preset to

```
HRF: Block
Initial Rest Images: 10
Epochs: 20
Task Images per Epoch: 15
Rest Images per Epoch: 15
SNR: 5
Task 1
CNR: 0.75
Phase: 0
```

A.2 SHAKER Structure and Objects

The SHAKER directory is structured in the following way:

```
main
  data                                : contains phantom and activation data
    Phantom96.mat                     : 4x[96x96x96] struct of digital phantom
    ActMap96.mat                       : [96x96x96] logical double
```

```

functions           : internal helper functions of SHAKER
kSpaceTrajectories  : folder of custom k-space trajectories
reconstructionMethods : folder of custom reconstruction methods
signalEquations     : folder of custom signal equations
statisticsFunctions : folder of statistical tests for the time series
ksim.mlapp          : main SHAKER application

```

As shown, the `data` folder contains the digital phantom and default activation map that were described in Sections 1.2.2 and 1.2.3. The `functions` folder contains `generateTimeSeries.m` and `generateCoils.m`. The `generateTimeSeries.m` file takes in a noiseless rest and task k -space image as well as the experimental design, then outputs the complex-valued time series array with size $[n_{k_x} \times n_{k_y} \times n_c \times n_{IMG}]$. Other “helper” functions will continue to be placed here.

The `kSpaceTrajectories` folder contains the trajectories described in Section 1.1.2. The `reconstructionMethods` folder contains all reconstruction method files (simple Cartesian IFT being the most commonly used). The `signalEquations` folder contains the signal equations described in Section 1.1.1. The `statisticsFunctions` folder contains statistical maps used to analyze the time series data. All four of these folders are where the relevant custom user files should be placed. There is a provided template in each of these folders to help support the creation of custom files. The input/output structure of each is given below. Note that while the templates expect each of the described inputs and custom files should be structured this way, they need not be used in the functions themselves. It is recommended for advanced statistical techniques and models that the simulated time series data be exported from data and examined in a more controlled setting.

```

template_kSpace IN
    MRI : struct object from SHAKER
template_kSpace OUT
    [kx, ky, timeMap] : array of single k-space measurement orders
                        kx      : real double of measured k-space x locations
                        ky      : real double of measured k-space y locations
                        timeMap  : real double of time (s) at measurement of k.s. points

template_reconstruction IN
    MRI      : struct object from SHAKER
    kSpace    : struct object from SHAKER
    kSpaceTimeSeries : [nkx x nky x nc x nimg] complex double of k-space
template_reconstruction OUT
    imageTimeSeries : [nkx x nky x nc x nimg] complex double of image space

template_sigEq IN
    M0      : [nx x ny] real double of M0 slice

```



```

    T1      : [nx x ny] real double of T1 slice
    T2star  : [nx x ny] real double of T2* slice
    deltaB  : [nx x ny] real double of dB slice
    kSpace  : struct object from SHAKER
    MRI     : struct object from SHAKER
template_sigEq OUT
    kspace  : [nkx x nky] complex double of single noiseless k-space array

template_stats IN
    TS      : [n(k)x x n(k)y x nc x nimg] complex double of img- or k-s t.s.
    kSpace  : struct object from SHAKER
    MRI     : struct object from SHAKER
template_stats OUT
    statsMap : [n(k)x x n(k)y x nc x nimg] real double of statistical map

```

The two main objects referenced in many of the functions and operability of SHAKER are the MRI and kSpace objects. These struct objects contain the relevant information about the MRI settings and the *k*-space settings, respectively. The objects are made up of the following variables:

```

MRI (struct)
    SignalEquation (int)      : value indicating selected sig. eq.
    Trajectory (int)         : value indicating selected trajectory
    AccelerationFactor(int)   : n_a acceleration factor
    FieldStrength (double)    : magnetic field strength
    EchoTime (double)         : TE stored in SI (seconds)
    RepetitionTime (double)   : TR stored in SI (seconds)
    EESP (double)            : effective echo spacing (milliseconds)
    FlipAngle (double)        : flip angle in deg
    NumberOfCoils (int)       : n_c number of coils
    CoilSensitivities (double) : [nx x ny x nc] coil sensitivities
    gamma (double)           : gyromagnetic ratio in MHz/T
    RelaxivityTable (double)  : table of real values
    IncludeB0Inhomogeneity (int): indicator of if dB was applied

kSpace (struct)
    kX (double)              : measured k-space x locations
    kY (double)              : measured k-space y locations
    timeMap (double)         : time (s) at measurement of k.s. points.

```

APPENDIX B: FMRI TIME SERIES PHASE CORRECTION AND FURTHER USES

B.1 Phase Correction in fMRI Time Series

As was mentioned in Sections 2.4.1 and 3.3.2, the spatial and temporal changes in magnetic field inhomogeneities causes a drift in phase over time. This is shown in Figure B1. The real and imaginary parts of a k -space value drift over time, so it is preferred to work in the magnitude and phase scheme where only the phase changes over time.




Figure B1: A depiction of phase drift in k -space measurements over time. For a given spatial frequency coefficient in steady state images the magnitude will remain approximately the same but the phase will drift over time due to changing magnetic fields.

To illustrate the phase correction process, a random voxel within the experimental brain data was selected, and its phase time series isolated. This original phase time series is the blue line labeled “1” in Figure B2. At about the 100th TR, the phase drift has reached the $[-\pi, \pi]$ boundary, and wraparound artifacts are present. The first step in correcting this is subtracting off the angular mean $\bar{\phi}$, resulting in the orange line labeled “2”. Next, the centered time series is unwrapped from the $[-\pi, \pi]$ boundary shown by the yellow line labeled “3”. A linear fit is modeled to the unwrapped phase then subtracted off, resulting in a linearly stable phase over time as seen in the purple line labeled “4”. A spatial model is then fit to the surface of the brain, smoothing out some noise in the phase over time. This subtle change is shown by the green line labeled “5”. The final step is to add the original angular mean $\bar{\phi}$ back to the time series resulting in a linearly stable and spatially smoothed phase time series (light blue line labeled “6”).




Figure B2: A single voxel's time series through the stages of phase-drift correction.

B.2 T_1 Map Estimation

It has been shown that the first image as well as an average steady-state image can be used to estimate a T_1 map of the object being imaged (Karaman, 2016). The signal produced after the first RF pulse is given by

$$M_1 = M_0 e^{-TE/T_2^*}. \quad (\text{B1})$$

The steady-state signal is given by

$$M_{SS} = M_0 \left(1 - e^{TR/T_1}\right) e^{-TE/T_2^*}. \quad (\text{B2})$$

So, a T_1 map estimate can be done by evaluating

$$T_1 = \frac{TR}{\ln\left(\frac{R}{R-1}\right)}, \quad (\text{B3})$$

where $R = M_1/M_{SS}$. This was done practically using the same data set in Section 2.4 by setting the first image in the time series to be M_1 , and averaging images 6-16 to make up M_{SS} . The results are shown in Figure B3 and are consistent with known T_1 values for the tissues that make up the human brain. The T_1 map from the digital phantom described in Section 1.2.1 as well as a T_1 estimation based on a simulated time series is also presented.




Figure B3: Left is the digital phantom T_1 map. To the right are the estimations of the T_1 maps using the first image of an fMRI time series in addition to steady-state images at TRs 6-16 for both a simulated and experimental time series. The units are in seconds and estimation is consistent with known T_1 values.

APPENDIX C: A FURTHER LOOK INTO PHASE-ONLY ACTIVATION

C.1 Power analysis of alternative distributions

In the hypothesis test described in Section 3.2 a significance level of $\alpha = 0.05$ was chosen. This represents the probability of committing a Type I error or false positive, which in this application would mean falsely claiming a voxel has some task-related phase change when in reality it does not. A result of fixing $\alpha = 0.05$ is the value of the power β of the hypothesis test. Statistical power is the ability of a hypothesis test to correctly reject the null hypothesis when it is false, or for this application, the probability of correctly identifying a voxel as having some task-related phase change. A depiction of statistical power is shown in Figure C1. In this figure, the maximum likelihood estimates of the parameters from an experimental voxel that was determined to be task-active are used to draw theoretical distributions. In black is the distribution for phase measurements under the null hypothesis that $\theta = \theta_0$. In red, blue, and green are the distributions for $\theta_0 + \theta_1$ under the alternative hypothesis that $\theta = \theta_0 + \theta_1 x_t$ using the mathematically correct Lathi distribution, the Normal approximation, and the von Mises approximation respectively. The power is represented by the similarly colored area under each $\theta_0 + \theta_1$ distribution curve, to the right of the critical value $\theta_{0\alpha}$. Here we can visually see that the Lathi distribution (red) and von Mises distribution (green) provide larger statistical power than the Normal




Figure C1: Theoretical distributions based on maximum likelihood estimates from an active voxel in Figure 3.9. The null hypothesis (black) assumes the Lathi distribution in Equation 3.6. The alternative hypothesis using the Lathi (red), Normal (blue), and von Mises (green) distributions. The power is the highlighted area to the right of the critical value for each alternative hypothesis distribution.

assumption (blue) due to the large difference in area to the right of the critical value. The Lathi and von Mises distributions perform similarly.

The example in Figure C1 is for a specific voxel with an estimated value for θ_0 and $\theta_0 + \theta_1$. One way to generalize to other voxels is to vary the effect size (i.e. compare different values of θ_1) and observe how the power changes. This was done without loss of generality by using the estimate for θ_0 from the same voxel that was used for the previous example and adjusting the estimate for θ_1 from 0 to 1 radian. The power for each effect size is then calculated and plotted, resulting in Figure C2. The red, blue, and green power curves are for the alternative hypothesis following a Lathi, Normal, and von Mises distribution respectively. For effect sizes where $\theta_1 \lesssim 0.5$ radians (about 30 degrees), the Lathi and von Mises distributions provide higher statistical power. The maximum likelihood estimate for θ_1 for the voxel from the active experimental voxel is marked by a vertical line at about 0.18. Measurements for θ_1 are observed to be at or below this level.




Figure C2: Power curves for Lathi (red), Normal (blue), and von Mises (green) alternative hypotheses under a null hypothesis that assumes the Lathi distribution in Equation 3.6. The vertical line at about 0.18 indicates where the maximum likelihood estimate $\hat{\theta}_1$ was located for this particular active voxel. The inset image is a zoomed in region of interest to highlight the difference in statistical power of the three distributions.

Note that this demonstrative example is for a one-tailed hypothesis test which assumes the phase change will be positive. The methods described in Section 3.2 support a two-tailed test, which is what is measured in practice in Section 3.3. This allows for the possibility of a negative change in phase, though it is observed that most changes are in the positive direction.

C.2 Grid search approach to parameter estimation

To validate the results produced from the maximum likelihood estimation, a simple grid search was performed to estimate the parameters θ_0 , θ_1 , and σ^2 for the null and alternative hypothesis. Search boundaries were tightened using prior knowledge of what the values for $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ were expected to be. The search interval for the baseline phase was $\hat{\theta}_0 \in [\bar{\phi} - \pi/24, \bar{\phi} + \pi/24]$, which assumes that the estimate $\hat{\theta}_0$ should be reasonably close to the mean value $\bar{\phi}$. The search interval for the additive phase was $\hat{\theta}_1 \in [-\pi/72, \pi/24]$ which assumes that the task-related phase change in a voxel is relatively small and likely

positive. Lastly, the search interval for the variance was set to $\hat{\sigma}^2 \in [\hat{\sigma}^2 - \hat{\sigma}^2/5, \hat{\sigma}^2 + \hat{\sigma}^2/5]$ which assumes that the variance is reasonably close to the previously estimated variance $\hat{\sigma}^2$ from the magnitude estimation as described in Section 3.2.2. Each interval had 10 linearly spaced search points, resulting in 100 possible combinations of $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ for the null hypothesis and 1,000 possible combinations for the alternative hypothesis. This is a very low precision grid search; best practice would involve a significantly denser search space. However, this small search took 3 hours to run on the same computer that computed the maximum likelihood estimates in less than 3 minutes.

Figure C3 shows the estimated parameters $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\sigma}^2$ under the null and alternative hypothesis in the top and bottom rows respectively. The baseline phase estimate $\hat{\theta}_0$ and variance estimate $\hat{\sigma}^2$ are comparable across the hypotheses. An interesting contrast between the grid search result for variance $\hat{\sigma}^2$ and $\hat{\theta}_1$ to the maximum likelihood result is the regions in space. While the estimations within the brain are similar, the grid search tends to estimate values near zero in space whereas the maximum likelihood estimates are much noisier across the domain.

The z -statistics produced from the grid search method, presented in Figure C4, are similarly filtered using the Benjamini-Hochberg procedure to control the false discovery rate at the $\alpha = 0.05$ level which resulted in a critical z -statistic of 4.27. The zoomed in region indicates the left motor cortex, the region that is expected to have a task-related signal increase as a result of right-hand finger tapping (the performed task for the experimental data).

When comparing the results from the grid search to the results from the maximum likelihood estimation in Figure 3.8 some interesting conclusions can be drawn. Perhaps a limitation of the search space, the grid search tends to return a value near 0 for $\hat{\sigma}^2$ and $\hat{\theta}_1$ in empty space whereas the maximum likelihood estimates are randomly distributed as one might expect. The grid search does return approximately the same region of task-related phase activation as the maximum likelihood estimate in Figure 3.9 with a slightly higher critical statistic of 4.27, eliminating some noise in the FDR filtered z -statistics (and possibly some task detection as well).




Figure C3: Grid search results for $\hat{\theta}_0$ (a) and $\hat{\sigma}^2$ (b) under the null hypothesis, $\theta_t = \theta_0$ with anatomical image (c). Grid search results for $\hat{\theta}_0$ (d), $\hat{\sigma}^2$ (e), and $\hat{\theta}_1$ (f) under the alternative hypothesis, $\theta_t = \theta_0 + \theta_1 x_t$.




Figure C4: Z-statistics from the likelihood ratio test statistic described in Equation 3.11, filtered by the Benjamini–Hochberg procedure controlling the false discovery rate at the $\alpha = 0.05$ level. The critical z-statistic was found to be 4.27, as indicated by the black lines on the color bar. The left motor cortex is enhanced to the left. This is the region of expected activation for right hand finger tapping experiments.

C.3 Experimental results for all seven slices

As mentioned in Section 3.3.2, the experimental data included seven axial slices. A particularly active slice was chosen for model demonstration in that section. Figure C5 shows a summary of results for all seven slices. The rows are organized by slice; slice 6 was examined in Section 3.3.2. The first column shows the maximum likelihood estimate for ρ for each slice. There is a small amount of Nyquist ghosting present in the magnitude estimation for each slice. The second and third columns show the estimates for θ_0 and θ_1 under the alternative hypothesis. Each baseline phase estimate $\hat{\theta}_0$ has some amount of phase wrapping that is known to cause issues when using simple linear estimation. The task-related phase change estimate $\hat{\theta}_1$ is visually similar for each slice when on the $[-\pi, \pi]$ scale. The fourth column is the variance estimation $\hat{\sigma}^2$. The final column are the Benjamini-Hochberg thresholded z-statistics superimposed onto anatomical magnitude images of the slices. The task-related phase change in the left motor cortex is most noticeable in slices 6 (critical $z = 3.89$) and 7 (3.88), with trace amounts of activation detected in slice 5 (3.90). Slices 1-4 exhibit no obvious signs of a task-related phase change.

Further, as described Section 3.4, the methods in Chapter 3 can be applied to the enhanced data from Chapter 2. Maximum likelihood estimation for phase-only task detection is performed on the MAP estimations formed as described in Chapter 2. Figure C6 shows a summary of results for all seven slices. The rows are organized by slice. The first column shows the maximum likelihood estimate for ρ for each slice. Similarly to the original data, there is a small amount of Nyquist ghosting present in the magnitude estimation for each slice. The magnitude estimation for each slice has, as expected, improved signal and contrast in comparison to the original data. The second and third columns show the estimates for θ_0 and θ_1 under the alternative hypothesis. The task-related phase change estimate $\hat{\theta}_1$ is again visually similar for each slice when on the $[-\pi, \pi]$ scale. However, the MLE estimates $\hat{\theta}_1$ the MAP data tend towards 0 in empty space, rather than uniform noise on $[-\pi, \pi]$. The fourth column is the variance estimation $\hat{\sigma}^2$. Similarly to $\hat{\theta}_1$, the estimates $\hat{\sigma}^2$ for the MAP data tend towards 0 in empty space. This could be a result of the denoising of data that results from the methods described in Chapter 2. The final column are the Benjamini-Hochberg

filtered z -statistics superimposed onto anatomical magnitude images of the slices. The task-related phase change in the left motor cortex is most noticeable in slices 6 (critical $z = 4.18$) and 7 (4.00). Slices 1-5 exhibit no obvious signs of a task-related phase change.

The linear regression estimates for all seven experimental slices are presented in Figure C7. The poor estimation of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ is clear in each slice. The banding artifact is present in each slice, which has severely negative implications when assessing task-active regions. The estimates in Figure C5 that use the Lathi distribution show significant improvements from the Normal assumption estimates. Figure C8 shows the likelihood ratio statistic λ , the χ^2 distributed statistic Λ , and an unfiltered z -statistic for each of the experimental slices. In some slices it is sufficiently clear in the likelihood ratio statistic alone that there is task-related activation in the region of interest (left motor cortex).




Figure C5: Summary images for all seven experimental slices (rows). The columns indicate magnitude estimation $\hat{\rho}$, baseline phase $\hat{\theta}_0$, task-related phase change $\hat{\theta}_1$, and variance $\hat{\sigma}^2$ under the alternative hypothesis, and FDR=0.05 filtered z -statistics. Slice 6 was explored further in Section 3.3.2.




Figure C6: Summary images for MAP estimates of all seven experimental slices (rows). The columns indicate magnitude estimation $\hat{\rho}$, baseline phase $\hat{\theta}_0$, task-related phase change $\hat{\theta}_1$, and variance $\hat{\sigma}^2$ under the alternative hypothesis, and $\alpha = 0.05$ thresholded z -statistics. Slice 6 was explored further in Section 3.4.




Figure C7: The left column is the average phase $\bar{\phi}$ in each voxel for each of the seven experimental slices. The two columns to the left are the linear regression estimates $\hat{\gamma}_0$ and $\hat{\gamma}_1$ that are formed under the Normal assumption. Each slice suffers from phase wrap-around banding artifacts that leads to poor estimation.




Figure C8: Resulting statistics for each of the seven experimental slices. The leftmost column is the likelihood ratio statistic, λ described in Section 3.2.2. This is followed by the the statistic λ which is known to have a χ^2 distribution with one degree of freedom. Lastly is the unfiltered z -statistic given by Equation 3.11.