

**Solving Inverse Problems in Medical Imaging Using
Statistical Signal Processing and Deep Learning**

**A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY**

Seyed Amir Hossein Hosseini

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy**

Prof. Mehmet Akçakaya, advisor

February, 2020

© Seyed Amir Hossein Hosseini 2020
ALL RIGHTS RESERVED

Acknowledgements

I wish to express my sincere gratitude to my advisor Prof. Mehmet Akçakaya for his constant guidance and support through my PhD journey. Prof. Akçakaya helped me to develop my understanding of biomedical signal processing and medical imaging, and to implement innovative ideas by employing the state-of-the-art machine learning tools. Without Prof. Akçakaya's supervision, it would have been infeasible to complete this thesis, especially in a timely manner. I also owe gratitude to Prof. Bin He for providing me with the opportunity to be part of BFINL lab where I developed my initial understanding of biomedical signal processing and medical imaging. I would like to extend my thanks to Prof. Emad Ebbini, Prof. Steen Moeller and Prof. Mingyi Hong for serving on my thesis committee.

I would like to especially thank my dear friends, Burhaneddin Yaman and Abbas Sohrabpour for all their invaluable contribution and kind support. I also wish to thank my other colleagues and friends, Chi Zhang, Burak demirel, Mehran Elyasi, Adel Elmahdy and Bhaskar Sen for their helpful feedback and suggestions. My kind thanks goes to my friends who were like my second family while being away, Alireza Sadeghi, Danial Panahandeh, Siavash Ghavami, Fazel Zare, Ali Moghimi and Mehdi Malmali.

Foremost, I would like to express my deepest appreciations to my parents and brother for their continuous, unparalleled and the most sincere love, inspiration and support. Maman, I have always been thrilled and motivated by your truly sincere and encouraging words of wisdom especially at difficult moments of life. Baba, your patience has always inspired me and your supporting incentives are my primary source of energy and optimism toward my aspirations. Ahmadreza, you are my best friend and the one who I can be myself with. I am so thankful that I have you my ever little brother.

Seyed Amir Hossein Hosseini, Minneapolis, February 2020

Dedication

Simply to my family...

Abstract

Solving inverse problems remain an active research area in various fields to study the cause of a phenomenon by observing the effects. In particular, such efforts are well grounded in medical imaging applications where inverse problems naturally arise due to the imaging target being either inaccessible or invisible to human eyes. Non-regularity or ill-conditioning is a major challenge in such situations which is a direct consequence of limited observations/measurements being available. Medical imaging applications have classically incorporated domain-specific knowledge about the forward encoding operator to regularize the pertaining inverse problem. Deep learning has recently received high interest as an alternative approach for data-driven regularization. In principle, such regularization can be achieved by acquiring more measurements. However, sufficient measurements are usually unavailable due to practical considerations. For instance, in electromagnetic source imaging (ESI) using electroencephalography (EEG), number of recording electrodes can not be excessively increased due to subsequent distortions that electrodes contacts can cause in the medium. This limitation adversely affects the spatial resolution of EEG-based ESI. Another example is magnetic resonance imaging (MRI) where increasing imaging resolution trades off with noise properties, temporal resolution and scan duration. High-resolution MRI can become prohibitively long especially when a sequence of images need to be obtained. Thus, accelerated MRI techniques are required to reconstruct high-resolution images using less data.

The present thesis contributes to solving regularization issues pertaining to inverse problems in medical imaging with a focus on high-resolution electromagnetic source imaging (ESI) using electroencephalography (EEG) and accelerated high resolution magnetic resonance imaging (MRI) with high fidelity. The ESI spatial resolution and the brain full coverage is shown to significantly improve when simultaneous scalp EEG and intracranial EEG recordings are jointly utilized. Additionally, a new ESI technique is proposed that incorporates forward modeling uncertainties into inverse problem formulations, considerably improving localization accuracy and recovering underlying source dynamics and extents when such uncertainties are properly determined and modeled. A

novel deep learning-based reconstruction technique is proposed for accelerated MRI using a recurrent neural network architecture that offers desired noise properties compared to conventional techniques, particularly at higher acceleration rates. Furthermore, a new deep learning methodology is developed that unrolls inverse problem optimization algorithms into dense recurrent neural networks for improved reconstruction performance without requiring extra computational power or increasing reconstruction time.

Contents

| | |
|---|------------|
| Acknowledgements | i |
| Dedication | ii |
| Abstract | iii |
| List of Tables | ix |
| List of Figures | x |
| 1 Introduction | 1 |
| 1.1 Motivation | 1 |
| 1.2 High-Resolution Electromagnetic Source Imaging | 2 |
| 1.3 High-Resolution Magnetic Resonance Imaging | 3 |
| 2 Electromagnetic Source Imaging Using Simultaneous Scalp EEG and Intracranial EEG | 5 |
| 2.1 Introduction | 5 |
| 2.2 Methods | 6 |
| 2.2.1 Electrode Setup | 8 |
| 2.2.2 Source Activity | 8 |
| 2.2.3 Forward Problem | 10 |
| 2.2.4 Inverse Problem | 11 |
| 2.2.5 Source Localization | 12 |
| 2.2.6 Connectivity Analysis | 13 |

| | | |
|----------|--|-----------|
| 2.3 | Results | 14 |
| 2.3.1 | Source Localization | 14 |
| 2.3.2 | Connectivity Analysis | 17 |
| 2.3.3 | DBS Electrode Leads | 18 |
| 2.4 | Discussion | 19 |
| 3 | Electromagnetic Source Imaging by Means of a Robust Minimum Vari- | |
| | ance Beamformer | 39 |
| 3.1 | Introduction | 39 |
| 3.2 | Methods | 41 |
| 3.2.1 | Linearly Constrained Minimum Variance (LCMV) | 41 |
| 3.2.2 | Robust Minimum Variance Beamformer (RMVB) | 44 |
| 3.2.3 | Uncertainty Region Estimation | 45 |
| 3.2.4 | Computer Simulation Protocol | 46 |
| 3.2.5 | Performance Measures | 48 |
| 3.3 | Results | 49 |
| 3.3.1 | Focal Sources | 49 |
| 3.3.2 | The Effect of Number of Active Nodes | 51 |
| 3.3.3 | Extended Sources | 53 |
| 3.4 | Discussion | 55 |
| 4 | Accelerated MRI Reconstruction with Database-Free Deep Learning | 58 |
| 4.1 | Introduction | 58 |
| 4.2 | Methods | 59 |
| 4.2.1 | Calibration | 60 |
| 4.2.2 | Reconstruction | 61 |
| 4.2.3 | Targeted Coronary MRI | 62 |
| 4.2.4 | Image Analysis | 63 |
| 4.2.5 | Whole-Heart Coronary MRI | 63 |
| 4.3 | Results | 64 |
| 4.4 | Discussion | 66 |

| | | |
|----------|---|------------|
| 5 | Database-Free Accelerated MRI Using Recurrent Neural Networks with Densely Connected Blocks | 71 |
| 5.1 | Introduction | 71 |
| 5.2 | Methods | 72 |
| 5.2.1 | Implementation Details | 73 |
| 5.2.2 | In Vivo Coronary MRI | 74 |
| 5.3 | Results | 75 |
| 5.4 | Discussion | 75 |
| 6 | Dense Recurrent Neural Networks for Inverse Problems: History-Cognizant Unrolling of Optimization Algorithms | 78 |
| 6.1 | Introduction | 78 |
| 6.2 | Theory | 80 |
| 6.2.1 | Forward Model and Inverse Problem | 80 |
| 6.2.2 | Physics-Based DL Reconstruction by Unrolling PGD | 81 |
| 6.2.3 | Proposed History-Cognizant PGD Unrolling with Dense Recurrent Neural Networks | 81 |
| 6.3 | Multi-Coil MRI Reconstruction Experiments and Results | 82 |
| 6.3.1 | Knee MRI Datasets | 82 |
| 6.3.2 | Implementation Details | 84 |
| 6.3.3 | Performance of HC-PGD versus PGD Unrolling | 85 |
| 6.3.4 | Insensitivity of Dense-RNN Improvement to the Regularization CNN Choice | 87 |
| 6.3.5 | Performance Evaluation for Random Undersampling | 87 |
| 6.3.6 | Extensions to Other Optimization Algorithms | 87 |
| 6.4 | Discussion | 89 |
| 7 | Conclusion | 97 |
| 7.1 | Thesis Summary | 97 |
| 7.2 | Limitations and Future Directions | 99 |
| | References | 101 |

| | |
|---|------------|
| Appendix A. Supporting Information | 126 |
| A.1 Supporting Information for Chapter 2 | 126 |
| A.2 Supporting Information for Chapter 3 | 126 |
| A.2.1 Convex Optimization | 126 |
| A.2.2 Time-courses of Activity | 127 |
| A.2.3 Practical Considerations | 127 |
| A.2.4 The Effect of Normalization on Bias Removal | 128 |
| A.2.5 The Effect of Correlated Activities | 129 |
| A.2.6 Beamformers with Diagonal Loading | 129 |
| A.3 Supporting Information for Chapter 4 | 130 |

List of Tables

| | | |
|-----|---|----|
| 5.1 | Quantitative NMSE (left) and SSIM (right) assessment of SPIRiT, sRAKI and sRAKI-RNN in reconstructing a right coronary artery MRI dataset at acceleration rates 2, 3, 4 and 5. sRAKI-RNN achieves the least NMSE and the highest SSIM values at all rates. | 77 |
| 6.1 | Average and standard deviation (in ms) of reconstruction time over all test slices for PGD, HC-PGD, VS and HC-VS in Coronal PD and Coronal PD-FS datasets. History cognizant unrolling has very slight impact on reconstruction time compared to its conventional counterpart (< 1.5%). Furthermore, VS methods, which utilize CG-based data consistency units, are considerably slower than PGD methods. | 91 |

List of Figures

| | | |
|-----|---|----|
| 2.1 | An overview of the simulation procedures | 24 |
| 2.2 | The position of the electrodes in the brain for this study. EEG electrodes are on the scalp (not shown here), while sEEG electrodes are inside the brain tissue in the left frontal lobe and part of the left temporal lobe. | 25 |
| 2.3 | The configuration of the simulated networks. (A) The time course assigned to the dipole in single-node sources, (B) the time courses assigned to dipoles in 3-node sources, (C) the connectivity links between nodes recovered by directed transfer function (DTF) at different frequencies in the range 1 – 50 Hz and (D) the recovered connectivity links averaged over all frequencies in the mentioned range. | 26 |
| 2.4 | Median localization error for 3-node network configurations. Error bars refer to the first and third quartiles (of localization error distribution). Two noise types, namely, additive white Gaussian noise (AWGN) and additive colored Gaussian noise (ACGN) with SNRs of 5 dB and 20 dB are simulated. The simulations are repeated for both EEG and sEEG electrodes and with a region of interest (ROI) assumption (refer to section 2 for more details). As it can be seen, the ROI assumption can improve the source localization, however it requires some extra information, which is not generally available. | 27 |

| | | |
|-----|---|----|
| 2.5 | Localization error and histogram of source distance to electrodes for 1-node networks. (A) The histogram of the source-to-nearest electrode distance for 250 single-node sources in the three recording schemes EEG, sEEG and EEG+sEEG. It is obvious that sources are generally further away from EEG electrodes. (B) The median localization error for the same source configurations and for white (AWGN) and colored (ACGN) noise. EEG recordings are subject to the noise with an SNR of 5 dB, while this is 20 dB for sEEG. These values are maintained in the combined version (EEG+sEEG), as well | 28 |
| 2.6 | An example of a 3-node network. (A) The simulated electrical potential recordings for a single 3-node source configuration in both EEG and sEEG setups. sEEG contains stronger spikes than EEG. (B) The accuracy of the source imaging for the same source using EEG+sEEG recordings. It can be seen that the time courses for each active node of the source are recovered precisely. Besides, the connectivity links are identified truly in the sense that source 1 is driving sources 2 and 3, as expected. | 29 |
| 2.7 | .Localization error and histogram of source distance to electrodes for 3-node networks. (A) The histogram of the distances between active nodes and the nearest electrode to them. This figure depicts the results of different electrode set-ups, i.e., EEG, sEEG and EEG+sEEG. (B) The median localization errors using EEG, sEEG and EEG+sEEG recordings under AWGN and ACGN noisy conditions. (C and D) The same plot only for the sources that lie within a distance of 0 – 15 mm and 15 – 30 mm from sEEG electrodes, respectively. | 30 |
| 2.8 | .Localization error versus distance to electrodes for 3-node network configurations. The plots display the median localization errors over five intervals (distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a polynomial of a degree determined by cross validation. | 31 |

| | | |
|------|---|----|
| 2.9 | .Localization error scatter plots. Localization error scatter plots versus distance to electrodes and inter-nodal space for (A) EEG, (B) sEEG and (C) EEG+sEEG recording set-ups under AWGN and ACGN noisy conditions. The horizontal axis and the vertical axis are respectively, the distance to electrodes for every single node and the distance to the nearest other active nodes in the same source configuration (inter-nodal space). The localization errors, the third axis (not shown), are represented with colors. | 32 |
| 2.10 | .Connectivity error and histogram of source distance to electrodes for 3-node networks. (A) The histogram of the average distances between active nodes and the nearest electrode to them. This figure depicts results of different electrode set-ups, i.e., EEG, sEEG and EEG+sEEG. (B) The median connectivity errors using EEG, sEEG and EEG+sEEG recordings under AWGN and ACGN noisy condition. (C and D) The same plot only for the sources that lie within a distance of 0 – 15 mm and 15 – 30 mm from sEEG electrodes respectively. | 33 |
| 2.11 | .Connectivity error versus distance to electrodes for 3-node network configurations. The plots display the median connectivity errors over five intervals (average distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG set-ups and for AWGN and ACGN noise types. In order to make the trends easier to detect, the scattered points are regressed with a line. | 34 |
| 2.12 | .Connectivity error scatter plots. Connectivity error scatter plots versus distance to electrodes and inter-nodal space for (A) EEG, (B) sEEG and (C) EEG+sEEG recording schemes under AWGN and ACGN noisy conditions. The horizontal axis and the vertical axis are respectively, the average distance to electrodes for every single node and the average inter-nodal space in the same source configuration. The connectivity errors, the third axis (not shown), are represented with colors. | 35 |

| | | |
|------|---|----|
| 2.13 | .Localization error versus distance to sEEG electrodes for sources distributed over the entire brain. The plots display the median localization errors over five ranges intervals (distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a line. | 36 |
| 2.14 | .The position of the DBS electrodes in the brain. EEG electrodes are on the scalp (not shown here), while DBS electrodes are implanted inside the brain tissue. The green area shows the subthlamic nucleus. | 37 |
| 2.15 | .Localization error versus distance to DBS electrodes for sources distributed over the entire brain. The plots display the median localization errors over five intervals (distance to the nearest electrode) for (A) EEG, (B) DBS and (C) EEG+DBS schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a line. | 38 |
| 3.1 | The Monte Carlo simulation statistics for the single-node source scenario. (A) The median DLE results for SNRs of 5 dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration. | 50 |
| 3.2 | An example of a 3-node network with focal activity. The source imaging results for a single 3-node source with focal activity using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution is thresholded with a cut-off value of 0.01. The blue rings of the plots mark the positions of the true source. | 51 |

| | | |
|-----|---|----|
| 3.3 | The Monte Carlo simulation statistics for the three-node uncorrelated source scenario. (A) The median DLE results for SNRs of 5 dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration. | 52 |
| 3.4 | An example of a three-node network with extended activity. The simulated three-node patches of activity as well as the source imaging results using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution of each technique is thresholded separately such that MCC is maximized for that technique. The blue regions and boundaries of the plots mark the true source extents. | 53 |
| 3.5 | An example of a three-node network with extended activity. The simulated three-node patches of activity as well as the source imaging results using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution is thresholded with a preset cut-off value of 0.01. The blue regions and boundaries of the plots mark the true source extents. | 54 |
| 3.6 | The Monte Carlo simulation statistics for the extended source scenario. The Average ROC curves for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN and for SNRs of 5 and 20 dB in the left and right panels, respectively. The corresponding AUC of all the curves are reported in the legend. | 55 |
| 3.7 | The Monte Carlo simulation statistics for the extended source scenario. The median MCC results along with the first and third quartile error bars for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN and for SNR of 20 dB. | 56 |

- 4.1 The CNN architecture to learn and enforce the coil self-consistency rule. The number of layer output channels is denoted by depth of blocks. All layers, except the last one, were followed by rectifier linear units (ReLU) as activation functions. The kernel sizes of the layers were 5×5 , 3×3 , 3×3 and 5×5 , respectively. Each layer had 16, 8, 16 and $2n_c$ output channels, respectively. The 3D k-space data was first inverse Fourier transformed along fully-sampled k_x dimension. Subsequently 2D convolutional kernels were jointly trained on the ACS region of resultant 2D slices of data to learn the self-consistency rule. 60
- 4.2 Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were retrospectively under-sampled at rates 2, 3, 4, and 5 in the $k_y - k_z$ plane and then reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI (top, middle and bottom rows). Fully-sampled images are also displayed in the first column as a reference for comparison. sRAKI is visually more robust to noise amplification and blurring artifacts at high acceleration rates compared to SPIRiT and l_1 -SPIRiT, respectively. (RCA: right coronary artery; AO: Aortic Root) 65
- 4.3 Reformatted right coronary artery (RCA) images from another 3D targeted coronary MRI dataset. This data was also retrospectively under-sampled at rates 2, 3, 4, and 5, and fully-sampled images are shown in the first column as reference. The difference between SPIRiT and sRAKI is visually evident at all acceleration rates for this subject with more apparent noise amplification. Furthermore, compared to l_1 -SPIRiT, sRAKI is more robust to blurring artifacts with increasing acceleration rates. (RCA: right coronary artery; AO: Aortic Root) 66

| | | |
|-----|---|----|
| 4.4 | (a) Mean normalized mean squared error (NMSE) and (b) quantitative normalized vessel sharpness measures for all reconstructions of rates 2 to 5. Error bars represent standard deviation across subjects. sRAKI outperforms SPIRiT and l_1 -SPIRiT at all rates for both metrics. The improvements in NMSE are statistically significant at all rates over both SPIRiT and l_1 -SPIRiT, whereas the improvements in vessel sharpness with sRAKI are significant at rates 2 and 3 over SPIRiT, and rates 3 – 5 over l_1 -SPIRiT. Red lines mark significant differences in the graphs. | 67 |
| 4.5 | Reformatted coronal image from a prospectively 5-fold accelerated whole-heart coronary MRI dataset. The results show similar characteristics to targeted coronary MRI, where sRAKI reduces blurring with respect to l_1 -SPIRiT, and noise amplification with respect to SPIRiT. | 68 |
| 5.1 | (a) The recurrent network architecture of sRAKI-RNN unrolled for n iterations to simultaneously apply self-consistency (SC) and data consistency (DC) and perform the iterative reconstruction. (b) A single SC unit with 3 dense blocks (DB), a single convolutional layer (CL) at the output, along with two skip connections to facilitate information flow through network. (c) A closer view of the DB. | 74 |
| 5.2 | A representative slice from a 3D right coronary artery MRI data of a healthy subject. The data were retrospectively undersampled at rates 2, 3, 4 and 5 in the $k_y - k_z$ plane and then reconstructed using techniques SPIRiT, sRAKI and sRAKI-RNN (top, middle and bottom rows). The fully-sampled images are provided in the first column as a reference for comparison. All techniques successfully remove aliasing artifacts, and the least reconstruction noise is achieved by sRAKI-RNN. (RCA: right coronary artery; AO: aortic root) | 76 |

| | | |
|-----|--|----|
| 6.1 | (a) The conventional unrolling for iterative algorithms for inverse problems and (b) the proposed history-cognizant unrolling leading to a dense recurrent neural network architecture. In both cases, the algorithms are unrolled for T iterations, where each unrolled iteration consists of proximal (or regularization) unit (P), and a data consistency unit (DC). The P unit is implemented implicitly via a neural network, while the DC unit corresponds to a gradient descent step for the proximal gradient descent algorithms and typically to conjugate gradient for variable splitting methods. For the proposed history-cognizant unrolling, an additional 1×1 convolutional layer combines all previous proximal operator outputs prior to the DC unit, where C denotes concatenation. | 83 |
| 6.2 | Summary of the training procedure in the proposed history-cognizant unrolling of inverse problem optimization algorithms with a Dense-RNN architecture for accelerated DL MRI reconstruction. Parameters of the unrolled network, which are dominantly characterized by proximal (P) operator units, are updated in each epoch based on a given loss metric using conventional end-to-end training. Data consistency (DC) units utilize the acquired data, coil sensitivity maps and undersampling mask in updates. The uniform undersampling mask utilized for training the network is also depicted. | 85 |
| 6.3 | A representative slice from the Coronal PD dataset with 4-fold acceleration rate using a uniform undersampling pattern, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images with respect to the reference fully-sampled image scaled by a factor of 10. Zerofilled image corresponds to the input to the network, $\mathbf{x}^{(0)}$. The green rectangle in the reference image mark the zoom-up area shown in the second and last rows for each method. Residual aliasing artifacts, which are visible in the PGD unrolling results, have been suppressed using the HC-PGD scheme. Improvements in the quantitative metrics, shown on the lower-left corner of the images, also align with these visual observations. | 92 |

| | | |
|-----|---|----|
| 6.4 | A representative slice from the Coronal PD-FS dataset with a 4-fold acceleration rate using a uniform undersampling pattern, reconstructed with conventional PGD and the proposed HC-PGD methods, along with the error images scaled by a factor of 10. The green rectangle in the reference image mark the zoom-up area. Proposed HC-PGD unrolling successfully removes some residual artifacts that are visible in the PGD. The quantitative metrics align with this observation. | 93 |
| 6.5 | Summary of the quantitative metrics for reconstruction quality, (a) NMSE and (b) SSIM in Coronal PD and Coronal PD-FS datasets using the conventional PGD and the proposed HC-PGD methods. The boxes mark the interquartile range and median of the metrics. The proposed history-cognizant unrolling of the PGD has visibly enhanced reconstruction, based on both metrics and in both datasets. Statistical analysis based on Wilcoxon signed rank test further confirm that all improvements are statistically significant. | 94 |
| 6.6 | Results of changing the proximal operator CNN, with a slice from the Coronal PD dataset with 4-fold uniform undersampling, reconstructed with DL-MRI reconstruction using the PGD and HC-PGD methods unrolled in a network with a U-Net design for the proximal operator units, along with Error images ($\times 5$), with the green rectangle marking the zoom-up area. Similar to the ResNet experiments, the proposed HC-PGD outperforms PGD in suppressing visible residual artifacts. | 94 |
| 6.7 | Results on random undersampling experiments from (a) the Coronal PD and (b) Coronal PD-FS datasets, at a 4-fold acceleration rate, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoom-up area. Regardless of the undersampling pattern, the history-cognizant unrolling approach, HC-PGD suppresses the residual aliasing visible in the conventional PGD unrolling, as marked by the arrows. The quantitative metrics align with these observations. | 95 |

| | | |
|-----|--|-----|
| 6.8 | Example slices from (a) the Coronal PD and (b) Coronal PD-FS datasets with 4-fold uniform undersampling, reconstructed with DL-MRI reconstruction based on the conventional VS and the proposed HC-VS unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoom-up area. The reconstructions from the history-cognizant unrolling approach, HC-VS suppresses the residual aliasing visible in the conventional VS unrolling, as marked by the arrow. Improvements are also observed in the quantitative metrics. | 96 |
| A.1 | Average localization error for 3-node network configurations using three different regularization parameter selection techniques, namely, the l-curve (LCV), the generalized cross validation (GCV) and the discrepancy principle (DCP). Two noise types, namely, additive white Gaussian noise (AWGN) and additive colored Gaussian noise (ACGN) with SNRs of 5 dB and 20 dB are simulated. The simulations are repeated for both EEG and sEEG modalities. | 131 |
| A.2 | The source space from which source locations are selected. The red spheres in this figure show the sEEG electrodes and blue spheres show the source space from which source locations are selected. As it can be seen, while sources are confined to the cortex, they can be very deep. . . | 131 |
| A.3 | The time-course assigned to each dipole in the single-node source scenario. | 132 |
| A.4 | The time-courses and correlation diagrams in the second and third scenarios. (A) The time-courses assigned to each node in the 3-node uncorrelated source scenario and (B) the corresponding correlation coefficients among the nodes. As it can be seen, the correlation coefficient between different nodes is small in this scenario. (C) The time-courses assigned to each node in the 3-node correlated source scenario and (D) the corresponding correlation coefficients among the nodes. Unlike the second scenario, nodes are highly correlated in this case. | 133 |

| | | |
|------|---|-----|
| A.5 | The histogram of the parameter β and its correlation with angle standard deviation within a region. (A) Plots the histogram of the parameter β and provides information regarding the distribution of the mentioned parameter for all voxels. Red vertical line marks the position of the cut-off value in this distribution. (B) Parameter β of all voxels versus the standard deviation (STDev) of the angle between the center and the neighbors' orientations in each uncertainty region. | 134 |
| A.6 | The DLE results versus depth for the three-node uncorrelated source scenario. The plots display the median DLE along with the error bars over five depth intervals for (A) LCMV, (B) RMVB, (C) LCMV-ND-DN and (D) RMVB-ND-DN. SNR in this figure is set to 5dB and depth is defined as the distance between each active node in a source configuration and the nearest electrode in a standard 128-channel BioSemi cap. | 135 |
| A.7 | The color-coded depth of different locations in the brain. Depth in this diagram is defined as the distance to the nearest electrode in a standard 128-channel BioSemi cap. | 136 |
| A.8 | The Monte Carlo simulation statistics for the three-node correlated source scenario. (A) The median DLE results for SNRs of 5dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration. | 136 |
| A.9 | The Monte Carlo simulation statistics for the three-node uncorrelated source scenario. (A) The median DLE results for SNRs of 5dB and 20dB and for diagonal loading beamformer with three factors of σ^2 , $10\sigma^2$ and $100\sigma^2$. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration. | 137 |
| A.10 | The Monte Carlo simulation statistics for the extended source scenario. The median MCC results along with the first and third quartile error bars for diagonal loading beamformer with three factors of σ^2 , $10\sigma^2$ and $100\sigma^2$ and for SNR of 20 dB. | 137 |

A.11 Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were uniformly undersampled retrospectively at rates 2×2 , 3×2 , 4×2 and 5×2 in $k_y - k_z$ plane, which are approximately equivalent to net acceleration rates 4, 5, 6 and 7 (including the ACS lines and an elliptical mask). These data were then reconstructed using GRAPPA, SPIRiT, l_1 -SPIRiT, RAKI and sRAKI (from top to bottom). Acceleration rate was set no higher than 2 for k_z dimension, since the size of data along this dimension was small (20 lines in total and 10 lines for ACS). For RAKI, a 3-layer network was designed with a kernel size of 2×2 (with dilations equaling acceleration rates to match the undersampled uniform pattern) for the first layer and a kernel size of 1×1 for subsequent layers. Note that this 2D undersampling is different from the original RAKI paper, thus the network architecture may be sub-optimal. The learning rate and number of epochs for RAKI were tuned to 0.05 and 2000 iterations, respectively. Fully-sampled images are also displayed in the first column as a reference for comparison. While RAKI is robust, GRAPPA is very sensitive to noise with increasing rates. In addition, RAKI outperforms SPIRiT, but RAKI and sRAKI perform comparatively, similar to the relationship between GRAPPA and SPIRiT. 138

A.12 Three $k_y - k_z$ undersampling patterns were tested. Poisson disc (top), uniform-density random (middle) and variable-density random (bottom) with 4-fold acceleration. 139

A.13 Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were retrospectively undersampled with the three different patterns shown in Supporting **Figure A.12**. These data were then reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI. The results show that sRAKI is more resilient to noise amplification compared with SPIRiT, regardless of undersampling pattern. 140

| | | |
|------|---|-----|
| A.14 | Noise sensitivity of the reconstruction methods are shown on reformatted right coronary artery images. Additive Gaussian noise was added to the datasets retrospectively, and the reported SNR was measured at aorta (signal power at aorta divided by noise power in an empty region of interest), with the original dataset having SNR of 50. The datasets were then retrospectively undersampled at rate 4 and reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI. sRAKI is more robust against noise of data compared with SPIRiT. However, noise amplification becomes evident with increasing levels of noise compared with l_1 -SPIRiT. . . . | 141 |
| A.15 | Normalized mean squared error (NMSE) of reconstruction across all subjects with 4-fold acceleration for the experiment setup described in Figure A.14 . Error bars represent standard deviation across subjects. . . . | 142 |

Chapter 1

Introduction

1.1 Motivation

Medical imaging seeks to provide visual representation of some desired aspects in internal body organs for clinical analysis or medical intervention. On a high-level classification, medical imaging can be divided into two major categories; functional imaging and structural imaging. Functional imaging aims to study functional properties of a living system such as the brain by monitoring its physiological activities. In functional brain imaging, function and dysfunction of the brain are monitored by imaging its functional dynamics over time. Such modalities include functional magnetic resonance imaging (fMRI) [1–3], positron emission tomography (PET) [4], electro-encephalography (EEG) [5–7], and magneto-encephalography (MEG) [8, 9]. Structural imaging, on the other hand, is more concerned with anatomical properties of a certain tissue or organ. Magnetic resonance imaging (MRI) [10–13] and computed tomography (CT) [14, 15] are two common structural imaging modalities. The present thesis aims to address challenges pertaining to inverse problems in high-resolution imaging with a focus on two medical imaging modalities namely, EEG (or similarly MEG) and MRI (or similarly CT) by employing statistical signal processing and the state-of-the-art deep learning methods.

1.2 High-Resolution Electromagnetic Source Imaging

Among functional brain imaging modalities, EEG is typically noninvasive, inexpensive, easy to set up and is readily available in most clinical settings [7]. Such properties make it an appropriate tool for clinical and research studies on human brain. Particularly, the high temporal resolution of EEG, makes it well suited to study brain dynamics [16, 17]. This is essential in studying the brain where the physiological states can vary at a high pace. Additionally, EEG directly measures the instantaneous potentials generated at the scalp due to neuronal excitation inside the brain, while for example fMRI detects the magnetization caused by changes in blood flow and blood oxygenation; an indirect consequence of neuronal activity [17, 18]. In contrast to these advantages, EEG has historically provided low-resolution and coarse information about the brain activities e.g. the brain hemisphere that is responsible for these activities. EEG has also been used to locally image regions of the brain that are close to electrodes. For high-resolution and global imaging of the brain using EEG, electromagnetic source imaging (ESI) has been developed. ESI projects the EEG (or MEG) signals back onto the source space, where the EEG signals are originated from [5, 6, 18, 19]. In other words, ESI attempts to find the underlying brain electrical activities using EEG (or MEG) signals, rather than only perceiving and interpreting the recordings over the sensor space, i.e. surface scalp recordings, with the purpose of improving imaging resolution.

Over the past decades, a number of approaches have been developed for high-resolution ESI. These methods show promises of more accurate localization and imaging of the brain activity and connectivity [18, 20]. However, these methods face challenges in localizing deep brain activities compared to sources that are more superficial. In such cases, recordings from intracranial EEG electrodes (iEEG), implanted in brain tissue, may provide more information about deeper sources [21]. More recently, implantation of deep electrodes in the brain for Parkinson's disease [22], as well as the responsive cortical stimulation systems for controlling medically intractable epilepsy [23] indicate that even long term implantation of such deep and cortical electrodes is clinically applicable. Deep brain stimulation (DBS) electrodes, which were once used for stimulation only, are now capable of recording and stimulating concurrently [24]. Stereo electroencephalography (sEEG) [25] is another example where the electrodes are placed in deep

and far-from-scalp brain tissues with the purpose of sensing and monitoring deeper brain activity. Chapter 2 investigates the performance, merits and limitations of source imaging using such intracranial recordings, individually or along with surface recordings on the scalp. Although intracranial recordings can provide auxiliary information for improved high-resolution ESI, invasive implementation of electrodes is not desired or even possible in many cases. In Chapter 3, we introduce a new technique that can improve high-resolution ESI without necessitating deep recordings from intracranial electrodes.

1.3 High-Resolution Magnetic Resonance Imaging

MRI is a noninvasive imaging technology with a wide range of applications in medical imaging, for example, in the diagnosis of coronary artery disease (CAD) which is the leading cause of death in the United States, accounting for one in seven deaths [26]. Coronary MRI provides a radiation-free diagnostic tool for CAD assessment [27], with a potential for repeated imaging. Nonetheless, data acquisition process in MRI is very slow which is imposed by the physics behind MRI. In order to image a desired volume, data acquisition process needs to be sequentially repeated, leading to a slow imaging modality. This problem becomes more pronounced in applications where a reasonably high-resolution image is required to capture crucial details of anatomy or where a sequence of images need to be acquired for future analysis. Numerous studies have attempted to accelerate MRI by initially acquiring less data from MRI scanners. These data, known as k-space, are originally in Fourier domain, requiring inverse Fourier transformation before reconstructing the final image. Parallel imaging with multi-coil data [28, 29] and compressed sensing [30, 31] are two major conventional approaches to reconstruct the image from under-sampled k-space data.

Recently, deep learning-based techniques [32–46] have also gained attention as a means to accelerate MRI acquisition. Numerous studies have designed neural network architectures that either establish an end-to-end nonlinear mapping from under-sampled k-space/distorted image to full k-space/undistorted image [33, 35, 36, 39, 41–43, 45] or decompose an iterative optimization problem into (recurrent) deep learning blocks that learn a data-specific regularization [34, 37, 38, 40, 46]. A number of these studies also show support for parallel imaging with multi-coil data [32, 34, 37, 39, 44]. While these

studies show promising results in accelerated MRI, there are limitations regarding the training phase of reconstruction. Primarily, large datasets are required for training the neural networks, which are not readily available in all situations. This challenge has been partially addressed by transfer learning approaches, which pre-train neural networks on available large datasets and then fine-tune them on smaller datasets of specific applications [36, 47]. However, transfer learning still requires training on fully-sampled data. Acquisition of fully-sampled training data in some applications e.g., in whole-heart coronary MRI, may be infeasible, since the scan time would become prohibitively long. Furthermore, training datasets may not include all pathologies of interest, which may lead to risks in generalizability for diagnosis [48]. These obstacles may hinder the clinical application of current transfer learning-based techniques to high-resolution MRI [48].

In Chapter 4, we introduce a scan-specific and database-free accelerated MRI reconstruction technique based on deep learning that trains neural networks using data from the same scan only. This technique does not require ground-truth fully-sampled data during training. An extension of this technique with higher effectiveness and efficiency in reconstructing under-sampled MRI data is introduced in Chapter 5. Our focus in these two chapters is on accelerating coronary MRI. However, the proposed techniques can essentially be applied to other MRI data as well. Finally, in Chapter 6, we introduce a particular neural network design for solving inverse problems of imaging applications using deep learning in general. We focus on accelerated MRI reconstruction as a specific application for this method and show that database-trained deep learning MRI reconstruction can be improved by utilizing this neural network design.

Chapter 2

Electromagnetic Source Imaging Using Simultaneous Scalp EEG and Intracranial EEG

2.1 Introduction

About 30% of the patients with epilepsy do not respond to any medication treatment [49–52]. These patients are potential candidates of receiving surgery for the resection of the epileptogenic zone with the hope of becoming seizure-free [53]. To minimize the risks and side effects of the resection and to improve surgical outcomes at the same time, clinicians need to come to certain and decisive conclusions about the exact origins and sites of the epileptic activity in the brain. To this end, patients might undergo multiple procedures before receiving the surgery. Occasionally and depending on the patient’s conditions, noninvasive imaging techniques may fail to achieve any conclusive results about the epileptogenic zone [54]. In such conditions, clinicians may opt for more invasive techniques; for example by implanting sEEG electrodes directly into the regions where they assume epileptic sources are situated with high probability [55,56]. It should be emphasized that in clinical practice, iEEG recordings are generally used as mapping devices where the activity recorded at an electrode is assumed to have been generated from the vicinity of the recorded electrode [55–57]. In reality, this might not be the case

and iEEG recordings can be due to sources located at further distances that will affect the electrode recordings through volume conduction. Thus applying ESI at recordings from iEEG, could potentially shed some light into the matter. There are some works in the literature where possibility of source imaging using intracranial recordings has been demonstrated [21, 58–64].

The goal of this chapter is to investigate the possibility, merits and limitations of source imaging using iEEG recordings (with a focus on sEEG and DBS lead recordings) in comprehensive computer simulations. To this end, we exploit distributed source models to solve the inverse problem, subsequently to determine the source locations, and to identify the inter-nodal connectivity associated with underlying brain networks. We conduct systematic simulations to investigate the efficacy of ESI techniques applied to iEEG and combined EEG and iEEG recordings to determine underlying brain networks (location of network nodes and their inter-nodal links and connections). Specifically, the network geometry, i.e. the distance of sources (nodes of the network) from electrodes and distance of nodes from each other, is studied in detail to determine the accuracy of estimated underlying brain networks. In brief, we investigate source imaging based on surface and intracranial measurements, individually or simultaneously.

2.2 Methods

Computer simulations were conducted to test the hypothesis that applying ESI to intracranial recordings will produce more accurate results for deep brain activity. To give an overview of the whole procedure, we start by explaining how the EEG and sEEG electrodes were positioned; a standard 128-channel EEG cap was used for simulating scalp recordings. The head model used in this study was the Montreal Neurological Institute Colin brain [65]. Ten electrode shanks each containing 13 electrode contacts were manually positioned in regions located in the left frontal and temporal lobes. The positioning of these electrodes was based on a real epileptic patient undergoing a sEEG study, adapted from a study performed by Nobili and colleagues [66]. To model the head volume conductor, a boundary element method (BEM) model [5, 67] was used for EEG simulations, consisting of three layers modelling the scalp, the skull and the brain with electrical conductivities of 0.33 S/m, 0.0165 S/m and 0.33 S/m, respectively [68–70].

An infinitely homogeneous model was used for sEEG simulations with a connectivity of 0.33 S/m. Computer simulations were performed using 500 different cortical source configurations. Two network configurations were modeled and tested in these simulations. One configuration consisted of networks with a single node (250 cases) and another configuration where networks had three nodes (250 cases). The simulated sources were modeled as point dipoles and placed randomly on the cortex with a fixed orientation, perpendicular to the cortical surface. The activities of the nodes in the 3-node source configuration were co-dependent and one node was chosen as the central node to drive and control the activities of the other two nodes through a multi-variate autoregressive (MVAR) model (2.1). The forward problem was then solved to obtain the potentials sensed by the scalp EEG and iEEG electrodes. In order to study the effect of noise on the process of estimating underlying brain networks, two noise models were used; an additive white Gaussian noise or an additive colored Gaussian noise. Different grid sizes were used for modeling the forward and inverse lead field matrices (1 mm and 2.5 mm equivalent to a total of $\sim 131,000$ and $\sim 28,000$ voxels, respectively). The noisy maps were then used to solve the inverse problem using the standardized low resolution electromagnetic tomography (sLORETA) [71] to estimate the underlying activities. The results of the inverse problem which are spatio-temporal distributions showing the activity of different brain regions over time, were then processed using a principal component analysis (PCA) technique to localize major nodes of activity among all source locations (to identify the location of the single node or the three nodes within the underlying networks). The PCA algorithm is efficient in separating incoherent components [72]. While the simulated sources are constructed through an MVAR model, the spatial separation of the simulated networks creates a distinct spatio-temporal pattern in the dipole source, which enables the PCA to determine the nodes of the network efficiently [73]. The location of the recovered major nodes of activity is then compared to the simulated nodes. A causality and connectivity analysis was further conducted on the time course of activity extracted from the estimated network nodes to determine how they are interconnected and driving each other. To this end, the dimension of the problem was primarily reduced to only include node locations estimated from the PCA algorithm to only extract time course activity of the major nodes. Then, a further step was taken to solve the inverse problem again using a least square (LS) operator on the

mentioned reduced size problem. The goal of the aforementioned step was to generate more accurate time course of the underlying activities, to later feed as input into the connectivity analysis [73]. A measure of connectivity error was also used (detailed explanations to follow) to make a comparison between recovered inter-nodal connectivity and those of the original simulated network. **Figure 2.1** provides an overview of the simulation procedure.

2.2.1 Electrode Setup

The EEG electrodes were positioned on the scalp based on a 128-channel configuration (similar to a BioSemi electrode cap). Unlike scalp EEG, sEEG does not benefit from any standardized technique to implant the electrodes in the brain tissue as it varies from case to case and patient to patient. Thus, electrode placement for sEEG is individualized and subject-dependent. Implantation strategy involves the pre-implantation hypothesis that takes into consideration the visible lesions in patient’s anatomical images such as magnetic resonance imaging (MRI) images, the more likely structure(s) of ictal onset, the early and late spread regions and the interactions with the functional networks [74]. Therefore, an adequate number of recording channels is required to have an informative sEEG system. On the other hand, excessive number of electrodes could increase the risks and expenses associated with implantation. It is usually rare to have more than 15 deep shanks implanted in the brain tissue [74]. Based on these considerations, we have simulated 10 sEEG electrode shanks mostly in the left frontal lobe and part of the left temporal lobe, following an example presented in [66]. Each shank consists of 13 contacts with a constant inter-contact distance of 3.5 mm. There is a typical distance of 3.5 mm between the centers of adjacent contacts on a sEEG electrode shank, which justifies our choice of electrode spacing. **Figure 2.2** illustrates the sEEG electrode shanks used in the present simulation study.

2.2.2 Source Activity

The location of the aforementioned sources was chosen randomly to be on the cortex within a distance of 3 to 30 mm from sEEG electrodes. The rationale behind this simulation setting is that sEEG electrodes are inserted in the areas with a higher probability

of activities and thus are closer to underlying sources. It should be noted that with this choice of source location, the simulated sources can still be situated in deep regions such as the thalamus. To mimic the behavior of a realistic inter-ictal spike, a time course activity was constructed by summing eight Gaussian time courses (see **Figure 2.3 (A)**) with an equivalent sampling frequency of 400 Hz. Although the assumption of a single dominant source was suggested to be reasonable in modeling epileptic activity in some cases [75], it may not be enough for all epilepsy cases. In a second attempt, 250 three-node source configurations were simulated to more realistically model epileptic activity. To this end, the general criteria of setting locations, orientations and frequency remained the same as the single-node sources. The time course of activity assigned to each node (see **Figure 2.3 (B)**) was generated through the MVAR model presented here:

$$\begin{aligned} x_1(t) &= \phi(t) + n_1(t), \\ x_2(t) &= 0.35x_1(t - 0.100) + 0.25x_2(t - 0.0075) + 0.1x_2(t - 0.0125) + n_2(t), \\ x_3(t) &= 0.60x_1(t - 0.200) + 0.10x_3(t - 0.0125) + n_3(t). \end{aligned} \quad (2.1)$$

In the model above, the first node is the central node, which controls the activity of the other two nodes, mimicking seizure onset in focal epilepsy. The activity of this node is driven by the Gaussian waveform ($\phi(t)$) which models the inter-ictal spike activity. The directed transfer function (DTF) values of the estimated networks were then calculated to perform the Granger causality analysis [76–79]. The statistical significance of the obtained DTF values was validated by a surrogating scheme, where surrogate time series with random phase were generated based on the estimated time series [80,81]. The DTF values were recalculated from these surrogate time series and served as a comparison to determine if the original calculated DTF values are statistically significant or not [76, 82, 83]. As it can be observed in **Figure 2.3 (C)**, the strength of connectivity links changes as a function of frequency depending on spectral properties determined by the MVAR model [79]. The outputs of the connectivity analysis can be averaged over different frequency bands to obtain a two-dimensional connectivity matrix as in **Figure 2.3 (D)**. It should be mentioned that while MVAR parameters dictate how sources are activated and consequently how they interact with each other, we have observed that slightly changing the parameters would not change the whole connectivity links

significantly. The methodological framework could be applied to any MVAR model regardless, and qualitatively no major difference is expected.

2.2.3 Forward Problem

Neuronal activities inside the brain can be modeled as current densities [18, 84]. According to this model, a dipole is assigned to each voxel with components along different Cartesian axes \hat{x} , \hat{y} and \hat{z} . The forward problem establishes the relation between the dipole current density and the potentials at the electrodes. It can be reasonably assumed that this relation is instantaneous and linear as the Maxwell's equations are being solved in a quasi-static regime [85]. More specifically, it can be shown that

$$\Phi_{n_e \times 1} = \mathbf{K}_{n_e \times (3n_v)} \times \mathbf{J}_{(3n_v) \times 1} + \mathbf{N}_{n_e \times 1}, \quad (2.2)$$

where n_e and n_v are the number of electrodes and voxels respectively, $\mathbf{J}_{(3n_v) \times 1}$ is the current density vector at each voxel, $\mathbf{K}_{n_e \times (3n_v)}$ is called the lead field matrix, $\Phi_{n_e \times 1}$ is the vector of generated or recorded potentials and $\mathbf{N}_{n_e \times 1}$ models measurement noise. In this study, two noise models are considered for both scalp and intracranial measurements. First, additive white Gaussian noise (AWGN) which models measurement noise and additive colored Gaussian noise (ACGN) which models the background neuronal activities as well as the measurement noise. To generate ACGN, a source of white Gaussian noise was assumed at every voxel, which was then projected onto the electrodes (by solving the forward problem) to form a correlated noise at the electrodes, which we refer to as the colored noise. Both the additive white noise at electrode level and the background source-space noise were used (with equal power at the electrodes) to create the ACGN noise model. The obtained noise (in AWGN or ACGN) was then scaled to get either a 5 or 20 dB signal to noise ratio (SNR) at the electrodes. SNR in this study is defined as the ratio of the average signal power (over channels and over time) over average noise power (expressed in a log-scale, hence the dB). It should also be emphasized that in the figures where the results of combined EEG and sEEG electrodes are presented, the same background activity is used (to model background noise in sEEG and EEG). In the combined sEEG and EEG simulations, the SNR level at the sEEG was set to 20 dB (the white and colored noise were present at the sEEG electrodes with equal power). For EEG, however, white noise was added in such a manner to achieve

a total SNR of 5 dB, since the same background noise was used for EEG and sEEG (thus the power of white and colored noise present in EEG electrodes were not equal for combined simulation cases).

2.2.4 Inverse Problem

The term ESI is used to denote the procedure of recovering underlying source activities from sets of electromagnetic measurements. In other words, ESI aims at finding current density vector \mathbf{J} by using the potential vector Φ . Since the number of the knowns in the inverse problem is generally much less than the number of the unknowns (less electrodes than dipoles), ESI is an underdetermined inverse problem, and has to be regularized [84]. Regularization is imposing an extra condition over the solution, which enforces some constraints on the solution based on some a priori information. In this study, a Tikhonov regularization technique was used to regularize the inverse problem [84,86,87]. Mathematically speaking, the following optimization problem was solved to obtain the solution:

$$\hat{\mathbf{J}} = \arg \min_{\mathbf{J}} (\Phi - \mathbf{KJ})^T \Sigma^{-1} (\Phi - \mathbf{KJ}) + \lambda \|\mathbf{J}\|_2^2, \quad (2.3)$$

where Σ is the covariance matrix of the noise estimated from baseline or noise-dominated segments, $(\Phi - \mathbf{KJ})^T \Sigma^{-1} (\Phi - \mathbf{KJ})$ is called the fitting term and determines how well the obtained solution can generate the recorded potentials, λ is the regularization parameter and $\|\mathbf{J}\|_2^2$ is the regularization term. To obtain more stable solutions, only the diagonal elements of Σ were kept. The covariance matrix can also be approximated with a unity matrix of appropriate size, if the baseline is not available [88]. However, when recordings are contaminated with different levels of noise (similar to the case of combined EEG and iEEG), estimating Σ from baseline improves results significantly. This fact can be demonstrated in the following formulation:

$$(\Phi - \mathbf{KJ})^T \Sigma^{-1} (\Phi - \mathbf{KJ}) = \|\Sigma^{-\frac{1}{2}} (\Phi - \mathbf{KJ})\|_2^2 = \|\Sigma^{-\frac{1}{2}} \mathbf{N}\|_2^2 \quad (2.4)$$

where \mathbf{N} is the noise vector. In other words, the matrix $\Sigma^{-\frac{1}{2}}$ is used to whiten the noise. The inverse problem is then solved using the covariance matrix of the whitened noise, which is now a unity matrix. The regularization parameter λ makes a balance between our a priori knowledge about the solution (regularization term) and what measurements bring in (fitting term). There are various types of regularization parameter

selection techniques such as the L-curve (LCV), generalized cross validation (GCV) and discrepancy principle (DCP) [84, 87]. Unfortunately, a single parameter selection technique with universally optimal performance does not exist. A selected technique’s performance may change from problem to problem and from condition to condition [86]. We have observed that for the problem in this study, DCP outperforms the other two methods (see **Figure A.1**), thus, it was used as a method for selecting λ in the rest of this study. Note that λ could be estimated either at each time point or once for an interval of interest. Since the former is more time consuming, λ was estimated based on the time point at which maximum SNR was achieved (at peak time). After current density vectors at different time points were obtained, a further step was taken and the solution was standardized according to the covariance matrix of the estimated current density vector [71]. This technique is known as sLORETA in the literature. Finally, it is necessary to mention that to accelerate the recovery procedure, only the initial 200 ms of the solutions were used to localize the network nodes.

2.2.5 Source Localization

Although monitoring the dynamics of the whole brain can be useful in understanding its networks, finding major active regions in the brain provides valuable information. To find these foci of activity, we followed a procedure based on PCA [73]. The outputs of the inverse algorithm were input to the PCA unit. After taking the principal components, only one component was kept in single-node sources and three components were kept in 3-node sources. This can be justified based on the fact that the singular values of the recordings generally demonstrate either one or three dominant singular values corresponding to the two simulated network configurations. The source locations with locally higher activities in the selected principal components are reasonable candidates for dominant sources, and thus the network node. Once these local maxima were found, the algorithm did not look for any other local maxima near such locations. The reason for such a strategy is twofold. Firstly, the activity of a single active node can leak to multiple components (in the PCA), and secondly, at each component the location might be jittered to slightly different positions in the vicinity of the local maxima while still pertaining to the same location. In this study, a radius of 12.5 mm was chosen as the threshold (to avoid selecting the same nodes multiple times in different components).

The recovered nodes were then matched up with original ones using the Hungarian assignment algorithm [89]. Consequently, the localization errors were calculated in terms of Euclidian distance, based on this matching.

2.2.6 Connectivity Analysis

In order to investigate the ability of the imaging algorithm in identifying brain networks, a connectivity analysis technique was applied to the temporal activities of dominant nodes [90]. In order to further improve the accuracy of the estimated time courses, the inverse problem was solved again (after the initial attempt, which determines nodes of major activity) only for the dipoles located at the identified nodes (from the PCA process). Effectively, such a procedure downsizes the problem to estimating activity at nodes of interest and major activity. The results of the inverse problem are basically dipole moments (volume vector) along the \hat{x} , \hat{y} and \hat{z} axis. To use them as inputs for connectivity analysis, one can simply calculate the amplitude of each dipole (norm 2 of the components along each axis). However, the information about the polarity will be lost in this manner (dipoles pointing inward and outward). To maintain the aforementioned information, which is particularly important in connectivity assessment using Granger causality, and to make the time series more compatible with the true time series (simulated), the polarity information was included in the amplitude. This was achieved by considering the angle between the dipole's orientation at each time point and its orientation at the time of the maximum activity. In simple words, the orientation of the dipole at peak time was set as the reference. Whenever the angle between a node's orientation and its reference orientation rose to more than 90 degrees, a negative polarity was assigned to its amplitude at that time point and a positive polarity otherwise. The revised time series were then used for connectivity analysis and calculation of DTF values. The Frobenius matrix norm of the difference between the recovered and the original connectivity matrices was calculated to measure the degree of similarity between estimated and simulated networks.

2.3 Results

Figure 2.4 depicts the median localization error from EEG and sEEG for the 3-node source configurations. While we could use standard deviations to present the uncertainties, we decided to report the second and third quartiles of data as error bars, since it is more consistent with our report of median for localization error. When solving for the inverse solution two approaches were considered. First, all the grids on the cortex were considered equally probable candidates of being active nodes. In the second scenario, we used only the cortical locations lying in the space residing within the distance of 3 to 30 mm away from sEEG electrodes. This is the space, from which original source nodes were randomly chosen. Although the second scenario, which is denoted by ROI (region of interest) in **Figure 2.4**, may result in a more accurate source localization, we do not report the results of this method in the figures presented afterwards, as it was observed that such an approach did not improve results much. Additionally, such an a priori knowledge is not generally available and one of the purposes of ESI is to estimate source location (not to assume it). **Figure 2.4** also suggests that imaging using EEG recordings is more accurate than sEEG. Although this is generally true, we will later show that under some conditions, sEEG can do better than EEG. However, a combination of sEEG and EEG will always outperform using each modality, individually.

2.3.1 Source Localization

In **Figure 2.5**, localization error for 250 single-node network configurations using EEG only electrodes, sEEG only electrodes and a combination of both (denoted as EEG+sEEG) is compared. **Figure 2.5 (A)** demonstrates the histogram of minimum source-to-electrode distance in each modality. In other words, while each source is in a certain distance from each electrode, we report the distance to the nearest electrode. For the combined EEG+sEEG scheme, we have both scalp and intracranial electrodes in the brain. Since scalp EEG electrodes are generally further away from sources, we expect the nearest electrodes to be among the intracranial electrodes. Thus, the minimum distance range (or the horizontal axis values) are very similar to the sEEG scheme. In

Figure 2.5 (B), the median localization error (with the first and third quartiles presented as error bars) are reported. Since signals need to pass through skull and skin to reach the scalp EEG electrodes, it is expected for these signals to be subject to more severe noise (in terms of SNR) compared to sEEG. Therefore, in **Figure 2.5 (B)** and all the figures hereafter, we have assumed an SNR of 5 dB for EEG signals recorded at scalp electrodes and 20 dB for sEEG signals recorded at deep electrodes. It should be emphasized that these settings were also applied to the EEG+sEEG setup. Besides, to generate ACGN, the same background activity has been simulated for both EEG and sEEG in all the figures to follow. As it can be seen in **Figure 2.5 (B)**, small error was perceived for EEG with greater variance, though, it should be mentioned that the sources were further away from EEG electrodes than sEEG electrodes and thus the signals recorded at the EEG electrodes were inherently weaker. However, source localization is generally less sensitive to the noise (including its level and type) in this case. This is due to the fact that single sources of activity can be detected pretty well and some previous literature [71,91] have claimed almost perfect results as also confirmed in this work.

Figure 2.6 (A) shows the simulated electrical potential recordings for a single 3-node source configuration in both EEG and sEEG setups. As the figure indicates, there are stronger spikes in sEEG recordings. This is because the sources can be close to the electrodes in sEEG. Although this is an advantage for closer nodes, it can mask the activity of nodes situated further away from sEEG electrodes (this will be further discussed in section 2.4). In **Figure 2.6 (B)**, the recovered activity using EEG+sEEG recordings is shown. By comparing the recovered time-courses to the simulated time-courses in **Figure 2.3 (B)**, it can be observed that the time courses for each active node are recovered accurately. Besides, the connectivity links estimated are identical to simulated networks, in the presented example (compare these to original connectivity links presented in **Figure 2.3 (D)**).

Figure 2.7 provides a statistical inference of the localization error with respect to distance, for the three node network configurations (250 cases). In **Figure 2.7 (A)**, the histogram of the distances between active nodes and the nearest electrode to each source is shown. This figure indicates how nodes are distributed in each different electrode setup, i.e. EEG, sEEG or EEG+sEEG. **Figure 2.7 (B)** plots the median localization

errors using EEG, sEEG and EEG+sEEG recordings under AWGN and ACGN noise conditions. While the colored noise models can generally impact source imaging more negatively, this plot implies that once noise covariance is dealt with appropriately, results of ACGN contaminated source imaging is as good as AWGN contaminated source imaging results. This plot also indicates that in spite of nodes being closer to sEEG electrodes, EEG-based localization yields even more accuracy generally. However, this is not true for the sufficiently close sources (within ~ 15 mm of sEEG electrodes). The next two plots (C) and (D) show that for the sources that lie within the distance of 0-15 mm from an sEEG electrode, using sEEG recordings can reduce the localization error significantly, as compared to EEG. However, for the sources located more than 15 mm away, the localization error increases very steeply. Note that distances reported in these two plots, are based on sEEG electrodes not EEG (the distance between sources and the closest sEEG electrode). Besides, it can be seen that combining EEG and sEEG reduces the localization error significantly in all cases.

Figure 2.8 depicts the localization error as a function of distance for the same data as in **Figure 2.7**. The plots display the median localization errors over five intervals (distance to the nearest electrode) for EEG, sEEG and EEG+sEEG schemes and for AWGN and ACGN noise types. In order to observe the trend more clearly, the scattered points are regressed with a polynomial of a degree determined using cross validation. Multiple observations can be made; first, the nodes are generally further away from EEG electrodes. This is a unique feature that sEEG electrodes have, as they are placed deep in the brain, supposedly closer to underlying sources. The second observation is that, source localization using sEEG recordings can be more accurate than EEG-based localization, if the active nodes are close enough to the electrodes (< 15 mm). The bar plots and the fitted regression lines for sEEG suggest that localization error grows very fast for nodes situated further away (> 15 mm). This demonstrates the high sensitivity of the sEEG-based imaging to source distance from electrode location. The third and the most important observation of this figure is that combining EEG and sEEG (EEG+sEEG) improves the source localization accuracy significantly on one hand and reduces the sensitivity to distance on the other. It seems that by combining EEG and sEEG, one can make a recording paradigm which benefits from the strengths of both modalities.

Generally, if the active nodes of a source configuration are closer to each other, their activities are more likely to interfere with each other (orientation of sources also plays an important role). To assess the effect of distance from electrodes and inter-nodal space on the localization accuracy, **Figure 2.9** depicts a set of 3D scatter plots. The horizontal axis in all the subplots of this figure is the distance from electrodes, while the vertical axis measures the distance of each node from the nearest other active nodes. In this figure, lighter colors imply higher localization errors, while darker colors denote lower errors. As it can be seen, the general trend is high localization errors at the lower right corner of the plots (higher node-to-electrode distance and lower inter-nodal space) to the low errors at the upper left corners (lower node-to-electrode distance and higher inter-nodal space). Therefore, the closer the active nodes are to the electrodes and the further they are from each other, the more accurate is the source localization.

2.3.2 Connectivity Analysis

In order to further evaluate the source imaging performance and investigate its ability in identifying brain networks, the connectivity links were analyzed. **Figures 2.10, 2.11** and **2.12**, which are in parallel with **Figures 2.7, 2.8** and **2.9**, focus on connectivity error rather than localization error as the measure of performance. In these figures, distance is defined for the whole source configuration. To this end, we simply averaged the node-to-nearest-electrode distances for all the active nodes in the network to obtain a single distance measure for each case. Same approach was followed to define a single inter-nodal space for every network, i.e. by averaging the inter-nodal distances of all the nodes in a network. Connectivity analysis is not possible for the cases where two nodes are missed, but if only a single node is lost, it is still possible to match them with the corresponding two nodes in the simulated network and compare the relevant links. **Figure 2.10** provides statistical analysis of the connectivity errors. Based on the data presented in this figure, it is observed that connectivity analysis based on EEG+sEEG recordings can improve the accuracy significantly as compared to only EEG or only sEEG based source imaging. The bar plots of **Figure 2.11** suggest that although the connectivity error increases with increasing distance, its overall trend is not similar to the localization error trend observed in **Figure 2.8**. In other words, it seems that while the dependency of the connectivity error on distance is not zero, it still is much less than

the dependency of localization error on distance. This phenomenon is also observed for inter-nodal space in **Figure 2.12**.

Based on the results provided so far, distance of ~ 15 mm seems to be a break point for recovering nodes in 3-node networks using sEEG recordings. This shows that the choice of 30 mm as the maximum distance between sources and sEEG electrodes has been sufficiently large. However, it is still interesting to observe how increasing the distance between sources and electrodes can affect the performance of imaging for single-node configurations. **Figure 2.13**, which is in parallel with **Figure 2.8** except that 250 single-node sources are distributed over the entire brain, shows that both EEG and sEEG are quite successful in localizing single dominant activities regardless of their distance to the recording sites. According to this figure, the median of localization error is less than 4 mm in all schemes and for all distances.

2.3.3 DBS Electrode Leads

The implantation of DBS (deep brain stimulation) electrodes in the brain for Parkinson’s disease and other brain disorders such as epilepsy, indicates that long term implantation of deep electrodes is clinically applicable [22, 92, 93]. DBS electrodes, which were once used for stimulation only, are now capable of recording and stimulating concurrently. Therefore, source localization can be performed on DBS electrode recordings. While DBS electrodes are capable of being implanted chronically (like in Parkinson’s disease), sEEG electrodes are used for short-term recordings to determine the seizure onset zone (SOZ). Thus, each modality is used in different settings and under different conditions; as a result, we wanted to study these two important sub-sets of iEEG recordings separately. To this end, two leads of electrodes were manually positioned in the subthalamic nucleus, which is marked with green in **Figure 2.14**. Each lead consisted of 64 channels including 16 rows of 4 equally spaced contacts. The diameter of each lead was set to be 1.27 mm spanning a length of ~ 12 mm based on a design proposed by [94]. The results for 250 single-node sources distributed over the entire brain are presented in **Figure 2.15**. Unlike sEEG, source imaging based on DBS electrodes is more sensitive to distance from electrodes for both single (compare **Figures 2.13 (B)** and **2.15 (B)**) and multiple-node sources. This is because the number of DBS leads is less than sEEG shanks, which is also more limited in spatial coverage, even though the numbers of

channels are almost the same for both modalities.

2.4 Discussion

We have investigated the possibility, merits and limitations of source imaging using intracranial EEG (iEEG) recordings (which includes sEEG and DBS schemes) and compared its accuracy to the results of EEG source imaging. In this work, the underlying sources were modeled as networks of inter-connected nodes of activity, and the localization error as well as the accuracy of the estimated networks in identifying connectivity patterns, were calculated to assess the performance of imaging brain networks. This feature is particularly important, since we know that physiological and pathological brain processes in the brain are thought to be network phenomenon [76, 95], specifically for epileptic activities [83, 96–101]. Besides, while most of the previous works have not been negligent about the effect of network geometry on source imaging results [58–60, 64], only a few of them have provided statistical analysis on how the mentioned parameters affect the accuracy [21, 102]. Moreover, the initial purpose in these studies has been to demonstrate the possibility of source imaging using intracranial recordings, not providing a benchmark to compare imaging based on deep recordings versus scalp recordings.

We initially placed the sources in the gray matter but finally, decided to place them on the cortical surface. Since, in this manner source orientation could be defined more easily (the source orientations were set to be perpendicular to the cortex surface). However, the results were very close in both cases. It should be noted, that this choice does not remove deep regions from the source space. To demonstrate this visually, we have prepared **Figure A.2**, which is included in the supplementary materials. The red spheres in this figure show the sEEG electrodes and the blue spheres show the source space from which source locations are selected. While sources are confined to the cortex, they can be very deep.

EEG is an effective tool to study the dynamics and temporal characteristics of the electrical activity in the human brain. Its high temporal resolution besides being noninvasive and easy to set up has made it among the popular techniques for human brain functional imaging. Despite the merits, EEG usually has difficulty to perform well in localizing deep brain activities due to low SNR compared to cortical source

localization. When the results of EEG source imaging are not decisive enough, clinicians may opt for an invasive method such as sEEG. To the best of our knowledge, sEEG recordings are not currently used with the purpose of performing source imaging, in the clinical practice [55, 56]. Clinicians use the electrodes only to locally monitor the activities in the pathological regions of interest, usually pre-determined before electrode placement based on patient history and other relevant clinical data and possible imaging studies.

Imaging based on sEEG recordings share most of the major properties of EEG. The major difference, as the name may suggest, is that the electrodes are distributed in the regions with a higher probability of seizure activity rather than over the scalp. In other words, the electrodes are closer to major active nodes, which can compensate for the low coverage of deep and far-from-scalp sources in EEG. Despite this advantage, there are several reasons leading the sEEG-based imaging not to perform as well as it may seem at the first glance. Although our results indicate that, it can perform better than EEG under some conditions (if there is only a single dominant activity or if sources are close enough to the sEEG electrodes). When there are multiple active nodes in the vicinity, the nodes that are closer to the electrodes leave a much stronger impact on the recording electrodes than the further ones. Therefore, the distant sources may be totally masked by the activity of the closer sources. This is indeed the same issue as the low coverage of deep and far-from-scalp sources in EEG, which may be even more severe in this case, since some sources might be located very close to the electrodes (depending on the electrode placement and the actual location of the underlying epileptic sources). The high sensitivity to distance is the direct consequence of the mentioned issue. Our simulation results demonstrate that the threshold for the break point is around 15 mm, implying that for the active nodes, which are closer than 15 mm to the electrodes, we can expect an accurate localization. While for the nodes located further than 15 mm to the closest electrode, there is rarely any merit using sEEG recordings, for imaging purposes.

Additionally, the spatial location of the scalp EEG electrodes allows them to see the sources with a larger angle, which leads to a better coverage of the brain. Besides, the lead field matrix obtained for sEEG modality may be a worse conditioned matrix thus making it more difficult to obtain stable solutions compared to EEG. This is

probably because the contacts of each electrode either collect redundant information about the source (depending on how similar their relative location and orientation with respect to the source is), or sense nothing due to far distance between source and electrode. However, the mentioned issue is highly dependent on the relative location of the electrodes, how distributed they are and how close the contacts of each electrode are to each other. Therefore, it is hard to make a general remark regarding the lead field matrix of sEEG electrodes compared to that of EEG electrodes, inexplicably. The interested readers are referred to [103] for a comprehensive discussion of which properties of a sensing matrix make it more appropriate and desired, from an inverse-problem point of view.

Another important contribution of this study to the literature is that we have shown how combining EEG and iEEG can improve the functional imaging accuracy in different ways, especially for source configurations of multiple nodes. In order to solve the problem with more confidence, more constraints, equations or information about the inverse solution are desirable. The combination of EEG and iEEG, can serve this purpose by increasing the number of measurements. The combination of information from EEG and iEEG also improves the performance by reducing the number of blind spots in the brain, which are not covered by either of the setups, individually. In the new scheme, one can really hope that regions in the brain, which are considered deep, relative to some surface electrodes, are positioned closely with respect to deeply implanted electrodes and vice versa. It should be emphasized that by combining EEG and iEEG, we do not imply that EEG and iEEG recordings are measured separately and then compared to each other, but rather, we refer to a unified inverse algorithm, which incorporates EEG and iEEG recordings measured simultaneously. **Figures 2.7** and **2.10** support this claim. According to **Figure 2.7**, we can expect median localization error of less than 5 mm, which seems to be quite favorable. **Figure 2.10** suggests that by combining EEG and sEEG, one can detect the connectivity links between network nodes much better, which is essential in studying brain networks.

In this study, we have assumed different noise levels for EEG and iEEG recordings to simulate a more realistic situation; more severe conditions for EEG recordings (5 dB SNR) and more favorable ones for iEEG recordings (20 dB SNR). While SNR can vary significantly depending on the location, orientation and strength of brain activities,

one can generally expect scalp recordings to be noisier than intracranial recordings, since electrical signals need to pass through skull, skin and brain tissue to reach scalp electrodes. Moreover, since the origins of deeper activities are further away from the scalp, EEG measurements are weaker at the scalp electrodes. Since the placement of deeper electrodes (by the physician) usually suggests that epileptic sources are not superficial (most probably), in most of our simulations we have placed the sources near the iEEG electrodes and thus have adopted this lower SNR for EEG recordings. It is obvious that if sources are more superficial and closer to scalp electrodes, then the EEG signal will have a better SNR and much improved results can be achieved for EEG source imaging. We have observed in simulations that by simply changing the depth (distance to closest scalp electrode) of sources, the SNR could change by 8 dB. It should be emphasized that the reported SNRs are average values and not necessarily the peak SNR. Therefore, in practice we have a range of SNR values at different times.

This study has not been limited to imaging based on sEEG recordings. The long term implantation of DBS electrodes in treatment of different brain disorders such as Parkinson’s disease and epilepsy [22,92,93] and the recording capability of current DBS electrodes, poses the question of whether brain functional imaging can be improved using DBS recordings or not. Based on our study, the answer to this question depends on different factors such as the number and relative location of nodes of activity in the brain. While imaging based on DBS recordings share most of the features of an sEEG-based scheme, the fewer number of electrodes and their distribution in a smaller area can make it a more vulnerable modality against sources located further away from DBS electrodes compared to sEEG. This might be a far-field phenomenon; to a far-field observer, where sources are situated far away from DBS electrodes, all the DBS electrodes might be perceived as a single super-electrode. This is foreseeable, as the distance between DBS electrodes is less than the distance between sEEG electrodes and DBS shanks are placed near each other and not as distributed as sEEG electrodes. However, if the placement of DBS electrodes were optimized for individual patients so that they are placed near the active sources, improvement can be expected.

The significance of these results stems from the fact that detecting the epileptogenic zone plays an important role in the treatment of about one third of the epilepsy patients, who do not respond to any type of medication. Therefore, receiving surgery for

removal of the epileptogenic zone could be justified for these patients. Any error in determining the epileptogenic zone may lead to either unsatisfactory surgery outcomes or lifetime irreversible side effects for the patients [74]. As already mentioned, clinicians combine the results of different imaging modalities before they finalize their decision on determining the epileptogenic zone. On the other hand, the simple use of intracranial recordings without applying ESI to it, may lead to overlooking of some active sites located further away. In these situations, combining EEG and iEEG recordings can achieve a more accurate imaging of the whole brain. Furthermore, as simulation results suggest, the sensitivity of imaging to source distance (from recording site) is reduced in the combined scheme. The conclusions regarding the combined modality can also extend to more realistic scenarios, where sources of activity are assumed to have an extent rather than being focal and encompassing just a few active nodes [104].

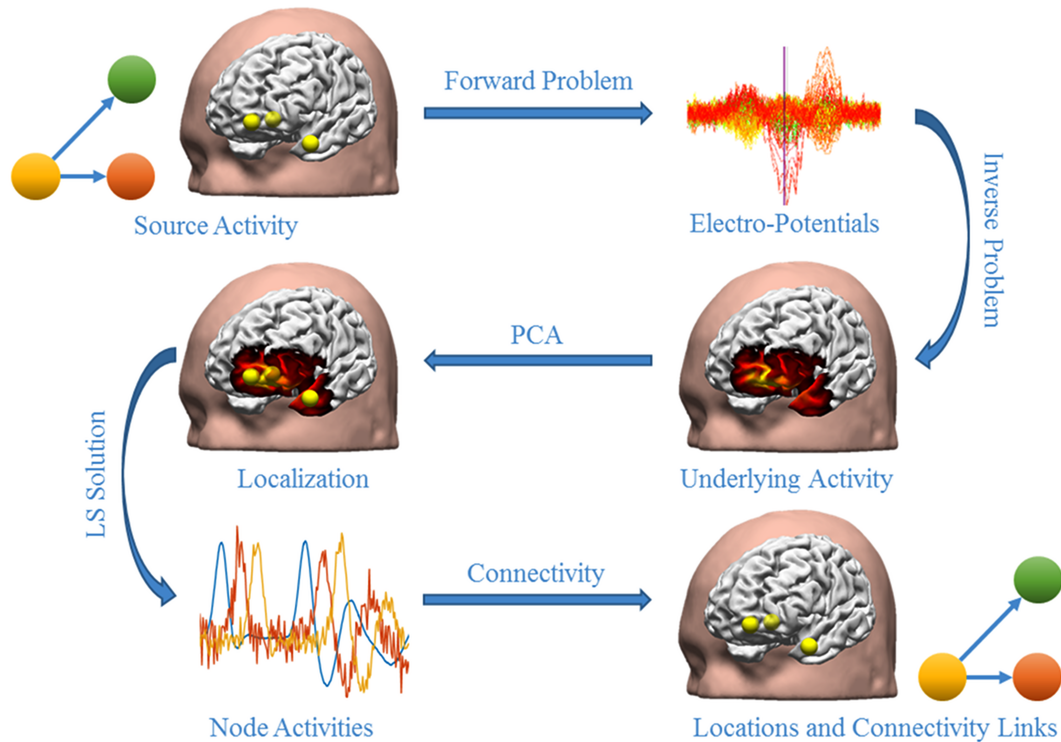


Figure 2.1: An overview of the simulation procedures. After EEG electrodes and iEEG electrodes are placed within a realistic head geometry, the location of the network nodes is randomly selected and a connectivity link is created among these nodes based on MVAR models (here node 1 is driving nodes 2 and 3). Then the forward problem is solved to calculate the potentials at the electrodes and noise (AWGN and ACGN) is added to these simulated potentials. The inverse problem is then solved and PCA is performed on the estimated source activities to determine the location of the simulated network nodes. Consequently, the activity of these nodes over time is extracted. These activities are then used to detect the connectivity links among the mentioned nodes.

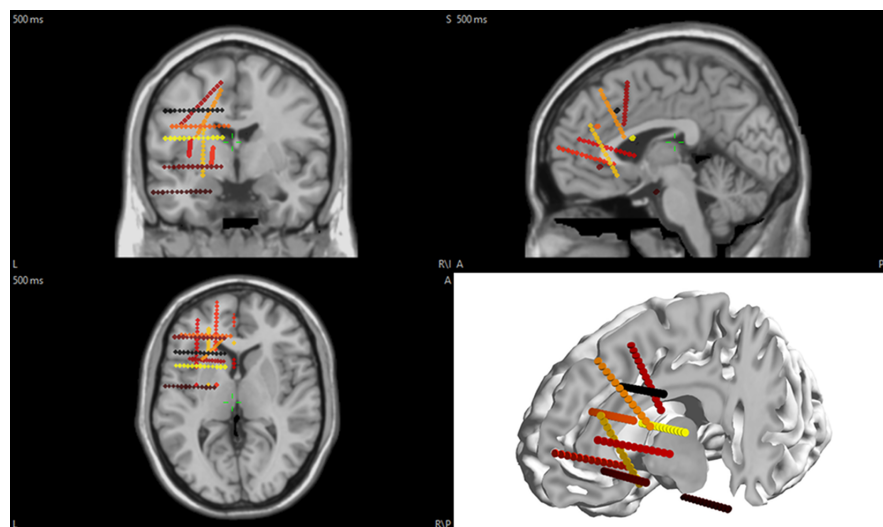


Figure 2.2: The position of the electrodes in the brain for this study. EEG electrodes are on the scalp (not shown here), while sEEG electrodes are inside the brain tissue in the left frontal lobe and part of the left temporal lobe.

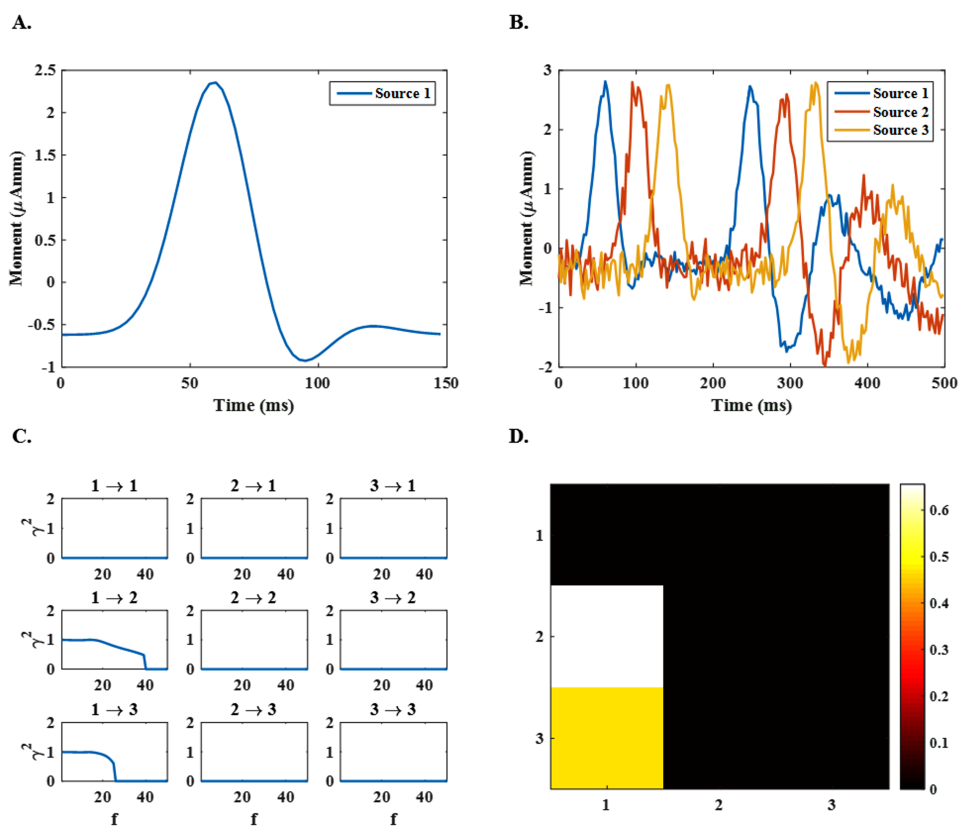


Figure 2.3: The configuration of the simulated networks. (A) The time course assigned to the dipole in single-node sources, (B) the time courses assigned to dipoles in 3-node sources, (C) the connectivity links between nodes recovered by directed transfer function (DTF) at different frequencies in the range 1 – 50 Hz and (D) the recovered connectivity links averaged over all frequencies in the mentioned range.




Figure 2.4: Median localization error for 3-node network configurations. Error bars refer to the first and third quartiles (of localization error distribution). Two noise types, namely, additive white Gaussian noise (AWGN) and additive colored Gaussian noise (ACGN) with SNRs of 5 dB and 20 dB are simulated. The simulations are repeated for both EEG and sEEG electrodes and with a region of interest (ROI) assumption (refer to section 2 for more details). As it can be seen, the ROI assumption can improve the source localization, however it requires some extra information, which is not generally available.

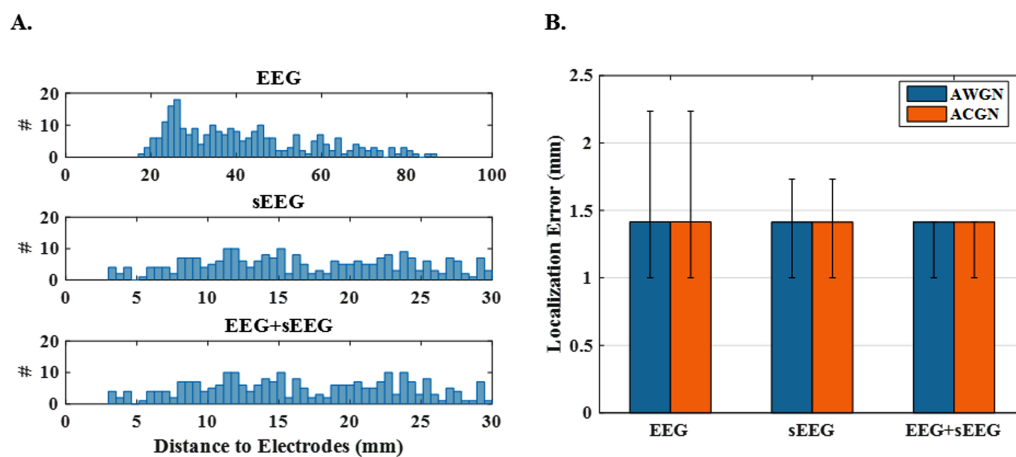


Figure 2.5: Localization error and histogram of source distance to electrodes for 1-node networks. (A) The histogram of the source-to-nearest electrode distance for 250 single-node sources in the three recording schemes EEG, sEEG and EEG+sEEG. It is obvious that sources are generally further away from EEG electrodes. (B) The median localization error for the same source configurations and for white (AWGN) and colored (ACGN) noise. EEG recordings are subject to the noise with an SNR of 5 dB, while this is 20 dB for sEEG. These values are maintained in the combined version (EEG+sEEG), as well

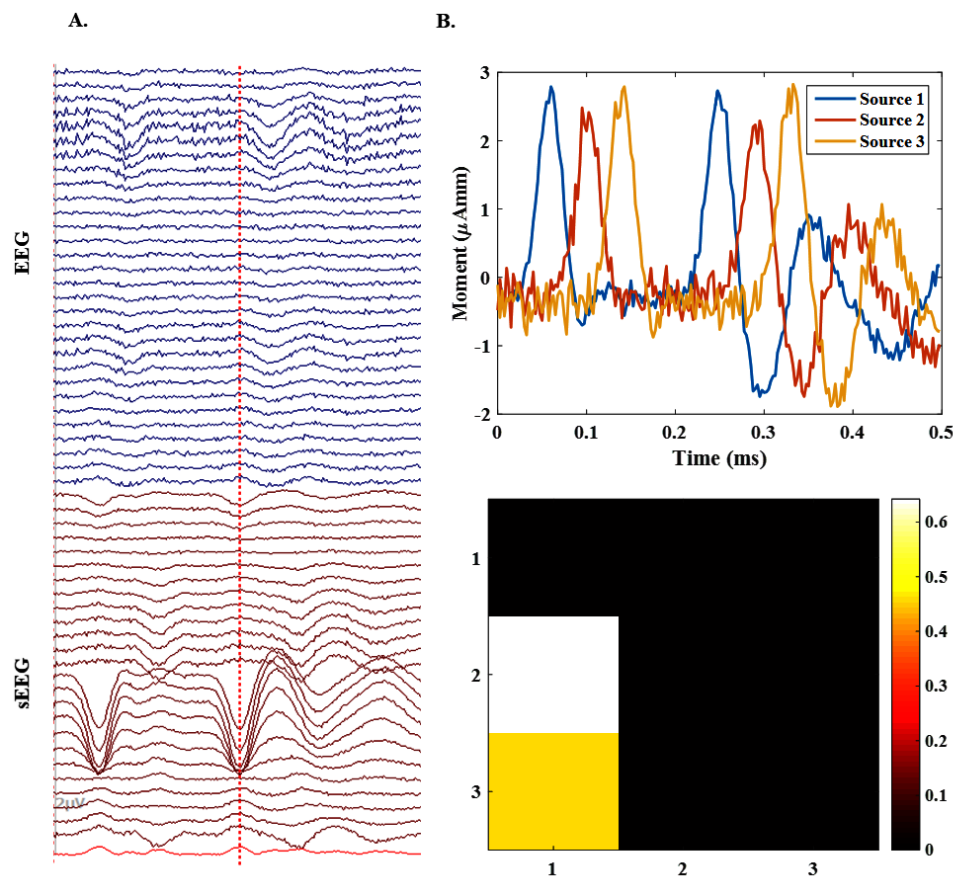


Figure 2.6: An example of a 3-node network. (A) The simulated electrical potential recordings for a single 3-node source configuration in both EEG and sEEG setups. sEEG contains stronger spikes than EEG. (B) The accuracy of the source imaging for the same source using EEG+sEEG recordings. It can be seen that the time courses for each active node of the source are recovered precisely. Besides, the connectivity links are identified truly in the sense that source 1 is driving sources 2 and 3, as expected.

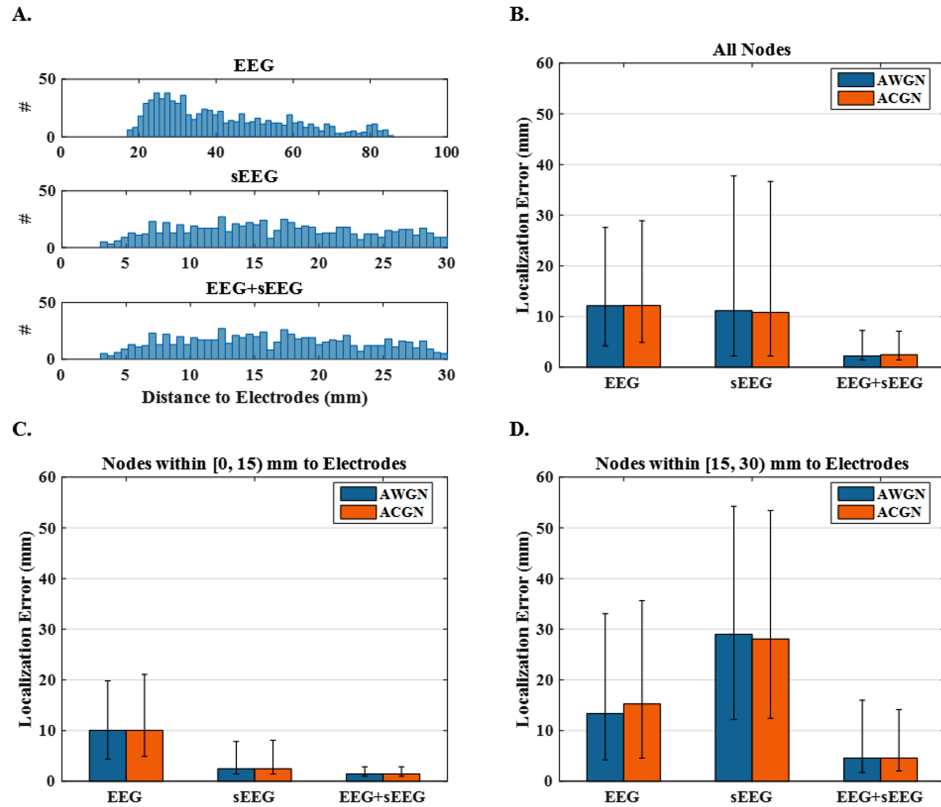


Figure 2.7: Localization error and histogram of source distance to electrodes for 3-node networks. (A) The histogram of the distances between active nodes and the nearest electrode to them. This figure depicts the results of different electrode set-ups, i.e., EEG, sEEG and EEG+sEEG. (B) The median localization errors using EEG, sEEG and EEG+sEEG recordings under AWGN and ACGN noisy conditions. (C and D) The same plot only for the sources that lie within a distance of 0 – 15 mm and 15 – 30 mm from sEEG electrodes, respectively.

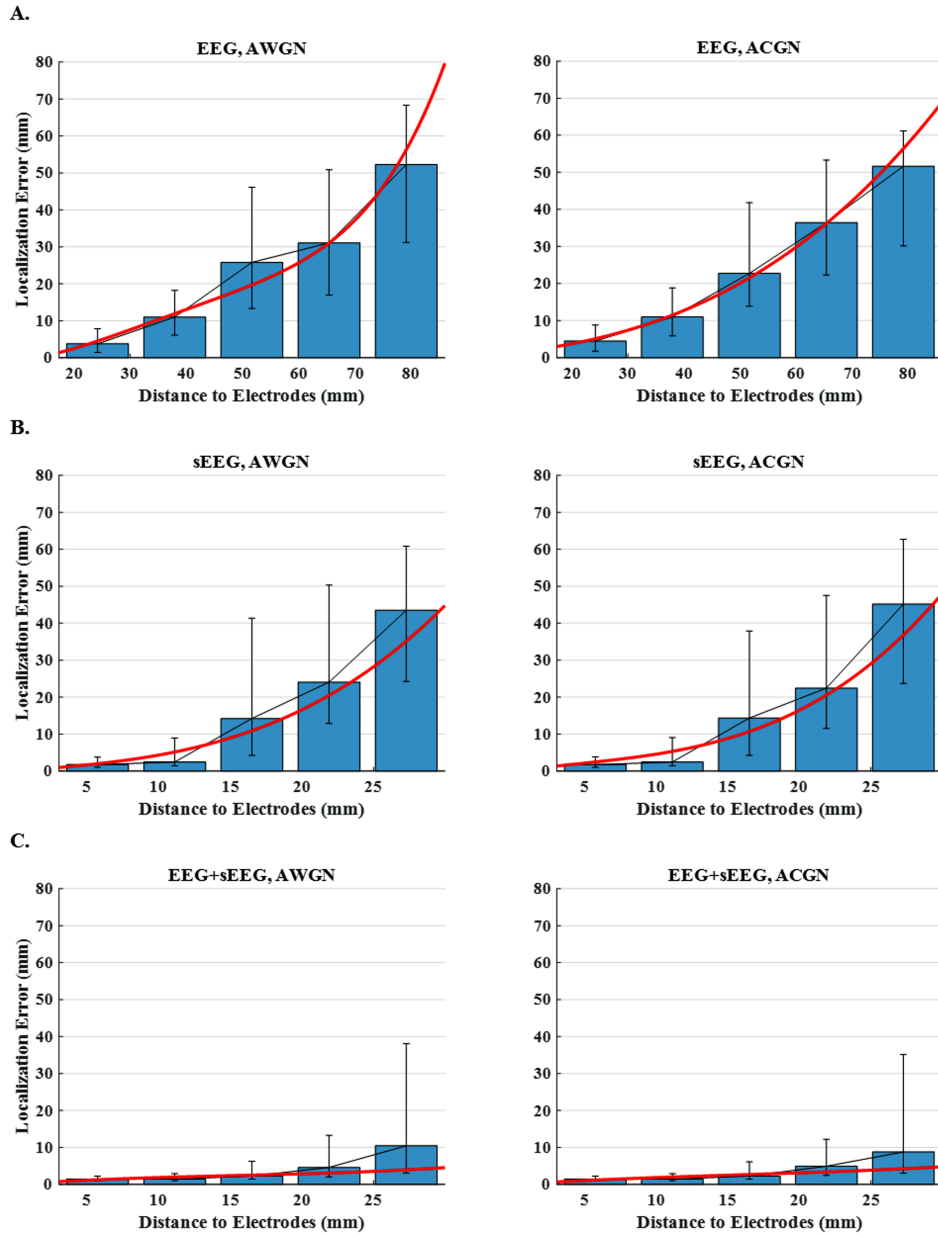


Figure 2.8: Localization error versus distance to electrodes for 3-node network configurations. The plots display the median localization errors over five intervals (distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a polynomial of a degree determined by cross validation.

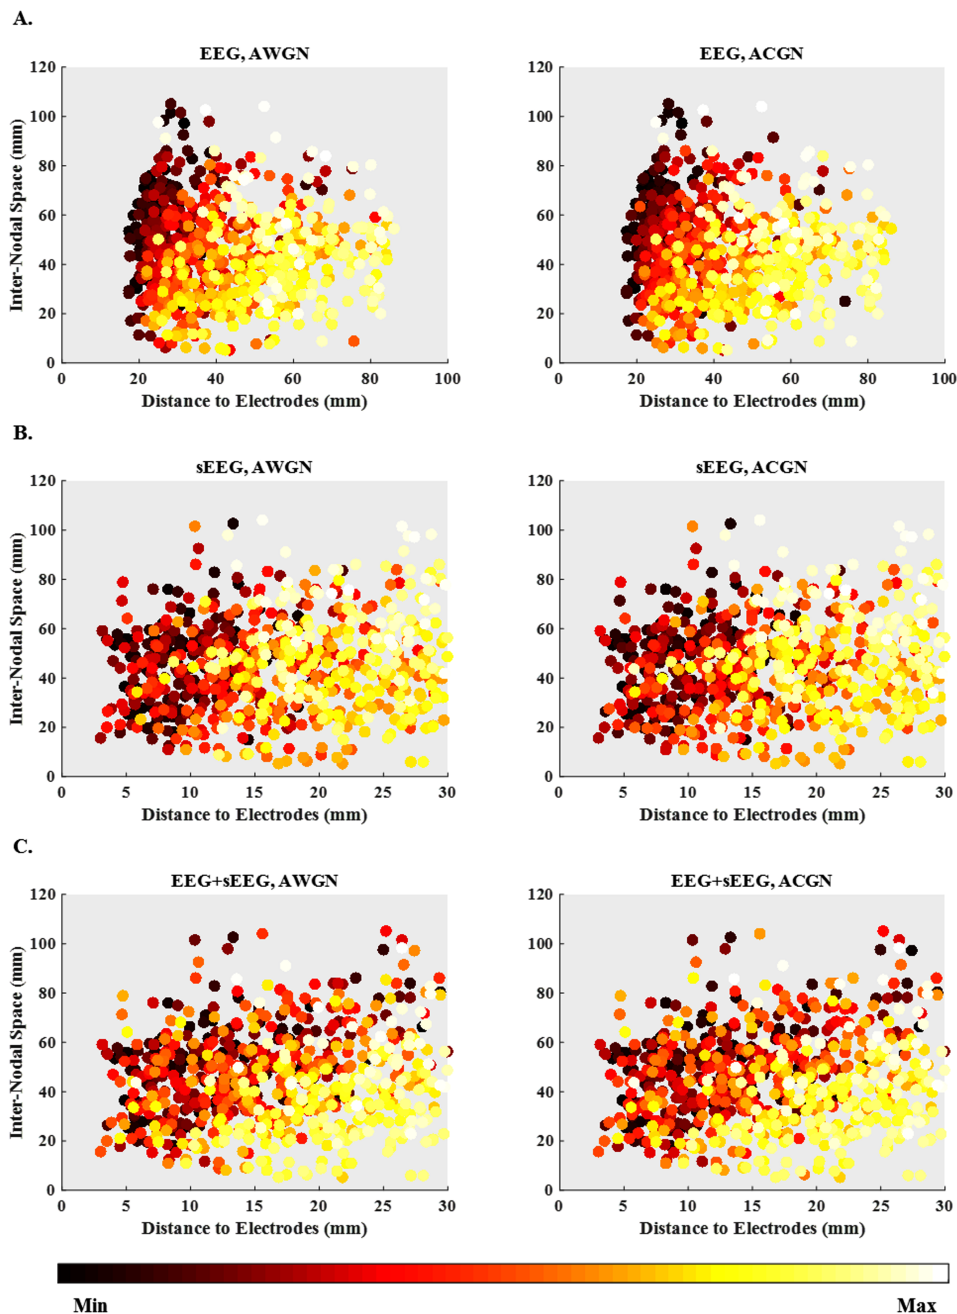


Figure 2.9: Localization error scatter plots. Localization error scatter plots versus distance to electrodes and inter-nodal space for (A) EEG, (B) sEEG and (C) EEG+sEEG recording set-ups under AWGN and ACGN noisy conditions. The horizontal axis and the vertical axis are respectively, the distance to electrodes for every single node and the distance to the nearest other active nodes in the same source configuration (inter-nodal space). The localization errors, the third axis (not shown), are represented with colors.

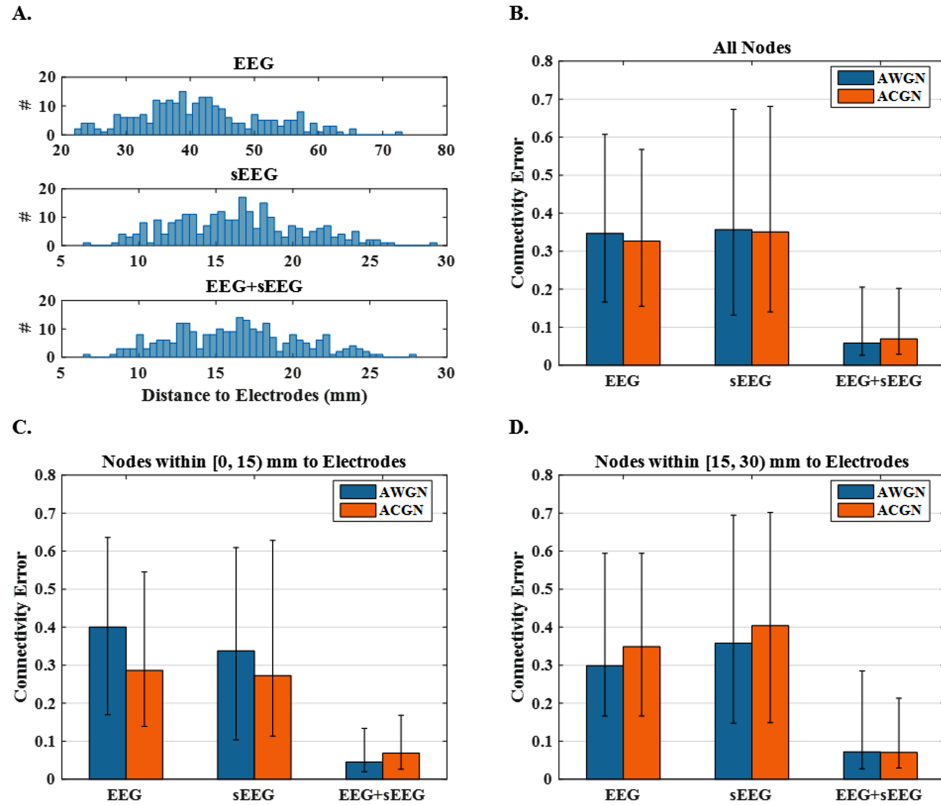


Figure 2.10: .Connectivity error and histogram of source distance to electrodes for 3-node networks. (A) The histogram of the average distances between active nodes and the nearest electrode to them. This figure depicts results of different electrode set-ups, i.e., EEG, sEEG and EEG+sEEG. (B) The median connectivity errors using EEG, sEEG and EEG+sEEG recordings under AWGN and ACGN noisy condition. (C and D) The same plot only for the sources that lie within a distance of 0 – 15 mm and 15 – 30 mm from sEEG electrodes respectively.

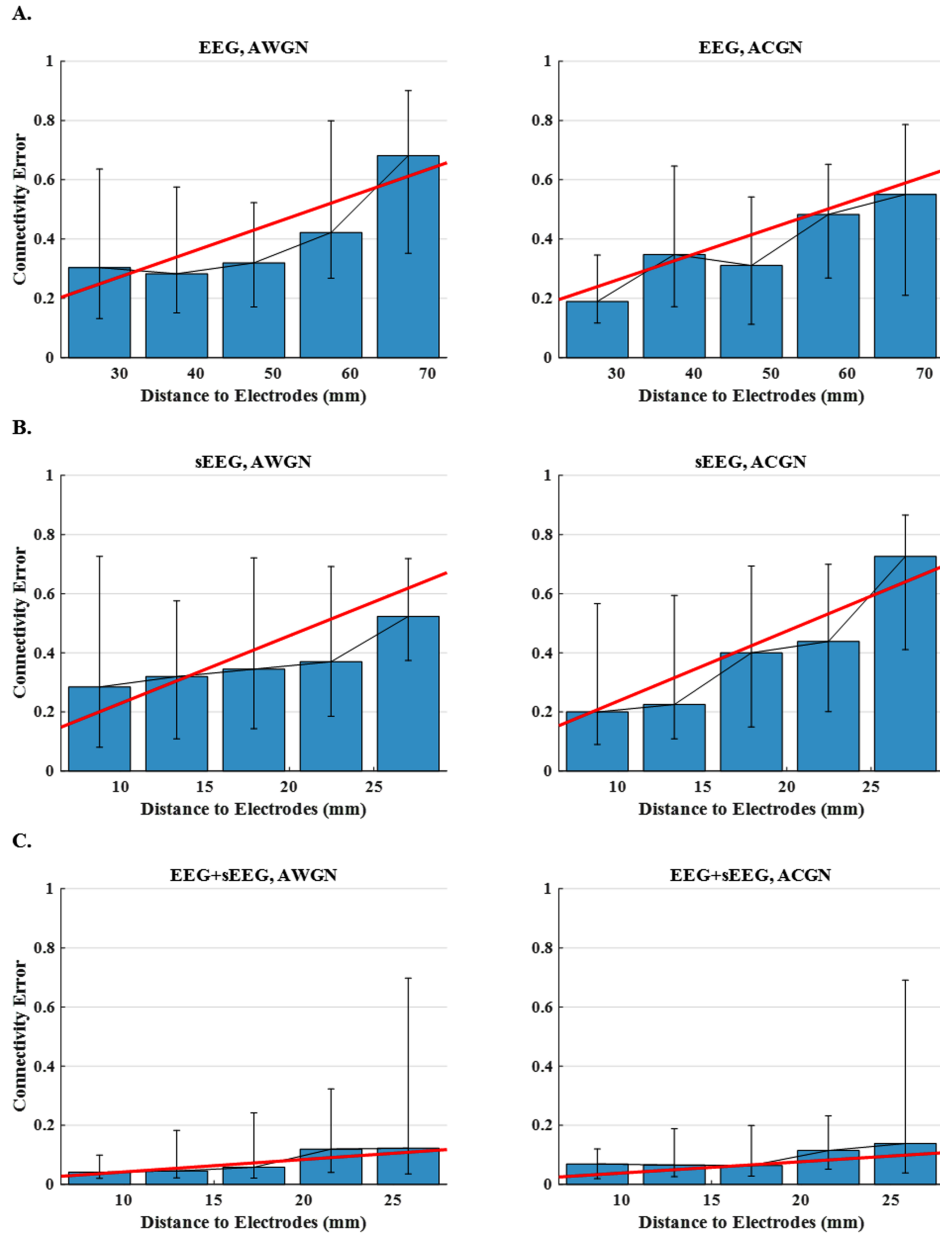


Figure 2.11: Connectivity error versus distance to electrodes for 3-node network configurations. The plots display the median connectivity errors over five intervals (average distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG set-ups and for AWGN and ACGN noise types. In order to make the trends easier to detect, the scattered points are regressed with a line.

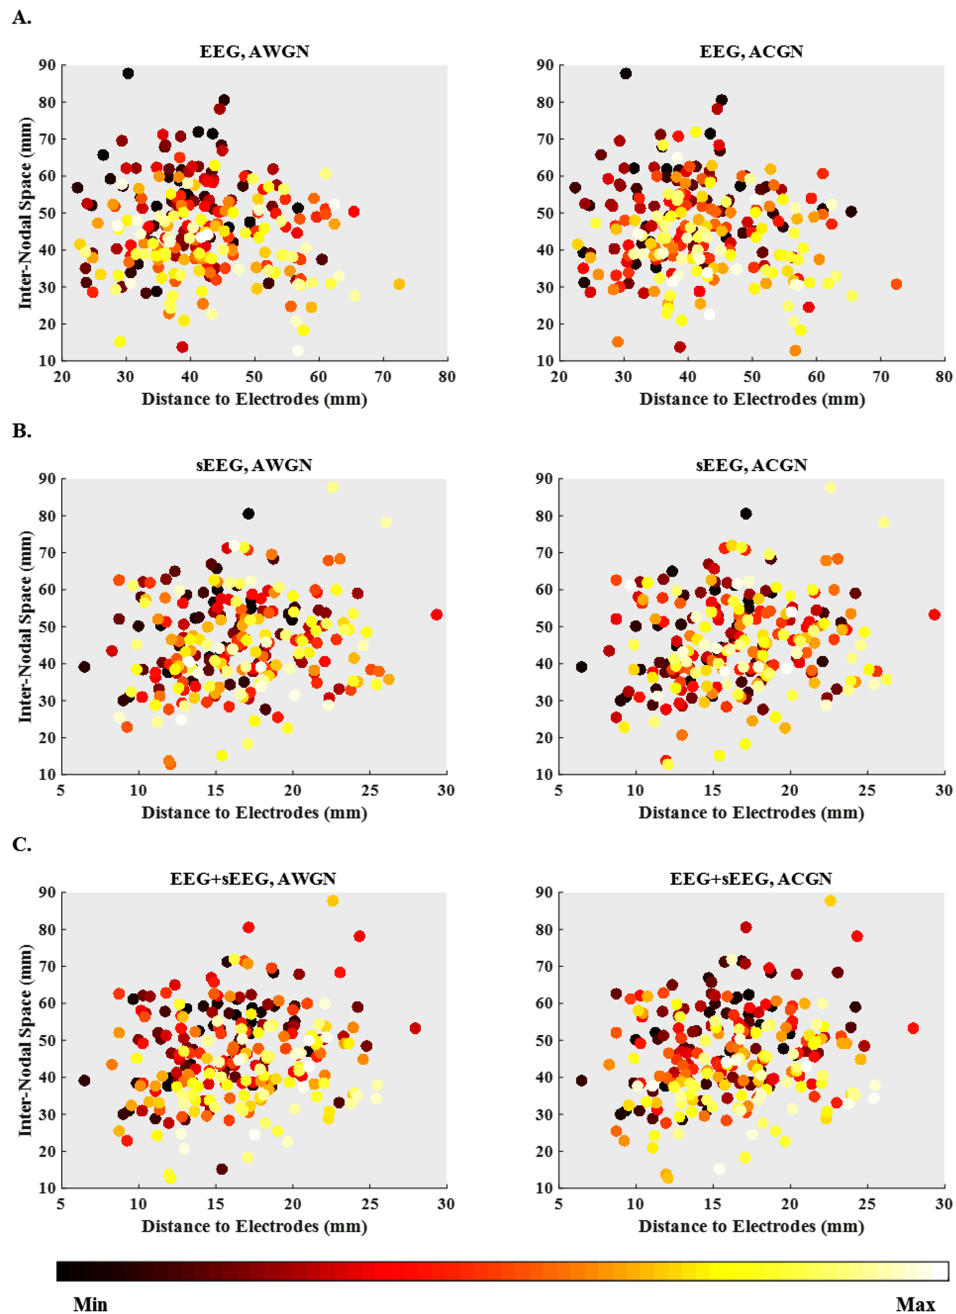


Figure 2.12: .Connectivity error scatter plots. Connectivity error scatter plots versus distance to electrodes and inter-nodal space for (A) EEG, (B) sEEG and (C) EEG+sEEG recording schemes under AWGN and ACGN noisy conditions. The horizontal axis and the vertical axis are respectively, the average distance to electrodes for every single node and the average inter-nodal space in the same source configuration. The connectivity errors, the third axis (not shown), are represented with colors.

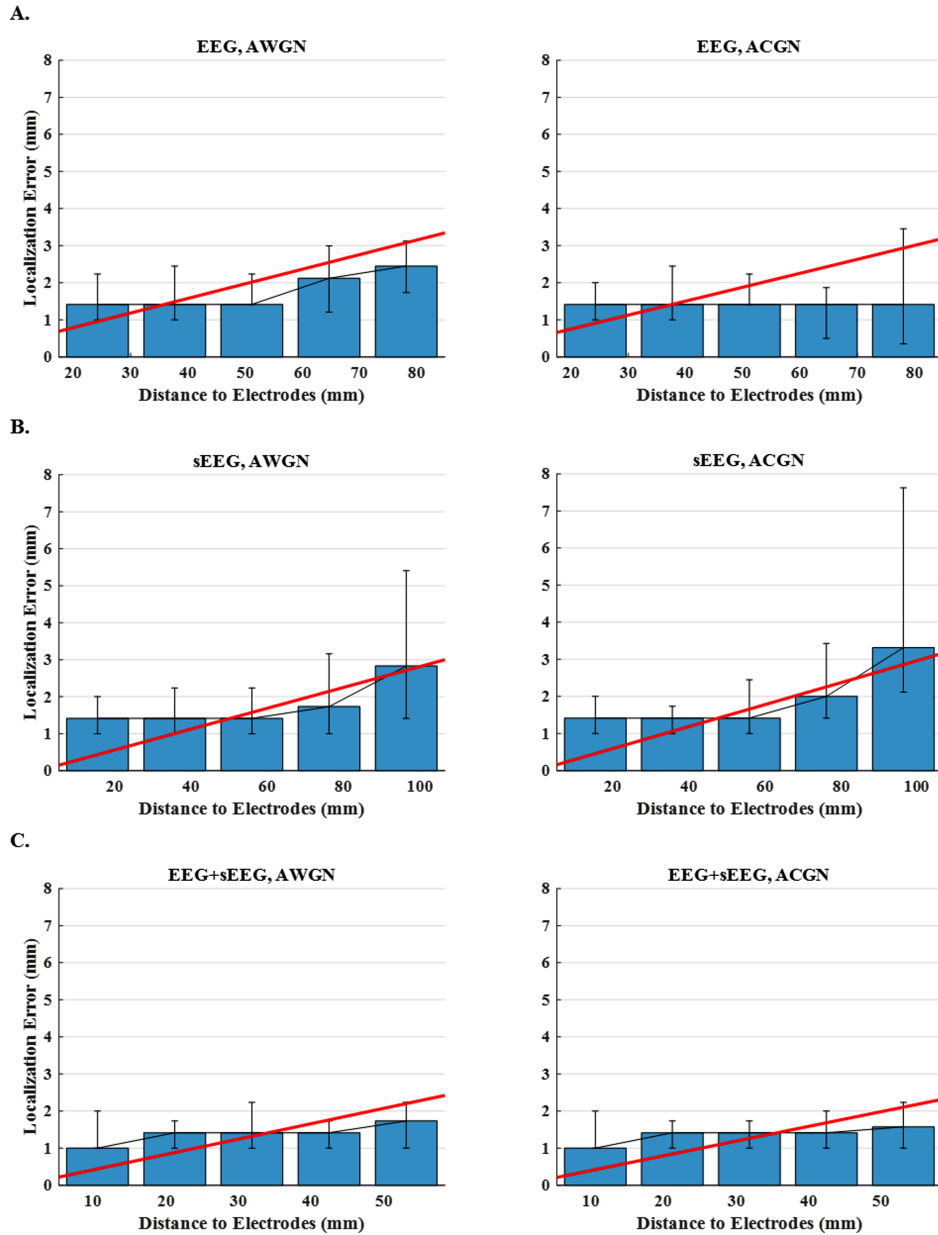


Figure 2.13: Localization error versus distance to sEEG electrodes for sources distributed over the entire brain. The plots display the median localization errors over five ranges intervals (distance to the nearest electrode) for (A) EEG, (B) sEEG and (C) EEG+sEEG schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a line.




Figure 2.14: .The position of the DBS electrodes in the brain. EEG electrodes are on the scalp (not shown here), while DBS electrodes are implanted inside the brain tissue. The green area shows the subthalamic nucleus.

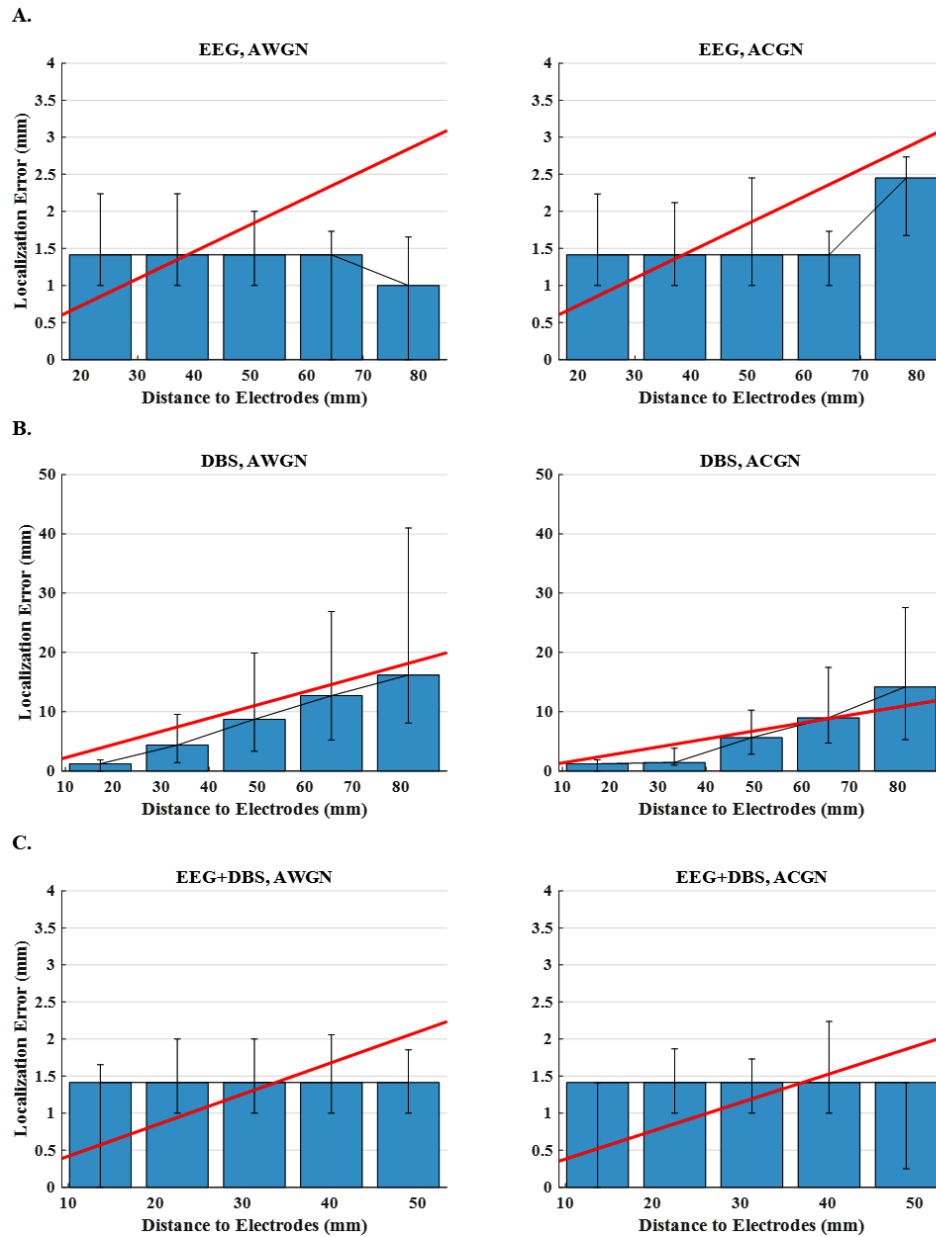


Figure 2.15: Localization error versus distance to DBS electrodes for sources distributed over the entire brain. The plots display the median localization errors over five intervals (distance to the nearest electrode) for (A) EEG, (B) DBS and (C) EEG+DBS schemes and for AWGN and ACGN noise types. In order to make the trends more detectable, the scattered points are regressed with a line.

Chapter 3

Electromagnetic Source Imaging by Means of a Robust Minimum Variance Beamformer

3.1 Introduction

Electromagnetic source imaging (ESI) [6, 18, 19] using electroencephalography (EEG) [5, 7] or magnetoencephalography (MEG) [8, 9] measurements, as discussed in Chapter 2, is an effective tool for mapping and imaging dynamic brain electrical activities. This functional imaging modality, which is typically non-invasive (see [21, 58, 105] for ESI using intracranial recordings), has been utilized by many researchers in clinical environments to study the brain function or dysfunction in various physiological or pathological (e.g., in epilepsy patients) conditions. Due to the ill-posed nature of the EEG/MEG source-imaging problem, many techniques have been proposed in the literature to regularize the problem. See [106] for review of ESI techniques. Adaptive beamformers [107, 108], which are the focus of this study, design spatial filters to selectively pass the signals associated with desired locations while suppressing the activities coming from the rest of the brain. The word adaptive corresponds to a feature by which the ultimate model used to solve the inverse problem depends on the measurements as well as the head volume conductor properties. In adaptive beamformers, adaptation

is achieved by incorporating the measurements and noise covariance matrices into the problem formulation.

Since first introduced to the brain source imaging community, adaptive beamformers have been more often used for MEG source imaging rather than EEG source imaging [19, 109]. This may in part be due to the fact that sources of uncertainty in the forward modeling of EEG is usually considered to be more extensive compared to MEG. In practice, electrical activities of the brain can be situated anywhere in the brain tissue, while current dipoles are assumed in fixed and pre-discretized locations (in a distributed model). In addition, there are levels of uncertainty in the electrical properties of the head volume conductor, such as the relative conductivity ratios of different tissues (the scalp, the skull and the brain, in boundary element model) and their inhomogeneity profiles. Furthermore, realistic head volume model of the subject may be unavailable in some cases. Unreliable estimation of the covariance matrices because of insufficient or noisy data is another source of error in practice. Consequently, availability of the true lead field is almost impossible in many cases. On the other hand, beamformers are generally sensitive to the errors in the forward models. More specifically, linearly constrained minimum variance (LCMV) beamformer, which is the base of all adaptive beamformers, is known to be highly sensitive even to slight mismatches between the true and estimated models [110, 111]. See e.g., [109] for a detailed beamformer analysis of sensitivity to the forward modeling mismatches.

To deal with these issues, the authors of [112, 113] have applied a technique called diagonal loading (DL) to the MEG source localization problem. DL replaces the covariance matrix of the measurements with a regularized version, by adding a constant factor of the unity matrix to the measurement covariance matrix. Although DL can reduce the sensitivity to some extent, it is understood that such regularizations lead to a trade-off between the output signal-to-noise ratio (SNR) and the spatial resolution of the LCMV [110]. More importantly, it is not clear how to determine the optimal value of the DL factor based on known levels of uncertainty in the lead field matrix [111]. Besides DL, eigenspace beamformer [110, 114] is another technique that can yield robustness against modeling errors as well as measurement noise. We sought to show that the performance of this method may be further improved if the uncertainty in the forward modeling is taken into consideration in the model. While modeling lead field

uncertainty has already been investigated in the signal processing community through a technique called robust minimum variance beamformer (RMVB) [111, 115, 116], this study, to the best of our knowledge, is the first to adapt the RMVB to functional brain imaging and source imaging [117]. Besides output SNR, which was originally used to compare the RMVB to its peers, we used three other criteria to investigate the merits and limitations of the RMVB and to compare its accuracy to the conventional adaptive beamformers for the purpose of ESI. To this end, we conducted comprehensive computer simulations to show the merits of a robust modeling.

3.2 Methods

Brain electrical activities can be modeled by current dipoles [18, 84]. Since Maxwell's equations are solved in a quasi-static regime, it can be assumed that the relation between these dipoles and the EEG/MEG potentials generated at the sensors is instantaneous and linear [85]. More specifically,

$$\Phi = \mathbf{K}\mathbf{J} + \mathbf{N}_0, \quad (3.1)$$

where Φ is the matrix of EEG/MEG potentials generated/recorded at the sensors at different time points, \mathbf{K} is the lead field matrix defining the linear relation between the current dipoles and the potentials, \mathbf{J} is the matrix of current dipoles over time and \mathbf{N}_0 models the noise at the sensors over time. Assuming M sensors, N current dipoles and T time points, Φ and \mathbf{N}_0 are the matrices of size $M \times T$, \mathbf{J} is a matrix of size $N \times T$ and \mathbf{K} is a matrix of size $M \times N$, which encompasses the geometrical and electrical properties of the medium through which brain signals propagate to reach to the sensors.

3.2.1 Linearly Constrained Minimum Variance (LCMV)

The goal of ESI techniques is to recover the underlying source activity (current dipole matrix \mathbf{J}) using a set of electromagnetic recordings (the measurements matrix Φ). Scanning techniques (LCMV, (RAP)-MUSIC, FINES, etc.) [108, 118, 119] use all the measurements to estimate only a single element of the current dipole vector located in a specific voxel, so in order to estimate the full current dipole vector, all predefined source space locations have to be scanned; hence the name scanning. LCMV, which lies in this

category, performs the task by designing a spatial filter for each voxel. This spatial filter is a linear operator, which once applied to the measurement matrix Φ , back-projects the activities of a desired voxel from the scalp measurements, while attempting to suppress the activities of all other voxels besides noise. This procedure is repeated until all current dipoles (at every voxel in the source space) are estimated. Speaking more mathematically, LCMV solves the following optimization problem:

$$\begin{aligned} \mathbf{w}_i^* &= \arg \min_{\mathbf{w}_i} \mathbf{w}_i^T \mathbf{C} \mathbf{w}_i \\ \text{s.t. } \mathbf{w}_i^T \mathbf{k}_i &= 1, \quad \forall i \in \{1, \dots, N\} \end{aligned} \quad (3.2)$$

where \mathbf{C} is the measurements covariance matrix of size $M \times M$ (see Practical Considerations section A.2.3 in the Supplementary materials for estimation details), \mathbf{k}_i is the i -th column of the lead field matrix \mathbf{K} corresponding to the current dipole at the i -th voxel and \mathbf{w}_i is a vector of size $M \times 1$, which yields the estimation of dipole i through the following equation:

$$\hat{\mathbf{J}}_i = \mathbf{w}_i^{*T} \Phi, \quad \forall i \in \{1, \dots, N\} \quad (3.3)$$

Using Lagrange multipliers method, the optimal solution of this optimization problem can be shown to be

$$\mathbf{w}_i^* = (\mathbf{k}_i^T \mathbf{C}^{-1} \mathbf{k}_i)^{-1} \mathbf{C}^{-1} \mathbf{k}_i, \quad \forall i \in \{1, \dots, N\} \quad (3.4)$$

To provide an intuitive interpretation of this procedure, one should notice that the objective function in the optimization problem (3.2) is in fact the output power of the linear filter (or the variance of the estimated current dipole at voxel i). By minimizing the filter output power (minimum variance (MV)) under the constraint $\mathbf{w}_i^T \mathbf{k}_i = 1$ (linearly constrained (LC)), LCMV ensures that to the best of its ability, the filter removes the contribution of all irrelevant activities to the measurements while keeping the desired signal intact.

Based on (3.4), it is clear that \mathbf{w}_i^* is inversely proportional to the norm of the i -th lead field column associated with the i -th voxel. Since the norm of the lead field vector is generally smaller for the voxels that are located further away from the electrodes, the filter coefficients become larger for deeper locations. This generates some bias towards deeper activities. Additionally, depending on the location and orientation of the dipoles,

noise may affect the sources differently. To compensate for the depth and asymmetric spatial distribution of the noise, LCMV is usually followed by a normalization step of the filter coefficients [108, 113, 114, 120]. To this end, the filter coefficients are normalized either by their ℓ_2 norm [114] or by a factor, which is a function of noise covariance matrix. This factor can be calculated in different ways [113], for example by finding the power of each voxel, if the filter coefficients are applied to pure noise data. In other words, the filter coefficients can be normalized according to

$$\mathbf{w}_i = \frac{\mathbf{w}_i^*}{\sqrt{\mathbf{w}_i^{*T} \mathbf{C}_N \mathbf{w}_i^*}}, \quad \forall i \in \{1, \dots, N\} \quad (3.5)$$

where \mathbf{C}_N is the noise covariance matrix estimated using baseline (see Practical Considerations section A.2.3 in the Supplementary materials for estimation details). The authors of [108] employed a slightly different strategy; they first customized the filter coefficients by solving the optimization problem (3.2) for noise-only segments and then followed the same strategy, i.e.,

$$\tilde{\mathbf{w}}_i = \frac{\mathbf{w}_i^*}{\sqrt{\tilde{\mathbf{w}}_i^{*T} \mathbf{C}_N \tilde{\mathbf{w}}_i^*}}, \quad \forall i \in \{1, \dots, N\} \quad (3.6)$$

where $\tilde{\mathbf{w}}_i^*$ is the solution of optimization problem (3.2) with the measurements covariance matrix \mathbf{C} replaced with noise covariance matrix \mathbf{C}_N . This strategy was adopted for normalization throughout this study. It should be mentioned that in some applications such as resting state analysis, the estimation of noise covariance matrix is not straightforward. In such situations, \mathbf{C}_N can be replaced with a unity matrix of appropriate size, which is equivalent to assuming pure white and identically distributed noise across all sensors.

Besides the normalization, LCMV can also benefit significantly from a denoising procedure in the end. This method, which is called the eigenspace beamformer, was first introduced by [110, 114] to the source imaging community. The eigenspace beamformer assumes that the number of sources is known a priori. It then exploits this information to separate the signal and noise subspace and project the filter coefficients vector \mathbf{w}_i onto the signal subspace as follows:

$$\bar{\mathbf{w}}_i = \mathbf{E}_s^T \mathbf{E}_s \mathbf{w}_i, \quad \forall i \in \{1, \dots, N\} \quad (3.7)$$

where \mathbf{E}_s is the truncated covariance matrix of the measurements \mathbf{C} . \mathbf{E}_s is calculated by keeping the principal components (associated with dominant singular values) obtained from singular-value decomposition (SVD) of matrix \mathbf{C} and setting the rest to zero (the small values are assumed to be due to noise). In practice, there are different approaches to determine the number of principal components. The number of components can be estimated by keeping the components that explain a preset level of variation starting from components with higher singular values, or those components that lie above the knee of the curve showing the sorted singular values. In this study, Kaiser’s rule [121] was used to select the number of principal components. More specifically, we kept the components associated with the singular values, which were greater than the average of all singular values. Based on our simulations, while both the normalization and denoising steps play an important role in the quality of the LCMV solutions, the role of denoising is more crucial. Furthermore, applying both steps in series improves the accuracy of solutions only slightly, compared to the denoising alone. Nonetheless, we decided to implement both steps in this study, since it was quite straightforward to apply the normalization step as well. This method will be referred to as “LCMV-ND-DN” (LCMV-normalized-denoised) in the rest of this chapter.

3.2.2 Robust Minimum Variance Beamformer (RMVB)

In order to more explicitly consider the uncertainty in the forward modeling, let $\mathbf{R}_i = \{\mathbf{z} | (\mathbf{z} - \mathbf{k}_i)^T \mathbf{P}_i^{-1} (\mathbf{z} - \mathbf{k}_i) \leq 1\} = \{\mathbf{A}_i \mathbf{u} + \mathbf{k}_i | \|\mathbf{u}\| \leq 1\}$ be an M -dimensional ellipsoid centered at \mathbf{k}_i , with $M \times M$ matrices \mathbf{P}_i and \mathbf{A}_i determining its size and shape ($\mathbf{P}_i = \mathbf{P}_i \mathbf{P}_i^T$). This ellipsoid is assumed to cover all possible values for the lead field column of the i -th voxel, namely the uncertainty region (spanned by \mathbf{z} and \mathbf{u} $M \times 1$ vectors). To consider the uncertainty of the lead field, one idea is to enforce the spatial filter to pass the activities associated with not only \mathbf{k}_i , but also all values in the uncertainty region \mathbf{R}_i . Hence, the LCMV optimization problem (3.2) can be reformulated as

$$\begin{aligned} \mathbf{w}_i^* &= \arg \min_{\mathbf{w}_i} \mathbf{w}_i^T \mathbf{C} \mathbf{w}_i \\ s.t. \quad &\mathbf{w}_i^T \mathbf{z} \geq 1, \quad \forall \mathbf{z} \in \mathbf{R}_i. \end{aligned} \tag{3.8}$$

The constraint in (3.8), which includes infinite number of linear equations, is equivalent to:

$$\mathbf{w}_i^T (\mathbf{A}_i \mathbf{u} + \mathbf{k}_i) \geq 1, \quad \forall \mathbf{u} \quad s.t. \quad \|\mathbf{u}\| \leq 1 \quad (3.9)$$

Inequality (3.9) holds if and only if it holds for \mathbf{u}^* that minimizes the term $\mathbf{w}_i^T \mathbf{A}_i \mathbf{u}$. Cauchy-Schwartz inequality and the constraint $\|\mathbf{u}\| \leq 1$ lead to $\mathbf{u}^* = -\mathbf{A}_i^T \mathbf{w}_i / \|\mathbf{A}_i^T \mathbf{w}_i\|$ [111,116]. By substituting this value and some manipulations, the optimization problem (3.8) can be expressed as

$$\begin{aligned} \mathbf{w}_i^* &= \arg \min_{\mathbf{w}_i} \mathbf{w}_i^T \mathbf{C} \mathbf{w}_i \\ s.t. \quad \mathbf{w}_i^T \mathbf{k}_i &\geq 1 + \|\mathbf{A}_i^T \mathbf{w}_i\|, \quad \forall i \in \{1, \dots, N\} \end{aligned} \quad (3.10)$$

which is a second order cone programming (SOCP) problem [122,123]. After solving this problem, the solution can be normalized such that $\mathbf{w}_i^{*T} \mathbf{k}_i = 1$. Henceforth, we refer to this problem as the robust minimum variance beamformer (RMVB). While the RMVB does not enjoy a closed-form solution as the conventional LCMV, it can still be solved efficiently using any convex optimization solver such as CVX [124,125] (See Supplementary materials for more details on convex optimization at A.2.1). Additionally, the normalization and denoising steps can also be applied to the solution of RMVB as before. This will be referred to as the RMVB-ND-DN, hereafter.

3.2.3 Uncertainty Region Estimation

Multiple sources of uncertainty have to be considered in order to find the uncertainty region, e.g., uncertainty in the location (due to discretization) and orientation of the current dipoles or uncertainty in the forward model parameters (due to insufficient information about the head volume geometry, conductivities, inhomogeneity, etc.). In any case, it is possible to estimate the uncertainty region empirically by sampling the surrounding source space for each voxel [58] or by using several forward models, for instance, constructed for electrical conductivities in a given range [126]. More specifically, assuming S_i neighbors and F forward models, the first and second order statistics of the uncertainty for the i -th voxel can be estimated as

$$\bar{\mathbf{k}}_i = \frac{1}{F} \sum_{f=1:F} \mathbf{k}_i^f \quad (3.11)$$

and

$$\mathbf{Q}_i = \frac{1}{S_i \times F} \sum_{s=1:S_i} \sum_{f=1:F} (\mathbf{k}_i^{(s, f)} - \bar{\mathbf{k}}_i)(\mathbf{k}_i^{(s, f)} - \bar{\mathbf{k}}_i)^T \quad (3.12)$$

where \mathbf{k}_i^f is the lead field column of the i -th voxel in the f -th forward model, $\mathbf{k}_i^{(s, f)}$ is the lead field column of the s -th sampled neighbor in the f -th forward model, $\bar{\mathbf{k}}_i$ is the average lead field (of all $\mathbf{k}_i^{(s, f)}$) to be used in the inverse problem and \mathbf{Q}_i is the covariance matrix of the uncertainty. The authors of [58] proposed to apply a projection on each $\mathbf{k}_i^{(s, f)}$ before using them in (3.12). This projection in a fixed-orientation model, which is the case for this study, translates to flipping the sign of $\mathbf{k}_i^{(s, f)}$ (or equivalently the orientation of dipoles), wherever this reduces the distance between $\mathbf{k}_i^{(s, f)}$ and $\bar{\mathbf{k}}_i$. This is helpful in avoiding overestimation of the uncertainty region size. \mathbf{P}_i , which defines the shape and size of the uncertainty ellipsoid, is indeed an inflated version of \mathbf{Q}_i ($\mathbf{P}_i = \alpha \mathbf{Q}_i$). The inflation factor α_i can be found such that the uncertainty region \mathbf{P}_i contain all the points $\mathbf{k}_i^{(s, f)}$ [116], which mathematically translates to

$$\alpha = \sup_{s, f} (\mathbf{k}_i^{(s, f)} - \bar{\mathbf{k}}_i) \mathbf{Q}_i^{-1} (\mathbf{k}_i^{(s, f)} - \bar{\mathbf{k}}_i)^T \quad (3.13)$$

See Practical Considerations section A.2.3 in the Supplementary materials for further details of tuning uncertainty ellipsoids sizes.

3.2.4 Computer Simulation Protocol

In order to investigate the performance of robust beamformer and compare its merits and limitations to the conventional beamformer, a series of computer simulations were conducted. To this end, a realistic head volume model was built upon the Montreal Neurological Institute Colin brain [65] consisting of three layers i.e., the scalp, the skull and the brain. To solve the forward problem and to simulate the EEG recordings, a standard 128-channel BioSemi cap was fitted to the Colin brain, and a boundary element method (BEM) model [5, 67] was then derived to find the lead field matrix, which linearly projects the current dipoles to the electrical potentials at the electrodes.

In order to avoid any form of inverse crime and to evaluate the capabilities of robust beamformer in dealing with model violations, different models were derived for the forward and inverse problems. While cortex was meshed very finely with a grid of 1 mm ($\sim 131,000$ elements on the cortex, in total) in the forward problem, a coarser

grid of 5 mm ($\sim 9,000$ elements on the cortex, in total) was used for the inverse problem. In the forward model, the electrical conductivity of the scalp, the skull and the brain (σ_{scalp} , σ_{skull} and σ_{brain}) were set to 0.33 S/m, 0.022 S/m and 0.33 S/m, respectively, while an average model based on what follows was built to solve the inverse problem. For estimation of uncertainty regions only the conductivity and discretization uncertainties were considered. To this end, 11 different models were built. The electrical conductivities of the scalp and the brain were fixed to 0.33 S/m in all the models, while for the skull it was set such that the conductivity ratio ($\sigma_{scalp}/\sigma_{skull}$) picked values incrementally with the steps of 1 in the physiological range of 15 – 25 [68–70]. The brain in each model was meshed with two different choices of grid sizes; a same coarser grid of 5 mm for the center of uncertainty regions as in the inverse problem and a very fine grid of 1.1 mm to estimate the variations for the nearest neighbor in the coarser mesh according to (3.12). In this setting, each point of the fine mesh was considered as a neighbor of one and only one point in the coarser mesh. Note that different fine meshes were used in the forward problem and estimation of uncertainty regions (grid of 1 mm versus 1.1 mm); as also in real situations the true source space is unavailable beforehand.

In this study, four different scenarios were simulated. The first scenario consisted of 100 point dipoles placed randomly on the cortex with a fixed orientation normal to the cortex surface. To mimic a realistic inter-ictal spike, each dipole was assigned a time course of activity sampled at the rate of 1 kHz (see **Figure A.3** in the Supplementary materials). In order to investigate the effect of number of active nodes on the imaging accuracy, in the second scenario, each configuration had three nodes of activity. The time-courses in this scenario, were such to yield a minimal correlation between the nodes (see **Figure A.4 (A)** and **(B)** in the Supplementary materials). In the third scenario, the time-courses had high correlation factors to consider the effect of correlation between nodes in 3-node networks (see **Figure A.4 (A)** and **(B)** in the Supplementary materials). Please refer to the Supplementary materials to review the results of this scenario. The last scenario was designed to evaluate the performance of robust and conventional beamformers in situations, where sources are not focal. The general criteria in this scenario remained the same as the second scenario with the exception that each node had an extent with a radius size roughly ranging from 10 mm to 30 mm (randomly

selected for each node of the network). In this study, all the voxels within the extent of each node had the same amplitude and the same time-course. After solving the forward problem, the generated potentials at the electrodes were contaminated by additive white Gaussian noise (AWGN) to simulate a more realistic condition. The effect of noise was further assessed by considering two different SNRs (calculated based on power) i.e., 5 and 20 dB representing low and high levels of noise. Finally, a fixed-orientation (normal to the cortex surface) model was used to solve the inverse problem. This choice can be justified by the fact that EEG signals are generated by pyramidal cells, which project their dendrites orthogonally to the cortex surface [127]. The simulation protocol in this study is similar to our previous works [73, 105].

3.2.5 Performance Measures

Performance in this study was assessed by different metrics, namely dipole localization error (DLE), signal to noise ratio (SNR), receiver operator characteristics (ROC) curve and Matthew’s correlation coefficient (MCC). The former two metrics were used for focal activity scenarios, while the latter two evaluated the conventional and robust beamformers in the extended source scenario. DLE measures the Euclidean distance between the simulated and recovered source locations, which are determined by estimating the foci of activity employing a PCA-based technique (dominant local maxima of the principal components as the foci of activity) [73, 105]. The SNR characterizes how well the temporal profile of the underlying sources can be recovered by ESI. Noise in this metric is defined as the difference between the simulated and recovered time courses. The power of the simulated signal over the power of noise in logarithmic scale yields the SNR in dB. In order to evaluate the concordance between temporal patterns, both the simulated and recovered signals were normalized by their power before SNR was calculated. The time course of the estimated locations, and not the true source locations, were considered in calculating the SNR. In order to compare the simulated and estimated source distributions in the extended source scenario, ROC curve was used to evaluate the trade-off between sensitivity and specificity of the estimation by plotting the true positive rate (TPR) versus false positive rate (FPR) [128] of the estimation. Any location in the solution beyond a threshold is considered a source and active, while any location with an amplitude below the threshold is inactive and not a

source. Thresholds are varied, and the FPR and TPR are calculated accordingly (each point of curve is at a different threshold). Area under the curve (AUC) is the area under the ROC curve; a metric for comparing different ROC curves. The closer the AUC of a ROC curve is to 1, the more accurate the estimation is [129]. MCC, which in essence measures the correlation between the simulated and estimated distributions, is another metric used in this study to assess the estimation quality [130]. MCC calculates the predictability of an observed and estimated classification, which in this case is the source extent, and is formulated as follows:

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (3.14)$$

where, TP, TN, FP and FN represent the true positive, true negative, false positive and false negative rates of the estimation, respectively. Since MCC works in a binary mode, it requires a thresholding mechanism to turn the estimated distribution into active and inactive regions. In this study, the threshold was swept and determined in such a way to yield the maximum MCC. We note that MCC takes values from -1 to $+1$, from the most disparate case to the most concordant case, respectively.

3.3 Results

The results presented in this section are to evaluate the effect of normalization and denoising procedures in addition to modeling the lead field uncertainty on the ultimate performance of beamformers. To this end, four types of beamformers namely, LCMV [108], RMVB [111, 116], LCMV-ND-DN [110, 114] and RMVB-ND-DN were tested.

3.3.1 Focal Sources

Figure 3.1 depicts the performance of the mentioned beamformers in situations, where the underlying source activity can be modeled with a single dominant current dipole. In order to find the effect of noisy measurements on the performance, the simulations were repeated for high and low noise levels associated with the SNRs of 5 dB and 20 dB, respectively. **Figure 3.1 (A)** depicts the median DLE for the 100 dipoles as simulated in the first scenario. To present the uncertainties in the results, the first and third quartiles of data are reported with error bars. These results indicate that the

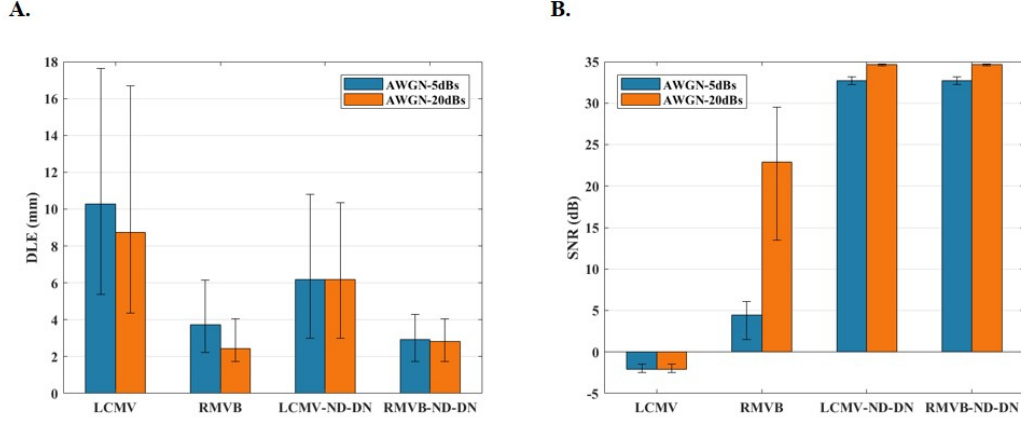


Figure 3.1: The Monte Carlo simulation statistics for the single-node source scenario. (A) The median DLE results for SNRs of 5 dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartile (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration.

robust adaptive beamformers, i.e., the RMVB and RMVB-ND-DN, outperform their conventional counterparts i.e., the LCMV and LCMV-ND-DN, respectively. Additionally, including the normalization and denoising procedures plays an important role in the quality of the results. **Figure 3.1 (B)** yields the evaluation in terms of output SNR. According to this plot, while both the LCMV-ND-DN and RMVB-ND-DN outperform the LCMV and RMVB, they both perform equally well in recovering the temporal profile of the simulated activity. It can also be observed that including the post-processing steps can considerably decrease the sensitivity to noise levels. In fact, the gap between the RMVB performance in low and high levels of noise compared to the methods with post-processing of the solution, can be justified with regard to its lack of appropriate measures or mechanisms to remove or counteract the effect of noise. Furthermore, since the LCMV does not take into consideration the depth and asymmetric spatial distribution of the noise on one hand and is not designed to be robust against modeling errors on the other hand, its performance is especially poor when estimating output SNRs.

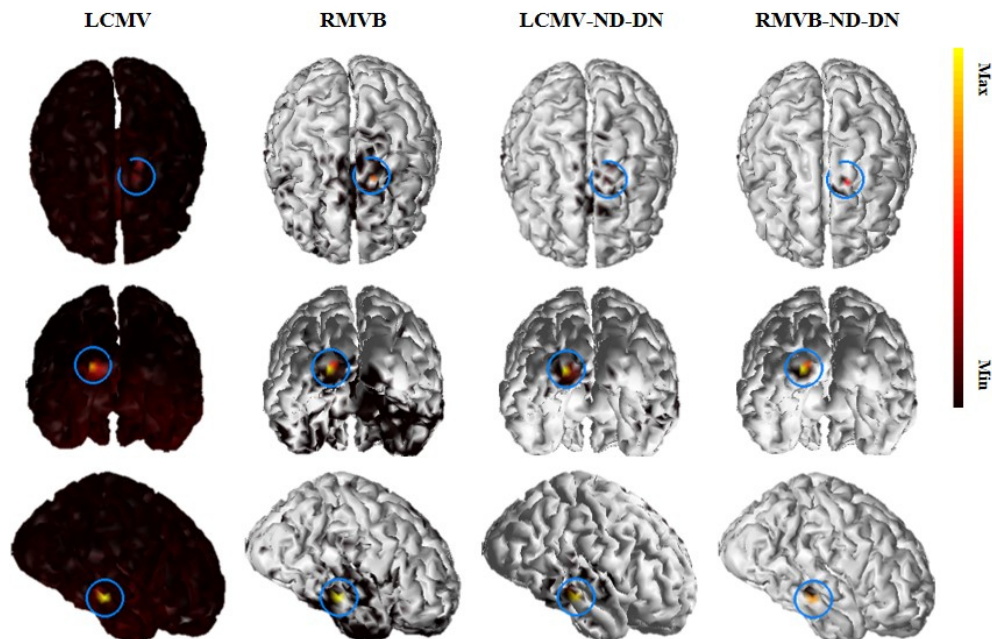


Figure 3.2: An example of a 3-node network with focal activity. The source imaging results for a single 3-node source with focal activity using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution is thresholded with a cut-off value of 0.01. The blue rings of the plots mark the positions of the true source.

3.3.2 The Effect of Number of Active Nodes

In the second scenario, 100 source configurations each containing three active nodes were considered. **Figure 3.2** shows an example, where the three simulated current dipoles are placed at the right temporal, left frontal and left occipital lobes, respectively (the center of the blue rings on the plots mark the exact positions). SNR in this example is set to 20 dB. Following the localization procedure discussed in the Computer Simulation Protocol section, it can be found that all versions except the LCMV are able to pinpoint the true location of the simulated dipoles in this example. However, the solution of RMVB-ND-DN is more focal and thus more consistent with the underlying dipole source, while the LCMV yields the smoothest solution, which is not as focal as the RMVB. It should be mentioned that, for display purposes, the threshold was set to 1% in this figure. This implies that e.g., in the solution of RMVB-ND-DN almost all falsely recovered dipoles

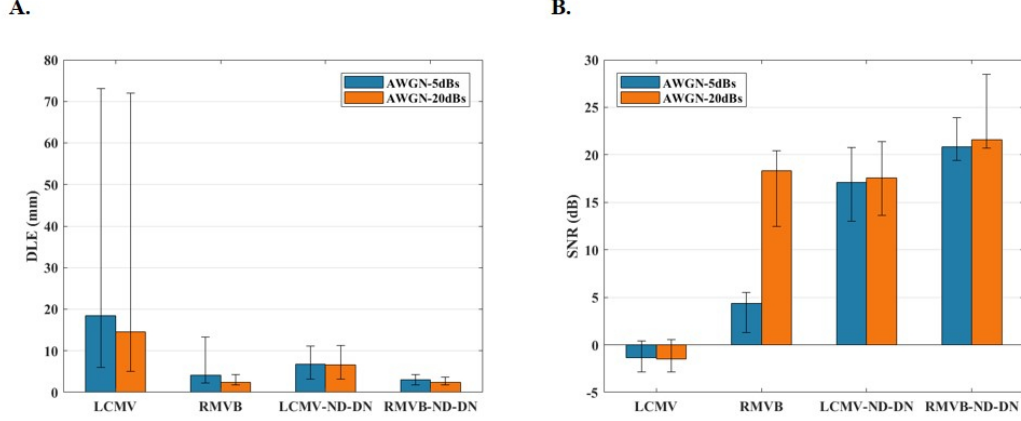


Figure 3.3: The Monte Carlo simulation statistics for the three-node uncorrelated source scenario. (A) The median DLE results for SNRs of 5 dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartile (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration.

are at least 100 times weaker than the strongest activity.

Figure 3.3 (A) and **(B)** plot the median DLE and SNR along with the corresponding first and third quartile error bars for two noise levels. The LCMV is sensitive to the number of sources in terms of DLE, while the other three techniques show less sensitivity. As in the earlier scenarios, normalization and denoising play a crucial role in the performance of both conventional and robust beamformers. However, robust beamformers are successful in modeling the uncertainty, even without these normalization and denoising steps. When such processing is applied, the results indicate that it is still better to use the robust version instead of the conventional beamformer, if computational resources are available. Furthermore, it can be seen that the RMVB is sensitive to input noise level especially when evaluating the output SNR. Please refer to the Supplementary materials to review the role of normalization in removing the inherent bias of adaptive beamformers toward deeper locations as well as the effect of correlated activities on beamformers performance.

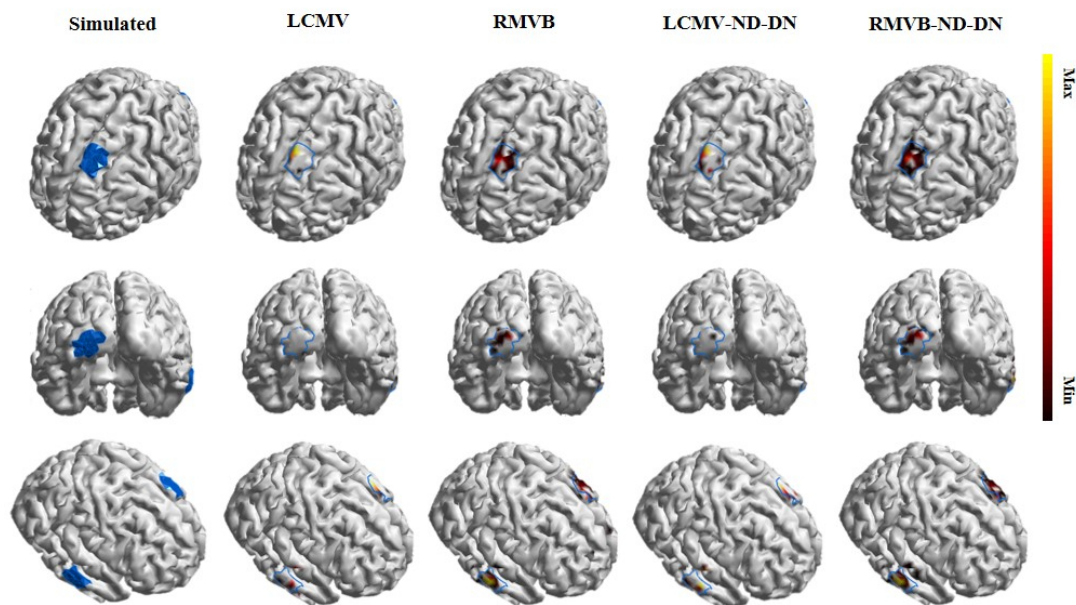


Figure 3.4: An example of a three-node network with extended activity. The simulated three-node patches of activity as well as the source imaging results using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution of each technique is thresholded separately such that MCC is maximized for that technique. The blue regions and boundaries of the plots mark the true source extents.

3.3.3 Extended Sources

The fourth scenario was designed to evaluate the aforementioned beamformers in determining the distribution and shape of the underlying activities. To this end, **Figures 3.4** and **3.5** are to simulate three patches of extended sources rather than point dipoles. The blue boundaries in these plots mark the true source extents. The area of each patch is roughly 750 mm^2 and again SNR is 20 dB. In **Figure 3.4**, the threshold of display is optimized for each technique separately such that MCC is maximized, while in **Figure 3.5**, it is preset to an arbitrary small value of 1%. The main observation is that in both cases, the RMVB-ND-DN captures the extent of the underlying activity while keeping number of falsely recovered dipoles small. Whereas, regular beamformers either underestimate the true source extent (**Figure 3.4**) or have high false positive




Figure 3.5: An example of a three-node network with extended activity. The simulated three-node patches of activity as well as the source imaging results using four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. SNR in this example is set to 20 dB and the solution is thresholded with a preset cut-off value of 0.01. The blue regions and boundaries of the plots mark the true source extents.

rates (**Figure 3.4**). Since the optimal value of the threshold to detect the source extent is unknown a priori, robustness to threshold is highly valuable in practice.

Figure 3.6 plots the average ROC curve for 100 source configurations each containing three extended nodes of activity for SNRs of 5 and 20 dB in the left and right panels, respectively. Each point on a ROC curve (TPR versus FPR) is associated with a cut-off threshold, the values below which were set to 0. The corresponding AUC of all the curves are reported in the legend of this figure. Obviously, the RMVB-ND-DN better estimates the extent of the underlying source compared to other versions of beamformers according to this metric. Moreover, the results of the RMVB and LCMV-ND-DN are close to each other (the RMVB is slightly superior to the LCMV-ND-DN). **Figure 3.7** reports the median MCC for the same data. The threshold for each source is determined such that MCC is maximized for that specific configuration. As it can be

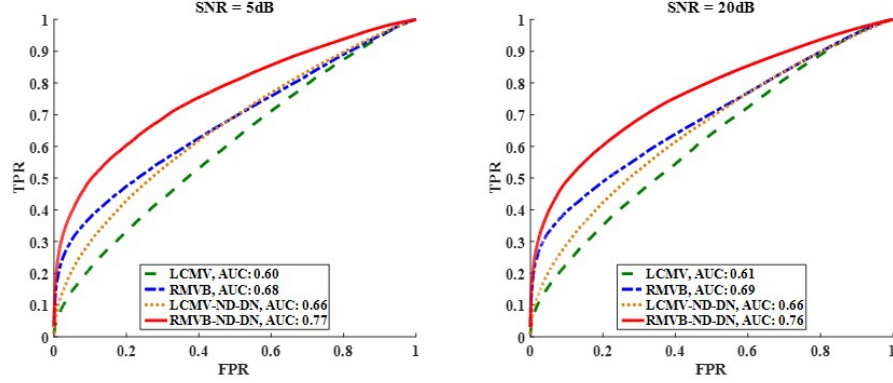


Figure 3.6: The Monte Carlo simulation statistics for the extended source scenario. The Average ROC curves for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN and for SNRs of 5 and 20 dB in the left and right panels, respectively. The corresponding AUC of all the curves are reported in the legend.

seen, the trend of this figure is the same as **Figure 3.6**, with the RMVB-ND-DN at the top followed by the RMVB, LCMV-ND-DN and LCMV, respectively. It should also be noted that the absolute values of MCC in this figure do not provide much information about the quality of the estimation, and that the goal of this figure is to provide only a benchmark to compare different beamformers' performance. Finally, interested readers are referred to **Figures A.9** and **A.10** in Supplementary materials for the results of beamformers with diagonal loading.

3.4 Discussion

In this study, we proposed a new technique named robust minimum variance beamformer (RMVB), which unlike traditional adaptive beamformers, explicitly takes into consideration the uncertainty of the forward models. In order to compare the performance of RMVB with the well-known LCMV adaptive beamformer, we performed Monte Carlo simulations in various scenarios representing different conditions for underlying source activity. We also simulated the modified versions of each technique by applying the post-hoc normalization/denoising step (i.e., LCMV-ND-DN and RMVB-ND-DN techniques). Based on simulation results, RMVB and RMVB-ND-DN can better estimate the underlying source activities compared to LCMV and LCMV-ND-DN respectively;




Figure 3.7: The Monte Carlo simulation statistics for the extended source scenario. The median MCC results along with the first and third quartile error bars for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN and for SNR of 20 dB.

although, LCMV-ND-DN yields favorable results in certain cases as well. As a general observation, both post-processing of the solutions (normalization and/or denoising steps) and modeling the uncertainty (the robust versions) enhance the performance. By combining both procedures in RMVB-ND-DN, our technique outperformed other versions based on all metrics that were used in this study. It was also observed that there is especially a clear benefit in using robust beamformers to estimate extended sources (as depicted in **Figures 3.6** and **3.7**).

Although the optimization problem (3.10) can be solved easily using general-purpose convex solvers such as CVX [124,125], RMVB can be much slower than the conventional LCMV, since the optimization problem has to be solved for many voxels. While CVX was used in this study, RMVB can be significantly accelerated by employing a specifically tailored algorithm for the optimization module. Fortunately, this algorithm, which is based on Lagrange multipliers method, exists and can solve the problem much more efficiently and faster than the CVX package [116]. Furthermore, parallel computing can also be exploited by virtue of the scanning scheme of the beamformers, yielding further acceleration (by estimating the weights for different voxels in parallel). Overall, it takes about 7 seconds to find the solution of RMVB using parallel computing (in MATLAB and on a server computer with 8x32 GB of RAM and a 2x Intel Xeon E5-2697v2 2.7GHz

processor) which is approximately 20 times slower than the conventional LCMV. This amount of time is reasonably short for applications where source imaging is performed off-line. Besides optimization, estimation of uncertainty regions can be computationally demanding as well, since it requires building several forward models depending on the source of uncertainty in the actual problem. However, these regions are computed only once before source imaging is performed. Furthermore, the uncertainty region estimation at each location is independent from other locations. Thus, these can be computed in a parallel fashion (as in solving the optimization problem), resulting in efficient implementations of beamformer approaches. Such reduction in computation times due to parallel computing makes robust beamformer approaches desirable and practically achievable.

Chapter 4

Accelerated MRI Reconstruction with Database-Free Deep Learning

4.1 Introduction

Coronary artery disease (CAD) is the leading cause of death in the United States, accounting for one in seven deaths [26]. Coronary MRI provides a non-invasive and radiation-free diagnostic tool for CAD assessment [27], with a potential for repeated use. It is typically acquired with electrocardiogram (ECG) triggering during diastolic quiescence, where $\sim 30 - 35$ k-space lines are sampled per R-R interval [131–133]. When imaging the right coronary artery in a targeted manner [131], this leads to a ~ 3 minute nominal scan time. Since this scan time necessitates a free-breathing acquisition [134, 135], respiratory motion compensation needs to be applied [132, 133], typically with navigator gating [133, 136], which further reduces the efficiency of the scans by $\sim 2 - 3$ fold, leading to a scan time of $\sim 6 - 10$ minutes. Alternatively, coronary MRI can be acquired with whole-heart coverage, which leads to a higher signal-to-noise ratio (SNR) [137, 138], albeit at a longer nominal acquisition time of $6 - 8$ minutes. The additional scan time overhead due to respiratory motion compensation often requires accelerated acquisitions, necessitating a trade-off with SNR [137, 139].

Several strategies have been used to accelerate coronary MRI acquisitions such as parallel imaging [140, 141], compressed sensing [142–144], and their combinations [145–151]. Recently, deep learning has emerged as an alternative means to accelerate MRI reconstruction [32–46]. In contrast to conventional reconstruction methods, deep learning requires large datasets of fully-sampled data which may be unavailable due to prohibitively long acquisition time e.g., in whole-heart coronary MRI. An alternative line of work considers subject-specific application of neural networks [32, 152]. In this approach, called robust artificial-neural-networks for k-space interpolation (RAKI), several convolutional neural networks (CNN) are calibrated from scan-specific autocalibrating signal (ACS) data for improved interpolation of missing k-space lines. Thus, this method extends the linear convolutions used in GRAPPA [29], and was shown to increase noise resilience for uniform undersampling patterns, especially in low-SNR and high-acceleration rate regimes [32]. However, previous work has shown the benefits of random undersampling in high-resolution three-dimensional (3D) coronary MRI, for instance in the setting of compressed sensing [145]. For such undersampling patterns, iterative self-consistent parallel imaging reconstruction (SPIRiT) [153] provides a k-space interpolation approach for multi-coil data. SPIRiT utilizes multi-coil information by including a self-consistency term that ensures the interpolated k-space is consistent with itself according to the calibration kernels, along with a data-consistency term in reconstruction. SPIRiT requires iterative processing in the reconstruction and is consequently more computationally-intensive than GRAPPA.

In this chapter, we exploit the notion of coil self-consistency in SPIRiT to enable RAKI with arbitrary undersampling. The proposed technique, called self-consistent RAKI (sRAKI), is evaluated in targeted and whole-heart coronary MRI, and compared with SPIRiT and ℓ_1 -SPIRiT at various acceleration rates.

4.2 Methods

All imaging protocols were approved by the University of Minnesota institutional review board. Written informed consent was obtained from all participants before each examination.




Figure 4.1: The CNN architecture to learn and enforce the coil self-consistency rule. The number of layer output channels is denoted by depth of blocks. All layers, except the last one, were followed by rectifier linear units (ReLU) as activation functions. The kernel sizes of the layers were 5×5 , 3×3 , 3×3 and 5×5 , respectively. Each layer had 16, 8, 16 and $2n_c$ output channels, respectively. The 3D k-space data was first inverse Fourier transformed along fully-sampled k_x dimension. Subsequently 2D convolutional kernels were jointly trained on the ACS region of resultant 2D slices of data to learn the self-consistency rule.

4.2.1 Calibration

For multi-coil k-space data with n_c coils, a k-space point in the j th coil, $x_j(k_x, k_y, k_z)$ can be estimated as a function of distinct k-space points from all coils $i \in \{1, \dots, n_c\}$ within a neighborhood region of (k_x, k_y, k_z) [29, 153]. In linear parallel imaging techniques, this function is modeled by a linear spatially shift-invariant convolution, and the convolutional kernels can be found by solving n_c linear least squares optimization problems [29, 153]. In particular, SPIRiT uses these linear convolutional kernels to define a coil self-consistency rule that connects all the k-space elements with neighboring elements across all coils. However, it has been noted that a nonlinear mapping may be advantageous from a noise reduction perspective due to two factors. First, the shape and size of the neighborhood is heuristically set in practice [32], which may not capture all the required dependencies. Second, in contrast to typical least squares optimization problems, both the target and source points for the kernels in calibration are contaminated with noise, and nonlinear functions have been shown to deal more effectively with such imperfections [32, 154]. Thus, we propose to utilize CNNs that are calibrated (the terminology used for finding the self-consistency rule using ACS data) on ACS data of a single scan only to nonlinearly model the self-consistency in multi-coil k-space data.

In this study, a 4-layer CNN architecture was employed to learn the self-consistency

rule among coils (**Figure 4.1**). In contrast to conventional RAKI, where separate CNNs were used for mapping to individual coils, a single CNN was used to map from all coils of multi-coil k-space onto itself, facilitating considerably reduced run time. For reduced computational complexity, 3D k-space data was first inverse Fourier transformed along fully-sampled k_x dimension. Subsequently 2D convolutional kernels were jointly calibrated on the resultant 2D slices of data [155]. The k-space data across all coils were normalized to have unit power as a preprocessing step to enable the use of a fixed learning rate. In addition, the complex k-space data was embedded to the real field, by concatenating the real and imaginary components of k-space along the coil dimension leading to $2n_c$ input and output channels. All layers, except the last one, were followed by rectifier linear units (ReLU) as activation functions. The kernel size at input and output layers was 5×5 , while the hidden layers used 3×3 kernels. The number of output channels of different layers was 16, 8, 16 and $2n_c$, respectively. The network was designed such that the middle layer narrowed down to fewer output channels [156], since the CNN is prone to learn a trivial identity mapping from identical input and output training data, otherwise. A zero-padding strategy was used at each layer to maintain the size of input at the output after convolution. The network was trained by minimizing a MSE objective function using Adam optimizer [157]. A learning rate of 0.01 and maximum number of iterations of 1000 were used in training.

4.2.2 Reconstruction

After calibrating the CNN on ACS data to learn the coil self-consistency rule, the following objective function is minimized to reconstruct k-space:

$$\arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{D}\mathbf{x}\|_2^2 + \beta \|\mathbf{x} - \mathbf{G}(\mathbf{x})\|_2^2, \quad (4.1)$$

where \mathbf{x} is the reconstructed k-space data across all coils, \mathbf{y} is the noisy acquired data, \mathbf{D} is the undersampling operator and $\mathbf{G}(\cdot)$ represents the calibrated CNN for self-consistency. The first term in the objective function in (4.1) ensures that the reconstructed k-space is consistent with acquired data. The second term enforces self-consistency in the reconstructed k-space according to the coil self-consistency rule that was learned by calibrating on the ACS data. The parameter β determines the balance between these two terms. Note that the main difference between sRAKI and RAKI is in

this phase, where RAKI performs a one-time application of calibration kernels to estimate the missing data, whereas sRAKI requires iterative optimization of Equation (4.1). Additional regularization terms can also be incorporated in (4.1), although this was not investigated in the current study to maintain the focus on multi-coil data processing.

The objective function in (4.1) was optimized using the Adam optimizer with a tuned learning rate of 2, for the same k-space normalization to unit power as before. We note that Adam only requires the gradient of the objective function with respect to the optimization variable \mathbf{x} , which is the network input rather than network parameters in the reconstruction phase. Similar to network parameters, gradients with respect to the input were efficiently calculated through back-propagation by using the deep learning package Keras. In order to avoid a heuristic tuning of β , consistency with data was strictly enforced as in SPIRiT [153]. This led to gradients being calculated for non-acquired elements only while the rest of k-space was directly replaced with acquired data at each iteration. For comparison, SPIRiT using a conjugate gradient reconstruction was implemented with a 5×5 kernel [153]. ℓ_1 -SPIRiT was also implemented with additional a Daubechies-wavelet regularization [153], where the thresholding parameter was empirically tuned to 0.0005 of the maximum absolute wavelet coefficient. The number of reconstruction iterations were tuned separately for each technique and was set to 50 for SPIRiT and sRAKI, and 15 for ℓ_1 -SPIRiT. We note that ℓ_1 -SPIRiT converges at a faster rate, thus necessitating fewer iterations. All hyper-parameters were selected to optimize the performance of each technique both qualitatively and quantitatively.

4.2.3 Targeted Coronary MRI

All imaging was performed on a 3T Siemens Magnetom Prisma (Siemens Healthineers, Erlangen, Germany) system with a 30-channel receiver body coil-array. Targeted right coronary artery (RCA) MRI was acquired on 6 healthy subjects (26.7 ± 2.9 years, 3 women). Scout images were followed by axial breath-hold cine bSSFP images to identify the quiescent period of the RCA, which was used for the trigger delay of coronary acquisitions. A low-resolution free-breathing ECG-triggered 3D coronary survey was acquired for slab orientation of the RCA imaging. Targeted RCA MRI was then acquired with a free-breathing ECG-triggered GRE sequence with imaging parameters: TR/TE = 3.4/1.5ms, flip angle = 20° , bandwidth = 601 Hz/pixel, field-of-view (FOV) = $300 \times$

$300 \times 48 \text{ mm}^3$, resolution = $1 \times 1 \times 3 \text{ mm}^3$, navigator window = 5 mm. The nominal scan time was 160 seconds at a heart rate of 60 bpm. T_2 -preparation and a spectrally-selective fat saturation were utilized for improved contrast.

The 3D k-space data was exported and retrospectively under-sampled with a Poisson disc pattern at acceleration rates 2, 3, 4, and 5 with a fully-sampled 40×10 ACS region in $k_y - k_z$ plane. These under-sampled data were then reconstructed using SPIRiT, ℓ_1 -SPIRiT and sRAKI for comparison, with the implementations detailed above. Final images were obtained using root-sum-squares combination of all coil images. All algorithms were implemented in Python, and processed on a workstation with an Intel E5-2640V3 CPU (2.6GHz and 256GB memory), and an NVIDIA Tesla V100 GPU with 32GB memory. Additional comparisons for uniformly under-sampled data are shown in Supporting Information **Figure A.11**, while different random undersampling patterns, and their reconstructions are depicted in Supporting Information **Figure A.12** and **Figure A.13**, respectively.

4.2.4 Image Analysis

Quantitative analysis of the reconstructions was performed using normalized mean square error (NMSE) with respect to the fully-sampled reference, as well as normalized vessel sharpness measurements. NMSE was calculated in image domain between a given reconstruction method and the fully-sampled reference, normalized by the energy of the reference. Vessel sharpness scores were calculated for both sides of the vessel using a Deriche algorithm [158]. Normalized vessel sharpness was calculated as the average score of both sides divided by the intensity at vessel center. A normalized vessel sharpness value closer to 1 represents a sharper vessel border. The NMSE and normalized vessel sharpness measurements of the different reconstructions were statistically compared across subjects using a nonparametric sign-rank test for each acceleration rate. A p-value of < 0.05 was considered significant.

4.2.5 Whole-Heart Coronary MRI

Prospectively under-sampled whole-heart coronary MRI was acquired on an additional subject (28 years, male) at an acceleration rate of 5 with a Poisson disc pattern. The

same sequence parameters were used with $\text{FOV} = 300 \times 300 \times 106 \text{ mm}^3$, resolution = $1.2 \times 1.2 \times 1.2 \text{ mm}^3$. The data were then reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI for comparison, with the same implementations described above. We note that this scenario poses a challenge for traditional machine learning algorithms that perform supervised learning on databases of fully-sampled data, as it is difficult to acquire high-quality fully-sampled whole-heart coronary MRI data. This is due to the long scan time of a fully-sampled acquisition, which leads to quality degradation due to drift and changes in the motion patterns. We also note that there have been some recent efforts to acquire fully-sampled whole-heart coronary MRI for this purpose, even though the acquisition time remains long [159].

4.3 Results

Figure 4.2 depicts reformatted images from a targeted coronary MRI dataset reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI techniques at retrospective acceleration rates 2, 3, 4, and 5. RCA is visualized at all rates for all methods. sRAKI has visibly less noise at high acceleration rates compared to SPIRiT and fewer blurring artifacts compared to l_1 -SPIRiT. The reformatted images from a second subject, are shown in **Figure 4.3** with similar results showing that sRAKI has visibly less noise at high acceleration rates. sRAKI demonstrates improved quality at higher acceleration rates, reducing noise amplification and blurring artifacts compared with other reconstruction methods.

Figure 4.4 Summarizes the mean and standard deviation of the NMSE and normalized vessel sharpness measurements for SPIRiT, l_1 -SPIRiT and sRAKI across all subjects. sRAKI improves mean NMSE by 34%, 30%, 39%, 44% compared to SPIRiT, and 18%, 21%, 21% and 21% compared to l_1 -SPIRiT for rates 2, 3, 4 and 5, respectively. Statistical analysis confirms that sRAKI significantly improves NMSE at all acceleration rates over both SPIRiT and l_1 -SPIRiT. In terms of normalized vessel sharpness, sRAKI provides 7%, 9%, 11%, 10% improvement compared to SPIRiT and 4%, 5%, 13% and 20% improvement compared to l_1 -SPIRiT for rates 2 to 5, respectively. The improvements over SPIRiT are statistically significant at rates 2 and 3, while improvements over l_1 -SPIRiT are statistically significant at rates 3 – 5.




Figure 4.2: Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were retrospectively under-sampled at rates 2, 3, 4, and 5 in the $k_y - k_z$ plane and then reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI (top, middle and bottom rows). Fully-sampled images are also displayed in the first column as a reference for comparison. sRAKI is visually more robust to noise amplification and blurring artifacts at high acceleration rates compared to SPIRiT and l_1 -SPIRiT, respectively. (RCA: right coronary artery; AO: Aortic Root)

Figure 4.5 depicts the results of a prospectively 5-fold accelerated whole-heart coronary imaging. sRAKI yields an improved visualization of both the RCA and the left circumflex artery (LCX) compared to SPIRiT and l_1 -SPIRiT. The normalized vessel sharpness measurements for this subject were 0.30, 0.31 and 0.33 for RCA and 0.25, 0.22, 0.28 for LCX with SPIRiT, l_1 -SPIRiT and sRAKI reconstructions.




Figure 4.3: Reformatted right coronary artery (RCA) images from another 3D targeted coronary MRI dataset. This data was also retrospectively under-sampled at rates 2, 3, 4, and 5, and fully-sampled images are shown in the first column as reference. The difference between SPIRiT and sRAKI is visually evident at all acceleration rates for this subject with more apparent noise amplification. Furthermore, compared to l_1 -SPIRiT, sRAKI is more robust to blurring artifacts with increasing acceleration rates. (RCA: right coronary artery; AO: Aortic Root)

4.4 Discussion

In this study, we proposed a novel reconstruction method called sRAKI to accelerate coronary MRI. sRAKI trained subject-specific CNNs to learn a nonlinear coil self-consistency rule for multi-coil k-space data. In the reconstruction phase, this learned self-consistency rule was enforced along with data-consistency constraints, similar to SPIRiT reconstruction. Thus, sRAKI enabled reconstruction with arbitrary under-sampling patterns, an extension to RAKI [32], which was designed to handle uniform undersampling patterns only. A nonlinear calibration may capture further dependencies for learning coil self-consistency rule more effectively, when the shape and size of the

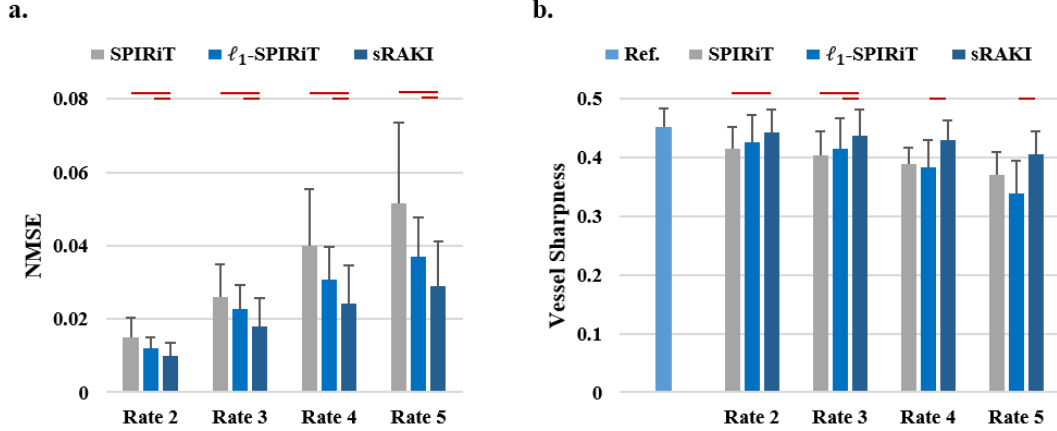


Figure 4.4: (a) Mean normalized mean squared error (NMSE) and (b) quantitative normalized vessel sharpness measures for all reconstructions of rates 2 to 5. Error bars represent standard deviation across subjects. sRAKI outperforms SPIRiT and ℓ_1 -SPIRiT at all rates for both metrics. The improvements in NMSE are statistically significant at all rates over both SPIRiT and ℓ_1 -SPIRiT, whereas the improvements in vessel sharpness with sRAKI are significant at rates 2 and 3 over SPIRiT, and rates 3 – 5 over ℓ_1 -SPIRiT. Red lines mark significant differences in the graphs.

neighborhood is heuristically set [32] and both the target and source points for the kernels in calibration are contaminated with noise [32, 154]. In this study, this translated to considerable reduction of reconstruction noise compared to SPIRiT. In contrast to the recent machine learning-based MRI techniques [33–46], which require large training datasets, sRAKI is trained on scan/subject-specific ACS data.

We note that there have been other methods for performing k-space interpolation using machine learning [32, 44, 160]. In DeepSPIRiT [44], multi-coil k-space data from a training database is first pre-processed with coil compression to yield a similar number of channels. Subsequently, CNNs are trained for different regions of k-space, which are then applied in a multi-resolution approach. This method was shown to reduce aliasing, while difficulty with high-resolution content was also noted. Because this method uses a training database, it still requires fully-sampled training data, which is difficult to apply in whole-heart coronary imaging, and thus differs fundamentally from the sRAKI approach. An alternative method that more closely matches sRAKI was proposed in [160]. This method, proposed independently after the initial presentation of our work




Figure 4.5: Reformatted coronal image from a prospectively 5-fold accelerated whole-heart coronary MRI dataset. The results show similar characteristics to targeted coronary MRI, where sRAKI reduces blurring with respect to l_1 -SPIRiT, and noise amplification with respect to SPIRiT.

[161–164], extends the AC-LORAKS approach to CNNs, in a manner similar to RAKI [32] and sRAKI. LORAKI has shown promising results in brain imaging, and has not been studied in the context of coronary MRI.

Several modifications were made to RAKI [32]. First, RAKI employed separate CNNs to learn nonlinear mapping functions from zero-filled multi-coil k-space data to missing data of individual coils. Therefore, $2n_c$ CNNs were trained to learn a full mapping function from multi-coil data to itself. In the new setting, we exploited a single CNN with more hidden layers to learn the coil self-consistency rule jointly, considerably reducing run time. Second, RAKI was examined in only 2D scenarios, whereas sRAKI was implemented for 3D datasets with two phase encoding dimensions. Another major difference is concerned with the reconstruction phase in which RAKI interpolates missing data with no iterations, but sRAKI optimizes an objective function to enforce data-consistency and self-consistency among coils. This procedure, which is similar to the reconstruction phase of SPIRiT, increases the computational burden by requiring first-order derivative calculation in each iteration. However, the extra complexity is not limiting. In this study, calibration on targeted right coronary artery datasets took ~ 20

seconds for SPIRiT and l_1 -SPIRiT, and ~ 40 seconds for sRAKI all on GPU implementations, although none of the implementations were fully optimized. In addition, the reconstruction phase on GPU took ~ 220 , 120 and 100 seconds for SPIRiT, l_1 -SPIRiT and sRAKI, respectively. sRAKI is also different from its previous version, in which a different reconstruction optimizer was used for 2D imaging [164]. In addition the 4-layer architecture of SPIRiT-RAKI in [161] is different from sRAKI, as the former applies 3D kernels on the whole 3D volume rather than training a single network with 2D kernels on 2D slices of the 3D volume. We observed that the latter further improves training by reducing the number of trainable parameters for the same amount of data.

In this study, the CNN parameters including the number of layers, the number of layer output channels and kernel sizes were empirically set to optimize the performance visually and quantitatively. Meanwhile, we noted that a simpler model would be more likely to generalize to future data. Other parameters such as learning rates were also tuned to achieve the best performance across the coronary MRI data sets. We note that this set of parameters may not yield the best performance for imaging all organs. Therefore, similar to other techniques, a parameter tuning procedure may need to be performed prior to using sRAKI in other applications for an optimal performance.

Similar to SPIRiT, regularization terms can be included in the sRAKI objective function, in order to incorporate additional prior information, such as sparsity in transform domains [142–144]. However, these regularization parameters often need to be carefully tuned to avoid residual artifacts [144]. On the other hand, sRAKI without transform domain regularization, whose objective function requires no additional parameter tuning, showed desirable noise properties. The noise improvement in sRAKI is learnt from the coil geometry, and does not inherently include any assumptions about compressibility in transform domains. A combination of sRAKI with advanced regularizers bears potential for improved reconstruction quality in certain lower SNR scenarios (additional experiments for these scenarios are shown in Supporting Information **Figures A.14** and **A.15**, but this was beyond the scope of this work, which emphasized the multi-coil aspect of the data at the native acquisition SNR.

In addition to explicit regularization terms, some noise amplification reduction can be achieved by limiting the number of iterations for both SPIRiT and sRAKI, at the cost of incomplete unaliasing. The effect of early termination of reconstruction is most

pronounced as residual blurring artifacts, which is particularly troublesome in the application of high-resolution coronary MRI. Thus, our main criterion for number of iterations was to assure that blurring artifacts were entirely removed before reconstruction noise started to be amplified. We further observed that in contrast to l_1 -SPIRiT and sRAKI, the transition between residual blurring artifacts and amplified reconstruction noise was particularly fast for SPIRiT, which is due to lack of a regularization mechanism in SPIRiT, consistent with the literature on iterative SENSE [165]. In terms of quantitative evaluation, since NMSE is captured in the loss function, additional evaluation of image quality was incorporated using the vessel sharpness measurements, as is standard in coronary MRI [139, 144, 149].

Finally, while this work showed the feasibility of using sRAKI for accelerating coronary MRI, we have not comprehensively evaluated the failure modes and the residual artifacts, and how these would affect diagnostic decisions. Further study of these effects is warranted in patient populations to establish diagnostic utility, and will be explored in future research.

Chapter 5

Database-Free Accelerated MRI Using Recurrent Neural Networks with Densely Connected Blocks

5.1 Introduction

As explained in Chapter 4, coronary MRI provides a non-invasive and radiation-free tool in diagnosis of coronary artery disease (CAD) [27], the leading cause of death in the United States [26]. Coronary MRI is typically acquired with electrocardiogram (ECG) triggering during diastolic quiescence and in free-breathing, leading to long scan durations and necessitating accelerated imaging [131–133]. Parallel imaging [140, 141], compressed sensing [142–144] and their combinations [145–151] are some of the classical methods that have been used for accelerated coronary MRI. Deep learning-based techniques [32–43] have also been proposed recently that mostly depend on large databases of fully-sampled datasets for training neural networks. However, in some MRI applications such as whole heart coronary MRI [137–139], reference fully-sampled data cannot be acquired due to impractically long scan times. We introduced sRAKI in Chapter 4 which trains CNNs on scan-specific ACS data, and thus addresses this issue. This

scan-specificity feature of sRAKI also ensures that inter-scan or inter-individual variability of training data does not adversely affect the generalizability of learning for reconstruction [48]. The notion of self-consistency in sRAKI is similar to that proposed in SPIRiT [153]. However, in sRAKI, a nonlinear mapping using CNNs is learnt as the self-consistency rule instead of the linear convolutions in SPIRiT. Similar to SPIRiT, sRAKI also performs the calibration and reconstruction in separate phases, which can degrade efficiency of reconstruction. In this study, we sought to combine these two steps by using recurrent neural networks (RNN) [37, 38, 40, 166] with multiple densely connected neural network blocks [167]. The new technique, called sRAKI-RNN is evaluated in coronary MRI and compared to SPIRiT and sRAKI for various acceleration rates.

5.2 Methods

Let \mathbf{x} denote the full k-space data across all coils and \mathbf{y} be the corresponding acquired undersampled noisy data from a multi-coil MRI system with n_c coils. The forward model for this system is formulated as:

$$\mathbf{y} = \mathbf{D}\mathbf{x} + \mathbf{n}, \quad (5.1)$$

where \mathbf{D} is the undersampling operator and \mathbf{n} represents acquisition noise. The following objective function is optimized to estimate the full k-space data \mathbf{x} from measurements \mathbf{y} by using the self-consistency notion of SPIRiT [153]:

$$\arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{D}\mathbf{x}\|_2^2 + \beta \|\mathbf{x} - \mathbf{G}(\mathbf{x})\|_2^2, \quad (5.2)$$

where $\mathbf{G}(\cdot)$ represents the self-consistency interpolation function. The first term in objective function (5.2) ensures that the reconstructed data remains consistent with the acquired portion of data. The second term enforces the self-consistency rule of $\mathbf{G}(\cdot)$ on the reconstructed data and β is a weight term. SPIRiT assumes a linear self-consistency rule [153], whereas sRAKI generalizes the interpolating function $\mathbf{G}(\cdot)$ to a nonlinear mapping that can be learned by CNNs [161–163]. In both approaches, the self-consistency rule is determined by training on ACS data prior to reconstruction. Subsequently, the objective function in (5.2) can be optimized iteratively by alternating between enforcing data consistency and self-consistency [153].

An alternative approach is to unroll iterations in an RNN [37, 38, 40, 166] with multiple CNN blocks each implementing the unrolled version of this iterative approach. For the n^{th} unrolled iteration, the self-consistency and data-consistency operations are applied as follows:

$$\begin{cases} \mathbf{z}_n = \mathbf{G}(\mathbf{x}_n), \\ \mathbf{x}_{n+1} = (\mathbf{I} - \mathbf{D}^T \mathbf{D}) \mathbf{z}_n + \mathbf{D}^T \mathbf{y}, \end{cases} \quad (5.3)$$

where \mathbf{x}_n is the reconstructed k-space data at iteration n , \mathbf{z}_n is the output of self-consistency unit and \mathbf{I} is the identity operator. Thus, the second equation in (5.3) enforces data-consistency with acquired data \mathbf{y} . This approach, called sRAKI-RNN, eliminates the need for a separate calibration of CNNs in sRAKI to specify the self-consistency rule $\mathbf{G}(\cdot)$. Instead, sRAKI-RNN trains an end-to-end network to concurrently learn the self-consistency rule and implement the iterative scheme of reconstruction in sRAKI. Hence, the reconstruction phase of sRAKI-RNN is a single forward propagation step in training neural networks. In addition, since the CNN parameters are shared across iterations, the training phase can be efficiently performed using the ACS data, as in sRAKI, maintaining scan-specificity.

5.2.1 Implementation Details

In this study, an RNN architecture unrolled for 5 iterations was employed to perform an end-to-end k-space reconstruction (**Figure 6.1**). Each iteration consists of a single self-consistency (SC) and data-consistency (DC) unit, connected to each other. For the self-consistency unit, 3 blocks of densely connected CNNs [167] followed by a single convolutional layer were employed with 2 outer skip connections (see **Figure 6.1**) to further facilitate the flow of information across blocks. Each dense block (DB) consisted of 3 convolutional layers with a growth factor of 8 channels. All layers in the dense blocks were followed by rectifier linear units (ReLU) as activation functions. The kernel size of all convolutional layers was 3×3 .

To generate training data, ACS region was retrospectively undersampled in $k_y - k_z$ plane using a single Poisson-disc pattern with a rate identical to the rest of data. The undersampled and full ACS data pairs were later used to train the network. To reduce the training computational complexity, 3D data were first inverse-Fourier transformed

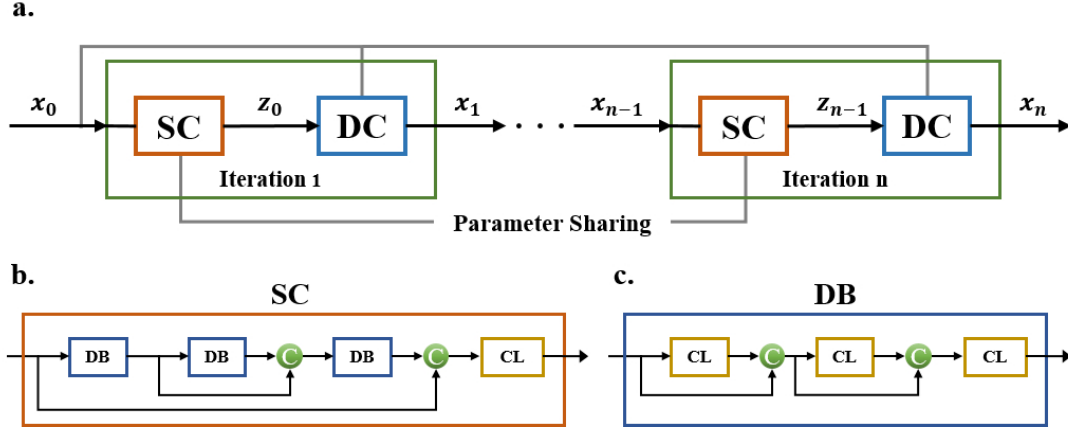


Figure 5.1: (a) The recurrent network architecture of sRAKI-RNN unrolled for n iterations to simultaneously apply self-consistency (SC) and data consistency (DC) and perform the iterative reconstruction. (b) A single SC unit with 3 dense blocks (DB), a single convolutional layer (CL) at the output, along with two skip connections to facilitate information flow through network. (c) A closer view of the DB.

along readout direction, k_x . The resultant 2D slices were then fed to the network in batch sizes of 32. The real and imaginary components of complex k-space data were concatenated before feeding them into the network, which led to $2n_c$ input and output channels. The network was trained by minimizing an MSE loss function using Adam optimizer with a learning rate of 0.001, iterated over 500 epochs. To standardize the hyper-parameters such as the learning rate, the k-space was scaled such that the maximum absolute value was 1. For comparison purposes, SPIRiT using a conjugate gradient reconstruction [153] and sRAKI with the parameters as previously described [161–163] were also implemented.

5.2.2 In Vivo Coronary MRI

Targeted right coronary artery (RCA) MRI was acquired on a 28-year old male subject using a 3T Siemens Magnetom Prisma (Siemens Healthineers, Erlangen, Germany) system with a 30-channel receiver body coil-array and an ECG-triggered GRE sequence with imaging parameters: TR/TE=3.4/1.5ms, flip angle=20°, bandwidth=601 Hz/pixel, field-of-view (FOV)=300 × 300 × 48mm³, resolution=1 × 1 × 3mm³, navigator

window=5mm. T_2 -preparation and a spectrally-selective fat saturation were utilized for improved contrast. The 3D k-space data was exported and retrospectively undersampled with a Poisson disc pattern at acceleration rates 2, 3, 4, and 5 with a fully-sampled 36×10 ACS region in $k_y - k_z$ plane. The images reconstructed using SPIRiT, sRAKI and sRAKI-RNN were quantitatively assessed by normalized mean square error (NMSE) and structural similarity index (SSIM) metrics with respect to the reference fully-sampled image. Final images were obtained using root sum-squares combination of all coil images. All algorithms were implemented in Python, and processed on a workstation with an Intel E5-2640V3 CPU (2.6GHz and 256GB memory), and an NVIDIA Tesla V100 GPU with 32GB memory.

5.3 Results

Figure 5.2 depicts a representative slice from the targeted RCA MRI dataset that was retrospectively undersampled at rates 2, 3, 4 and 5, and reconstructed using techniques SPIRiT, sRAKI and sRAKI-RNN. All techniques successfully remove aliasing artifacts and allow visualization of the RCA at all rates. Nonetheless, reconstruction noise is amplified by SPIRiT, especially at higher rates. Both sRAKI and sRAKI-RNN demonstrate desirable noise properties, with sRAKI-RNN being superior in suppressing reconstruction noise at all rates.

A quantitative analysis is provided in Table 5.1 for the NMSE and SSIM metrics. These values, which are consistent with the visual evaluation of **Figure 5.2**, further indicate that sRAKI-RNN improves accelerated coronary MRI reconstruction over SPIRiT and sRAKI.

5.4 Discussion

We proposed an accelerated MRI reconstruction technique, sRAKI-RNN that uses an RNN with densely connected blocks to enforce self-consistency among multi-coil MRI datasets and data-consistency with arbitrarily undersampled data. sRAKI-RNN further improves the desirable noise properties of our previous technique, sRAKI [161–163]. In contrast to most deep learning-based accelerated MRI reconstruction methods that need




Figure 5.2: A representative slice from a 3D right coronary artery MRI data of a healthy subject. The data were retrospectively undersampled at rates 2, 3, 4 and 5 in the $k_y - k_z$ plane and then reconstructed using techniques SPIRiT, sRAKI and sRAKI-RNN (top, middle and bottom rows). The fully-sampled images are provided in the first column as a reference for comparison. All techniques successfully remove aliasing artifacts, and the least reconstruction noise is achieved by sRAKI-RNN. (RCA: right coronary artery; AO: aortic root)

large amounts of data to train neural networks [33–43], sRAKI-RNN inherits the scan-specificity feature of the RAKI and sRAKI methods. Consequently, sRAKI-RNN relies on ACS data from the same scan to train the neural network, which is particularly beneficial for applications where long scans impede acquiring fully-sampled data, such as whole-heart imaging [137–139].

Table 5.1: Quantitative NMSE (left) and SSIM (right) assessment of SPIRiT, sRAKI and sRAKI-RNN in reconstructing a right coronary artery MRI dataset at acceleration rates 2, 3, 4 and 5. sRAKI-RNN achieves the least NMSE and the highest SSIM values at all rates.

| NMSE | R2 | R3 | R4 | R5 | SSIM | R2 | R3 | R4 | R5 |
|-----------|--------|--------|--------|--------|-----------|-------|-------|-------|-------|
| SPIRiT | 0.0119 | 0.0205 | 0.0334 | 0.0449 | SPIRiT | 0.907 | 0.855 | 0.792 | 0.748 |
| sRAKI | 0.0090 | 0.0162 | 0.0215 | 0.0253 | sRAKI | 0.945 | 0.906 | 0.875 | 0.853 |
| sRAKI-RNN | 0.0079 | 0.0145 | 0.0197 | 0.0238 | sRAKI-RNN | 0.954 | 0.919 | 0.893 | 0.874 |

By using an unrolled RNN scheme with parameter sharing across iterations, sRAKI-RNN combines the calibration and reconstruction phases of sRAKI into a single training phase. This compaction reduced the total running time of sRAKI-RNN to approximately 60 seconds, which is twice faster than sRAKI. In addition, fewer hyper-parameters (e.g. learning rate of gradient descent) are involved in the whole reconstruction process, which further facilitates training. Another difference with sRAKI is that sRAKI-RNN utilizes a densely connected neural network design for the self-consistency units. This configuration with outer skip connections considerably improved convergence of training by facilitating the flow of information across blocks. We note that such a design for sRAKI leads the network to learn a trivial identity mapping for the self-consistency rule due to the skip connections, since input and target data are identical in calibration phase of sRAKI.

We have compared the performance of sRAKI-RNN with SPIRiT and sRAKI in a targeted RCA dataset. Visual and quantitative evaluation of the results indicate that sRAKI-RNN considerably reduces amplification of reconstruction noise in SPIRiT, particularly at higher acceleration rates. In addition, sRAKI-RNN improves noise properties of sRAKI at all acceleration rates.

Chapter 6

Dense Recurrent Neural Networks for Inverse Problems: History-Cognizant Unrolling of Optimization Algorithms

6.1 Introduction

Inverse problems have been extensively utilized in many imaging modalities, such as magnetic resonance imaging (MRI) [31, 165, 168–172], computed tomography (CT) [173–178], and optical imaging [179–182]. These inverse problems are directly driven from the physics of data acquisition, known as the forward operator, which incorporates domain-specific knowledge. Such inverse problems are typically ill-conditioned and thus regularizers are incorporated into the respective objective functions. Subsequently, iterative algorithms are employed to solve the regularized inverse problem by alternating between enforcing physics-driven data consistency and regularization using pre-selected priors.

Recently, deep learning (DL) has emerged as an alternative state-of-the-art technique for solving such inverse problems in imaging applications [32–40, 45, 183–206]. Several of the earlier works concentrated on training convolutional neural networks (CNN) to

solve the problem without incorporating explicit knowledge of the acquisition physics, typically requiring re-training when acquisition parameters are modified [33,35,36,39,45, 183–190]. Another line of work focuses on a physics-driven formulation, where the known forward model is incorporated into training, utilizing domain knowledge to solve the corresponding inverse problem [34,37,38,40,191–202,207]. In physics-driven approaches, a conventional iterative optimization problem, which lands itself to a regularized least-squares formulation for modalities such as MRI and CT, is unrolled for a fixed number of iterations. This unrolled network alternates between data consistency and regularization steps similar to the conventional algorithm, where the regularizers are implemented via neural networks. The unrolled network is trained end-to-end to learn the network parameters, which predominantly characterize the regularization units.

There are numerous approaches for solving regularized least squares problems, which have been explored for algorithm unrolling in physics-driven DL reconstruction methods [34, 37, 38, 40, 191–202, 207]. These methods include gradient descent (GD) [34], proximal gradient descent (PGD) [40,191–194], variable-splitting (VS) [37,38,195–199] and primal-dual (PD) [200,201] methods. GD, PGD and PD unrolling typically use a gradient descent step for data consistency [34,37,38,40,191–201], while for VS-based methods, either a gradient descent [196] or conjugate gradient (CG) data consistency can be used, with the latter being utilized when the least squares problem involving the forward encoding operator becomes computationally expensive [37,38,195,197–199]. Data consistency updates with gradient descent step is generally computationally inexpensive. However, these optimization methods have slower convergence, which necessitates higher number of unrolled iterations for improved performance, which in turn can lead to numerical issues such as gradient vanishing and memory constraints.

In this study, we seek to develop a new methodology for unrolling inverse problem optimization algorithms into dense recurrent neural networks (RNN) for improved reconstruction performance [208]. Motivated by accelerated PGD approaches, which calculate the gradient descent at a combined history of previous iterations rather than the immediate previous iteration only [209–215], we propose a history-cognizant unrolling of the optimization algorithms with skip connections across unrolled iterations, leading to a dense RNN architecture. Skip connections have been previously shown to facilitate information flow through neural networks in architectures such as ResNet [216]

and DenseNet [217]. Different than these previous neural network design problems, we propose to utilize such connections in an unrolled RNN architecture for solving inverse problems, which readily comes with its domain-specific design considerations. We compare the conventional unrolling and the proposed history-cognizant unrolling of the PGD algorithm in accelerated multi-coil MRI reconstruction using the fastMRI knee dataset [218], showing both visual and quantitative improvement in reconstruction quality with the proposed methodology. We also extend the same unrolling strategy to other optimization algorithms and investigate the performance improvement.

6.2 Theory

6.2.1 Forward Model and Inverse Problem

Let \mathbf{y} denote a vector of noisy measurements from an imaging system and \mathbf{x} be the corresponding underlying image. The forward model for this system is generally given as $\mathbf{y} = \mathcal{A}(\mathbf{x}) + \mathbf{n}$, where \mathcal{A} is called a forward operator and \mathbf{n} is measurement noise. In a number of medical imaging scenarios, such as MRI and CT reconstruction, the forward model can be simplified to a linear equation:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \quad (6.1)$$

where \mathbf{A} is a linear forward operator and \mathbf{n} is measurement noise, which will be the focus of the rest of the study. The corresponding inverse problem is typically ill-conditioned. Thus, regularization is frequently used, leading to a regularized least-squares method:

$$\arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (6.2)$$

where the first term enforces consistency with measurement data and $\mathcal{R}(\cdot)$ denotes a regularizer. In classical reconstruction approaches, regularizers have included Tikhonov [165, 168], edge-preserving functions [219] or sparsity-promoting terms in pre-defined domains [31, 169, 171, 172]. A variety of algorithms can iteratively solve the optimization problem in (6.2) such as PGD [220], VS [221–223] and PD [224, 225] methods.

6.2.2 Physics-Based DL Reconstruction by Unrolling PGD

Several physics-based DL methods have unrolled a PGD-based method [40,191–194] for a fixed number of iterations. Such PGD-based unrolling alternates between a proximal operation that is implicitly defined by a learned neural network, and a gradient descent step to enforce data consistency, as follows:

$$\begin{aligned} \mathbf{z}^{(i)} &= \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) = \\ &= \arg \min_{\mathbf{z}} \|\mathbf{x}^{(i-1)} - \mathbf{z}\|_2^2 + \mathcal{R}(\mathbf{z}) \end{aligned} \quad (6.3a)$$

$$\mathbf{x}^{(i)} = \mathbf{z}^{(i)} + \mu_i \mathbf{A}^H (\mathbf{y} - \mathbf{A} \mathbf{z}^{(i)}) \quad (6.3b)$$

where μ_i is the gradient descent step size at unrolled iteration i , and the proximal operation in (6.3a) is performed by a neural network. This leads to the schematic in **Figure 6.1a**, where the unrolled RNN consists of multiple blocks of proximal (P) and data consistency (DC) units. The former is a neural network and the latter implements a gradient descent step based on the known forward encoding operator. The unrolled RNN is then trained end-to-end to learn the mapping from the acquired sensor-domain data to reference images for medical imaging reconstruction.

6.2.3 Proposed History-Cognizant PGD Unrolling with Dense Recurrent Neural Networks

One of the major appeals of the PGD unrolling scheme is its computational simplicity. However, PGD algorithm itself has a slow convergence rate [219], potentially requiring a large number of iterations (i.e., a large number of blocks in **Figure 6.1a**) when unrolled into a neural network for sufficient reconstruction quality. In turn, this increases memory requirements and impedes training complexity. On the other hand, several methods have been proposed to accelerate convergence rate of PGD [209–215]. The underlying theme for these methods is to calculate the gradient step not only at the most recent proximal operation output, but at a linear combination of all past such outputs.

We propose to adapt this general accelerated PGD approach for network unrolling.

In this setting, the accelerated PGD methods involve the following steps:

$$\mathbf{z}^{(i)} = \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) \quad (6.4a)$$

$$\mathbf{v}^{(i)} = \mathbf{F}(\mathbf{z}^{(i)}, \dots, \mathbf{z}^{(1)}) \quad (6.4b)$$

$$\mathbf{x}^{(i)} = \mathbf{v}^{(i)} + \mu_i \mathbf{A}^H (\mathbf{y} - \mathbf{A} \mathbf{v}^{(i)}), \quad (6.4c)$$

where $\mathbf{F}(\cdot)$ forms a linear combination of the previous proximal operator outputs. As a special case, Nesterov’s method uses the previous two outputs [210], while more general versions have also been studied, in which all the past iterations are used [211, 215].

For our unrolled network, we propose to learn $\mathbf{F}(\cdot)$ during training instead of specifying a pre-determined form for the linear combination of previous iterates $\mathbf{F}(\cdot)$ as in Nesterov’s method [210]. This history-cognizant unrolling of the accelerated PGD algorithms leads to a dense recurrent neural network architecture (Dense-RNN) with skip connections across all unrolled iterations, as depicted in **Figure 6.1b**. In addition to the P and DC units to perform proximal operation (or regularization) and data consistency updates, each unrolled iteration also consists of a 1×1 convolutional layer implementing $\mathbf{F}(\cdot)$ on the concatenation (denoted by C) of all past proximal operator outputs prior to the DC unit. We refer to this method as history-cognizant PGD (HC-PGD) in the rest of this study.

6.3 Multi-Coil MRI Reconstruction Experiments and Results

The proposed history-cognizant algorithm unrolling approach was compared to conventional algorithm unrolling in multi-coil MRI reconstruction. In a multi-coil MRI system with n_c coils, \mathbf{y} is the sub-sampled noisy data from all coils and \mathbf{x} is the underlying image to be recovered. The forward linear operator $\mathbf{A} : \mathbb{C}^{M \times N} \rightarrow \mathbb{C}^P$ in this system includes sensitivities of the receiver coil array and a partial Fourier matrix for sub-sampling in k-space, i.e. spatial Fourier domain [165].

6.3.1 Knee MRI Datasets

Knee MRI datasets were obtained from the New York University (NYU) fastMRI initiative database [218] to investigate accelerated MRI reconstruction using conventional

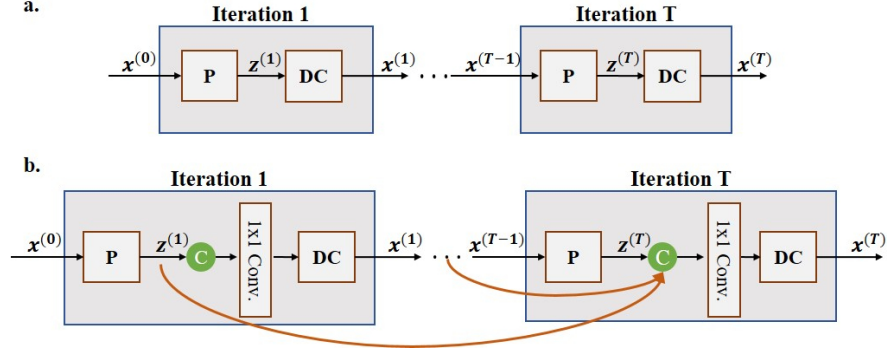


Figure 6.1: (a) The conventional unrolling for iterative algorithms for inverse problems and (b) the proposed history-cognizant unrolling leading to a dense recurrent neural network architecture. In both cases, the algorithms are unrolled for T iterations, where each unrolled iteration consists of proximal (or regularization) unit (P), and a data consistency unit (DC). The P unit is implemented implicitly via a neural network, while the DC unit corresponds to a gradient descent step for the proximal gradient descent algorithms and typically to conjugate gradient for variable splitting methods. For the proposed history-cognizant unrolling, an additional 1×1 convolutional layer combines all previous proximal operator outputs prior to the DC unit, where C denotes concatenation.

unrolled RNN and Dense-RNN architectures. Data were acquired without any acceleration on a clinical 3T system (Magnetom Skyra, Siemens, Germany) with a 15-channel knee coil. Data acquisition was performed using 2D turbo spin-echo sequences in coronal orientation with proton-density (Coronal PD) and proton-density weighted with fat suppression (Coronal PD-FS) weightings. Relevant imaging parameters were: resolution = $0.49 \times 0.44\text{mm}^2$, slice thickness = 3mm, matrix size = 320×368 for both datasets.

All data were uniformly sub-sampled at an acceleration rate of 4 by keeping the central 24 lines using the mask provided in the fastMRI database [218] retrospectively. Uniform undersampling for 2D acquisitions is typically considered a more challenging setup for DL reconstruction than random undersampling [34]. A 24×24 central window was used to generate coil sensitivity maps using ESPIRiT [226]. The training sets consisted of 300 slices from 15 subjects for each sequence. Testing was performed on 10 different subjects for both sequences.

6.3.2 Implementation Details

The optimization algorithms were unrolled for $T = 10$ iterations. The proposed HC-PGD algorithm was implemented using the Dense-RNN architecture in **Figure 6.1b**. Proximal operator unit outputs are first concatenated and then combined using convolutional layers with a 1×1 kernel size and 2 output channels (real and imaginary components) before being inputted to the subsequent data consistency unit. On the other hand, the PGD algorithm was implemented using a conventional architecture that only includes P and DC units, without skip connections, concatenation of previous iterates or 1×1 convolutions.

The same CNN was used to implement the proximal operation for both PGD and HC-PGD. A residual network (ResNet) that has been established in a different regression problem was utilized [227]. This CNN consisted of 15 sub-residual blocks each having two convolutional layers of kernel size = 3×3 and output channels = 64. The first and second convolutional layers in each sub-residual block are followed by a ReLU activation function and a scaling factor of 0.1, respectively. In addition to sub-residual blocks, two input and output convolutional layers without any nonlinear activation function are to match number of desired channels. For both algorithm unrolling schemes, the proximal units shared parameters across unrolled iterations [37], leading to a total of 592,138 and 592,358 parameters for PGD and HC-PGD, respectively.

All networks were trained end-to-end, as summarized in **Figure 6.2**, by using Adam optimizer to minimize

$$\min_{\boldsymbol{\theta}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathbf{x}_{ref}^i, f(\mathbf{y}^i, \mathbf{A}^i; \boldsymbol{\theta})), \quad (6.5)$$

where \mathbf{x}_{ref}^i , \mathbf{y}^i and \mathbf{A}^i are the reference SENSE-1 image, undersampled data and forward encoding operator for slice i , respectively and N is the number of training slices. $f(\mathbf{y}^i, \mathbf{A}^i; \boldsymbol{\theta})$ denotes the network output for slice i with the corresponding measurements and forward encoding operator, while $\boldsymbol{\theta}$ contains the trainable parameters of the networks. A normalized $\ell_1 - \ell_2$ loss [205] was used for $\mathcal{L}(\cdot, \cdot)$:

$$\mathcal{L}(\mathbf{u}, \mathbf{v}) = \frac{\|\mathbf{u} - \mathbf{v}\|_2}{\|\mathbf{u}\|_2} + \frac{\|\mathbf{u} - \mathbf{v}\|_1}{\|\mathbf{u}\|_1}, \quad (6.6)$$

with \mathbf{u} and \mathbf{v} representing reference fully-sampled and network output images, respectively. Training was performed over 100 epochs. All algorithms were implemented using

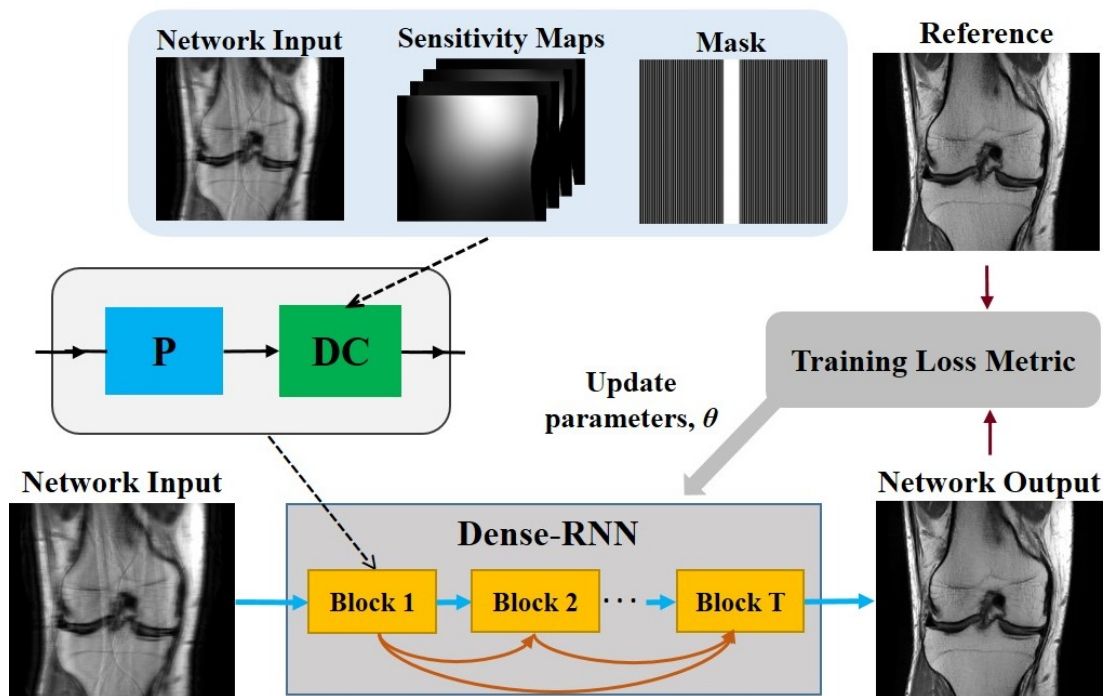


Figure 6.2: Summary of the training procedure in the proposed history-cognizant unrolling of inverse problem optimization algorithms with a Dense-RNN architecture for accelerated DL MRI reconstruction. Parameters of the unrolled network, which are dominantly characterized by proximal (P) operator units, are updated in each epoch based on a given loss metric using conventional end-to-end training. Data consistency (DC) units utilize the acquired data, coil sensitivity maps and undersampling mask in updates. The uniform undersampling mask utilized for training the network is also depicted.

TensorFlow, and processed on a workstation with an Intel E5-2640V3 CPU (2.6GHz and 256GB memory), and two NVIDIA Tesla V100 GPUs with 32GB memory.

6.3.3 Performance of HC-PGD versus PGD Unrolling

Experiments were conducted on both Coronal PD and PD-FS datasets to compare the reconstruction quality of the conventional PGD unrolling and the proposed HC-PGD unrolling with the Dense-RNN architecture. A total of 380 slices from 10 subjects (excluded from training) were utilized for testing. Reconstructed testing data were

evaluated quantitatively using normalized mean square error (NMSE) and structural similarity index (SSIM) metrics with respect to the fully-sampled reference image.

Figure 6.3 depicts a representative slice from the Coronal PD dataset reconstructed using DL-MRI with PGD and HC-PGD unrolling methods. The first row shows the reference fully-sampled image, as well as the zero-filled acquired image, which serves as the input $\mathbf{x}^{(0)}$ to the network. Difference images with respect to the reference image, scaled by a factor of 10, are also provided to facilitate comparison by showing errors. DL-MRI based on the proposed HC-PGD unrolling was able to suppress the strong aliasing artifacts successfully, while residual aliasing is visible on the DL-MRI reconstructions based on conventional PGD unrolling. For both methods, there is a slight degree of blurring in the final images, consistent with previous reports on DL-MRI methods with similar training loss functions [34]. Quantitative evaluations, shown in the lower left corner of the images, also confirm the improvement using the proposed history-cognizant unrolling with Dense-RNN architecture compared to the conventional algorithm unrolling.

The reconstructed images for a representative slice from the coronal PD-FS dataset in addition to the difference images scaled by a factor of 10, as well as the reference and zero-filled images are shown in **Figure 6.4**. There are slight residual aliasing artifacts in the DL-MRI reconstruction with PGD unrolling. The proposed HC-PGD unrolling successfully removes these artifacts further. The quantitative metrics also show improvement for the reconstruction results in moving from PGD to HC-PGD unrolling.

Figure 6.5 summarizes the quantitative analysis of the DL-MRI reconstructions using NMSE and SSIM metrics for the PGD and HC-PGD unrollings for both the coronal PD and PD-FS datasets. The box plots depict the interquartile range with the center at the median of the metric for the corresponding reconstruction methods. Both metrics are visibly improved over the whole test sets when using the proposed history-cognizant algorithm unrolling over the conventional unrolling approach. Further statistical tests were performed using the Wilcoxon signed rank test with a significance level of $P < 0.05$, indicating that all differences are statistically significant based on both metrics and in both datasets.

6.3.4 Insensitivity of Dense-RNN Improvement to the Regularization CNN Choice

In order to verify that the improvements between PGD and HC-PGD unrolling are not limited to the choice of CNN used for the proximal operation, the two approaches in **Figure 6.1** were also implemented using a U-Net architecture for the regularization CNN. This was based on the U-net used in [218], which includes contracting and expanding paths up to 256 channels without batch normalization. **Figure 6.6** depicts the results of this experiment on the coronal PD datasets, where the proposed history-cognizant unrolling outperforms the conventional unrolling methodology both visually and quantitatively. This highlights that the improvement from the use of Dense-RNN is not limited to a specific choice of network architecture for the proximal operation.

6.3.5 Performance Evaluation for Random Undersampling

In order to evaluate performance of the proposed history-cognizant unrolling approach in random undersampling, we utilized an observation from a previous study [34, 228], showing that physics-based networks trained on uniformly undersampled data can be transferred to random undersampling patterns at the same rate directly, even though the opposite is not necessarily true. Thus, we employed the networks trained on uniformly undersampled data, as detailed in Section 6.3.2, to reconstruct test slices which were randomly undersampled at a similar 4-fold rate. The corresponding results, shown in **Figure 6.7** indicate that the improvement trend remains similar to the uniform sampling scenario with the proposed HC-PGD method outperforming the conventional PGD in suppressing residual artifacts as observed in both Coronal PD and PD-FS datasets, which also align with the quantitative evaluation metrics in the images.

6.3.6 Extensions to Other Optimization Algorithms

While PGD unrolling is commonly used for DL-MRI reconstruction [40, 191–194], alternative algorithms have also been utilized [37, 38, 195–199]. One such method for solving the objective function in Equation (6.2) is VS with quadratic penalty, which decouples data consistency and regularization terms into two blocks of variables [219] by introducing auxiliary variables to formulate a constrained objective function. This is then

relaxed to an unconstrained problem by enforcing similarity between these variables via a quadratic penalty:

$$\arg \min_{\mathbf{x}, \mathbf{z}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{z}) + \beta \|\mathbf{x} - \mathbf{z}\|_2^2, \quad (6.7)$$

where \mathbf{z} is an auxiliary variable and β is the quadratic penalty parameter. When this algorithm is unrolled as a neural network [37, 38, 196–199], reconstruction alternates between a proximal operation that is again learned by a neural network and data consistency, similar to PGD unrolling, as follows:

$$\mathbf{z}^{(i)} = \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) \quad (6.8a)$$

$$\mathbf{x}^{(i)} = (\mathbf{A}^H \mathbf{A} + \mu \mathbf{I})^{-1} (\mathbf{A}^H \mathbf{y} + \beta \mathbf{z}^{(i)}) \quad (6.8b)$$

where regularization update in (6.8a) is performed by a neural network and β was absorbed into the regularizer for notational convenience. The data consistency update in (6.8b) is computationally challenging in many practical scenarios, such as multi-coil MRI reconstruction. Thus, this data consistency update can be solved iteratively via conjugate gradient (CG) to avoid matrix inversion [37].

As an extension to the HC-PGD approach, VS with quadratic relaxation can also be unrolled in a history-cognizant manner using the Dense-RNN in **Figure 6.1**, building on the optimization approaches in [213, 214]. This will be referred to as the history-cognizant VS method (HC-VS). We note that the main difference to HC-PGD is that the DC units in HC-VS now implement an unrolled CG algorithm for a different least squares problem instead of a gradient step. To test the applicability of the history-cognizant unrolling scheme to different baseline algorithms, both VS and HC-VS were also implemented on the coronal PD and PD-FS datasets, using the setup described in Section 6.3.2. In each data consistency unit, CG was implemented with 10 internal iterations.

Figure 6.8 depicts representative slices from the Coronal PD and Coronal PD-FS datasets reconstructed using DL-MRI with the VS and HC-VS unrolling methods along with the error images scaled by a factor of 10 with respect to the reference fully-sampled image, as well as the reference and zero-filled images. The HC-VS approach successfully removes residual artifacts apparent in the VS results as shown by the yellow arrows. Improvements are also observed in the quantitative metrics.

6.4 Discussion

In this study, we proposed a history-cognizant method for unrolling optimization problems to solve regularized inverse problems in medical imaging using physics-based DL reconstruction. Our unrolling method builds on optimization approaches that utilize a history of the previous iterates via their linear combinations in updating the current estimate [209–214]. While these optimization procedures have been traditionally proposed to improve convergence rates or provide stronger theoretical guarantees, in the context of algorithm unrolling for physics-based DL reconstruction, they are used to enhance reconstruction quality without incurring additional computational complexity. Our method was pre-dominantly motivated by accelerated PGD approaches [209–212] and was shown to improve upon conventional PGD unrolling for multi-coil MRI reconstruction for multiple scenarios. We also explored how the unrolling can be applied to VS with quadratic relaxation for solving a regularized least squares problem, and showed that it can improve over the conventional unrolling approach.

The proposed history-cognizant unrolling methodology leads to several skip connections across iterations, leading to a densely-connected unrolled neural network. Dense neural networks have been shown to improve training by facilitating information flow through the network in forward propagation and by relieving issues concerned with gradient vanishing in backward propagation [217]. One of the advantages of these dense connections in the proposed unrolling is that the only computational overhead compared to conventional unrolling is the calculation of the 1×1 convolutions from the previous iterations. Thus, once trained, the computational time is increased $<1.5\%$ compared to the conventional counterpart, as depicted in Table I.

One of the most commonly used accelerated PGD approaches is based on Nesterov acceleration [210]. In this case, the gradient descent is evaluated at a linear combination of the previous two iterates weighted by a step size μ and $1 - \mu$ respectively. Theoretically, this approach achieves the optimal convergence rate for first-order methods [210]. As such, Nesterov-type unrolling for DL reconstruction was readily explored in [194] for coded-illumination microscopy, where the unrolled accelerated PGD network used the history of the two previous outputs of the proximal units. However, the main focus

of this work was developing an illumination pattern design algorithm for a fixed reconstruction algorithm. On the other hand, when focusing on the reconstruction aspect for inverse problems, we found that the Nesterov-type unrolling was outperformed by the full history-cognizant Dense-RNN architecture in our initial experiments (not shown). Furthermore, the full history-cognizant unrolling subsumes Nesterov-type as a special case. Thus, we focused on this full history-cognizant case throughout our study to maximize the performance benefits.

In our experiments, the history-cognizant unrolling visibly improved the performance of PGD-based methods compared to its conventional counterpart for all datasets. In the VS setting, especially for the coronal PD datasets, the improvement is less pronounced. Coronal PD datasets have higher SNR compared to PD-FS datasets [34]. In this higher SNR case, the conventional unrolling of the VS algorithm readily provides a high-quality reconstruction. Thus, additional improvement with HC-VS is less apparent than in the PGD setting.

Our experiments in multi-coil MRI reconstruction pre-dominantly utilized uniform undersampling patterns. While random undersampling patterns are commonly used in literature for training and testing DL-based reconstruction, uniform undersampling remains the most clinically used approach, and was considered to be a harder problem compared to random undersampling in [34, 228], as detailed in Section 6.3.5.

To the best of our knowledge, the Dense-RNN architecture that arises from the proposed history-cognizant unrolling have not been used for regression problems in general, and inverse problems in particular. However, Dense-RNNs have been independently proposed for different problems, such as scene parsing or labeling [229] where semantic information of an image segment is to be extracted. In such applications, RNNs are used for each image unit to receive the dependencies from other units through recurrent information forwarding between adjacent units, which decays for further units. In this setting, Dense-RNNs improve labeling by capturing long-range semantic dependencies among image units, highlighting a similar memory dependence as in this study.

MRI data were obtained from the NYU fastMRI initiative database [218] which was acquired with the relevant institutional review board approvals as detailed in [218]. NYU fastMRI investigators provided data but did not participate in analysis or writing. fastmri.med.nyu.edu provides listing of NYU fastMRI investigators, subject to

| | PGD | HC-PGD | VS | HC-VS |
|---------------|------------------|------------------|------------------|------------------|
| Coronal PD | 197 ± 4.2 ms | 197 ± 7.2 ms | 340 ± 6.8 ms | 325 ± 6.1 ms |
| Coronal PD-FS | 197 ± 5.6 ms | 198 ± 6.6 ms | 339 ± 5.6 ms | 325 ± 6.2 ms |

Table 6.1: Average and standard deviation (in ms) of reconstruction time over all test slices for PGD, HC-PGD, VS and HC-VS in Coronal PD and Coronal PD-FS datasets. History cognizant unrolling has very slight impact on reconstruction time compared to its conventional counterpart ($< 1.5\%$). Furthermore, VS methods, which utilize CG-based data consistency units, are considerably slower than PGD methods.

updates.

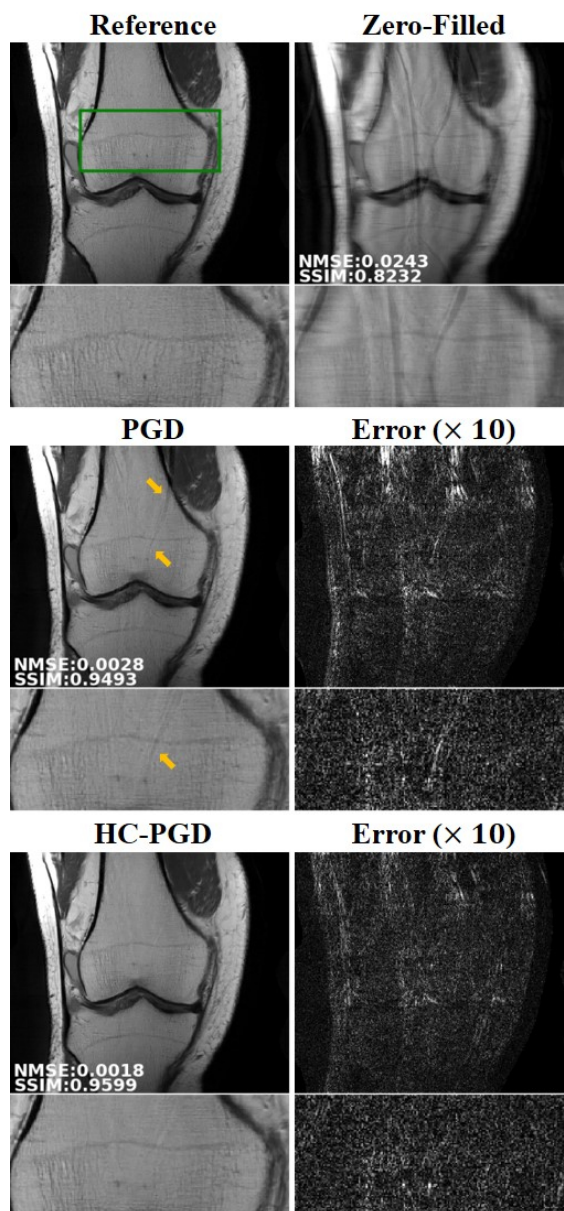


Figure 6.3: A representative slice from the Coronal PD dataset with 4-fold acceleration rate using a uniform undersampling pattern, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images with respect to the reference fully-sampled image scaled by a factor of 10. Zerofilled image corresponds to the input to the network, $\mathbf{x}^{(0)}$. The green rectangle in the reference image mark the zoom-up area shown in the second and last rows for each method. Residual aliasing artifacts, which are visible in the PGD unrolling results, have been suppressed using the HC-PGD scheme. Improvements in the quantitative metrics, shown on the lower-left corner of the images, also align with these visual observations.




Figure 6.4: A representative slice from the Coronal PD-FS dataset with a 4-fold acceleration rate using a uniform undersampling pattern, reconstructed with conventional PGD and the proposed HC-PGD methods, along with the error images scaled by a factor of 10. The green rectangle in the reference image mark the zoom-up area. Proposed HC-PGD unrolling successfully removes some residual artifacts that are visible in the PGD. The quantitative metrics align with this observation.

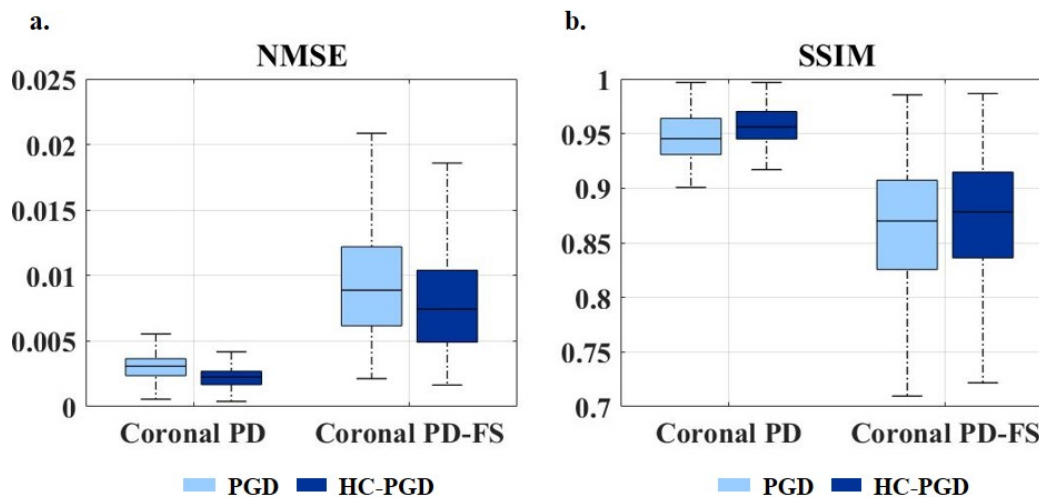


Figure 6.5: Summary of the quantitative metrics for reconstruction quality, (a) NMSE and (b) SSIM in Coronal PD and Coronal PD-FS datasets using the conventional PGD and the proposed HC-PGD methods. The boxes mark the interquartile range and the median of the metrics. The proposed history-cognizant unrolling of the PGD has visibly enhanced reconstruction, based on both metrics and in both datasets. Statistical analysis based on Wilcoxon signed rank test further confirm that all improvements are statistically significant.

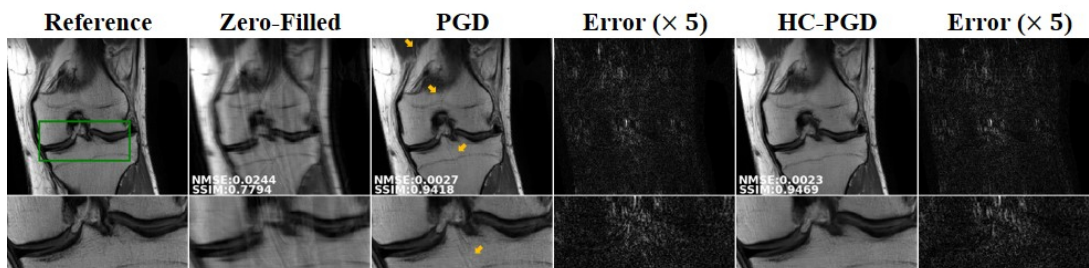


Figure 6.6: Results of changing the proximal operator CNN, with a slice from the Coronal PD dataset with 4-fold uniform undersampling, reconstructed with DL-MRI reconstruction using the PGD and HC-PGD methods unrolled in a network with a U-Net design for the proximal operator units, along with Error images ($\times 5$), with the green rectangle marking the zoom-up area. Similar to the ResNet experiments, the proposed HC-PGD outperforms PGD in suppressing visible residual artifacts.

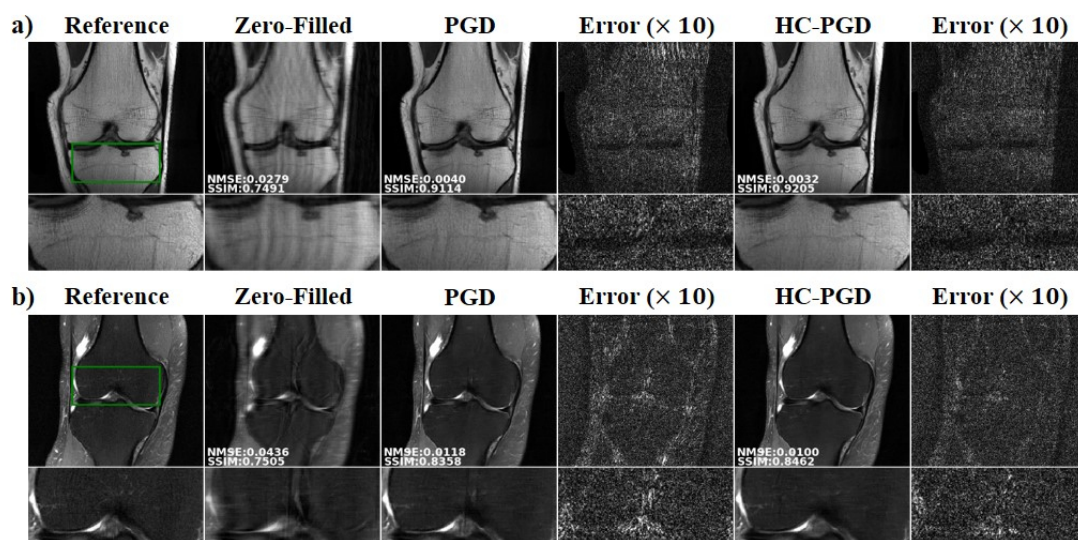


Figure 6.7: Results on random undersampling experiments from (a) the Coronal PD and (b) Coronal PD-FS datasets, at a 4-fold acceleration rate, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoom-up area. Regardless of the undersampling pattern, the history-cognizant unrolling approach, HC-PGD suppresses the residual aliasing visible in the conventional PGD unrolling, as marked by the arrows. The quantitative metrics align with these observations.

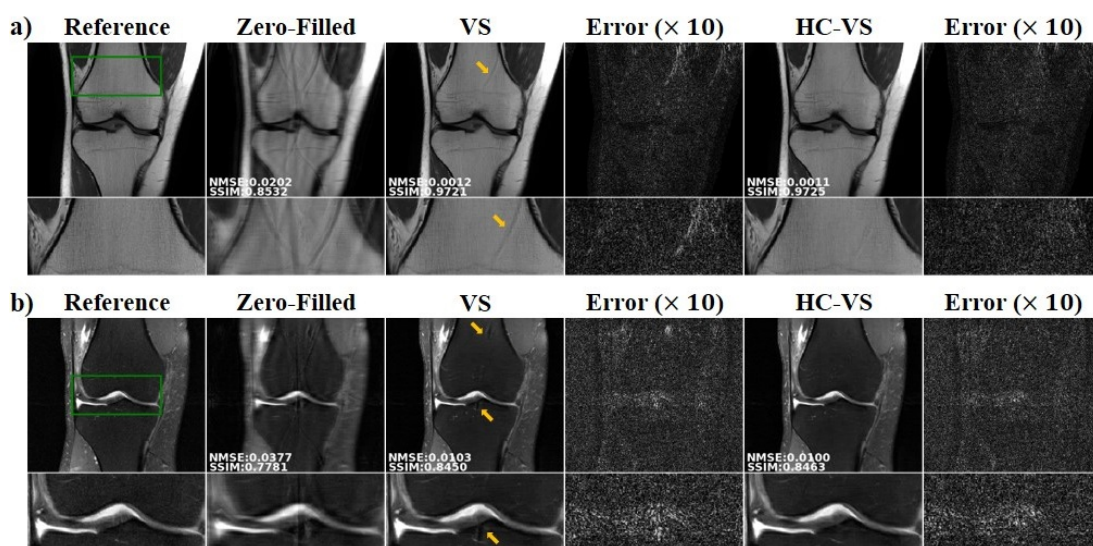


Figure 6.8: Example slices from (a) the Coronal PD and (b) Coronal PD-FS datasets with 4-fold uniform undersampling, reconstructed with DL-MRI reconstruction based on the conventional VS and the proposed HC-VS unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoom-up area. The reconstructions from the history-cognizant unrolling approach, HC-VS suppresses the residual aliasing visible in the conventional VS unrolling, as marked by the arrow. Improvements are also observed in the quantitative metrics.

Chapter 7

Conclusion

7.1 Thesis Summary

In this thesis, we mainly attempted to address the challenge of high-resolution imaging in two medical imaging modalities, ESI using EEG signals and MRI by applying statistical signal processing and deep learning methods. Although higher resolution is desired to allow more details to be captured during imaging, it is usually accompanied by side effects. High resolution ESI is biased in favor of superficial sources, while deeper sources are likely to be blocked by the brain activities that are closer to recording electrodes, usually placed on the scalp. High resolution MRI, on the other hand, trades off with SNR, provided that long scan duration would not dismiss acquisition feasibility. In addition, accelerated MRI reconstruction methods that require less data to be acquired for imaging, can amplify noise or be unable to fully remove residual artifacts from reconstructed images. This thesis has aimed to deliver solutions for these challenges to a possible extent.

In Chapter 2, we analyzed source imaging using intracranial EEG recordings and compared its accuracy to the results of EEG source imaging. Accuracy was measured both by determining the location and inter-nodal connectivity of underlying brain networks. Systematic computer simulation studies were conducted to evaluate iEEG-based source imaging vs. EEG-based source imaging, and source imaging using both EEG and iEEG. Results of this study showed that source imaging based on iEEG recordings may improve the accuracy in localization, depending on the number and location of active

nodes relative to iEEG electrodes and to other nodes within the network. However, results indicated that combining EEG and iEEG modalities (simultaneous scalp and intracranial recordings) can improve the imaging accuracy significantly. Thus, while iEEG source imaging is useful in estimating the exact location of sources near the iEEG electrodes, combining EEG and iEEG recordings can achieve a more accurate imaging due to the high spatial coverage of the scalp electrodes and the added near field information provided by the iEEG electrodes. These results suggest the feasibility of localizing brain electrical sources from iEEG recordings and improving EEG source localization using simultaneous EEG and iEEG recordings to cover the whole brain. The hybrid EEG and iEEG source imaging may assist clinicians when unequivocal decisions about determining the epileptogenic zone cannot be reached using a single modality.

In Chapter 3, we proposed and investigated robust beamformers for the purpose of ESI. The main advantage of robust beamformers over conventional adaptive beamformers, was that the RMVB and RMVB-ND-DN were more robust to mismatches between the forward and inverse modeling, which are inevitable in practice. The robust beamformers outperformed the conventional beamformers in terms of localization error, recovering source dynamics and estimation of the underlying source extents, when uncertainty in the lead field matrix was properly determined and modeled. This justifies the robust implementation of beamformers for applications, where source imaging is performed to study normal and pathological brain networks such as in epilepsy.

Chapter 4 was dedicated to sRAKI, an accelerated MRI reconstruction technique based on deep learning. While deep learning has gained significant attention for accelerated MRI reconstruction, most of current studies require databases of fully-sampled data for training the neural networks. We showed that in contrast to these methods, sRAKI performed a scan-specific training using data from the same scan only. Hence it is database-free and is particularly useful in applications where fully-sampled data are difficult or infeasible to obtain. Additionally, sRAKI was not limited to a specific under-sampling pattern which further facilitates its application in different scenarios. Results of this chapter indicated that sRAKI can effectively improve accelerated coronary MRI in terms of reducing reconstruction noise amplification and blurring artifacts over conventional methods.

We introduced sRAKI-RNN in Chapter 5 that combined the calibration and reconstruction phases of sRAKI to jointly learn the self-consistency rule and perform iterative reconstruction using recurrent neural networks. Similar to sRAKI, sRAKI-RNN supported arbitrary undersampling patterns and was a database-free technique. However, sRAKI-RNN reduced the whole reconstruction time of sRAKI and involved fewer hyper-parameters further facilitating reconstruction. Densely connected blocks were used in each iteration of the recurrent neural networks to improve the convergence during the learning phase.

In Chapter 6, we proposed a history-cognizant approach to unroll a given optimization algorithm for solving inverse problems in medical imaging using physics-driven deep learning. The proposed approach led to an unrolled dense recurrent neural network architecture. From a network-design perspective, this consequent architecture enhanced training and performance by facilitating information flow through the network. Results of this chapter, which were tested in accelerated MRI reconstruction experiments, revealed that the proposed unrolling approach may considerably enhance reconstruction quality without substantially modifying computational complexity in physics-driven deep learning.

7.2 Limitations and Future Directions

Head volume conduction modeling (for iEEG recordings) is one of the limitations of the study in Chapter 2, which can be improved in future works. Currently, we used an infinitely homogeneous model for the deep iEEG electrodes. Caune et al have demonstrated [59] that an infinitely homogeneous modeling for deep electrodes can lead into reasonable results. In spite of that, one can always expect to improve accuracy and predictability of the simulations for real clinical data by taking into consideration the geometry and electrical conductivity of the cranial tissues in more sophisticated and realistic models such as finite element method (FEM). Additionally, by assuming point-sensing contacts, the effect of iEEG electrodes located in the medium is not considered in this study. This assumption was made based on the fact that perturbation produced in the intracranial electric potential distributions due to the presence of the depth electrodes was negligible, if source of electrical activity was located more than 1 mm away

from the electrode [230]. To comply fully with this finding, the active nodes in this study were located such that the minimum distance from the electrodes is 3 mm. Validation of the conclusions stated in this study on real clinical data, is a natural extension of this work and the next step of this research study.

In Chapter 3, a fixed-orientation model was used to solve the inverse problem. Although this choice can be justified through the anatomical location and orientation of pyramidal neuron cells generating EEG/MEG signals, estimating the orientations through solving the inverse problem rather than assuming the orientation a priori may be beneficial. For instance, segmentation of the cortex may not be accurate enough, and the normal orientations estimated from such segmentation results may be inaccurate, especially if the segmentation is coarse. Estimating the orientation of dipoles while solving the inverse problem may be achieved by considering free orientations for the current dipoles along different axis e.g., \hat{x} , \hat{y} and \hat{z} in a Cartesian coordinate. Extension of this work to a vector robust beamformer, which is based on rotational models, will be the topic of a future study.

Finally, while we showed the feasibility of using sRAKI and sRAKI-RNN for accelerating coronary MRI, we have not comprehensively evaluated the failure modes and the residual artifacts, and how these would affect diagnostic decisions. Further study of these effects in patient populations will be explored in future research to establish diagnostic utility.

References

- [1] P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky, and J. S. Hyde. Time course EPI of human brain function during task activation. *Magnetic resonance in medicine*, 25(2):390–397, 1992.
- [2] K. K. Kwong, J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, and R. Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, 89(12):5675–5679, 1992.
- [3] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89(13):5951–5955, 1992.
- [4] M. M. Ter-Pogossian, M. E. Phelps, E. J. Hoffman, and N. A. Mullani. A positron-emission transaxial tomograph for nuclear imaging (PETT). *Radiology*, 114(1):89–98, 1975.
- [5] B. He, T. Musha, Y. Okamoto, S. Homma, Y. Nakajima, and T. Sato. Electric dipole tracing in the brain by means of the boundary element method and its accuracy. *IEEE Transactions on Biomedical Engineering*, (6):406–414, 1987.
- [6] C. M. Michel and B. He. EEG mapping and source imaging. In *Niedermeyer's Electroencephalography: Basic Principles, Clinical Applications, and Related Fields: Sixth Edition*, pages 1179–1202. Wolters Kluwer Health Adis (ESP), 2012.

- [7] E. Niedermeyer and F. L. da Silva. *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins, 2005.
- [8] D. Cohen. Magnetoencephalography: detection of the brain's electrical activity with a superconducting magnetometer. *Science*, 175(4022):664–666, 1972.
- [9] M. Hämäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa. Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of modern Physics*, 65(2):413, 1993.
- [10] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince. *MRI from Picture to Proton*. Cambridge university press, 2017.
- [11] P. C. Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *nature*, 242(5394):190–191, 1973.
- [12] Z.-P. Liang and P. C. Lauterbur. *Principles of magnetic resonance imaging: a signal processing perspective*. SPIE Optical Engineering Press, 2000.
- [13] M. A. Brown, R. C. Semelka, and B. M. Dale. *MRI: basic principles and applications*. John Wiley & Sons, 2015.
- [14] G. T. Herman. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Science & Business Media, 2009.
- [15] J. T. Hathcock and R. L. Stickle. Principles and concepts of computed tomography. *Veterinary Clinics of North America: Small Animal Practice*, 23(2):399–415, 1993.
- [16] B. He, L. Yang, C. Wilke, and H. Yuan. Electrophysiological imaging of brain activity and connectivity—challenges and opportunities. *IEEE transactions on biomedical engineering*, 58(7):1918–1931, 2011.
- [17] B. J. Edelman, N. Johnson, A. Sohrabpour, S. Tong, N. Thakor, and B. He. Systems neuroengineering: Understanding and interacting with the brain. *Engineering*, 1(3):292–308, 2015.

- [18] B. He and L. Ding. Electrophysiological mapping and neuroimaging. In *Neural engineering*, pages 499–543. Springer, 2013.
- [19] C. M. Michel, M. M. Murray, G. Lantz, S. Gonzalez, L. Spinelli, and R. G. de Peralta. EEG source imaging. *Clinical neurophysiology*, 115(10):2195–2222, 2004.
- [20] M. Ding and B. He. Exploring functional and causal connectivity in the brain. In *Neural engineering*, pages 545–564. Springer, 2013.
- [21] Y. Zhang, W. van Drongelen, M. Kohrman, and B. He. Three-dimensional brain current source reconstruction from intra-cranial ECoG recordings. *Neuroimage*, 42(2):683–695, 2008.
- [22] A. L. Benabid, S. Chabardes, J. Mitrofanis, and P. Pollak. Deep brain stimulation of the subthalamic nucleus for the treatment of parkinson’s disease. *The Lancet Neurology*, 8(1):67–81, 2009.
- [23] M. J. Morrell. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology*, 77(13):1295–1304, 2011.
- [24] S. Stanslaski, P. Afshar, P. Cong, J. Giftakis, P. Stypulkowski, D. Carlson, D. Linde, D. Ullestad, A.-T. Avestruz, and T. Denison. Design and validation of a fully implantable, chronic, closed-loop neuromodulation device with concurrent sensing and stimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):410–421, 2012.
- [25] J. Talairach, J. Bancaud, G. Szikla, A. Bonis, S. Geier, and C. Vedrenne. New approach to the neurosurgery of epilepsy. stereotaxic methodology and therapeutic results. 1. introduction and history. *Neuro-chirurgie*, 20:1, 1974.
- [26] E. J. Benjamin, P. Muntner, and M. S. Bittencourt. Heart disease and stroke statistics-2019 update: a report from the american heart association. *Circulation*, 139(10):e56–e528, 2019.

- [27] W. Y. Kim, P. G. Danias, M. Stuber, S. D. Flamm, S. Plein, E. Nagel, S. E. Langerak, O. M. Weber, E. M. Pedersen, M. Schmidt, et al. Coronary magnetic resonance angiography for the detection of coronary stenoses. *New England Journal of Medicine*, 345(26):1863–1869, 2001.
- [28] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*, 42:952–962, 1999.
- [29] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*, 47(6):1202–1210, 2002.
- [30] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing MRI. *IEEE Sig Proc magazine*, 25(2):72, 2008.
- [31] M. Lustig, D. Donoho, and J. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med*, 58:1182–1195, 2007.
- [32] M. Akçakaya, S. Moeller, S. Weingärtner, and K. Ugurbil. Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging. *Magn Reson Med*, 81(1):439–453, 2019.
- [33] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang. Accelerating magnetic resonance imaging via deep learning. In *Proc IEEE ISBI*, pages 514–517, 2016.
- [34] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magn Reson Med*, 79:3055–3071, 2018.
- [35] D. Lee, J. Yoo, S. Tak, and J. C. Ye. Deep residual learning for accelerated MRI using magnitude and phase networks. *IEEE Trans Biomed Eng*, 65(9):1985–1995, 2018.

- [36] Y. Han, J. Yoo, H. H. Kim, H. J. Shin, K. Sung, and J. C. Ye. Deep learning with domain adaptation for accelerated projection-reconstruction MR. *Magn Reson Med*, 80:1189–1205, 2018.
- [37] H. K. Aggarwal, M. P. Mani, and M. Jacob. Modl: Model-based deep learning architecture for inverse problems. *IEEE Trans Med Imaging*, 38(2):394–405, 2018.
- [38] C. Qin, J. Schlemper, J. Caballero, A. N. Price, J. V. Hajnal, and D. Rueckert. Convolutional recurrent neural networks for dynamic MR image reconstruction. *IEEE Trans Med Imaging*, 38:280–290, 2018.
- [39] K. Kwon, D. Kim, and H. Park. A parallel MR imaging method using multilayer perceptron. *Medical Physics*, 44(12):6209–6224, 2017.
- [40] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Trans Med Imaging*, 37(2):491–503, 2017.
- [41] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, et al. DAGAN: deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction. *IEEE Trans Med Imaging*, 37(6):1310–1321, 2017.
- [42] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo. Deep learning for undersampled MRI reconstruction. *Physics in Medicine & Biology*, 63(13):135007, 2018.
- [43] T. Eo, Y. Jun, T. Kim, J. Jang, H. Lee, and D. Hwang. KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images. *Magn Reson Med*, 80(5):2188–2201, 2018.
- [44] J. Cheng, M. Mardani, M. Alley, J. Pauly, and S. Vasanawala. Deepspirit: generalized parallel imaging using deep convolutional neural networks. In *Proceedings of the Joint Annual Meeting of ISMRM-ESMRMB, Paris, France*, page 570, 2018.
- [45] Y. Han, L. Sunwoo, and J. C. Ye. k-space deep learning for accelerated MRI. *IEEE Trans Med Imaging*, 2019.

- [46] H. K. Aggarwal, M. P. Mani, and M. Jacob. Multi-shot sensitivity-encoded diffusion mri using model-based deep learning (modl-mussels). In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1541–1544. IEEE, 2019.
- [47] S. U. H. Dar, M. Özbey, A. B. Çatlı, and T. Çukur. A transfer-learning approach for accelerated MRI using deep neural networks. *preprint arXiv:1710.02615*, 2017.
- [48] Y. C. Eldar, A. O. Hero III, L. Deng, J. Fessler, J. Kovacevic, H. V. Poor, and S. Young. Challenges and open problems in signal processing: Panel discussion summary from ICASSP 2017. *IEEE Sig Proc Magazine*, 34:8–23, 2017.
- [49] J. F. Annegers, W. A. Hauser, and L. R. Elveback. Remission of seizures and relapse in patients with epilepsy. *Epilepsia*, 20(6):729–737, 1979.
- [50] O. C. Cockerell, J. Sander, Y. M. Hart, S. D. Shorvon, and A. Johnson. Remission of epilepsy: results from the national general practice study of epilepsy. *The Lancet*, 346(8968):140–144, 1995.
- [51] W. A. Hauser and L. T. Kurland. The epidemiology of epilepsy in rochester, minnesota, 1935 through 1967. *Epilepsia*, 16(1):1–66, 1975.
- [52] P. Kwan and J. Sander. The natural history of epilepsy: an epidemiological view. *Journal of Neurology, Neurosurgery & Psychiatry*, 75(10):1376–1381, 2004.
- [53] J. Engel. Surgical treatment for epilepsy: too little, too late? *Jama*, 300(21):2548–2550, 2008.
- [54] F. Rosenow and H. Lüders. Presurgical evaluation of epilepsy. *Brain*, 124(9):1683–1700, 2001.
- [55] M. Gavaret, A. Trébuchon, F. Bartolomei, P. Marquis, A. McGonigal, F. Wendling, J. Regis, J.-M. Badier, and P. Chauvel. Source localization of scalp-EEG interictal spikes in posterior cortex epilepsies investigated by HR-EEG and SEEG. *Epilepsia*, 50(2):276–289, 2009.

- [56] A. McGonigal, F. Bartolomei, J. Régis, M. Guye, M. Gavaret, A. T.-D. Fonseca, H. Dufour, D. Figarella-Branger, N. Girard, J.-C. Péraгут, et al. Stereoelectroencephalography in presurgical assessment of MRI-negative epilepsy. *Brain*, 130(12):3169–3183, 2007.
- [57] G. Alarcon, C. Binnie, R. Elwes, and C. Polkey. Power spectrum and intracranial EEG patterns at seizure onset in partial epilepsy. *Electroencephalography and clinical neurophysiology*, 94(5):326–337, 1995.
- [58] S. Le Cam, R. Ranta, V. Caune, G. Korats, L. Koessler, L. Maillard, and V. Louis-Dorr. SIEEG dipole source localization based on an empirical bayesian approach taking into account forward model uncertainties. *NeuroImage*, 153:1–15, 2017.
- [59] V. Caune, R. Ranta, S. Le Cam, J. Hofmanis, L. Maillard, L. Koessler, and V. Louis-Dorr. Evaluating dipolar source localization feasibility from intracerebral SIEEG recordings. *NeuroImage*, 98:118–133, 2014.
- [60] N. Chang, R. Gulrajani, and J. Gotman. Dipole localization using simulated intracerebral EEG. *Clinical neurophysiology*, 116(11):2707–2716, 2005.
- [61] M. Dümpelmann, J. Fell, J. Wellmer, H. Urbach, and C. E. Elger. 3D source localization derived from subdural strip and grid electrodes: a simulation study. *Clinical neurophysiology*, 120(6):1061–1069, 2009.
- [62] C. Baumgartner, N. Alexopoulos, S. Taylor, and R. L. Rogers. MEG and ECoG localization accuracy test. *Electroencephalography and clinical neurophysiology*, 94:109–114, 1995.
- [63] G. Ramantani, D. Cosandier-Rimélé, A. Schulze-Bonhage, L. Maillard, J. Zentner, and M. Dümpelmann. Source reconstruction based on subdural EEG recordings adds to the presurgical evaluation in refractory frontal lobe epilepsy. *Clinical Neurophysiology*, 124(3):481–491, 2013.
- [64] B. Yvert, C. Fischer, O. Bertrand, and J. Pernier. Localization of human supratemporal auditory areas from intracerebral auditory evoked potentials using distributed source models. *Neuroimage*, 28(1):140–153, 2005.

- [65] C. J. Holmes, R. Hoge, L. Collins, R. Woods, A. W. Toga, and A. C. Evans. Enhancement of MR images using registration for signal averaging. *Journal of computer assisted tomography*, 22(2):324–333, 1998.
- [66] L. Nobili, I. Sartori, M. Terzaghi, F. Stefano, R. Mai, L. Tassi, L. Parrino, M. Cossu, and G. L. Russo. Relationship of epileptic discharges to arousal instability and periodic leg movements in a case of nocturnal frontal lobe epilepsy: a stereo-EEG study. *Sleep*, 29(5):701–704, 2006.
- [67] M. S. Hamalainen and J. Sarvas. Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE transactions on biomedical engineering*, 36(2):165–171, 1989.
- [68] Y. Lai, W. Van Drongelen, L. Ding, K. Hecox, V. Towle, D. Frim, and B. He. Estimation of in vivo human brain-to-skull conductivity ratio from simultaneous extra-and intra-cranial electrical potential recordings. *Clinical neurophysiology*, 116(2):456–465, 2005.
- [69] T. F. Oostendorp, J. Delbeke, and D. F. Stegeman. The conductivity of the human skull: results of in vivo and in vitro measurements. *IEEE transactions on biomedical engineering*, 47(11):1487–1492, 2000.
- [70] Y. Zhang, W. Van Drongelen, and B. He. Estimation of in vivo brain-to-skull conductivity ratio in humans. *Applied physics letters*, 89(22):223903, 2006.
- [71] R. D. Pascual-Marqui et al. Standardized low-resolution brain electromagnetic tomography (sloreta): technical details. *Methods Find Exp Clin Pharmacol*, 24(Suppl D):5–12, 2002.
- [72] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- [73] A. Sohrabpour, S. Ye, G. A. Worrell, W. Zhang, and B. He. Noninvasive electromagnetic source imaging and granger causality analysis: an electrophysiological connectome (eConnectome) approach. *IEEE Transactions on Biomedical Engineering*, 63(12):2474–2487, 2016.

- [74] S. Alomar, J. Jones, A. Maldonado, and J. Gonzalez-Martinez. The stereo-electroencephalography methodology. *Neurosurgery Clinics*, 27(1):83–95, 2016.
- [75] L. Koessler, C. Benar, L. Maillard, J.-M. Badier, J. P. Vignal, F. Bartolomei, P. Chauvel, and M. Gavaret. Source localization of ictal epileptic activity investigated by high resolution EEG and validated by SEEG. *Neuroimage*, 51(2):642–653, 2010.
- [76] L. Ding, G. A. Worrell, T. D. Lagerlund, and B. He. Ictal source analysis: localization and imaging of causal interactions in humans. *Neuroimage*, 34(2):575–586, 2007.
- [77] C. W. Granger. Testing for causality: a personal viewpoint. *Journal of Economic Dynamics and control*, 2:329–352, 1980.
- [78] M. Kamiński, M. Ding, W. A. Truccolo, and S. L. Bressler. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biological cybernetics*, 85(2):145–157, 2001.
- [79] M. J. Kaminski and K. J. Blinowska. A new method of the description of the information flow in the brain structures. *Biological cybernetics*, 65(3):203–210, 1991.
- [80] M. Palus and D. Hoyer. Detecting nonlinearity and phase synchronization with surrogate data. *IEEE Engineering in Medicine and Biology Magazine*, 17(6):40–45, 1998.
- [81] J. Theiler, B. Galdrikian, A. Longtin, S. Eubank, and J. D. Farmer. Testing for nonlinearity in time series: the method of surrogate data. Technical report, Los Alamos National Lab., NM (United States), 1991.
- [82] Y. Lu, L. Yang, G. A. Worrell, and B. He. Seizure source imaging by means of FINE spatio-temporal dipole localization and directed transfer function in partial epilepsy patients. *Clinical Neurophysiology*, 123(7):1275–1283, 2012.

- [83] C. Wilke, L. Ding, B. He, et al. Estimation of time-varying connectivity patterns through the use of an adaptive directed transfer function. *IEEE transactions on biomedical engineering*, 55(11):2557–2564, 2008.
- [84] B. He and J. Lian. Electrophysiological neuroimaging. In *Neural engineering*, pages 221–261. Springer, 2005.
- [85] P. Malmivuo, J. Malmivuo, and R. Plonsey. *Bioelectromagnetism: principles and applications of bioelectric and biomagnetic fields*. Oxford University Press, USA, 1995.
- [86] F. Bauer and M. A. Lukas. Comparing parameter choice methods for regularization of ill-posed problems. *Mathematics and Computers in Simulation*, 81(9):1795–1841, 2011.
- [87] P. C. Hansen. Regularization tools: a matlab package for analysis and solution of discrete ill-posed problems. *Numerical algorithms*, 6(1):1–35, 1994.
- [88] D. A. Engemann and A. Gramfort. Automated model selection in covariance estimation and spatial whitening of MEG and EEG signals. *NeuroImage*, 108:328–342, 2015.
- [89] H. W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [90] B. He, Y. Dai, L. Astolfi, F. Babiloni, H. Yuan, and L. Yang. eConnectome: A MATLAB toolbox for mapping and imaging of brain functional connectivity. *Journal of neuroscience methods*, 195(2):261–269, 2011.
- [91] R. P. Marqui, C. M. Michel, and D. Lehmann. Low-resolution electromagnetic tomography—a new method for localizing electrical activity in the brain. *International Journal of psychophysiology*, 18:49–65, 1994.
- [92] R. Fisher, V. Salanova, T. Witt, R. Worth, T. Henry, R. Gross, K. Oommen, I. Osorio, J. Nazzaro, D. Labar, et al. Electrical stimulation of the anterior nucleus of thalamus for treatment of refractory epilepsy. *Epilepsia*, 51(5):899–908, 2010.

- [93] M. D. Johnson, H. H. Lim, T. I. Netoff, A. T. Connolly, N. Johnson, A. Roy, A. Holt, K. O. Lim, J. R. Carey, J. L. Vitek, et al. Neuromodulation for brain disorders: challenges and opportunities. *IEEE Transactions on Biomedical Engineering*, 60(3):610–624, 2013.
- [94] H. Martens, E. Toader, M. Decré, D. Anderson, R. Vetter, D. Kipke, K. B. Baker, M. D. Johnson, and J. L. Vitek. Spatial steering of deep brain stimulation volumes using a novel lead design. *Clinical neurophysiology*, 122(3):558–566, 2011.
- [95] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.
- [96] D. J. Mogul and W. van Drongelen. Electrical control of epilepsy. *Annual review of biomedical engineering*, 16:483–504, 2014.
- [97] C. O. Oluigbo, A. Salma, and A. R. Rezai. Deep brain stimulation for neurological disorders. *IEEE reviews in biomedical engineering*, 5:88–99, 2012.
- [98] C. J. Stam. Modern network science of neurological disorders. *Nature Reviews Neuroscience*, 15(10):683–695, 2014.
- [99] C. Wilke, G. Worrell, and B. He. Graph analysis of epileptogenic networks in human partial epilepsy. *Epilepsia*, 52(1):84–93, 2011.
- [100] C. Wilke, W. Van Drongelen, M. Kohnman, and B. He. Neocortical seizure foci localization by means of a directed transfer function method. *Epilepsia*, 51(4):564–572, 2010.
- [101] C. Wilke, W. Van Drongelen, M. Kohnman, and B. He. Identification of epileptogenic foci from causal analysis of ECoG interictal spike activity. *Clinical Neurophysiology*, 120(8):1449–1456, 2009.
- [102] Y. Zhang, L. Ding, W. van Drongelen, K. Hecox, D. M. Frim, and B. He. A cortical potential imaging study from simultaneous extra-and intracranial electrical recordings by means of the finite element method. *NeuroImage*, 31(4):1513–1524, 2006.

- [103] K. Bryan and T. Leise. Making do with less: An introduction to compressed sensing. *Siam Review*, 55(3):547–566, 2013.
- [104] A. Sohrabpour, Y. Lu, G. Worrell, and B. He. Imaging brain source extent from EEG/MEG by means of an iteratively reweighted edge sparsity minimization (IRES) strategy. *NeuroImage*, 142:27–42, 2016.
- [105] S. A. H. Hosseini, A. Sohrabpour, and B. He. Electromagnetic source imaging using simultaneous scalp EEG and intracranial EEG: An emerging tool for interacting with pathological brain networks. *Clinical Neurophysiology*, 129(1):168–187, 2018.
- [106] B. He, A. Sohrabpour, E. Brown, and Z. Liu. Electrophysiological source imaging: A noninvasive window to brain dynamics. *Annual review of biomedical engineering*, 20:171–196, 2018.
- [107] M. E. Spencer, R. M. Leahy, J. Mosher, and P. Lewis. Adaptive filters for monitoring localized brain activity from surface potential time series. In *ASILOMAR CONFERENCE ON SIGNALS SYSTEMS AND COMPUTERS*, pages 156–156. COMPUTER SOCIETY PRESS, 1992.
- [108] B. D. Van Veen, W. Van Drongelen, M. Yuchtman, and A. Suzuki. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on biomedical engineering*, 44(9):867–880, 1997.
- [109] O. Steinsträter, S. Sillekens, M. Junghoefer, M. Burger, and C. H. Wolters. Sensitivity of beamformer source analysis to deficiencies in forward modeling. *Human brain mapping*, 31(12):1907–1927, 2010.
- [110] K. Sekihara, S. S. Nagarajan, D. Poeppel, A. Marantz, and Y. Miyashita. Application of an MEG eigenspace beamformer to reconstructing spatio-temporal activities of neural sources. *Human brain mapping*, 15(4):199–215, 2002.
- [111] S. A. Vorobyov, A. B. Gershman, and Z.-Q. Luo. Robust adaptive beamforming using worst-case performance optimization: A solution to the signal mismatch problem. *IEEE transactions on signal processing*, 51(2):313–324, 2003.

- [112] J. Groß and A. Ioannides. Linear transformations of data space in MEG. *Physics in Medicine & Biology*, 44(8):2081, 1999.
- [113] S. Robinson. Functional neuroimaging by synthetic aperture magnetometry (SAM). *Recent advances in biomagnetism*, 1999.
- [114] K. Sekihara, S. S. Nagarajan, D. Poeppel, A. Marantz, and Y. Miyashita. Reconstructing spatio-temporal activities of neural sources using an MEG vector beamformer technique. *IEEE Transactions on Biomedical Engineering*, 48(7):760–771, 2001.
- [115] A. B. Gershman, N. D. Sidiropoulos, S. Shahbazpanahi, M. Bengtsson, and B. Ottersten. Convex optimization-based beamforming. *IEEE Signal Processing Magazine*, 27(3):62–75, 2010.
- [116] R. G. Lorenz and S. P. Boyd. Robust minimum variance beamforming. *IEEE transactions on signal processing*, 53(5):1684–1696, 2005.
- [117] S. A. H. Hosseini, A. Sohrabpour, M. Akçakaya, and B. He. Electromagnetic brain source imaging by means of a robust minimum variance beamformer. *IEEE Transactions on Biomedical Engineering*, 65(10):2365–2374, 2018.
- [118] J. C. Mosher and R. M. Leahy. Source localization using recursively applied and projected (RAP) MUSIC. *IEEE Transactions on signal processing*, 47(2):332–340, 1999.
- [119] X.-L. Xu, B. Xu, and B. He. An alternative subspace approach to EEG dipole source localization. *Physics in Medicine & Biology*, 49(2):327, 2004.
- [120] M. X. Huang, J. Shih, R. Lee, D. Harrington, R. Thoma, M. Weisend, F. Hanlon, K. Paulson, T. Li, K. Martin, et al. Commonalities and differences among vectorized beamformers in electromagnetic source imaging. *Brain topography*, 16(3):139–158, 2004.
- [121] H. F. Kaiser. The application of electronic computers to factor analysis. *Educational and psychological measurement*, 20(1):141–151, 1960.

- [122] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [123] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear algebra and its applications*, 284(1-3):193–228, 1998.
- [124] M. C. Grant and S. P. Boyd. Graph implementations for nonsmooth convex programs. In *Recent advances in learning and control*, pages 95–110. Springer, 2008.
- [125] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [126] S. T. Hansen, S. Hauberg, and L. K. Hansen. Data-driven forward model inference for EEG brain imaging. *NeuroImage*, 139:249–258, 2016.
- [127] S. Baillet, J. C. Mosher, and R. M. Leahy. Electromagnetic brain mapping. *IEEE Signal processing magazine*, 18(6):14–30, 2001.
- [128] S. M. Kay. *Fundamentals of statistical signal processing*. Prentice Hall PTR, 1993.
- [129] C. Grova, J. Daunizeau, J.-M. Lina, C. G. Bénar, H. Benali, and J. Gotman. Evaluation of EEG localization methods using realistic simulations of interictal spikes. *Neuroimage*, 29(3):734–753, 2006.
- [130] B. W. Matthews. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451, 1975.
- [131] M. Stuber, R. M. Botnar, P. G. Danias, D. K. Sodickson, K. V. Kissinger, M. Van Cauteren, J. De Becker, and W. J. Manning. Double-oblique free-breathing high resolution three-dimensional coronary magnetic resonance angiography. *Journal of the American College of Cardiology*, 34(2):524–531, 1999.
- [132] J. N. Oshinski, L. Hoffland, S. Mukundan Jr, W. T. Dixon, W. J. Parks, and R. I. Pettigrew. Two-dimensional coronary mr angiography without breath holding. *Radiology*, 201(3):737–743, 1996.

- [133] M. V. McConnell, V. C. Khasgiwala, B. J. Savord, M. H. Chen, M. L. Chuang, W. J. Manning, and R. R. Edelman. Prospective adaptive navigator correction for breath-hold MR coronary angiography. *Magn Reson Med*, 37(1):148–152, 1997.
- [134] R. R. Edelman, W. Manning, D. Burstein, and S. Paulin. Coronary arteries: breath-hold mr angiography. *Radiology*, 181(3):641–643, 1991.
- [135] P. A. Wielopolski, R. Van Geuns, P. De Feyter, and M. Oudkerk. Breath-hold coronary mr angiography with volume-targeted imaging. *Radiology*, 209(1):209–219, 1998.
- [136] D. Li, S. Kaushikkar, E. M. Haacke, P. K. Woodard, P. J. Dhawale, R. M. Kroeker, G. Laub, Y. Kuginuki, and F. R. Gutierrez. Coronary arteries: three-dimensional mr imaging with retrospective respiratory gating. *Radiology*, 201(3):857–863, 1996.
- [137] T. Niendorf, C. J. Hardy, R. O. Giaquinto, P. Gross, H. E. Cline, Y. Zhu, G. Kenwood, S. Cohen, A. K. Grant, S. Joshi, et al. Toward single breath-hold whole-heart coverage coronary mra using highly accelerated parallel imaging with a 32-channel mr system. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 56(1):167–176, 2006.
- [138] O. M. Weber, A. J. Martin, and C. B. Higgins. Whole-heart steady-state free precession coronary artery magnetic resonance angiography. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 50(6):1223–1228, 2003.
- [139] A. Etienne, R. M. Botnar, A. M. Van Muiswinkel, P. Boesiger, W. J. Manning, and M. Stuber. “soap-bubble” visualization and quantitative analysis of 3d coronary magnetic resonance angiograms. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 48(4):658–666, 2002.
- [140] P. Hu, J. Chan, L. H. Ngo, J. Smink, B. Goddu, K. V. Kissinger, L. Goepfert, T. H. Hauser, N. M. Rofsky, W. J. Manning, et al. Contrast-enhanced whole-heart

coronary mri with bolus infusion of gadobenate dimeglumine at 1.5 t. *Magnetic resonance in medicine*, 65(2):392–398, 2011.

- [141] X. Bi, J. C. Carr, and D. Li. Whole-heart coronary magnetic resonance angiography at 3 tesla in 5 minutes with slow infusion of gd-bopta, a high-relaxivity clinical contrast agent. *Magn Reson Med*, 58(1):1–7, 2007.
- [142] M. Akçakaya, T. A. Basha, R. H. Chan, H. Rayatzadeh, K. V. Kissinger, B. Goddu, L. A. Goepfert, W. J. Manning, and R. Nezafat. Accelerated contrast-enhanced whole-heart coronary MRI using low-dimensional-structure self-learning and thresholding. *Magn Reson Med*, 67(5):1434–1443, 2012.
- [143] L. Feng, S. Coppo, D. Piccini, J. Yerly, R. P. Lim, P. G. Masci, M. Stuber, D. K. Sodickson, and R. Otazo. 5D whole-heart sparse MRI. *Magn Reson Med*, 79(2):826–838, 2018.
- [144] M. Akçakaya, T. A. Basha, B. Goddu, L. A. Goepfert, K. V. Kissinger, V. Tarokh, W. J. Manning, and R. Nezafat. Low-dimensional-structure self-learning and thresholding: regularization beyond compressed sensing for MRI reconstruction. *Magn Reson Med*, 66(3):756–767, 2011.
- [145] M. Akçakaya, T. A. Basha, R. H. Chan, W. J. Manning, and R. Nezafat. Accelerated isotropic sub-millimeter whole-heart coronary MRI: compressed sensing versus parallel imaging. *Magn Reson Med*, 71(2):815–822, 2014.
- [146] D. Piccini, L. Feng, G. Bonanno, S. Coppo, J. Yerly, R. P. Lim, J. Schwitter, D. K. Sodickson, R. Otazo, and M. Stuber. Four-dimensional respiratory motion-resolved whole heart coronary MR angiography. *Magn Reson Med*, 77(4):1473–1484, 2017.
- [147] C. Forman, D. Piccini, R. Grimm, J. Hutter, J. Hornegger, and M. O. Zenge. High-resolution 3D whole-heart coronary MRA: a study on the combination of data acquisition in multiple breath-holds and 1D residual respiratory motion compensation. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 27(5):435–443, 2014.

- [148] T. Correia, G. Ginami, G. Cruz, R. Neji, I. Rashid, R. M. Botnar, and C. Prieto. Optimized respiratory-resolved motion-compensated 3 DC artesian coronary MR angiography. *Magn Reson Med*, 80(6):2618–2629, 2018.
- [149] A. P. Aitken, M. Henningsson, R. M. Botnar, T. Schaeffter, and C. Prieto. 100% efficient three-dimensional coronary MR angiography with two-dimensional beat-to-beat translational and bin-to-bin affine motion correction. *Magn Reson Med*, 74(3):756–764, 2015.
- [150] G. Cruz, D. Atkinson, M. Henningsson, R. M. Botnar, and C. Prieto. Highly efficient nonrigid motion-corrected 3D whole-heart coronary vessel wall imaging. *Magn Reson Med*, 77(5):1894–1908, 2017.
- [151] C. Munoz, G. Cruz, R. Neji, R. M. Botnar, and C. Prieto. Motion corrected water/fat whole-heart coronary MR angiography with 100% respiratory efficiency. *Magn Reson Med*, 82(2):732–742, 2019.
- [152] C. Zhang, S. A. H. Hosseini, S. Weingärtner, K. Uğurbil, S. Moeller, and M. Akçakaya. Optimized fast GPU implementation of robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction. *PloS one*, 14(10), 2019.
- [153] M. Lustig and J. M. Pauly. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic resonance in medicine*, 64(2):457–471, 2010.
- [154] Y. Chang, D. Liang, and L. Ying. Nonlinear grappa: A kernel approach to parallel mri reconstruction. *Magnetic resonance in medicine*, 68(3):730–740, 2012.
- [155] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, and M. Lustig. Fast ℓ_1 -spirit compressed sensing parallel imaging MRI: Scalable parallel implementation and clinically feasible runtime. *IEEE transactions on medical imaging*, 31(6):1250–1262, 2012.
- [156] C.-Y. Liou, W.-C. Cheng, J.-W. Liou, and D.-R. Liou. Autoencoder for words. *Neurocomputing*, 139:84–96, 2014.

- [157] D. Kingma and J. Ba. Adam: A method for stochastic optimization, in proceedings of the 3rd international conference for learning representations, 2015.
- [158] R. Deriche. Fast algorithms for low-level vision. *IEEE transactions on pattern analysis and machine intelligence*, 12(1):78–87, 1990.
- [159] N. Fuin, A. Bustin, R. Botnar, and C. Prieto. A variational neural network for whole-heart coronary MRA reconstruction. Bellevue, WA USA, February 2019. SCMR 22nd Annual Scientific Sessions.
- [160] T. Kim, P. Garg, and J. Haldar. LORAKI: Reconstruction of undersampled k-space data using scan-specific autocalibrated recurrent neural networks. In *Proc. Int. Soc. Magn. Reson. Med*, page 4647, 2019.
- [161] S. A. H. Hosseini, S. Moeller, S. Weingärtner, K. Uğurbil, and M. Akçakaya. Accelerated coronary MRI using 3D SPIRiT-RAKI with sparsity regularization. In *Proc. IEEE ISBI*, pages 1692–1695, 2019.
- [162] S. A. H. Hosseini, S. Moeller, S. Weingartner, K. Ugurbil, and M. Akcakaya. Accelerated targeted coronary MRI using sparsity-regularized SPIRiT-RAKI. Montreal, QC Canada, May 2019. ISMRM 27th Annual Meeting and Exhibition.
- [163] S. A. H. Hosseini, S. Moeller, S. Weingartner, K. Ugurbil, and M. Akcakaya. Accelerated coronary MRI using SPIRiT-RAKI with gradient descent-based self-consistency. Bellevue, WA USA, February 2019. SCMR 22nd Annual Scientific Sessions.
- [164] S. Hosseini, S. Moeller, S. Weingartner, K. Uğurbil, and M. Akçakaya. SPIRiT-RAKI: Scan-specific self-consistency neural networks for reconstruction arbitrary k-space. In *ISMRM Workshop on Machine Learning Part II*, 2018.
- [165] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med*, 46(4):638–651, 2001.

- [166] K. Zhang, W. Zuo, S. Gu, and L. Zhang. Learning deep CNN denoiser prior for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3929–3938, 2017.
- [167] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708, 2017.
- [168] B. P. Sutton, D. C. Noll, and J. A. Fessler. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE Trans Med Imaging*, 22(2):178–188, 2003.
- [169] K. T. Block, M. Uecker, and J. Frahm. Undersampled radial MRI with multiple coils. iterative image reconstruction using a total variation constraint. *Magn Reson Med*, 57(6):1086–1098, 2007.
- [170] J. A. Fessler. Model-based image reconstruction for MRI. *IEEE Sig Proc Magazine*, 27(4):81–89, 2010.
- [171] M. Akçakaya, S. Nam, P. Hu, M. H. Moghari, L. H. Ngo, V. Tarokh, W. J. Manning, and R. Nezafat. Compressed sensing with wavelet domain dependencies for coronary MRI: a retrospective study. *IEEE Trans Med Imaging*, 30(5):1090–1099, 2010.
- [172] F. Knoll, K. Bredies, T. Pock, and R. Stollberger. Second order total generalized variation (TGV) for MRI. *Magn Reson Med*, 65(2):480–491, 2011.
- [173] G. Wang, D. L. Snyder, J. A. O’Sullivan, and M. W. Vannier. Iterative deblurring for ct metal artifact reduction. *IEEE Trans Med Imaging*, 15:657–664, 1996.
- [174] J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh. A three-dimensional statistical approach to improved image quality for multislice helical CT. *Medical physics*, 34(11):4526–4544, 2007.
- [175] P. Jin, C. A. Bouman, and K. D. Sauer. A model-based image reconstruction algorithm with simultaneous beam hardening correction for X-ray CT. *IEEE Trans Comp Imaging*, 1(3):200–216, 2015.

- [176] L. Pfister and Y. Bresler. Model-based iterative tomographic reconstruction with adaptive sparsifying transforms. In *Comp Imaging XII*, volume 9020, page 90200H, 2014.
- [177] Z. Chang, D. H. Ye, S. Srivastava, J.-B. Thibault, K. Sauer, and C. Bouman. Prior-guided metal artifact reduction for iterative X-ray computed tomography. *IEEE Trans Med Imaging*, 38:1532–1542, 2018.
- [178] A. Tuysuzoglu, Y. Khoo, and W. C. Karl. Variable splitting techniques for discrete tomography. In *IEEE Int Conf Image Proc*, pages 1764–1768, 2016.
- [179] L. Tian and L. Waller. Quantitative differential phase contrast imaging in an LED array microscope. *Optics Express*, 23(9):11394–11403, 2015.
- [180] R. A. Claus, P. P. Naulleau, A. R. Neureuther, and L. Waller. Quantitative phase retrieval with arbitrary pupil and illumination. *Optics Express*, 23(20):26672–26682, 2015.
- [181] Y.-Z. Lin, K.-Y. Huang, and Y. Luo. Quantitative differential phase contrast imaging at high resolution with radially asymmetric illumination. *Optics letters*, 43(12):2973–2976, 2018.
- [182] S. Sreehari, S. V. Venkatakrishnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman. Plug-and-play priors for bright field electron tomography and sparse interpolation. *IEEE Trans Comp Imaging*, 2(4):408–423, 2016.
- [183] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler. Image reconstruction is a new frontier of machine learning. *IEEE Trans Med Imaging*, 37:1289–1296, 2018.
- [184] S. H. Dar and T. Çukur. Transfer learning for reconstruction of accelerated MRI acquisitions via neural networks. Paris, France, 2018. Proc ISMRM.
- [185] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang. Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE Trans Med Imaging*, 36:2524–2535, 2017.

- [186] Y. H. Yoon, S. Khan, J. Huh, and J. C. Ye. Efficient b-mode ultrasound image reconstruction from sub-sampled RF data using deep learning. *IEEE Trans Med Imaging*, 38(2):325–336, 2018.
- [187] M. U. Ghani and W. C. Karl. Fast enhanced CT metal artifact reduction using data domain deep learning. *IEEE Trans Comp Imaging*, 2019.
- [188] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang. Low-dose CT image denoising using a generative adversarial network with wasserstein distance and perceptual loss. *IEEE Trans Med Imaging*, 37(6):1348–1357, 2018.
- [189] E. Kang, W. Chang, J. Yoo, and J. C. Ye. Deep convolutional framelet denoising for low-dose CT via wavelet residual network. *IEEE Trans Med Imaging*, 37(6):1358–1369, 2018.
- [190] Z. Wu, Y. Sun, A. Matlock, J. Liu, L. Tian, and U. S. Kamilov. Simba: Scalable inversion in optical tomography using deep denoising priors. *preprint arXiv:1911.13241*, 2019.
- [191] M. Mardani, H. Monajemi, V. Pappyan, S. Vasanawala, D. Donoho, and J. Pauly. Recurrent generative adversarial networks for proximal learning and automated compressive image recovery. *preprint arXiv:1711.10046*, 2017.
- [192] J. Zhang and B. Ghanem. ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing. In *Proc IEEE CVPR*, pages 1828–1837, 2018.
- [193] E. Kobler, M. Muckley, B. Chen, F. Knoll, K. Hammernik, T. Pock, D. Sodickson, and R. Otazo. Variational deep learning for low-dose computed tomography. In *Proc IEEE ICASSP*, pages 6687–6691, 2018.
- [194] M. Kellman, E. Bostan, N. Repina, and L. Waller. Physics-based learned design: Optimized coded-illumination for quantitative phase imaging. *IEEE Trans Comp Imaging*, 2019.

- [195] Y. Yang, J. Sun, H. Li, and Z. Xu. Deep ADMM-Net for compressive sensing MRI. In *Advances in neural information processing systems*, pages 10–18, 2016.
- [196] D. H. Ye, S. Srivastava, J.-B. Thibault, K. Sauer, and C. Bouman. Deep residual learning for model-based iterative CT reconstruction using plug-and-play framework. In *Proc IEEE ICASSP*, pages 6668–6672, 2018.
- [197] J. Duan, J. Schlemper, C. Qin, C. Ouyang, W. Bai, C. Biffi, G. Bello, B. Statton, D. P. O’Regan, and D. Rueckert. VS-Net: Variable splitting network for accelerated parallel MRI reconstruction. In *Proc MICCAI*, pages 713–722, 2019.
- [198] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya. Self-supervised physics-based deep learning MRI reconstruction without fully-sampled data. *IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, 2020.
- [199] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya. Self-supervised learning of physics-based reconstruction neural networks without fully-sampled reference data. *arXiv preprint arXiv:1912.07669*, 2019.
- [200] J. Adler and O. Oktem. Learned Primal-Dual Reconstruction. *IEEE Trans Med Imaging*, 37(6):1322–1332, 06 2018.
- [201] J. Cheng, H. Wang, L. Ying, and D. Liang. Model learning: Primal dual networks for fast MR imaging. In *Proc MICCAI*, pages 21–29, 2019.
- [202] E. Bostan, U. S. Kamilov, and L. Waller. Learning-based image reconstruction via parallel proximal algorithm. *IEEE Sig Proc Letters*, 25(7):989–993, 2018.
- [203] S. A. H. Hosseini, C. Zhang, S. Weingärtner, S. Moeller, M. Stuber, K. Uğurbil, and M. Akçakaya. Accelerated coronary MRI with sRAKI: A database-free self-consistent neural network k-space reconstruction for arbitrary undersampling. *preprint arXiv:1907.08137*, 2019.

- [204] S. A. H. Hosseini, C. Zhang, K. Uğurbil, S. Moeller, and M. Akçakaya. sRAKI-RNN: accelerated MRI with scan-specific recurrent neural networks using densely connected blocks. In *SPIE Wavelets and Sparsity XVIII*, page 111381B, 2019.
- [205] F. Knoll, K. Hammernik, C. Zhang, S. Moeller, T. Pock, D. K. Sodickson, and M. Akcakaya. Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. *IEEE Signal Processing Magazine*, 37(1):128–140, 2020.
- [206] G. Wang. A perspective on deep imaging. *IEEE Access*, 4:8914–8924, 2016.
- [207] U. S. Kamilov and H. Mansour. Learning optimal nonlinearities for iterative thresholding algorithms. *IEEE Sig Proc Letters*, 23(5):747–751, 2016.
- [208] S. A. H. Hosseini, B. Yaman, S. Moeller, M. Hong, and M. Akçakaya. Dense recurrent neural networks for inverse problems: History-cognizant unrolling of optimization algorithms. *arXiv preprint arXiv:1912.07197*, 2019.
- [209] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- [210] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [211] Y. Nesterov and M. Florea. Gradient methods with memory. Technical report, 2019.
- [212] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imag Sci*, 2(1):183–202, 2009.
- [213] A. Chambolle and T. Pock. A remark on accelerated block coordinate descent for computing the proximity operators of a sum of convex functions. 2015. technical report.
- [214] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo. Iteration complexity analysis of block coordinate descent methods. *Mathematical Programming Series A*, 2016.

- [215] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. 2008. preprint.
- [216] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [217] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proc IEEE CVPR*, pages 4700–4708, 2017.
- [218] J. Zbontar, F. Knoll, A. Sriramand M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. Geras, J. Katsnelson, H. Chandarana, et al. fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1), 2020.
- [219] J. A. Fessler. Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms. *IEEE Signal Processing Magazine*, 37(1):33–40, 2020.
- [220] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [221] T. Goldstein and S. Osher. The split bregman method for L1-regularized problems. *SIAM journal on imaging sciences*, 2(2):323–343, 2009.
- [222] S. Ramani and J. A. Fessler. Parallel MR image reconstruction using augmented lagrangian methods. *IEEE Trans Med Imaging*, 30(3):694–706, 2010.
- [223] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.
- [224] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.

- [225] T. Valkonen. A primal–dual hybrid gradient method for nonlinear operators with applications to MRI. *Inverse Problems*, 30(5):055012, 2014.
- [226] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn Reson Med*, 71(3):990–1001, 2014.
- [227] R. Timofte, E. Agustsson, L. Van Gool, M.-H. Yang, and L. Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proc IEEE CVPR*, pages 114–125, 2017.
- [228] F. Knoll, K. Hammernik, E. Kobler, T. Pock, M. P. Recht, and D. K. Sodickson. Assessment of the generalization of learned image reconstruction and the potential for transfer learning. *Magn Reson Med*, 81:116–128, 2019.
- [229] H. Fan, P. Chu, L. J. Latecki, and H. Ling. Scene parsing via dense recurrent neural networks with attentional selection. In *Proc IEEE WACV*, pages 1816–1825, 2019.
- [230] N. von Ellenrieder, L. Beltrachini, and C. H. Muravchik. Electrode and brain modeling in stereo-EEG. *Clinical neurophysiology*, 123(9):1745–1754, 2012.
- [231] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Mathematical programming*, 95(1):3–51, 2003.

Appendix A

Supporting Information

A.1 Supporting Information for Chapter 2

See **Figures A.1** and **A.2**.

A.2 Supporting Information for Chapter 3

A.2.1 Convex Optimization

Convex optimization solvers, including CVX, employ a multitude of algorithms to solve optimization problems formulated as being convex. Interior point methods are famous algorithms used to solve convex optimization problems in general and second-order cone programs (SOCP) in particular [122]. These methods define a so-called “barrier function” that basically transforms constrained optimization problems into unconstrained ones. These methods impose a condition on the goal function of the transformed unconstrained optimization problem, that it cannot result in solutions that do not satisfy the constraints of the original optimization problem by penalizing solutions that get close to the constraint boundary (in limit penalizing solutions on the constraint boundary by infinity), hence creating a barrier that does not let the solution to pass which ensures the satisfaction of the constraint. How SOCP problems are tackled can be found in great details in [231]. The details of these algorithms have been well studied and practiced in the literature and given the limited space we have in the paper and our primary focus on the source imaging applications in this paper, and the wide use of these techniques

and solvers in the optimization field, we refer the interested reader to the optimization literature for further reading and investigation [122, 231].

A.2.2 Time-courses of Activity

Refer to **Figures A.3 and A.4**.

A.2.3 Practical Considerations

In this section, some considerations regarding the implementation of the beamformers in general and its robust version in particular will be discussed. The accuracy of the covariance estimation plays a crucial role in the ultimate performance of the beamformers. On the other hand, the empirical estimation can poorly represent the true covariance matrix, if the number of data points is not sufficiently large. To address this issue, many techniques have been developed to regularize and purify the estimation [88]. In this study, a technique based on shrinkage models [88] was used to estimate the covariance matrices of the measurements, noise and uncertainty regions.

Another issue that requires attention is the size of the uncertainty region for each voxel with respect to the size of the lead field column. The uncertainty ellipsoid should not be too large to include the origin; otherwise, the constraint in (3.8) becomes infeasible, which in turn makes the problem unsolvable. On the other hand, a nonconservative inflation of the uncertainty ellipsoid can degrade the performance. This issue is pronounced even more in a fixed orientation model as is the case of this study, since in such a model orientations may vary drastically within each region depending on the size and location of the uncertainty region. Therefore, a constricting mechanism is required to ensure that each uncertainty ellipsoid does not grow larger than necessary. To this end, we defined the parameter β_i for each voxel as the ellipsoid radius along the direction $\mathbf{k}_i/\|\mathbf{k}_i\|$ normalized by $\|\mathbf{k}_i\|$. This parameter, which measures the size of the ellipsoid with respect to the distance of its center from the origin, can be calculated as:

$$\beta_i = \frac{1}{\sqrt{\mathbf{k}_i^T \mathbf{P}_i^{-1} \mathbf{k}_i}}, \quad \forall i \in \{1, \dots, N\} \quad (\text{A.1})$$

Each uncertainty ellipsoid can then be re-scaled such that its β_i parameter is no larger than a cut-off value. The threshold can be determined either upfront or in an

automated manner such as the mean of all β_i parameters that are less than 1 (of the ellipsoids that do not contain the origin). In this study, the latter approach was used. **Figure A.5 (A)**, which plots the histogram of the β_i parameters, provides information regarding the distribution of the mentioned parameter for all voxels as well as the position of the cut-off value in this distribution. In **Figure A.5 (B)**, the parameter β_i of all voxels versus the standard deviation (STDev) of the angles within the uncertainty region (not corrected by β_i is plotted). This plot shows that the higher the variance in the orientation of the dipoles within the uncertainty region (due to the convoluted cortex), the larger its size will be. When there is a higher degree of variance among the lead field vectors in the neighborhood of a point, the uncertainty region has to be large enough to encompass all such points; which inadvertently results in a higher parameter β_i . Thus, to relieve the situation different remedies can be applied; one is to look at smaller neighborhoods around any given point to calculate the covariance. This will result in less variance but the sample size will also decrease, and consequently the estimation will be less accurate. Another approach is to use a rotational model and not fix the dipole orientation, which makes the problem size larger by a factor of three and can deviate from biophysical models regarding the generation of M/EEG signals from pyramidal cells. Finally, the objective approach can be taken (as we have here), which is to make sure for each constraint in the optimization problem to be feasible by limiting the value of β_i (and hence the inflation) with a statistical means.

A.2.4 The Effect of Normalization on Bias Removal

As explained in section 3.2.1, normalization acts as a tool to remove the inherent bias of adaptive beamformers toward deeper locations. In order to investigate this issue, **Figure A.6** plots the DLE for the aforementioned beamformers with respect to the depth of source activity (SNR is set to 5 dB). Depth in this figure is defined as the distance between each active node in a source configuration and the nearest electrode (see **Figure A.7**, which provides a sense of depth in different brain regions). The plots in **Figure A.6** display the localization error over 5 depth intervals. According to this figure, while error is generally expected to increase as activities become deeper (since SNR decreases with depth) the LCMV and RMVB act worse in localizing superficial

activities for high levels of noise. Although, the trend for LCVM-ND-DN and RMVB-ND-DN is as expected (higher errors for deeper sources). This observation reflects the mentioned inherent bias of the LCMV and RMVB towards deeper sources and the fact that normalization is quite successful in dealing with this depth-bias issue. Furthermore, it can be observed that once superficial sources are removed from the set of simulated sources, the RMVB results become slightly superior to the LCMV-ND-DN and comparable to the RMVB-ND-DN.

A.2.5 The Effect of Correlated Activities

Correlation between distinct sources is known to affect the performance of ESI techniques and particularly beamformers in a negative way. In order to investigate this issue, the third scenario was designed with 100 three node source configurations, which were highly correlated. **Figure A.8 (A)** and **(B)** display the results of this scenario in terms of DLE and output SNR, respectively. It can be seen that correlated activity can degrade the performance mostly in terms of SNR. Although, its effect on DLE is not so severe. Besides, it can be seen that neither normalization and denoising nor the robust modeling of the uncertainty can help to alleviate this condition. Therefore, dealing with correlated activity calls for a separate strategy. It should also be mentioned that the spatially distinct sources are less likely to be maximally correlated [108]. Besides, based on the results of this study, the perfect correlation between the sources that are also spatially coherent (i.e., the dipoles that have the same time-course of activity within the extended sources simulated in the fourth scenario), do not negatively affect the performance in a severe manner. Given the scope of this work, the full investigation of this issue can be the topic of a separate study.

A.2.6 Beamformers with Diagonal Loading

In beamformers with diagonal loading, the covariance matrix of the measurements is replaced with a regularized version in which a constant factor λ of the unity matrix is added to the measurements covariance matrix. While λ can determine the ultimate quality of reconstruction, its optimal value is unknown. It is proposed in [122] that λ be set to $10\sigma^2$ where σ^2 is the variance of noise. In this study, we have simulated LCMV

with diagonal loading (LCMV+DL) using this factor as well as two other values, namely, σ^2 and $100\sigma^2$. The variance of noise in our simulation is calculated by averaging over the diagonal elements of the noise covariance matrix which is determined using baseline. **Figure A.9** and **Figure A.9**, which show the statistical results of LCMV+DL in the second and fourth scenarios, suggest that diagonal loading can improve the reconstruction compared to conventional LCMV. Although, the performance is overall worse than other beamformers introduced in this study. More specifically, it can be observed that RMVB outperforms LCMV+DL based on all metrics and in almost all situations. The only exception is when the noise level is high (SNR = 5 dB) and the brain activities are focal. In this case, by setting the diagonal loading factor to a high value ($100\sigma^2$), one can improve the performance in the sense of output SNR. Although, RMVB still works better than LCMV+DL according to other metrics and in all other situations. It should also be noted that the mentioned improvement of output SNR is obtained at the cost of losing spatial resolution in the solution of LCMV+DL. Furthermore, RMVB-ND-DN outperforms LCMV+DL in all situations and based on all metrics.

A.3 Supporting Information for Chapter 4

See **Figures A.11, A.12, A.13, A.14** and **A.15**.




Figure A.1: Average localization error for 3-node network configurations using three different regularization parameter selection techniques, namely, the l-curve (LCV), the generalized cross validation (GCV) and the discrepancy principle (DCP). Two noise types, namely, additive white Gaussian noise (AWGN) and additive colored Gaussian noise (ACGN) with SNRs of 5 dB and 20 dB are simulated. The simulations are repeated for both EEG and sEEG modalities.




Figure A.2: The source space from which source locations are selected. The red spheres in this figure show the sEEG electrodes and blue spheres show the source space from which source locations are selected. As it can be seen, while sources are confined to the cortex, they can be very deep.




Figure A.3: The time-course assigned to each dipole in the single-node source scenario.

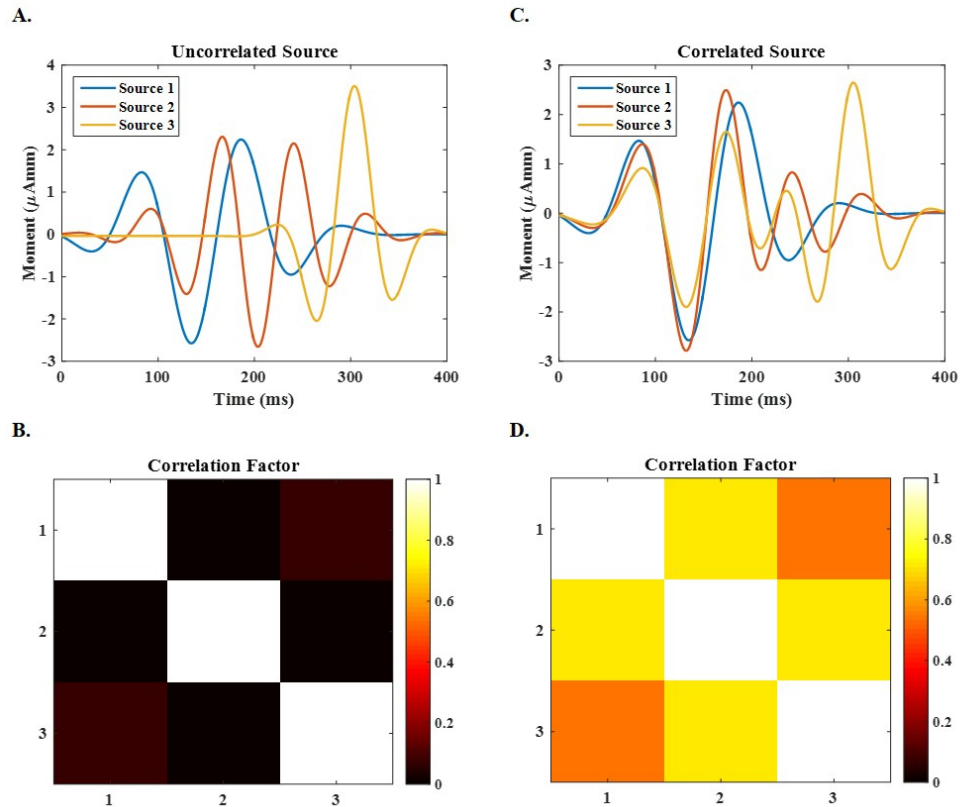


Figure A.4: The time-courses and correlation diagrams in the second and third scenarios. (A) The time-courses assigned to each node in the 3-node uncorrelated source scenario and (B) the corresponding correlation coefficients among the nodes. As it can be seen, the correlation coefficient between different nodes is small in this scenario. (C) The time-courses assigned to each node in the 3-node correlated source scenario and (D) the corresponding correlation coefficients among the nodes. Unlike the second scenario, nodes are highly correlated in this case.

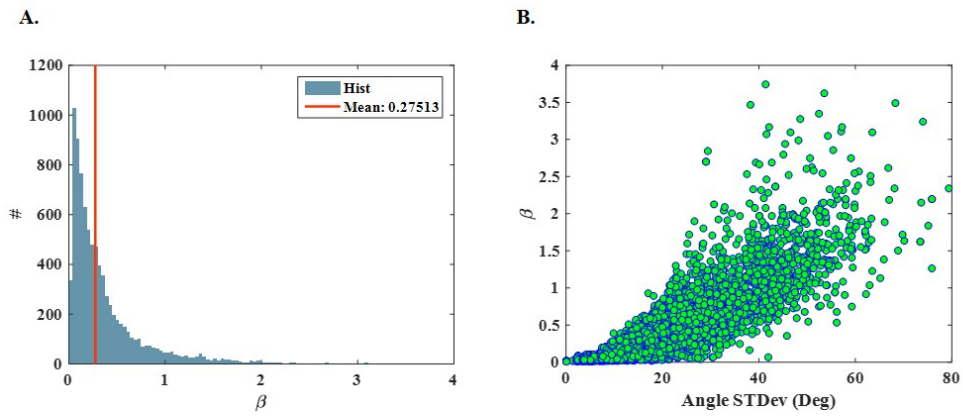


Figure A.5: The histogram of the parameter β and its correlation with angle standard deviation within a region. (A) Plots the histogram of the parameter β and provides information regarding the distribution of the mentioned parameter for all voxels. Red vertical line marks the position of the cut-off value in this distribution. (B) Parameter β of all voxels versus the standard deviation (STDev) of the angle between the center and the neighbors' orientations in each uncertainty region.

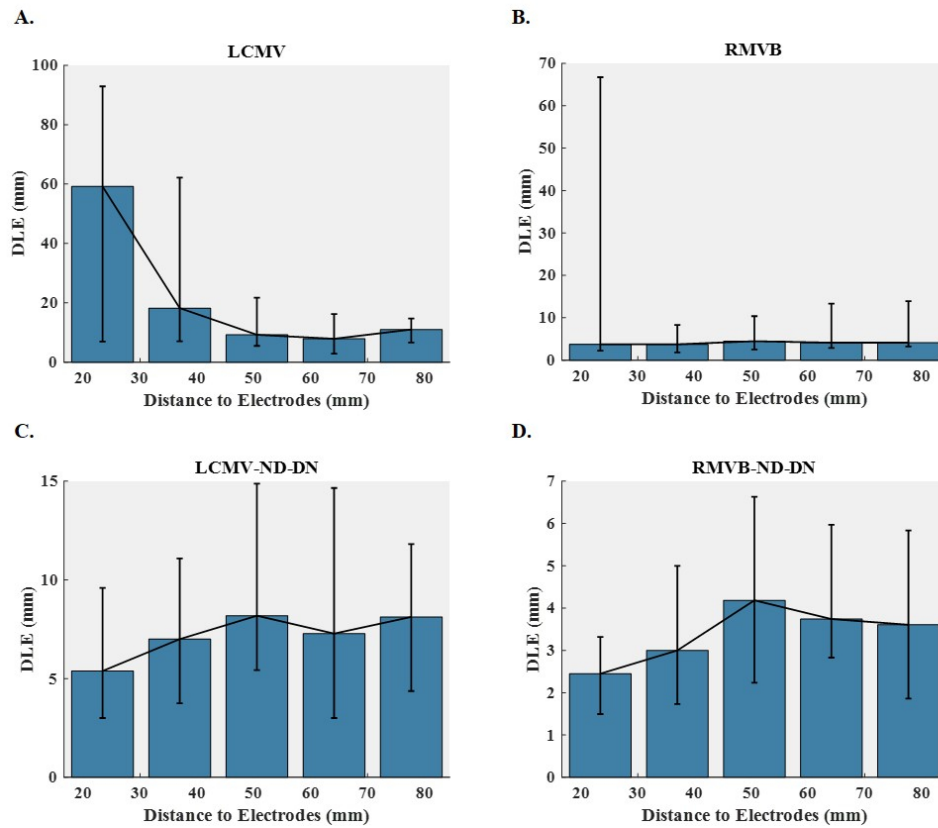


Figure A.6: The DLE results versus depth for the three-node uncorrelated source scenario. The plots display the median DLE along with the error bars over five depth intervals for (A) LCMV, (B) RMVB, (C) LCMV-ND-DN and (D) RMVB-ND-DN. SNR in this figure is set to 5dB and depth is defined as the distance between each active node in a source configuration and the nearest electrode in a standard 128-channel BioSemi cap.




Figure A.7: The color-coded depth of different locations in the brain. Depth in this diagram is defined as the distance to the nearest electrode in a standard 128-channel BioSemi cap.

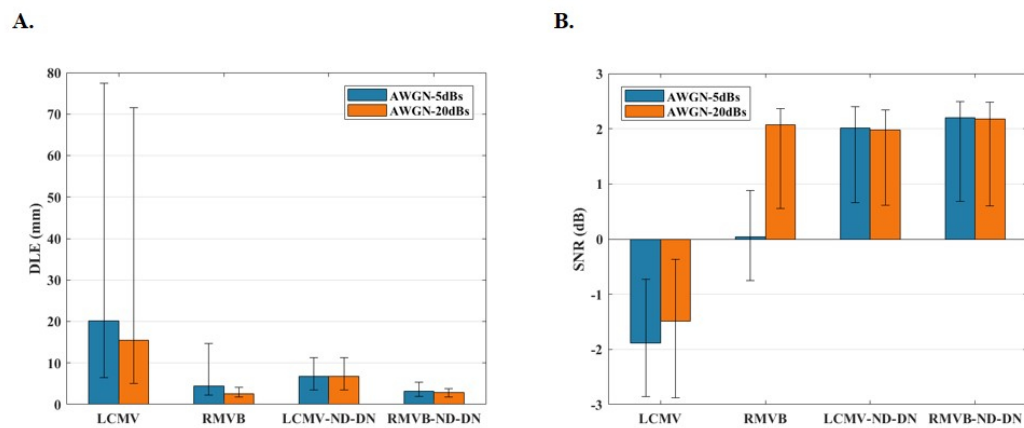


Figure A.8: The Monte Carlo simulation statistics for the three-node correlated source scenario. (A) The median DLE results for SNRs of 5dB and 20 dB and for four types of beamformers namely, LCMV, RMVB, LCMV-ND-DN and RMVB-ND-DN. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration.

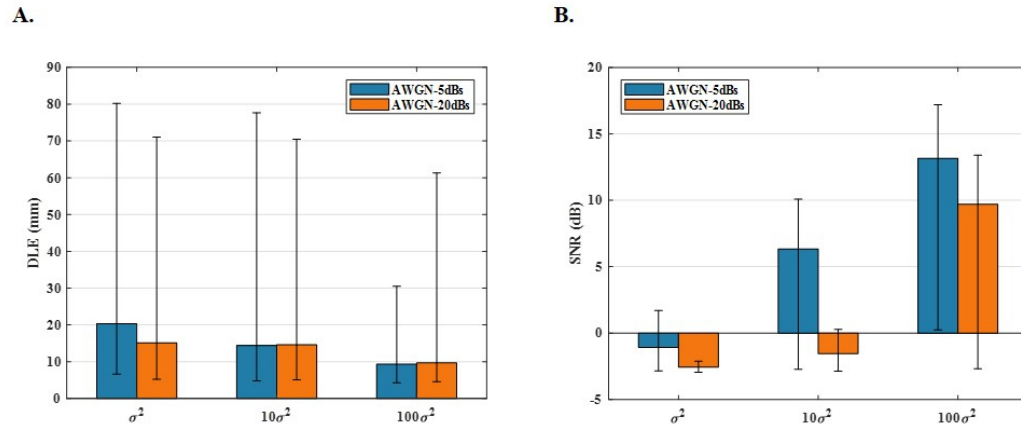


Figure A.9: The Monte Carlo simulation statistics for the three-node uncorrelated source scenario. (A) The median DLE results for SNRs of 5dB and 20dB and for diagonal loading beamformer with three factors of σ^2 , $10\sigma^2$ and $100\sigma^2$. The error bars mark the first and third quartiles (of DLE distribution). (B) The median output SNR results along with the first and third quartile error bars for the same configuration.




Figure A.10: The Monte Carlo simulation statistics for the extended source scenario. The median MCC results along with the first and third quartile error bars for diagonal loading beamformer with three factors of σ^2 , $10\sigma^2$ and $100\sigma^2$ and for SNR of 20 dB.




Figure A.11: Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were uniformly undersampled retrospectively at rates 2×2 , 3×2 , 4×2 and 5×2 in $k_y - k_z$ plane, which are approximately equivalent to net acceleration rates 4, 5, 6 and 7 (including the ACS lines and an elliptical mask). These data were then reconstructed using GRAPPA, SPIRiT, l_1 -SPIRiT, RAKI and sRAKI (from top to bottom). Acceleration rate was set no higher than 2 for k_z dimension, since the size of data along this dimension was small (20 lines in total and 10 lines for ACS). For RAKI, a 3-layer network was designed with a kernel size of 2×2 (with dilations equaling acceleration rates to match the undersampled uniform pattern) for the first layer and a kernel size of 1×1 for subsequent layers. Note that this 2D undersampling is different from the original RAKI paper, thus the network architecture may be sub-optimal. The learning rate and number of epochs for RAKI were tuned to 0.05 and 2000 iterations, respectively. Fully-sampled images are also displayed in the first column as a reference for comparison. While RAKI is robust, GRAPPA is very sensitive to noise with increasing rates. In addition, RAKI outperforms SPIRiT, but RAKI and sRAKI perform comparatively, similar to the relationship between GRAPPA and SPIRiT.



Figure A.12: Three $k_y - k_z$ undersampling patterns were tested. Poisson disc (top), uniform-density random (middle) and variable-density random (bottom) with 4-fold acceleration.




Figure A.13: Reformatted right coronary artery (RCA) images from a 3D targeted coronary MRI dataset. The data were retrospectively undersampled with the three different patterns shown in Supporting **Figure A.12**. These data were then reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI. The results show that sRAKI is more resilient to noise amplification compared with SPIRiT, regardless of undersampling pattern.




Figure A.14: Noise sensitivity of the reconstruction methods are shown on reformatted right coronary artery images. Additive Gaussian noise was added to the datasets retrospectively, and the reported SNR was measured at aorta (signal power at aorta divided by noise power in an empty region of interest), with the original dataset having SNR of 50. The datasets were then retrospectively undersampled at rate 4 and reconstructed using SPIRiT, l_1 -SPIRiT and sRAKI. sRAKI is more robust against noise of data compared with SPIRiT. However, noise amplification becomes evident with increasing levels of noise compared with l_1 -SPIRiT.

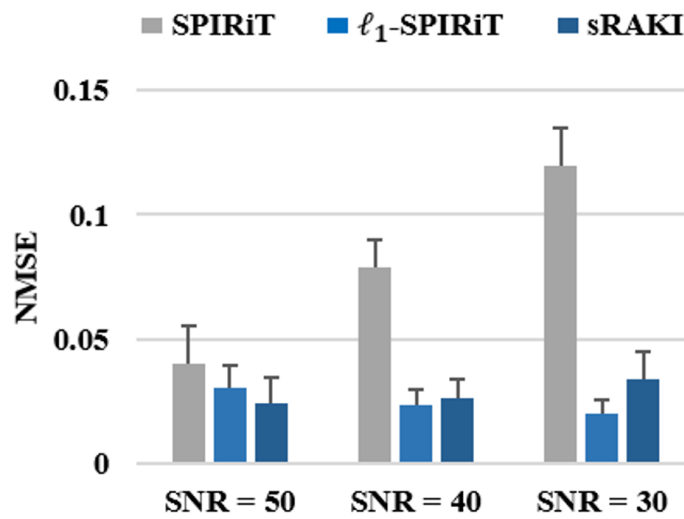


Figure A.15: Normalized mean squared error (NMSE) of reconstruction across all subjects with 4-fold acceleration for the experiment setup described in **Figure A.14**. Error bars represent standard deviation across subjects.